

Time-Series Properties of Inflation Expectations

Felipe Schwartzman *Economics Department
Fisher Hall, Princeton University
Email: fschwart@princeton.edu

March 26, 2006

1 Introduction

At least since Keynes, economists have acknowledged that a completely specified macroeconomic theory must take a stand on how agents form their expectations concerning the future behavior of economic variables. This point was made most forcefully by Lucas (1976), who showed in his critique how an inadequate theory of expectations formation can lead to bad policy advice. Lucas's proposal was that any macroeconomic model to be used in policy evaluation should take expectations to be rational, as proposed by Muth (1961). From that time on, rational expectations have formed a fundamental building block of virtually all macro models.

The Rational Expectations Hypothesis (hence REH) has largely been kept as maintained hypothesis in most macro models in spite of the fact that it has fared badly in the many different empirical implementations performed over a large span of time. On the one hand, the simplest rational expectations based models have had trouble explaining basic facts about consumption, inflation and asset pricing.¹ In effect, much of the theoretical research done in macroeconomics and finance from the 90's onwards could be described as attempts to reconcile rational expectations based models with these empirical regularities. In macroeconomics, the failure of rational expectations based models has shown itself in the difficulty of accounting

*I thank Ricardo Reis for providing me with the data and to Chris Sims for the detailed and useful advice on the econometric techniques. Thanks also to Dante Amengual for careful comments

¹For surveys that document the difficulties of such models see respectively Deaton (1992), Mankiw (2002), and Shleifer (2002).

for “stickiness” of various variables without appeal for ad-hoc adjustment costs or fixes in the preferences of agents (Sims, 1998).

Apart from these indirect problems, the REH has not fared well in more direct empirical tests using survey data on expectations. Such tests typically take the average or the median forecast made by a panel of interviewees that may or may not be specialists, and test whether this behaves as rational expectations should. In particular, expectations should be an unbiased predictor of the corresponding variable and it should also be a sufficient statistic at the time of the prediction, meaning that the prediction error should not be correlated with any variable known at the time the prediction was made or before. Section 2 includes a description of how these tests are typically done as well as a description of the typical results.

More recently, a number of theories have been put forward to substitute for the REH. These include theories that assume that agents have limited information processing capacity (Sims, 2003), that they update information only infrequently (Reis, 2003), that agents deceive themselves into being overoptimistic, thus reaping “antecipatory utility” (Brunnemeier and Parker, 2004), or that even if their average expectation is rational, the existence of disagreement means that, given strategic complementarities and lack of common knowledge about the state of the economy, agents will react more sluggishly to shocks (Woodford 2001). These theories are mostly justified based on their a priori reasonableness and in their ability to take account in an economical way of facts that one can only explain with REH if one is willing to greatly complicate one’s model. Given the recent interest in formulating alternative theories of expectations formation, it seems that taking a closer look on how expectations taken from survey data behave may be a good idea. The literature using survey data to test the REH is of little help, as it’s main thrust is to test the maintained hypothesis of rational expectations, rather than checking what other empirical facts could be found that could inform alternative theories of expectations formation. The latter is what this paper purposes to do. It does not present any formal test of a specific model, rather the work is exploratory in spirit. The aim here is to find stylized facts about expectations that could serve to inform future theoretical research as well as inspire more precise tests the research done to date.

The data used is the Michigan Survey of inflation expectations. This survey is interesting because it contains monthly data since 1978 and a fairly large panel of interviewees. While in other studies the use of this survey has been criticised because the forecasters are not professionals, and hence do not have incentives to make the best possible forecast they can, arguably for

the purposes of this paper this is rather a strength than a weakness. What one wants is to know how agents actually form their expectations, not how they would form if they were putting the best effort they could. In effect, as pointed out above, the fact that agents will not put the best effort in forming their expectations is a basic tenant of many of the new approaches. Moreover, professional forecasts may be subject to other sources of bias due to agency problems (see Lamont 2002), as professional will be evaluated based on how well their forecasts fare in comparison to the forecasts of other professionals.

The approach taken is to estimate a Bayesian VAR adding survey expectations data to the variables usually included in the analysis of monetary policy. VAR's present a flexible econometric framework for the analysis of multivariate time-series data with minimal use of theoretical assumptions. They have been widely used in the analysis of macro-economic data and monetary policy. Prominent examples include Bernanke and Blinder (1992), Bernanke and Mihov (1998). The best example of the use of a Bayesian VAR for the analysis of monetary policy is Leeper, Sims and Zha (1996). These studies have been successful in finding widely agreed upon stylized facts about the impact of monetary policy on real variables.

This paper expands on this previous research by including the survey expectations variables. In particular, I included median expectations and the cross-sectional inter-quartile range. Median expectations are typically the statistic used in tests of the REH. This is because by its very nature, REH does not allow for heterogeneity in expectations, but, as will be discussed below, under certain assumptions average expectations can be taken to be equivalent to the expectation of the representative agent. The use of the median instead of the mean is advisable because the cross-section of the survey-data includes a small number of people with very high inflation expectations that cannot be reasonably assumed to follow from their best calculations. These outliers tend to distort the mean, so that median expectations become a better measure of the central tendency of the cross-sectional distribution of expectations. I included the inter-quartile range as a measure of the dispersion of these expectations. So far macro theories have little predictions concerning how this dispersion may affect macro-economic variables, but given that cross-sectional variability of expectations is a feature of the data not typically accounted for by traditional REH, inquiring on whether and how it interacts with other macro-variables is an interesting exercise.

The VAR can also be used to evaluate directly in which ways expectations deviate from REH. If all variables included in the VAR are observed

at some time t , the s periods ahead inflation predicted by the VAR should always coincide with s period ahead inflation expectation quoted at time t . This follows from the fact that inflation expectations should be a sufficient statistic for the predicted inflation. By calculating the difference between inflation expectations and inflation predicted by the VAR one has a measure of the deviation from rational expectations.²

2 Tests of Rational Expectations

In its simplest form, the REH as proposed by Muth (1961) postulates that if all public information available at time t is Ω_t , then expectations have to be such that

$$\pi_{t+s|t}^e = E_t [\pi_{t+s} | \Omega_t]$$

where $\pi_{t+s|t}^e$ denotes the expectation formed as of time t concerning for inflation at time $t + s$, and E_t is the mathematical expectations operator.³ The underlying assumptions are that the unpredictable behavior of the economy can be described by a probabilistic model, that this model only involves variables observed by all agents, that agents know this model and use it in their forecasts.⁴

When testing rational expectations most (but not all) of the literature relies on aggregate expectational data. Typically one will use the average or median expectation of some survey. Under REH in its simplest form, stated above, using such a statistic would be equivalent to taking the expectations of a single agent, as agents should all have the same expectations. Of course this is typically not the case as one can readily see by inspecting the dispersion of expectations in any such survey. Alternatively, one can admit that agents have private information but that on average the effect of those will cancel out so that average expectations coincide with REH, but it may be difficult to justify such a restriction in the information set of the agents on a priori grounds. In any case, given that most REH based macro models and

²In practice, of course, not all variables at time t are observed at that time, but they will become known in one or two months. This fact is also taken into account in the empirical implementation below.

³The theoretical discussion here follows mainly Pesaran and Weale (2005). The reader is referred to that paper for further details and variations on the basic model. For a survey of the literature testing REH with use of survey data, see Thomas (1999).

⁴Another more technical and less easily interpretable assumption is that, when reporting their forecasts, agents face a quadratic loss function, that is, they want to minimize the square distance between their forecast and the realized value. This is necessary for their forecast to be equal to the expected value.

econometrics still think as the “expectations” of the agents as something one can calculate from publicly observable variables understanding how the central tendency of expectations deviate from the REH is an interesting exercise.

As stated above, often one will consider the median expectations instead of average expectations. This can make a difference given that often averages can be affected by a small number of extreme expectations. This seems to be the case with the Michigan survey, used in this paper. Because the purpose of this paper is not so much test the validity of REH but provide a description of how expectations behave, the focus on median forecast and the way it deviates from the expectations one would have based on an econometric model based on publicly available variables is still interesting on its own right.

Most of tests attempt to evaluate whether the forecasts are biased and whether they are efficient. Biasedness can be tested by running the following OLS regression:

$$\pi_{t+s} - \pi_{t+s|t}^e = \alpha + \varepsilon_t$$

Where ε_t is the error term. Expectations are considered unbiased if $\alpha = 0$. Note however, that while one may not be able to reject bias for the whole sample, the forecasts can be biased over smaller parts of the sample. Therefore it is not a very powerful test, as it only imposes a minimum, that is, that over a large enough sample, forecasters should not make persistent mistakes in a same direction. It turns out that whether or not one finds bias in the forecasts depends very much in the period taken into consideration. While when using samples stretching from the 60’s or 70’s to the late 90’s one finds little or no evidence for bias, Thomas (1999) points out that inflation forecasts tend to be biased downwards in sub-periods where inflation is rising and biased upwards when they are falling, implying some kind of inertia in the forecasts.

The second set of tests typically done in the literature is whether forecasts are informationally efficient, that is, whether they incorporate all information available at the time the forecast is done. This is typically done by running the OLS regression below:

$$\pi_{t+s} - \pi_{t+s|t}^e = \alpha + \beta\pi_{t+s|t}^e + \gamma\mathbf{y}_t + \varepsilon_t$$

Where \mathbf{y}_t is a vector of variables known at time t and γ is a conformable vector of coefficients. A particular strong rejection of efficiency occurs if $\beta \neq 0$, as this implies that not even the information available in the forecast itself is efficiently used by the forecasters. Efficiency also implies that $\gamma = 0$.

This kind of test typically rejects REH. More interesting for the purposes of this paper is that as pointed out by Ball and Croushore (1995) the sign of γ combined with knowledge of the direction in which different variables predict inflation can indicate whether agents are overreacting or under-reacting to information. While Ball and Croushore's results seem to point rather to a pattern of underreaction, the results in the literature tend not to be very conclusive, with the signs varying depending on the time sample, series and specification used. The contradictory results are not surprising given that typically this kind of test are performed by regressing forecast errors on a limited and varying set of variables, typically with a single lag for each of them. While under the null all variables included in the \mathbf{y}_t should have zero coefficients irrespective of which variables are included, once the null is rejected and one wants to do inference based on the results, one is vulnerable to omitted variable bias problems. This paper attempts to minimize this problem by running a full VAR system.

A further test inquires on whether forecast errors are correlated. More specifically, given an MA representation for inflation, it must be the case that

$$\pi_{t+12} - E_t\pi_{t+12} = \sum_{i=1}^{12} a_{12-i}\epsilon_{t+i}$$

Where ϵ are the iid innovations to the inflation process. This implies that

$$E [\pi_{t+12} - E_t\pi_{t+12}] [\pi_{t+24} - E_{t+12}\pi_{t+24}] = 0$$

This implication of the REH is tested among others by Mankiw, Reis and Wolfers (2004) and the typical finding is that forecasts errors are positively correlated, implying that expectational errors are persistent. This points to an excessive sluggishness of expectations, that will be further explored in subsequent sections.

Also, this paper includes an analysis of the time series property of the dispersion of expectations. This aspect of the data is not typically taken into account by most studies of survey expectations, or is taken as a proxy for the uncertainty associated with the forecasts. Such an interpretation is of course fragile, as individuals could, for example have wide disagreements while still being very confident on their own opinions. Mankiw et al. (2004) provide a survey of this literature and find that the findings are not particularly robust.

There is very little more recent work on dispersion of expectations. An important exception is Mankiw et al. (2004) itself. In this paper the authors evaluate how a specific alternative model of inflation expectations fits

survey data and look specifically at the dispersion of expectations with this purpose. The authors find among other things a strong relationship between dispersion in expectations and the level of inflation, as well as the size of changes in inflation. They also find a strong link between the dispersion in expectations and the dispersion in the change in prices of different products. Finally, they find a somewhat weaker relation with measures of real activity.

3 Data

The data used are monthly US series of the Federal Funds rate, unemployment, M1, CPI, Industrial Production, a Commodity Price Index, SP-500 index as well as the median and the Interquartile range of the Michigan Survey of Consumer Attitudes and Behavior.⁵ The sample period goes from 1/1978 to 8/2001. The VAR is estimated with 12 lags, so that the first 12 months taken as initial conditions.

The choice of the Michigan Survey over other possibilities as the Survey of Professional Forecasters and the Livingston Survey was based on two sets of considerations. First, the Michigan Survey is the only one which is available with monthly frequencies. Given the nature of the identifying restrictions to be discussed below, working with monthly data seems a more natural thing to do. The other reason is that arguably consumer expectations are the more relevant for the understanding of macroeconomic dynamics and monetary policy. This is true first because consumers form a very large part of the relevant decision makers and second because expectations made by professional forecasters can be subject to biases and "herding effects" caused by agency considerations (Lamont 2002).

One criticism often made about the use of the Michigan Survey is that because consumers are not normally concerned about forecasting inflation they are less likely to give a well thought answer. Arguably the question is not so much whether they would be able to make better forecasts given more incentives to think through, but whether they actually base their decisions on an inflation outlook compatible with their answers. This is not an easy question to answer but my working hypothesis is that they do. In particular, as shown by Thomas (1999), the median forecast of the Michigan Survey does a very good job in predicting average inflation, which probably means that it is not artificially detached from Rational Expectations.

⁵I am grateful to Ricardo Reis for providing me with the latter two series.

4 Preliminary Analysis of data

Figure 1 below portrays the survey forecast series together with actual inflation in that month and the VAR 12 step ahead inflation forecasts.⁶ The most noticeable feature is that the actual inflation expectations do a much poorer job in predicting the fast rise in inflation in the early part of the sample than either of the VAR based forecasts. It is interesting to note that survey expectations become less volatile after the '90s, perhaps reflecting the increased credibility of monetary policy.

Figure 2 portrays the log spectral density plots for the series shown in Figure 1. Spectral density plots show what portion of the variation of the data can be assigned to cycles of different frequencies.⁷ One can readily confirm first that most forecast series have less variance in almost all frequencies than the actual series with exception of the very high frequencies.⁸ This reduced variation is to be expected, as the actual series can be understood as the forecasts with extra innovations added to it. Another observation is that the variance of the median forecast is below that of the rational expectations forecast in almost all frequencies, perhaps a sign that the median forecast does not incorporate new information as efficiently.

Given the estimates one can construct an "irrational expectations" series as the deviation between actual survey expectations and the 12 period ahead forecast made by the VAR. The irrational expectations series is plotted below. They present mostly high frequency variation and are largely unbiased.

Of course one can explain the existence of this series without having to abandon the REH completely. For one thing one can admit that agents form their expectations rationally in the sense that they understand how the economy works and make the best use of the information available to them, but that they have access to private information and in such a way that their median expectations need not coincide with what one would expect based purely on publicly available data. While this is certainly a possibility, macroeconomic models are typically based on the more naive version that all relevant data is publicly available, so that deviations from this paradigm are worthwhile understanding.

⁶The forecasts used in this section were all done using the maximum posterior estimate of the Bayesian VAR. Maximum likelihood estimates generate similar results.

⁷These and all other spectral density estimates were estimated with a Bartlett window with width equal to half of the square root of the length of the series.

⁸VAR forecasts have excess variance on very high frequencies. This is a feature of the prior used, and disappears under a flat prior.

A second objection is that the econometrician has an advantage with respect to the survey respondents, in that some of the data that is available now was not available at the time. For example, price data are only released the month after shocks to prices are realized etc. This interpretation allows for deviations from naive REH but they would die out quickly. In particular, if they are released with k months of delay, the deviations from naive REH should be MA(k). Given the MA representation of the VAR system, the discrepancy would follow the following process.

$$E_t \pi_{t+12} - \pi_{t+12|t}^e = \sum_{j \in J} (a_{12,j} - a_{11,j}) \epsilon_{j,t} + \sum_{j \in J} (a_{13,j} - a_{12,j}) \epsilon_{j,t-1}$$

Where $\pi_{t+12|t}^e$ is the reported inflation expectation, J is the set of variables that consumers only observe with delay, $\epsilon_{j,t}$ are the shocks to variable j at time t , and the $a_{t,j}$ are the impulse response coefficients relative to shock j on prices at t periods after the impulse is given. The differences appear because the VAR provides forecasts for prices, but the survey data is on inflation. From this representation it is straightforward to get to the spectral density, which would be:

$$\widetilde{E_t \pi_{t+12} - \pi_{t+12|t}^e}(\omega) = \sum_{j \in J} (a_{12,j} - a_{11,j})^2 + \sum_{j \in J} (a_{13,j} - a_{12,j})^2 + 2 \sum_{j \in J} (a_{12,j} - a_{11,j})(a_{13,j} - a_{12,j}) \cos\left(\frac{2\pi\omega}{T}\right)$$

With T denoting the size of the sample size. Figure 3 presents the spectral density of the irrational expectations series and the one that would come out of a model of rational expectations in which consumers only had access to information about variables in the production sector and the dispersion of expectations after two periods, but then use them fully to make their forecasts. Such a delay is not allowed for interest rate, exchange rate, stock prices and commodity prices, as those are always available in real time.⁹

It is immediately visible that the empirical and the theoretical spectral densities are very different. First of all, the empirical irrational expectations series has its spectral density everywhere above the theoretical one, meaning that even after allowing for a delay of two months in data absorption there is too much variance in the deviations from rational expectations. Also, note that the predicted spectral densities have a concave shape, meaning that variation tends to be spread out fairly evenly through different frequencies, falling only slowly at higher frequencies. The empirical spectral density is convex, implying a much larger concentration of variation in low frequencies. This is a sign of the excessive sluggishness in expectations formation. The question of how sluggish expectations are and which shocks tend to be

⁹Allowing those to be observed only with delay keeps the qualitative results unchanged.

absorbed more slowly will be discussed more at length later on with the analysis of the impulse response plots.

5 Granger Causality Tests

A useful first way of looking at the data is to inquire whether certain variables are “Granger Causally Prior” (GCP) to others.¹⁰ In spite of the name, Granger Causality is not about causality per-se but about whether a certain group of variables can help to predict variables from another group beyond what lagged values of the predicted variables are able to do by themselves. In other words, it tests whether including new variables in the VAR system improves the fit of the equations for the initial set of variables. Failure of a variable to be GCP to another has much more to do with how much information the first carries about future values of the second than anything else, and is system specific as the controls will change as the system changes. For the purposes of this paper, performing such tests is interesting, as its main contribution is exactly to add extra variables to an otherwise standard VAR.

Table 1 below presents the result of bayesian model comparison and frequentist χ^2 likelihood ratio tests.¹¹ The results on the odds ratio tests are presented with different weights on the Minnesota priors, as the results are very sensitive to those.¹² The likelihood ratio was calculated with estimation under a flat prior, consistent with frequentist practice.

The Likelihood ratio test in general rejects the null of non-GCP in both directions. The bayesian odds ratio is very sensitive to the tightness parameter of the Minnesota prior. This prior states that the econometrician believes that all the variables in the system are very close to a random walk, meaning that a good forecast of the value of that variable at any point in the future is the current value of that variable. As discussed with more detail in the appendix, this and other priors are used to state that, given typical

¹⁰For a discussion of Granger Causality and Granger Causality tests see Hamilton (1994) and Sims (2005).

¹¹Frequentist tests tend to reject the null of no GCP too often in small samples. A metric that can be justified asymptotically from a Bayesian point of view for whether the likelihood ratio is large enough to justify the use of one or another model is Schwarz’s Bayesian Information Criterion (B.I.C.). However, for the application at hand, it turns out that this criterion is too stringent, pointing to lack of GCP in all directions even for pairs of variables that are widely believed to have a GCP relationship such as broad money (M1) and consumer prices. For this reason, the comparisons with this criterion were not reported.

¹²See Appendix for a detailed discussion of the different priors used.

initial conditions, the econometrician should not easily predict the variables in the system to have very complicated trend-like behavior. The tightness parameter regulates how strong the belief in that prior is. It turns out that the more one believes in these priors, the lower are the odds of no GCP in either direction. In particular, for intermediate values of these parameter the evidence points to the typical variables being GCP to the survey variables but not vice versa.^{13,14} This result means that the common practice in macroeconomic VARs of not including survey variables does not significantly change the result. It does not mean necessarily that expectations do not matter. It could mean that they do not have much variation which cannot be explained by usual macro time-series, so that adding them does not add much new information. The evidence of the usual macro variables being GCP to expectational variables supports this interpretation to some extent.

6 Impulse Response Functions

Another useful way of evaluating how variables interact overtime is to plot orthogonalized impulse response functions. These plot how different variables react to independent sources of variation in the vector of residuals. Absent further a priori restrictions, the correct interpretation is not necessarily structural, meaning that these sources of variation do not necessarily correspond to exogenous shifts in variables. But they do show patterns of behavior that are suggestive and allow for such interpretation given auxiliary assumptions. Also, impulse response plots are a somewhat less stringent concept of the way a variable impacts another than GCP. This is because impulse response plots allow for indirect effects. For example median inflation expectations may have little predictive power about stock prices, but may affect them through its effect on interest rates, prices and so on. Also, the effect can be non-linear, so that small initial impacts on few variables

¹³Therefore, if one is uncertain about what the best value for that parameter but thinks that it is probable that 1 is a good value, the results seem to indicate that there is some slight evidence in favor of predicting survey data with use of macroeconomic and financial data, but not vice versa. One could make this statement more formal by assigning a distribution to the precision hyper-parameter and then integrating over that distribution.

¹⁴A very similar result was found when the survey expectations variables were removed and GCP was evaluated for the interest rate vis a vis the other variables. This seems to indicate that if one is willing to interpret from these results that expectations have little effect over the other macro variables, than one should also have a similar view about the interest rate.

may amplify overtime as they feed back through the system. The disadvantage is that their interpretation is less clear-cut absent a priori restrictions on the interactions between different variables.

The orthogonalized impulse response functions were found by use of a triangular specification scheme. Given an ordering of the variables in the system, the first variable is regressed against past values of itself and of the other variables, the second variable is regressed against those same past variables plus the current value of the first variable and so on. So the second equation contains all observed variables on both sides of the first equation, which means that the error term of that equation must necessarily be independent to the error term of the first equation etc. By setting the error term of any equation to a value equal to one standard deviation in its variation, one can track how a typical orthogonal shock to some variable feeds back into the system.

To give such a specification a structural interpretation, one must believe that indeed some variables are not affected in a contemporaneous way by changes in other variables. The ordering is therefore very important. The way I did it was to follow the monetary policy VAR literature in putting first variables which reflect decisions made by the production side of the economy and which can be reasonably be expected to react slowly to changes in other variables, such as prices, industrial output, unemployment and, to some extent, broad money (M1). Second I put financial variables, which one can reasonably assume to react fast to other variables in the system, such as commodity prices, foreign exchange and SP-500 stock price index. The use of this set of variables is typically recommended because they are very rich in information about economic conditions, so that one can more reasonably expect the bulk of the relevant publicly available information to be included into the model. After those variables, I followed Bernanke and Blinder (1992) in putting the interest rate set by the Fed afterwards. This decision is not uncontroversial, given that one can reasonably expect financial markets to react fast to interest rates, but in general the impulse response plots seem to be fairly robust to that particular choice (see for example Leeper et al. (1996) for a comparison of different specification schemes). I put the survey expectations variables last, and in particular I put the median expectations last of all. This choice was made so that shocks to these variables can be reasonably interpreted as orthogonal to relevant information that is publicly observed. In other words this choice was made so that one can interpret shocks to median expectations as deviations from

rational expectations.¹⁵

The impulse response plots together with 66% probability bands¹⁶ are presented in figures 4, 5 and 6 and 7. Figure 4 presents the impulse response functions for the subset of variables that is typically included in analysis of monetary policy. A reader which is familiar with this literature will readily recognize that the impulse response functions have the typically expected behavior.¹⁷ This finding suggests that, consistent with the finding in the previous section, the survey expectations variables have little direct effect over the behavior of the other variables, so that shutting down possible feedbacks of the shocks to non-expectational variables through survey expectations does not present an important source of bias.

Figure 5 shows the interactions between shocks to the central tendency and to the dispersion of expectations. The shock to the dispersion generates a significant increase in the central tendency, but shocks to the central tendency do not have a very discernible effect on the dispersion. In any case, any interpretation of these results should be careful, since the ordering of the two is arbitrary and the shock to the dispersion is potentially correlated with changes in the central tendency.

Figure 6 is the most interesting for the purposes of this section. It shows how different variables react to expectational shocks. Shocks to median inflation seem to have an expansionary effect, with reduction of unemployment and a slight increase in industrial output. This is consistent with a model with wage rigidities, so that an expectation of higher inflation for given nominal wages is associated with more hiring. The effect on monetary policy is contractionary as one would expect, with an increase in the federal funds rate. M1 responds accordingly by a reduction. The reduction in M1 could also be credited to a Cagan style reduction of money demand due to increased inflation expectations. In line with the expansionary effect, stock prices increase. The nominal exchange rate responds with a depreciation,

¹⁵A similar reasoning motivated Bernanke and Blinder's (1992) decision to put the interest rate less, as they wanted to make sure that shocks to the interest rate were not motivated by new information available to the Central Bank, but instead represented exogenous variation in policy.

¹⁶Those bands imply that there is a 66% chance that a given point in the impulse response function will be inside the band. Interpreting them as stating that the whole impulse response function should be within the band with 66% probability is not correct absent further investigation. See Sims and Zha (1999) for a discussion on error bands for impulse-response functions.

¹⁷I am using the Broad Trade Weighted nominal exchange rate index provided at St. Louis Fed website, which is constructed in such a way that higher exchange rates correspond to appreciations.

which is also to be expected given some version of long run Purchasing Power Parity relationship.

Shocks to the dispersion of inflation expectations are also very interesting. They have a clearly contractionary effect, with increase in unemployment and decrease in industrial production. This effect cannot be easily associated with a supply or demand shock, given that prices do not respond in a significant way. Interestingly, shocks to the dispersion of inflation expectations are associated with a clearly discernible positive impact on stock prices. To the extent that the dispersion in inflation expectations reflects disagreement about stock returns, this effect is in line with Miller (1977) model of stock prices in the presence of divergence of opinion and short sales constraints.

Figure 7 shows how shocks to the different variables impact the dispersion and central tendency of inflation forecasts. The dispersion of inflation forecasts seems to respond positively to shocks to the price level, consistent with the finding in Mankiw et al. (2004). It responds negatively to shocks in the SP-500. Good news for the financial market seem to be associated with, at least temporarily, more agreement about inflation forecasts. The effect of the federal funds is to first create greater dispersion in inflation expectations, but then this seems to become permanently lower. Perhaps the initial dispersion could be associated with an initial learning phase in which agents are trying to figure out how to interpret the policy shock. If this is so, then learning takes around a year.

About the median expectation, the reaction has the sign one would typically expect given rational expectations, so a more careful analysis will have to compare it with these. Such an analysis will be taken in the next section.

7 Deviations Between Median Expectations and Rational Expectations

Figure 8 shows the maximum posterior estimate of the impulse response plots for median expectations overlaid with the impulse response plots for the REH 12 month inflation calculated using the VAR. Under REH both should coincide. This is simply a consequence of the law of iterated expectations as

$$E_t [E_{t+s} [p_{t+s+12} - p_{t+s}]] = E_t [p_{t+s+12} - p_{t+s}]$$

A quick glance at the graphs show that the responses seem to be quite distinct. In most cases the reaction of survey expectations is in the right

direction, but the basic pattern is one of under-reaction. Exceptions are the case of shocks to industrial output and to stock prices, where expectations move in the opposite direction. The only shock which is largely associated with a correct update in expectations are shocks to the nominal exchange rate. Shocks to M1 have very little impact over consumer forecasts, which is also an interesting fact. As for interest rates, while at first there is some under-reaction as in most other cases, on a longer run there is over-reaction, that is, after inflation has already settled down back to its trend people are still expecting the monetary policy shock to have effects. Also, as one would expect, there is over-reaction to the shock in median inflation expectations, supporting the interpretation of that component as reflecting a deviation from REH.

Figure 9 provides a more formal assessment of the deviation from rational expectations. It portrays the difference between the impulse response function for inflation and for inflation forecasts together with 66% and 95% error bands. Most deviations seem to be significantly different from zero at least in single points at a 95% level and most of them are significantly different from zero at the 66% level. It is also interesting to note that deviations tend to be more pronounced in the first 12 months after the shock, subsiding afterwards.

8 Discussion

The first interesting, and somewhat comforting result, is that survey variables do not seem to be important in explaining how the set of variables usually included in monetary policy VAR's interact among themselves. This is indicative that while expectational variables may have important indirect effects, they operate mainly through their effects in variables more traditionally taken into account by the literature. Also, this may indicate that variation in survey expectations are to a large extent endogenous to variation in the other variables, thus allowing one to think of theories to explain their variation over-time.

Also, to the extent that the identification strategy used is informative, it seems that exogenous shocks to both location and dispersion of expectations matter. The impact of shocks to the central tendency of inflation expectations behave like Keynesian demand shocks, leading to increase in prices, reduction in unemployment and a contractionary reaction from monetary policy. Shocks to the dispersion of expectations seem to have an important impact but are more difficult to interpret. They have a contractionary im-

impact over unemployment and industrial output, a positive impact on stock prices and no impact on prices or exchange rate. These results are certainly interesting, but should be taken with care given the fragility of the identifying assumptions.

Perhaps, the main pattern that emerged from the analysis is that inflation expectations seem to respond slowly to new relevant information. In some sense, there seems to be an "excess smoothness" in expectations. This result is interesting as many of the new approaches to expectations formation such as Sims (2002) and Mankiw and Reis (2002) attempt to explain stickiness in macroeconomic variables exactly through some form of stickiness in expectations formation.

9 Appendix 1: Econometric Framework

The basis of all the analysis done in this paper is the estimation of a Bayesian VAR with monthly data and 12 lags, and the simulation of its posterior so as to have a measure of the uncertainty related to the different estimates.

One of the main issues that arises in the estimation of standard VAR systems is that the model takes the initial conditions as given (Sims 2000). This leads to a model that does not evaluate whether these conditions are themselves compatible with the estimates provided. For example, in a univariate AR this can show up as the model interpreting a trend-like behavior as a regression to a mean which is very distant from the initial condition. In more complicated multivariate systems with multiple lags taking initial conditions as given may mean that when using the estimates to construct a forecast based on the initial conditions the model does an unrealistically good job in forecasting trends, reversals and the like. The strategy normally adopted by the Bayesian VAR literature is to use priors that imply that either the initial conditions are typical of the rest of the sample or else that the data has non-stationary behavior. This is done with the use of dummy observation priors, that is, with the introduction of "mental observations" in the regression data that express these beliefs. Specifically, with y^* being the observation in the left hand side and x_s^* being the observation in the right hand-side "lagged" s periods:

$$y^* = \lambda \bar{y}$$

$$x_s^* = [\lambda \bar{y}, \lambda \bar{y}, \dots, \lambda \bar{y}, 1]$$

and

$$y^* = \mu \bar{y}_i$$

$$x^* = [0, 0, \dots, 0, \mu \bar{y}_i, 0, \dots, 0, 0]$$

Where the barred variables are averages over the initial conditions. The first set of dummy observations express the belief that if all variables remain constant near their initial values, then in the absence of shocks they should not change in any way. With the coefficient on the constant being estimated to 0 this can be interpreted as saying that the variables are cointegrated and close to their steady state relationship, and with that same coefficient different from 0 this means that their initial values should not be far away from their unconditional means. The second set of dummy variables expresses a similar belief, but without allowing for cointegration or stationarity around a constant different from zero. The constants λ and μ are weights that express the strength of the beliefs and are typically set to $\lambda = 5$ and $\mu = 2$. In my estimates, I kept the standard values for the weights.

For the GCP tests I also included a full set of so called "Minnesota priors". Let $t_{k,s}^*$ be the artificial date for the dummy observation corresponding to the s^{th} lag of variable k. Then it has the form.¹⁸:

$$y(t_{k,1}^*) = \begin{bmatrix} 0 \\ \dots \\ \pi \sigma_k \\ \dots \\ 0 \end{bmatrix} = y(t_{k,1}^*)$$

$$y(t_{k,s}^* - s) = \begin{bmatrix} 0 \\ \dots \\ \pi s^\mu \sigma_k \\ \dots \\ 0 \end{bmatrix}$$

and

$$y(t_{k,s}^* - v) = 0 \text{ for } v \neq s > 1 \text{ and for } v > s = 1$$

These are needed because model comparison requires that the prior be proper, that is, that it be integrable with respect to all parameters in the model. With dummy observation priors, this requires at least one dummy observation for each parameter. The latter priors are a form of doing that consistent with the notion that the variables in the system should be approximately random walks. I set the σ_k terms equal to the variance of the initial

¹⁸This definition, like the others above was taken from Sims 2005.

conditions and the decay parameter μ equal to 0.03, implying some degree of mean-reversion in the forecasts. This decay factor makes the choice of lags in the model less of an issue, as more distant lags are assumed to be less important. The parameter π is the one to which the results of the GCP test are sensitive, so I set it to different values when performing those tests.

10 References

- Ball, L. and D. Croushore (1995) "Expectations and the Effects of Monetary Policy, NBER Working Paper 5344.
- Bernanke, B. and A. Blinder(1992), "The federal funds rate and the channels of monetary transmission, *American Economic Review*, vol. 82 (4).
- Bernanke, B. and I. Mihov(1998), "Measuring Monetary Policy". *Quarterly Journal of Economics*, vol 113, 3, pp 860-902
- Brunnermeier, M. and J. Parker (2004), Optimal Expectations. NBER Working Paper 10707
- Deaton, A. (1992) *Understanding Consumption*, Oxford University Press
- Hamilton, J. (1994) *Time Series Analysis*, Princeton University Press, pp. 302-309
- Lamont, O. (2002), "Macroeconomic Forecasts and Microeconomic Forecasters" *Journal of Economic Behavior and Organization*, forthcoming.
- Leeper, E, C. Sims and T. Zha (1996) "What does Monetary Policy Do?". *Brookings Papers on Economic Activity*, vol. 2
- Lucas R.E., (1976). "Econometric Policy Evaluation: A Critique" in , *The Phillips curve and Labor Markets*, eds K. Brunner and A.H. Meltze, eds.
- Mankiw, G. (2001) "The inexorable and mysterious tradeoff between inflation and unemployment, *Economic Journal*, vol. 111.
- Mankiw, G, Reis, R and J. Wolfers(2003), "Disagreement About Inflation Expectations". NBER Working Paper 9796.
- Miller, Edward, 1977, "Risk, uncertainty and divergence of opinion", *Journal of Finance* 32, 1151-1168.
- Muth, John F. (1961) "Rational Expectations and the Theory of Price Movements," *Econometrica* 29(2) 315-335
- Pesaran, M. H. and M. Weale, "Survey Expectations". IEPR Working Paper No. 05.30
- Reis, R. (2003) "Inattentive Producers, Princeton mimeo

- Shleifer A. *Inefficient Markets An Introduction to Behavioral Finance*, Oxford University Press, 2000
- Sims, C. (1998) “Stickiness, *Carnegie-Rochester Conference Series on Public Policy*, vol. 49.
- Sims, C. (1999) “Dummy Observation Priors for VARs”. Princeton mimeo. URL: <http://eco-072399b.princeton.edu/yftp/TimesF99/DummyObs2.pdf>
- Sims, C. A. (2000): “Using a Likelihood Perspective to Sharpen Econometric Discourse: Three Examples *Journal of Econometrics*, vol. 95, issue 2, pages 443-462
- Sims, C. (2003) “Implications of rational inattention, *Journal of Monetary Economics*, vol. 50 (3).
- Sims C. (2005) “Causal Orderings and Exogeneity”. Princeton mimeo. URL: sims.princeton.edu/yftp/Times05/GCP05.pdf
- Sims, C and T. Zha (1999), “Error Bands for Impulse Responses”, *Econometrica*, p. 1113-1156
- Thomas, L. (1999), “Survey Measures of Expected U.S. Inflation”. *The Journal of Economic Perspective*, vol. 13(4), pp. 125-144
- Woodford, Michael (2001) Imperfect common knowledge and the effects of monetary policy, in *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, eds. Aghion, Frydman, Stiglitz, and Woodford.



Figure 1: Time Series of Inflation, VAR Forecast and Survey Expectations

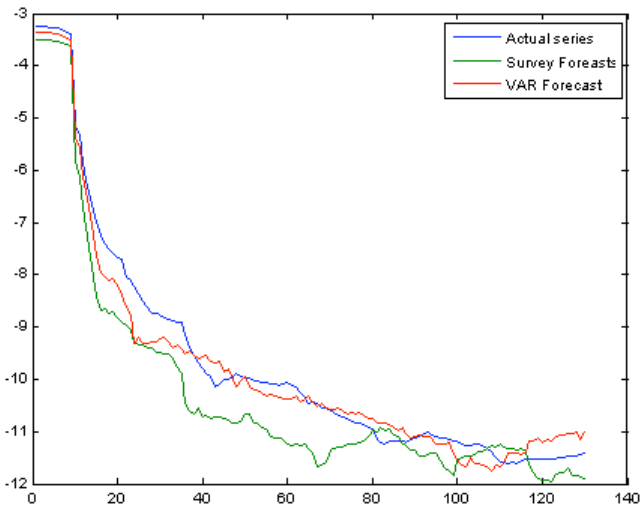


Figure 2: Spectral Density of Inflation, VAR Forecast and Survey Expectations

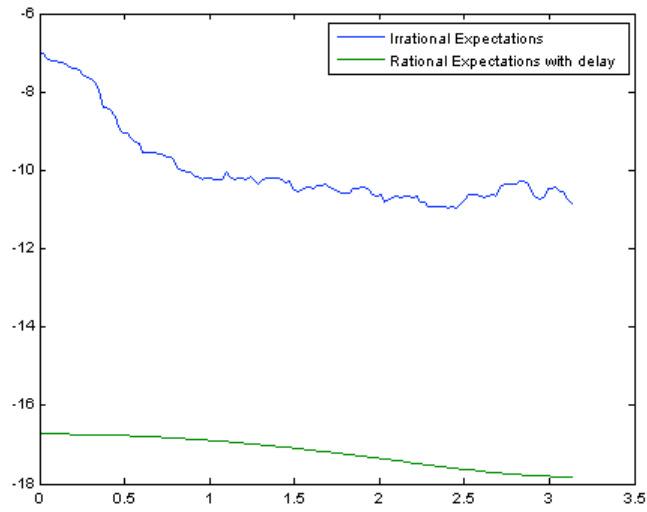


Figure 3: Spectral Density of “Irrational Expectations” and theoretical spectral density of deviation from VAR forecast with delay in absorption of information.

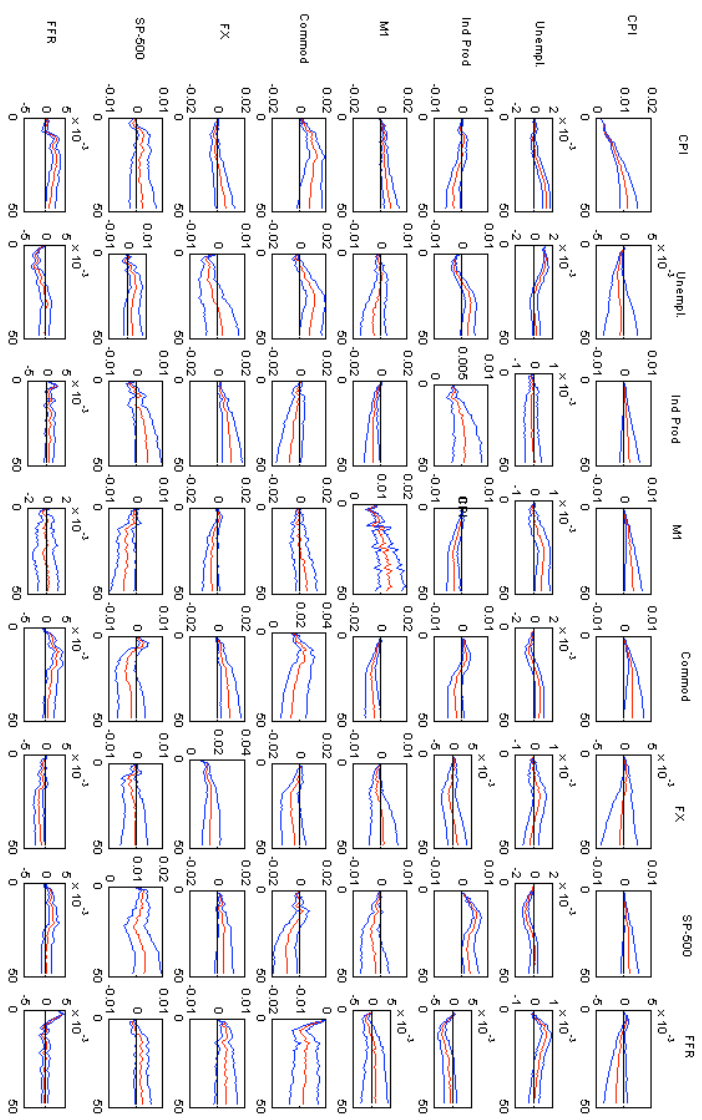


Figure 4: Response of non-survey variables to shocks to the non-survey variables (shocks in the columns, responses in the lines)




Figure 5: Impulse Response Plots for survey expectations variables (shocks in the columns, responses in the lines)

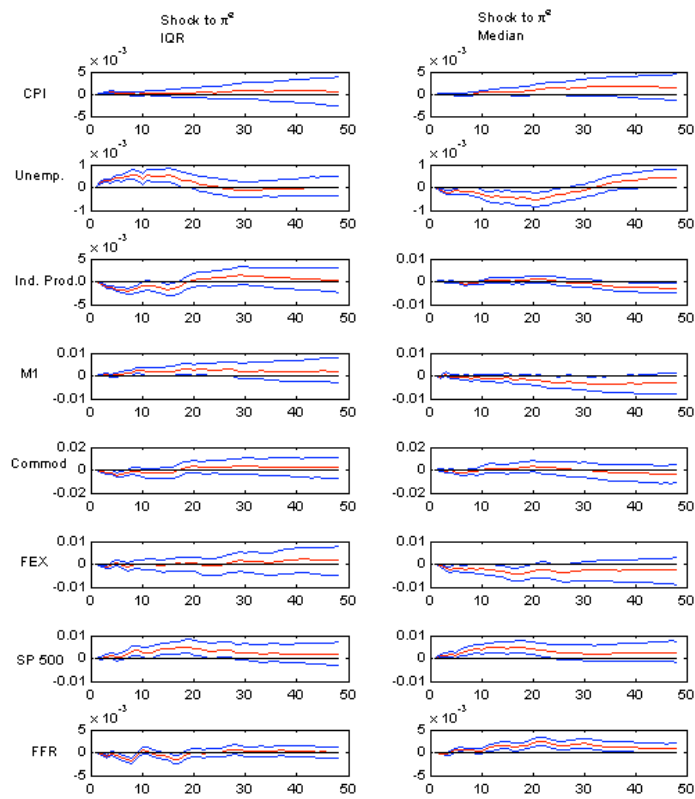


Figure 6: Response of non-survey variables to shocks to the survey variables (shocks in the columns, responses in the lines)

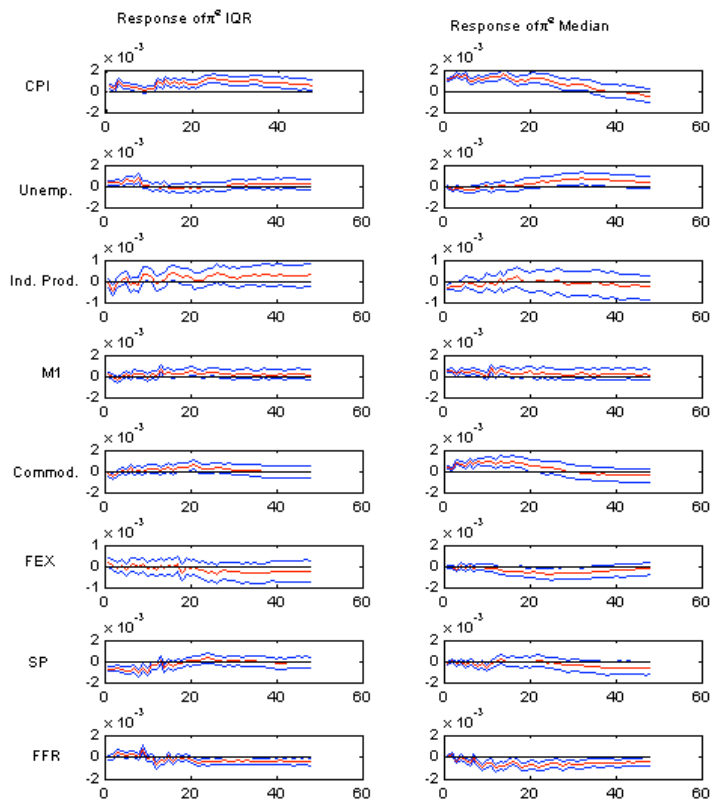


Figure 7: Response of survey variables to shocks to the non-survey variables (shocks in the lines, responses in the columns)

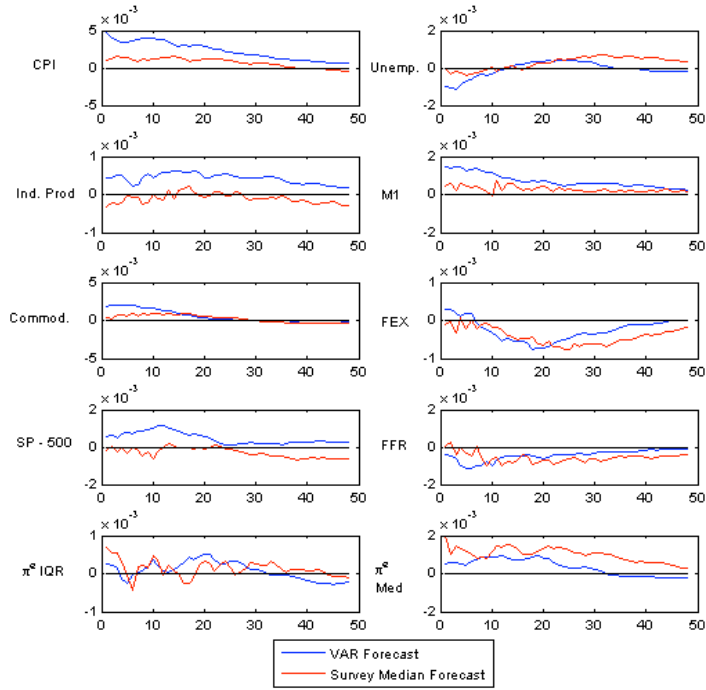


Figure 8: Response of 12-month inflation and of inflation forecasts to shocks

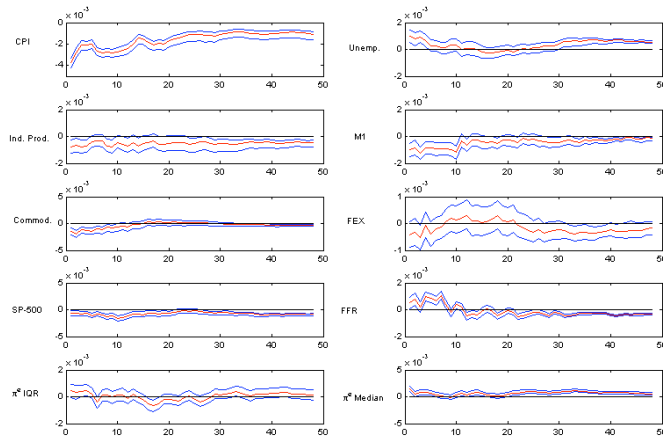


Figure 9: Response of the difference between 12-month inflation and 12-month inflation forecast to shocks