

The Macroeconomic and Fiscal Implications of Inflation Forecast Errors

Harris Dellas
University of Bern
Bank of Greece

Heather D. Gibson
Bank of Greece

Stephen G. Hall
University of Leicester
Bank of Greece

George S. Tavlas^{*}
Bank of Greece
University of Leicester

August 2017

ABSTRACT

The accuracy of inflation forecasts has important implications for macroeconomic stability and real interest rates in economies with nominal rigidities. Erroneous forecasts destabilize output, undermine the conduct of monetary policy under inflation targeting and affect the cost of both short and long-term government borrowing. We propose a new method for forecasting inflation that combines individual forecasts using time-varying-coefficient estimation along with an alternative method based on neural nets. Its application to forecast data from the US and the euro area produces superior performance relative to the standard practice of using individual or linear combinations of individual forecasts, especially during periods marked by structural changes.

^{*} Corresponding author: George Tavlas, Bank of Greece, 21 E Venizelos Ave, Athens, 10250, Greece, Tel. no. +30 210 320 2370; Fax. no. +30 210 320 2432; Email address: gtavlas@bankofgreece.gr.

1. Introduction

Forecasts of inflation play a key role in at least three areas of macroeconomics -- the determination of economic activity, the formation of asset prices, and the conduct of monetary policy. Specifically, inflation forecasts form the basis of price-wage setting so that forecast errors lead to a Phillips-curve relationship. They also matter for nominal asset prices and rates of return when agents are risk averse; the volatility of inflation forecasts impacts on the inflation risk-free rate while the cyclical nature of forecast errors affects inflation risk premia. These effects on real interest rates matter for savings and investment decisions and affect, among other things, the cost of short-term and long-term government borrowing and, thus, the level and dynamics of public debt. Finally, inflation forecasts provide a key input for monetary policy under inflation targeting; forecast errors by the central bank can lead to sub-optimal policies, creating inefficient variations in aggregate economic activity and inflation.

In this paper, we use a standard model to illustrate -- and to offer a rough quantification of -- the aforementioned effects of inflation forecasts on macroeconomic stability and asset prices. Having demonstrated the mechanisms through which forecast errors work in the macroeconomy, we then propose a new empirical strategy for improving the accuracy of inflation forecasts. We use the forecasting results from this methodology to quantify the potential importance of forecast errors on aggregate quantities, prices and welfare.

It has long been recognized that it is hard to forecast inflation, especially during periods marked by high inflation variability and/or structural change.¹ To appreciate this point, consider the following example. The European Central Bank (ECB), whose main objective is to achieve a (headline) inflation rate of close to, but below, 2 percent in the medium term, announces inflation projections on a quarterly basis. During the 12 quarters from 2012:Q1 through 2014:Q4, the ECB's 2-year ahead (that is, for the period 2014:Q1 through 2016:Q4) inflation forecast showed an average projected inflation rate of 1.5 percent, close to the ECB's objective.² The actual outcome, however, was an average inflation rate of 0.2 percent, far below the ECB's objective. Clearly, had the ECB access to more accurate inflation forecasts, its monetary policy stance would have been quite different from that which, in fact, prevailed.

The empirical literature has utilized numerous methods to forecast inflation.³ In what follows, we focus on two branches of the recent literature -- that which focuses on single-equation models and allows these models to incorporate structural breaks (typically, a

¹ Within the context of the late-1970s and early-1980s, a period marked by high inflation variability, Tobin (1981, p. 391) observed: "We have not done well in modeling the inflation process." More recently, González-Rivera (2013, p. 185) noted: "In fact, inflation rates are notoriously difficult to predict."

² The inflation projections varied within the narrow range of 1.3 percent to 1.6 percent.

³ See Stock and Watson (2008) for a comprehensive evaluation of several prototype inflation forecasting models that pay special attention to Phillips-curve inflation forecasts; see, also, Faust and Wright (2013).

single shift in the parameters), and that which emphasizes the linear combination of forecasts based on the forecasts of others.

Rolling windows and structural breaks. The aim of the forecasting literature on structural breaks is to identify periods of instability in the data.⁴ To do so, researchers often use fixed, rolling windows -- for example, 20 quarterly observations -- that move through the data, re-estimating the parameters of the model at each window. The intuition here is that, if a break has occurred in the data, forecasts that use only pre-break data will provide inferior forecasts than forecasts that utilize post-break data. A problem that is encountered is the following: if the window is too short, the parameter estimates will be inefficient; however, if the window is too long so that it incorporates a large data sample before the break, the estimates will be inconsistent. Therefore, researchers have aimed to use small amounts of data prior to the break to minimize the loss of efficiency. The balance between inefficiency and inferior forecast performance has been the subject of considerable discussion about the optimal forecast window.

Forecast combinations. Beginning with the work of Granger (1969), who argued that forecast accuracy can be improved by using a covariance method of combining forecasts rather than using any individual forecast, and the work of Granger and Ramanathan (1984), who proposed a regression-based methodology for combining forecasts, a large literature has established the overall superiority of combining linear forecasts.⁵ The intuition underlying this finding is straightforward: combining forecasts achieves diversification gains -- combinations of forecasts based on, say, different information sets, pools together

⁴ See Rossie (2013b) for a survey of this literature.

⁵ For literature reviews, see Timmermann (2006) and Elliott and Timmermann (2008).

different sources of information and, therefore, should result in lower expected loss (for example, lower expected error variance). Also, given it is likely that some models will be misspecified over certain periods of time -- for example, some models may adjust faster to regime changes than other models -- combining forecasts from a number of models may offer some insurance against “breaks” or other unknown sources of misspecification (Elliott and Timmermann, 2008, p. 42; González-Rivera, 2013, Chapter 9).

In what follows, we aim to extend both of the above strands of the literature. We argue that combinations of forecasts themselves are likely to be subject to structural breaks, and we propose extending forecast combinations to a *nonlinear* context; In particular, we show that combinations of forecasts based on weights derived using time-varying-coefficient (TVC) forecasts produce superior forecasts than combinations derived using a linear framework. As an alternative nonlinear method we also propose using a simple neural net as a method of combining forecasts. We combine forecasts of inflation for (1) the euro area and (2) the United States based on TVC and compare those combinations with combinations generated with linearly-derived weights. The results indicate that forecast accuracy is generally improved using the nonlinear combinations, further that the TVC model works better when there is signs of structural instability while the neural net seems to work better when the system is a stable nonlinear one.

The reason underlying the improvement in performance is that TVC can accommodate structural change, both in the actual inflation process and the forecasting process (that is, forecasters may change the *way* they forecast). The former arises from changes in the macroeconomic environment (monetary policy conduct, changes in the transmission mechanism, globalization, etc.); the latter arises from changes in the identity of the

individual forecasters, within the combination as well as their forecasting tools. As shown below, the weights of the forecast combinations in our sample exhibit significant instability, suggesting the presence of either one or both of these sources of instability.

The remainder of the paper consists of five sections and an Appendix. Section 2 provides a theoretical model that shows that inflation forecast errors lead to inefficient variation in macroeconomic activity, with the welfare losses rising with the inaccuracy of inflation predictions. Section 3 describes our data and the models used to forecast inflation. Section 4 provides the empirical forecasting results. Based on our theoretical model, that section uses the empirical forecasting results to quantify the effects of the TVC procedure on (i) the standard deviation of business-cycle variation, (ii) welfare, in terms of steady-state consumption, and (iii) risk premia. Section 5 presents conclusions. The Appendix presents a detailed description of the model.

2. The Model

In this section we employ a workhorse macroeconomic model in order to demonstrate the effects of forecast errors on economic activity, asset prices (real interest rates) and welfare. The model is standard (see, for instance, related versions in the textbooks by Walsh, 2010, and Wickens, 2015).

Without any loss of generality, we abstract from capital as well as from shocks other than a shock to the supply of money;⁶ the supply of money is assumed to evolve exogenously in a stochastic fashion. Money matters for real economic activity due to existence of

⁶ It is not essential to include TFP shocks in our model as our focus is on the variation in output relative to its natural level.

nominal wage rigidities. We assume that the nominal wage for period t is set in period $t-1$, before the realization of the money supply shock. Furthermore, following standard practice, we assume that the nominal wage is set at a level that is *expected* to produce a wage that would equate labor demand and labor supply in period t (the flexible price real wage).⁷ Under these assumptions, the model leads to the Lucas supply curve:

$$y(t) - y_{(t)}^* = \frac{\alpha}{1-\alpha} (\pi(t) - E_{t-1}\pi(t)) \quad (1)$$

where $y(t)$, $y_{(t)}^*$, $\pi(t)$ and $E_{t-1}\pi(t)$ are the logarithm of the actual level of output, the logarithm of the natural level of output, the actual rate of inflation and the expected – as of period $t-1$ – inflation rate, respectively. Thus, inflation forecast errors lead to inefficient variation in macroeconomic activity.

How large is this inefficient variation? And how much does it reduce economic welfare?

From (1) we have that $\sigma(y) = ((\frac{\alpha}{1-\alpha})^2 RMSE(\pi(t)))$, where $\sigma(y)$ is the standard deviation of output, α is the share of labor in the production function and RMSE is the Root Mean Square Error. In Section 5 we utilize our results from the empirical section to quantify the effect of the improvement in forecasting accuracy arising from the TVC method on the size of inefficient business cycles.

2.1 Fiscal Implications of Inflation Forecast Errors

An underappreciated dimension of inflation unpredictability concerns asset pricing. For assets with a fixed nominal rate of return (such as government bonds), unexpected

⁷ That is, the contract wage, W^c satisfies $E_{t-1} \frac{W(t)^c}{p(t)} = Z_{(t)}^*$ where $Z_{(t)}^*$ is the flexible price real wage, $p(t)$ is the price level and E_{t-1} is the expectation formed in period $t-1$.

movements in inflation (inflation forecast errors) generate movements in the realized rate of return. Positive surprises in inflation depress, and negative surprises boost, realized returns on such assets. These realized gains and losses are priced by risk averse investors into the *ex ante* rates of return demanded. In general, the ex-ante real rate of interest on a nominal asset can be decomposed into a risk free (or inflation indexed) rate and a risk premium (Canzoneri and Dellas (1998)). Forecast errors matter for both components: by influencing the stability of consumption (through the Phillips curve), they affect the inflation indexed rate; and by influencing the correlation between the price level and consumption (in the CCAPM, or the market portfolio return in the CAPM), they affect the risk premium. Consequently, inflation forecast errors have implications, not only for macroeconomic stability, but also for the average cost of borrowing of the private and public sectors. We now elaborate on how the cost of borrowing by the government is affected by the quality of inflation forecasting by market participants.

The standard Euler equation takes the form:

$$\frac{c_t^{-\gamma}}{P_t} = \beta I_t E_t \frac{c_{t+1}^{-\gamma}}{P_{t+1}} \quad (2)$$

where c is consumption, I is the nominal interest rate on a one period nominal bond, P is the price level and γ is the coefficient of relative risk aversion.

Let us consider the benchmark case in which inflation forecast errors do not matter for asset pricing because the nominal rate is indexed to inflation. In this case, the Euler equation takes the form:

$$C_t^{-\gamma} = \beta r_t E_t C_{t+1}^{-\gamma} \quad (3)$$

where r_t is the real rate of return on a one period, inflation indexed bond.

The *ex ante* real rate of return, R_t , on a one period bond that carries a fixed nominal interest rate, I_t is given by

$$R_t = I_t E_t \frac{P_t}{P_{t+1}} \quad (4)$$

Combining the last two equations with equation 2 gives

$$R_t^{-1} = \frac{\beta E_t \frac{C_{t+1}^{-\gamma} P_t}{C_t^{-\gamma} P_{t+1}}}{E_t \frac{P_t}{P_{t+1}}} = r_t^{-1} + \frac{\beta Cov_t \left(\frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}, \frac{P_t}{P_{t+1}} \right)}{E_t \frac{P_t}{P_{t+1}}} \quad (5)$$

where the last fraction in (5) represents the risk premium on nominal assets when inflation cannot be perfectly forecasted.

The sign of the risk premium depends on the covariance between prices and consumption. A positive co-movement between inflation and consumption means that unexpectedly high inflation (and, thus, capital losses) on holdings of nominal assets occur when consumption is unexpectedly high (a movement along the short-run Phillips curve). In this case, government bonds provide insurance to investors and carry a negative inflation premium. The opposite is true for productivity shocks that are not accommodated by monetary policy. Such shocks induce a negative covariance between inflation and output. Bond holders incur a capital loss (inflation is higher than expected) when output and consumption are low. As a result, they will demand a positive inflation risk premium when they face this type of shocks.

2.2 Central Bank Inflation Forecasts and the Conduct of Monetary Policy

We have argued above that the accuracy of central-bank forecasts of inflation is important for the efficient conduct of policy, especially for an inflation targeting central bank. Forecast errors will lead to inappropriate policy choices, destabilizing output and inflation relative to target. A simple way to demonstrate this point in the context of our model is to endow the central bank with an inflation (price targeting) objective, \bar{p} , and to assume that the central bank sets its monetary policy instrument (M) based on its forecast of inflation (which depends on the forecast of the productivity shock, A_t). In this case, $E_{t-1}p_t = \bar{p}$ in the Lucas supply curve (equation 1), and to the extent that p_t differs from \bar{p} due to the inappropriate choice of M_t (which is due to the inaccurate forecast of A_t , by the central bank), a Lucas-supply-curve relationship emerges solely on the basis of central bank forecast errors.

3. The Data and Empirical Methodology

In recent years a considerable literature has focused on the issue of structural breaks in economic forecasting.⁸ It is obvious that using pre-break data can lead to biased parameters and, therefore, poor economic forecasts (Clark and McCracken (2009)). An important recent research theme has, therefore, been to use rolling windows to help overcome this problem, the idea being that the correct window length will effectively write-off most of the earlier (pre-break) regime. The choice of optimal window size, however, is difficult since different window sizes may give very different results (Inoue and Rossi, 2012). Pessaran and Timmerman (2007) suggest that a trade-off exists between forecast bias and efficiency, so that use of only a post-break window may lead to a loss in

⁸ See, for example, Clements and Hendry (1996, 1998), Hendry (2000), Pesaran and Timmerman (2002), Goyal and Welch (2003), Pessaran, Pettenuzzo and Timmerman (2006), Koop and Potter (2007), Castle Clements and Hendry (2013) and Rossie (2013a, b).

efficiency due to the potentially small number of observations; those authors suggest that using some data from the pre-break regime may be desirable to ameliorate this problem. Inoue, Jin and Rossi (2017) have proposed a feasible solution to the choice of the window size based on approximating the forecasters' quadratic loss function. This literature, however, rests on the assumption that the breaks are relatively few (often only one break) and discrete, while, in fact, breaks may be endemic and continuous.

As mentioned, a separate branch of the literature on economic forecasting posits that combinations of individual forecasts typically outperform any of the individual forecasts that comprise the combinations (Bates and Granger (1969), Granger and Ramanathan (1984) and Elliott and Timmermann (2008)). A related finding -- both empirically and theoretically -- is that simple forecast averaging (equal weights on each forecaster) often outperforms more complex combination techniques (Timmermann (2006)). The existing literature, however, has, by-and-large, focused on forecast combinations within a linear framework. Moreover, it has not taken account of the above-mentioned literature on structural breaks. In what follows, we combine the above two branches of the literature. Specifically, we argue that the optimal forecast combination may itself be subject to structural breaks. Our procedure for dealing with structural breaks is the following.

1. We will not use fixed-length rolling windows since we posit that the optimal weights change frequently and often in a relatively-smooth way. Instead, we use a random walk time-varying-coefficient model since this is a direct generalization of the standard OLS fixed window method. Specifically, when the variances of the random coefficients are zero, the filtered estimates give recursive OLS, and the smoothed estimates give full-sample OLS estimates. Once the variances are non-

zero, the filtered estimates become a backward exponentially weighted average of past regimes, where the weights decline at a rate governed by the size of the variances (that is, the larger the variances, the faster the weights decline to zero and, correspondingly, the less weight assigned to history). This procedure seems intuitively more appealing than a fixed-window length, which weighs the past equally up to its fixed length and then assigns zero to all observations beyond this length⁹.

2. The above interpretation is agnostic as to the cause of the structural breaks. An alternative view, however, is that structural breaks are a manifestation of an incorrect functional form and that a non-linear function, in fact, exists that is structurally stable. We do not, of course, know what that functional form might be, but any unknown functional form can be represented as a linear model with time-varying coefficients; this important theorem was first established by Swamy and Mehta (1975), and subsequently confirmed by Granger (2008). Thus, we allow for nonlinear combinations of forecasts using TVCs. In this same vein we will also propose a forecast combination based on a simple neural net as a way of capturing unknown stable nonlinear structures.

Our conjecture is that TVC¹⁰-based combinations will produce forecasts that are at least as good as linear combinations since the TVC technique always has OLS nested within it.

⁹ This is akin to the move from the fixed length ARCH models to the more general GARCH models which give an infinite-order exponentially declining set of weights to the past squared errors.

¹⁰ It would, of course, be possible to apply other nonlinear estimation techniques, such as neural nets or non-parametric estimators to forecast combinations. We have not done this as we want to stay within a close generalisation of the rolling window estimators, described above. Doing so

If the variables under examination have been subjected to structural change, we would expect a considerable improvement. Similarly we would expect a neural net to also perform as well or better than as a linear combination.

To examine the foregoing conjecture, we use quarterly forecasts of headline inflation for both the euro area and the United States. For the euro area, our data set covers the quarterly period 2000:Q1 through 2016:Q4 and is comprised of predictions made by forecasters included in the ECB's quarterly report, *Survey of Professional Forecasters*. The *Survey* provides one input in the ECB's assessment of the future evolution of inflation.¹¹ The *Survey*, however, does not contain a complete set of predictions for each of the forecasters during the entire estimation period.¹² There are gaps in the sets of predictions contained in the successive issues of the *Survey* reflecting the fact that individual participating forecasters did not consistently provide responses, while, over time, some forecasters moved in, and others moved out of, the *Survey*. For the U.S., data on forecasts of inflation are taken from the Federal Reserve Bank of Philadelphia, *Survey of Professional Forecasters*. Our sample period for the U.S. runs from 1993:Q2 through 2016:Q3. The forecasts are available from some 583 individual forecasters.

would also open up a difficult area of forming comparisons between very different estimation techniques, all of which can fit a data set perfectly with enough training.

¹¹ The ECB does not regard inflation forecasts as providing full information on the economy and the future path of inflation. The ECB also uses other sources of information, including on money and credit growth, to determine the path of inflation. Thus, the ECB does not follow a strict inflation-targeting framework (Issing, 2008, pp. 91-92).

¹² The *Survey* presently includes predictions made by about 40 forecasters. The number of forecasters has not been constant; it has been rising over time.

Our aim was to provide a reasonably complete set of forecasts made by individual forecasters over the estimation period. For the case of the euro area, this procedure yielded predictions made by six forecasters; for the United States we constructed a data set comprised of four forecasters. Even the forecasts provided by these forecasters did not yield an entirely complete set of forecasts as there were individual quarters of missing observations for some forecasters.¹³ In order to deal with these missing observations we interpolated the data set using a cubic spline interpolation method. For the euro area, the specific inflation forecasts are those produced in the corresponding quarter of the previous year -- that is, the forecasts are four-quarters ahead. For the U.S., the specific forecasts are one-quarter ahead. The difference in forecast horizons is due to the fact that the euro-area *Survey of Professional Forecasters* provides forecasts beginning with four-quarters ahead while for the U.S. we had access to forecasters who provided one-quarter ahead forecasts.

In what follows, we describe our empirical methodology in two steps. First, we provide a description of the combination methods used. Second, we then provide a general description of our proposed methodology.

3.1 The Empirical Models

¹³ For the six forecasters chosen for the euro area, there are nine missing observations out of a total of 408 observations; most of these were for periods for which the forecasts were the same both before and after the missing observation. For the four U.S. forecasters, there are 45 missing observations out of 272 observations; most of the missing observations occurred early in the sample (when inflation was relatively stable).

We have n forecasters who each produce a forecast for variable y at a particular forecast horizon $t+h$ at time t , Thus, $\hat{y}_{t+h,t}^i$ $i = 1 \dots n$. The following estimation methods are used to generate forecast combinations.

- (1) *Simple averages*. We take the average of the forecasts of each of the forecasters - six in the case of euro-area inflation and four in the case of United States inflation.

That is $\hat{y}_{\alpha t+h,t} = \frac{\sum_{i=1}^n \hat{y}_{t+h,t}^i}{n}$, where α is the average of the forecasters, h is the

number of periods ahead being forecast, and n is the number of forecasters.

- (2) *OLS*. We run an OLS regression of the actual inflation rate (the dependent variable) on the forecast of inflation -- again, six for the euro area and four for the United States -- for the entire sample period. This procedure provides a set of OLS weights,

$\hat{y}_{OLS t+h,t} = \sum_{i=1}^n \beta_i \hat{y}_{t+h,t}^i$ where β_i are the weights given by full-sample OLS

estimation. Since this method incorporates future values of variables to generate forecasts, it is not available as a real time forecast; however, we retain this method as it has often been used as a combination technique.

- (3) *Recursive Least Squares (RLS)*. To produce a feasible real-time forecast, we use RLS, beginning with an initial sample of data and re-estimating the coefficients period-by-period as well as adding additional observations until the sample is exhausted:

$\hat{y}_{RLS t+h,t} = \sum_{i=1}^n \chi_{it} \hat{y}_{t+h,t}^i$ where χ_{it} are the recursive weights on each forecaster

for each period.

- (4) *TVC*. To capture structural breaks and/or nonlinearities, we use two TVC models.

(i) One model is what we call TVC1, under which the TVCs are determined by a random walk. This is given by the following state space form using the Kalman filter formulation:

$$y_{t+h,t} = \gamma_{0t} + \gamma_{1t}\hat{y}_{t+h,t}^1 + \dots + \gamma_{nt}\hat{y}_{t+h,t}^n + V_t \quad (6)$$

which is the measurement equation for the state space form. There are a number of possibilities for the structure of the state equations. The simplest is the following:

$$\gamma_{it} = \gamma_{i,t-1} + \varepsilon_{it} \quad i = 0, \dots, n \quad (7)$$

that is, the parameters are specified as a simple random walk, where $\varepsilon_i \sim N(0, \sigma_i^2)$; if all $\sigma_i^2 = 0$, then the filtered parameters will yield recursive OLS estimates of the weights and the smoothed estimates will be constant at the full sample OLS values. If $\sigma_i^2 > 0$ the filtered parameters will be an exponentially weighted average of the past parameters; the larger the variance the more rapid will be the rate of decay. The smoothed parameters will then be a two sided exponentially weighted average involving both the past and the future values. This specification allows the parameters to change in an unrestricted way; however, the forecast for the future values of the parameters will always be a constant value at last period's estimate.

ii) The above discussion motivates our second TVC model, TVC2. If the parameters have been changing in a steady or trend-like way, we would want to incorporate this trend behavior in our forecast. This can be done in the following way:

$$\begin{aligned}\gamma_{it} &= \gamma_{it-1} + \gamma_{k+1+it-1} + \varepsilon_{it} \\ \gamma_{k+1+it} &= \gamma_{k+1+it-1} + \varepsilon_{k+1+it} \quad i = 0, \dots, k\end{aligned}\tag{8}$$

This procedure puts a local stochastic drift term into each of the parameters, allowing each parameter to continue changing into the future.

In the cases of both TVC1 and TVC2, there is an added complication when $h > 1$ since we have to forecast the parameters h periods ahead. When $h = 1$, the Kalman filter does this automatically before it updates the parameters; otherwise, when $h > 1$, we have to use the standard Kalman filter forecasting equations to derive an h -step ahead forecast of the weights.

(5) *Neural Net*. A neural net is essentially a complex combination of simple nonlinear functions, as such a sufficiently complex net can approximate any unknown nonlinear function with any desired degree of accuracy (White 1992, Swanson and White(1997)). A good introduction may be found in Kuan and White(1994). The simplest feed forward neural net can be represented algebraically as;

$$Y = h\left(\sum_{k=1}^K \alpha_k g\left(\sum_{j=1}^J \beta_{jk} X_j\right)\right)$$

Where X is a vector of observed inputs, exogenous variables in econometric terms, and Y is the output, or endogenous variable. α and β are parameters and h and g are simple nonlinear functions, in practice h and g are usually the same function. In the application here both h and g are the logistic function, $g(u) = h(u) = 1/(1 + e^{-u})$. This is called a feedforward model as X affects Y but Y does not affect X . This is a simple one layer neural

net with J inputs and K neurons. A more complex one would involve further summations over more nonlinear transformations.

From an econometric perspective as the net becomes more complicated the parameters α and β grow in numbers rapidly and are generally not identified or unique. The process of arriving at estimates of these parameters is termed learning and it is generally undertaken by some form of hill climbing technique which attempts to minimize the squared deviations between Y and the forecast from the net. Typical procedures are back propagation, simulated annealing or genetic algorithms. It is generally recognized that a complex net will need a large number of observations to train it effectively. Given the number of observations we have available here we will restrict the net to one hidden layer and four neurons.

In most of the neural network literature, applications divide the sample into a training sub sample and a forecasting sub sample. This does not represent how a neural net would be used in actual forecasting as in a practical setting we would use all available data to train the net and then forecast into the future. For this reason we will implement a recursive neural net, training it first over an initial small sample and then recursively moving through the sample forecasting and re training in a manner analogous to recursive OLS.

It is important to realize that the net is based around the assumption of the existence of a stable nonlinear structure. Of course a sufficiently complex net can approximate a changing structure but this may have to be excessively complex. Essentially the point is that α and β are being trained over the whole sample up to period t and so information from the past is not being written of here as it is with the TVC model above.

Our proposed methodology is as follows;

1. First we must determine if there is structural instability in the weights in question. This can be done in many ways. If we were considering only a unique break point, we could use, for example, the bootstrap procedure of Nankervis and Savin (1996) or the semi-parametric bootstrap of Davidson and MacKinnon (2006). In the case of a single break we could also use formal test procedures such as Bai and Perron (1998, 2003); for multiple breaks we could use a test such as that proposed by Altissimo and Corradi (2003). Single breakpoint tests, however, would seem to be inappropriate if (as we believe) the world is subject to many continuous breaks and gradual structural change. Even the multiple break tests are not well suited to dealing with such a situation. We, therefore, begin with simple recursive estimation of the parameters. We then inspect these parameters with the conventional CUSUM squared test.
2. If there is no sign of instability, we would expect the conventional combination techniques to work well. However, if there are signs of instability in the weights we would expect the TVC based methods to outperform the fixed-weight methods, that is averaging or OLS based combinations. If the recursive estimates of the weights seem to be trending, we would expect TVC2 to outperform TVC1; if not, then TVC1 would likely be the better performer. If there is little signs of instability then we would still expect the neural net to outperform the linear combination methods, although as noted above the neural net may not cope well with instability.

4. Empirical Results

To evaluate forecast accuracy, we use three criteria: (i) root mean square error, which is the square root of the mean of the squared differences between predicted and actual values; (ii) mean absolute error (MAE), which is the mean of the absolute errors (the RMSE and MAE are different because the squaring of the differences used to calculate the RMSE gives larger weights to large errors); and (iii) mean absolute percent error (MAPE), which is the mean of the absolute percent errors -- taking percent errors weighs the errors differently for different inflation rates (for example, a 1 percentage point error is a large percent error when actual inflation is zero, and small when inflation is high).

4.1 Results for the euro area

We begin by considering the recursive estimation of the weights which combine the six euro-area forecasters. Figure 1 shows six sets of recursive coefficients, that is, the weights assigned to each forecaster based on recursive estimation. It is clear from this figure that there is coefficient variation; the weights assigned to the individual forecasters in generating forecast combination exhibit considerable instability. Clearly, there is a fairly discrete break around 2008, at the start of the financial crises but there is also quite considerable instability before that time. Indeed, as shown in Figure 2, the CUSM of squares test is continuously outside of its 95 per cent confidence bands from 2005 to 2009¹⁴. The recursive parameters also trend over the second half of the period, although this is not necessarily an indication that structural change is continuing over this period since these parameters are not unbiased estimates of the underlying optimal weights.

¹⁴ This test is well known to have relatively low power.

Table 1 reports the RMSEs, the MAEs, and the MAPEs results for euro-area inflation based on simple averages of forecast (equal weights), forecasts based on OLS weights, and forecasts based on recursive OLS weights¹⁵. The table also presents the RMSEs, the MAEs, and the MAPEs results for the six individual euro-area forecasters used in this study. As shown in the table, each of the individual forecasters performs similarly. For example, individual-forecaster RMSEs fall within the narrow range of .942 (forecaster 2) to 1.029 (forecaster 1). The equally-weighted combined forecasts show little improvement compared with the individual forecasts. The OLS combination, however, represents a substantial improvement; for example, with OLS the RMSE falls to .839, which is less than all of the individual-forecaster RMSEs. However, the OLS weights are derived on the basis of the entire sample period and so, at each point in time, they use (future) actual inflation to derive the weights for the current period -- information that was not available to the individual forecasters. Finally, based on the RMSE and MAE criteria the recursive weights perform worse than the simple average and than each of the individual forecasters. This result might seem surprising, but it should be recalled that the recursive parameters are not consistent estimates of the underlying time-varying weights; given the degree of structural instability shown in Figure 2, they perform poorly in approximating these optimal weights.

The results using the two TVC models and the neural net are presented in Table 2. The TVC results represent a considerable improvement over both the recursive OLS results and the equally-weighted results. The RMSE and the MAE under the TVC methods are

¹⁵ Since in this case we are forecasting 4 periods ahead when evaluating the RLS weights, we use the weights dated at t-4.

about one-half of those yielded by the other methods, while the MAPE is about one-third of those other methods. Thus, there appears to be considerable forecasting gains in using time-varying weights. The neural net however presents little improvement over recursive least squares and actually does worse than the equal weights. This may at first sight be surprising but it must be remembered that the coefficients in the net are being chosen based on the whole sample up to the period being forecast and so, although this is a nonlinear model, it is forecasting on the basis of fixed parameters which do not allow for any structural change.

TVC2 does not perform as well as the simpler TVC1 method. The intuitive reason underlying the superior performance of TVC1 relative to TVC2 is that the combination weights are relatively stable and mean-reverting, so that projecting them as a level process at their previous value (that is, a pure random walk), as under TVC1, produces a better forecasting result than projecting their recent trend movement into the future (that is, a random walk with drift), as under TVC2. In other words, since the parameters do not trend, it is better to project them as a random walk than as a random walk with draft.

4.2 Results for the United States

As in the case of the euro area, for the United States we begin by considering the recursive combination weights. These are shown in Figure 3. In contrast to the euro area, there is not a strong and obvious break in 2008, when the financial crises hit, but there is substantial variation during the entire period up to 2009. Until that year, the CUSUM of squares are near the 5 per cent boundary, indicating instability; subsequently, the weights are fairly stable. As shown in Figure 4, the CUSUM of squares test does not breach its

confidence interval, although it is very close to the lower bound until around 2010. It would seem fair to conclude that there are signs of instability, but not nearly as strong as in the case of the euro area.

We now compare the inflation performances of the individual forecasters with those derived on the basis of the standard combination methods. The results are displayed in Table 3. Based on the three forecast criteria, the simple average method generally performs better than three of the individual forecasts (the exception is forecaster 2), while the OLS-based combination performs better than all four individual forecasts. However, since the OLS results are based on the entire sample, we again turn to recursively-estimated coefficients. In this case, RLS generally performs better than the simple average, under both the RMSE and the MAE criteria, while the simple average is better than RLS under the MAPE criterion. An individual forecaster (number 2), however, outperforms RLS under each of the forecast criteria. Given the relatively greater degree of stability in the recursive parameters, this result is not unexpected.¹⁶

We now turn to the forecasts based on the nonlinear methods. These results are reported in Table 4. The TVC1 results outperform (by a considerable margin) both recursive OLS, the method based on equal weights, and the method that uses full-sample OLS weights.

¹⁶ If the parameters were perfectly stable, OLS would always perform better than both the individual forecasters and the average of the individual forecasters since OLS could always pick the best individual forecast. Moreover, if the parameters were perfectly stable, RLS would, in the limit, provide the same weights as OLS. In the case above, the parameters are stable, but not perfectly stable; thus, an individual forecaster outperforms RLS.

There is not a great deal of overall difference between the two TVC techniques, although the three measures of forecast accuracy differ somewhat as to which of the two TVC models performs the best. This is probably due to a single period when actual inflation fell to just about zero and hence the mean absolute percent error is dominated by the performance at this single observation (as we express the error as a percentage of near zero, a huge weight is put on that one error).¹⁷ In this case the neural net outperforms both TVC models by a small margin. This reinforces the idea that the TVC methods are particularly useful in the presence of significant structural breaks while the neural net is good at capturing stable non-linear structures as in the case of the US the evidence of structural instability is much weaker.

To sum up, our results for inflation confirm two key findings provided in the previous literature on forecasting, namely, that (1) forecasts based on forecast combinations generally outperform the individual forecasts comprising the combination, and (2) simple forecast averaging (equal weights on each forecaster) sometimes outperforms more complex linear combination techniques such as OLS and recursive estimation. Our results also suggest, however, that a nonlinear technique based on TVC estimation can provide superior inflation forecasts in the presence of structural instability and that in the presence of a stable non-linear relationship the neural net does well. We conjecture that this TVC does particularly well during periods of relatively-high instability and the neural net does better in a stable regime.

4.3. Quantifying the Effects of Forecast Errors

¹⁷ The single observation occurred in 2009q1, at which point inflation was minus 0.1 per cent.

We now use the forecasting results to demonstrate the effects of forecast errors on economic activity, welfare, and asset prices (that is, real interest rates). Table 5 shows how much the improvement in forecasting quality can limit the size of inefficient business cycles. The table utilizes the RMSEs for both the euro area and the United States provided by the methods based on equal weights and TVC1, respectively. It will be recalled that for the euro area the method based on equal weights and TVC1 produced RMSEs of .98 and .49, respectively. For the United States the method of equal weights produced an RMSE of 1.14, while TVC1 produced an RMSE of .89. As reported in Table 5, the improvement in forecast accuracy reduces the standard deviation of the output gap from .014 to .007 for the euro area and from 0.021 to 0.013 for the United States.¹⁸

We can also use the model to produce estimates of the welfare cost of forecasting inaccuracy (that is, inefficient business-cycle variation). In the appendix we show that, using the solution of the model, the utility function of the representative agent, u , can be expressed as a function of inflation forecast errors, namely, $u(\pi(t)/E_{t-1}\pi(t))$. By doing so, we can compute the welfare cost of output instability in terms of the percentage of steady state consumption an agent would be willing to sacrifice in order to avoid such instability.¹⁹ Table 6 presents the results derived on the basis of the information on forecast errors associated with the alternative forecasting methods considered in the empirical section. In particular, starting from a situation where we use the linear combination with equal weights, it reports how much the agents in the economy would be

¹⁸ Typical estimates of the standard deviation of the inefficient output gap (actual minus natural level) range from 1 to 2%. Our results indicate that much of this volatility can be attributed to unanticipated inflation, a finding in line with the Keynesian model and also the imperfect information model of Lucas, 1973.

¹⁹ This is a standard method for measuring the welfare costs of business cycles pioneered by Lucas.

willing to pay, in terms of steady state consumption, in order to prevent a worsening, or to achieve an improvement, in forecasting quality. For instance, agents would be willing to pay three-quarters of a percentage point (0.76) of steady state consumption in order to achieve perfect forecasting. Similarly, agents would be willing to pay one-sixth of one percent (0.16) in order to prevent a deterioration in forecasting quality (in terms of the RMSE) by 20%.²⁰

Finally, we turn to the effect of forecast errors on asset prices. Without knowledge of the *sign* of the conditional covariance term in equation (5), it is not possible to tell whether better forecasting improves (decreases) or worsens (increases real interest rates on government debt) the government budget. Note also from equation (5) that the *size* of the covariance depends on the variance of the forecast errors. In order to illustrate this effect, we utilize information from the Philadelphia FED forecast series on consumption and inflation forecast errors to compute the contribution of forecasting accuracy to the size of the risk premia and real interest rates.

In the data, the covariance is positive (as suggested by the Phillips curve). As greater forecasting accuracy decreases the size of the covariance, an improvement in the quality of inflation forecasts in financial markets is bad news for the treasuries since it increases interest rates. Nonetheless, the quantitative effect is quite small, with a change of 10% in the accuracy of inflation forecasting implying a change in the real interest rate by about 1 basis point.

5. Conclusions

²⁰ These numbers are large and comparable to those found in the literature regarding the cost of business fluctuations due to nominal rigidities.

The ability to provide accurate forecasts of inflation plays a crucial role in macroeconomics. We have established analytically that nonlinear combinations of inflation forecasts are superior to linear combinations. Applying our time-varying-coefficient estimation method to the euro area and the United States confirmed that our technique performs better, especially during periods marked by structural changes.

The improvement in forecasting performance is substantial; by about 50% in the euro area and 20% in the US.²¹ Such improvement would have significant effects on the size of inefficient business cycle variation, welfare and real interest rates. In particular, we find that an improvement in forecasting quality of the size achieved by the use of our TVC forecasting method instead of the standard linear combination of forecasts lowers the standard deviation of inefficient business cycle variation by 0.7 of one percentage point and increases welfare by an amount equivalent to increasing steady state consumption by between one sixth and one third of steady state consumption.

²¹ The improvement of the euro area, which exhibits considerable coefficient variation, is larger than that of the United States, where coefficient variation is smaller. This result is in line with our expectations that time-varying-coefficient methods do well in periods of structural instability.

Table 1: Euro-area: Forecasts based on simple averages OLS and RLS weights

	1	2	3	4	5	6	Equal weights	OLS weights	RLS weights
RMSE	1.029	0.942	1.008	1.019	0.943	0.951	0.982	0.839	1.305
MAE	0.850	0.740	0.827	0.845	0.762	0.767	0.787	0.616	1.089
MAPE	240.1	177.4	258.2	288.2	240.1	216.5	240.7	135.5	153.4

Sample 2002q4-2016q3

Table 2: Euro area: non-linear results

	TVC1	TVC2	Neural Net
RMSE	0.488	0.574	1.275
MAE	0.357	0.460	0.958
MAPE	79.1	119.1	157.6

Sample 2002q4-2016q3. Both TVC models use the Kalman filter forecast equations to predict the parameters four periods ahead.

Table 3: United States: Forecasts based on simple averages and OLS weights

	1	2	3	4	Equal weights	OLS weights	RLS Weights
RMSE	1.291	1.045	1.285	1.588	1.143	0.931	1.110
MAE	0.983	0.837	0.987	1.307	0.947	0.689	0.909
MAPE	271.0	163.7	239.8	261.9	136.6	124.1	170.1

Sample 2002q4-2016q3.

Table 4: United States: nonlinear results

	TVC 1	TVC2	Neural net
RMSE	0.897	0.997	0.832
MAE	0.620	0.703	0.546
MAPE	60.0	50.0	61.88

Sample 2002q4-2016q3. In this case, as we are only forecasting one period ahead, the Kalman filter predicted parameters may be used.

Table 5: Inflation forecast errors and macroeconomic instability

	RMSE EW	RMSE TVC	$\sigma(y)$ EW	$\sigma(y)$ TVC
Euro area	0.98	0.49	0.014	0.007
US	1.14	0.89	0.021	0.013

Note: $\alpha=0.6$. RMSE EW is the forecasting method that uses the equally-weighted forecast combinations, $\sigma(y)$ is the standard deviation of the output gap.

Table 6: Welfare changes in terms of steady state consumption from improving/worsening inflation forecasting quality

Change in forecasting quality	20% worse	50% worse	Improve to perfect forecasting
Consumption change	-0.16%	-0.34%	+0.76

Note: The reference point is the model with linear combination of forecasters with equal weights. The entries have been derived under the following parameterization of the model: alpha (labor share in production) = 0.6, sigma (inverse Frisch elasticity of the labor supply) = 1, gamma (coefficient of relative risk aversion)= 2. The comparison involves TVC1 against the method of equal weights.

Figure 1: Euro area: Recursive OLS combination weights

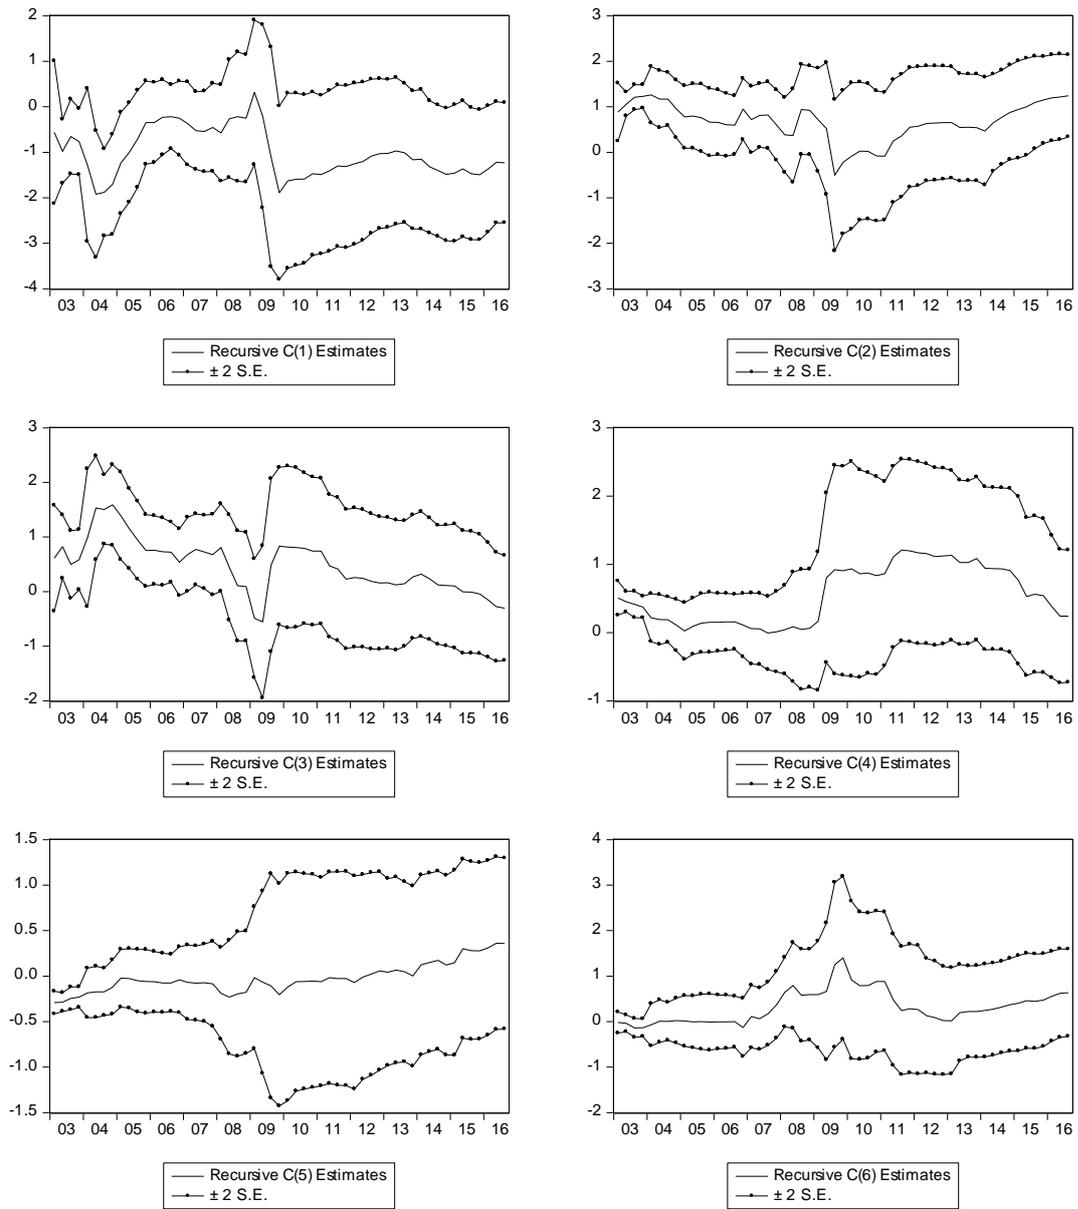


Figure 2: CUSUM Squared test for the euro area

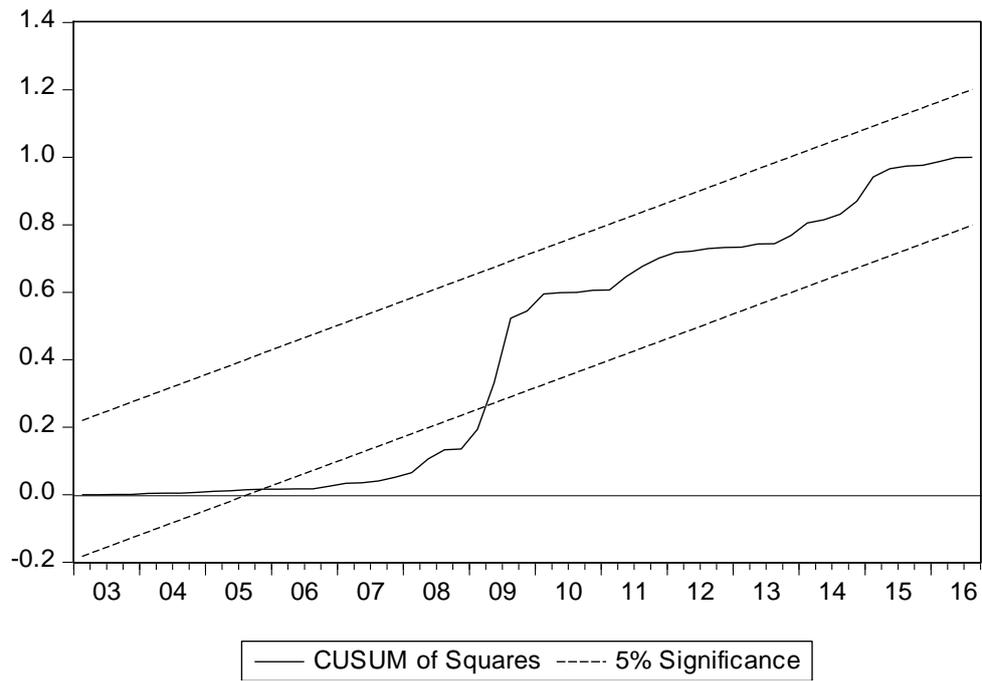


Figure 3: United States: Recursive OLS weights

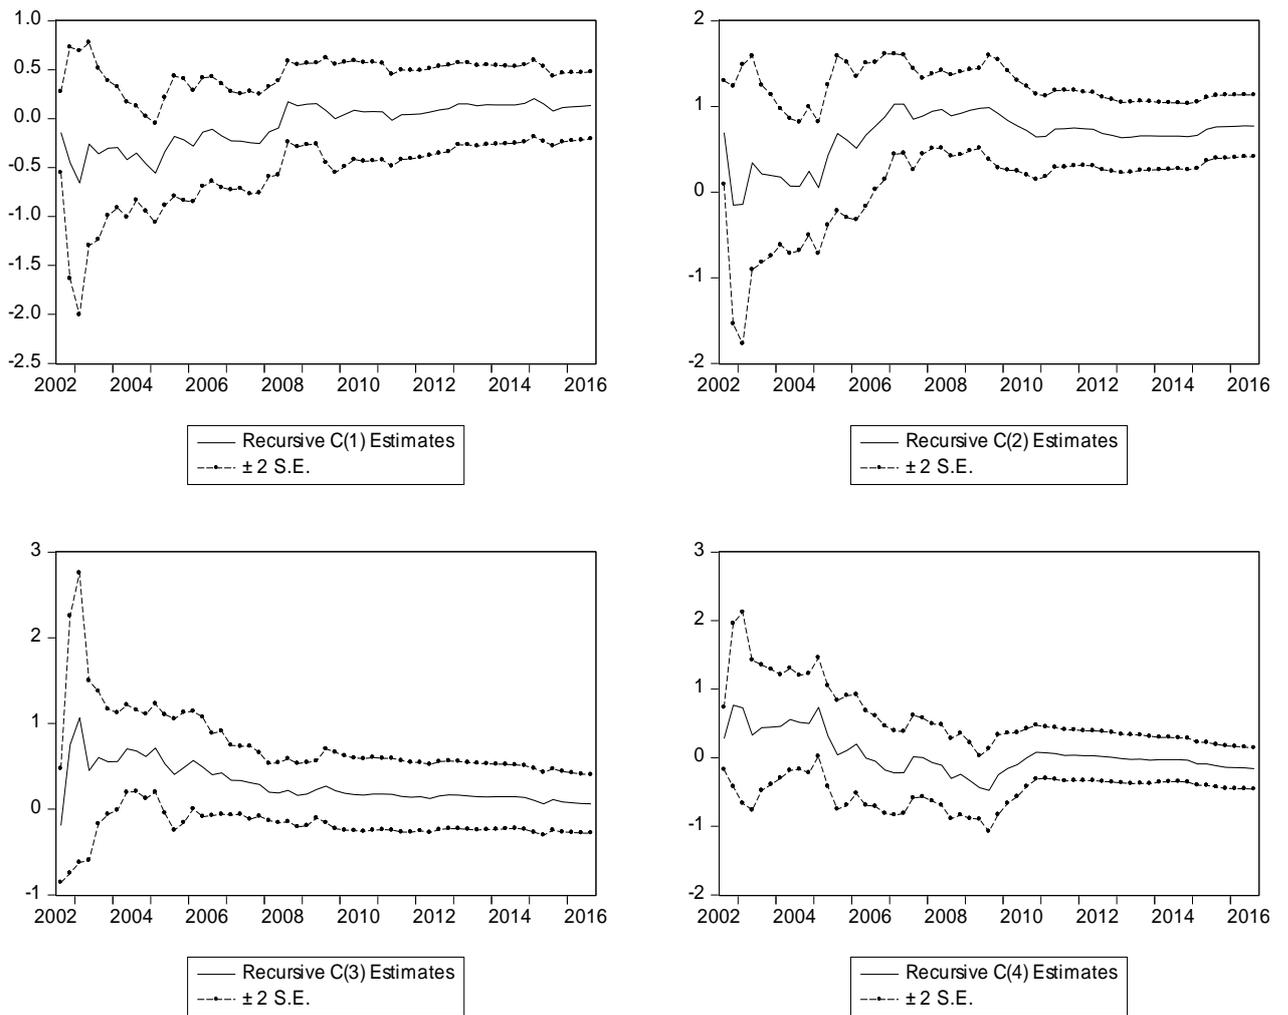
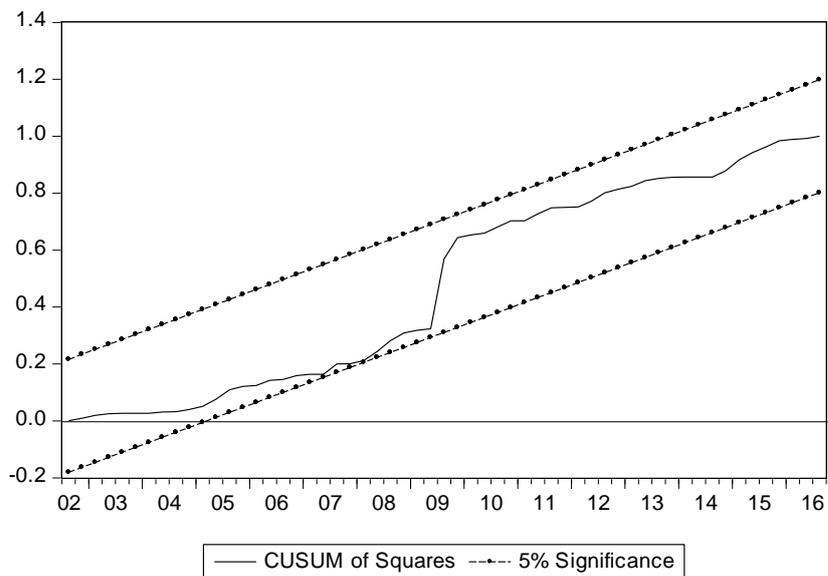


Figure 4: The CUSUM of squares test for the USA



References

Bates John and Granger, Clive, 1969. The Combination of Forecasts. *Operations Research Quarterly* 20(4), 451-468.

Bai J and Perron P, 1998. Estimating and Testing Linear Models with multiple structural Changes. *Econometrica*, 47-78.

Bai J and Perron P, 2003. Computation and Analysis of Multiple Structural Change Models. *Journal of Applied Econometrics*, 18, 1-22.

Campbell, John and Shiller, Robert, 1996. A Scorecard for Indexed Government Debt. Chapter in *NBER Macroeconomics Annual 1996, Volume 11 (1996)*, Ben S. Bernanke and Julio J. Rotemberg, Eds, 155 - 208.

Canzoneri, Matthew and Dellas, Harris, 1998. Real Interest Rates and Central Bank Operating Procedures. *Journal of Monetary Economics* 42(3), 471-494.

Castle J., Clements M.P. and Hendry D.F., 2013. Forecasting by Factors, by Variables, by Both or Neither. *Journal of Econometrics*, 177(2), 305-319.

Clark T.E. and McCracken M.W., 2009. Improving Forecast Accuracy by Combining recursive and rolling forecasts. *International Economic Review*, 50(2), 363-395.

Collard, Fabrice and Dellas, Harris, 2002. Exchange rate systems and macroeconomic stability. *Journal of Monetary Economics* 49(3), 571-599.

Davidson, R. and MacKinnon, R.G., 2006. Bootstrap methods in Econometrics. In Palgrave Handbook of Econometrics, ed C.M. Terence and Patterson K. vol 1 pp812-838, Palgrave Macmillan.

Elliott, Graham and Timmermann, Allan, 2008. Economic Forecasting. *Journal of Economic Literature* 46(1), 3-56.

Faust, J. and Wright, J. H., 2013. Forecasting inflation. In *Handbook of Economic Forecasting*, eds G. Elliot and A. Timmermann, vol 2A, Elsevier.

González-Rivera, Gloria, 2013. *Forecasting for Economics and Business*. The Pearson Series in Economics, Pearson/Addison-Wesley.

Goyal A. and Welch I., 2003. Predicting the equity Premium with Dividend Ratios. *Management Science*, 49(5), 639-654.

Granger, Clive, 1969. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* 37(3), 424-38.

Granger, Clive, 2008. Non-linear Models: Where Do We Go Next – Time Varying Parameter Models. *Studies in Nonlinear Dynamics and Econometrics* 12, 1-10.

Granger, Clive and Ramanathan, Ramu, 1984. Improved methods of combining forecasts. *Journal of Forecasting* 3(2), 197-04.

Harvey, Andrew C., 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press: Cambridge.

Hendry D.F., 2000. On detectable and non-Detectable structural Change. *Structural Change and Economic Dynamics*, 11(1), 639-654.

Inoue A., Jin L. and Rossi B, 2017. Rolling window selection for out-of-sample forecasting with time varying parameters. *Journal of Econometrics*, 196(1), 55-67.

Inoue A. and Rossi B., 2012. Out-of-Sample Forecast Tests robust to the choice of window size. *Journal of Business and Economic Statistics*, 30(3), 432-453.

Issing Otmar, 2008. *The Birth of the Euro*. Cambridge University Press: New York.

Koop G. and Potter S.M., 2007. Estimation and Forecasting in Models with multiple breaks. *The Review of Economic Studies*, 74(3), 763-789.

Kuan C.M. and White H. 1994, *Artificial Neural Networks: An Econometric Perspective*. *Econometric Reviews* 13, 1-91

Nankervis J.C. and Savin N., 1996. The level and Power of the Bootstrap t test in the AR(1) model with Trend. *Journal of Business and Economic Statistics*, 14(2), 161-168.

Pesaran M.H., Pettenuzzo D. and Timmermann A., 2006. Forecasting Time Series subject to Multiple Structural Breaks. *The Review of Economic Studies*, 73(4), 1057-1084.

Pesaran M.H. and Timmermann A., 2002. Market Timing and Return Prediction under Model Instability. *Journal of Empirical Finance*, 9(5), 495-510.

Pesaran M.H. and Timmermann A. 2007. Selection of Estimation Window in the presence of Breaks. *Journal of Econometrics*, 137(1), 134-161.

Rossi B., 2013a. Exchange rate Predictability. *Journal of Economic Literature*, 51(4), 1063-1119.

Rossi B., 2013b. Advances in Forecasting under Instability. In *Handbook of Economic Forecasting*, eds G. Elliot and A. Timmermann, vol 2B, Elsevier.

Stock James and Watson Mark, 2008. Phillips Curve Inflation Forecasts. NBER Working Paper No. 14322.

Swamy P. A. V. B. and Mehta Jatinder, 1975. Bayesian and non-Bayesian Analysis of Switching Regressions and a Random Coefficient Regression Model. *Journal of the American Statistical Association*, 70(351), 593–02.

Swanson N.R. and White H. A Model Selection Approach to Real Time Macroeconomic Forecasting Using Linear Models and Artificial Neural Networks. *Review of Economic and Statistics* 79, 540-50

Timmermann Allan, 2006. Forecast Combinations. *Handbook of Economic Forecasting* 1, 135–196.

Tobin James, 1981. Comment on 'On a Theoretical and Empirical Basis of Macroeconometric Models'. In eds J. Kmenta and J. Ramsey, *Large-Scale Macroeconometric Models*, Amsterdam: North Holland, 391-94.

White H, 1992, *Artificial Neural Networks: Approximation and Learning*, Cambridge, Blackwell

Appendix

Model description

The economy is populated by a large number of identical, infinitely-lived households and firms that are owned by the households. The firms produce a single, homogeneous, perishable good.¹

2.1 The Household

Household preferences are characterized by the lifetime utility function:²

$$E_t \sum_{T=0}^{\infty} \beta^T U(C_{t+T}, H_{t+T}) \quad (1)$$

where $0 < \beta < 1$ is a constant discount factor, C denotes consumption and H hours worked.

The utility function, $U(C, H) : \mathbb{R}_+ \times [0,1] \rightarrow \mathbb{R}$ is increasing and concave in its arguments; \mathbb{R}_+ is real space and $[0,1]$ is the domain of hours worked.

In each and every period, the representative household faces two constraints: a budget constraint and a cash in advance constraint (CIA). The budget constraint takes the form:

$$B_t + M_t \leq I_{t-1}B_{t-1} + M_{t-1} - P_{t-1}C_{t-1} + N_t + \Pi_t + W_t H_t \quad (2)$$

and the CIA constraint

¹ The model can easily accommodate capital and investment. We abstract from them as they do not bring anything of interest to the results.

² $E_t(\cdot)$ denotes mathematical conditional expectations. Expectations are conditional on information available at the beginning of period t .

$$P_t C_t \leq M_t \quad (3)$$

where B_t and M_t are nominal bonds and money acquired during period t , respectively, P_t is the nominal price of the good, I_t is the gross nominal interest rate and W_t is the nominal wage rate. N_t is a nominal transfer of money the household receives from (or pays to) the government and Π_t is its share of firm profits. The household determines its consumption, bond, money holdings and hours worked plans by maximizing its utility (1) subject to the sequence of budget constraints (2) and (3).

Let us assume that the instantaneous utility function takes the form:

$$U(C_t, H_t) = \frac{1}{1-\gamma} C_t^{1-\gamma} - \frac{1}{1+\sigma} H_t^{1+\sigma} \quad (4)$$

where γ, σ are positive constants.

The first order conditions are then:

$$H_t^\sigma = C_t^{-\gamma} \frac{W_t}{P_t} \quad (5)$$

$$\frac{C_t^{-\gamma}}{P_t} = \beta I_t E_t \frac{C_{t+1}^{-\gamma}}{P_{t+1}} \quad (6)$$

2.2 The Firm

Output, Y_t , is produced using labor. The production function takes the form:

$$Y_t = A_t H_t^a \quad (7)$$

where A is a random variable capturing productivity. The producers are assumed to behave competitively and to determine their demand for labor by maximizing their profits,

Π :

$$\Pi_t = P_t A_t H_t^\alpha - W_t H_t \quad (8)$$

The first order condition is:

$$\alpha P_t A_t H_t^{\alpha-1} = W_t \quad (9)$$

2.3 The monetary authorities

The monetary authorities determine the money supply, M^S , according to the exogenous³ rule:

$$M_t^S = \bar{M} e^{g_t} \quad (10)$$

where g_t is a stationary process.

2.4 The equilibrium

We first solve the model under flexible prices in order to have the relevant benchmark. Assuming that the nominal interest rate is positive so that the CIA constraint is satisfied with equality, we can determine the equilibrium of the economy as the solution to

³ We could have instead written the money supply rule in a form that makes the growth rate of the money supply rather than the money stock stationary.

equations 5–6, 9, the production function, 7, the money supply rule, 10, and the following conditions:

$$C_t = Y_t \quad (11)$$

$$M_t = M_t^s \quad (12)$$

$$M_t^s = P_t C_t \quad (13)$$

$$B_t = 0 \quad (14)$$

It is straightforward to show using equations 5, 7, 9 and 11 that under flexible prices, Y_t , C_t , H_t , and $\frac{W_t}{P_t}$ are invariant to the supply of money. Let us use Z_t^* , H_t^* , Y_t^* and C_t^* to represent the flexible price equilibrium values of the real wage, hours worked, output and consumption respectively.

2.5 Sticky Wages

Now suppose that the nominal wage for period t is set in period $t-1$, that is before the realization of the money supply shock. Furthermore, following standard practice, assume that it is set at a level that is *expected* to produce a wage that would equate labor demand and labor supply in period t (the flexible price real wage). In other words, the contract wage, W^c satisfies:

$$E_{t-1} \frac{W^c}{P_t} = Z_t^* \quad (15)$$

The realized real wage in period t , $\frac{W^c}{P_t}$ will differ from the expected real wage, $E_{t-1} \frac{W^c}{P_t} = W^c E_{t-1} \frac{1}{P_t}$ when the price (inflation) forecast of the nominal wage setters turns out to be wrong. This will have implications for the equilibrium values of real economic activity. In particular, forecast errors will lead to fluctuations of economic activity and employment away from their flexible price counterparts. In order to derive the implication of sticky wages, we drop the labor supply equation (5) and use the wage contract equation in the labor demand curve (equation 9). Using small letters to denote logs and approximating around the flexible price equilibrium gives:

$$h_t - h_t^* = \frac{1}{1-\alpha} (p_t - E_{t-1} p_t) \quad (16)$$

which combined with the production function leads to the celebrated Phillips curve

$$y_t - y_t^* = \frac{\alpha}{1-\alpha} (p_t - E_{t-1} p_t) \quad (17)$$

Thus, inflation forecast errors lead to inefficient variation in macroeconomic activity.

We can also use the solution of the model to compute the level of welfare associated with alternative forecasting methods and thus to access the gains from adopting our proposed method. Using the solution of the model in the utility function produces the following expression for welfare

$$U = \frac{1}{1-\gamma} \alpha^{\frac{\alpha(1-\gamma)}{1+\sigma+\alpha(\gamma-1)}} \left(\frac{\Pi(t)}{E_{t-1}\Pi(t)} \right)^{\frac{\alpha(1-\gamma)}{1-a}} - \frac{1}{1+\sigma} \alpha^{\frac{1+\sigma}{1+\sigma+\alpha(\gamma-1)}} \left(\frac{\Pi(t)}{E_{t-1}\Pi(t)} \right)^{\frac{1+\sigma}{1-a}} \quad (18)$$

If inflation forecast errors are zero (or, equivalently, if wages are flexible then this expression reduces to (18) with $\Pi(t) = E_{t-1}\Pi_t$

By plugging into these expressions the time series for the actual and predicted inflation under the alternative forecasting methods from the empirical sections we can compute the associated levels of (expected) welfare. We can then utilize these numbers to compute the percentage of steady state consumption that the agents are willing to pay in order to have access to superior forecasting methods (as proposed by Lucas).