WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER NO. 10-6
EXPECTATIONS AND ECONOMIC FLUCTUATIONS: AN ANALYSIS USING SURVEY DATA

Sylvain Leduc
Federal Reserve Bank of San Francisco
and
Keith Sill
Federal Reserve Bank of Philadelphia

February 2010

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA
Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

Expectations and Economic Fluctuations: An Analysis Using Survey Data

Sylvain Leduc[1]
Federal Reserve Bank of San Francisco

Keith Sill
Federal Reserve Bank of Philadelphia

February 2010

[1]Corresponding Author: Sylvain Leduc, Telephone 415-974-3059, Fax 415-974-2168. Email address: sylvain.leduc@sf.frb.org. We thank Chris Gust and Robert Vigfusson for helpful comments.

Abstract

Using survey-based measures of future U.S. economic activity from the Livingston Survey and the Survey of Professional Forecasters, we study how changes in expectations, and their interaction with monetary policy, contribute to fluctuations in macroeconomic aggregates. We find that changes in expected future economic activity are a quantitatively important driver of economic fluctuations: a perception that good times are ahead typically leads to a significant rise in current measures of economic activity and inflation. We also find that the short-term interest rate rises in response to expectations of good times as monetary policy tightens. Our results provide quantitative evidence on the importance of expectations-driven business cycles and on the role that monetary policy plays in shaping them.

JEL classification:

Keywords: Expectations, survey data, economic fluctuations

1 Introduction

The idea that changes in expectations of future economic activity can be important drivers of economic fluctuations has received increased attention with the unfolding of boom-bust cycles around the world over the past 20 years. The experiences of Japan in the late 1980s, East Asia in the late 1990s, and the United States in 2001 and 2007 suggest that optimism about future growth prospects may help fuel booms and that subsequent downward revisions in expectations may help precipitate busts. In addition, these episodes have served to generate debate about the importance of the role played by monetary policy in boom-bust cycles: the episodes were often accompanied by heightened criticism of central banks for fueling the booms by keeping monetary policy too easy for too long.[1]

Although boom-bust cycles are interesting events that suggest the importance of expectations for economic fluctuations, there has been relatively little empirical analysis that attempts to formally quantify the role played by expectations for the cyclical behavior of the economy. In this paper, we add to this literature by introducing survey-based measures of future U.S. economic activity into simple empirical models to measure how changes in expectations, and their interaction with monetary policy, contribute to fluctuations

[1] For instance, see Taylor (2008) for a criticism of the Federal Reserve's policies under Chairmen Greenspan and Bernanke, and Lionel Robbin's (1934) argument that the Federal Reserve kept interest rates below the natural rate for too long during the late 1920s. See also Okina et al (2001) concerning reasons that Japanese monetary policy may have been too loose in the period leading up to the burst of the stock-market bubble at the beginning of 1990.

in macroeconomic aggregates. We take expectations measures compiled by the Livingston Survey and the Survey of Professional Forecasters and introduce them into an otherwise conventional vector autoregression framework. Survey data remain relatively unused in empirical work on economic fluctuations, despite the fact that surveys are closely monitored by policymakers, who view them as important indicators of market participants' perceptions of future economic activity. A benefit of using survey data is that they provide an independent source of information about agents' perceptions of future economic activity. Consequently, one need not impose modeling assumptions to back out those expectations.

We exploit the timing of the surveys' construction to help identify structural shocks to expectations. To circumvent difficult issues surrounding the use of ex post revised data in assessing the quantitative role of changes in expectations, our VARs are estimated using data that are not subject to revision over time. In particular, we use the unemployment rate as a measure of economic activity.

Our main finding is that changes in expected future economic activity are a quantitatively important driver of economic fluctuations: a perception that good times are ahead typically leads to a significant rise in current measures of real economic activity and inflation. In response, the short-term interest rate rises as monetary policy tightens. The results are robust across the two surveys and to the inclusion of different forward-looking financial variables in the empirical models. Moreover, we show that our results are substantively unchanged when we measure expectations using data from the University of Michigan survey of households rather than survey measures from professional forecasters. Although the Michigan survey asks respondents only whether they expect the unemployment rate to go up, down, or stay unchanged, it remains an important check on our findings, since the coverage of the survey is much more extensive.

Our results shed some light on the role played by monetary policy in fueling boom episodes in the U.S. The conventional wisdom, as embodied in Bernanke and Gertler (2000), is that an inflation-targeting central bank will naturally act as a stabilizing force with respect to boom-bust cycles as it contracts monetary policy in response to factors that raise expected inflation and lower output gaps. In contrast, Christiano, Motto, and Rostagno (2006) point out that this need not be the case if nominal wages are sticky. Expectations of higher future productivity growth put upward pressure on the real wage. To the extent that the nominal wage is sticky, this is necessarily accomplished by a fall in prices, to which an inflation- targeting central bank responds by lowering the short-term nominal interest rate, thus feeding the boom. Our empirical results are more consistent with the conventional view: upward revisions to expectations about future economic performance are accompanied by a rise in current activity and inflation, and a concomitant rise in the short-term interest rate, which tends to stabilize the economy. Though our findings are not derived exclusively from boom-bust episodes in the data, they nevertheless suggest that during times of "economic optimism" the Federal Reserve tended not to run an expansionary monetary policy that amplified fluctuations.[2]

[2] Christiano, Motto, and Rostagno (2006) focus their analysis solely on boom-bust episodes, of which they identify three over the period 1870 to 2006: one that began in 1920 and ended at the start of the Great Depression, one that began in the mid 1950s and

Our findings are somewhat at odds with the predictions of the standard neo-classical business cycle model. In that model, expectations that good times are ahead, usually modeled as an anticipated increase in productivity, lead to a current period recession, as the positive wealth effect of the anticipated productivity increase induces an increase in leisure today. However, the standard model can be modified so that expectations of good times can generate business cycle booms, as shown by Beaudry and Portier (2006,2007), Christiano, Motto, and Rostagno (2006), and Jaimovich and Rebelo (2009). Typically, the modification involves adding complementarities in the production technology or adding certain types of adjustment costs, of which a labor-matching friction would be an example.

Indeed, our finding that the economy expands in response to an upward revision in expectations of future activity squares well with the predictions of standard labor matching models with respect to the impact of changes in expectations. In that framework, expectations that good times are ahead increase the marginal benefit of a match and lead to a fall in the current unemployment rate as more vacancies are being posted. Studying the effects of an anticipated increase in productivity in a labor search model, den Haan and Kaltenbrunner (2009) find that it induces entrepreneurs to increase investment in new projects and post vacancies early and so induce an economic expansion.

Our empirical analysis is not geared to identifying specific factors that map into expectations shocks. Nonetheless, our view is that such factors are likely to include news revelations about future economic developments, the

ended in the 1970s, and one that began in the mid 1990s and ended in the early 2000s.

possibility of future labor strikes, new technological developments, or pre-announced monetary policy actions, etc. that are observed by survey participants, but that are difficult to adequately capture in a small-scale VAR. The empirical of work of Beaudry and Portier (2006) suggests that such "news shocks" explain about 50 percent of business cycle fluctuations. Similarly, Schmitt-Grohé and Uribe (2008) estimate the contribution of news shocks to business cycles in a real business cycle model and find that anticipated shocks account for close to 70 percent of aggregate fluctuations.[3] Similarly, our estimates suggest a quantitatively important role for expectations shocks in economic fluctuations.

The rest of the paper is organized as follows. Section 2 describes the baseline VAR and how we measure expectations and economic activity. Sections 3 and 4 then present empirical results on the baseline model and some extensions to the baseline. Section 5 investigates the robustness of the results, and Section 6 concludes.

2 A Small Structural VAR

We are interested in quantifying the extent to which changes in agents' expectations about the future may affect current economic variables. To measure expectations we use data from two sources: the Survey of Professional Forecasters (SPF) and the Livingston Survey (LS). Both surveys collect pre-

[3] On the other hand, Sims (2008) uses a different identification scheme in a larger-scale VAR and argues that shocks to expectations about future productivity are not an important source of business cycles and that the responses to such shocks are similar to what is predicted by the standard real business cycle model.

dictions from professional forecasters (typically about 40 to 50 respondents per survey) and both are conducted by the Federal Reserve Bank of Philadelphia. The SPF is a quarterly survey that dates from 1968, at which time it was conducted by the American Statistical Association and the National Bureau of Economic Research (the Federal Reserve Bank of Philadelphia took over the survey in 1990). The SPF is released at the end of the second month of each quarter (or early in the next month). Survey participants provide forecasts of variables such as CPI inflation, the unemployment rate, real GDP growth, and nonfarm payroll growth over a 5- quarter horizon. The LS, which was initiated in 1946, is conducted twice a year. Survey questionnaires go out in May and November and the survey's results are made public in the second week of June and December. The survey started compiling forecasts of the unemployment rate in 1961 and it covers a somewhat broader set of macroeconomic variables than the SPF.

We use survey forecasts of the unemployment rate to proxy expectations about future economic activity. The unemployment rate has the advantage that it is subject to only minor revisions, which are limited to changes in seasonal factors. By using forecasts of the unemployment rate we can bypass difficult questions about real-time data and subsequent data revisions. For example, the use of expected and actual real GDP growth (measured using the latest vintage of data) in our VARs would be problematic because real GDP revisions may incorporate information that is unavailable to forecasters at the time their forecasts were being made. Since the unemployment rate series is unrevised, we can include expected and actual unemployment in a VAR and not otherwise have to account for the possibility that the VAR conditioning set contains more information than forecasters had when making their predictions.

Figure 1 shows the 2-quarter-ahead SPF forecast of the unemployment rate and the realized unemployment rate. Forecasters were able to predict the unemployment rate reasonably well, notwithstanding its large increase to more than 10 percent in 1982 from 4 percent in the early 1970s, and its quick decline coming out of the 1981-82 recession.[4] Overall, there doesn't appear to be a systematic bias in the unemployment rate forecasts. This is in contrast to inflation forecasts, which tended to be consistently below the inflation rate during the inflation runup in the 1970s and consistently above the inflation rate during the Volcker disinflation of the early 1980s (see, for instance, the discussion in Leduc, Sill, and Stark (2007)).

We first consider the implications of a baseline VAR model composed of the 6-month-ahead expected unemployment rate, the realized unemployment rate, the realized CPI inflation rate, and the realized nominal 3-month Treasury bill rate.[5] The key specification issue for investigating the consequences of shifts in expectations is how to identify expectations shocks. We use a recursive identification scheme that places the expected unemployment rate first in the ordering, followed by the actual unemployment rate, CPI inflation, and the nominal interest rate. Consequently, this ordering assumes that there is no contemporaneous response of expected unemployment to shocks to the other variables in the system.

[4] A similar picture emerges for longer-horizon forecasts and for forecasts from the Livingston survey.

[5] Note that the CPI index is generally not revised over time nor is the measured Treasury bill rate.

Following Leduc, Sill, and Stark (2007), the placement of expected unemployment first in the recursive identification is motivated by the timing of the surveys and the way we have aligned the other data in the VAR. The timing of the survey is critical in that it allows us to put expected unemployment first, since when making forecasts at time t, the information set on which agents condition their forecasts does not include, by construction, the time t realizations of the unemployment rate and the other variables in our VAR.

To elaborate, take the case of the SPF. The response deadline is generally the third week of the second month of the quarter (although the deadline does vary a bit over the survey sample period).[6] Based on the survey's timing, we redefine quarters of the year so that the first month of a quarter is the month that survey responses are filled out. Thus, the redefined first quarter is February, March, April. The second quarter becomes May, June, July, and so on. With this timing convention and associated data definition, the SPF is by construction conducted at the start (generally the second or third week) of each quarter. Consequently, the data are aligned so that agents have past values of the unemployment rate, inflation and interest rates in their information set when the surveys are filled out, but they do not yet have the official data releases telling them contemporaneous quarter values of the unemployment rate and inflation. However, they do have some information on quarterly interest rates (the first two or three weeks of the quarterly realization).

[6]For example, in 1995 the Q1 SPF respondents had to return the survey questionnaire by February 21. For the 1995Q2, survey responses were due by May 22. For 1995Q3, survey responses were due by August 22, and for 1995Q4 survey responses were due by November 20.

We employ a similar strategy when constructing the data set for the VARs that use the Livingston Survey measure of expectations. Now, half-years are defined based on the timing of the Livingston Survey to mitigate the influence that contemporaneous realizations of the unemployment rate, inflation, and interest rates can have on forecasters' decisions about future unemployment rates. Since in this case the survey questionnaire is due back in May and November, we redefine half-years as running from April to October and from October to April. As with the SPF, this data alignment implies that the survey is conducted at the start of each period: May and November. The remaining variables in the VAR are then measured as the average monthly value of the corresponding six-month period.

It could of course be the case that agents condition their forecasts on variables that are omitted from the VAR and which provide important information about the within-period realization of the unemployment rate, interest rates, and inflation. We address this concern by expanding the set of variables in the baseline VAR to include additional financial market variables that are likely to be influenced by potentially important omitted variables. In addition, we introduce further controls for oil and fiscal policy shocks. As shown below, these modifications of the baseline structure do not qualitatively change the results (though there is some small change in the quantitative responses).

To summarize then, the baseline VARs contain four dynamic variables and use a recursive identification scheme that orders expectation variables first. The variables are ordered as the expected unemployment rate, the unemployment rate, inflation, and the nominal interest rate. The data used in the VAR are largely unrevised over time, and the definitions of quarters or half-years and the measurements of quarterly and biannual realizations are consistent with placing the expectations variable above the other model variables in the recursive ordering. We now turn to an analysis of the baseline VAR results.

3 Results from a Baseline VAR

Our interest focuses on the economy's response to an unanticipated shock to expectations of the future unemployment rate. We interpret a shock to the expected unemployment rate as news received by agents that leads them to reassess their beliefs regarding future prospects for the economy. A negative shock to the expected unemployment rate (i.e., lower expected unemployment) then has the interpretation of news that agents get at the start of a quarter that leads them to become more optimistic about future economic conditions. The models are estimated over the sample period 1961H1 to 2007H1 for the Livingston Survey VAR (because the LS is biannual, we denote the first survey observation in a year as H1) and 1968Q4-2007Q2 for the SPF VAR. For our baseline results, we ended the sample before the onset of the current financial crisis. We did this to avoid misspecification issues related to the zero lower bound on nominal interest rates and to the unconventional tools used to conduct monetary policy over the post-2007 period. However, we verify that similar results obtain when post-2007 data are included in the sample.

Figure 2 shows the impulse responses to a normalized one unit negative shock to the 6-month-ahead expected unemployment rate for the baseline VARs. The panel on the left shows the impulses from a VAR that uses the Livingston Survey measure of expectations, while the panel on the right shows the responses from a VAR that uses the SPF expectations measure. Each panel of the figure shows the response of a variable to the expectations shock, as well as 68 and 90 percent confidence intervals that are generated using Kilian's (1998) bootstrap-within-bootstrap method.

On impact, a negative innovation to the expected unemployment rate six months ahead leads to a fall in the current unemployment rate, an increase in inflation, and an increase in the 3-month Treasury bill rate. All responses are significantly different from zero at the 90 percent confidence level. The unemployment rate is significantly below zero for about 2 years, while the increase in the nominal interest rate is significant for 2 to 3 years. The rise in inflation is somewhat more significant and persistent in the Livingston VAR than in the SPF VAR. Expected unemployment is quite persistent, staying significantly below zero for roughly 2 years.

The two VARs largely tell the same story. Unexpected good news that leads agents to revise down their forecasts for future unemployment rates brings about a current boom in economic activity and a concomitant tightening of monetary policy. Note that while the implied short-term real interest rate falls slightly on impact, it is rising and above zero by the second period after the shock and so is consistent with a somewhat delayed policy tightening. On balance, the estimated monetary policy response to a news shock supports the story in Bernanke and Gertler (2000). More optimistic expectations of a future boom (in the form of a lower future unemployment rate) coincide with an anticipatory monetary policy tightening. In this respect, the baseline specification suggests that, on average, monetary policy over the sample period did not serve to amplify expectations-driven fluctuations. Rather, policymakers appear to have responded to anticipated booms and the concomitant higher near-term inflation by raising the short-term interest rate. The response is consistent with the view that a monetary policy that responds aggressively to changes in inflation serves to dampen economic fluctuations.

The finding that expectations of good times in the future lead to good times today is consistent with the view in Beaudry and Portier (2006), Den Haan and Kaltenbrunner (2009), and Jaimovich and Rebelo (2009). These authors argue that upward revisions to expectations of future productivity growth can lead to business cycle booms in suitably modified business cycle models. While our VAR-based findings do not provide direct evidence on the effects of revisions to expected productivity growth, the movement in the VAR's unemployment rate in response to revisions to expectations of future activity is qualitatively consistent with the hours/unemployment dynamic response in the aforesaid papers.[7]

3.1 Controlling for Fiscal Policy and Oil Shocks

To accurately assess the role of expectations shocks for economic fluctuations it is important to control for shocks that may play a significant role in

[7] Since 1991, the SPF asks survey respondents about their expectations of productivity growth 10 years out. However, the question is asked only in the first quarter of each year and so leaves us relatively few data points to conduct a meaningful analysis.

the system's dynamic behavior. Two obvious candidates for such shocks are oil price movements and fiscal stimulus/contractions. There is evidence that large upward movements in oil prices are associated with economic downturns (see Hamilton (2003) and the references therein). To control for exogenous, unanticipated increases in oil prices, we employ the quantitative dummy variable developed by Hamilton (2003). The quantitative dummy variable captures the disruptions in the oil market due to political events in the Middle East that are plausibly exogenous to developments in the U.S. economy. Hamilton identifies the following dates as being associated with exogenous declines (in parenthesis) in world oil supply: November 1956 (10.1%), November 1973 (7.8%), December 1978 (8.9%), October 1980 (7.2%), and August 1990 (8.8%). Three of these episodes fall within the sample period of our baseline model: December 1978, October 1980, and August 1990. The quantitative dummy takes a value equal to the drop in world production during the period in which the episodes occur and is otherwise zero.

To identify exogenous fiscal shocks, we appeal to the narrative approach of Ramey and Shapiro (1998) and its extension in Ramey (2009). They identify four exogenous fiscal shocks in the postwar U.S. data: 1950Q3, associated with the Korean War; 1965Q1, associated with the Vietnam War; 1980Q1 associated with the Carter-Reagan military buildup; and 2001:Q3, associated with terrorist attack on September 11. Of these shocks, only the 1980Q1 and 2001Q3 episodes fall within our estimation period.

Figure 3 shows the impulse responses to an innovation in the expected unemployment rate in the baseline VARs that have been modified to include the oil and fiscal dummy variables (we maintain the same recursive ordering as in Figure 2). Comparing the impulse responses to those in Figure 2, controlling for exogenous oil and fiscal shocks has little effect on the dynamic response of the unemployment rate, inflation, and the nominal interest rate. It remains the case that a fall in the expected unemployment rate leads to a contemporaneous decline in the current unemployment rate, a rise in the inflation rate, and an increase in the short-term interest rate.

4 Extending the Baseline Model

It is conceivable that the expectations shocks in the VAR models are contaminated by omitted variables that convey important information to forecasters about the current and future states of the economy. To mitigate this empirical concern we add additional financial variables to the baseline VARs: stock returns as measured by the S&P500 and long-term bond returns as measured by the yield on ten-year Treasury notes. Presumably, important news about the future economy would be reflected in such financial asset prices and so conditioning on them is a straightforward, although somewhat crude, way to add omitted information to the VAR analysis. The financial data enter the VARs as period averages of daily data. Consequently, news that arrives in the time interval between the last observation of a variable and the date of the survey can potentially affect financial asset prices and be reflected in the VARs, so the measure is not without some concern vis-à-vis the identification assumption.

Figure 4 shows the impulse responses to an expectations shock in the baseline VARs augmented by the equity return and long bond yield series.

We again use a recursive identification scheme with the variables ordered as: expected unemployment rate, unemployment rate, inflation, equity returns, long-term interest rate, and fed funds rate. The responses for the unemployment rate and inflation are similar to those in Figures 2 and 3. Equity returns (Q) and long-term interest rates (LR) increase in response to a negative innovation to unemployment rate expectations, and both responses are significant on impact in the VAR that uses the Livingston Survey measure of expectations. The equity return response is only marginally significant in the VAR that uses the SPF expectations measure. The federal funds rate response is somewhat weaker on impact in Figure 4 compared to the VARs without equity returns and long-term interest rates, but continues to show a significant monetary policy tightening in response to expectations of a lower unemployment rate. On balance, the addition of the financial market variables does not change much the qualitative or quantitative results from the baseline VAR in Figure 2.

4.1 Variance Decompositions and Prediction

The importance of innovations to expectations for the dynamics of the VAR system is indicated as well by variance decompositions. Table 1 shows the contribution of shocks to 6-month-ahead unemployment rate expectations for the unemployment rate and inflation in the baseline 4-variable and 6-variable system. In the first two columns of the table, we report the 1-year-ahead and 5-year-ahead forecast error variance contributions for both the Livingston Survey and SPF VARs. The range of estimated contributions of expectations shocks to the forecast error variance of the actual unemployment rate is wide. In part, this variation is due to the fact that observations on the expectations measures, and thus the VARs, start at different dates. Clearly, though, the table indicates that shocks to expectations are important for economic fluctuations. At the one-year horizon, these shocks account for more than 35 percent of the forecast error of the unemployment rate. While the contribution falls at the 5-year horizon, it remains substantial. These results are broadly in line with the findings of Beaudry and Portier (2006), who use VARs, and those of Schmitt-Grohé and Uribe (2008), who estimate a structurel model, in that expectations shocks account for more than 50 percent of aggregate fluctuations.

We also find that shocks to expectations are important for the forecast error variance of the inflation rate, although that contribution is smaller than for the unemployment rate. For instance, Table 1 indicates that the 1-year-ahead contributions range from 4 to 20 percent. This range is about the same for the 5-year-ahead contributions.

We also consider how well the variables in the VARs predict expected unemployment rates. Granger-causality tests indicate that the variables in both the 4-variable VAR and the 6-variable VAR Granger-cause the 6-month-ahead expected unemployment rate. The P-values on all the tests were essentially zero for the null hypothesis of no causality. We also examined how much the variance of residuals from a regression of the expected unemployment rate on a constant and four of its own lags falls when we add inflation, the unemployment rate, and short-term interest rate to the regression. We find that the variance of the residual drops on the order of 30 to 40 percent for the Livingston Survey measure and for the SPF measure of expectations, respectively. If the set of regressors includes long-term interest rates and equity returns, the variance of the residual falls 43 percent for the SPF measure of expectations, and 52 percent for the Livingston Survey measure of expectations (the baseline remains a regression of the expected unemployment rate on itself only). Consequently, we are confident that the variables in our VARs are capturing important information forecasters use when making projections of the unemployment rate. [8]

4.2 Discussion

Given that we emphasize the unemployment rate as a measure of real activity in the empirical analysis, a natural framework with which to interpret the findings is a Mortenssen-Pissarides style labor search model. Den Haan and Kaltenbrunner (2009) use such a model to study the impact of positive news about future productivity growth on the business cycle and argued that, because of the matching friction central to these frameworks, such news can lead to co-movements in macro variables that resemble typical business cycles. As a result of the matching friction, firms post more vacancies in anticipation of better times ahead, which increases today's employment rate and lowers the number of unemployed workers. Our results are consistent

[8] Because the unemployment rate tends to be very persistent, movements in the expected unemployment rate may be capturing the past more than future movements in the unemployment rate. To mitigate this concern, we also considered an alternative version of our baseline model with the actual and expected unemployment rates in first difference. In response to a sudden fall in expected unemployment, the results continue to show a fall in the actual unemployment rate and a rise in inflation and interest rates. To preserve space, we did not include the figure, but the results are available upon request.

with this model's predictions: expectations of a downward movement in the unemployment rate leads to an immediate drop in the unemployment rate, and thus a pickup in economic activity.[9]

The responses of the short-term interest rate and the inflation rate to expectations shocks that we uncover with our VAR analysis is also in line with a simple monetary version of the labor search model. Expectations of better times ahead would lead to an increase in current demand, which would push marginal costs upward, resulting in higher prices and a rise in the inflation rate. Assuming that the central bank follows an interest-rate rule of the type estimated by Clarida, Gali, and Gertler (2000), the short-term interest rate would rise as a result of both the higher inflation rate and a positive output gap.

Of course, our model documents the average response of monetary policymakers to expectations shocks over the entire sample period. This is not to say that in particular instances monetary policymakers may not have acted more slowly than usual when tightening interest rates before the economy heated up. Indeed, some commentators argue that the low inflation rates in the United States during the late 1990s and during the housing-market boom kept policymakers from raising interest rates in a proactive manner to stem incipient booms (see Taylor (2008)). On balance, though, our results point to a tighter monetary policy following waves of optimism.

[9] As already mentioned, our analysis does not identify specific factors that map into expectation shocks. As a result, although the results suggest that expectations of better times ahead lead to a current rise in economic activity, the identification scheme doesn't necessarily ascribe this effect to revisions in expectations of future productivity growth, as is the case in the work of Den Haan and Kaltenbrunner (2009).

5 Robustness Checks

5.1 Additional Measure of Real Activity

To provide more evidence of the effect of an innovation to the expected unemployment rate on real economic activity, we introduced another measure of real activity to the empirical model. Since we use unrevised, or real-time, data in the VAR, the options for real variables are limited. We use the Institute for Supply Management (ISM) activity index as an additional gauge of real activity. The ISM index is a composite index based on surveys of purchasing managers in the manufacturing sector. An index value above 50 generally indicates the manufacturing sector is expanding, while an index value below 50 indicates contraction. The index is not revised over time and is widely believed to reflect future movements in real output.

Figure 5 shows the impulse responses to an expectations shock when the ISM index is ordered last in the VAR. (The ordering of the other variables is the same as in Figure 4.) In response to a negative innovation in the expected unemployment rate, the ISM index rises on impact and remains significantly above zero for about 6 months before exhibiting a hump-shaped pattern that takes the index significantly below zero for about 12 months, which is consistent with a tighter monetary policy. On balance, though, the finding that the ISM index rises in response to a positive innovation in the expected unemployment rate is consistent with our finding that the current unemployment rate falls. Note as well that the introduction of the ISM in the SPF VAR leads to generally less significant results in comparison to the Livingston Survey VAR and to the other specifications. However, the qualitative pattern of the responses is similar to that of the aforesaid VARs.

5.2 Longer Horizon Forecast and Sample Size

We also consider how the results might change if forecasts of the future unemployment rate were made for a longer horizon. The Livingston Survey and SPF both have measures of expected unemployment 12 months ahead. As can be seen in Figure 6, using 12-month-ahead expectations in the VARs leads to virtually the same results as in the baseline that has 6-month-ahead expectations. The story told by the VARs remains that expected good times ahead lead to a current fall in the unemployment rate, an increase in inflation, and a more restrictive monetary policy.

We also verify that the VAR results are stable over a sample period that includes the current financial crisis. We re-estimated the 6-variable VARs over the period 1960M1-2009M4 for the Livingston Survey VAR and 1968Q1-2009Q2 for the SPF VAR. The results are shown in Figure 7. Overall, extending the sample to include the current crisis has little effect on the impulse responses.

5.3 Measure of Households' Expectations

To this point, the empirical results are conditioned on using measures of unemployment expectations from surveys of professional forecasters. One possible issue with the use of such surveys is that their coverage, in terms of participants, is relatively small: the number of respondents in the LS or in the SPF is typically on the order of 40 to 50. An additional issue, noted, for example, in Mankiw, Reis, and Wolfers (2003) is that professional and household forecasts can differ quite substantially along some important dimensions. For example, they find that consumer forecasts of inflation tend to be less efficient than professional forecasts: forecast errors of consumers are predictable based merely on past forecasts, while such is not the case for professional forecasts.

To investigate the robustness of our results to an alternative measure of survey expectations data, we re-ran our VARs using data from the Michigan survey of households. The Michigan survey asks roughly 500 households about their assessment of the economy and their expectations of a large number of economic variables. One drawback of the survey, and an important reason that we didn't use it in our baseline model, is that it asks respondents only whether they think the unemployment rate will go up, down, or stay unchanged over the next 12 months.[10] Contrary to the LS or the SPF, the University of Michigan survey doesn't ask respondents to provide a point estimate for the unemployment rate.

Nevertheless, one measure of the survey tracks the changes in the unemployment rate quite well: the difference between the percentage of households who thought the unemployment rate would increase minus the percentage who thought it would decline, normalized to 100. Figure 8 shows that this measure tracks the broad movements in the unemployment rate and tends to lead the changes in the actual unemployment rate. It appears then that this measure of household expectations has informational content that can

[10] More specifically, the survey asks respondents the following question: "How about people out of work during the coming 12 months — do you think that there will be more unemployment than now, about the same, or less?"

be used in our empirical models.

Figure 9 shows the results from our baseline and 6-variable models when we use the Michigan survey of household expectations instead of the professional forecasts. Since the survey of consumers asks respondents about the expected change in the unemployment rate over the next 12 month, we also modify the VARs so that they include the yearly change in the actual unemployment rate, rather than its level. The other variables in the VAR are left unchanged. The survey of consumers has compiled quarterly data on expectations of unemployment changes since 1968, and at a monthly frequency since 1978. We use the monthly data, since they allow us more flexibility in lining up the data to make the results more comparable to those with the SPF. Since the response deadline for the SPF is generally the third week of the second month of the quarter, we also use the survey of consumers conducted in that month.

Qualitatively, using household expectations instead of professional forecasts makes little difference to the results. A sudden drop in expected unemployment is followed by a drop in the actual unemployment rate, a rise in inflation, and a tightening of monetary policy. In the 6-variable VAR, a downward revision to expected unemployment also leads to an increase in stock prices and long-term interest rates, just as was the case when expectations of professional forecasters were used. Quantitatively, the variance decomposition reported in Table 1 shows that whether we use households forecasts or professional forecasts, the contribution of shocks to expectations for the variance of the unemployment and the inflation rates is substantial. For example, at the 5-year horizon, shocks to expected unemployment account for from 15 to 48 percent of the forecast error variance of unemployment. The contribution to the forecast error variance of inflation is lower, at 3 to 17 percent. In general, the results from the Michigan survey line up more closely with the Livingston Survey than with the Survey of Professional Forecasters, especially at the longer horizon. On balance, though, the results in Table 1 point to a significant role for expectations shocks in accounting for variation in inflation and unemployment.

6 Conclusion

Expectations play a key role in the determination of dynamic paths for economic variables in cogent equilibrium models of the economy. While recent applied theoretical work has suggested that expectations of future events, events that are not part of current fundamentals, may play an important role in economic fluctuations, the empirical evidence on how important expectations are for business cycles remains somewhat sparse. Existing studies have generally used economic data, such as asset price data, to infer expectations about the future. In contrast, we have examined this issue using actual expectations measures from surveys and used them to investigate how unanticipated shifts in expectations can influence movements in economic variables.

We find that changes in expectations of future economic activity are a quantitatively important driver of economic fluctuations. An anticipation of good times ahead leads to a fall in current unemployment, a rise in inflation, and a tighter monetary policy. These impulse responses hold across a variety of expectations measures, from professional forecasters to households, and a variety of VAR specifications. In this respect, our empirical evidence generally supports the findings of a recent generation of business cycle models that imply that expectations of good times in the future lead to current-period booms, rather than busts. Our results also suggest that during these times of "economic optimism," the Federal Reserve tended not to run an expansionary monetary policy that amplified fluctuations.

With policymakers' and market participants' interests in survey data increasing, existing surveys have expanded and new surveys have also been introduced. For instance, since 1992 the SPF tracks expectations for 10-year-ahead productivity once a year, while in 1999 the ECB introduced a survey of European forecasters similar to the SPF. As enough data become available, analyzing waves of optimism using these relatively newer data sources will be an interesting avenue for future research

7 References

References

[1] Beaudry, P., Portier, F., 2006. "Stock prices, news, and economic fluctuations." *American Economic Review* 96, pp. 1293-1307.

[2] Beaudry, P., Portier, F., 2007. When can changes in expectations cause business cycle fluctuations in neo-classical settings? *Journal of Economic Theory* 135, pp. 458-77.

[3] Bernanke, B., Gertler, M., 2000. "Monetary policy and asset price volatility." National Bureau of Economic Research Working Paper No. 7559.

[4] Christiano, L., Motto, R., Rostagno, M., 2006. "Monetary policy and stock market boom-bust cycles." Manuscript.

[5] Clarida, R., Gaili, J., Gertler, M., 2000. "Monetary policy rules and macroeconomic stability: evidence and some theory." *Quarterly Journal of Economics* 115, pp. 147-80.

[6] Den Haan, W.J., Kaltenbrunner, G., 2009. "Anticipated growth and business cycles in matching models." *Journal of Monetary Economics* 56, pp.309-27.

[7] Hamilton, J.D., 2003. "What is an oil shock?" *Journal of Econometrics* 113, pp. 363-98.

[8] Jaimovich, N., Rebelo, S., 2009. "Can news about the future drive the business cycle?" *American Economic Review* 99, pp. 1097-1118. Papers: 5877.

[9] Kilian, L., 1998. "Small sample confidence intervals for impulse responses." *Review Economics and Statistics* 80, pp. 281-30.

[10] Leduc, S., Sill, K., Stark, T., 2007. "Self-fulfilling expectations and the inflation of the 1970s: evidence from the Livingston Survey." *Journal of Monetary Economics* 54, pp. 433-459.

[11] Mankiw, G., Reis, R., Wolfers, J., 2003. Disagreement about inflation expectations NBER Working Paper No. 9796.

[12] Okina, K., Shiratsuka, S., 2001. Asset price bubbles, price stability, and monetary policy: Japan's experience.IMES discussion paper series, no. 2001-E-16.

[13] Ramey, V.A. Shapiro, M.D., 1998. "Costly capital reallocation and the effects of government spending." *Carnegie-Rochester Conference Series on Public Policy* 48, 145-194.

[14] Ramey, V.A, 2009. "Identifying government spending shocks: it's all in the timing." Manuscript.

[15] Robbins, L., 1934. "The great depression. "

[16] Schmitt-Grohé, S., Uribe, M., 2008. "What's news in business cycles." Working Paper No. 14215, National Bureau of Economic Research.

[17] Sims, E.R., 2008. "Expectations driven business cycles: an empirical evaluation." Manuscript.

[18] Taylor, J.B., 2008. The financial crisis and the policy responses: an empirical analysis of what went wrong.manuscript.

Table 1. Contribution of Expectation Shocks to the Variances of Unemployment and Inflation

	LS	*SPF*	*Michigan*
Unemployment			
1-year ahead			
4-variable system	74.9	54.6	47.6
6-variable system	79.1	36.5	41.2
5-year ahead			
4-variable system	40.2	14.6	48.5
6-variable system	47.2	11.1	30.7
Inflation			
1-year ahead			
4-variable system	13.8	4.9	13.2
6-variable system	12.6	2.2	13.0
5-year ahead			
4-variable system	16.6	3.2	13.2
6-variable system	13.1	3.3	11.5

All entries are in percentage term.

Figure 1. Expected and Realized Unemployment Rates

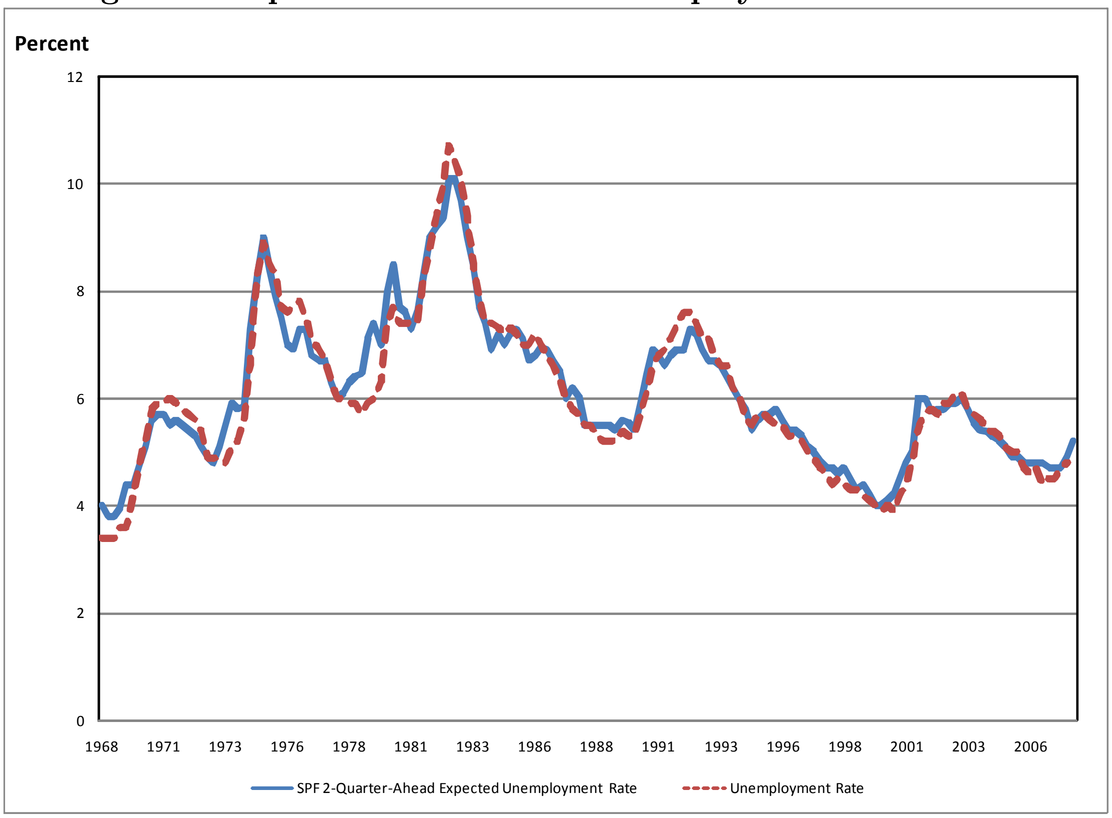

Figure 2. Responses to a Shock to Expected Unemployment

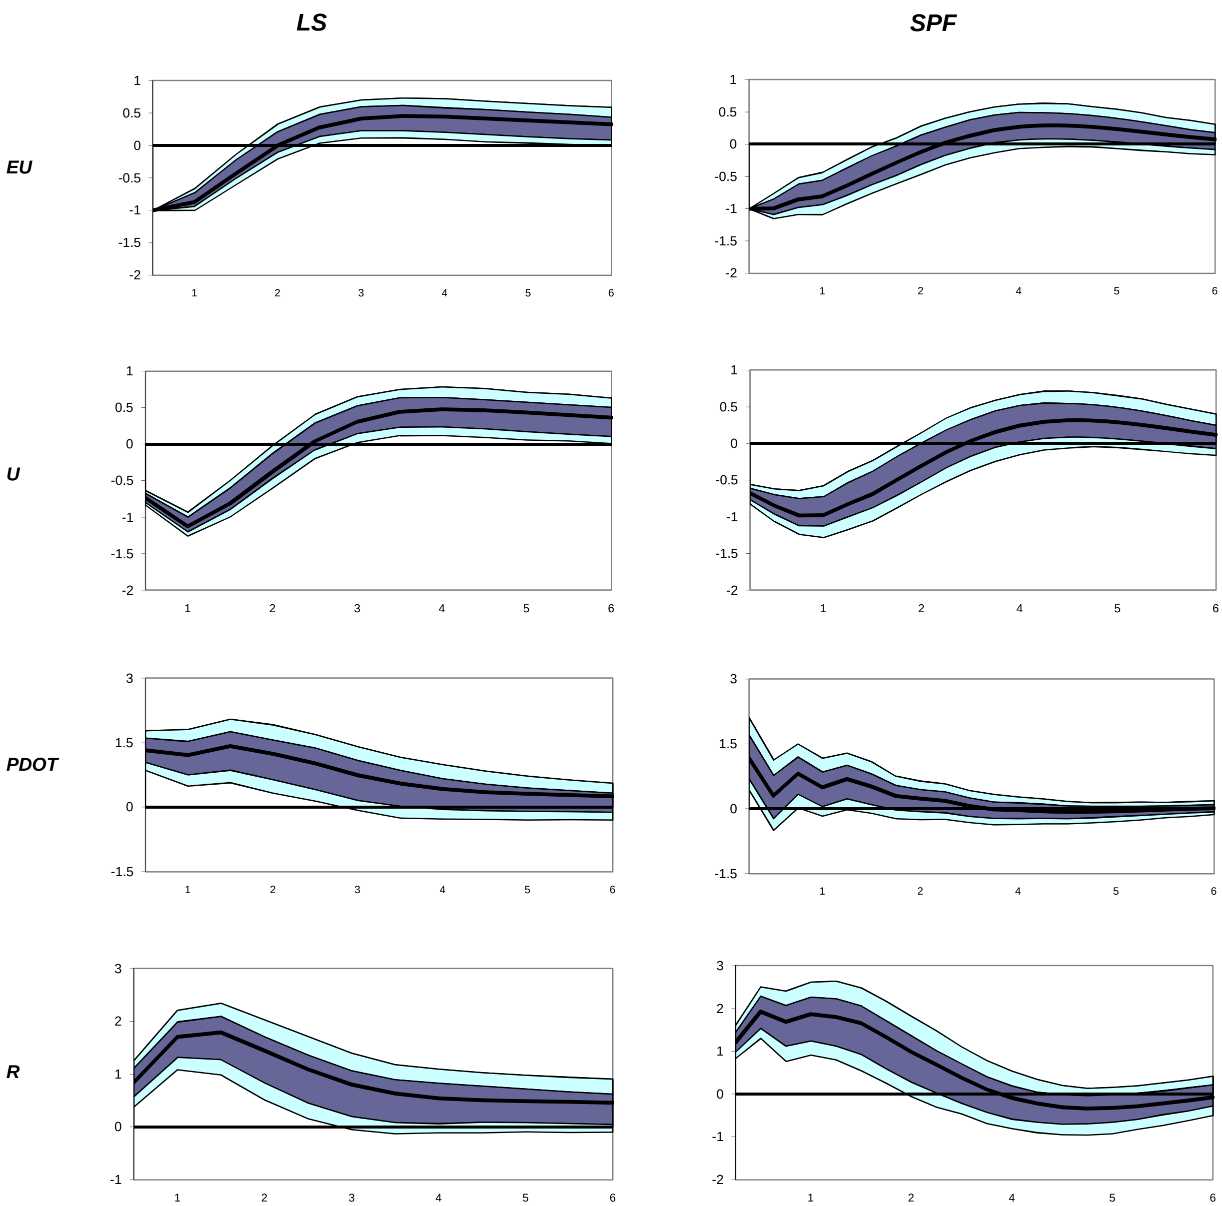

The responses were generated from a VAR with expected unemployment (EU), actual unemployment (U), inflation (PDOT), and the three-month T-Bill rate (R).
All the responses are expressed in percentage terms. The x-axis denotes years. In each chart, the darker area represents the 68% confidence interval, while the sum of the darker and lighter areas denote the 90% confidence interval. The system with the LS survey is estimated over the period 1960H1-2007H1, while that using the SPF survey is estimated over the period 1968Q4-2007Q2.

Figure 3. Responses to a Shock to Expected Unemployment: Controlling for Oil and Fiscal Shocks

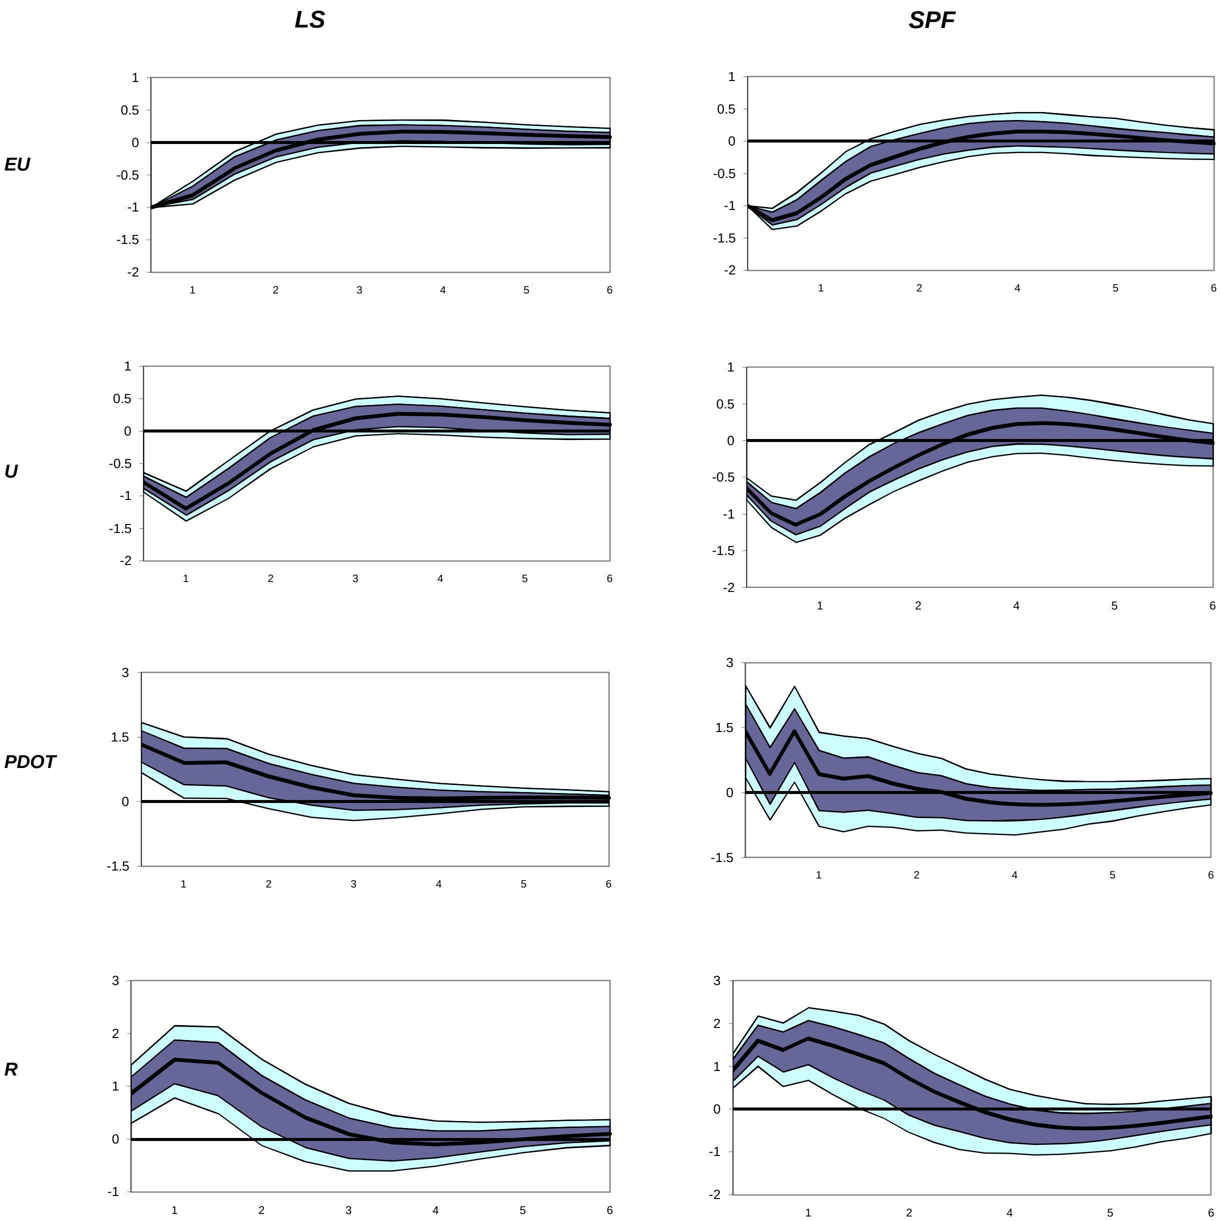

The responses were generated from a VAR with expected unemployment (EU), actual unemployment (U), inflation (PDOT), and the 3-month T-Billt rate (R), ar for oil and fiscal shocks. All the responses are expressed in percentage terms. The x-axis denotes years. In each chart, the darker area represents the 68% co while the sum of the darker and lighter areas denote the 90% confidence interval. The system with the LS survey is estimated over the period 1960H1-2007H1, while that using the SPF survey is estimated over the period 1968Q4-2007Q2.

Figure 4. Responses to a Shock to Expected Unemployment: Larger System

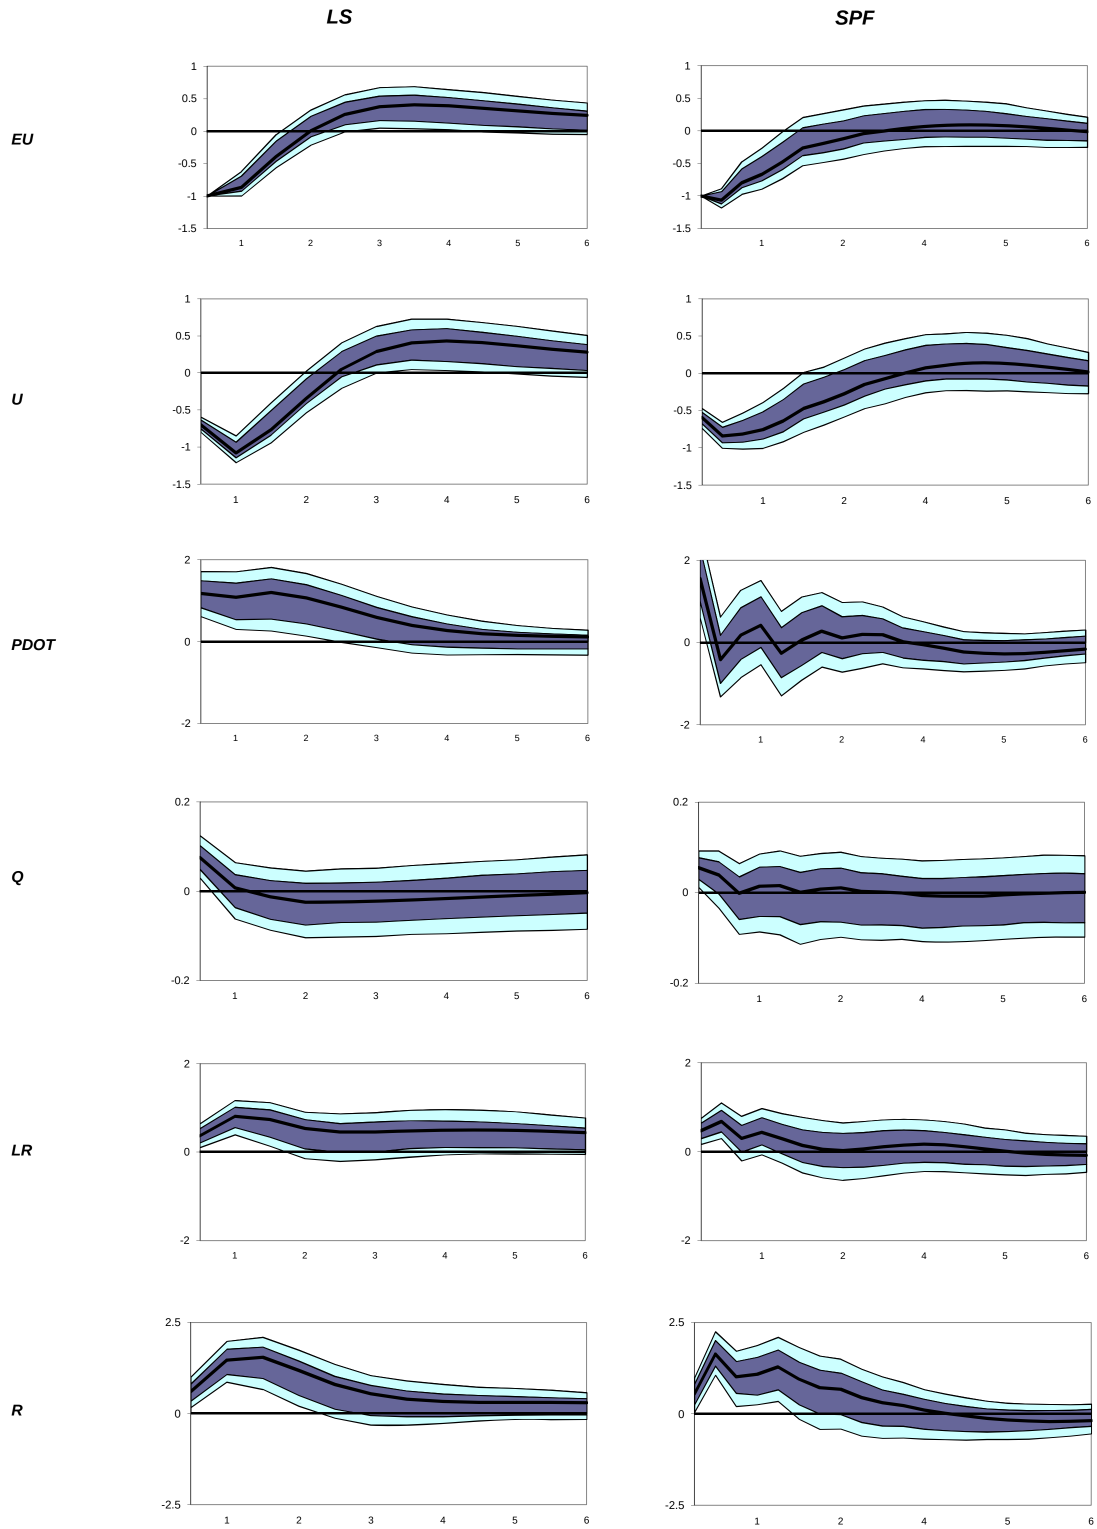

The responses were generated from a VAR with expected unemployment (EU), actual unemployment (U), inflation (PDOT), equity prices (Q), the 10-year T-Bill rate (LR), and the 3-month T-Bill rate (R). All the responses are expressed in percentage terms. The x-axis denotes years. In each chart, the darker area represents the 68% confidence interval, while the sum of the darker and lighter areas denote the 90% confidence interval.
The system with the LS survey is estimated over the period 1960H1-2007H1, while that using the SPF survey is estimated over the period 1968Q4-2007Q2.

Figure 5. Responses to a Shock to Expected Inflation: Additional Activity Indicators

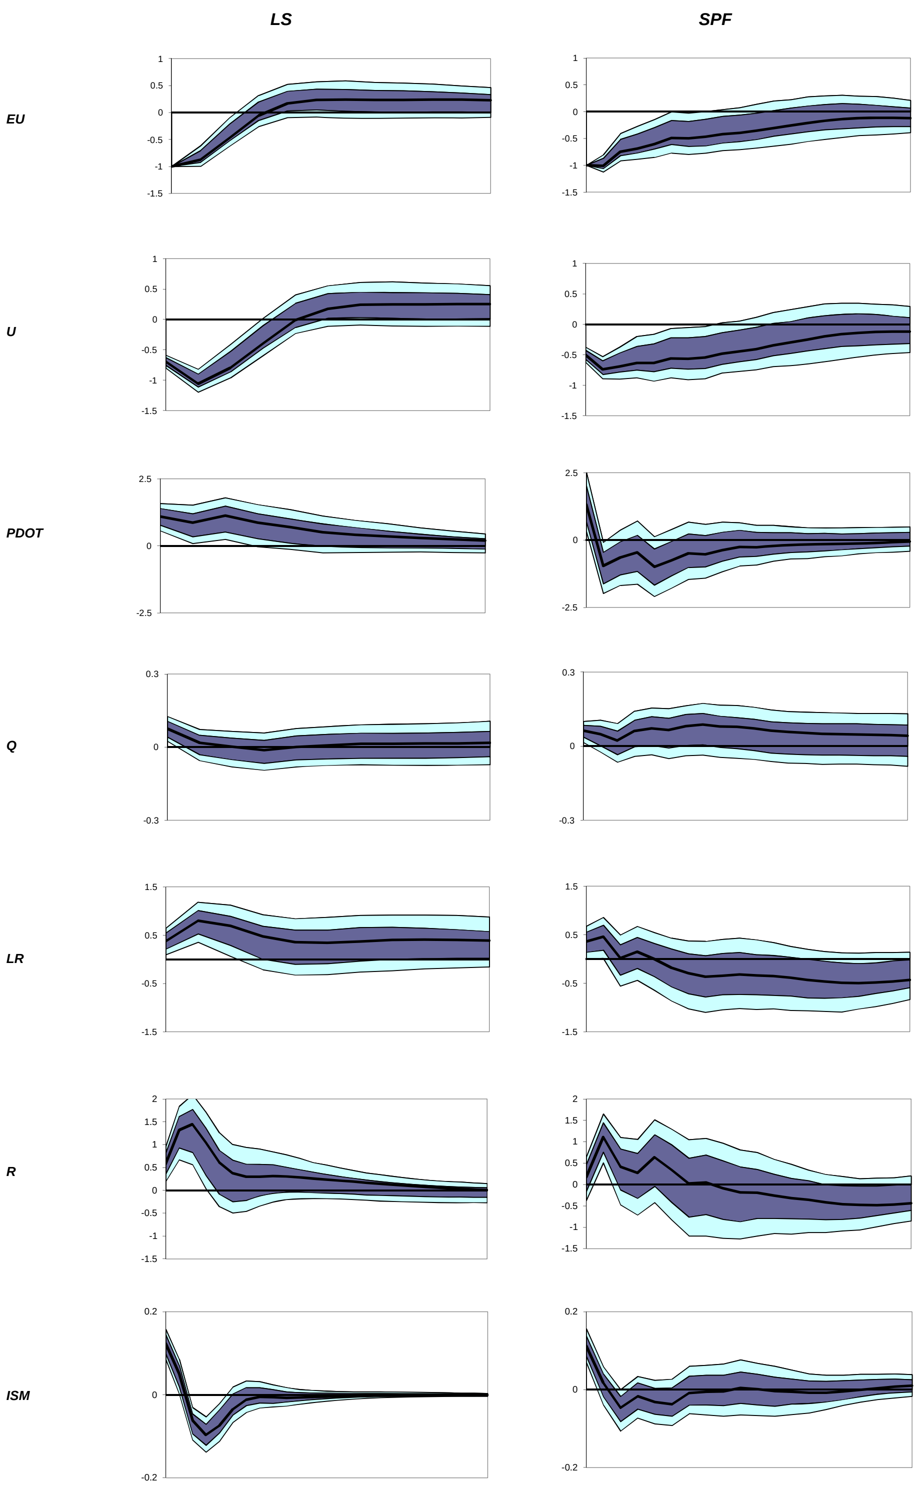

The responses were generated from a VAR with expected unemployment (EU), actual unemployment (U), inflation (PDOT), equity prices (Q), the 10-year T-Bill rate, the 3-month T-Bill rate, and the ISM survey (ISM). All the responses are expressed in percentage terms.
The x-axis denotes years. In each chart, the darker area represents the 68% confidence interval, while the sum of the darker and lighter areas denote the 90% confidence interval. The system with the LS survey is estimated over the period 1960M1-2007M4, while that using the SPF survey is estimated over the period 1968Q1-2007Q2.

Figure 6. Responses to a Shock to Expected Unemployment: Longer Horizon Forecasts

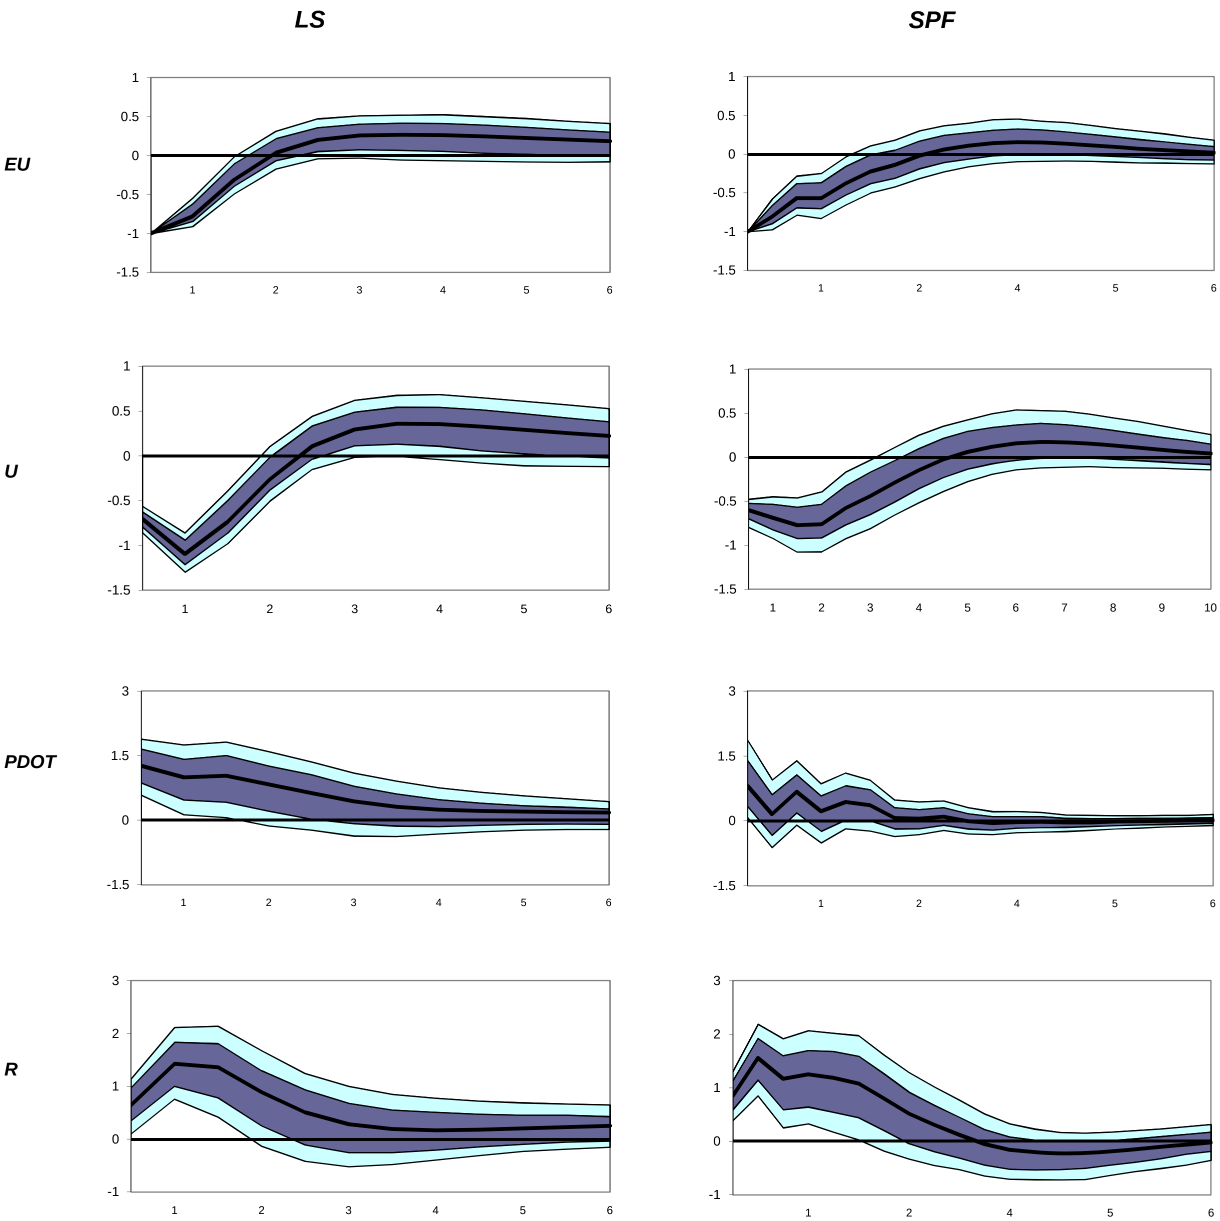

The responses were generated from a VAR with expected unemployment (EU), actual unemployment (U), inflation (PDOT), and the 3-month T-Bill rate. All the responses are expressed in percentage terms. The x-axis denotes years. In each chart, the darker area represents the 68% confidence interval, while the sum of the darker and lighter areas denote the 90% confidence interval. The system with the LS survey is estimated over the period 1960H1-2007H1, while that using the SPF survey is estimated over the period 1968Q4-2007Q2.

Figure 7. Responses to a Shock to Expected Unemployment (Extended sample)

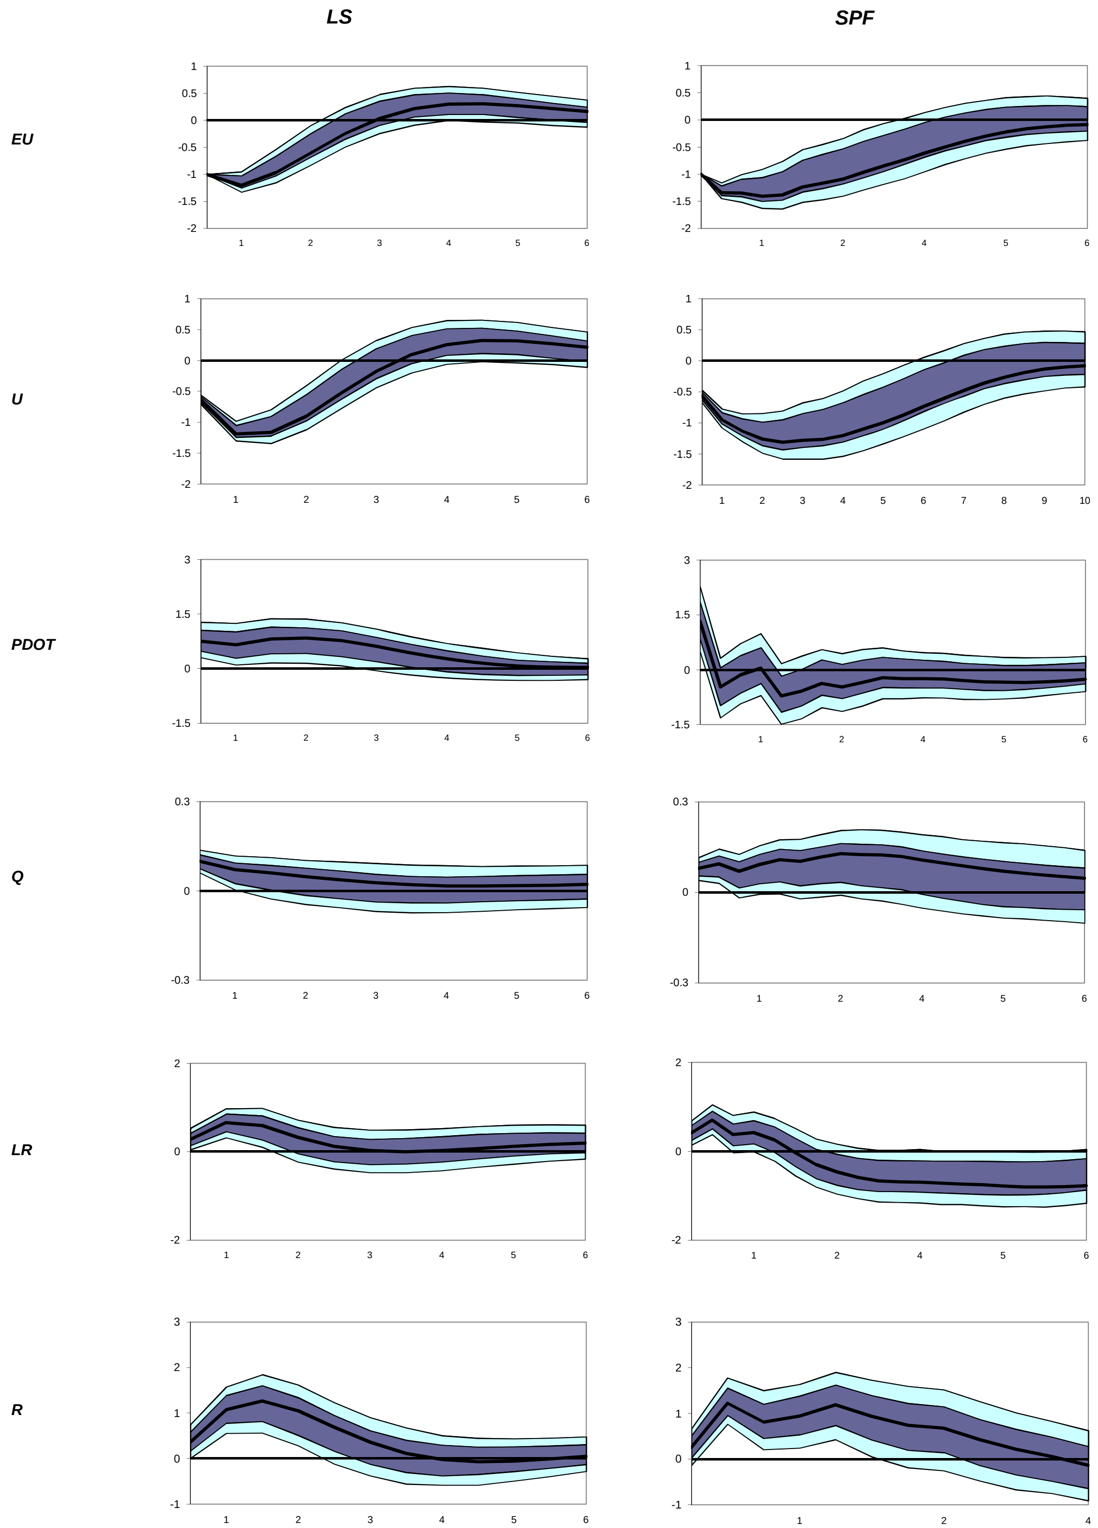

The responses were generated from a VAR with expected unemployment (EU), actual unemployment (U), inflation (PDOT), equity prices (Q), the 10-year T-Bill rate (LR), and the 3-month T-Bill rate (R). All the responses are expressed in percentage terms. The x-axis denotes years. In each chart, the darker area represents the 68% confidence interval, while the sum of the darker and lighter areas denote the 90% confidence interval.
The system with the LS survey is estimated over the period 1960H1-2009H1, while that using the SPF survey is estimated over the period 1968Q4-2009Q1.

Figure 8. Michigan Survey Unemployment Expectations vs Changes in the Unemployment Rate

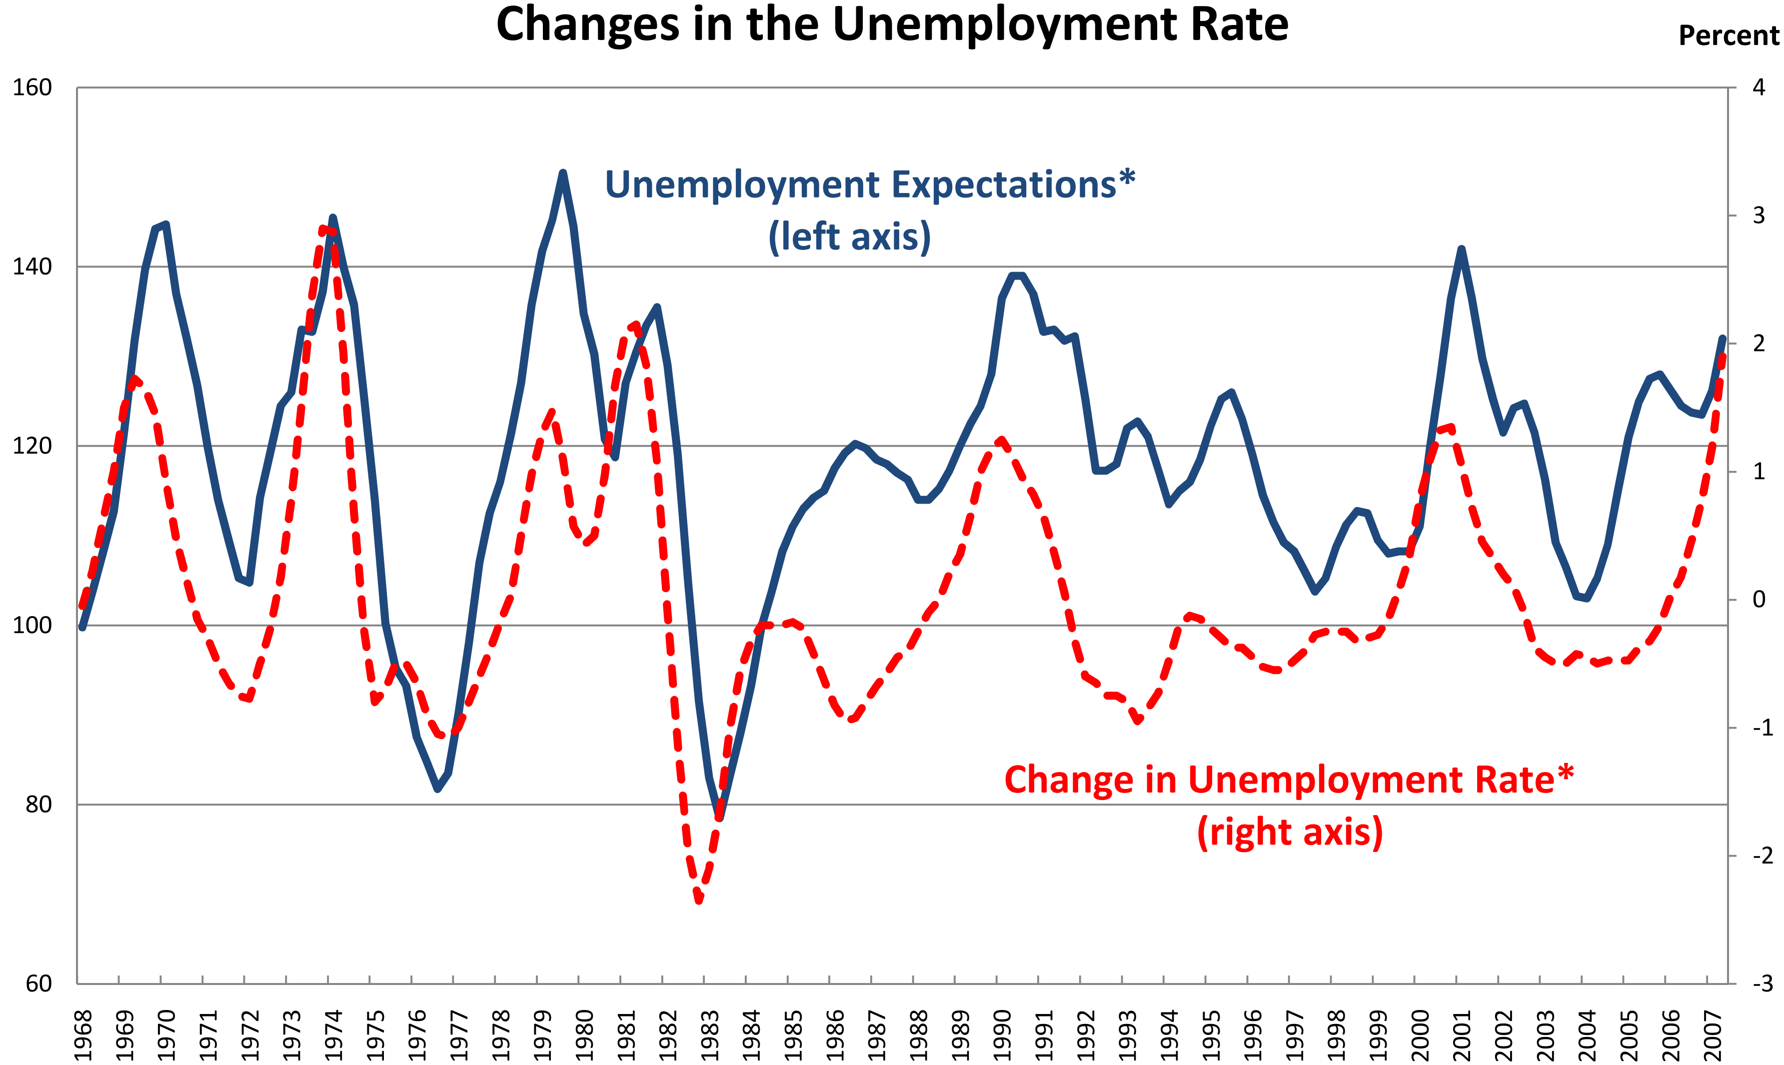

* 4-quarter moving average

Figure 9. Responses to a Shock to Expected Unemployment: Michigan Survey

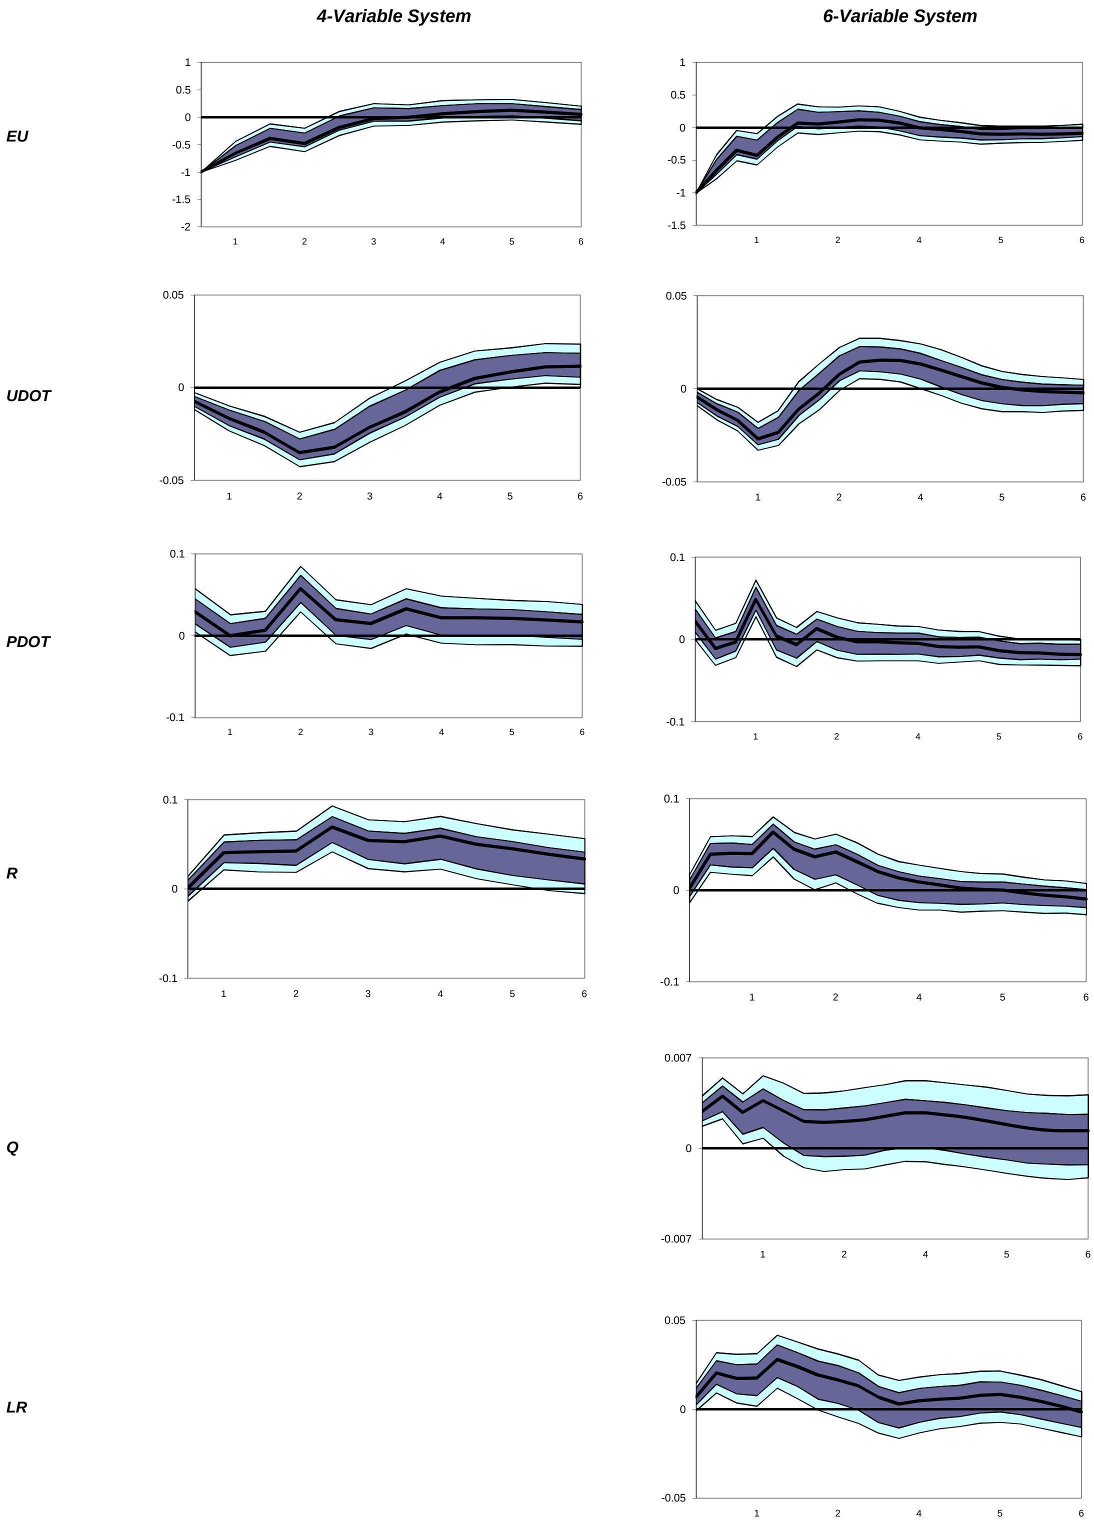

The responses for the 4-variable system were generated from a VAR with expected unemployment (EU), the year-over-year change in unemployment (UDOT), inflation (PDOT), equity prices (Q), the 10-year T-Bill rate (LR), and the 3-month T-Bill rate (R). Equity prices (Q) and the 10-year T-Bill rate (LR) were added to those variables for the 6-variable system. All the responses are expressed in percentage terms. The x-axis denotes years. In each chart, the darker area represents the 68% confidence interval, while the sum of the darker and lighter areas denote the 90% confidence interval. The systems were estimated over the period 1978Q1-2007Q2.