Adaptive Learning and Survey Expectations of Inflation

Sergey Slobodyan*and Raf Wouters
CERGE-EI and National Bank of Belgium

preliminary draft (please do not circulate without permission)

June 17, 2016

Abstract

The use of survey information on inflation expectations as observable in a DSGE model can refine substantially the identification of the shocks that drive the inflation process. An optimal integration of the survey information improves the model forecast for both inflation and the other macrovariables. Models with expectations based on an Adaptive Learning setup can exploit the survey information more efficiently than their Rational Expectations counterparts. The resulting time-variation in the perceived inflation target, persistence and shock sensitivity provides a rich description of the joint dynamics in realized and expected inflation.

JEL Classification: C5, D84, E3

Keywords: Inflation, Expectations, Survey data, Adaptive Learning, DSGE models

*This work is supported by the EU 7th framework collaborative project "Integrated Macro-Financial Modelling for Robust Policy Design (MACFINROBODS)" grant no. 612796.

The objective of this paper is to illustrate how the observation of survey evidence on future inflation expectations - as measured by the Survey of Professional Forecasters - affects the estimation outcomes of the standard new-Keynesian DSGE model. After documenting that the survey expectations deviate substantially from the implied model expectations in particular during periods with important trend changes in the inflation dynamics, we propose a simple re-specification of the price and wage mark-up shocks in the Smets Wouters model that allows to reconcile the survey and the model forecasts. In particular, the survey data on inflation expectations help to identify the innovations in the persistent component of the inflation process: these innovations constitute only a small fraction of the high frequency volatility in inflation and are therefore hard to distinguish without additional and timely information that is summarized in the survey expectations. By exploiting efficiently this information and making the model expectations process more consistent with the survey information, we can substantially improve the fit and the forecasts of the model. Models with expectations based on an Adaptive Learning setup can exploit the survey information more efficiently than their Rational Expectations counterparts. The resulting time-variation in the perceived inflation target, persistence and shock sensitivity provides a rich description of the dynamics in realized and expected inflation.

Inflation expectations have become the most important indicator for monetary policy. Since the financial crisis, realised inflation has been below target in most advanced countries. Recently, declining oil and commodity prices have further pushed headline inflation down and caused repeated negative surprises in the inflation process. In this situation, monetary policy attention has been concentrated on inflation expectations that are observed in surveys or distilled from financial yields. Policy makers stress that it is absolutely necessary that the inflation expectations remain anchored around the long run inflation objective. The following quote from a speech by C. Evans (2016) illustrate the the crucial role of inflation expectations for monetary policy makers: "However, there are some downside risks to this forecast. We might see further declines in energy prices or greater appreciation of the dollar. In addition, undershooting our 2 percent inflation target for as long as we have invites the risk of the public beginning to expect persistently low inflation in the future. If this mindset becomes embedded in decisions regarding wages and prices, then getting inflation back to 2 percent will be that much more difficult. Here, I find it troubling that the compensation for prospective inflation built into a number of financial market asset prices has drifted down considerably over the past two years. More recently, some survey-based measures of inflation expectations, which had previously seemed unmovable, have also edged down. So to achieve our inflation target — and to provide a buffer against downside risks — it is appropriate that we follow a gradual path to policy normalization."

The behaviour of inflation expectations is also central in the criticism of the low interest rate policies, such as e.g. Cochrane (2015), arguing that prolonged low nominal interest rates and the central bank promises to maintain such policy, may bring about lower inflation rather than higher. Evans et al (2016) and Garcia Schmidt and Woodford (2015) explain how these outcomes depends on the way agents formulate their expectations and perceive the policy reaction function. Understanding the way inflation expectations are formed and how they develop over time is crucial for evaluating the risk of convergence to a low inflation-low interest rate steady state. Survey expectations might contain relevant signals to interpret how agents formulate their beliefs and which solution path they select.

Inflation expectations also play a crucial role in modern macro models: the forward looking New Keynesian Phillips curve is the central equation in the monetary DSGE models. Inflation expectations drive actual price and wage setting and via the expected real return and the intertemporal substitution, they have also the potential to steer aggregate demand in particular when the nominal rate is constraint by the zero lower bound. Recently, inflation expectations were also crucial to understand the behaviour of inflation over the Great recession and the recent recovery period. Gorodnichenko and Coibon (2015) suggest that stability in expectations, in particular on the consumer side, was crucial to understand the relative stable inflation realisations during the Great recession. As consumers' inflation expectations are more responsive to oil prices than those of professional forecasters, the increase in oil prices between 2009 and 2012 contributed to prevent the onset of deflationary dynamics. Del Negro, Giannoni and Schorfheide (2015) illustrate how a high degree of price stickiness is necessary in the DSGE model to understand the inflation dynamics during the Great Recession period. As the expected real marginal costs over a long forward horizon is relative stable, downward price pressure remained modest during the major downturn.

Survey expectations on inflation are very informative data. The review of various inflation forecasting models and survey data by Ang, Bekaert and Wei (2007) has documented the superior forecasting performance of the survey expectations for inflation. The survey expectations probably reflect a large amount of information that is processed in an efficient way with sufficient flexibility to adjust over time. Survey expectations about future inflation have the advantage that they provide unambiguous information on agents' forecasts. Financial market evidence on inflation compensation on the other hand requires further processing to correct for risk premiums. In practice, central bank forecasts combine model forecasts and judgemental forecasts with the latter heavily influenced by survey or market evidence. By integrating survey evidence in the standard dataset that is processed by the model, the need for additional judgemental adjustments to the model forecast will become less pressing.

So, there are many important reasons to include data on inflation expectations in the standard datasets on which our models are estimated. The dynamics of these inflation expectations should be analysed together with realized inflation data to pin down more precisely the transmission mechanism of the various shocks. The overall information should result in a consistent estimate of the state of the economy and the inflation expectation process in particular. Rational or model-consistent expectations is one hypothesis for explaining the way expectations are formed, but this hypothesis should be evaluated against alternatives that allow for more flexibility and that provide more insight in how the agents formulate their beliefs and how they adjust them when confronted with new data and changing environments.

We therefore include the SPF inflation expectations in our model dataset. By treating the survey expectations as additional observables, we can interpret the information in these survey expectations in terms of the fundamental shocks. The surveys contain news about future innovations that are not readily distilled from other contemporaneously observed macro data. The overall precision of the model estimate and the filtering of the underlying shocks should therefore improve with the observation of this information.

Given the strong arguments for introducing survey data in the information set on which our models are estimated, relatively few papers have actually implemented this approach.[1] A paper that comes closest to ours is Ormeno and Molnar (2015). These authors use survey data on inflation expectations as an observable for the model expectations via an additional measurement equation. Their results illustrate that the survey contains information not present in macro data and this can improve the model forecast. They also show how an adaptive learning approach based on small forecasting models is more flexible to exploit the information than a fully rational expectations model. But the paper does not explain what type of information is revealed by the survey and what changes in the model specification can optimize the integration of the survey data in the model. Eusepi and Del Negro (2011) introduce inflation target shocks in their RE-model to improve the match between model and survey forecasts. This exogenous shock captures the additional information provided by the survey and transmits this to the model forecast. While improving the model fit, the approach does not succeed in completely bridging the gap between model and survey forecasts. Similarly, De Graeve et al (2009) illustrate how an inflation target shock is necessary to match inflation expectations in the yield structure. Del Negro and Schorfheide (2013) use long run inflation expectations as the observable and an inflation target shock as the modelling device. Eusepi et al (2015), on the other hand, use short term survey expectations and learning about the long run inflation target to model the inflation expectations. These examples illustrate how rational expectation models typically resort to exogenous shocks in the inflation target to match the survey expectations while adaptive learning models can explain the long run drift in expectations by the

[1] There is broader literature that investigates the what type of price setting models are consistent with the inflation expectations in survey data: f.e. models with sticky information or heterogeneous beliefs (Mankiw and Reis 2002, Branch 2007).

updating of agents expectation process and their beliefs about the inflation target in particular. We will review and compare these alternative approaches systematically and come up with an optimal modelling device. We will concentrate on the nature of the shocks that drive the expectation as in Milani and Rajbhandari (2012): they use survey data on various macrovariables to identify a set of news shocks in addition to the contemporaneous innovations and show that these news shocks explain a major part of business cycle, but they do not focus on inflation.[2]

The structure of the paper is as follows. First, we document the discrepancy between inflation expectations in the surveys and the expectations implied in standard macromodels. We also illustrate that including survey expectations in the model by simply adding a measurement equation is not sufficient. In section two, we explain how the introduction of two markup shocks, one i.i.d. shock and one persistent shock, is extremely helpful for an efficient integration of the survey data in the model. We show this first in a standard Rational Expectations model and discuss the remaining issues in this context. Then we present briefly our Adaptive Learning approach and show how the updating of beliefs accounts well for the time-varying properties of the joint dynamics in realised and expected inflation. Finally, we illustrate the robustness of our results.

1 Survey expectations versus model expectations

First, we document that the survey expectations deviate substantially from the expectations that are implicitely present in the DSGE model. This applies for models that are estimated with rational or model-consistent expectations (RE) as in Smets and Wouters (2007) and for models estimated under adaptive learning (AL) as developed in Slobodyan and Wouters (2012).[3] This model is using an adaptive learning setup in which agents update their perceived forecasting models over time as new data become available with a Kalman filter learning scheme. In SW2012, it was shown that beliefs based on simple AR(2) forecasting models capture well the time-varying persistence in the inflation process. We document the difference between the various forecasts: we plot the forecasts implied by these models against the SPF forecasts and we document the statistic properties of the various forecasts errors. The SPF typically outperforms the model forecasts. Then, we re-estimate the models using the Survey expectations as observable for the model expectations and allowing for measurement error in the observation equations. This results in substantial and systematic

[2] See Monti (2010) and Smets, Warne and Wouters (2014) for previous examples on how survey information can be related to structural shocks.

[3] We refer to the original articles for the detailed model specification and estimation results. We provide more information on the learning setup in section 3.

measurement errors and the model forecasts are only marginally improved in this approach. Clearly, the original model specification is missing the flexibility to exploit efficiently the information that is available in the survey forecasts

1.1 Comparing model expectations and SPF forecasts at different horizons.

In this context, it is important to note that we re-estimated our models on real-time data. Including SPF-forecasts in the model requires to specify the model in real-time so that model expectations are based on information that was also available for survey participants in real-time. As illustrated in Table 1, over the sample since 1971q1, the revision in the second release for the inflation rate in the quarter-to-quarter GDP-deflator relative to the first release has a standard deviation of 0.12 and the final revision has a standard deviation of 0.24. The magnitude of these revisions is of the same order of magnitude than the forecast error in the SPF which has a standard error of 0.27 for the one quarter ahead forecast. The magnitude of the data revisions is important and therefore this real-time data issue cannot be ignored when including survey forecasts in the model.

Table 1: Statistical properties of the inflation revisions and SPF forecasts errors

	1971q1-2015q3			1996q1-2015q3		
Revisions	bias	mad	rmse	bias	mad	rmse
π_r2_t to π_r1_t	-0.03	0.08	0.11	-0.02	0.05	0.07
π_rf_t to π_r1_t	-0.02	0.17	0.22	-0.04	0.12	0.16
π_rf_t to π_r2_t	0.01	0.16	0.23	-0.02	0.10	0.13
SPF statistics						
$\pi\_SPF_{t+1\|t} - \pi\_r1_{t+1}$	0.03	0.21	0.26	0.03	0.17	0.21
$\pi\_SPF_{t+1\|t} - \pi\_r2_{t+1}$	0.01	0.21	0.27	0.01	0.16	0.19
$\pi\_SPF_{t+1\|t} - \pi\_rf_{t+1}$	0.01	0.18	0.24	-0.01	0.15	0.20
SPF for longer horizons						
$\pi\_SPF_{t+2\|t} - \pi\_r1_{t+2}$				0.04	0.19	0.23
$\pi\_SPF_{t+3\|t} - \pi\_r1_{t+3}$				0.07	0.19	0.23
$\pi\_SPF_{t+4\|t} - \pi\_r1_{t+4}$				0.07	0.21	0.25

Note: π_r1 is the first release for inflation, π_r2 the second release and π_rf the final release. π_SPF is the SPF forecast. Period 1971q1-2015q3 starts with the period for which the inflation and GDP forecasts are available in the Survey of Professional Forecasters. The sample 1996q1-2015q3 is the typical period that we use in the out-of-sample model forecast tests presented in the paper.

The models are re-estimated with real time data for inflation in the GDP-deflator and the GDP growth rates.[4] For the other five observables (real growth in consumption, investment and real wages, total hours worked and the Fed-funds rate) we still use the final data because there are no exact real-time counterparts available and survey forecasts for these components starts later as well.[5] Agents in the model are assumed to observe the first and the second release of these series: the second release is considered here as the 'true' measure for inflation and GDP growth. The first release is assumed to contain a simple i.i.d. measurement error ξ. For inflation the measurement equations become:

$$\begin{aligned} \pi\_r1_t &= \overline{\pi} + \widetilde{\pi}_t + \xi_t^{\pi r} \\ \pi\_r2_t &= \overline{\pi} + \widetilde{\pi}_{t-1} \end{aligned}$$

and similarly for GDP growth:

$$\begin{aligned} dy\_r1_t &= \overline{\gamma} + \widetilde{y}_t - \widetilde{y}_{t-1} + \xi_t^{yr} \\ dy\_r2_t &= \overline{\gamma} + \widetilde{y}_t - \widetilde{y}_{t-1} \end{aligned}$$

When the agents in the model form their expectations for quarter t+1, the information set will include the first release of the data for quarter t and the second release for quarter t-1. This timing assumption is an approximation for the information structure available for the SPF participants: survey forecasts for t+1 is collected after the first release of data for quarter t is published. Of course data processing and publication takes time, and surveys are collected when quarter t+1 is already ongoing, more precisely during the first half of the second month of quarter t+1. That is why the forecast for quarter t+1 is also called the nowcast. Nowcasts can reflect information that became available only after quarter t was realized. This timely nature of the information set available to SPF-participants might contribute to the excellent forecasting performance of the surveys when compared to model forecasts.[6] Table 1 also provide descriptive statistics for the forecasting performance of the SPF. The rmse forecasting error of 0.26 for the complete sample (1971q1-2015q3) and 0.21 for our out-of-sample prediction sample (1996q1-2015q3) will be important benchmarks for the model forecasts later on.

The estimated parameters for the real-time model versions are documented in Table A1. We will refer to these models as the 9obs RE and AL version:

[4] Our exercise is based on inflation expectations as measured by the GDP-deflator series. This choice corresponds with the data used in the original SW2007 and SW2012 models. We test the robustness of our results with CPI and PCE expectations.

[5] Real-time data and SPF data are downloaded from the Philadeplphia Fed web-site https://www.philadelphiafed.org/research-and-data/real-time-center/

[6] In the robustness exercise in section 4, we use an alternative timing assumption to illustrate the sensitivity of our results to this issue.

relative to the original versions with seven observables, these versions include two additional real-time data series as observables. The estimated parameter values are relative comparable to the original SW2007 and SW2012 models despite the use of real-time data and the update in the sample. The estimated standard errors for the measurement errors in the first release data of inflation and output is almost identical under RE and under AL: 0.12 for GDP inflation and 0.19 for GDP growth.

Figure 1 and 2 illustrate the inflation expectations that are implicitly present in these real-time models. The forecasts for t+1 are of particular interest as they appear directly in the first-order conditions that describe the decision rules of the agents. The plots present the projection trajectories at each point in time for the next 4 quarters. These are out-of-sample forecasts as the underlying models are estimated on the datasets that are available at the moment that the forecast is made. Each forecast starts from the last available observation for the first-release and in that way the figure provides also information on the ex-post realisation. In the RE-version, these model forecasts are consistent with agents expectations. In the AL-version, the model forecasts, that are the forecast consistent with the Actual Law of Motion process (ALM-forecasts), are plotted and these are in general not equivalent to the agents' perceived forecasts (PLM forecasts).[7]

The lower panel of the figures display the cumulative deviations between the model and SPF forecasts. Here, we can observe that there are often large deviations between the two forecasts. Both models, but RE-model in particular, tend to predict higher inflation than the SPF-forecast for most of the period between 1996 and 2002, and again from 2004 and 2007. Since the start of the Great recession, both models, and the AL most outspoken, produce lower forecasts than the survey. Also at the end of the sample the models tend to produce forecasts below the SPF. The deviations between model and survey forecasts are very persistent for most of the projection trajectories but also across projections as time proceed

[7] With our solution procedures, ALM forecast are easily rolled forward for longer horizons. Depending on the specification of the belief models, it is less evident how to produce longer horizon PLM forecasts in the models.

Figure 1: Model and SPF forecasts for 1 to 4 quarters ahead versus realized inflation (first release)

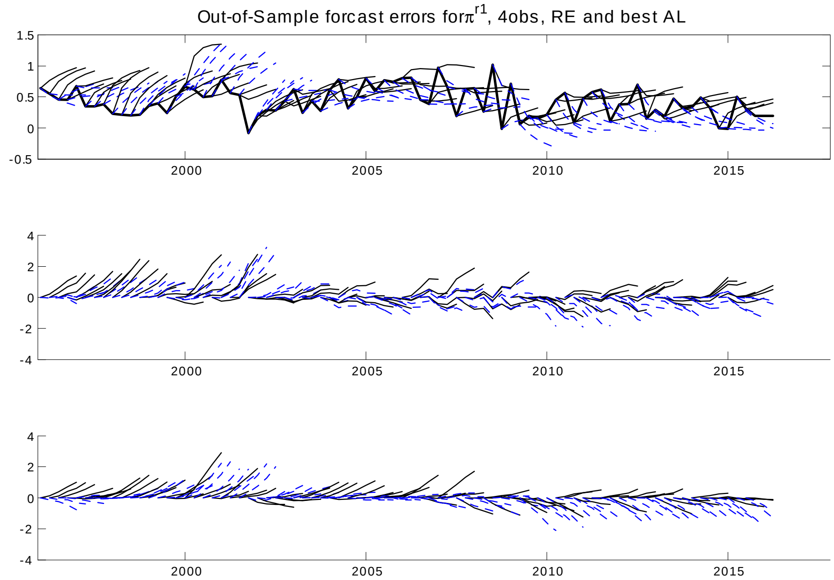

Note: The graph in the middle reproduce the cumulative forecast errors of the two models relative to the realized inflation. The figure at the bottom contains the cumulative deviations of the two model forecasts relative to the SPF-survey forecasts. The black line corresponds with the RE-forecast, the blue dashed line represents the AL-forecast.

The statistics reported in Table 2 further document the inflation forecasting performance of the 9obs-RE-SW2007 model re-estimated with real-time data. The forecast errors of the model are large compared to the SPF forecasts on all three criteria that are considered (bias, amd and rmse). The bias increases systematically with the forecast horizon in the out-of-sample forecasts that covers the more recent sub-period since 1996. The inflation target that is estimated in the RE models is probably biased upward by the first part of the sample. Important at this point is that the RE forecast deviates substantially from the SPF forecast. The standard deviation between the two forecasts for quarter t+1 is 70% of the SPF-forecast error. So the difference between the two forecasts is almost as big as the forecast error itself. Testing the equivalence of the two forecasts by the Diebold-Mariano test clearly rejects the hypothesis that the two forecasts are equivalent for horizon from one to four. The SPF significantly outperforms the model forecast.

Table 2: Statistical properties of the 9obs-RE-SW2007 model forecast errors and comparison to SPF

	1971q1-2015q3			1996q1-2015q3		
t+1 forecast	bias	mad	rmse	bias	mad	rmse
$\pi\_RE_{t+1\|t}-\pi\_r1_{t+1}$	0.02	0.27	0.34	0.06	0.23	0.28
$\pi\_RE_{t+1\|t}-\pi\_r2_{t+1}$				0.05	0.23	0.27
$\pi\_RE_{t+1\|t}-\pi\_rf_{t+1}$				0.02	0.21	0.25
longer horizon						
$\pi\_RE_{t+2\|t}-\pi\_r1_{t+2}$	0.02	0.29	0.39	0.12	0.25	0.31
$\pi\_RE_{t+3\|t}-\pi\_r1_{t+3}$	0.02	0.32	0.41	0.17	0.27	0.35
$\pi\_RE_{t+4\|t}-\pi\_r1_{t+4}$	0.019	0.33	0.45	0.21	0.29	0.36
RE versus SPF	rel.rmse%		DM-test	rel.rmse%		DM-test
$horizon = 1$	70.86		4.61	?		4.46
$horizon = 2$	56.70		3.10			3.07
$horizon = 3$	47.22		2.37			3.28
$horizon = 4$	44.04		1.22			2.29

Note: Rel rmse is defined as $\frac{stdev(\pi_RE_{t+1|t}-\pi_SPF_{t+1|t})}{stdev(\pi_SPF_{t+1|t}-\pi_r1_{t+1})}$x100. DM-test is the Diebold-Mariano test for equal accuracy between RE and SPF forecast.

Table 3 repeats the same statistics for the 9obs-AL-SW2012 model re-estimated with real-time data. Also in this model, the forecast errors for inflation are large and worse than the SPF in terms of amd and rmse. The standard deviation between the two forecasts is up to 78% of the SPF-forecast error. The Diebold-Mariano test indicates the superiority of the SPF forecasts over all forecast horizons considered. The longer forecast horizons deteriorate more under AL than under RE: the structure imposed on the RE forecasts seems to pay off for long term forecasts while the flexibility of the AL-beliefs can become costly for longer term forecasts. These results apply for the ALM-forecasts in the AL-model. But in this model, the PLM forecasts are also relevant. It are these PLM expectations that enter into the agents decision rule when they make the actual price decision. The PLM based on the small forecasting model does a good out-of-sample forecasting job in this model at least compared to the ALM-forecast.

Table 3: Statistical properties of the 9obs-AL-SW2012 model forecasts errors and comparison to SPF

	1971q1-2015q3			1996q1-2015q3		
t+1 horizon	bias	mad	rmse	bias	mad	rmse
$\pi\_AL\_PLM_{t+1\|t}-\pi\_r1_{t+1}$				0.02	0.19	0.26
$\pi\_AL\_ALM_{t+1\|t}-\pi\_r1_{t+1}$	-0.02	0.26	0.35	-0.01	0.23	0.29
$\pi\_AL\_ALM_{t+1\|t}-\pi\_r2_{t+1}$				-0.03	0.22	0.28
$\pi\_AL\_ALM_{t+1\|t}-\pi\_rf_{t+1}$				-0.05	0.21	0.27
longer horizon						
$\pi\_AL\_ALM_{t+2\|t}-\pi\_r1_{t+2}$	-0.03	0.30	0.41	-0.00	0.27	0.33
$\pi\_AL\_ALM_{t+3\|t}-\pi\_r1_{t+3}$	-0.04	0.33	0.47	0.01	0.31	0.37
$\pi\_AL\_ALM_{t+4\|t}-\pi\_r1_{t+4}$	-0.04	0.37	0.53	0.03	0.34	0.42
AL versus SPF	rel.rmse%		DM-test	rel.rmse%		DM-test
$horizon = 1$	78.56		3.94	?		3.95
$horizon = 2$	70.80		2.82			2.10
$horizon = 3$	69.20		2.70			2.21
$horizon = 4$	69.77		2.03			2.14

Note:

1.2 re-estimation including SPF-data as observables with measurement error

Here we present results for the two models re-estimated with the SPF-forecast for t+1 as additional observable. We integrate the SPF survey data as observable for the expected variable as follows:

$$\pi\_f_{t+1|t} = \overline{\pi} + E_t\widetilde{\pi}_{t+1} + \xi_t^{\pi f1} \tag{1}$$

where $\xi_t^{\pi f1}$ is an i.i.d. measurement error (ME) between the observed SPF forecast $\pi\_f_{t+1|t}$ and the model forecast $E_t\widetilde{\pi}_{t+1}$ which is expressed in deviation from the inflation target $\overline{\pi}$.

The estimated parameters for these 10obs-ME-RE and AL version are available in Table A2. The i.i.d. m.e. for SPF expectations has a standard error of 0.1773 under RE and 0.1810 under AL-AR(2). Including the survey data in the model in this elementary way does change some of the estimated parameters and shocks. The most striking changes are the higher degree of nominal stickiness. In particular, the Calvo-probability for prices is significantly higher with observations on the SPF expectations both under RE and under AL.[8]

[8] Smith (2009) and Nunes (2010) have discussed the information content of SPF survey data as indicators for agents expectations in the estimation of the New Keynesian Phillips-Curve. These econometric exercises are performed in a single equation context.

Tables 4 and 5 illustrate that the models with the survey observable perform better on all statistics related to the inflation forecast. By minimizing the m.e. on the SPF forecasts, the forecasts are forced to resembles more to the survey forecasts and in this way the forecast performance measured against the ex-post realisations improves also. The model inflation forecasts can clearly benefit from the excellent prediction potential of the survey data.

Table 4: Forecast Statistics for the RE-SW2007 with SPF observable

	1971q1-2015q3			1996q1-2015q3		
t+1 forecast	bias	mad	rmse	bias	mad	rmse
$\pi\_RE_{t+1\|t} - \pi\_r1_{t+1}$	0.02	0.25	0.32	0.05	0.20	0.25
$\pi\_RE_{t+1\|t} - \pi\_r2_{t+1}$				0.04	0.20	0.23
$\pi\_RE_{t+1\|t} - \pi\_rf_{t+1}$				0.01	0.18	0.23
longer horizon...						
$\pi\_RE_{t+2\|t} - \pi\_r1_{t+2}$	0.02	0.28	0.37	0.09	0.22	0.27
$\pi\_RE_{t+3\|t} - \pi\_r1_{t+3}$	0.02	0.30	0.40	0.12	0.24	0.30
$\pi\_RE_{t+4\|t} - \pi\_r1_{t+4}$	0.01	0.32	0.43	0.15	0.26	0.32
RE versus SPF	rel.rmse%		DM-test	rel.rmse%		DM-test
$horizon = 1$	56.64		4.26			3.35
$horizon = 2$	45.16		2.51			2.56
$horizon = 3$	38.21		1.44			2.82
$horizon = 4$	38.09		0.66			1.18

Note:

Table 5: Forecast Statistics for the AL-SW2012 model with SPF observable

	1971q1-2015q3			1996q1-2015q3		
t+1 horizon	bias	mad	rmse	bias	mad	rmse
$\pi\_AL\_PLM_{t+1\|t} - \pi\_r1_{t+1}$				0.06	0.19	0.24
$\pi\_AL\_ALM_{t+1\|t} - \pi\_r1_{t+1}$	0.02	0.26	0.34	0.06	0.20	0.25
$\pi\_AL\_ALM_{t+1\|t} - \pi\_r2_{t+1}$				0.04	0.20	0.24
$\pi\_AL\_ALM_{t+1\|t} - \pi\_rf_{t+1}$				0.02	0.17	0.23
longer horizon						
$\pi\_AL\_ALM_{t+2\|t} - \pi\_r1_{t+2}$	0.02	0.30	0.40	0.09	0.23	0.27
$\pi\_AL\_ALM_{t+3\|t} - \pi\_r1_{t+3}$	0.02	0.33	0.45	0.12	0.23	0.28
$\pi\_AL\_ALM_{t+4\|t} - \pi\_r1_{t+4}$	0.02	0.35	0.49	0.14	0.24	0.29
AL versus SPF	rel.rmse%		DM-test	rel.rmse%		DM-test
$horizon = 1$	63.11		4.53			3.03
$horizon = 2$	57.39		2.57			2.51
$horizon = 3$	50.82		1.85			2.12
$horizon = 4$	50.12		1.73			1.16

Note:

Figure 2 plots the smoothed estimates of the time series for the measurement error in the 10obs-ME-RE and AL models. The measurement error is still substantial, illustrating that the models have a hard time to produce a forecast that resembles the survey forecast. During several episodes, the measurement error is also highly persistent with deviations between the two forecasts repeatedly in the same direction for more than a year.[9] The measurement errors for the AL are relatively large during the first half of the sample, but, on average, smaller during the second subsample. This picture confirms the statistics in reported in Table 5 and Figure 3 with the out-of-sample forecast plots. Compared to the RE-model, the AL model does relatively a good job in the out-of-sample prediction exercise over the period since 1996. The out-of-sample inflation forecasts of the AL model are more in line with the survey forecast according to the relative rmse criteria but their forecast accuracy is still outperformed by the survey according to the DM-test. Also in terms of longer horizon inflation forecasts, the AL-model is now superior to the RE-model while the opposite was holding for the 9obs-models.

Figure 2: Model forecast with SPF observed and measurement error

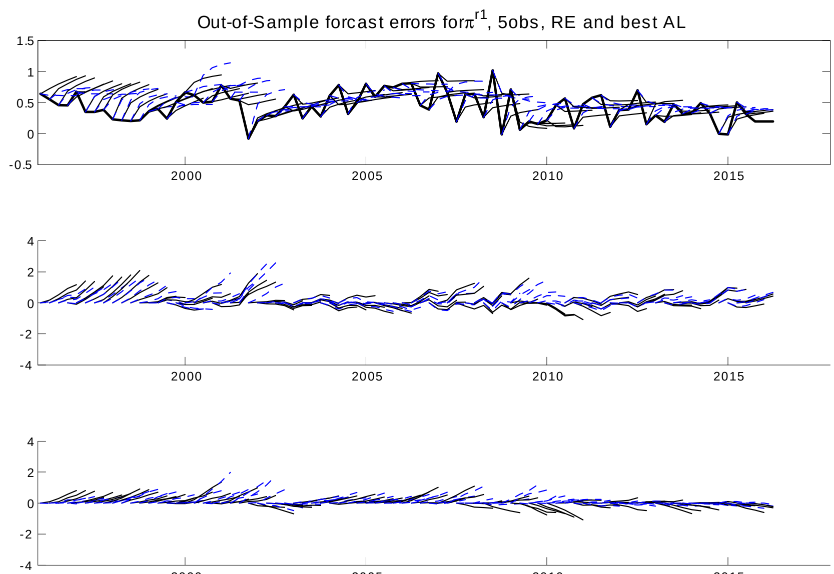

Note: see Figure 1 (ferr_pir1_hair_5obs.eps)

[9] Add discussion on relation between forecast errors and measurement error and on correlation among shocks

Figure 3: Estimated measurement error in SPF-data in the RE and AL Model

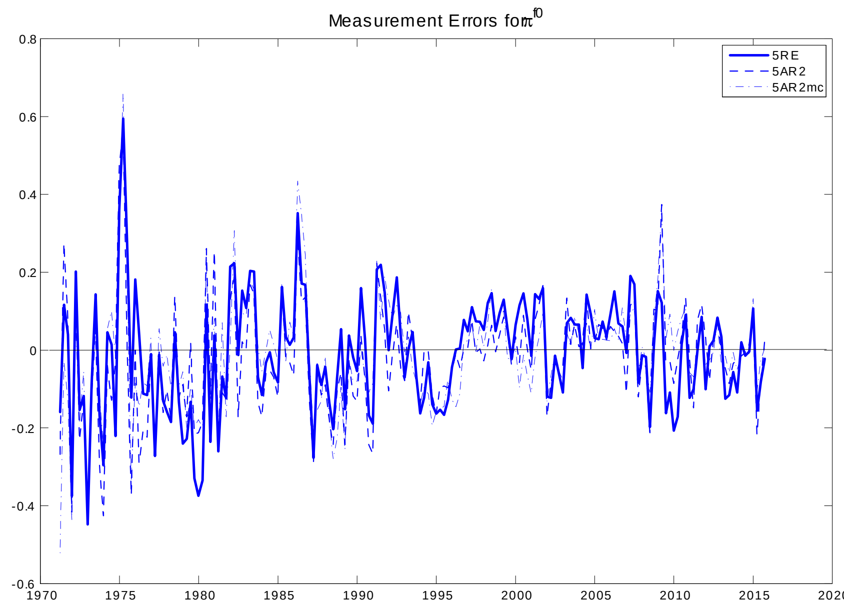

Note: pexp_dif_pif0_3_5obsme Reverse the sign in these graphs ?

This relative succes of the AL-model in capturing the overall dynamics in the inflation process is confirmed in the marginal likelihood comparison summarized in Table 6. Overall the AL-models with 9 and 10 observables have a better score than the RE models. By relaxing the RE-restrictions and assuming that expectations are based on small belief models that are updated over time depending on new realisations, the AL-models have some extra flexibility which is useful for forecasting. This flexibility is particular helpful in replicating the survey forecast observable in the 10obs model. This becomes most obvious when we calculate the marginal likelihood for the original 9 variables implied by the 10obs model. For the AL-models, the marginal likelihood for this common block remains more or less the same relative to the 9obs model (-941 versus -943) while for the RE-models, including the survey forecast in the model, deteriorates the marginal likelihood of the common 9 observables (-977 versus -965). The AL models are more flexible to deliver forecasts consistent with the survey forecasts and retain their overall good forecasting performance, while the RE-model is losing in overall performance when it has to comply with the survey forecasts.

AL-models with simple AR(2) belief models do a reasonable good job in mimicking the survey expectations was also suggested in SW2012. This observation is consistent with experimental evidence about expectations forecasting.[10] Small forecasting models provide a good respresentation of how agents formulate their expectations. However, it would be surprising that small models are able to reproduce the potentially rich information set that is underlying the SPF forecasts. To illustrate the role of the belief specification under AL, we also considered slightly more complicated belief models in which the AR(2) specification was augmented with the marginal cost variable. This Phillips Curve based specification can capture the basic relation between inflation and its underlying macroeconomic determinants. While this augmented belief model (AR2+MC) is not producing any gain for the standard 9obs model, again consistent with SW2012, it becomes informative for the model with observed survey forecasts. The estimated standard error for the m.e. on SPF-expectations drops from 0.18 to 0.15. The marginal likelihood of the AL model with the marginal cost (MC) in the beliefs improves by 30 relative to the model basic AR(2) specification without MC in the beliefs. The PC-relation seems to have some relevance in forecasting the relative smooth inflation expectation variable while it was not informative for forecasting the highly volatile realized inflation process.

Table 6: Marginal likelihood of alternative model specifications

	RE			AL		
	71q1-15q3		96q1-15q3	71q1-15q3		96q1-15q3
	9obs	10obs		9obs	10obs	
9obs	-965.22		-361.25			
PLM=AR2				-943.42		-340.96
PLM=AR2+MC				x		
10obs ME	-977	-910.87	-302.28			
PLM=AR2				x	-885.94	-284.38
PLM=AR2+MC				-941	-856.06	-266.46
10obs 2MU	-944	-840.90	-267.06			
PLM=AR2				x	-899.87?	-287.42
PLM=AR2+MC						
PLM=AR2+MC+UC				-920	-790.51	-226.70

Note:

We can conclude from this section that just adding SPF as observable for the expectations in these models is not producing strong gains. The discrepancy between the SPF and the model forecasts leads to some interesting changes in the estimated parameters and the inflation forecasts improve. But the measurement errors remain large and persistent (and they are correlated with other structural innovations in the model). There is no evidence that the additional

[10] See C. Hommes and M. Zhu (2014) for a review of arguments in favor of small forecasting models

observables leads to better identification of shocks or parameters that could improve the overall model performance. Some further adjustment in the model specification is required in order to bridge the gap between model and survey forecasts more efficiently.

2 Reconciling model and survey expectations

To reconcile the model expectations with the survey expectations, we need more flexibility in the specification of the inflation dynamics. In the RE-SW model, the price and wage markup shocks are modelled as an ARMA process. For the price markup this process is written as:

$$\mu_t^p = \rho_\mu^p \cdot \mu_{t-1}^p - \theta_\mu^p \cdot \varepsilon_{t-1}^p + \varepsilon_t^p$$

This specification implies that the same innovation ε_t^p is driving the volatile high frequency MA-component on the one hand and the persistent low frequency AR-component on the other hand. This ARMA specification works well to capture the complex exogenous shock process in the actual price and wage dynamics.[11] As long as the dataset is limited to the standard seven macrovariables, there is no need to distinguish between separate innovations driving the high and the low frequency component. These innovations are simply not identified individually by the standard seven observables. In the AL-version of SW2012, the price and wage markup shocks reduces to an i.i.d. process, $\mu_t^p = \varepsilon_t^p$ and the time-varying AR(2) beliefs generate the required dynamics to match the observed price and wage persistence. But also in this setup, one exogenous innovation is sufficient to describe the exogenous shock process.

Observing the survey expectations allow precisely to distinguish the pure i.i.d. component and the persistent component in the markup process. The survey forecasts which are, most likely, based on a broader and a more timely information set, makes it possible to identify these two innovations. Therefore, we specify the price and wage mark-up process as a combination of a persistent AR process (μ_t^{par}) and a separate i.i.d. shock process (μ_t^{piid}) each with their own innovation:

$$\begin{aligned} \mu_t^p &= \mu_t^{par} + \mu_t^{piid} \\ \mu_t^{par} &= \rho_\mu^p \cdot \mu_{t-1}^{par} + \varepsilon_{t-1}^{par} \\ \mu_t^{piid} &= \varepsilon_t^{piid} \end{aligned}$$

We assume that the innovation to the persistent shock process (ε_{t-1}^{par}) is already publicly observed in the quarter prior to its actual impact on price

[11] See also Ang, Bekart and Wei (2007) and Stock and Watson (2007) for more evidence supporting the ARMA specification for forecasting the inflation dynamics.

setting. This assumption is not crucial for the results that are presented, but the model fit improves when using the "news" assumption instead of the standard contemporaneous innovation assumption. Examples of this type of events are oil shocks or other commodity shocks that are observed in world prices before they actually enters into the retail price, or announced changes in regulated prices or taxes that are communicated in advance of the actual application, etc. The same dual process is introduced for the wage markup shock to remain symmetric.

We use the same measurement equation for the SPF survey data as in equation (1). We discuss the implication of this new specification first for the RE-setup of SW2007 and in the next section in the AL-setup of Slobodyan Wouters 2012.[12]

2.1 integrating survey data in the augmented RE-SW2007 model

Table A3 summarize the estimated parameter for this RE model with two markup shocks estimated on ten observables including the SPF forecast (10obs-2MU-RE model). The estimated standard deviation of the measurement error for the SPF-forecast reduces to 0.04, against 0.18 in the model specification with measurement error only (10obs-ME_RE). Also the nominal stickiness in prices becomes extremely high. This tendency to more stickiness was already present in the RE-model with measurement error only. Under RE, high price stickiness seems important to match the survey expectations. In the wage setting, the wage markup shocks becomes close to a random walk process while nominal stickiness is relatively low.

Table 7 documents the inflation forecast performance of this model. The inflation forecast of this model improves on all dimensions. The model forecast resembles now very closely the SPF-forecast and the forecast statistics are therefore also very similar to these of the SPF. The rmse forecast error for the one quarter ahead inflation rate approximates the benchmark SPF performance. The DM-test confirms that the two forecasts are not significantly different, except for a minor deviation in the beginning of the in-sample test. Also at longer horizons, the model forecast remains very similar to the SPF forecast although these are not observed in the model. Figure 4 provide the corresponding plot of the cumulative forecast deviations for the out-of-sample forecasts. For the RE-forecast there are almost no obvious deviations from the SPF forecast since 1996

[12] Note that this identification of two separate markup innovations is not feasible as long as the survey expectations are not included in the data file. The marginal likelihood of the model with this additional shock is identical to the model with the ARMA structure both under RE and under AL. The filtered innovations are highly correlated. The smoothed innovations are only partially identified.

Table 7: Forecast Statistics for the 10obs-2MU-RE-SW2007 with SPF observable

	1971q1-2015q3			1996q1-2015q3		
t+1 forecast	bias	mad	rmse	bias	mad	rmse
$\pi\_RE_{t+1\|t} - \pi\_r1_{t+1}$	0.03	0.21	0.26	0.03	0.17	0.21
$\pi\_RE_{t+1\|t} - \pi\_r2_{t+1}$				0.02	0.16	0.19
$\pi\_RE_{t+1\|t} - \pi\_rf_{t+1}$				-0.01	0.15	0.20
longer horizon...						
$\pi\_RE_{t+2\|t} - \pi\_r1_{t+2}$	0.04	0.25	0.34	0.04	0.19	0.23
$\pi\_RE_{t+3\|t} - \pi\_r1_{t+3}$	0.04	0.28	0.38	0.05	0.20	0.24
$\pi\_RE_{t+4\|t} - \pi\_r1_{t+4}$	0.04	0.30	0.41	0.06	0.21	0.25
RE versus SPF		rel.rmse%	DM-test		rel.rmse%	DM-test
$horizon = 1$		4.14	2.47			1.36
$horizon = 2$		29.28	0.50			0.45
$horizon = 3$		25.06	-0.75			1.16
$horizon = 4$		25.84	0.95			0.06

Note:

Figure 4: Model forecast with SPF observed and two markup shocks

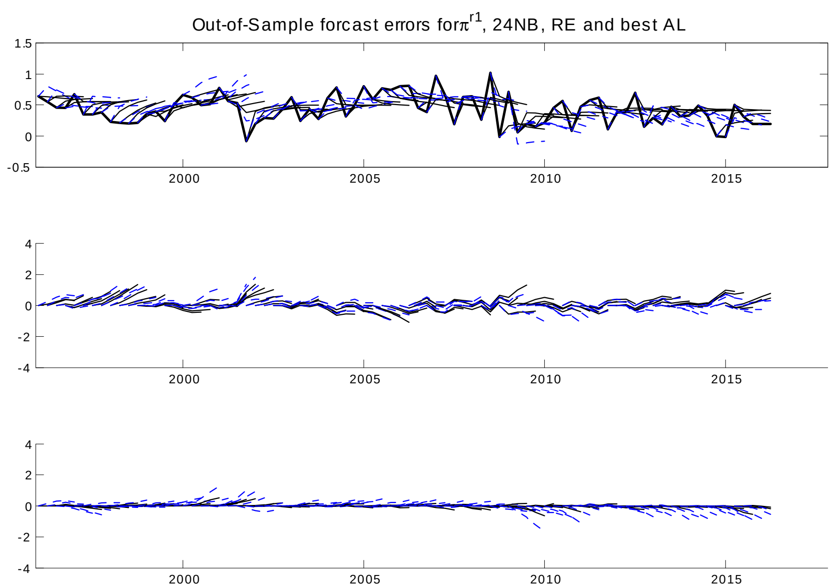

Note:

The new model structure with two markup shocks gives the model precisely the flexibility necessary to fit jointly the realised inflation process and the survey forecasts. The highly volatile i.i.d. markup shocks with a standard deviation of 0.24 explain the volatile high-frequency component in actual inflation. As illustrated in the impulse response function in Figure 5, inflation is only affected on impact and returns to its pre-shock level in the next period with a very small negative correction afterwards. This implies that the shock is almost irrelevant for the inflation forecast for t+1 and the spillover effects to the real economy are also minimal. On the other hand, the persistent autoregressive markup shocks in prices and wages with a standard error of respectively 0.026 and 0.007 and persistence of 0.78 and 0.997, are crucial to capture the innovations in the survey forecasts. These "news" shocks have already an impact on actual prices at time t as well, consistent with the forward-looking nature of the price setting problem. The magnitude of this impact effect is substantial smaller than the iid shock: one half for the persistent price shock and one third for the wage shock. Over time, the wage shock starts to dominate as it is a quasi-permanent shock.

As also indicated in the conditional covariance decomposition in Table 8, the three markup shocks have each their own role in the inflation process. The iid price markup explains the one quarter ahead forecast error in realised inflation (77%) but is irrelevant for expectations. The persistent price markup shock is crucial for the short term forecast error in the survey expectations (65%), and consistent with that, also for realized inflation over the medium term horizon of one or two years ahead. The role of the wage markup shock builds up only gradually but is dominant over the very long horizon and explains at that frequency 78% of the inflation expectations and 60% of the realized inflation variance. Note that this wage shock has minor effects on the short term wage developments and in that sense it is not surprising, that the precise timing of the innovations to this wage shock is difficult to identify. In fact, the persistent wage and price innovations are highly correlated (0.79) among themselves but not with the iid shocks. This observation is important as it raises questions about the correct interpretation of the persistent wage markup shock.

Table 8: Conditional variance decomposition for the 10obs-2MU-RE model

	ε^a	ε^b	ε^g	ε^{qs}	ε^m	ε^{piid}	ε^{wiid}	ε^{par}	ε^{war}	$\xi^{\pi f1}$
1 quarter horizon										
π_f0	2.44	0.14	0.69	0.00	0.52	0.82	0.29	64.66	24.55	5.89
π_r1	0.65	0.04	0.17	0.00	0.13	76.73	0.09	14.06	6.12	0.00
w	0.05	0.86	0.10	0.05	0.84	9.63	86.60	1.59	0.29	0.00
y	18.71	41.39	20.33	5.04	13.24	0.80	0.33	0.01	0.15	0.00
1 year horizon										
π_f0	3.46	0.18	1.06	0.00	0.73	0.36	0.31	53.97	37.68	2.25
π_r1	1.64	0.09	0.48	0.00	0.35	47.69	0.18	32.65	16.93	0.00
w	0.22	3.74	0.83	0.52	5.47	6.07	72.73	6.16	4.25	0.00
y	17.07	39.17	6.53	10.96	24.28	1.19	0.14	0.45	0.21	0.00
10 year horizon										
π_f0	4.26	0.08	1.54	0.04	0.42	0.11	0.11	14.56	78.33	0.54
π_r1	3.41	0.07	1.22	0.03	0.35	19.07	0.11	15.19	60.56	0.00
w	2.57	5.06	7.78	3.65	16.69	2.50	17.49	10.47	33.78	0.00
y	34.13	15.30	3.20	11.01	24.65	1.14	0.04	4.84	5.68	0.00

Note:

To illustrate this interpretation problem, we considered a model that has also a quasi-permanent inflation target shock. This shock is able to substitute almost perfectly for the persistent wage markup shock. In this alternative setup, the long term inflation trend, common to expectations and realisations, is explained by an exogenous inflation target shock. The wage markup interpretation, on the other hand, is more consistent with a severe trade-off problem for monetary policy typically for cost-push shock situations. Without further information from additional labour market variables (as in Gali Smets Wouters 2014) and/or about monetary policy objectives, the RE has a hard time to differentiate among these interpretations. Given our list of observables, the model filter delivers a relative precise estimate of the long run inflation expectations component but no unique interpretation for why expectations were not anchored more precisely.

Figure 5: IRF function for iid and ar price markup, wage markup and productivity

	inflation expectations	inflation	output

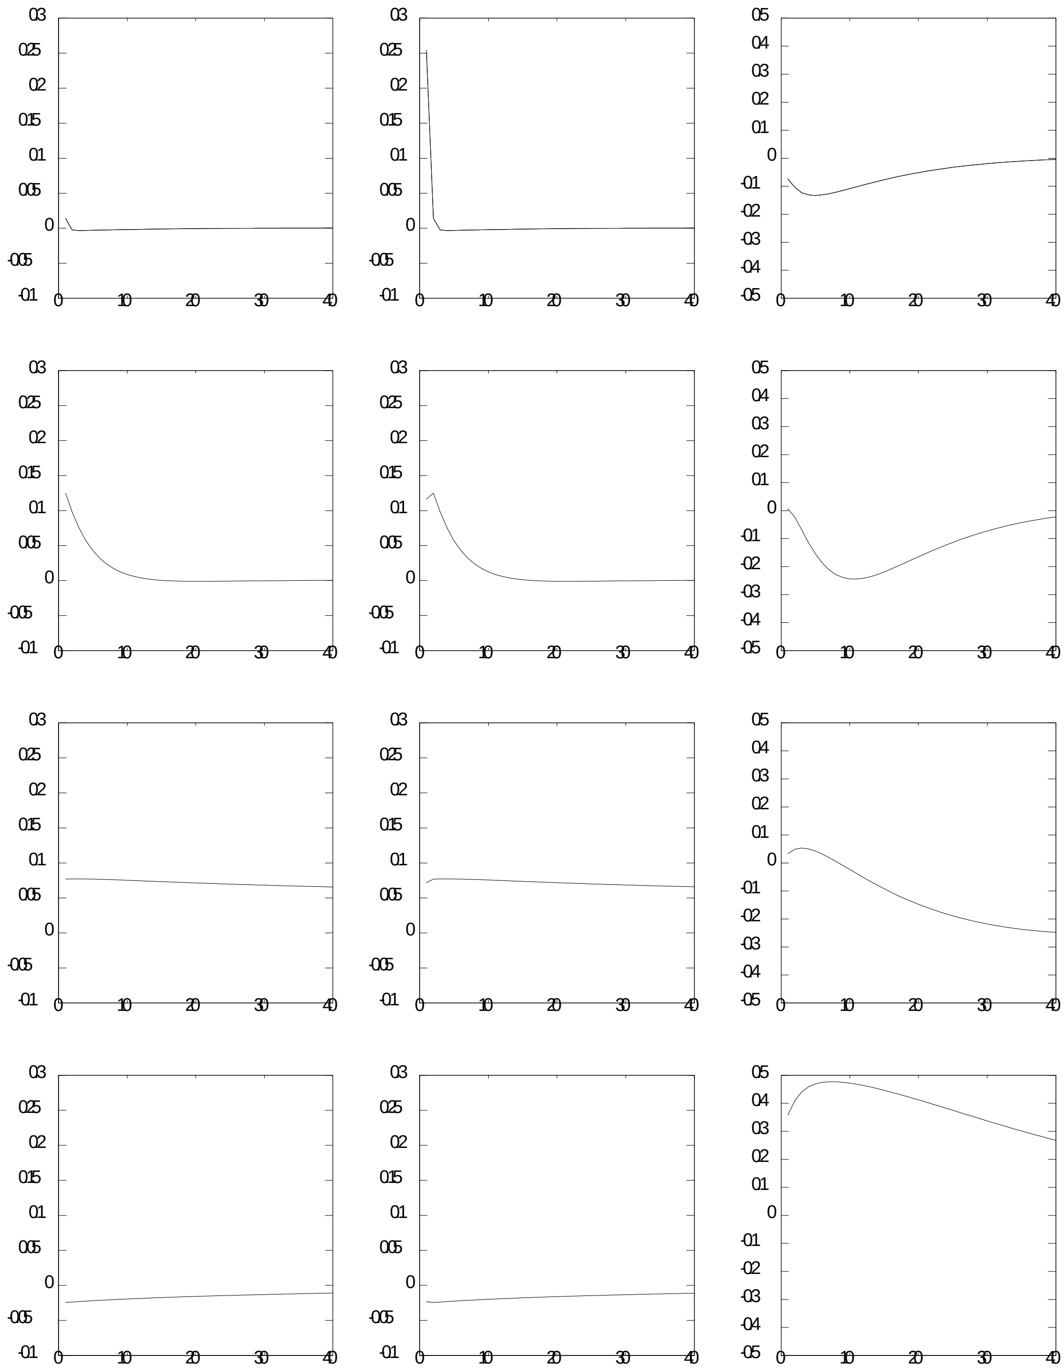

ε^{piid}

ε^{pap}

ε^{wap}

ε^{a}

Note:

The next important issue is then how this model performs on other dimensions. Allowing for the two separate markup innovations boosts the log marginal likelihood of the 10obs model by 60 relative to the 10obs-ME-RE model with measurement error only. Of course the improved fit of inflation expectations delivers an important contribution to this gain. But when we look at the marginal likelihood of the original 9 observables in this model, we observe that also on this dimension the model outperforms the original 9obs-RE model, and the 10obs-ME_RE model with measurement error only. So the overall performance of the model is improved under this specification and the information from the survey expectations helps to predict the other variables in the economy as well. This result is further documented in Table 9 with the forecasting results for the individual variables. Over the entire sample, the in-sample rmse indicates that the main gains are concentrated in the inflation block. Similarly In the out-of-sample forecasts over the recent period, the gains are concentrated in the price, wage and interest rate block but the forecast on other real variables (investment, output and hours) deteriorate slightly.

Table 9: Forecast performance of the 10obs-2MU-RE model

	π_r1	π_r2	π_f0	dy_r1	dy_r2	dc	$dinve$	$hours$	dw	r
1971q1-2015q3										
bias	0.03	-0.02	0.01	0.19	-0.01	-0.07	-0.21	0.04	0.01	0.03
amd	0.21	0.08	0.10	0.48	0.17	0.47	1.27	0.43	0.57	0.15
rmse	0.26	0.12	0.15	0.63	0.21	0.63	1.71	0.55	0.78	0.23
1996q1-2015q3										
bias	0.03	-0.02	0.01	0.29	-0.00	0.03	0.33	0.30	-0.09	0.08
emd	0.17	0.05	0.07	0.47	0.17	0.40	1.17	0.47	0.75	0.11
rmse	0.21	0.07	0.10	0.57	0.21	0.56	1.62	0.60	0.98	0.14
log lik score	0.89	2.00	1.56	-0.09	1.05	0.05	-0.99	0.01	-0.57	1.22
improvement over RE-model with measurement error										
rmse	0.25	0.07	0.17	0.54	0.21	0.57	1.58	0.52	1.01	0.15
log lik score	0.77	1.99	1.14	-0.08	1.05	0.04	-0.98	0.13	-0.59	1.16
improvement relative to original RE										
rmse	0.28	0.06		0.54	0.21	0.57	1.59	0.57	0.99	0.14
log lik score	0.65	1.99		-0.06	1.04	0.03	-0.98	0.05	-0.67	1.17

Note: Results reported for the complete sample 1971q1-2015q3 are calculated as in-sample results, while the result for the sample 1996q1-2015q3 are out-of-sample forecasts with re-estimation

In sum, in this RE-model with two markup shocks, the nominal block performs very well. The information of the survey is exploited efficiently to improve on forecast of the inflation and nominal variables But the RE model imposes a constant structure on the data, while we know from reduced form exercises that the nature of the inflation process has changed over time and the question therefore is whether a structure with two shocks with different persistence but constant variance is optimal. The AL-setup could precisely improve on that dimension.

2.2 integrating survey data in the AL-SW2012 model

We repeat here first the main steps of our Kalman filter based AL-algorithm. Then we discuss the assumptions that we make on the forecasting models that represent the beliefs of the agents in the AL-model. The simple AR2 specification that we retained in SW2012 is not able to exploit optimally the rich information structure from the survey. We will reformulate the forecasting/belief models so that there is a role for the expectation signals in the agents beliefs. the estimation results are presented in the third section.

2.2.1 Adaptive Learning Set-up based on the Kalman filter

As in Evans & Honkapohja (2001), we assume that the economic agents do not have perfect knowledge of the reduced form parameters of the model when forming expectations about the future. Therefore, they forecast future values of the forward variables in the model (y^f) with a linear functions of endogenous model variables.[13] The general logic of adaptive learning works as follows.

The model is represented as

$$\overline{\alpha} + A_0 y_{t-1} + A_1 y_t + A_2 E_t y_{t+1} + B\epsilon_t = 0, \tag{2}$$

where y_t is a vector of endogenous and exogenous model variables. The RE solution of this system is presented as a VAR(1) process,

$$y_t = \mu + T y_{t-1} + R\epsilon_t.$$

Under adaptive learning, agents assume that the forward-looking variables are linear combinations of some variables in the vector y_{t-1}. This assumption is known as a Perceived Law of Motion, or PLM:

$$y_t^f = \alpha_{t-1} + \beta_{t-1}^T y_{t-1}. \tag{3}$$

By rolling forward the PLM, we obtain the agents' expectations of forward-looking variables as

$$E_t y_{t+1}^f = \alpha_{t-1} + \beta_{t-1}^T y_t.$$

[13]Our adaptive learning models are realized in a specialized DYNARE toolbox. Therefore, we follow the Dynare notation in our formula's.

These expectations are then plugged into the model representation (2), and the resulting purely backward-looking model is solved to produce the Actual Law of Motion, or ALM:

$$y_t = \mu_t + T_t y_{t-1} + R_t \epsilon_t. \tag{4}$$

The model transmission mechanism (μ, T, and R) thus becomes a time-varying function of coefficients in the agents' forecasting equations (α and β), called *beliefs*.

The beliefs could be updated using any convenient adaptive algorithm. In the literature, Recursive Least Squares (RLS) and Kalman filter proved to be the most popular.

In a previous paper (Slobodyan and Wouters 2012b) we utilized Kalman filter learning over a set of small forecasting models, where the set of variables the agents use for forming their forecasts is much smaller than the Minimum State Variable (MSV) set that is needed to achieve Rational Expectations.

The precise learning procedure is defined as follows. Agents estimate the forecasting model at each point in time given the information set available at that time. We assume that they use an efficient Kalman filter updating mechanism[14]. They believe that the coefficients β (a vector obtained by stacking all β_j) follow a vector autoregressive process around $\overline{\beta}$ (which will be specified later): $vec\left(\beta_t - \overline{\beta}\right) = F \cdot vec\left(\beta_{t-1} - \overline{\beta}\right) + v_t$, where F is a diagonal matrix with $\rho \leq 1$ on the main diagonal[15]. Errors v_t are assumed to be *i.i.d.* with variance-covariance matrix V.

We can write the forecasting model in the following SURE format[16]:

$$\begin{bmatrix} y^f_{1t} \\ y^f_{2t} \\ \vdots \\ y^f_{mt} \end{bmatrix} = \begin{bmatrix} X_{1,t-1} & 0 & \dots & 0 \\ 0 & X_{2,t-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & X_{m,t-1} \end{bmatrix} \begin{bmatrix} \beta_{1,t-1} \\ \beta_{2,t-1} \\ \vdots \\ \beta_{m,t-1} \end{bmatrix} + \begin{bmatrix} u_{1,t} \\ u_{2,t} \\ \vdots \\ u_{m,t} \end{bmatrix}, \tag{5}$$

The errors $u_{j,t}$ depend on a linear combination of the true model innovations ϵ_t and therefore they are likely to be correlated, making the variance-covariance matrix non-diagonal: $\Sigma = E\left[u_t \cdot u_t^T\right]$. With the above notation, the Kalman

[14] Sargent and Williams (2005) showed that even if Kalman filter and constant gain learning are asymptotically equivalent on average, their transitory behaviour may differ a lot. In particular, Kalman filter tends to result in much faster adjustment of agents' beliefs. With faster adjustment of beliefs, we are able to better understand whether the initial beliefs or time–varying coefficients matter more for the improved model fit.

[15] ρ is restricted to be the same for the seven variables that are forecasted. Allowing for a variable specific autocorrelation provides some extra flexibility but also larger parameter uncertainty.

[16] The SURE format and the corresponding GLS estimator are necessary to get an efficient estimator of the complete forecasting model because the variables appearing on the RHS in each equation are not identical.

filter updating and transition equations for the belief coefficients and the corresponding covariance matrix are given by

$$\beta_{t|t} = \beta_{t|t-1} + P_{t|t-1}X_{t-1}\left[\Sigma + X_{t-1}^T P_{t|t-1} X_{t-1}\right]^{-1} \times \left(y_t^f - X_{t-1}^T \beta_{t|t-1}\right), \quad \text{with } (\beta_{t+1|t} - \overline{\beta}) = F \cdot (\beta_{t|t} - \overline{\beta}). \tag{6a}$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1}X_{t-1}\left[\Sigma + X_{t-1}^T P_{t|t-1} X_{t-1}\right]^{-1} \times X_{t-1}^T P_{t|t-1}, \quad \text{with } P_{t+1|t} = F \cdot P_{t|t} \cdot F^T + V. \tag{6b}$$

These best estimates for the beliefs ($\beta_{t|t-1}$) are then substituted for β_t in (5) to generate expectations of forward-looking variables, $E_t y_{t+1}^f$. Plugging these expectations into (2), we obtain a purely backward-looking representation of the model (4).[17] The resultant time-dependent matrices μ_t, T_t, and R_t replace the constant equivalents in the RE solution. These matrices depend now on both the structural parameters of the decision problem (Θ) and on the best estimates of the forecasting model ($\beta_{t|t-1}$), and contain all necessary information to describe the dynamics and the propagation of the shocks in the model under learning. In terms of adaptive learning literature, the equation (4) represents the Actual Law of Motion (ALM) of the model.

In order to initialize this Kalman filter for the belief coefficients, we need to specify $\beta_{1|0} = \overline{\beta}$, $P_{1|0}$, Σ, and V. In our baseline approach, all these expressions are derived from the correlations between the model variables implied by the RE Equilibrium evaluated for the corresponding structural parameter vector Θ. In other words, the initial beliefs are assumed to be model consistent.[18]

Using the fact that $\widehat{\beta}_{OLS} = \left(X^T X\right)^{-1} X^T y$ is unbiased, we use the theoretical moment matrices $E\left[X^T X\right]$ and $E\left[X^T y\right]$ from the RE solution and set $\beta_{1|0} = (E\left[X^T X\right])^{-1} \cdot E\left[X^T y\right]$. Given $\beta_{1|0}$, we calculate Σ as

$$\Sigma = E\left[\left(y_t^f - X_{t-1}^T \beta_{1|0}\right)\left(y_t^f - X_{t-1}^T \beta_{1|0}\right)^T\right],$$

again using the RE theoretical moments. Finally, $P_{1|0}$, the initial guess about the mean square forecast error of the belief coefficients, and V, the variance–covariance matrix of shocks v_t to these coefficients, are both taken to be proportional to $\left(X^T \Sigma^{-1} X\right)^{-1}$:[19] $P_{1|0} = \sigma_0 \cdot \left(X^T \Sigma^{-1} X\right)^{-1}$, and $V = \sigma_v \cdot \left(X^T \Sigma^{-1} X\right)^{-1}$. This initialization leaves just three parameters, σ_0, σ_v, and ρ, to fully describe

[17] Note that we expand the state vector y in this representation with additional lags that occur in the forecasting models.

[18] An alternative approach would be to derive the initial beliefs and the underlying moment matrixes from the restricted expectations equilibrium. Given our under-parameterized beliefs, this equilibrium deviates from the REE and requires the solution of the underlying ODE. Computationally this procedure was not feasible in the estimation context.

[19] $\left(X^T \Sigma^{-1} X\right)^{-1}$ is equal to $Var\left[\widehat{\beta}_{GLS}\right]$ where $\widehat{\beta}_{GLS} = \left(X^T \Sigma^{-1} X\right)^{-1} X^T \Sigma^{-1} y$, which gives an efficient estimator for the SURE model. Given knowledge of theoretical moments and of Σ, the matrix $\left(X^T \Sigma^{-1} X\right)^{-1}$ could be readily calculated.

the learning dynamics, but in practice, we can keep σ_0, σ_v fixed and optimize over ρ only.

2.2.2 Specification of the belief models

In SW2012, an AR(2) belief specification was proposed for all forward variables. This very simple belief specification turned out to produce a good model fit.[20] The forecasting equation for inflation was of the form:

$$\left[\begin{array}{c} \pi_t^f \end{array}\right] = \left[\begin{array}{ccc} 1 & \pi_{t-1}^f & \pi_{t-2}^f \end{array}\right] \beta_{\pi,t-1} + \left[u_{\pi,t}\right], \tag{7}$$

We have already mentioned that including the marginal cost as an additional regressor in the prediction model is very useful once we observe the survey expectations in the 10obs model.

In order for the AL-models to exploit the information from the survey data more efficiently, we must include additional variables in the belief models. A simple AR(2) belief model cannot capture the rich information structure of the survey data that we observed in the analysis of the RE-model. Therefore, we consider a specification for the beliefs that include all independent determinants that affect the inflation dynamics. This means that we have to include in the belief specification not only the lags of inflation and the marginal cost but also the unobserved innovations in the markup process. It is precisely by including these innovations in the beliefs that identification of the separate markup disturbances becomes possible.[21].

$$\left[\begin{array}{c} \pi_t^f \end{array}\right] = \left[\begin{array}{cccccc} 1 & \pi_{t-1}^f & \pi_{t-2}^f & mc_{t-1} & \varepsilon_{t-1}^{par} & \varepsilon_{t-2}^{piid} \end{array}\right] \beta_{\pi,t-1} + \left[u_{\pi,t}\right], \tag{8}$$

To keep symmetry, a similar belief model was used for beliefs about the wage process (with the marginal rate of substitution replacing the marginal cost):

$$\left[\begin{array}{c} w_t^f \end{array}\right] = \left[\begin{array}{cccccc} 1 & w_{t-1}^f & w_{t-2}^f & mrs_{t-1} & \varepsilon_{t-1}^{war} & \varepsilon_{t-2}^{wiid} \end{array}\right] \beta_{w,t-1} + \left[u_{w,t}\right] \tag{9}$$

Agents will use the available information about the markup shocks at time t and t-1 when forming their expectations for t+1. This information is disclosed in the model through the observation of the survey forecasts and this will affect decisions contemporaneously as well. However, these consequences for the traditional macrovariables are modest and not sufficiently strong for the identification of the precise nature of the shocks in the 9obs models.

[20] See also Hommes and Zhu (2014) for more evidence supporting the use of simple forecasting rules

[21] Note that we must include the iid innovation with two lags in order to secure independence among the RHS-regressors and to avoid singularity in the covariance matrix.

2.2.3 Estimation results with Survey data

The estimated parameters of this 10-obs-2MU-AL-model are standard (see Table A3): the stickiness in both prices and wages is high but not too extreme. The persistent markup processes have a small standard deviation, 0.037 for prices and 0.026 for wages, and reasonable persistence of respectively 0.78 and 0.65. The measurement error in the inflation expectation is further reduced to 0.011. The model forecast for t+1 almost perfectly match with the survey expectations. Of course in this setting the PLM coefficients are also crucial to understand the inflation dynamics both in terms of persistence and volatility. In an AL-context, the transmission via the endogenous belief coefficients is more important for the inflation dynamics than the exogenous persistence in the shocks which is crucial under RE

The one-quarter ahead inflation forecast of this augmented AL model is almost identical to the equivalent SPF inflation forecast (see Figure 4 and Table 10). The accuracy of the two forecasts is not significantly different according to the DM-test. For longer horizons, the quality of the inflation forecast is less impressive both relative to the SPF, although the difference is not significant, and relative to the RE-model. As suggested before, this can result from the flexibility of the AL coefficients and the lack of parameter restrictions on the belief model. We can solve this problem by incorporating longer horizon forecasts in the list of observables as we illustrate further below.

The marginal likelihood of this model is superior to all previous models. The improvement relative to the AL-model with measurement error only (10-obs-ME-AL) is of the order of 66 and with respect to the best RE-model (10obs-2MU-RE) the improvement is 50. The model does also an excellent job for the 9-common variables: here the improvement is of the order of 20 to 24. Note that the augmented belief equation for inflation is crucial for this excellent marginal likelihood result: the marginal likelihood for the AR(2) and AR(2)+MC belief models is much smaller (if any? to be completed in the Table). The information about the nature of the markup shocks must be incorporated in the belief equations. In this way we provide the agents in the model with the same timely information than the survey participants. When observing the survey forecast for time t+1 in the course of time t, the agents in the model can identify correctly the nature of the markup shocks. This information about the persistence of the shocks determines their contemporaneous actions and their expectations for next period. As in the RE-model, the survey data are extremely informative to distinguish the more persistent markup shocks from the i.i.d. component in the inflation dynamics. This result confirms the usefulness of adding the survey data to the set of observables as we stressed in the introduction

These impressive gains in marginal likelihood are to a large extent explained by the time-varying volatility that is produced by the updating of the beliefs.

We will document this in various ways. But the standard out-of-sample prediction statistics of this model (bias, amd and rmse) are also excellent. This applies to both the inflation variables and the real variables this time. The AL-model outperforms the RE-model in rmse on all variables except consumption growth. Compared to simpler AL-model, there is an overall gain except for investment. Given the time-varying covariance structure, the likelihood score of the forecasts becomes informative as it weights the forecast errors by their conditional variances. The impact of this correction on the forecast score is the largest for the inflation expectation variable. While in terms of rmse, the results are marginally worse than for the RE-model, in terms of log likelihood score the AL forecast dominates by far the RE outcome.

In order to illustrate the impact of the time-varying covariance matrix on the likelihood/posterior evaluation, we did the following experiment. We used the time-varying covariance matrixes for the one-period ahead forecast errors from the AL-model to evaluate the likelihood of the RE-prediction errors. The log posterior value of the RE-model, evaluated at the parameters corresponding with the RE-mode, improves with this correction for time variation in the prediction uncertainty from -718 to -673. This value can be compared with the log posterior of the AL-model at the mode of -664. Re-evaluating the AL-models with the fixed covariance structure of the RE model results in a deterioration of the log posterior to -716 which is still slightly better than the log posterior mode of the RE-model. Clearly, a large fraction of the improvement in the log posterior value can be attributed to the time-varying covariance matrix that is implied by the updating of the beliefs over time. One could expect that most of this gain is realized in the beginning of the sample which is characterized by large variation in the covariance matrix as will be illustrated below. However, the same result is confirmed when repeating the exercise over the 1996q1-2015q3 period: the log posterior of the RE-model for that sub-sample improves from -268 to -237 when evalated with the time-varying AL-covariance structure. On the other hand, the AL model deteriorates from -241 to -265 when the constant covariance structure of the RE-model is used. Time-variation in the covariance structure is very important, but the AL model still outperforms the RE-model even when evaluated with the same constant covariance structure.

Table 10: Forecast Statistics for the augmented-AL-SW2012 model with SPF observable

	1971q1-2015q3			1996q1-2015q3		
t+1 horizon	bias	mad	rmse	bias	mad	rmse
$\pi\_AL\_PLM_{t+1\|t} - \pi\_r1_{t+1}$				0.03	0.17	0.21
$\pi\_AL\_ALM_{t+1\|t} - \pi\_r1_{t+1}$	-0.05	0.21	0.27	0.04	0.17	0.21
$\pi\_AL\_ALM_{t+1\|t} - \pi\_r2_{t+1}$				0.02	0.16	0.19
$\pi\_AL\_ALM_{t+1\|t} - \pi\_rf_{t+1}$				-0.00	0.15	0.19
longer horizon						
$\pi\_AL\_ALM_{t+2\|t} - \pi\_r1_{t+2}$	0.04	0.25	0.24	0.04	0.20	0.24
$\pi\_AL\_ALM_{t+3\|t} - \pi\_r1_{t+3}$	0.04	0.27	0.37	0.04	0.20	0.26
$\pi\_AL\_ALM_{t+4\|t} - \pi\_r1_{t+4}$	0.03	0.30	0.43	0.04	0.22	0.28
AL versus SPF	rel.rmse%		DM-test	rel.rmse%		DM-test
$horizon = 1$	38.58		0.75			-0.12
$horizon = 2$	55.45		0.13			1.57
$horizon = 3$	57.42		-0.29			1.26
$horizon = 4$	59.52		0.13			0.77

Note:

Table 11: Forecast performance of the Augmented AL-SW model

	π_r1	π_r2	π_f0	dy_r1	dy_r2	dc	$dinve$	$hours$	dw	r
1971q1-2015q3										
bias	-0.05	0.02	0.00	-0.17	0.01	0.05	0.21	-0.04	0.02	-0.03
amd	0.21	0.08	0.11	0.50	0.16	0.53	1.24	0.42	0.55	0.13
rmse	0.27	0.12	0.16	0.68	0.21	0.72	1.68	0.54	0.75	0.22
1996q1-2015q3										
bias	0.04	-0.01	-0.01	0.15	-0.01	0.08	-0.36	0.14	0.00	0.04
emd	0.17	0.05	0.08	0.39	0.16	0.42	1.15	0.36	0.73	0.09
rmse	0.21	0.07	0.11	0.50	0.21	0.61	1.51	0.46	0.94	0.12
log lik score	0.87	1.99	1.71	-0.08	1.06	0.00	-0.92	0.20	-0.49	1.27
improvement over AL-model with measurement error										
rmse	0.25	0.07	0.13	0.54	0.21	0.61	1.47	0.51	0.97	0.13
log lik score	0.69	2.01	1.14	-0.09	1.07	0.03	-0.88	0.13	-0.52	1.20
improvement relative to original AL-AR(2)										
rmse	0.29	0.07		0.52	0.20	0.62	1.37	0.46	0.98	0.11
log lik score	0.68	1.96		-0.09	1.07	-0.07	-0.84	0.15	-0.66	1.23

Note:

Figure 6: Conditional Variance for inflation forecast in the RE and AL Model

inflation expectations

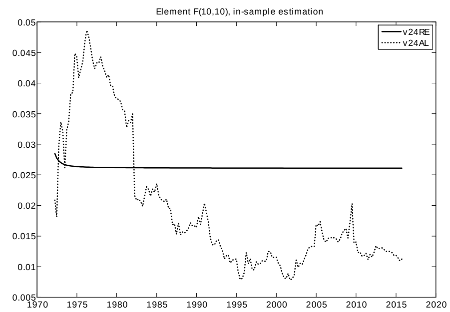

consumption growth

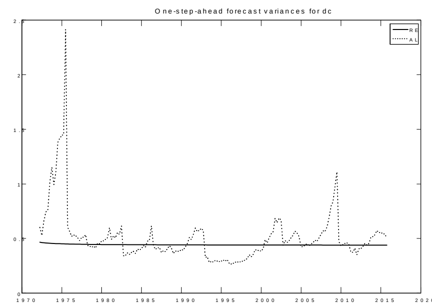

realized inflation

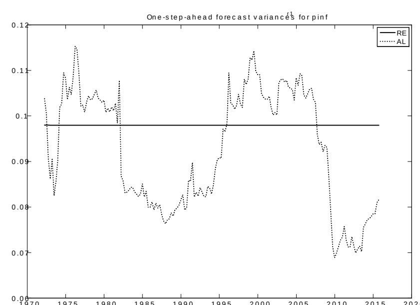

investment growth

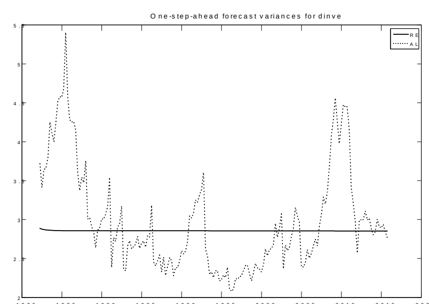

Note: pexp_dif_pif0_3_5obsme, Sill Philadelhia org 2014

2.2.4 Analysis of the time-variation in the AL model

Figure 6 illustrates the time profile of the conditional variance for a selected number of variables. Crucial variation is observed in the variance of the one-period forecast error in the inflation expectations. This variance was up to three times higher in the seventies than in the period since 1995. Interesting is that the variance increased again slightly from 2005 onwards with a peak in 2010, but a decline since then. Uncertainty in inflation expectations has reduced recently despite the remaining uncertainty about the speed of the recovery and corresponding monetary policy reaction. The profile in this uncertainty is also consistent with other indicators of the forecast uncertainty: for instance the squared IQ-dispersion among individual forecasts in the SPF survey follows a similar historical development. It also resembles the stochastic volatility process of the variance for the persistent unobserved component in the Stock and Watson UC-SV model for inflation (Stock and Watson 2007-JMCB) as illustrated in Figure 7.

The conditional variance in actual inflation realisations follows a more complex profile: it inherits the uncertainty peak in the seventies from the inflation expectations component, but it has an additional peak during the period 1995-2007. The conditional variance in output-growth is representative for all other real variables and this profile is strongly affected by a cyclical updating process with a positive outlier in the mid seventies.

Figure 7: Conditional Variance for inflation forecast in the RE and AL Model

squared IQ-range - spf dispersion | Stock-Watson (2007) volatility of permanent inflation component

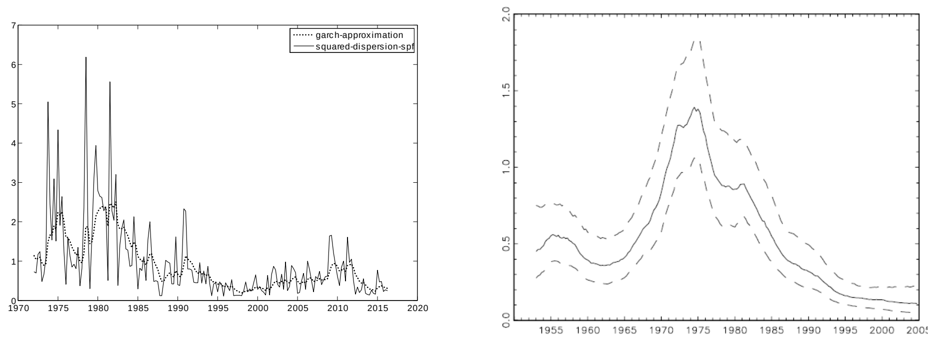

Note: pexp_dif_pif0_3_5obsme, Sill Philadelhia org 2014

This time-variation in the conditional variances is explained by the updating in the belief models. In Figure 8, we plot the time variation in the belief coefficients of the forecasting model for inflation, consumption and investment. Starting with the inflation beliefs, the persistence parameter and the constant follow the same profile as documented in SW2012. The updating in the constant follows systematically the surprise in the inflation realisation: unexpected higher inflation leads to a positive updating in the constant and vice-versa. The updating in the constant is very important for the long run inflation trend and the scale of this coefficient in Figure 8 is therefore misleading The updating in the persistence parameter (the sum of ρ_1 and ρ_2, the coefficients of the two lagged inflation terms in the beliefs) reacts in a slightly more complicated way as it depends on the level of inflation: in periods where inflation is high, a positive inflation surprise will generate an upward adjustment in the perceived persistence of inflation. However, when inflation is low, a positive surprise in realised inflation leads to lower perceived persistence. Note also the opposite adjustment in the first (ρ_1) and second (ρ_2) autocorrelation coefficient: ρ_1 is particular important for the impact effect of all shocks including the highly volatile iid markup shock. The coefficients to the markup shocks are also adjusting in a similar direction: the coefficient to the persistent markup mimics the updating in the constant, while the changes in the i.i.d. markup innovation are somewhere in between the updating in the constant and the persistence.

Clearly big surprises in inflation and repeated surprises in the same direction affect the belief coefficient substantially. Understanding this time-variation in the long run perceived inflation target and the perceived persistence and shock sensitivity of inflation is highly relevant for the correct interpretation of inflation expectations in the monetary policy analysis.

Figure 8: Time-varying belief coefficients for inflation, consumption and investment

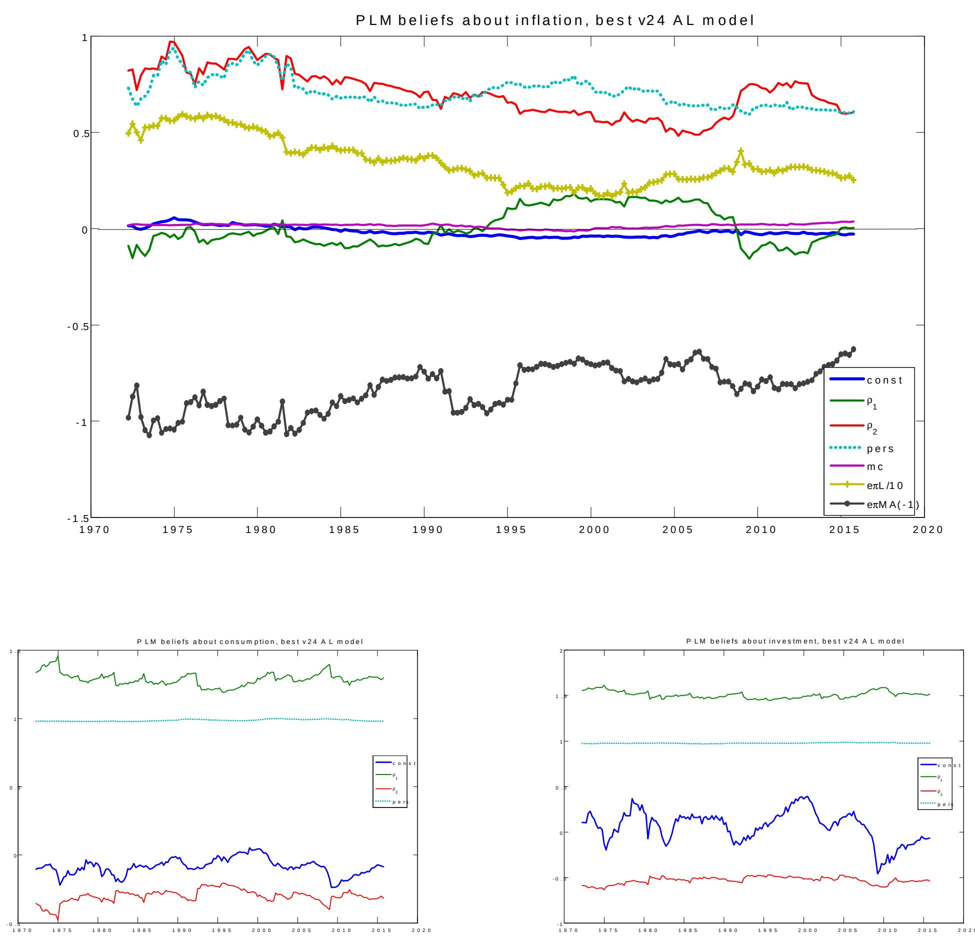

Note:PLM_beliefs_pinf_v24AL

The adjustment in the simple beliefs for consumption and investment are also very interesting: these variables are perceived as almost unit root processes with a highly time-varying drift in the "constant" and a time-varying autocorrelation term in the growth process. These beliefs generate a strong cyclical and skewed accelerator process in investment and consumption. In booms, both the constant drift coefficient and the first order autocorrelation coefficient in the growth rate tend to adjust positively reflecting higher confidence and optimism in the expectations. Once a negative shock interrupts the growth cycle, the beliefs about the constant growth rate decreases rapidly but the first order persistence parameter adjust only with a delay which means that the negative shocks are perceived as relatively persistent shocks and their impact is extrapolated into the future. Negative shocks in the beginning of the recession are therefore amplified and contributing to the asymmetry in the growth rate over the cycle.[22] The first order autoregressive parameter in the consumption beliefs is of particular importance in that it interacts with the habit coefficient in the consumption Euler equation. When this coefficient approaches the habit parameter, habits and growth rate extrapolation become reinforcing mechanisms that make consumption extremely sensitive to interest rate fluctuations. This explains the peak in the conditional variance of consumption in the mid seventies. Note also how the Great Recession has a huge impact on the drift growth factor in consumption and investment beliefs and how long it took for these beliefs to adjust again in the recovery.

The time-varying impulse response functions in Figure 9 confirm this amplifying or attenuating effect of the belief coefficients on the transmission mechanism of the various shocks over the cycle

[22] These implications for consumption and investment expectations should be verified by the survey expectations as well. We plan to do this in a follow up paper.

Figure 9: Time-varying impulse response functions in the AL-model: inflation expectations, inflation and output growth to shocks in persistent and iid price markups, persistent wage markup and productivity

inflation expectations	inflation	output

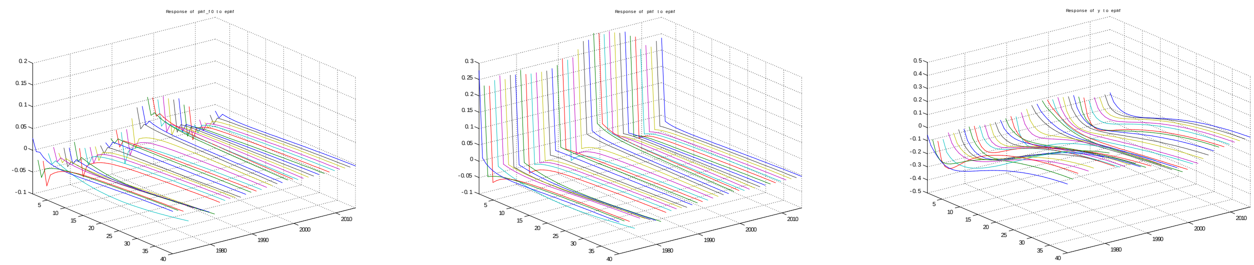

ε^{piid}

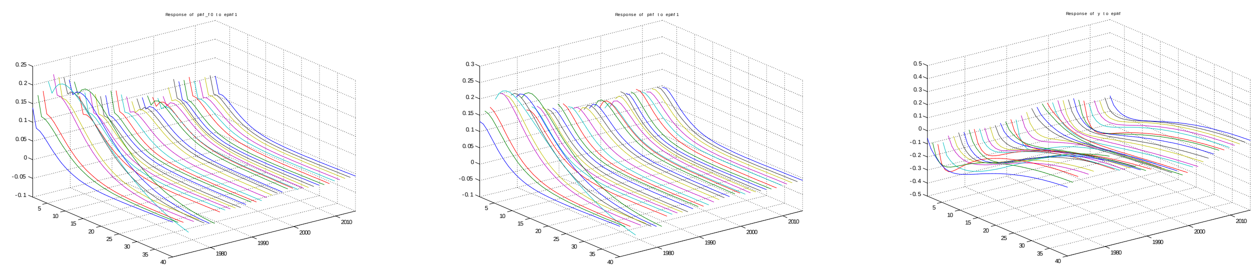

ε^{pap}

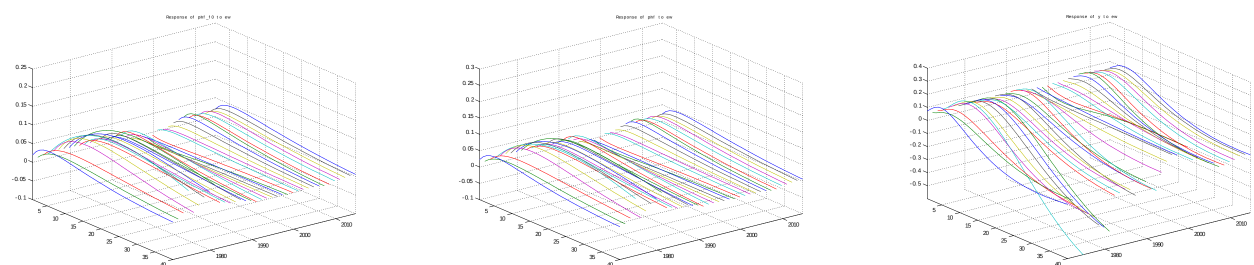

ε^{wap}

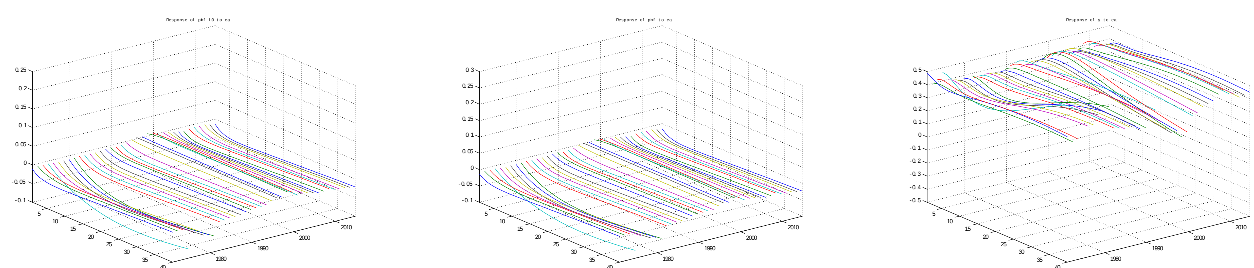

ε^{a}

Note:

The time variation in the irfs of the various shocks on inflation expectations and inflation realisation follow a similar time profile. The updating in the perceived inflation persistence and in the impact coefficients of the persistent markup innovations - which are updating in the same direction - are crucial for these dynamics. High impact effects and high persistence in the inflation belief models explain the high sensitivity and persistence in the seventies and the gradual moderation in the response later on. Inflation expectations and the actual inflation response are generally consistent with each other which is not guaranteed automatically in an AL-context. The time-profile in the irf of the iid markup shock deviates from the other shocks. The belief coefficient on this shock and the first order autoregressive coefficient (ρ_1) are responsible for this specific time variation. This shock, which is important for high frequency inflation fluctuations, explains the second peak between 1995 and 2005, in the variance of the one-period ahead inflation uncertainty.

Table 12: Conditional variance decomposition for the 10obs-2MU-AL model

	ε^a	ε^b	ε^g	ε^{qs}	ε^m	ε^{piid}	ε^{wiid}	ε^{par}	ε^{war}	$\xi^{\pi f1}$
1 quarter horizon										
π_f0	0.65	0.08	0.00	0.00	0.02	3.72	1.93	92.94	0.03	0.60
π_r1	0.19	0.02	0.00	0.00	0.00	81.18	0.57	18.03	0.01	0.00
w	0.02	0.34	0.02	0.00	0.07	15.33	82.67	0.06	1.48	0.00
y	15.77	63.96	2.20	0.66	13.26	0.21	0.35	3.59	0.01	0.00
1 year horizon										
π_f0	14.04	1.53	0.00	0.00	0.39	1.63	10.59	81.07	0.47	0.26
π_r1	1.65	0.63	0.00	0.00	0.16	53.60	4.32	39.44	0.20	0.00
w	0.06	3.64	0.04	0.02	0.93	11.73	79.20	0.51	3.89	0.00
y	6.40	71.15	0.81	0.37	18.85	0.62	0.36	1.42	0.02	0.00
10 year horizon										
π_f0	17.29	19.27	0.42	0.44	8.05	3.27	17.55	31.89	1.75	0.08
π_r1	11.66	12.99	0.29	0.30	5.42	29.03	11.84	27.29	1.18	0.00
w	0.55	40.76	0.06	0.03	17.04	6.17	29.00	3.64	2.75	0.00
y	11.19	58.44	2.25	0.14	23.20	0.95	1.13	2.61	0.11	0.00

Note:

The impact of the wage markup shocks on inflation and inflation expectations are highly reduced relative to the RE-model. Most important, the effects of these shocks are now transitory while they were responsible for the long term inflation trend in the RE-model. In fact, in this AL model, the long term inflation trend is no longer explained by exogenous shocks. It is the learning about the constant inflation rate in the belief equation that explains this long term trend. This means that all shocks can contribute to this long term inflation expectations depending on how the updating in the inflation beliefs is affected. This type of non-linear interactions is not taken into account in the conditional variance decomposition reported in Table 12. as it is constructed based on the initial beliefs only.[23] Still this decomposition illustrates that for the short and medium term dynamics the AL and the RE model are giving a similar interpretation in terms of shock contributions, except that the wage shocks are less important. At the long forecast horizon, all shocks contribute now to the inflation variance and there is also a non-negligible role for demand shocks such as the risk premium and the monetary policy shocks. But the learning responses must be added on top of this static analysis and this becomes a highly non-linear process.[24]

In sum, the AL model provides a more informative analysis of the inflation dynamics than the RE model. Instead of explaining the long term inflation trend by exogenous shocks, it are now the perceptions and the updating in the beliefs that are crucial for inflation anchoring. This basic AL-result is robust across various specifications. For instance, adding an inflation target shock in this AL-model does not change the results as it did for the RE-model.

The AL model does relatively worse than the RE on forecasting inflation over multiple quarters. This might suggest that there is too much flexibility in the belief specification and updating. One solution for overcoming that problem is to add survey expectations about future quarters to the list of observables. Up to now, we only experimented with one additional observable for inflation two quarters ahead. The beliefs can either contain two separate belief equations for each of these forecasts (unrestricted mode) or the two-horizon forecast can be written as a restricted belief equation where the restriction imposes consistency with the one period ahead forecast coefficients.[25] In both cases, the results are promising in that they improve on the inflation forecast without distorting the other implications of the model. In future work, we also plan to add long term inflation expectations so as to further discipline the updating in the belief coefficients. By doing so, long run inflation expectations will no longer be purely backward looking via the updating process but there might be left some role for forward-looking expectation shocks related for instance to changes in the monetary or fiscal policy context

3 robustness exercises and extensions

- change the timing of the survey observable: use the lagged SPF-forecast for two periods ahead instead of the current SPF-forecast for one period

[23] As explained in learning setup, these beliefs are consistent with the RE-equilibrium for the given structural parameters. The first line in the IRF figure is reproducing exactly this situation. Evaluating the decomposition around this point is still representative as an average approximation.

[24] In SW2012, we illustrated how the learning response react to various shocks depending on the state of the economy.

[25] The RHS-variable in the two period ahead forecast can be substituted consistent with the other PLM equations for π_t^f, w_t, r_{Kt} (both enter via mc_t) and the exogenous shocks processes (for productivity and price markup shocks)

ahead. This current SPF-forecast can have an information advantage as it is collected while the quarter is already ongoing. Most results are still valid and there are still large gains from including survey data in the model with two markup shocks but the magnitude of these gains are smaller. Our positive results are not only explained by the timely information contained in the nowcasts.

- results are robust when using a broader real-time dataset, other price indicators, longer horizon expectations.
- evaluate systematically the interpretation of the low frequency shock: markup versus expectations or policy shock.
- illustrate how the 10obs-2MU-Re and AL models interpret the inflation data during the Great Recession, the recovery and most recent period: are the shocks similar in the two models, how is updating affected by the big shock in the recession and the repeated negative innovations in the recovery, what is the impact of the time-varying conditional variance on the predictive distribution.

4 Conclusions

- A proper integration of survey expectations - as measured by the SPF - in a DSGE model makes it possible to identify separately the transitory and the persistent shocks in inflation.
- By improving the efficiency of the model filter, the forecasts improve both for inflation and for other macrovariables.
- Under AL, the updating of the belief models that are augmented with the information signals from the survey data, generate time-varying estimates of the perceived inflation target, persistence and sensitivity to shocks.
- In this way, the model captures the joint dynamics in the first and second moments of realized and expected inflation.

Remaining questions:

- what is the appropriate specification of the belief models in AL: small but informative, gradual updating versus model switching?

- are expectations fully determined by backward looking updating or is there a role for forward looking expectation shocks: exploit term structure of survey/market expectations?

- analyze policy implications of this learning process: what happens if zlb is imposed on policy rule?

- test AL dynamics against reduced form models with time-varying coefficients and volatility.

- extend with other survey expectations on real side: consumption, investment, wages etc.

References

Ang, Andrew & Bekaert, Geert & Wei, Min, 2007. "Do macro variables, asset markets, or surveys forecast inflation better?," Journal of Monetary Economics, Elsevier, vol. 54(4), pages 1163-1212, May.

Branch, W.A., (2007): "Sticky information and model uncertainty in survey data on inflation expectations". Journal of Economic Dynamics and Control, 31, 245–276.

Cochrane J. (2016), "Do Higher Interest Rates Raise or Lower Inflation?", mimeo,

De Graeve, Ferre & Emiris, Marina & Wouters, Raf, 2009. "A structural decomposition of the US yield curve," Journal of Monetary Economics, Elsevier, vol. 56(4), pages 545-559, May.

Del Negro, Marco & Eusepi, Stefano, 2011. "Fitting observed inflation expectations," Journal of Economic Dynamics and Control, Elsevier, vol. 35(12), pages 2105-2131.

Del Negro M., M. Giannoni and F.Schorfheide (2015), "Inflation in the Great Recession and New Keynesian Models," American Economic Journal: Macroeconomics, American Economic Association, vol. 7(1), pages 168-96, January

Del Negro, Marco & Schorfheide, Frank, 2013. "DSGE Model-Based Forecasting," Handbook of Economic Forecasting, Elsevier.

Eusepi S, E.Moench, B.Preston and C.Carvalho (2015à, "What drives long-run inflation expectations?," 2015 Meeting Papers 1228, Society for Economic Dynamics.

Evans C. (2016), "Connecting the Dots on Monetary Policy, Federal reserve bank of Chicago, speech delivered on January 13, 2016

Evans, G., and S. Honkapohja (2013): "Learning as a Rational Foundation for Macroeconomics and Finance", Ch. 2 in Rethinking Expectations: The Way Forward for Macroeconomics, Roman Frydman and Edmund S. Phelps (eds.), Princeton University Press.

Evans, G. and B. McGrough (2016), "Interest Rate Pegs in New Keynesian Models", April 2016.

Evans, G, K. Mitra and S. Honkapohja (2016), "Expectations, Stagnation and Fiscal Policy", May 2016

García-Schmidt M. and M. Woodford, 2015. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," NBER Working Papers 21614, National Bureau of Economic Research, Inc

Gorodnichenko Y and O.Coibion (2015),"Is The Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation" American Economic Journal – Macroeconomics 7(2015), 197-232

Hommes, Cars & Zhu, Mei, 2014. "Behavioral learning equilibria," Journal of Economic Theory, Elsevier, vol. 150(C), pages 778-814.

Mankiw, N. G., and R. Reis, (2003): "Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve." Quarterly Journal of Economics 117 (4), 1295-1328.

Milani, F. (2011): "Expectation Shocks and Learning as Drivers of the Business Cycle", The Economic Journal, 121 (May), 379–401.

Milani F. and A.Rajbhandari (2012) "Observed Expectations, News, and the Business Cycle", mimeo.

Monti, F. (2010): “Combining Judgment and Models”, Journal of Money, Credit, and Banking, Volume 42, Issue 8.

Nunes, R. (2010), "Inflation Dynamics: The Role of Expectations", Journal of Money, Credit and Banking, Vol. 42, No. 6

Ormeno A. and K. Molnar (2015), "Using Survey Data of Inflation Expectations in the Estimation of Learning and Rational Expectations Models", Journal of Money, Credit and Banking, Vol. 47, No. 4

Slobodyan S. and R.Wouters (2012), "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," American Economic Journal: Macroeconomics, American Economic Association, vol. 4(2), pages 65-101, April.

Smets, F., Warne, A. and R. Wouters (2014): "Professional forecasters and real-time forecasting with a DSGE model," International Journal of Forecasting, Elsevier, vol. 30(4), pages 981-995

Smets F. and R. Wouters (2007), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," American Economic Review, American Economic Association, vol. 97(3), pages 586-606, June.

Smith, G.W. (2009), "Pooling forecasts in linear rational expectations models", Journal of Economic Dynamics & Control, 33(2009), p.1858–1866.

Stock J. and M. Watson (2007), "Why Has U.S. Inflation Become Harder to Forecast?," Journal of Money, Credit and Banking, Blackwell Publishing, vol. 39(s1), pages 3-33, 02.

Woodford, M. (2013): “Macroeconomic Analysis without the Rational Expectations Hypothesis,” Annual Review of Economics.

Appendix

Table A1: Prior and posterior distributions 9obs models.

Parameter		Prior distribution			9obs-RE Metropolis Chain			SW2007 Posterior	9obs-AL Metropolis Chain			SW2012 Posterior
		type	mean	std.dev.	mean	5%	95%	mode	mean	5%	95%	mode
Calvo prob. wages	ξ_w	B	0.50	0.10	0.78	0.68	0.88	0.73	0.80	0.75	0.84	0.84
Calvo prob. prices	ξ_p	B	0.50	0.10	0.70	0.62	0.78	0.65	0.74	0.68	0.80	0.65
Indexation wages	ι_w	B	0.50	0.15	0.56	0.33	0.77	0.59	0.25	0.12	0.38	0.21
Indexation prices	ι_p	B	0.50	0.15	0.16	0.06	0.25	0.22	0.46	0.27	0.65	0.19
Gross price markup	ϕ_p	N	1.25	0.12	1.58	1.46	1.71	1.61	1.53	1.40	1.66	1.56
Capital production share	α	N	0.30	0.05	0.19	0.16	0.22	0.19	0.18	0.15	0.22	0.17
Capital utilization cost	ψ	B	0.50	0.15	0.71	0.56	0.87	0.54	0.56	0.34	0.76	0.56
Investment adj. cost	φ	N	4.00	1.50	4.25	2.58	5.82	5.48	3.24	2.11	4.33	3.23
Habit formation	$\varkappa$	B	0.70	0.10	0.65	0.51	0.79	0.71	0.64	0.53	0.76	0.68
Inv elast of subst.cons.	σ_c	N	1.50	0.37	1.51	1.08	1.97	1.59	1.76	1.24	2.16	1.58
Labor supply elast.	σ_l	N	2.00	0.75	1.77	0.97	2.57	1.92	2.22	1.58	2.89	1.77
Log hours worked in S.S.	$\bar{l}$	N	0.00	2.00	1.27	-1.01	3.57	-0.10	2.69	1.04	4.54	0.83
Discount factor	$100(\beta^{-1}-1)$	G	0.25	0.10	0.18	0.08	0.28	0.16	0.18	0.07	0.28	0.17
Quarterly Growth in S.S.	$\bar{\gamma}$	N	0.40	0.10	0.32	0.29	0.35	0.43	0.40	0.35	0.45	0.41
Stationary tech. shock	ρ_a	B	0.50	0.20	0.98	0.97	0.99	0.95	0.99	0.98	1.00	0.99
Risk premium shock	ρ_b	B	0.50	0.20	0.47	0.23	0.70	0.18	0.58	0.38	0.74	0.55
Invest. spec. tech. shock	ρ_i	B	0.50	0.20	0.84	0.77	0.92	0.71	0.48	0.38	0.58	0.51
Gov't cons. shock	ρ_g	B	0.50	0.20	0.99	0.99	1.00	0.97	0.99	0.99	1.00	0.97
Price markup shock	ρ_p	B	0.50	0.20	0.96	0.93	0.99	0.90				
Wage markup shock	ρ_w	B	0.50	0.20	0.94	0.90	0.99	0.97				
Response of g_t to ε_t^a	ρ_{ga}	B	0.50	0.20	0.59	0.43	0.74	0.52	0.64	0.49	0.78	0.54
Stationary tech. shock	σ_a	G	0.20	0.15	0.44	0.39	0.48	0.45	0.44	0.40	0.50	0.46
Risk premium shock	σ_b	G	0.20	0.15	0.18	0.11	0.25	0.24	0.20	0.15	0.24	0.15
Invest. spec. tech. shock	σ_i	G	0.20	0.15	0.37	0.31	0.43	0.45	0.43	0.37	0.48	0.45
Gov't cons. shock	σ_g	G	0.20	0.15	0.53	0.48	0.58	0.52	0.51	0.46	0.56	0.50
Price markup shock	σ_p	G	0.20	0.15	0.18	0.14	0.21	0.14	0.20	0.19	0.23	0.15
MA(1) price markup shock	ϑ_p	B	0.50	0.20	0.82	0.73	0.90	0.74				
Wage markup shock	σ_w	G	0.20	0.15	0.40	0.35	0.44	0.24	0.36	0.32	0.41	0.23
MA(1) wage markup shock	ϑ_w	B	0.50	0.20	0.89	0.83	0.97	0.88				
Quarterly infl. rate. in S.S.	$\bar{\pi}$	G	0.62	0.10	0.76	0.59	0.93	0.81	0.59	0.44	0.73	0.64
Inflation response	r_π	N	1.50	0.25	1.59	1.35	1.84	2.03	1.55	1.19	1.87	1.75
Output gap response	r_y	N	0.12	0.05	0.05	0.03	0.08	0.08	0.10	0.05	0.15	0.15
Diff. output gap response	$r_{\Delta y}$	N	0.12	0.05	0.17	0.13	0.20	0.22	0.12	0.09	0.15	0.14
Mon. pol. shock std	σ_r	G	0.20	0.15	0.24	0.21	0.26	0.24	0.22	0.20	0.24	0.22
Mon. pol. shock pers.	ρ_r	B	0.50	0.20	0.09	0.02	0.16	0.12	0.12	0.03	0.19	0.10
Interest rate smoothing	ρ_R	B	0.75	0.10	0.81	0.78	0.85	0.81	0.93	0.89	0.96	0.89
m.e. π_r1	$\sigma_{\pi r1}$	G	0.20	0.15	0.11	0.10	0.12		0.11	0.10	0.12	
m.e. dy_r1	σ_{yr1}	G	0.20	0.15	0.21	0.19	0.22		0.21	0.19	0.23	
Learning persistence	φ	U	0.00	1.00					0.99	0.98	1.00	0.97
Log marginal likelihood					MCMC	−965.22			MCMC	-943.42		

Note: models are evaluated over the period 1971Q1 - 2015Q3 using the first four observations as presample.

Table A2: Prior and posterior distributions 10obs-ME models.

Parameter		Prior distribution			10obs-RE Metropolis Chain			10obs-AL Metropolis Chain		
		type	mean	std.dev.	mean	5%	95%	mean	5%	95%
Calvo prob. wages	ξ_w	B	0.50	0.10	0.80	0.74	0.87	0.84	0.78	0.89
Calvo prob. prices	ξ_p	B	0.50	0.10	0.87	0.80	0.93	0.76	0.69	0.83
Indexation wages	ι_w	B	0.50	0.15	0.50	0.29	0.71	0.26	0.10	0.41
Indexation prices	ι_p	B	0.50	0.15	0.13	0.05	0.21	0.47	0.32	0.63
Gross price markup	ϕ_p	N	1.25	0.12	1.50	1.37	1.63	1.49	1.36	1.62
Capital production share	α	N	0.30	0.05	0.17	0.14	0.20	0.17	0.14	0.21
Capital utilization cost	ψ	B	0.50	0.15	0.65	0.47	0.82	0.52	0.27	0.77
Investment adj. cost	φ	N	4.00	1.50	3.86	2.12	5.82	2.65	1.66	3.84
Habit formation	$\varkappa$	B	0.70	0.10	0.58	0.44	0.71	0.60	0.51	0.69
Inv elast of subst.cons.	σ_c	N	1.50	0.37	1.05	0.73	1.41	1.77	1.45	2.11
Labor supply elast.	σ_l	N	2.00	0.75	1.60	0.79	2.42	1.96	1.24	2.69
Log hours worked in S.S.	$\bar{l}$	N	0.00	2.00	0.18	-2.22	2.73	2.87	1.20	4.62
Discount factor	$100(\beta^{-1}-1)$	G	0.25	0.10	0.18	0.08	0.28	0.18	0.07	0.28
Quarterly Growth in S.S.	$\bar{\gamma}$	N	0.40	0.10	0.35	0.31	0.38	0.40	0.36	0.43
Stationary tech. shock	ρ_a	B	0.50	0.20	0.98	0.97	0.99	0.99	0.98	1.00
Risk premium shock	ρ_b	B	0.50	0.20	0.91	0.86	0.99	0.65	0.54	0.76
Invest. spec. tech. shock	ρ_i	B	0.50	0.20	0.69	0.54	0.84	0.38	0.23	0.54
Gov't cons. shock	ρ_g	B	0.50	0.20	0.99	0.98	1.00	0.99	0.99	1.00
Price markup shock	ρ_p	B	0.50	0.20	0.93	0.87	0.98			
Wage markup shock	ρ_w	B	0.50	0.20	0.99	0.98	1.00			
Response of g_t to ε_t^a	ρ_{ga}	B	0.50	0.20	0.61	0.46	0.76	0.62	0.47	0.76
Stationary tech. shock	σ_a	G	0.20	0.15	0.35	0.31	0.38	0.45	0.41	0.49
Risk premium shock	σ_b	G	0.20	0.15	0.08	0.04	0.10	0.17	0.13	0.22
Invest. spec. tech. shock	σ_i	G	0.20	0.15	0.38	0.31	0.45	0.42	0.36	0.47
Gov't cons. shock	σ_g	G	0.20	0.15	0.52	0.47	0.57	0.51	0.47	0.55
Price markup shock	σ_p	G	0.20	0.15	0.21	0.19	0.23	0.24	0.21	0.26
MA(1) price markup shock	ϑ_p	B	0.50	0.20						
Wage markup shock	σ_w	G	0.20	0.15	0.43	0.39	0.47	0.37	0.34	0.41
MA(1) wage markup shock	ϑ_w	B	0.50	0.20						
Quarterly infl. rate. in S.S.	$\bar{\pi}$	G	0.62	0.10	0.66	0.52	0.81	0.68	0.55	0.79
Inflation response	r_π	N	1.50	0.25	1.73	1.45	2.04	1.69	1.40	1.97
Output gap response	r_y	N	0.12	0.05	0.05	0.00	0.09	0.09	0.06	0.13
Diff. output gap response	$r_{\Delta y}$	N	0.12	0.05	0.19	0.15	0.22	0.13	0.09	0.16
Mon. pol. shock std	σ_r	G	0.20	0.15	0.24	0.22	0.27	0.22	0.20	0.24
Mon. pol. shock pers.	ρ_r	B	0.50	0.20	0.05	0.01	0.09	0.11	0.03	0.19
Interest rate smoothing	ρ_R	B	0.75	0.10	0.83	0.79	0.87	0.89	0.86	0.93
m.e. π_r1	$\sigma_{\pi r1}$	G	0.20	0.15	0.11	0.10	0.12	0.11	0.10	0.12
m.e. dy_r1	σ_{yr1}	G	0.20	0.15	0.21	0.19	0.23	0.21	0.19	0.22
m.e. π_f1	$\sigma_{\pi f1}$	G	0.20	0.15	0.15	0.14	0.17	0.18	0.16	0.19
Learning persistence	φ	U	0.00	1.00				0.98	0.97	1.00
Log marginal likelihood					MCMC	-910.87		MCMC	-885.94	

Note:

Table A3: Prior and posterior distributions 10obs-2MU models.

Parameter		Prior distribution			10obs-RE Metropolis Chain			10obs-AL Metropolis Chain		
		type	mean	std.dev.	mean	5%	95%	mean	5%	95%
Calvo prob. wages	ξ_w	B	0.50	0.10	0.79	0.73	0.86	0.89	0.85	0.92
Calvo prob. prices	ξ_p	B	0.50	0.10	0.91	0.87	0.94	0.83	0.78	0.88
Indexation wages	ι_w	B	0.50	0.15	0.38	0.18	0.57	0.25	0.10	0.39
Indexation prices	ι_p	B	0.50	0.15	0.07	0.03	0.11	0.06	0.03	0.10
Gross price markup	ϕ_p	N	1.25	0.12	1.45	1.33	1.56	1.50	1.38	1.61
Capital production share	α	N	0.30	0.05	0.20	0.17	0.23	0.18	0.15	0.21
Capital utilization cost	ψ	B	0.50	0.15	0.69	0.453	0.85	0.72	0.57	0.87
Investment adj. cost	φ	N	4.00	1.50	1.85	1.00	2.67	1.72	1.24	2.17
Habit formation	$\varkappa$	B	0.70	0.10	0.46	0.35	0.57	0.52	0.45	0.59
Inv elast of subst.cons.	σ_c	N	1.50	0.37	1.35	1.03	1.66	1.43	1.26	1.59
Labor supply elast.	σ_l	N	2.00	0.75	1.58	0.75	2.39	1.83	1.06	2.60
Log hours worked in S.S.	$\bar{l}$	N	0.00	2.00	0.54	-1.15	2.25	3.37	2.37	4.37
Discount factor	$100(\beta^{-1}-1)$	G	0.25	0.10	0.17	0.07	0.27	0.18	0.09	0.27
Quarterly Growth in S.S.	$\bar{\gamma}$	N	0.40	0.10	0.36	0.32	0.39	0.39	0.35	0.42
Stationary tech. shock	ρ_a	B	0.50	0.20	0.98	0.97	0.99	0.99	0.98	1.00
Risk premium shock	ρ_b	B	0.50	0.20	0.83	0.74	0.94	0.80	0.73	0.88
Invest. spec. tech. shock	ρ_i	B	0.50	0.20	0.86	0.77	0.95	0.39	0.29	0.49
Gov't cons. shock	ρ_g	B	0.50	0.20	0.99	0.98	0.99	0.99	0.99	1.00
Price markup shock	ρ_{par}	B	0.50	0.20	0.77	0.59	0.93	0.79	0.68	0.91
Wage markup shock	ρ_{war}	B	0.50	0.20	1.00	0.99	1.00	0.43	0.09	0.77
Response of g_t to ε_t^a	ρ_{ga}	B	0.50	0.20	0.65	0.50	0.80	0.65	0.50	0.79
Stationary tech. shock	σ_a	G	0.20	0.15	0.45	0.40	0.49	0.44	0.39	0.48
Risk premium shock	σ_b	G	0.20	0.15	0.11	0.08	0.15	0.12	0.11	0.14
Invest. spec. tech. shock	σ_i	G	0.20	0.15	0.44	0.40	0.54	0.32	0.27	0.24
Gov't cons. shock	σ_g	G	0.20	0.15	0.52	0.47	0.57	0.51	0.46	0.55
Price markup shock-iid	σ_{piid}	G	0.20	0.15	0.24	0.22	0.26	0.26	0.23	0.28
Price markup shock-ar	σ_{par}	G	0.20	0.15	0.03	0.01	0.05	0.03	0.02	0.04
Wage markup shock-iid	σ_{wiid}	G	0.20	0.15	0.43	0.39	0.47	0.39	0.35	0.43
Wage markup shock-ar	σ_{wae}	G	0.20	0.15	0.01	0.00	0.01	0.04	0.01	0.08
Quarterly infl. rate. in S.S.	$\bar{\pi}$	G	0.62	0.10	0.62	0.46	0.77	0.60	0.50	0.69
Inflation response	r_π	N	1.50	0.25	1.46	1.17	1.73	1.65	1.37	1.93
Output gap response	r_y	N	0.12	0.05	0.13	0.08	0.17	0.05	0.02	0.08
Diff. output gap response	$r_{\Delta y}$	N	0.12	0.05	0.23	0.19	0.27	0.15	0.12	0.19
Mon. pol. shock std	σ_r	G	0.20	0.15	0.24	0.22	0.27	0.22	0.20	0.24
Mon. pol. shock pers.	ρ_r	B	0.50	0.20	0.06	0.01	0.10	0.11	0.03	0.19
Interest rate smoothing	ρ_R	B	0.75	0.10	0.87	0.84	0.91	0.90	0.87	0.94
m.e. π_r1	$\sigma_{\pi r1}$	G	0.20	0.15	0.11	0.10	0.12	0.11	0.10	0.12
m.e. dy_r1	σ_{yr1}	G	0.20	0.15	0.21	0.19	0.22	0.21	0.19	0.22
m.e. π_f1	$\sigma_{\pi f1}$	G	0.20	0.15	0.04	0.01	0.06	0.02	0.01	0.04
Learning persistence	φ	U	0.00	1.00				0.97	0.96	0.98
Log marginal likelihood					MCMC	-840.90		MCMC	-790.51	

Note: