AARHUS UNIVERSITY
BUSINESS AND SOCIAL SCIENCES
DEPARTMENT OF ECONOMICS AND BUSINESS

CREATES
Center for Research in Econometric Analysis of Time Series

Is the Quantity Theory of Money Useful in Forecasting U.S. Inflation?

Markku Lanne, Jani Luoto and Henri Nyberg

CREATES Research Paper 2014-26

Department of Economics and Business
Aarhus University
Fuglesangs Allé 4
DK-8210 Aarhus V
Denmark

Email: oekonomi@au.dk
Tel: +45 8716 5515

Is the Quantity Theory of Money Useful in Forecasting U.S. Inflation?†

Markku Lanne
University of Helsinki and CREATES

Jani Luoto
University of Helsinki

Henri Nyberg*
University of Helsinki

Abstract

We propose a new simple model incorporating the implication of the quantity theory of money that money growth and inflation should move one for one in the long run, and, hence, inflation should be predictable by money growth. The model fits postwar U.S. data well, and beats common univariate benchmark models in forecasting inflation. Moreover, this evidence is quite robust, and predictability is found also in the Great moderation period. The detected predictability of inflation by money growth lends support to the quantity theory.

Keywords: Money growth, transfer function model, low-pass filter

JEL classification: C22, E31, E40, E51

† Financial support from the Academy of Finland is gratefully acknowledged. The first author also acknowledges financial support from CREATES (DNRF78) funded by the Danish National Research Foundation, while the second author thanks the OP-Pohjola Group Research Foundation. Part of this research was done while the second author was visiting the Bank of Finland, whose hospitality is gratefully acknowledged.

* Corresponding author. Department of Political and Economic Studies, University of Helsinki, P.O.Box 17 (Arkadiankatu 7), FIN–00014 University of Helsinki, Finland, e-mail: henri.nyberg@helsinki.fi.

1 Introduction

Allegedly forecasting inflation has recently become more difficult. For instance, Stock and Watson (2007) show that improving upon univariate forecasts of U.S. inflation has become hard since the mid-1980s, and Lenza (2006) provides similar evidence for euro area inflation since the early 2000s. Money growth, one of the most often considered predictors of future inflation, has also been reported to have lost its predictive power, and typically this loss in predictability has been attributed to monetary policy. In particular, Sargent and Surico (2011) point out that the aggressive response to inflation by the central bank prevents the movements in money growth to show up in inflation, and, hence, it has had marginal predictive power only during periods with no credible commitment to fight inflation. In accordance with this, D'Agostino and Surico (2012) find little evidence in favor of the predictability of inflation by money growth in the U.S. since the mid-1980s, i.e., in the Great moderation period.

The rationale behind forecasting inflation by money growth is the quantity theory of money, according to which these variables should move one for one in the long run, and diminished predictability can thus be seen as weakened evidence in favor of that theory. This conclusion has quite commonly been reached with U.S. and international data in recent research, where the quantity theory has been taken as the starting point. Teles and Uhlig (2013) find that it has since 1990 become more difficult to establish a long-run relationship between money growth and inflation. D'Agostino and Surico (2009), on the other hand, claim that the information in U.S. money growth concerning inflation is embedded in measures of global liquidity. Berger and Österholm (2008), and Assenmacher-Wesche and Gerlach (2007) indeed do find predictability of U.S. inflation by money growth only when accounting for the impact of other variables, such as interest rates. The results of Amisano and Fagan (2013), and Amisano and Colavecchio (2013), in turn, suggest that nonlinearities need to be taken into account in order to uncover the predictability of inflation by money growth.

In much of the previous literature cited above, frequency-domain methods or relatively complicated nonlinear models, often containing additional variables besides money growth and inflation, have been employed. In contrast, in this paper, we put forth a new simple econometric model involving only money growth and inflation and incorporating the central implication of quantity theory that money growth and inflation should move one for one in the long run. In oder to capture the low-frequency (or long-run) component of money growth, we employ an exponentially weighted moving average (EWMA) filter. It turns out that the model (discussed in detail in Section 2 below) can be written as a classical transfer function (TF) model with restrictions on the parameters and an autocorrelated error term. Hence, it can be estimated in a straightforward manner by the method of maximum likelihood (ML), and the properties of the ML estimator (under mild regularity conditions) are well known.

We apply the new model to quarterly postwar U.S. consumer price and GDP deflator inflation, and it seems to fit the data quite well. We also find a clear improvement in inflation predictability by the low-frequency component of money growth over and above the own history of inflation. In particular, our model beats the benchmark univariate autoregressive model as well as the naïve (Atkeson and Ohanian (2001)) and random walk forecasts often entertained in the inflation forecasting literature. Moreover, with the exception of a short period in the latter half of the 1980s, the new model is incessantly superior in the entire forecasting period starting in 1970. The latter finding is remarkable in that breaks in predictability, often attributed to monetary policy, have in the previous literature typically been detected, and the existing evidence in favor of predictability during the Great moderation is weak. Hence, our results lend quite robust support to the quantity theory of money.

The plan of the paper is as follows. In Section 2, we describe the new model, termed EWMA-TF model, and discuss estimation, inference and forecasting. Section 3 contains the estimation results based on consumer price inflation, while in Section 4, we report the (pseudo) out-of-sample forecasting results based on a recursive scheme. To gauge the robustness of the results, in Section 5, we report results based on the GDP deflator inflation and a rolling forecasting scheme. These robustness checks do not overturn our main findings. Finally, Section 6 concludes.

2 EWMA-TF Model

In this section, we introduce our econometric model and discuss its estimation and use in forecasting. The model incorporates the core implication of the quantity theory of money that inflation and money growth move one-for-one, but following Lucas (1980), we assume this relationship to hold only in the long run. Therefore, instead of including money growth in the model as such we concentrate on its low-frequency component, obtained with an exponentially weighted moving average (EWMA) filter. As shown below, our model can be written in the form of a standard restricted transfer function (TF) model with an autocorrelated error term assumed to follow an AR(p) process (see, e.g., Box and Jenkins, 1976, Chapter 10) whose statistical properties are well known, Thus, hereafter we refer to the model as the EWMA-TF(p) model.

2.1 Model

In line with the quantity theory of money, our model postulates a one-for-one relationship between inflation, π_t, and the long-run component of money growth, x_t:

$$\pi_t = \mu + x_t + u_t, \quad t = 1, \ldots, T, \tag{1}$$

where μ is an intercept term, u_t is an error term following an autoregressive process of order p (an AR(p) process), and inflation is computed as $\pi_t \equiv \log(P_t) - \log(P_{t-1})$, with P_t the aggregate price level. Specifically, $\phi(L)\, u_t = \epsilon_t$ where ϵ_t is a sequence of independently and identically distributed random variables with mean zero and variance σ_ϵ^2, L denotes the usual lag operator, and the polynomial

$\phi(z) \equiv 1 - \phi_1 z - \ldots - \phi_p z^p$ has its zeros outside the unit circle. The latter specification is assumed to capture any high-frequency autocorrelation in inflation π_t not accounted for by the long-run component x_t, and according to the results in Section 3 this assumption seems adequate in our empirical application.

The long-run component of money growth x_t is obtained as an exponentially weighted moving average (EWMA) as

$$x_t = x_{t-1} + \alpha(m_t - x_{t-1}), \tag{2}$$

where $m_t \equiv \log(M_t) - \log(M_{t-1})$ with M_t the money supply. The properties of the one-sided filter (2) are determined by the parameter α ($0 < \alpha \leq 1$), and it has large power at low frequencies when the value of this smoothing parameter lies close to zero. This is illustrated in Figure 1 that depicts its gain function with α equal to our estimate from quarterly U.S. data in Section 3 ($\alpha = 0.035$). The gain is large only at frequencies corresponding to periods greater than 16 quarters, or four years (frequencies lower than $\pi/8$), while it is virtually zero at higher frequencies. Hence, our filtered component indeed manages to capture only the long-run movements of money growth.

2.2 Estimation and inference

Our econometric model is actually a standard transfer function model with restrictions on the parameters and an autocorrelated error term. To see this, write (2) as

$$x_t = \frac{\alpha}{1 - (\alpha - 1) L} m_t$$

and plug x_t into (1) to obtain

$$\pi_t = \mu + \frac{\alpha}{1 - (\alpha - 1) L} m_t + \phi(L)^{-1} \epsilon_t. \tag{3}$$

The parameters of the model can thus be consistently estimated by the method of (conditional) maximum likelihood (see, e.g., Harvey (1981b, Chapter 7)). Specifically, assuming normality of ϵ_t, the log-likelihood function conditional on the first

$p+1$ observations can be expressed as

$$l_T(\theta) = -\frac{T-p-1}{2}\log\left(2\pi\sigma_\epsilon^2\right) - \frac{1}{2\sigma_\epsilon^2}\sum_{t=p+2}^{T}\epsilon_t^2, \tag{4}$$

where ϵ_t is obtained recursively for $t = p+2, p+3, ..., T$ by setting $\epsilon_{p+1} = 0$ and writing (3) as

$$\epsilon_t = \left[1 - (\alpha - 1) L\right] \phi(L)\pi_t - (2-\alpha)\,\phi(1)\mu - \alpha\phi(L)m_t + (\alpha - 1)\,\epsilon_{t-1}.$$

Under mild regularity conditions, the maximum likelihood estimator is asymptotically normally distributed, and standard inference, including conventional diagnostic checks of the adequacy of the model, can be employed.

Box and Jenkins (1976, Chapter 11) and subsequently, among others, Liu et al. (2010) (see also the references therein) discuss model selection in transfer function models, and put forth a strategy involving cycles of identification, estimation, and diagnostic checking. When there is little a priori knowledge on the relations between the variables, these procedures are sensible, while the structure of our model is dictated by the quantity theory, and hence, model selection reduces to finding the correct order p of the lag polynomial $\phi(z)$ in the EWMA-TF(p) model (3). To that end we employ information criteria and standard diagnostic checks to guard against remaining error autocorrelation and conditional heteroskedasticity.

2.3 Forecasting

Optimal forecasts of π_t from (1) or, equivalently, from (3) h periods ahead at time point T are obtained recursively. As x_t also enters the model contemporaneously, for forecasting, a univariate model (with an error term independent of ϵ_t) must be specified for it (see, Harvey (1981b, Chapter 7)), and we assume that its dynamics are adequately captured by an autoregressive process. Forecasts are thus obtained in three stages. First, we fix the initial value of x_t at the average of m_t over the entire sample, and compute the sequence of smoothed money growth x_t from (2) for $t = 1, ..., T$. Second, we estimate an AR(s) process for m_t (using observations up to period T), and plugging its h-period forecast into Eq. (2), compute the h-period forecast $x_{T+h|T}$. Third, based on the forecast $x_{T+h|T}$, we obtain $\pi_{T+h|T}$, the h-period forecast of inflation π_t as

$$\pi_{T+h|T} = \widehat{\mu} + x_{T+h|T} + u_{T+h|T}, \tag{5}$$

where $\widehat{\mu}$ is the ML estimate of μ, and $u_{T+h|T}$ is the h-period forecast of $u_T \equiv \pi_T - \widehat{\mu} - x_T$ computed recursively (starting with the one-period forecast) from the AR(p) model $\widehat{\phi}(L)\, u_t = \epsilon_t$ with $\widehat{\phi}$ the ML estimate of the lag polynomial coefficients.

Our interest concentrates on the out-of-sample forecasting performance of the low-frequency component of money growth for inflation. In particular, we are interested in finding out about the marginal predictive ability of the money growth component over and above that of the own history of inflation. The relevant univariate benchmark for π_t is an adequate autoregressive moving average (ARMA) model because, under our assumption that m_t follows an AR(s) process, (3) is a sum of two independent autoregressive processes, and, hence, the univariate process of π_t is an ARMA process (see, e.g., Harvey (1981a, Chapter 2)). In our empirical analysis in Section 4, an AR model is deemed adequate for U.S. inflation.

In order to assess the forecast performance of our model vis-à-vis the univariate benchmark model, we conduct a (pseudo) out-of-sample forecasting excercise in Section 4. Specifically, we recursively compute forecasts from (5) and a univariate AR(p) model for inflation by expanding the estimation sample by one period at a time until the end of the data series to obtain sequences of h-quarter forecasts. We then measure forecast accuracy by the mean squared forecast error (MSFE) criterion. If the value of this criterion for forecasts computed from the EWMA-TF model lies below that of the univariate AR model, following Box and Jenkins (1976, Chapter 11), we call the smoothed money growth a leading indicator of inflation, and conclude that the quantity theory of money is useful in forecasting inflation.

3 Estimation Results

Our quarterly data set covers the period from 1947:Q1 to 2013:Q4. Following Lucas (1980) and Chan et al. (2013), among others, we estimate the models for the U.S. inflation computed as $\pi_t = 400 \times \log(P_t/P_{t-1})$, where P_t is the consumer price index (CPI) for all urban consumers.[1] Money supply is measured by the M2 money stock, and the money growth series is analogously computed as $m_t = 400 \times \log(M2_t/M2_{t-1})$. Both the CPI and M2 series are seasonally adjusted. The source of all the data is the FRED database of the Federal Reserve Bank of St. Louis, with the exception of the M2 money stock for the period from 1947:Q1 to 1958:Q4 that is constructed by Sargent and Surico (2011), whose work, in turn, builds upon Balke and Gordon (1986).

Although our main objective is to examine the predictive power of the smoothed money growth for inflation, and we mostly concentrate on (pseudo) out-of-sample forecasting, we start by reporting estimation results from the entire sample period. These results should provide some evidence on the fit of the proposed model and yield information on the statistical properties of U.S. inflation.

In Table 1, we present the estimation results of the AR(4) and EWMA-TF(4) models. The lag length is in each case selected by a sequential testing procedure, starting with a fifth-order model, and sequentially proceeding to the model, where the coefficient of the longest lag is significant at the 5% level. The fourth-order model is also selected by the Bayesian information criterion. The EWMA-TF(p) model reduces to the usual AR(p) model, when the low-frequency component of m_t (i.e. x_t) is excluded from the model. It is seen from (3) that this is the case when α equals zero, and, hence, our main interest focuses on the smoothing parameter α. In the fourth-order EWMA-TF model, the ML estimate of α equals 0.035, suggesting that it is indeed the permanent component of inflation that x_t captures (with $\alpha = 0.035$ the EWMA filter (2) has most of its power at low frequencies, as indicated

[1] As a robustness check, we also consider the GDP deflator as a measure of inflation in Section 5. Overall, the main conclusions remain intact irrespective of the inflation series.

by the gain function in Figure 1). This is also visible in Figure 2, which depicts the evolution of x_t with the time series of m_t superimposed. Comparison of the estimates of the autoregressive coefficients of the AR and EWMA-TF models also lends support to this claim: the sums of the estimated autoregressive coefficients of these models are 0.78 and 0.66, respectively, indicating that after the inclusion of x_t, the portion of the variation of inflation left to be explained by its own lags exhibits substantially reduced persistence.

The EWMA-TF model has higher values of both the maximized log-likelihood function and the Bayesian information criterion than the AR model. However, not too much emphasis should be laid on these in-sample comparisons that implicitly assume a one-period-ahead forecast horizon, as the EWMA-TF model is designed for long-horizon forecasting, and evidence on the superiority of the EWFA-TF model vis-à-vis the AR model in forecasting is provided in Section 4. According to the diagnostic checks, there is little evidence of remaining autocorrelation or conditional heteroskedasticity in the residuals of either model. In fact, the inclusion of x_t tends to somewhat reduce remaining conditional heteroskedasticity. All in all, it seems fair to state that the EWMA-TF(4) model is a statistically adequate specification for U.S. inflation.[2]

4 Out-of-Sample Forecasting

In this section, we report results of out-of-sample forecasting exercises with the EWMA-TF model estimated in Section 3, with emphasis on its multistep forecast performance compared to that of the benchmark AR model. As discussed in Section 3, the AR model is a natural benchmark because model (1) reduces to an AR(MA) model when excluding the long-run component x_t (under the assumption that

[2] One obvious extension of model (1) is obtained by not restricting the coefficient of x_t equal to unity in accordance with the quantity theory. However, with our data, this restriction cannot be rejected (the p-value in the likelihood ratio test equals 0.96), which lends further support to the adequacy of the restricted EWMA-TF(4) model and the quantity theory.

money growth is not hepful in forecasting inflation). In addition to these two models, we consider the naïve forecast of Atkeson and Ohanian (2001) and the random walk forecast popular in the previous inflation forecasting literature. For the most part, the forecast period runs from the first quarter of 1970 until the last quarter of 2013 (1970:Q1–2013:Q4). However, we also report results for the shorter forecast period starting in the first quarter of the year 1990 (the Great Moderation period). The results in this section are based on the expanding window approach in updating the parameters of the models. In Section 5, we also report forecast results based on the rolling window approach, but the general conclusions remain unchanged.

The estimation sample starts at 1947:Q1, and it is extended by one quarter at a time. Model selection is also performed recursively, selecting the lag length p in the EWMA-TF(p) and AR(p) models by the Bayesian information criterion (BIC) with a maximum lag length of five. More or less the same results are obtained with the Akaike information criterion (AIC).[3] In fact, perhaps somewhat surprisingly, both criteria typically select the fourth-order model ($p = 4$) for both models, with very few exceptions for the EWMA-TF model. In line with the previous inflation forecasting literature, we consider forecast horizons of four, eight, and twelve quarters, i.e., of one, two and three years. These horizons were also examined by D'Agostino and Surico (2012), albeit with emphasis on the eight-quarter horizon, and they are the most interesting ones, as we expect the smoothed money growth to have predictive power in the intermediate and long forecast horizons.

In Table 2, we report the MSFE statistics and the results of the test for equal predictive accuracy of Diebold and Mariano (1995) and West (1996) related to them for the full out-of-sample period. Following the common practice in the inflation forecasting literature, forecasts are computed for average inflation over the h-period forecast horizon. The EWMA-TF model clearly outperforms the univariate AR model. In addition to the MSFE criterion, following Chan et al. (2013), we

[3] These results are not reported to save space, but they are available upon request.

study the qualitative differences in predictive ability between the models by computing the fraction of quarters, when the EWMA-TF model is more accurate, and also according to this criterion, the EWMA-TF is the winner. However, the tests results are mixed. While the qualitative differences are statistically significant at the 1% level, by the MSFE criterion, there is a significant difference between the models (at the 10% level) only at the four-quarter horizon. A likely explanation to this discrepancy are the large differences in the MSFEs in favor of the AR model in the short period from the mid-1980s to the beginning of the 1990s (see Figure 3 and the related discussion below). Nevertheless, the smoothed M2 predictor x_t appears to have substantial long lasting out-of-sample predictive power.

Finally, Table 2 also reports the MSFEs of the random walk and naïve forecasts of inflation. Especially the latter has often been found superior in the inflation forecasting literature, and it indeed turns out to be more accurate than the AR model. However, the fact that it is beaten by our EWMA-TF model at all horizons considered further reinforces the evidence in favor of the marginal predictive ability of the smoothed money growth.

A long line of previous research (see, for instance, Stock and Watson (2007), Clark and Doh (2011), D'Agostino and Surico (2012), and Chan et al. (2013), and the references therein) points at breakdowns in the one-for-one movement of money growth and inflation implied by the quantity theory as a potential explanation to changes in the forecast performance of commonly employed models. Because also our model in (1) directly utilizes this central implication, we next explore whether the relative forecast performance of the EWMA-TF model compared to the conventional AR model is stable over time. To this end, following D'Agostino and Surico (2012), we consider the smoothed mean squared forecast errors (SMSFE) using a rolling window of 31 quarters to examine the predictive performance of the models.[4] To facilitate comparisons, we report the relative SMSFE as a ratio of the SMSFEs from the EWMA-TF model to the benchmark AR model. Hence,

[4] Stock and Watson (2009) used a window of 15 observations (quarters). The figures in that case are similar as those in Figure 3.

the values of the relative SMSFE below unity indicate that the forecasts of the former model with dependence on the smoothed money growth are on average more accurate than those of the AR model.

Figure 3 depicts the relative SMSFE of the EWMA-TF model relative to that of the univariate AR model for the forecast horizons of four, eight, and 12 quarters. Values above unity are rare and concentrated on a relatively short time period between the mid-1980s and the beginning of the 1990s, where deviations from unity also tend to be large, and the differences between the models increase with the forecast horizon. With the exception of this period, the smoothed money growth thus has clear predictive power and can be considered a leading indicator of inflation. While also D'Agostino and Surico (2012) pointed out this period as exceptional, our general conclusions deviate from theirs in that they found only very few periods, all related to specific monetary policy regimes, where the money growth is a useful predictor.

Based on the findings in Figure 3, it seems that since the beginning of the 1990s, the smoothed money growth has (again) had predictive power for inflation in the spirit of the quantity theory. To study this further, we repeat the analysis of Table 2 for the period of the Great Moderation and beyond (1990:Q1–2013:Q4) in Table 3. The results are essentially the same as those for the whole out-of-sample period above. However, in the shorter period also the differences in the absolute magnitudes of the squared forecast errors between the models are statistically significant (at the 10%, 5% and 1% levels of significance at the four eight and 12-quarter horizons, respectively). In the same way as in the full out-of-sample period, the EWMA-TF model yields superior forecasts also compared with the naïve (and random walk) forecasts, with even wider margins. Especially the fact that the EWMA-TF model outperforms the naïve forecasts in the Great Moderation period, is noticeable. As concluded by Stock and Watson (2007), inter alia, the naïve forecast has turned out hard to beat in this period.

5 Robustness checks

In this section, we check whether our findings in Section 4 are sensitive to the particular inflation series considered or choices made in producing the results. In particular, we compare the results based on the rolling window approach to those above in the out-of-sample forecasting exercises and report results based on the GDP deflator inflation. Overall, our conclusions turn out to be quite robust.

In Section 4, we used the expanding window approach when updating the parameters (and selecting the specification) in the forecasting period. The results of an alternative, rolling window, approach are reported in Table 4 for the entire out-of-sample period as well as the subsample period considered in Table 3. The size of the estimation window is fixed at 91 observations which is the size of the first estimation sample also in the expansive window case. Results based on model selection by the BIC are reported, but like with the expanding window approach, the AIC leads to no substantial changes. The results are essentially the same as obtained with the expansive window approach: The EWMA-TF model including the smoothed money growth clearly outperforms the AR model for inflation. In the rolling window case, the statistical significance of the differences in forecast accuracy are tested using the (unconditional) test of Giacomini and White (2006) which coincides with the test employed in Section 4 in this case. Note that the naïve and random walk forecasts are the same as in the expanding window case and therefore they are not reproduced in Table 4.

The main analysis in Section 4 concerns the CPI inflation. In Table 5, we report the results for inflation based on the GDP deflator for both forecast periods considered. As above, the results based on using the BIC to select the lag length are reported, but the AIC yields similar conclusions. As far as the relative performance of the EWMA-TF model vis-à-vis the univariate AR model is concerned, the results are essentially the same as those for the CPI inflation, albeit with a narrower margin. In other words, the EWMA-TF model yields superior forecasts also for the GDP inflation, suggesting that the smoothed money has statistically significant predictive power for inflation. However, while the qualitative differences in forecast accuracy are again statistically significant, this does not apply to the quantitative differences even in the more recent forecast period. The only major difference is that for the GDP deflator, the naïve and random walk forecasts are much more accurate than in the case of the CPI inflation, beating even the EWMA-TF model.

6 Conclusions

In this paper, we have proposed a new model for capturing the long-run comovement of inflation and money growth. The model is inspired by the classical quantity theory of money, and it incorporates its central implication that money growth and inflation should move one for one in the long run. Because the long-run (or low-frequency) component of inflation is obtained by an exponentially weighted moving average filter and the model can be written in the form of a transfer function model, we call it the EWMA-TF model. Its parameters can be consistently estimated by the method of maximum likelihood, and (under regularity conditions) standard asymptotic inference applies.

We applied the EWMA-TF model to modeling and forecasting postwar U.S. inflation. Diagnostically the model turned out to fit both the CPI and GDP inflation series well, and in forecasting it manages to beat a number of univariate benchmark models entertained in the literature. In particular, for the CPI inflation, its forecasts are more accurate than those of the univariate AR model (obtained as its special case by excluding the effect of money growth on inflation) and the naïve forecasts of Atkeson and Ohanian (2001). With the exception of a relatively short period in the late 1980s, the EWMA-TF model is superior to univariate alternatives, indicating that money growth indeed has predictive power for inflation in accordance with the quantity theory. The fact that this is the case also in the period of the Great moderation is quite remarkable in that in the previous literature money growth has been found a significant predictor of inflation only occasionally. Typically periods exhibiting such predictability have been attributed to weak monetary policy, captured by allowing for nonlinearities or augmenting the model with additional variables. Our findings, based on a relatively simple econometric model compared to those previously entertained, can be interpreted as quite robust support to the quantity theory of money.

In this paper, we concentrate on the U.S., but much of the previous empirical literature on the predictability of inflation by money growth is related to the euro area, and in future work it might be interesting to study the performance of our model in that as well as other geographic areas. Also, while the EWMA-TF model seems adequate for the U.S. data, this may not be the case in general, given that relatively complicated nonlinear models have widely been considered in the recent related literature. The development of nonlinear extensions may therefore be called for. Specifically, for some other data sets, allowing for switching between (monetary policy) regimes may be needed.

References

Amisano, G., and R. Colavecchio (2013). Money growth and inflation: Evidence from a Markov switching Bayesian VAR. DEP Discussion Papers Macroeconomics and Finance Series 4/2013.

Amisano, G., and G. Fagan (2013). Money growth and inflation: A regime switching approach. *Journal of International Money and Finance* 33, 118–145.

Assenmacher-Wesche, K., and S. Gerlach (2007). Money at low frequencies. *Journal of the European Economic Association* 5, 534–542.

Atkeson, A., and L.E. Ohanian (2001). Are Phillips curves useful for forecasting inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* 25, 2–11.

Balke, N.S., and R.J. Gordon (1986). Appendix B: Historical Data. In *The American Business Cycle: Continuity and Change*. National Bureau of Economic Research Conference on Research in Business Cycles, ed. Robert J. Gordon, 781–850. Chicago, University of Chicago Press.

Berger, H., and P. Österholm (2008). Does money matter for US inflation? Evidence from Bayesian VARs. IMF Working Papers, no. 08-76.

Box, G.E.P., and G.M. Jenkins (1976). *Time Series Analysis: Forecasting and Control*. Revised Edition. San Francisco, Holden-Day.

Chan, J.C., G. Koop, and S.M. Potter (2013). A new model of trend inflation. *Journal of Business and Economic Statistics* 31, 94–106.

Clark, T.E., and T. Doh (2011). A Bayesian evaluation of alternative models of trend inflation. Working paper 11-34, Federal Reserve Bank of Cleveland.

D'Agostino, A., and P. Surico (2009). Does Global liquidity help to forecast U.S. inflation? *Journal of Money, Credit and Banking* 41, 479–489.

D'Agostino, A., and P. Surico (2012). A century of inflation forecasts. *Review of Economics and Statistics* 94, 1097–1106.

Diebold, F.X., and R.S. Mariano (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 134–144.

Giacomini, R., and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74, 1545–1578.

Harvey, A.C. (1981a). *Time Series Models*. Oxford, Philip Allan.

Harvey, A.C. (1981b). *The Econometric Analysis of Time Series*. Oxford, Philip Allan.

Lenza, M. (2006). Does money help to forecast inflation in the euro area?, mimeo, European Central Bank.

Liu, J.M., R. Chen, and Q. Yao (2010). Nonparametric transfer function models. *Journal of Econometrics* 157, 151–164.

Lucas, R.E.Jr. (1980). Two illustrations of the quantity theory of money. *American Economic Review* 70, 1005–1014.

Sargent, T., and P. Surico (2011). Two illustrations of the quantity theory of money: Breakdowns and revivals. *American Economic Review* 101, 109–128.

Stock, J., and M. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking* 39, 3–33.

Stock, J., and M. Watson (2009). Phillips curve inflation forecasts. In Fuhrer J., Kodrzycki Y., Little J. and Olivei G. (eds) *Understanding Inflation and the Implications for Monetary Policy*. Cambridge, MIT Press, pp. 99–202.

Teles, P., and H. Uhlig (2013). Is quantity theory still alive? ECB working paper 1605.

West, K.D. (1996). Asymptotic inference about predictive ability. *Econometrica* 64, 1067–1084.

Tables and Figures

Table 1: Estimation results of the AR(4) and EWMA-TF(4) models.

Model	AR(4)	EWMA-TF(4)
μ	3.380	-2.573
	(0.596)	(0.373)
ϕ_1	0.634	0.600
	(0.061)	(0.060)
ϕ_2	-0.006	-0.034
	(0.069)	(0.068)
ϕ_3	0.331	0.307
	(0.069)	(0.068)
ϕ_4	-0.177	-0.214
	(0.060)	(0.060)
α		0.035
		(0.010)
σ_ϵ^2	4.392	4.227
	(0.384)	(0.369)
Log-likelihood	-565.616	-560.605
BIC	582.239	579.999
Ljung-Box test(8)	0.040	0.039
Ljung-Box test(12)	0.116	0.144
McLeod-Li test(8)	0.163	0.286
McLeod-Li test(12)	0.451	0.622

Notes: The sample period is 1947:Q1–2013:Q4. The standard errors of the estimated coefficients are given in parentheses. The p-values of the Ljung-Box and McLeod-Li tests (with the number of lags used in parentheses) for the null hypotheses of no remaining autocorrelation and conditional heteroskedasticity in the residuals are reported, respectively.

Table 2: MSFEs of the out-of-sample forecasts for the full out-of-sample period 1970:Q1–2013:Q4.

Model	Forecast horizon		
	$h = 4$	$h = 8$	$h = 12$
MSFE AR	5.125	6.497	6.965
MSFE EWMA-TF	3.926	4.035	3.672
Relative MSFE	0.766*	0.621	0.527
Fraction	0.618***	0.679***	0.709***
	Benchmark forecasts		
MSFE random walk	5.097	6.095	7.012
MSFE naïve	4.146	5.156	5.628

Notes: The MSFEs are computed for average inflation. The expanding window approach is employed, and the lag lengths in the AR and EWMA-TF models are at each step selected using the BIC information criterion. The relative MSFEs are obtained as ratios between the MSFEs of the EWMA-TF model and the AR model. Fraction is the percentage share of forecasts with the EWMA-TF model yielding a smaller squared forecast error than the AR model at a given forecast horizon. The statistical significance of the differences in the mean squared errors and fractions is tested using the test of Diebold and Mariano (1995) and West (1996). In the table, *, **, and *** denote the rejection of the null hypothesis of equal predictive performance at 10%, 5% and 1% significance levels, respectively. In the bottom panel, the MSFEs of the random walk and naïve forecasts of Atkeson and Ohanian (2001) are reported.

Table 3: MSFEs of the out-of-sample forecasts for the period 1990:Q1–2013:Q4.

Model	Forecast horizon		
	$h = 4$	$h = 8$	$h = 12$
MSFE AR	3.026	1.833	1.531
MSFE EWMA-TF	2.399	1.135	0.845
Relative MSFE	0.793*	0.619**	0.552***
Fraction	0.635**	0.682***	0.753***
	Benchmark forecasts		
MSFE random walk	5.265	4.482	4.379
MSFE naïve	2.559	1.934	1.706

Notes: See the notes to Table 2.

Table 4: Out-of-sample forecasting results based on the rolling window approach.

Model	Forecast horizon		
	$h = 4$	$h = 8$	$h = 12$
Period 1970:Q1–2013:Q4			
MSFE AR	3.636	4.958	5.958
MSFE EWMA-TF	3.273	3.855	4.212
Relative MSFE	0.900**	0.778*	0.707*
Fraction	0.606***	0.655***	0.618***
Period 1990:Q1–2013:Q4			
MSFE AR	2.007	1.645	1.771
MSFE EWMA-TF	1.859	1.401	1.399
Relative MSFE	0.926*	0.852	0.790
Fraction	0.659***	0.682***	0.600*

Notes: See the notes to Table 2.

Table 5: Out-of-sample forecasting results for the GDP deflator.

Model	Forecast horizon		
	$h = 4$	$h = 8$	$h = 12$
Period 1970:Q1–2013:Q4			
MSFE AR	2.322	3.313	3.841
MSFE EWMA-TF	1.920	2.362	2.503
Relative MSFE	0.827	0.713	0.652
Fraction	0.588**	0.600**	0.618***
MSFE random walk	1.348	2.031	2.455
MSFE naïve	1.526	2.097	2.356
Period 1990:Q1–2013:Q4			
MSFE AR	0.572	0.846	1.091
MSFE EWMA-TF	0.523	0.672	0.784
Relative MSFE	0.913	0.795	0.719
Fraction	0.624**	0.624**	0.682***
MSFE random walk	0.553	0.665	0.752
MSFE naïve	0.447	0.550	0.614

Notes: See the notes to Table 2.

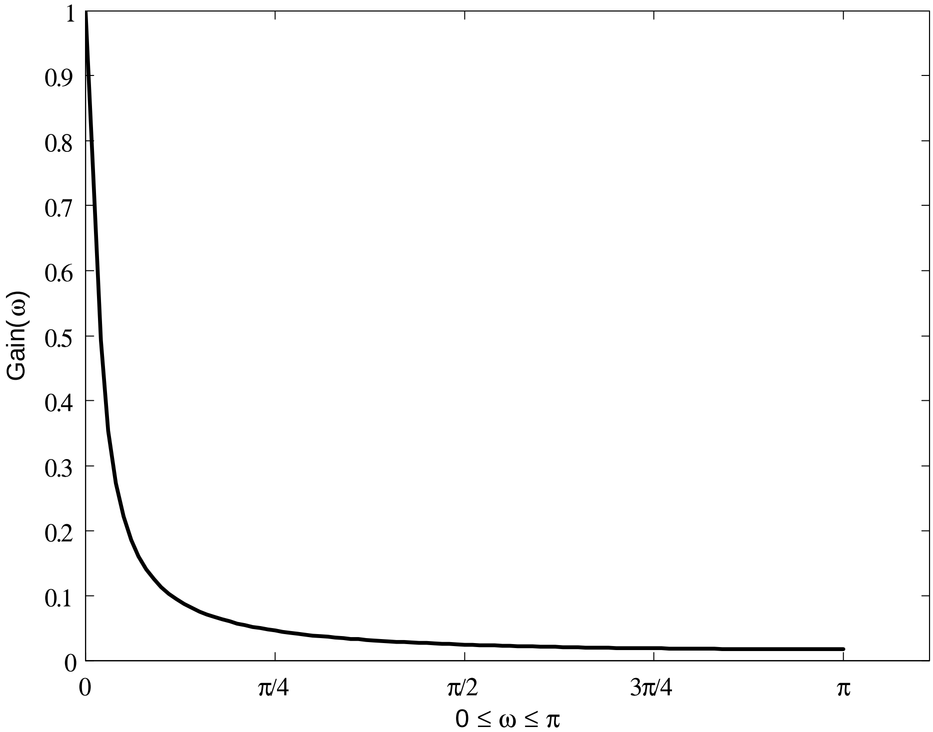

Figure 1: The gain function of the EWMA filter (2) with $\alpha = 0.035$.

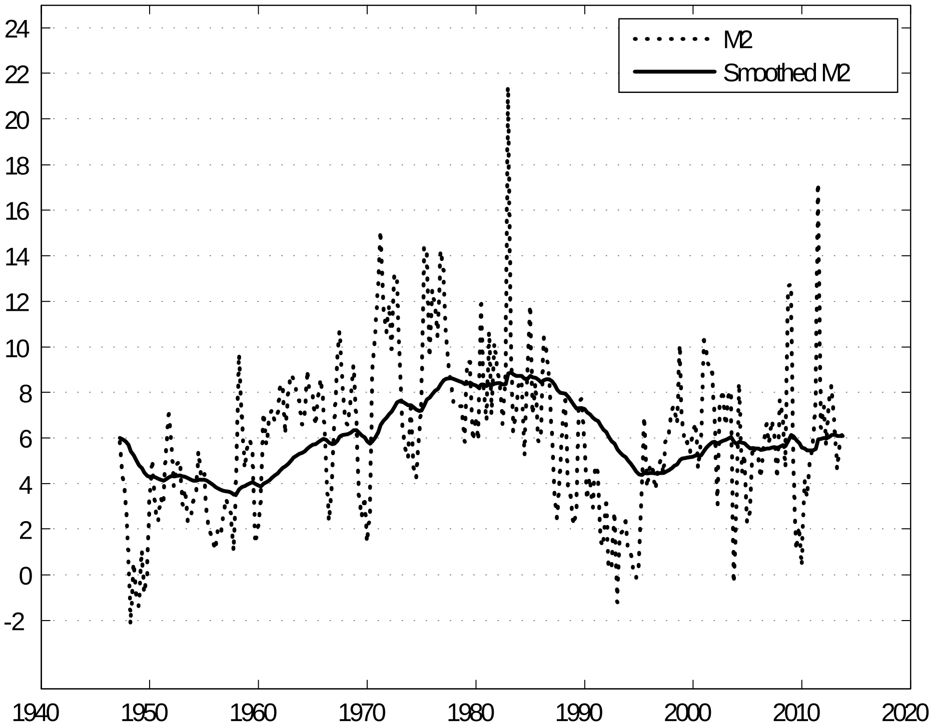

Figure 2: M2 growth rate (m_t, dotted line) and the smoothed M2 predictor (x_t, solid line) when $\alpha = 0.035$ (see Table 1).

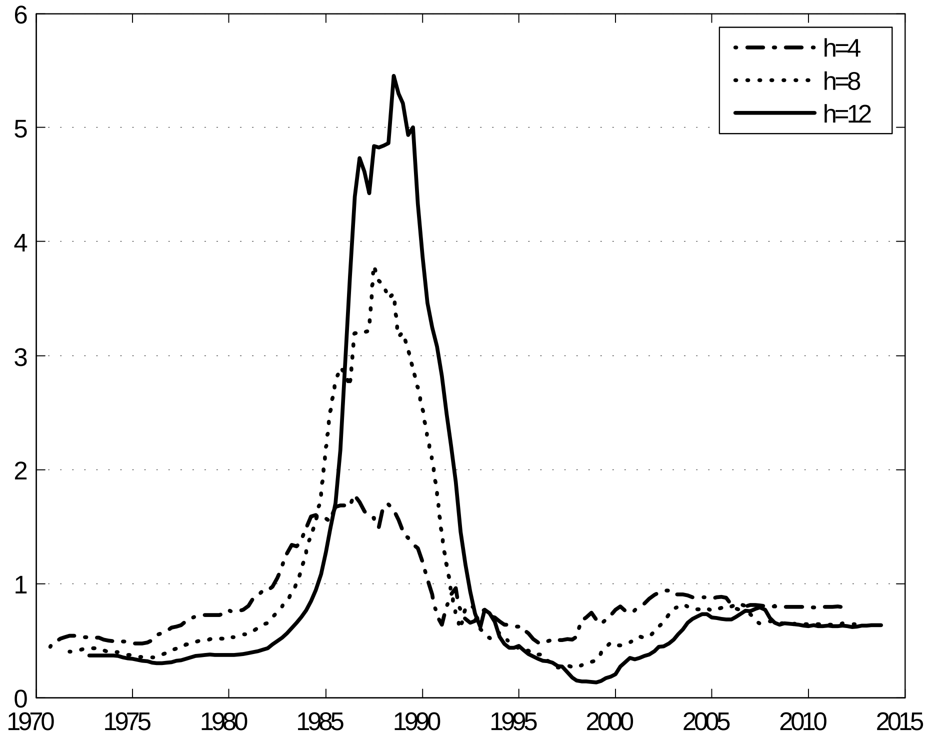

Figure 3: Relative predictability: The smoothed MSFE of the EWMA-TF model relative to that of the AR model. Values less than unity indicate that the EWMA-TF model outperforms the AR model. The forecast horizon is four ($h = 4$), eight ($h = 8$) and 12 ($h = 12$) quarters.

Research Papers 2013

2014-10: Marcelo Fernandes and Cristina M. Scherrer: Price discovery in dual-class shares across multiple markets

2014-11: Yukai Yang: Testing Constancy of the Error Covariance Matrix in Vector Models against Parametric Alternatives using a Spectral Decomposition

2014-12: Stefano Grassi, Nima Nonejad and Paolo Santucci de Magistris: Forecasting with the Standardized Self-Perturbed Kalman Filter

2014-13: Hossein Asgharian, Charlotte Christiansen and Ai Jun Hou: Macro-Finance Determinants of the Long-Run Stock-Bond Correlation: The DCC-MIDAS Specification

2014-14: Mikko S. Pakkanen and Anthony Réveillac: Functional limit theorems for generalized variations of the fractional Brownian sheet

2014-15: Federico Carlini and Katarzyna Łasak: On an Estimation Method for an Alternative Fractionally Cointegrated Model

2014-16: Mogens Bladt, Samuel Finch and Michael Sørensen: Simulation of multivariate diffusion bridges

2014-17: Markku Lanne and Henri Nyberg: Generalized Forecast Error Variance Decomposition for Linear and Nonlinear Multivariate Models

2014-18: Dragan Tevdovski: Extreme negative coexceedances in South Eastern European stock markets

2014-19: Niels Haldrup and Robinson Kruse: Discriminating between fractional integration and spurious long memory

2014-20: Martyna Marczak and Tommaso Proietti: Outlier Detection in Structural Time Series Models: the Indicator Saturation Approach

2014-21: Mikkel Bennedsen, Asger Lunde and Mikko S. Pakkanen: Discretization of Lévy semistationary processes with application to estimation

2014-22: Giuseppe Cavaliere, Morten Ørregaard Nielsen and A.M. Robert Taylor: Bootstrap Score Tests for Fractional Integration in Heteroskedastic ARFIMA Models, with an Application to Price Dynamics in Commodity Spot and Futures Markets

2014-23: Maggie E. C. Jones, Morten Ørregaard Nielsen and Michael Ksawery Popiel: A fractionally cointegrated VAR analysis of economic voting and political support

2014-24: Sepideh Dolatabadim, Morten Ørregaard Nielsen and Ke Xu: A fractionally cointegrated VAR analysis of price discovery in commodity futures markets

2014-25: Matias D. Cattaneo and Michael Jansson: Bootstrapping Kernel-Based Semiparametric Estimators

2014-26: Markku Lanne, Jani Luoto and Henri Nyberg: Is the Quantity Theory of Money Useful in Forecasting U.S. Inflation?