scottish institute for research in economics

SIRE DISCUSSION PAPER

SIRE-DP-2011-19

Estimating Phillips Curves in Turbulent Times using the ECB's Survey of Professional Forecasters

Gary Koopy
University of Strathclyde

Luca Onorante
European Central Bank

www.sire.ac.uk

Estimating Phillips Curves in Turbulent Times using the ECB's Survey of Professional Forecasters*

Gary Koop†
University of Strathclyde
gary.koop@strath.ac.uk

Luca Onorante
European Central Bank
luca.onorante@ecb.int

February 2011

Abstract: This paper uses forecasts from the European Central Bank's Survey of Professional Forecasters to investigate the relationship between inflation and inflation expectations in the euro area. We use theoretical structures based on the New Keynesian and Neoclassical Phillips curves to inform our empirical work. Given the relatively short data span of the Survey of Professional Forecasters and the need to control for many explanatory variables, we use dynamic model averaging in order to ensure a parsimonious econometric specification. We use both regression-based and VAR-based methods. We find no support for the backward looking behavior embedded in the Neoclassical Phillips curve. Much more support is found for the forward looking behavior of the New Keynesian Phillips curve, but most of this support is found after the beginning of the financial crisis.

Keywords: inflation expectations, survey of professional forecasters, Phillips curve, Bayesian

JEL classification codes: E31, C53, C11

*This working paper should not be reported as representing the views of the ECB. The views expressed are those of the authors and do not necessarily reflect those of the ECB.

†Financial support from the ESRC under grant RES-062-23-2646 is gratefully acknowledged. Gary Koop is a Fellow at the Rimini Centre for Economic Analysis.

1 Introduction

Expectations are essential for determining economic outcomes and for policymakers. Economic theory suggests that inflation, in particular, should be heavily influenced by expectations of future inflation. Producers' and consumers' decisions are taken on the basis of inflation expectations and from these decisions prices and real outcomes are determined. Central banks are very interested in expectations: anchoring long-term inflation expectations is crucial for maintaining price stability. And inflation expectations are an essential component to be monitored, providing essential information when conducting monetary policy in a forward-looking manner. This is because they determine the overall effect of exogenous shocks on prices, both in size and length, and the sacrifice in terms of output necessary to bring inflation to target once it has drifted away.

Quantitative information on the uncertainty surrounding the macroeconomic outlook is also an important part of the macroeconomic information set used by economic policy-makers and central banks (see, for example, the discussions in Tay and Wallis 2000, Wallis 2003, or Sims 2002), and may be an additional factor in driving financial market variables, households' and firms' decision-making and ultimately inflation outcomes. Indeed, uncertainty plays a role in several models of consumer and investor behavior (e.g. Lahiri et al. 1988, Giordani and Söderlind 2003, and D'Amico and Orphanides, 2008).

This paper focuses on inflation and inflation expectations in the euro area. We adopt specifications motivated by macroeconomic theory and, in particular, the New Keynesian and Neoclassical Phillips curves. We analyze: 1) whether inflation expectations have a sizable and measurable effect on inflation in the euro area; 2) whether this effect can be considered as robust evidence, considering the strong amount of uncertainty about how to precisely model inflation and the heterogeneous results found in the literature. Furthermore, and following the recent crisis, we try to assess 3) whether the link between inflation and expectations has changed in recent years; and 4) whether the increased inflation uncertainty observed during the crisis is likely to play a relevant role in determining future inflation.

To address these issues, we use direct measures of inflation expectations provided in the European Central Bank's Survey of Professional Forecasters and model these jointly with inflation itself as well as a standard set of predictors. In order to econometrically estimate our set of models, we require

a modelling framework which allows for many predictors as well as allowing for model change. However, such a framework risks being over-parameterized since the SPF was only launched in 1999 and, thus, the sample is quite short. Accordingly, our econometric methods use dynamic model averaging (DMA, see Raftery, Karny and Ettler 2010). DMA is a relatively new statistical method which, in a regression context, lets the set of explanatory variables switch over time. This contrasts with conventional nonlinear time series methods (e.g. time varying parameter, TVP, or Markov switching models) which typically allow for coefficients to change within a single model. In a macroeconomic forecasting exercise involving a large set of predictors, Koop and Korobilis (2009a) find that DMA forecasts well (much better than comparable TVP models). By switching between different highly parsimonious models, it avoids being over-parameterized, but at the same time allowing for the regime changes which have occurred in the macroeconomy. These considerations suggest DMA is an appropriate method for use in the present application. The previous literature uses DMA in the context of one equation regression models. In the present paper, we extend these methods for use with VARs.

Our empirical work involves both regression-based and VAR-based DMA methods. Both empirical exercises suggest that forward looking inflation expectations are an important predictor for current inflation. This finding is in line with the New Keynesian Phillips curve. However, most support for the New Keynesian Phillips curve is found near the end of our sample, after the beginning of the financial crisis. We never find backward looking inflation expectations to be important.

2 The Theoretical Relationship Between Inflation and Inflation Expectations

The existence of a strong relationship between inflation and inflation expectations is nowadays a standard component of any theoretical framework. Abstracting from model specificities, this relationship is generally modeled through some version of the Phillips curve. Three formulations of the Phillips curve are popular in the literature.

In the Neoclassical Phillips curve (Phelps 1968, Friedman 1968, Woodford 2003) a fraction of prices are set one period in advance, while the others are

fully flexible. Thus, inflation is influenced by the past expectation of current inflation and the Phillips curve can be formulated as:

$$\pi_t = \alpha \pi^e_{t|t-h} + \beta X_t \tag{1}$$

where π_t is inflation at time t, $\pi^e_{t|t-h}$ is inflation at time t as expected in $t-h$, X_t is a vector of explanatory variables including a measure of capacity utilization (e.g. the output gap or unemployment) as well as other regressors and α and β are parameters to be estimated. In the empirical section of this paper we have $h = 4$, consistent with quarterly data and expectations one year ahead.

In the New Keynesian Phillips curve, described for instance in Calvo (1983), only a fraction of firms can adjust prices in each period. These firms choose to adjust them optimally, taking into account expected future price developments and the probability that they may not be able to adjust their prices in the following periods. A general formulation for the New Keynesian curve is:

$$\pi_t = \alpha \pi^e_{t+h|t} + \beta X_t \tag{2}$$

where current expectations of future prices $\pi^e_{t+h|t}$ determine the current inflation rate π_t.

Neither the Neoclassical or the New Keynesian Phillips curves performed very well in empirical terms. To solve this problem,the hybrid Phillips curve (Gali and Gertler, 1999) introduces lagged inflation in order to explain persistence. This formulation can be derived on the hypothesis that at every period a fraction of firms sets prices optimally, as in Calvo, while the rest use a backward-looking rule. The hybrid New Keynesian Phillips curve can be formulated as:

$$\pi_t = \alpha \pi^e_{t+h|t} + (1-\alpha)\, \pi_{t-1} + \beta X_t \tag{3}$$

In this form of the curve both current expectations of future prices and past inflation determine current inflation. The intrinsic persistence of inflation in the hybrid curve contrasts with the previous formulations, where any persistence is stemming from expectations. The existence of different versions of the curve is a first element of complexity we must take into account when trying to bring the theory to the data.

A second important factor is the relationship with inflation uncertainty. The idea can be traced back to Friedman (1977), who suggested that higher

average inflation could be associated to higher inflation uncertainty. This idea was developed by Ball (1992), who proposes a model in which higher inflation leads to increasing uncertainty over the monetary policy stance. The opposite view of a negative effect of inflation on its uncertainty is taken by Pourgerami and Maskus (1987), who point out that in an environment of accelerating inflation agents may invest more resources in inflation forecasting, thus reducing uncertainty (see also Ungar and Zilberfarb, 1993). Causality from inflation uncertainty to inflation is a property of models based on the Barro–Gordon setup, such as the one due to Cukierman and Meltzer (1986). A more recent model that has uncertainty at its heart is the sticky-information model proposed by Mankiw and Reis (2002).

3 Which expectations?

In addition, macroeconomists have diverging opinions about an essential element of the curve itself: the expectations. Expectations, if measured correctly, are probably the most important variable required in the estimation of the Phillips curve. Among theorists, there is disagreement on how to model and measure expectations. Until the mid-1970s, expectations were generally modeled as adaptive processes, i.e. backward looking. Backward looking expectations naturally imply persistence in inflation, thus explaining what seems to be a feature of inflation developments. However, backward looking expectations also imply systematic expectation errors. The introduction of rational expectations rectified this inconsistency, assuming that expectation errors are non-systematic. However, modelling inflation assuming rational expectations has not performed well in empirical testing. Today, mainstream models featuring rational expectations are more and more complemented by frameworks featuring alternative formulations. Recent research has focused on the formation of expectations through learning processes, where agents attempt to continuously improve their knowledge of the economy. Like rational expectations, the formation of expectations through learning is consistent with forward looking behavior, but learning also gives rise to persistence in expectations, which is a desirable feature.

On the empirical side, following Gali and Gertler (1999), most econometric estimates of the New Keynesian Phillips curve simply estimate under the assumption of rational expectations. Taking advantage of the fact that under rational expectations forecasting errors are assumed to be unpredictable,

these estimates use data on future inflation outcomes. GMM estimation methods are often used. This technique confronts the problem that finding an instrumental variable which is at the same time outside the Phillips curve but has good predictive power for future inflation is difficult in practice. Additional difficulties of identification are discussed in Pesaran (1987), Ma (2002), Mavroeidis (2005) and Kleibergen and Mavroeidis (2009). For instance, Kleibergen and Mavroedis (2009) find parameters of the hybrid New Keynesian Phillips curve to be weakly identified and argue that this partially accounts for the conflicting estimates reported in the literature. However, they find enough identification occurs so as to make some empirical conclusions (e.g. they find that forward-looking dominates backward-looking behavior). Finally, assuming rational expectations sidesteps without solving it the issue of the relationship between expectations and inflation and is more useful for exploring the output gap parameters.

The empirical literature just discussed treats inflation expectations as being unobserved (thus, requiring the rational expectations assumption and clever justification of instrumental variables, etc.). If inflation expectations are observed, then the concerns expressed in the previous paragraph do not occur. This has inspired a recent literature which tries to obtain direct measures of inflation expectations. The present paper falls in this tradition.

Inflation expectations are typically taken from surveys. Particularly in the US, surveys of experts or the public at large have been used to empirically proxy inflation expectations. Examples include Ang, Bekaert and Wei (2007), Leduc, Sill and Stark (2007) and Mehra and Herrington (2008). The first of these is a forecasting paper which finds inflation surveys to be better predictors of inflation than a range of other alternatives. The latter use inflation surveys within a structural VAR model of the US economy to address issues of important for macroeconomic policy. Clark and Davig (2008) surveys the US literature and presents results from a tri-variate structural VAR for inflation, short- term and long-term inflation expectations. It finds that shocks to long-term inflation expectations have an effect on inflation, but shocks to short-term expectations have much less of an effect.

Focusing on the euro area, Paloviita (2005) compares different specifications of the Phillips curve. Paloviita and Virén (2005) estimate a VAR model of inflation, inflation expectations and output gap that allows for an analysis of the interrelationship between these variables.

A similar shift towards the use of survey data can be observed in relation to studies involving inflation uncertainty. Early work involved estimates

of uncertainty based on ARCH models (see Caporale, Onorante and Paesani 2010, for a recent treatment). This literature has been progressively complemented by studies using disagreement in survey data as a measure of inflation uncertainty (e.g. Davis and Kanago 1996, Bomberger 1996, Hayford 2000 and Giordani and Söderlind, 2003).

In this paper, we use inflation expectations as reported in the Survey of Professional Forecasters (SPF). The SPF, launched by the European Central Bank (ECB) in the last quarter of 1998, immediately following the establishment of the single currency, asks a panel of approximately 75 forecasters located in the European Union (EU) for their short- to longer-term expectations for euro area inflation, growth and unemployment (see Garcia 2003, for details). In this paper, we focus on inflation expectations. An important feature of the SPF is that it provides information for the euro area on a rolling window basis (in other words, always the same number of quarters ahead). Note that each of the up to 75 forecasters is asked to produce a distribution of the forecast for each variable. We use as our inflation expectations variable the average (across forecasters) of the mean forecasts of inflation one year ahead (expressed as an annualized rate). As for our variable reflecting inflation uncertainty, the SPF allows for the computation of several indicators:

- Using the mean forecast of each forecaster, the disagreement among forecasters can be calculated, as measured by the variance of these point forecasts.
- The variance of the forecast distribution of each forecaster can be calculated. Then this can be averaged across forecasters as a measure of individual uncertainty.
- Finally, the forecast distributions provided by the forecasters can be aggregated into one forecast distribution and the variance of this aggregate distribution used as a measure of uncertainty. It can be shown that this variance is equal to the average variance of the individual forecasters' distributions plus the variance of the point estimates. Thus, this measure combines the previous two measures, taking into account both individual uncertainty and disagreement.

These measures of uncertainty are strongly correlated. Our empirical results use the first measure, based on disagreement, and use the other two as robustness checks.

Not only does the survey of professional forecasters provide important information relevant to the conduct of monetary policy, it also has important information content for future development of inflation. The statistical properties of the survey of professional forecasters have been analyzed by Bowles et al. (2007) and in Garcia and Manzanares (2007). The authors find that the one year ahead inflation forecast slightly underestimated inflation, a result that they attribute to the fact that euro area has been affected by a number of upward shocks to prices (for example, forecast errors are correlated with oil and food price movements) and to a corresponding slight overestimation of the unemployment rate. They also find that the SPF forecasts for inflation are systematically more accurate than naïve forecasts and contain information about the future beyond what is already contained the most recent data.

However, and in contrast to the growing amount of literature related to the Survey of Professional Forecasters in the United States, to our knowledge there is no literature using the EU SPF and going beyond the analysis of the papers quoted above. Indeed, a current drawback with this survey is that it has been available for a relatively short time span. And this time span has covered a time of a great deal of macroeconomic instability, including the recent financial crisis. This consideration necessitates the use of econometric methods which allow for regime change, but at the same time ensuring parsimony so as to try and obtain relatively precise inferences with a short data set. It is to this we now turn.

4 Econometric Methods

4.1 General Considerations

All of our econometric methods will be extensions of regressions or VARs. Univariate, regression-based specifications are simpler and more parsimonious, allowing us to directly estimate equations motivated by the hybrid New Keynesian and Neoclassical Phillips curves. A bivariate specification, modelling inflation and an expectation of inflation jointly, will additionally allow us to investigate whether there are feedbacks from inflation to expectations (see, e.g., Clark and Davig, 2008). In this paper, we consider both univariate and bivariate specifications. It is worthwhile from the outset to make clear the dependent (y_t) and explanatory variables we use. In our re-

gressions, the dependent variable is inflation, $y_t = \pi_t$. Our extensions of VARs will be bivariate ones. Motivated by the Neoclassical Phillips curve in (1), some of our VARs involve $y_t = \left(\pi^e_{t|t-4}, \pi_t\right)'$. Motivated by New Keynesian Phillips curve (2), the remainder of our VARs use $y_t = \left(\pi^e_{t+4|t}, \pi_t\right)'$. We use annualized HICP inflation (the year-on-year percentage change in the HICP) as the SPF relates to this inflation definition.

Our explanatory variables will differ by specification, but motivated by a wish to investigate whether forward or backward looking expectations are important for inflation, our regressions include $\pi^e_{t+4|t}$ and $\pi^e_{t|t-4}$ as explanatory variables and our VARs will include appropriate lagged dependent variables as explanatory variables. In addition, we include other exogenous explanatory variables (labelled X_t in equations 1, 2 and 3). The literature on the so-called generalized Phillips curve is too voluminous to survey here (see, e.g., Stock and Watson, 2008), but a message coming out of this literature is that there are other explanatory variables which have power for inflation. Any model which omits such explanatory variables risks mis-specification and risks misunderstanding the relationship between inflation and inflation expectations. Another message (see, e.g., Koop and Korobilis 2009a, for an application using US data) is that the predictors for inflation can change over time. This is especially important in light of the recent financial crisis. As a consequence, we want to allow for the relationship between inflation and its expectation to potentially change over time. This will allow us to investigate whether the recent financial crisis can be considered a structural change in the Phillips curve. The unprecedented size and length of the recent recession has not been matched by a corresponding decrease in inflation, raising the possibility of a shift in the Phillips curve in the euro area. Indeed, a shift has been also found in recently estimated models with changing coefficients (Caporale et al., 2010). As far as the relationship between inflation and expectations is concerned, the first years of the EMU were characterized by inflation and expectations solidly anchored at about 2%. However, during the current downturn, with headline inflation turning negative for the euro area, inflation expectations, even though lower than their long-run average, remained close to historical values. Whether this is also the signal for a change in structure of the relationship between the two is an interesting question for our empirical investigation.

Accordingly, including a rich set of predictors for inflation, but using an

econometric method which allows for change over time is essential. In this paper, we use the following set of eight euro area predictors which includes a wide range of variables thought to be important by theory or found to be important in other studies:

1. ULC: unit labour costs (annual percentage change), seasonally adjusted, ESA95 National Accounts.

2. GAP: output gap, Economic Outlook (OECD).

3. UNEMP: standardized unemployment rate in percentage of civilian workforce, seasonally adjusted, Eurostat.

4. SPF_VAR: inflation uncertainty as measured by disagreement among forecasters (see Section 3).

5. POIL: oil price (annual percentage change, Brent crude).

6. ISHORT: Euribor 3-month interest rate, historical close.

7. ILONG: euro area 10-year government benchmark bond yield.

8. STOX: Dow Jones Eurostoxx 50 Index (percentage change), historical close, provided by Reuters.

The relatively short data span for the SPF poses problems for developing statistical methods which achieve the goals we have just set out. If we use a very flexible specification which includes many predictors and allows for breaks (or other sorts of coefficient change), it will be difficult to obtain reasonable estimates: the number of parameters will simply be too large relative to the number of observations. On the other hand, if we work with too simple models (e.g. with few explanatory variables and/or coefficients that are constant over time), then we risk working with mis-specified models and failing to address our research questions of interest. What we do in this paper is use a method called Dynamic Model Averaging (DMA) which, from a flexible specification, allows us to uncover parsimonious specifications which can change over time.

4.2 Dynamic Model Averaging

DMA is developed in Raftery, Karny and Ettler (2010) and used in Koop and Korobilis (2009a) and the reader is referred to these papers for complete details (see also the appendix to this paper). Here we describe the basic ideas in the context of the regression model where $X_t = (x_{1t}, .., x_{k_z t})'$ contains k explanatory variables:

$$y_t = \beta_0 + \sum_{i=1}^{k} \beta_i x_{it} + \varepsilon_t$$

where ε_t is i.i.d. $N(0, \sigma^2)$, Within this general model, we can define $K = 2^k$ restricted models which contain subsets of the k explanatory variables. The dynamic aspect of DMA arises since it allows for a different model to hold at each different time period. Let $L_t \in \{1, 2, .., K\}$ denote which model holds at time t and $q_{t|s,j} = \Pr(L_t = j|y^s)$ be the probability that model j holds at time t given information through time s. DMA is a recursive algorithm which allows for the calculation of $q_{t|t,j}$ and $q_{t|t-1,j}$ for $j = 1, .., K$. In an estimation exercise such as the one in this paper (e.g. estimating coefficients or impulse responses) $q_{t|t,j}$ can be used to carry out model averaging in a time varying fashion. $q_{t|t-1,j}$ can be used in a similar fashion when forecasting y_t given information through time $t-1$.

Note that, since K can be large and DMA allows for a different model to hold in every time period, the computational burden can be enormous. That is, with T observations, 2^{Tk} possible combinations of models at various times can exist. Exhaustive evaluation of all these combinations is computationally infeasible unless k and T are both small. The contribution of Raftery et al. (2010) was to develop a clever approximation, involving a so-called forgetting factor, α, which reduces the computational burden enormously (in essence, it only requires recursive estimation of K models).

To explain DMA in a bit more detail, we introduce notation where $y^s = (y_1, .., y_s)'$ and, thus, $p_k(y_t|y^{t-1})$ is the predictive density for model k. Note that the predictive density in the regression model has a familiar form that can easily be evaluated. The predictive density appears in the model updating equation of:

$$q_{t|t,s} = \frac{q_{t|t-1,s} p_k(y_t|y^{t-1})}{\sum_{l=1}^{K} q_{t|t-1,l} p_l(y_t|y^{t-1})}. \quad (4)$$

If we knew $q_{t|t-1,s}$ then, starting with $q_{0|0,s}$ we could recursively calculate the key elements of DMA: $q_{t|t,j}$ and $q_{t|t-1,j}$ for $j = 1, .., K$. Raftery et al. (2010) provide this missing link by using the approximation:

$$q_{t|t-1,s} = \frac{q_{t-1|t-1,s}^{\alpha}}{\sum_{l=1}^{K} q_{t-1|t-1,l}^{\alpha}}. \tag{5}$$

A detailed justification for why this is a sensible approximation is given in Raftery et al. (2010). Suffice it to note here that it implies:

$$\begin{aligned} q_{t|t-1,s} &\propto \left[q_{t-1|t-2,s} p_s\left(y_{t-1}|y^{t-2}\right)\right]^{\alpha} \\ &= \prod_{i=1}^{t-1} \left[p_s\left(y_{t-i}|y^{t-i-1}\right)\right]^{\alpha^i}. \end{aligned}$$

Thus, model s will receive more weight at time t if it has fit well in the recent past (where fit is measured by the predictive likelihood, $p_s\left(y_{t-i}|y^{t-i-1}\right)$). The interpretation of "recent past" is controlled by the forgetting factor, α. Thus, if $\alpha = 0.99$ (our benchmark value and also the value used by Raftery et al., 2010), forecast performance five years ago receives 80% as much weight as forecast performance last period (when using quarterly data). If $\alpha = 0.95$, then forecast performance five years ago receives only about 35% as much weight. These considerations suggest that we focus on the interval $\alpha \in (0.95, 0.99)$.

In our short data set, the potential advantages of DMA are clear. We can include a large number of explanatory variables, but DMA can attach weight to more parsimonious models, lessening the problems caused by our short data span.[1] Furthermore, DMA allows for model change. It can capture cases where certain explanatory variables are important in certain periods, but not in others. Given our application, which spans the time from the introduction of the euro through the recent financial crisis, allowing for such chance is likely crucial. In short, DMA is likely to meet our needs for a statistical framework involving many explanatory variables and allowing for change in an effective and parsimonious way.

We have described DMA in terms of the regression model. We use this in our univariate empirical exercises. In our multivariate empirical work, we

[1] See Koop and Korobilis (2009a) for evidence that DMA can effectively find very parsimonious models.

extend the existing literature by developing DMA methods for multivariate models such as VARs. Details are provided below and in the appendix.

We note also that, in the past, DMA has been done in the context of time-varying parameter (TVP) models where the coefficients evolve as $\beta_{it} = \beta_{i,t-1} + u_t$. Given our short data span and need to keep the model as parsimonious as possible, we do not consider this extension. Furthermore, in the empirical application of Koop and Korobilis (2009a), it is found that allowing for models to switch over time is of greater empirical benefit than allowing for coefficients to evolve in a TVP fashion.

4.3 Regression-Based Methods

We begin by investigating issues relating to the Neoclassical and New Keynesian Phillips curves in the following regression specification:

$$\pi_t = \alpha_0 + \alpha_1 \pi^e_{t|t-4} + \alpha_2 \pi^e_{t+4|t} + \sum_{i=1}^{8} \beta_i x_{it} + \sum_{i=1}^{p} \gamma_i \pi_{t-i} + \varepsilon_t \qquad (6)$$

where x_{it} for $i = 1, .., 8$ are the eight explanatory variables listed in Section 4.1. The coefficients α_1 and α_2 are of most interest since they shed light on the backward or forward looking nature of the Phillips curve. Note that including lags of the dependent variable does not shorten our data span (i.e. it is only the SPF variables for which the time span of the data is short). Given the timing convention of the SPF variables in (6) all of our regressions use data from 1998Q4 to 2010Q2, regardless of choice of lag length.

We begin by presenting some OLS estimates of (6) for $p = 4$. BIC's indicate that this is the preferred lag length (although the BIC for $p = 1$ is only very slightly higher). However, given our desire to allow for model change (and below we will present evidence of its importance), these OLS results (which assume constant coefficients) should be interpreted with caution.

Table 1: OLS Results		
Variable	estimate	t-stat
$\pi^e_{t\|t-h}$	-0.551	-1.102
$\pi^e_{t+h\|t}$	1.123	2.619
π_{t-1}	0.444	3.103
π_{t-2}	-0.244	-1.368
π_{t-3}	0.253	1.342
π_{t-4}	-0.601	-3.377
ULC	0.201	1.617
GAP	0.092	0.881
UNEMP	-0.202	-0.559
SPFVAR	-0.953	-0.497
POIL	0.002	0.887
ISHORT	-0.090	-0.482
ILONG	-0.274	-1.590
STOX	-0.009	-2.594

The OLS results strongly support forward looking expectations as being important in the Phillips curve, while backward looking expectations are insignificant. The other explanatory variables are mostly of theoretically-sensible sign (and where not, as with the GAP variable, this is due to correlation between related explanatory variables), and are typically found to be insignificant. This is as one would expect if we are including too many correlated predictors and the predictor set is changing over time. Accordingly, we turn to DMA results which can ensure more parsimony as well as switching between different sets of predictors.

We implement DMA based on Koop and Korobilis (2009a) and equation (6). The specification of a DMA requires: i) a prior for the parameters, ii) a prior over the initial model probabilities and iii) a choice for the forgetting factor. We use i) a noninformative prior over the parameters, ii) make the noninformative choice of $q_{0|0,j} = \frac{1}{K}$ for $j = 1, .., K$ and iii) set $\alpha = 0.99$. We initialize with OLS results (including all the predictors) using the initial three years of data, before doing DMA on the remainder of the sample. Posterior inference in the regression model using a noninformative prior involves standard textbook formulae (e.g. Koop, 2003, pages 36-38). The predictive density, which plays a crucial role in DMA via (4), also has a textbook formula (e.g. Koop, 2003, page 46).

Our regression model has 14 predictors (i.e. the 8 explanatory variables listed in Section 4.1, four lags of the dependent variable and forward and backward looking inflation expectations) and, in light of our small sample size, a first question of interest is whether DMA is being parsimonious. Figure 1 sheds light on this. It plots the expected size of the models used by DMA at each point in time where

$$E\left(Size_t\right) = \sum_{k=1}^{K} q_{t|t-1,k} Size_k,$$

and $Size_k$ is the number of predictors in model k. It can be seen that DMA is being fairly parsimonious. That is, it could have included 14 predictors but tends to be choosing roughly half of them. Furthermore, there is some variation over time.

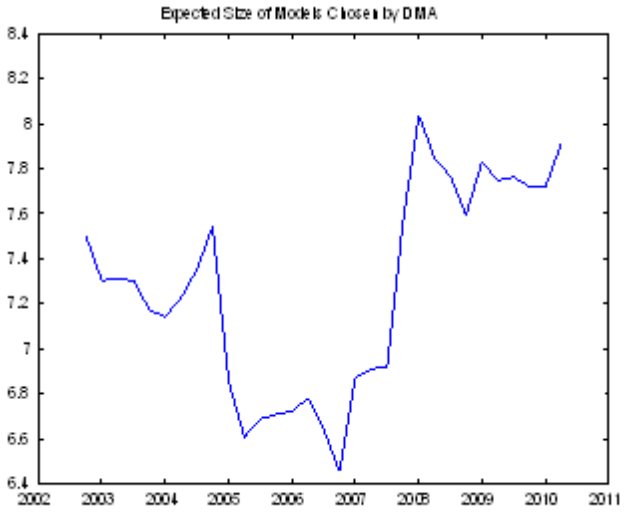

Figure 1

In order to see which predictors are being chosen by DMA, Figures 2 through 5 plot time-varying posterior inclusion probabilities associated with every explanatory variable. These are the probabilities allocated by DMA to models which contains each predictor (i.e. for the j^{th} predictor this is

$\sum_{i \in j} q_{t|t,i}$ where the summation is over all models which include the j^{th} predictor). Note that, near the beginning of the sample, there will be little information to estimate the posterior inclusion probabilities and, hence, it is possible that they differ little from their prior inclusion probability of $\frac{1}{2}$ (i.e. a priori, we assume each explanatory variable is equally likely to be included as excluded from the model, which is a standard noninformative choice). Similar considerations hold for any other model feature (e.g. such as regression coefficients) which will tend to be imprecisely estimated early in the sample. In Figures 2 through 5 we do find most (but not all, see results for UNEMP) posterior inclusion probabilities to begin near $\frac{1}{2}$. However, after this initial period, there is a wide variation in behaviors across predictors. In some cases, the weight attached by DMA is gradually changing over time. But in other cases, probabilities are switching very abruptly as DMA decides to include/exclude a variable at a particular point in time. For instance, for most of the sample, variables measuring changes in wages and oil prices (ULC and POIL) are relatively unimportant. However, in 2008 they switch to becoming important before abruptly switching to becoming unimportant again in 2009. In general, there is uncertainty over which predictor is important at each point in time. A method such as DMA is an appropriate method of dealing with this uncertainty.

Given our interest in the Phillips curve, Figure 2 is the most important graph. It can be seen that backward looking inflation expectations, although not completely excluded by DMA, are relatively unimportant. This finding is consistent with our OLS results of Table 1. The posterior inclusion probability of the forward looking inflation expectations variable is also consistent with the OLS results. However, it adds an important refinement in our understanding of the role of inflation expectations in the Phillips curve. That is, DMA is finding $\pi^e_{t+4|t}$ to be a very important predictor of inflation, but only near the end of the sample. Until 2008, we are finding it to be even less important than the backward looking expectations variable. It is only in 2008 to the present that it is a very strong predictor.

It is also interesting to note that our measure of uncertainty, SPFVAR, does not seem to have important explanatory power for inflation.

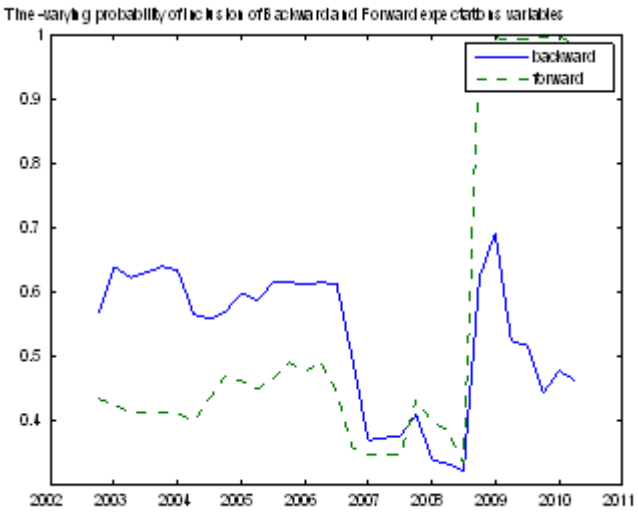

Figure 2

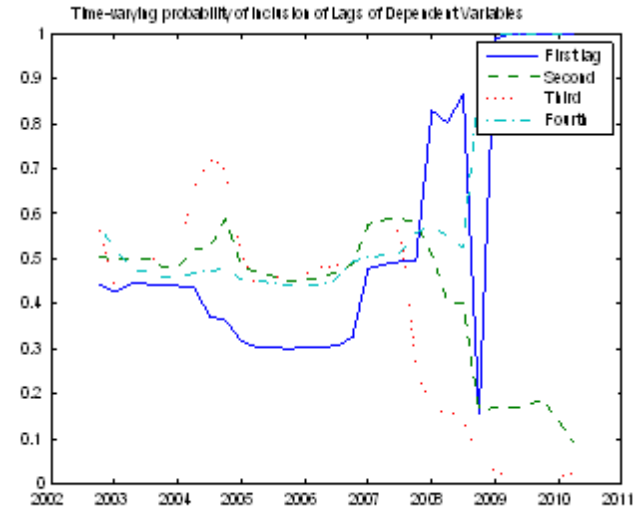

Figure 3

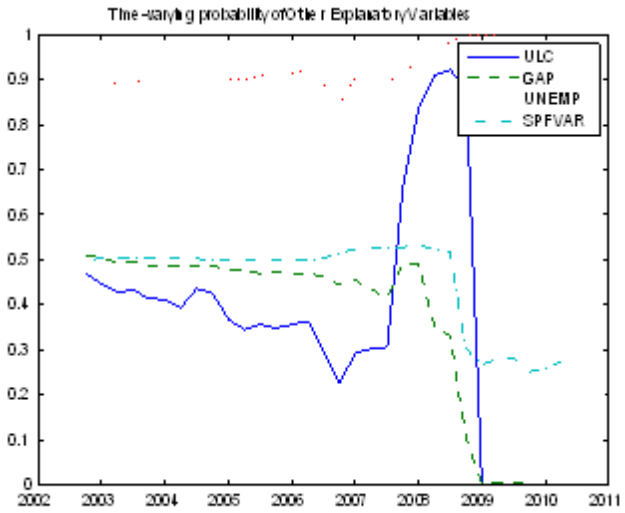

Figure 4

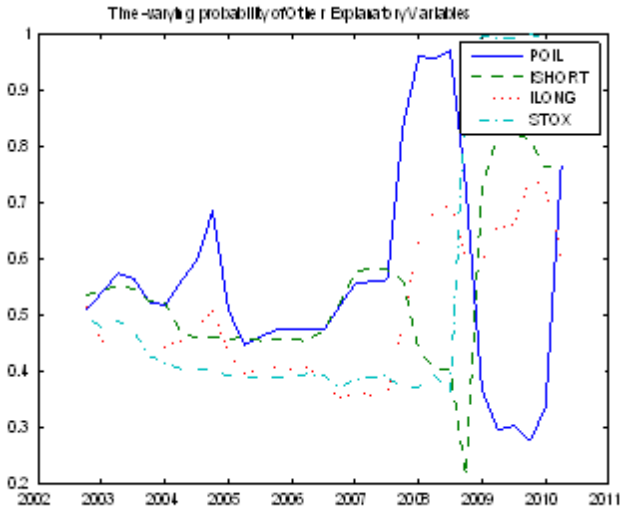

Figure 5

An examination of Figures 6 and 7 reinforces these findings relating to the coefficients on $\pi^e_{t|t-4}$ and $\pi^e_{t+4|t}$. These plot the posterior mean and +/-

two posterior standard deviation intervals for each of these coefficients. The posterior mean for the backward looking coefficient is slightly positive but near zero (which is more reasonable than the negative OLS estimate of Table 1). The +/- two posterior standard deviation interval narrows over time as the incorporation of new data allows for more precise estimation. However, at all points in time zero is well within this interval. Until 2008, Figure 7 shows a similar pattern for the forward looking coefficient. But in 2008 the posterior mean abruptly jumps to being approximately one (as we found with the OLS estimate) and the +/- two posterior standard deviation interval does not include zero.

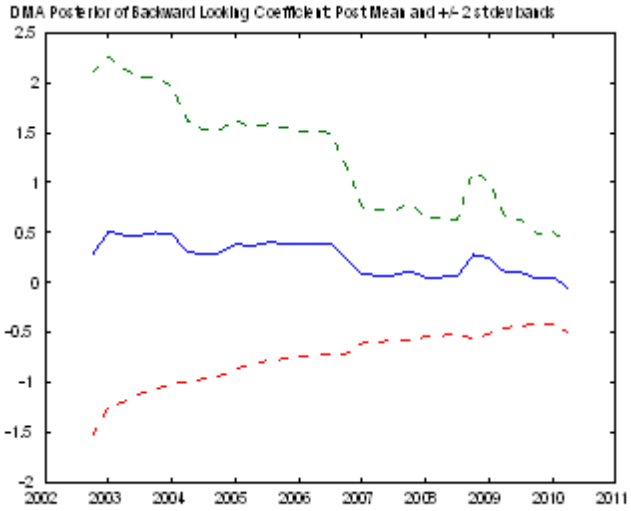

Figure 6

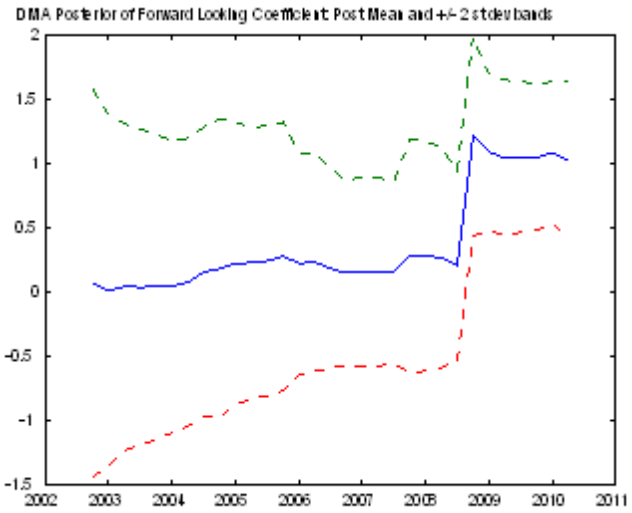

Figure 7

The regression results presented in this sub-section tell a consistent story. It is forward looking expectations which are important in the Phillips curve. Our DMA results also suggest that the importance of forward looking expectations is mostly at the end of our sample, at the time of the financial crisis and ensuing recession.

The change in importance of different regressors provides, in our opinion, a coherent explanation of the determinants of inflation dynamics before and during the crisis. First of all, the oil price shocks affecting prices as of 2006 are clearly visible in the increased importance of the oil variable in the period 2006/2008. In a period of strong and unpredictable shocks, sizeable errors in the agents expectations reduce their inclusion probability. At the same time, and given the more unpredictable inflation developments, the importance of the first lag of inflation increases, leading to models somewhat closer to a random walk. When the global crisis hits in 2008, this time in the context of coherent macroeconomic developments, the regressors representing fundamentals seem to be again more important, either as direct explanatory variables (short and long interest rates and financial markets developments) or through forward-looking expectations, which are computed by the agents keeping it into account a much wider information set than the one included in our regressions. At all times, the unemployment rate remains the most

important and reliable indicator of economic conditions, confirming the argument discussed by Stock and Watson (1999) and Amisano and Giacomini (2007) that the unemployment rate has strong advantages in terms of forecast accuracy in a model for inflation.

We next turn to VAR-based methods as a way of obtaining a deeper understanding of this process.

4.4 Multivariate VAR-based Methods

For the reasons outlined in Section 4.1 we are also interested in modelling inflation and an expectations variable jointly. To do this, we use DMA methods applied to VARs with exogenous explanatory variables:

$$y_t = a + \sum_{j=1}^{p} A_j y_{t-j} + BX_t + \varepsilon_t, \tag{7}$$

where $y_t = \left(\pi^e_{t+4|t}, \pi_t\right)'$ or $y_t = \left(\pi^e_{t|t-4}, \pi_t\right)'$ and X_t is the vector of eight explanatory variables listed in Section 4.1, B is a matrix of regression coefficients and ε_t is independent $N(0, \Sigma)$. To our knowledge, DMA has never been used with VARs. Accordingly, the development of some new econometric techniques is required. These are explained in the appendix in detail. It suffices to note here that the DMA is done in a similar way as with the regression model. Equation (7) defines an unrestricted VAR with exogenous variables X_t which appear in every equation. Our model space contains restricted versions of this model defined by whether explanatory variables are included/excluded. In terms of X_t, we allow for this to be done one equation at a time. That is, it would be simplest if we only considered restricted models defined by whether an element of X_t is included/or excluded in all equations. We do not adopt this simple strategy, but consider the full possible set of restricted models (i.e. we also allow for restricted models where an explanatory variable is included in one equation, but not the other). However (to keep the computational burden manageable), we do not do DMA on the VAR part of the model (i.e. all models include the intercept and VAR lags). As with our regression methods, we use a noninformative prior for the parameters of each model, set $\alpha = 0.99$ and initialize using OLS results (in a VAR including all the predictors) using the initial three years of data, before doing DMA on the remainder of the sample. DMA is based on predictive

likelihoods (see Section 4.2 or the appendix) which can be used as a model selection device. We use the product of the predictive likelihoods over all periods (except for the initial three years) to choose lag length and found evidence in favor of $p = 1$ for both definitions of the dependent variables and we adopt this choice in this sub-section. Given the need to include one lag of the SPF variables in the VAR, we estimate using one less observation than in our regressions and the data begins in 2000Q1.

To motivate our VAR specifications, note that structural VARs are usually written as:

$$C_0 y_t = c_0 + \sum_{j=1}^{p} C_j y_{t-j} + DX_t + u_t$$

where u_t is i.i.d. $N(0, I)$. A choice of C_0 identifies the model and allows for a structural impulse response analysis. If $y_t = \left(\pi^e_{t+4|t}, \pi_t\right)'$ then a specification motivated by the hybrid New Keynesian Phillips curve in (2) arises if C_0 is lower triangular. That is, the second equation in the structural VAR is a Phillips curve relationship between π_t and $\pi^e_{t+h|t}$ (including additional lags as in the hybrid version and extra exogenous variables as in the generalized Phillips curve). The fact that $\pi^e_{t+4|t}$ is ordered first in the identification scheme is motivated by the fact that the SPF is carried out in the first month of the quarter (so that π_t will not have been observed when forecasts $\pi^e_{t+4|t}$ are made). Clark and Davig (2008) adopt the same identification scheme in an application involving US surveys. In a similar fashion, when investigating the Neoclassical Phillips curve we use a VAR with $y_t = \left(\pi^e_{t|t-4}, \pi_t\right)'$, but the same identification scheme (i.e. C_0 is lower triangular with $\pi^e_{t|t-4}$ ordered first) is used.

We will refer to the VAR with dependent variables $\pi^e_{t+4|t}$ and π_t the New Keynesian VAR and the VAR with dependent variables $\pi^e_{t|t-4}$ and π_t the Neoclassical VAR. Given our recursive estimation procedure, impulse responses will change over time. Accordingly, we have different impulse responses in every time period. Figures 8 through 11 plot the one year impulse responses for the New Keynesian VAR against time. Figures 12 through 14 do the same for the Neoclassical VAR.

The most important Figures are 9 and 13 which show how inflation responds to forward and backward inflation expectations, respectively. These figures show the same pattern we found when doing DMA in the single equa-

tion regression. As implied by the New Keynesian Phillips curve, forward looking inflation expectations do have an important impact on inflation. However this impact only becomes substantial after the financial crisis begins in 2008. Backward looking inflation expectations never have an appreciable impact on inflation.

Figures 10 and 14 allow us to investigate the reverse effect: whether inflation has an impact on inflation expectations. Remember that our structural VAR is identified by assuming inflation has no immediate impact of inflation expectations. The New Keynesian impulse response in Figure 10 is basically zero in all time periods. The same finding is obtained using the Neoclassical VAR with one exception. The exception is in early 2009 where this impulse response briefly becomes negative. It is worth noting that 2009 was the most unusual year in our sample. Inflation was less than one percent in every quarter of the year and was actually negative in 2009Q3. In no other quarter in our sample was inflation ever less than one percent. The backward looking expectations were way off: the one year ahead forecasts made in the four quarters of 2008 were 2.13, 2.4, 1.95 and 1.44, respectively. Thus, in 2008 the SPF forecasters only gradually lowered their inflation forecasts and, in 2009, they turned out to have been much too high. The forward forecasts were much better. That is, in the four quarters of 2009 the one year-ahead forecasts were 1.17, 1.21, 1.29 and 1.37 which were much closer to both the contemporaneous 2009 inflation realizations and the 2010 realizations.

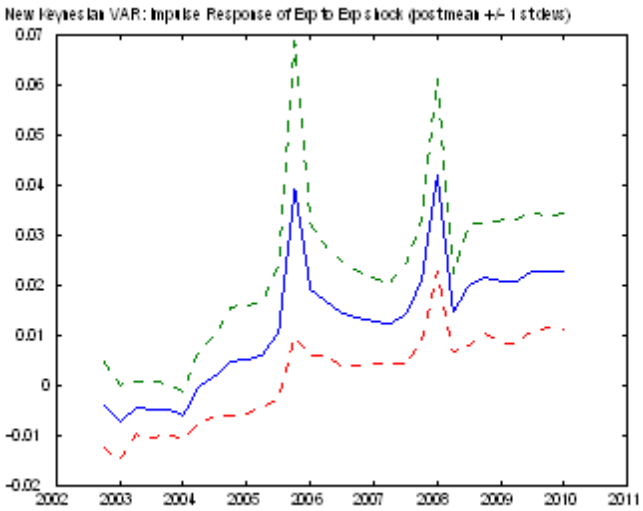

Figure 8

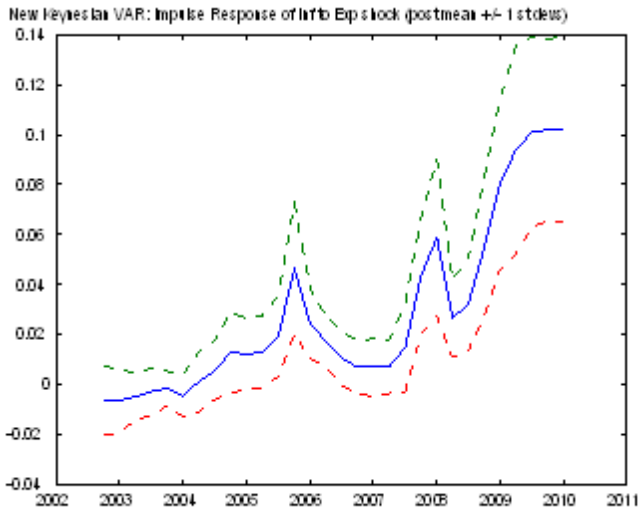

Figure 9

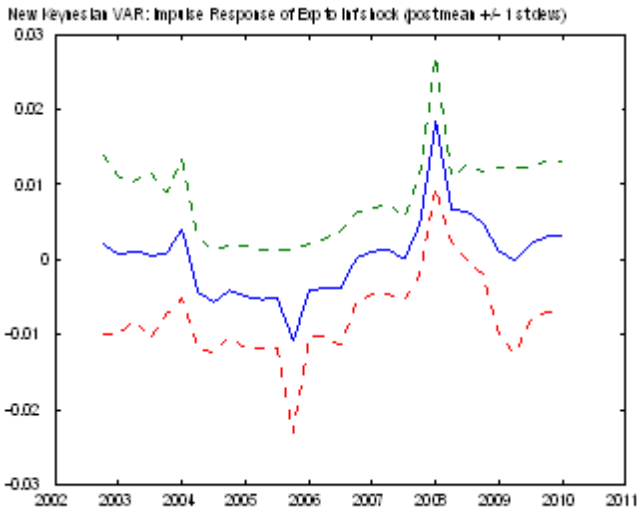

Figure 10

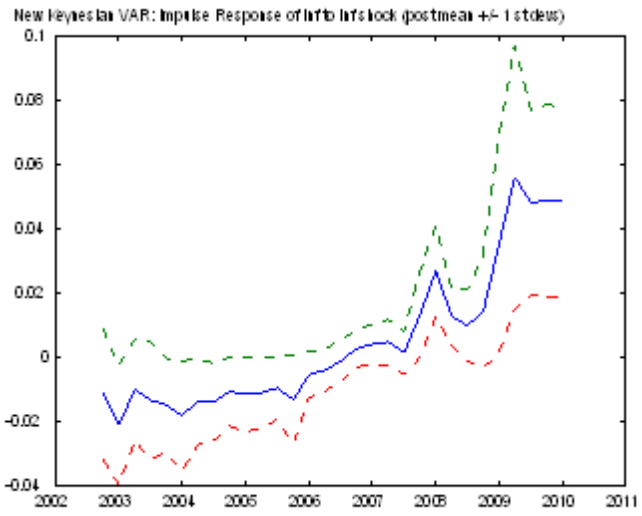

Figure 11

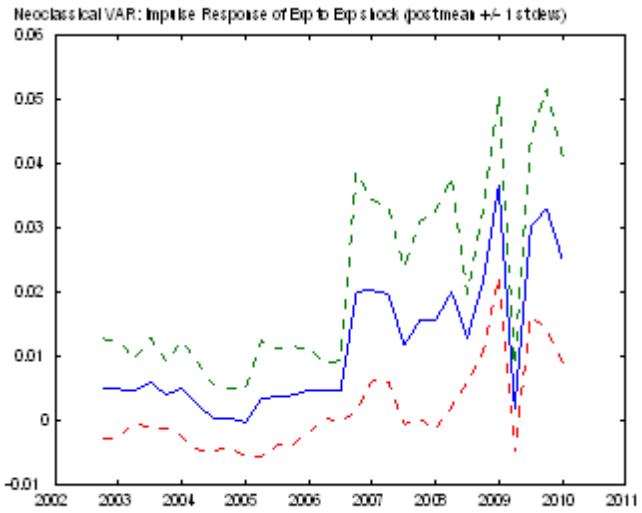

Figure 12

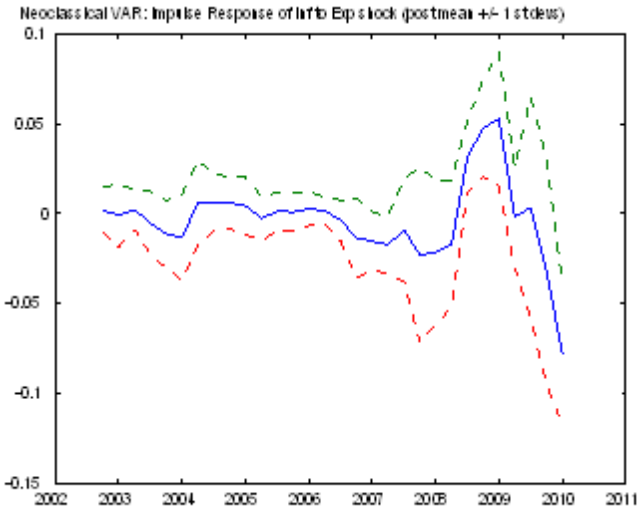

Figure 13

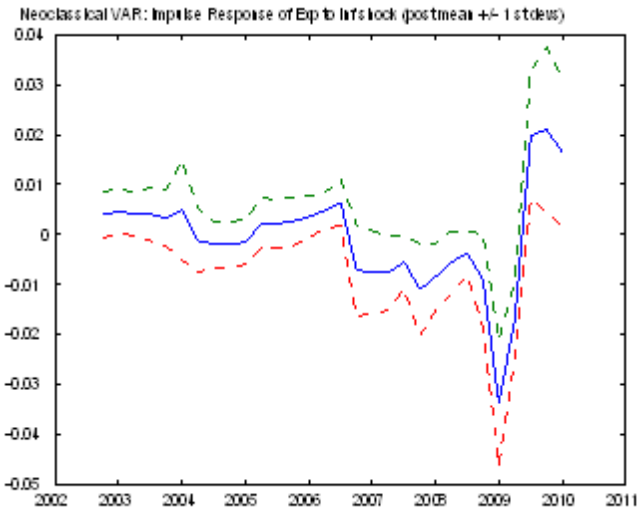

Figure 14

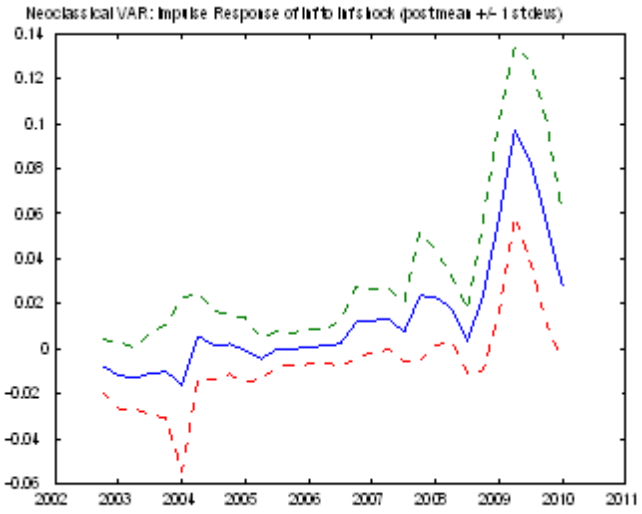

Figure 15

The New Keynesian and Neoclassical VAR cannot be directly compared using standard model comparison methods since they have different depen-

dent variables. However, if we consider only the inflation equation, then the models are comparison. For the New Keynesian and Neoclassical VARs, the sums of log predictive likelihoods (obtained from the DMA exercise) for only inflation equation are 116.87 and 56.10, respectively. This provides strong evidence that the New Keynesian VAR is much better at explaining inflation than the Neoclassical one.

To gain more insight on when to superior performance of the New Keynesian Phillips curve is occurring, Figure 16 present cumulative sums of log predictive likelihoods for the inflation equation. In most of the sample the New Keynesian VAR clearly per- forms much better than the Neoclassical VAR. As an exception, around the middle of 2008, when commodity prices were more important than expectations, the Neoclassical VAR performed almost as well than the New Keynesian VAR. However, after the financial crisis begins, the New Keynesian VAR pulls ahead and obtains a much higher log predictive likelihood.

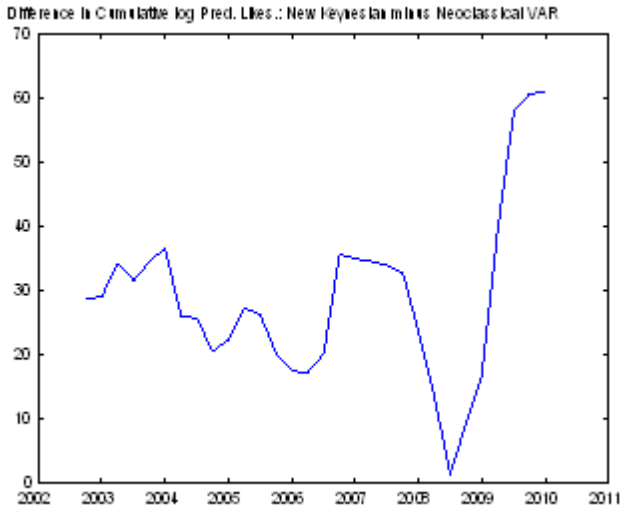

Figure 16

5 Conclusions

In this paper, we have used the forecasts from the ECB's SPF as a proxy for inflation expectations. We have used this proxy in several different empirical exercises as a way of investigating the relationship between inflation and inflation expectations in the euro area. We find that forward looking inflation expectations do have an important impact on current inflation (but only late in the sample period), thus supporting the theory which underlies the New Keynesian Phillips curve. However, we find that backward looking inflation expectations do not have an impact on current inflation.

The SPF is a relatively new data source and, thus, econometric methods designed to deal with the short data span were required. However, our empirical exercises suggest that, even with this short data span, there is information in the SPF than can be productively used by macroeconomists and policymakers.

References

Amisano G. and Giacomini, R. (2007). "Comparing density forecasts via weighted likelihood ratio tests," *Journal of Business and Economic Statistics*, 25, 177-190.

Ang, A. Bekaert, G. and Wei, M. (2007). "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics* 54, 1163-1212.

Ball, L. (1992). "Why does higher inflation raise inflation uncertainty?" *Journal of Monetary Economics*, 29, 371–378.

Bomberger, W. A. (1996). "Disagreement as a measure of uncertainty," *Journal of Money, Credit and Banking*, Blackwell Publishing, 28, 381-92.

Bowles, C., Friz, R., Genre, V., Kenny, G., Meyler, A. and Rautanen, T. (2007). "The ECB survey of professional forecasters (SPF) - A review after eight years' experience," Occasional Paper Series 59, European Central Bank.

Calvo, G. (1983). "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383-398.

Caporale, G., Onorante, L. and Paesani, P. (2010). "Inflation and inflation uncertainty in the euro area," Working Paper Series 1229, European Central Bank.

Clark, T. and Davig, T. (2008). "An empirical assessment of the relationships among inflation and short- and long-term expectations," Research Working Paper 08-05, Federal Reserve Bank of Kansas City.

Cukierman, A. and Meltzer, A. (1986). "A theory of ambiguity, credibility and inflation under discretion and asymmetric information," *Econometrica*, 54, 1099-1128.

D'Amico, S. and Orphanides, A. (2008). "Uncertainty and disagreement in economic forecasting," Finance and Economics Discussion Series 2008-56, Board of Governors of the Federal Reserve System.

Davis G., and Kanago, B. (1996). "On measuring the effect of inflation uncertainty on real GNP growth," *Oxford Economic Papers*, 48, 163-175.

Friedman M. (1968). "The role of monetary policy," *American Economic Review*, 58, 1–17.

Friedman, M. (1977). "Nobel lecture: inflation and unemployment," *Journal of Political Economy*, 85, 451–472.

Gali, J. and Gertler, M. (1999). "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 44, 195-222.

Garcia, J. (2003). "An introduction to the ECBs survey of professional forecasters," Occasional Paper Series 8, European Central Bank.

García, J. and A. Manzanares, A. (2007). "What can probability forecasts tell us about inflation risks?", ECB Working Paper 825.

Giordani, P. and Söderlind, P. (2003). "Inflation forecast uncertainty," *European Economic Review*, 47, 1037-1059.

Hayford, M. D. (2000). "Inflation uncertainty, unemployment uncertainty and economic activity", *Journal of Macroeconomics* 22, 315-329.

Kleibergen, F. and Mavroeidis, S. (2009). "Weak instrument robust tests in GMM and the new Keynesian Phillips curve," *Journal of Business and Economic Statistics*, 27, 293-311.

Koop, G. (2003). Bayesian Econometrics. Chichester: Wiley.

Koop, G. and Korobilis, D. (2009a). "Forecasting inflation using dynamic model averaging," Working Paper Series 34_09, Rimini Centre for Economic Analysis.

Koop, G. and Korobilis, D. (2009b). "Bayesian multivariate time series methods for empirical macroeconomics," *Foundations and Trends in Econometrics*, 3, 267–358.

Lahiri, K., Teigland, C. and Zaporowski, M. (1988). "Interest rates and the subjective probability distribution of inflation forecasts," *Journal of Money, Credit and Banking*, 20, 233-248.

Leduc, S, Sill, K. and Stark, T. (2007). "Self-fulfilling expectations and the inflation of the 1970s: evidence from the Livingston Survey," *Journal of Monetary Economics*, 54, 433–59.

Ma, A. (2002). "GMM estimation of the new Phillips curve," *Economics Letters*, 76, 411-417.

Mankiw, N. G. and Reis, R. (2002). "Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve," *Quarterly Journal of Economics*, 117, 1295-1328.

Mavroeidis, S. (2005). "Identification issues in forward-looking models estimated by GMM, with an application to the Phillips curve," *Journal of Money, Credit and Banking*, 37, 421-48.

Mehra, Y. and Herrington, C. (2008). "On the sources of movements in inflation expectations: A few insights from a VAR Model," *Federal Reserve Bank of Richmond Economic Quarterly*, 94, 121-146.

Paloviita , M. (2005). "Comparing alternative Phillips curve specifications: European results with survey-based expectations," Research Discussion Paper 22/2005, Bank of Finland.

Paloviita, M. and Virén, M. (2005). "The role of expectations in the inflation process in the euro area," Research Discussion Papers 6/2005, Bank of Finland.

Pesaran, M. H. (1987). The Limits to Rational Expectations, Oxford: Blackwell Publishers.

Phelps, E. (1968). "Money-wage dynamics and labor-market equilibrium," *Journal of Political Economy* 76, 678–711.

Pourgerami, A. and Maskus, K. (1987). "The effects of inflation on the predictability of price changes in Latin America: some estimates and policy implications," *World Development*, 15, 287–290.

Raftery, A., Karny, M. and Ettler, P., (2010). "Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill," *Technometrics*, 52, 52-66.

Sims, C. (2002). "The role of models and probabilities in the monetary policy process," *Brookings Papers on Economic Activity*, 2, 1-62.

Stock, J. and Watson, M. (1999). "Forecasting inflation," *Journal of Monetary Economics* 44, 293-335.

Stock, J. and Watson, M. (2008). "Phillips Curve Inflation Forecasts," NBER Working Paper No. 14322.

Tay, A. and Wallis, K. (2000). "Density forecasting: a survey," *Journal of Forecasting*, 19, 235-254.

Ungar, M. and Zilberfarb, B. (1993). "Inflation and its unpredictability – theory and empirical evidence," *Journal of Money, Credit, and Banking*, 25, 709-720.

Wallis, K. (2003), "Chi-squared tests of interval and density forecast, and the Bank of England's fan charts", *International Journal of Forecasting*, 19, 165-175.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

Appendix: Dynamic Model Averaging in VARs

In this paper, we consider M different models which are all VARs with k_x exogenous explanatory variables. Formally, each model can be written, for $t = 1, .., T$ and $m = 1, .., M$, as

$$y_t = Z_t^{(m)} \theta^{(m)} + \varepsilon_t^{(m)}, \tag{8}$$

where y_t is an $N \times 1$ vector of dependent variables, $\varepsilon_t^{(m)} \overset{ind}{\sim} N\left(0, \Sigma^{(m)}\right)$,

$$Z_t^{(m)} = \begin{pmatrix} z_{1t}^{(m)} & 0 & \cdots & 0 \\ 0 & z_{2t}^{(m)} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & z_{Nt}^{(m)} \end{pmatrix},$$

$z_{nt}^{(m)} = \left(1, y_{t-1}', .., y_{t-p}', x_{nt}^{(m)}\right)$ for $n = 1, ..N$ and $x_{nt}^{(m)}$ is a row vector of explanatory variables in the n^{th} equation of model m and h is the forecast horizon. $x_{nt}^{(m)}$ will be a subset of k_x potential explanatory variables. In the paper, these are the eight explanatory variables listed in Section 4. 1.

Note that our models all have a basic VAR(p) structure, but allow for different for different equations to have different explanatory variables among those included . Allowing for each of the k_x potential explanatory variables to be included in each of N equations means our model space contains $M = 2^{Nk_x}$ models.

In this paper, we wish to investigate whether the model used for explaining y_t is changing over time. Thus, we use a restricted version of the dynamic model averaging (DMA) methods developed by Raftery, Karny and Ettler (2010) and used by Koop and Korobilis (2009a). Ours is a restricted version in the sense that Raftery et al. (2010) allow for time variation in the coefficients (i.e. they replace our constant $\theta^{(m)}$ by a state equation $\theta_{t+1}^{(m)} = \theta_t^{(m)} + \eta_t^{(m)}$). Over our relatively short data span, estimating time variation in coefficients seemed too much to ask from our short data set. We estimate our model recursively which allows for some time variation in the coefficients and to some extent model change can be a substitute for parameter change. Note, though, that previous work with DMA has only been for the regression case ($N = 1$) and our use of DMA methods with VARs pushes DMA in a new direction.

To two key components required by DMA are: i) formula for Bayesian inference in each model and, ii) formula for the predictive density $p_k\left(y_t | y^{t-1}\right)$

so as to calculate (4). With these, recursive updating using (4) and (5) can be used to produce $q_{t|t,m}$ for $m = 1, .., M$ exactly as for regression case.

A slight complication arises since our models are restricted VARs and, hence, an analytical formula for posterior and predictive density does not exist and MCMC methods are typically required (see, e.g., Koop and Korobilis, 2009b). Given the computational demands required by working with $M = 2^{Nk_x}$ models, it is not possible to use MCMC methods for evaluating each model. Hence, we use approximate methods. We replace Σ by the usual OLS estimate. An analytical formula for the posterior of the VAR coefficients, conditional on an estimate of Σ, is available (see page 284 of Koop and Korobilis, 2009b, for the exact formula). The predictive density (conditional on Σ) is a Normal distribution with mean and covariance taking the formula on page 278 of Koop and Korobilis (2009b).

In summary, this appendix defines our model space which is a set of restricted VARs. DMA is done as described in the body of the paper, however involving VARs and relevant formulae for the predictive densities. These details have been provided in this appendix.