

Unemployment Expectations and the Business Cycle*

Daniel L. Tortorice[†]

Brandeis University

April 2010

Abstract

I compare unemployment expectations from the Michigan Survey of Consumers to VAR forecastable movements in unemployment. I document three key facts. First, one-half to one-third of the population expects unemployment to rise when it is falling at the end of a recession even though the VAR is able to predict the fall in unemployment. Second, more people expect unemployment to rise when it is falling at the end of a recession than expect it to rise when it is rising at the beginning of a recession even though these movements are predictable with the VAR. Finally, the lag change in unemployment is as important as the VAR prediction of the future unemployment change in predicting the fraction of the population that expects unemployment to rise. Least squares learning or real time expectations do little to help explain these facts. However, delayed updating of expectations can addresses some of these puzzles and extrapolative expectations addresses these puzzles the best. Individuals with higher income or education are only slightly less likely to make these expectational errors and those who makes these errors are 8-10 percent less likely to believe it is a good time to make a major purchase.

Keywords: consumer sentiment, rational expectations, business fluctuations; cycles.

JEL Code: E32, E37

*I thank Alberto Alesina, Laurence Ball, Christopher Carroll, George Hall, David Laibson, Hong Li, N. Gregory Mankiw, Ricardo Reis, Andrei Shleifer and participants of the Harvard macroeconomics lunch, Brandeis faculty workshop, and the 2010 Midwest Macroeconomics Conference for helpful comments and suggestions. I thank Richard Curtin for providing the data. This research was supported by the Jacob K. Javits Fellowships Program funded by the United States Department of Education. Sole responsibility for any and all errors in this article belongs to the author. Comments and questions should be sent to tortoric@brandeis.edu.

[†]Department of Economics and International Business School, Brandeis University, tortoric@brandeis.edu

1 Introduction

In modern macroeconomic theory, expectations about the future drive current economic activity. In general, models assume rational expectations: agents know the true model and use it to form expectations. However, this assumption is not without controversy. Sargent (2001) and Evans & Honkapohja (2001) advocate models where agents learn the true model over time. Mankiw & Reis (2002) argue for models with agents who use outdated information to form their expectations. Given this plurality of views, work testing models with micro-level expectations data can provide important insight into the formation of expectations.

Consequently, I provide a detailed analysis of unemployment expectations from the Survey of Consumers conducted by the University of Michigan Survey Research Center. The measurement of unemployment expectations stems from 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?" Answers fall into five categories: more unemployment, about the same, less unemployment, do not know, did not answer.

I compare the sign of the unemployment change that the respondents expect with forecastable movements in unemployment (predicted by a four variable VAR). There are important differences between the consumers' expectations and forecastable movements in unemployment. The three key discrepancies are: 1. In the six months after a recession's end—when unemployment is falling—one-half to one-third of the population expects unemployment to *rise* even though the VAR predicts the *fall* in unemployment. 2. At a recession's beginning—when unemployment is rising—fewer people expected unemployment to rise than at the recession's end – when unemployment is falling— even though the rise and fall of unemployment is forecastable with the VAR. 3. Controlling for the predictions of the VAR, the lag change in unemployment is as important as the VAR prediction of the future unemployment change in predicting the percent of the population that expects unemployment to rise.

Given the discrepancies between household expectations and VAR forecasts, I examine the ability of various models of expectations to match these facts. A least squares learning

model and a model that uses real time data fail to resolve these puzzles. A model with delayed updating of expectations can help explain why there are so few pessimists (defined as individuals who expect unemployment to rise) at the beginning of a recession but not why there are so many pessimists at the end of a recession. An extrapolative expectations model where agents partially form expectations by extrapolating current trends into the future can address all the puzzles.

Additionally, those with more education or a higher income are less likely to make expectational errors: however the difference is not economically significant. Expectational errors are not confined to an economically unimportant fraction of the population, but evenly distributed across income and education groups. Making expectational errors affects attitudes concerning whether it is a good time to make a major purchases, e.g.: a house, a car, or a durable good. Specifically, when an individual makes a pessimistic expectational error, that is expecting unemployment to rise when in fact the VAR predicts it will fall, they are 8-10% less likely to think that it is a good time to make a major purchase.

Carroll (2003) and Mankiw *et al.* (2003) use the Michigan data to evaluate models of expectations. However, these papers test only one expectations model, where individuals infrequently update their information set. Additionally, these papers mainly study inflation expectations. While the highest quality data available are for inflation expectations, there are serious limitations of focusing on inflation expectations. First, inflation expectations are most important in models of price setting. Since this is a firm's decision, it is not clear that the populations they study correspond to individuals choosing prices in macroeconomic models¹. Second, for the survey population where expectational errors are most likely and empirically are most significant, households from the Michigan survey, macroeconomic models suggest that expectations about future consumption, or the future state of the economy, are more important in determining household decisions than expectations about inflation.

Souleles (2004) analyzes the survey's questions concerning expectations of future business

¹They study households from the Michigan Survey of Consumers, professional forecasters from the Livingston Survey, and professional forecasters from the Survey of Professional Forecasters

conditions, financial positions, and household income and finds the forecast errors exhibit excess sensitivity, do not average out to zero over the twenty year sample period, and are correlated with demographic variables. My paper differs from the work of Souleles by focusing on unemployment expectations and business cycle induced changes in unemployment.

Carroll (2003) focuses on inflation expectations but also studies the Michigan unemployment expectations index: the percent of individuals who expect less unemployment minus the percent that expect more unemployment. He shows the dynamics of the index are well modeled by an equation that puts one-third of its weight on the professional forecast and two-thirds weight on the lag value of the index. Curtin (2003) shows that the same index is correlated with future unemployment changes. However, when he regresses the unemployment change on changes in the index, he finds the residuals are autocorrelated.

Like Carroll and Curtin I find evidence suggestive of serial correlation in the expectational errors of household though my work differs from theirs in substantial ways. First, the facts concerning excessive pessimism at the end of a recession and insufficient pessimism at the beginning of a recession are new. Second, I test the ability of a large number of theories to account for these facts. Third, I use the methodology of Carlson & Parkin (1975) to provide a rigorous mapping from models of expectation formation to equations that relate macro aggregates to aggregates of qualitative expectations like the indexes in the work of Carroll and Curtin. This methodology generates important insights. I show that expectations based solely on a distribution around the VAR expectation underestimate by a factor of 15 the importance of lag unemployment in predicting the fraction of people who expect unemployment to rise. Finally, contrary to Carroll, I find features of the data unaccounted for by the delay model: for example the large number of individuals expecting unemployment to rise at the end of a recession. Due to the mean reverting nature of unemployment, the delay model is unable to generate these pessimistic predictions while they are consistent with extrapolation.

The rest of the paper proceeds as follows. Section 2 discusses the data while section 3

establishes the three key empirical facts outlined above. Section 4 examines the ability of four different models of expectation formation to address these facts. Section 5 examines data from simulated, heterogenous, structural expectations models. Section 6 shows that a univariate forecast can correctly predict the sign of unemployment changes. Section 7 examines the ability of individual characteristics to predict expectational errors and the influence of expectations on buying attitudes. Section 8 concludes.

2 Data

The data come from the Survey of Consumers conducted by the University of Michigan Survey Research Center. It is a monthly survey used to calculate the index of consumer sentiment. Observations begin in 1978 and end in February of 2004². There are about 500 respondents per month. The survey asks about respondents' demographic characteristics, expectations of inflation, unemployment and interest rates, views on the current economic state, and attitudes towards the purchase of economically significant items.

I study answers to 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?" Answers fall into five categories: more unemployment, about the same, less unemployment, do not know, did not answer. Only about 1.5% respond that they do not know or do not answer the question. On average, half expect unemployment to stay the same, 35% expect it to get worse, while only 15% are optimistic and expect less unemployment. (Full summary statistics for the data are available in the appendix: table 1A)

This question is not ideal. It is not quantitative. Additionally, the meaning of "about the same" may differ among respondents and it is not clear what measure of unemployment the respondent will use in her answer. However the question does refer to a precise forecast

²I am primarily interested in expectations around recession dates. Preliminary analysis of the data from the current recession confirms one of the main points of this paper. Namely, households seem surprised by the rise in unemployment at the beginning of the recession. However, until a full 12 month period of recovery comes it is not possible to test if agents are also surprised by the recovery.

window. Additionally, two of the responses reveal the sign of the respondent's unemployment expectation making it possible to test the accuracy of their prediction about the sign of future unemployment changes. This analysis is the paper's primary concern.

The data contain information on the respondent's: sex, age, education, marital status, income and race. Mean household income in year 2000 dollars is 48,510. The mean age is 45, mean education is 13 years, and 31% have graduated college. The percent of the sample that gave a valid answer to the questions is high, ranging from 93% for income to 100% for if the respondent was a college graduate, which speaks to the reliability of these data. Additionally, the appendix reports deciles of the income and education distribution which are consistent with that of the U.S. population as a whole. Finally, 85% of the sample is White, 9% Black, 4% Hispanic, 1% Asian, and 1% Native American.

The survey asks the respondents about economically significant purchases: houses, cars and durable goods. The first question is: "Generally speaking, do you think now is a good time or a bad time to buy a house?" The answer is grouped into five categories: good, pros and cons, bad, doesn't know, or did not answer. The durable goods question is: "About the big things people buy for their homes – such as furniture, a refrigerator stove, television, and things like that, generally speaking, do you think now is a good or a bad time for people to buy major household items?" The groupings of responses are the same as for the house question. Finally, the interviewer queries for the respondent's attitude about a future car purchase: "Speaking now of the automobile market – do you think the next 12 months or so will be a good time or a bad time to buy a car?" with the same groupings of responses as with the items above.

Individuals have upbeat buying attitudes. They think it is a good time to buy a house or durable goods about two-thirds of the time and think it is a good time to buy a car more than half the time. Most individuals hold an opinion. Only 3% do not know or give no answer on whether or not it is a good time to buy a house and only 6-7% respond without an opinion to the queries about buying a car and buying durable goods.



3 Empirical Facts

Figure 1 plots the unemployment rate and pessimism (the fraction of people who expect unemployment to rise over the next year). Vertical lines mark the peak and trough of unemployment. At the peaks, around 35% of the population expects unemployment to rise, even though it is about to *fall* sharply. At the troughs a similar percent of the population expects unemployment to rise, even though it is about to *rise* sharply. During the recession of 2000, pessimism was 10 percentage points lower at the trough than at the peak of unemployment. If unemployment changes were not forecastable then pessimism should not be sensitive to the level of unemployment. However, unemployment is strongly mean reverting and we would expect future unemployment changes to be forecastable. This plot then suggests that, relative to a statistical forecast, pessimism levels may be too high at the end of a recession and too low at the beginning of a recession. It also suggests that individual expectations may be, in part, backward looking. This section establishes these facts more formally.

To examine pessimism at the recession's end, when unemployment is falling, and the

extent to which these movements are predictable I use the following three regressions³:

$$Y_{it} = X_t\beta + \varepsilon_{it} \quad (1)$$

Y_{it} equals P_{it} , ΔU_t , or $E\Delta U_t$. P_{it} is one if the individual expects there to be more unemployment at time $t - 12$ and zero otherwise⁴, ΔU_t , the change in the unemployment rate, is $u_t - u_{t-12}$ ⁵ and $E\Delta U_t$ is the VAR prediction of ΔU_t . I use a four variable (log GDP, CPI inflation, the fed funds rate, and unemployment), four lag VAR beginning in 1956^{6,7}.

X_t is a vector of twelve dummy variables each indicating a date corresponding to a specific number of months (zero to eleven) after the first time ΔU_t is negative at the end of a recession. (Since future unemployment changes become negative before the unemployment peak, these dates correspond to unemployment changes beginning a few periods before the unemployment peak dates.) I have data on expectations beginning in 1978, and treat the two recessions of the 1980s as one recession⁸, so each of these dummy variables equals one at exactly three dates. The coefficients on X_t give the mean level of the dependent variable across these three dates⁹.

³This regression is related to the analysis of inflation expectations in Coibion & Gorodnichenko (2008). They draw conclusions concerning different expectations models from the response of inflation expectations to various shocks. I show that the level of pessimism that exists after a recession is also useful for drawing conclusions concerning different expectational models.

⁴Estimating the first equation by collapsing the cross section into a single estimate for the percent of pessimists and using the resulting time series gave similar results.


⁵It is possible that the survey question elicits expectations not about the strict change in unemployment but the average level of unemployment over the next 12 months minus the current level of unemployment. Repeating the analysis in the paper with this definition did not change the main conclusions.

⁶More complicated VARs could be considered. But I will show that even this simple VAR is able to predict unemployment changes and individuals perform much more poorly than this simplistic, potentially misspecified, VAR.

⁷GDP is available only quarterly and the unemployment expectations are measured monthly. I require a procedure to impute monthly expectations from the quarterly VAR. Figure 1A in the appendix explains how I do this. The procedure results in VAR based expectations that are slightly lagged. Since I find that individuals look as if they are forecasting lagged unemployment changes relative to the unemployment changes they are asked to forecast, this procedure pushes the VAR closer to the data and strengthens the conclusion that the VAR does not fully represent the expectations contained in the data. Section 6 examines the impact of this assignment procedure and finds it to not be significant.

⁸I do this because there is no sustained unemployment recovery from the first recession of the 1980s.

⁹The starting points for where ΔU_t is first negative are: 1983 month 7, 1993 month 2 and 2003 month 12. The first element of X_t equals one for these dates and zero otherwise. The second element of X_t equals one on the dates 1983 month 8, 1993 month 3 and 2004 month 1 and zero otherwise, and so on.




The x axis denotes months since the unemployment change is first positive (negative) at the beginning of (after) a recession.

I repeat these regressions to estimate the levels of pessimism, actual unemployment changes, and average forecasts before the recession as well, when unemployment begins to rise. For this calculation, I replace X_t with a vector of dummy variables that mark months after unemployment begins to rise at the beginning of a recession¹⁰.

Figure 2 plots the coefficients from the pessimism regressions. (Full regression results with standard errors are given in tables 2A and 3A in the appendix.) The solid line gives the average level of pessimism when unemployment begins to fall (month 0) up until 11 months into the recovery. The dashed line gives the average level of pessimism when unemployment begins to rise (month 0) up until 11 months into the recession. The levels of pessimism when unemployment begins to fall are high. One half to one third of the population is expecting unemployment to rise, when in fact it is falling. On the other hand, only a third of the population expects unemployment to rise when it is actually rising, at the beginning of the recession. In fact, there is substantially more pessimism at the beginning of the recession (when unemployment is falling) than at the end of the recession (when unemployment is rising).

Figure 3 plots the coefficients for the actual unemployment change (solid line) and VAR forecasted unemployment change (dashed line) regressions. The coefficients at the beginning

¹⁰The starting points for when ΔU_t is first positive are: 1980 month 1, 1990 month 7 and 2001 month 1.



Note: Solid line is the actual unemployment change, the dashed line is the change forecasted from the VAR, the x line is the change forecasted from the LSL model, and the o line is the forecast from the real time data mode. The x axis denotes months since the unemployment change is first positive (negative) at the beginning of (after) a recession. The series above zero are the forecasts at the beginning of the recession. The series below zero are the forecasts after the recession.

of the recession are the lines which are above zero, indicating that unemployment is rising and the VAR correctly forecasts the sign of these changes. The coefficients at the end of the recession are the lines below zero indicating that unemployment is falling and the VAR forecasts these changes as well.

To test the statistical significance of these results I estimate the following three regressions: $Y_{it} = X_t\beta + Y_t\gamma + \varepsilon_{it}$ where $Y_{it} = \{P_{it}, \Delta U_t, E\Delta U_t\}$, X_t is a vector of dummy variables indicating zero months to five months after unemployment begins to fall at the end of a recession, and Y_t is a vector of dummy variables indicating zero months to five months after unemployment begins to rise at the beginning of a recession. This regression, results in Table 1, shows the pattern of more pessimism during the recovery than during the recession's beginning while the signs of the VAR prediction would have us expect the opposite.

This table tests if the coefficients are significantly different from each other. The first three coefficients on X_t are significantly different from the corresponding coefficients on Y_t indicating statistically significantly more pessimism after the recession than during the beginning of the recession. Even the results for month four and five are puzzling. We see a

statistically significant, 1.6 to 2 percentage point difference in unemployment forecasts but no statistically significant difference in the levels of pessimism. (I have omitted results for optimism – the percent of the population that expects unemployment to fall – but they support the main results here. There is three times as much pessimism than optimism at the end of a recession and the levels of optimism at the beginning and end of a recession are similar.)

Expectations around recession turning points appear partially backward looking. To see if lag unemployment changes influence expectations across the whole sample I run the following regression:

$$P_t = \alpha + \beta E_t(u_{t+12} - u_t) + \gamma(u_t - u_{t-12}) + \varepsilon_t \quad (2)$$

P_t is the fraction of the population at time t that expects unemployment to rise. $E_t(u_{t+12} - u_t)$ is the VAR forecast of the future unemployment change and $u_t - u_{t-12}$ is the lag change in unemployment. Table 2 contains the results from this regression¹¹. The coefficient on the VAR prediction is positive – when the VAR predicts unemployment will rise more, more people expect unemployment to rise. Nevertheless, controlling for the VAR prediction, the coefficient on the lag change in unemployment is positive and significant. Indeed, it is larger than the coefficient on the VAR expectation. The lagged unemployment change is at least as important as the VAR prediction in predicting the number of people who expect unemployment to rise at a given time.

If the VAR is an imperfect forecast then the lag unemployment rate may be correlated with actual unemployment changes even controlling for the VAR. Households may use the lag change to refine the VAR expectation. To see if this is the case, I run the regression above including the actual change in unemployment. If lagged unemployment changes are correlated with actual unemployment changes even conditional on the VAR expectation, the

¹¹I correct the OLS standard errors for autocorrelation in the residuals using a Newey-West procedure (Newey & West (1987)).

coefficient on lagged unemployment should fall. However, the coefficient on lagged unemployment does not change. The significant coefficient on lagged unemployment represents a distortion of expectations not a more accurate refinement of them.

Now it is difficult to relate a point forecast of unemployment to a population distribution of expectations. In section 5, I will consider a model with an explicit distribution of expectations around the VAR forecast. I show that there is a direct mapping in this model between the VAR forecast and the population distribution. However, this model underestimates the coefficient on lagged unemployment changes by a factor of 15, failing to explain the influence of lagged unemployment on the number of pessimists.

Taken together these facts show that individuals' expectations differ in important ways from the predictions of a VAR. They tempt us to ask what type of models can explain these beliefs. Answering this question is the aim of the next sections. First I consider different models of expectation formation then I consider models with an explicit distribution of expectations around the different model forecasts.

4 Alternative Models of Expectation Formation

4.1 Least Squares Learning

In the previous section I calculated the VAR on the full data sample. (This assumption mimics a rational expectations assumption where the agent knows the true model and calculates her expectations according to that model.) An alternative approach, is 'least squares learning' (Sargent (2001) and Evans & Honkapohja (2001)). Here the agent does not have access to the estimated VAR on all of the data, instead at each date she estimates the VAR on the data up until that date and uses this equation to forecast unemployment.

Figure 3 reports the predictions of the least squares learning model alongside the actual changes and VAR forecasts. The least squares learning (LSL) model line (x) is calculated in the same manner as the VAR line except the VAR forecast is replaced with the least

squares learning forecast in the regressions. After the recession, during the recovery, a least squares learner would expect unemployment to fall. In fact, the prediction is on average more negative than the prediction of the VAR. This result is not surprising. Recall, the actual unemployment recovery from the recessions of 1991 and 2001 was much slower than would have been predicted based on past data alone. Therefore, the least squares learner expects more of a recovery than an individual who uses the VAR equation. Figure 3 also shows the least squares learning expectation before the recession. The least squares learner does, on average, expect unemployment to rise. However, as shown in the appendix tables 2A and 3A the standard errors are fairly large and we can not reject a mean of zero.

On the one hand least squares learning provides a partial solution to one puzzle. It is hard based on only past data to predict the onset of a recession. So while on average a least squares learner correctly predicts the sign of the unemployment change there is a large amount of variance in the forecast. On the other hand, a least squares learner would be even more optimistic than a VAR forecaster during the recession recovery, deepening the puzzle.

4.2 Real Time Data

In approximating expectations with the VAR, I assume agents have access to future data revisions. Instead, for the most recent data, they have access to only the first or second release, not the final revision. As argued by Orphanides (2001) and Orphanides & van Norden (2002) the use of real time data can change measurements of output gaps and the optimal choice of monetary policy. Inflation (measured with the CPI), the fed funds rate and unemployment are subject to only minor revisions, but GDP is often substantially revised. To explore this issue, I make use of the real time data set available from the Philadelphia Federal Reserve Bank (Croushore & Stark (1999)). I assume that individuals have access to the VAR equations estimated on the whole sample, but only have access to GDP data available at the time when they made their expectation.

Figure 3 contains the results. The real time line of figure 3 (line o) is calculated like the

VAR line of figure 3 except the VAR forecasts are replaced with the real time data forecasts in the regressions. The use of real time data does not help to address the puzzles previously outlined. The expectations are similar to, and most importantly, of the same sign as the VAR. Therefore, it appears that neither the use of least squares learning nor real time data is fully responsible for the observed pattern of pessimism.

4.3 Delayed Updating of Expectations

Several authors including Reis (2004), Mankiw & Reis (2002), Gabaix & Laibson (2001) and Carroll (2003) argue it is unreasonable to expect that consumers update their information instantaneously. They argue that to do so requires a cognitive cost and therefore, to economize on this cost, the agent will update his information infrequently.

I examine the ability of this model to account for the observed puzzles. Following Mankiw & Reis (2002) and Carroll (2003), I assume that each agent has a fixed probability λ of updating her expectation in any given period. Therefore the percent of the population that has expectations based on information n periods old is: $\lambda(1 - \lambda)^n$. This formula implies that $\lambda(1 - \lambda)^n$ percent of the population has expectations that come from the following expression: $E_{t-n}[u_{t+12} - u_t]$. Here the expectation is calculated using the VAR equations estimated on the whole sample and information at time $t - n$. Importantly, this expression involves forecasting not only the future rate of unemployment but also the time t rate of unemployment when information is not completely up to date.

To implement this theory empirically, I assume that information is at most N periods old.¹² I then rescale the percent of the population with information n periods old (so that they sum to one) using the following formula: $\frac{\lambda(1-\lambda)^n}{1-(1-\lambda)^{N+1}}$.

To calibrate λ I solve the following problem: $\min_{\lambda} x'Wx$, where x is a vector of two observations per month for every month in the data set: the survey estimates of the percent of people who respond that they expect unemployment to rise (pessimists) and the survey


¹² N is chosen so that $1 - (1 - \lambda)^{N+1}$ (the percent of the population that would have information more than N periods old) is less than 5%.

estimates of the percent that expect unemployment to fall (optimists); each percent is subtracted from the model's prediction for these variables. W is a diagonal weighting matrix. For the weights I use the inverse of the variance of the survey estimates of the monthly number of pessimists and optimists.

However, the model's prediction for these percents depends on what cutoffs constitute unemployment staying the same, rising, and falling. Since the delay model will only tell me the respondent's prediction for the change of unemployment, I need a procedure to assign the quantitative prediction to the qualitative categories "more" "less" and "stay the same." I estimate c_l the lower cutoff, the point at which if the individual expects unemployment to fall by more than c_l they are classified as an optimist (expecting less unemployment) and c_u the upper cutoff, the point at which if the individual expects unemployment to rise by more than c_u they are classified as a pessimist (expecting more unemployment). Everyone else is classified as expecting unemployment to be about the same. I solve the minimization problem multiple times, each time for a different pair of cutoffs. To find the overall minimum I take the minimum across the solutions to the minimization problem at these different pairs of cutoffs¹³.

The minimization occurs at an upper cutoff of 0.1, a lower cutoff of -0.6 , and a value of λ of 0.09. A value of 0.09 implies that the agent updates her expectation roughly once every 10 quarters or once every two and a half years. This is considerably more delay than is assumed, or estimated, in the literature. Carroll (2003) estimates and Mankiw & Reis (2002) take λ to be 0.25. This implies that the agent updates her expectations once every year. However more recent work, Reis (2004) has found optimal updating to be every 8 quarters. Note also that the data want to make it comparatively easier to be pessimistic (the upper cutoff is 0.1) than optimistic (the lower cutoff is -0.6). This fact is due to the

¹³It would be better to minimize over the three parameters, λ and the two cutoffs, simultaneously. However, without adding noise, the function is not sufficiently well behaved in the cutoffs to make this a simple task. Because I want to first focus on the endogenous heterogeneity of the model, not exogenous heterogeneity from noise, I omit noise and therefore, I use the two step procedure described. In section 5, I explore estimation with noise.



tendency of respondents, on average, to more likely respond with pessimism than optimism.

Figure 4 contains the prediction of the delay model for the average number of pessimists. While the model does have some pessimists at the end of the recession, the prediction of the model (solid line with box) is much smaller than in the data (22% versus 48% at the beginning of the recovery) and the number of pessimists predicted by the model quickly falls towards zero. Turning to the model prediction at the beginning of the recession (dashed line with box), the model does a much better job matching the data, overestimating pessimism by only a few percentage points. Recall that if the agent had completely up to date information he would expect unemployment to rise, however, by a combination of the delay mechanism and the cutoffs, the model is able to suppress the number of pessimists at the beginning of the recession. One more virtue of the delay model is evident. It creates a fairly smooth endogenous distribution of expectations. One can see that the number of pessimists varies fairly smoothly from month to month.

Why does the delay model fail to match the percent of pessimists at the end of the recession? To answer this question I examined forecasts of unemployment changes at each point in time based on information which gets progressively older. (Table 4A in the appendix gives the details of this calculation.) The calculation showed that even with outdated information

the forecast of the unemployment change at the recession's end is almost always negative. To understand the intuition for this result it is useful to think about two cases.

First think about a situation where the agent has very old information. He would predict more or less no change in the unemployment rate. This is intuitive. Most economists would not have an expectation about the change in the unemployment rate six years into the future from five years into the future. They would predict more or less no change in unemployment. So very old information does not make someone expect unemployment to rise.


Second, consider a situation where the agent has only slightly old information. For example, the agent knows that unemployment rose substantially a few periods ago and does not know what happened since. An agent using a statistical forecast would not expect that past unemployment change to persist indefinitely into the future. Since he knows that unemployment changes eventually flatten out and then become negative – when forecasting what happens at the end of the recession he expects close to no change in unemployment over the next year or even a substantial fall in unemployment.

It is instructive to contrast this case with the case of forecasting inflation. Since inflation is close to a random walk, a forecast of inflation today, based on information from two years ago, will not differ much from the level of inflation two years ago. In this case, he will look fairly extrapolative, using past inflation rates to forecast today's inflation. However, unemployment *changes* are certainly not a random walk. And forecasts of unemployment changes today, based on old information, are not the same as unemployment changes from a few periods in the past.

4.4 Extrapolative Expectations

Next I consider a model where a percent of the population form their expectations by taking a weighted average of past changes. These agents form expectations according to:

$$\hat{E}_t[u_{t+12} - u_t] = \sum_{n=0}^{11} \frac{\lambda(1-\lambda)^n}{1-(1-\lambda)^{12}} [u_{t-n} - u_{t-n-12}] \quad (3)$$



where u_t is the unemployment rate at time t and λ is a parameter that controls how much the recent past is weighted relative to the less recent past.

To choose the percent of extrapolators, e , I solve the following problem: $\min_e x'Wx$ where, as in the previous section, x is a vector of two observations per year: the percent of people who respond that they expect unemployment to rise and the percent that expect unemployment to fall; each percent is subtracted from the model's prediction for these variables. W is a diagonal weighting matrix with the inverse of the variance of the data estimates on the diagonal (as in the previous section). I calibrate $\lambda = .25$. The estimate of e and the model's predictions were not very sensitive to the choice of λ . Again, the model's prediction depends on what cutoffs constitute unemployment staying the same, rising, and falling so I calculated e for different values of these cutoffs and then choose the overall minimum. The objective function is minimized at the values $e = 0.48$, the upper cutoff = 0.4 and the lower cutoff = -1.5.

Figure 5 shows the model's prediction for the average number pessimists. Extrapolation captures two features that the other models have not. First, the model predicts the large and lasting pessimistic predictions of the recovery (solid line). The extrapolators are pessimistic at the end of the recession. (The level of pessimism does not equal 0.48 for all months at the end of the recession because for the last recession unemployment was flat for a while before

falling.) Secondly, it matches the fact that there is more pessimism after the recession than at the beginning of the recession. Before the recession we have increasing pessimism from the rational agents but no pessimism from the extrapolators. (The level of pessimism does not equal 0.52 for all months at the beginning of the recession because for the first recession the rational forecast is less positive than the cutoff.)

5 Heterogeneity

In this section I relate a point forecast of an unemployment change to a distribution of unemployment expectations. I estimate models where there is a distribution of expectations around the predictions of the different expectational models.

I show that estimation of the VAR, least squares learning, and real time expectations models leads to a large variance of the distribution and unreasonable cutoffs for assigning individuals into the optimist or pessimist categories. Estimates of the delay parameter indicate substantial delay and estimates of the percent of extrapolators indicate significant extrapolation. Finally, only the extrapolation model comes close to explaining why the lag unemployment change helps predict the number of pessimists (table 2) and why there is more pessimism after the recession than before the recession (figure 2).

For the VAR model, the least squares learning model, and the real time data model I assume that at every point in time the distribution of expectations in the population is normally distributed with the mean being the expectation from the VAR model, the least squares learning model, or the real time data model, and with variance σ^2 which is a parameter to be estimated¹⁴. One interpretation is that individuals observe the VAR forecast with noise and the noise has variance σ^2 . Furthermore, I estimate c_l the lower cutoff, the

¹⁴Another method would follow the approach of Howrey (2001) and calculate the probability of unemployment rising at each point in time by sampling from the VAR residuals. I could use this probability to predict the percentage of pessimists at each point in time. I do not follow his approach for two reasons: 1) It is unclear why some individuals would expect a certain path of VAR residuals when they have an expected value of zero and 2) this method puts a restriction on the variance of the distribution around the VAR expectation. I prefer to make the variance a free parameter to give the model the best chance of matching the data. Even then I will find severe limitations of the VAR model.

point at which if the individual expects unemployment to fall by more than c_l they are classified as an optimist and c_u the upper cutoff, the point at which if the individual expects unemployment to rise by more than c_u they are classified as a pessimist. Otherwise, they are classified as expecting unemployment to stay the same. To estimate the parameters of the model I solve the following problem:

$$\min_{\sigma, c_u, c_l} x'Wx \quad (4)$$

where x is a vector of two observations per month for every month in the data set: the percent of people who respond that they expect unemployment to rise and the percent that expect unemployment to fall each subtracted from the model's prediction for these variables and W is a diagonal weighting matrix with the inverse of the variance of the data estimates on the diagonal (see section 4.3).

For the delay model, the estimation procedure is similar. First the population is broken up into individuals who base their expectations on information n periods old, for $n = 1, \dots, N$ ($N = 87$). The percent of the population in each group is given by the same formula as in section 4.3. Then within these groups the expectations are assumed to be distributed normally with mean equal to the expectation based on information n periods old and variance σ^2 which is estimated along with the upper and lower cutoffs. The exact problem solved is (4) with the predictions of the delay model replacing the predictions of the VAR model and also optimizing over the delay parameter λ .


Finally, for the extrapolation model, I assume that at each point in time a fraction $1 - e$ of the population has expectations which are normally distributed around the VAR expectation with variance σ^2 . Another fraction, e of the population has expectations normally distributed around the extrapolation equation (3) with the same variance. To estimate the parameters I solve (4) optimizing over e and λ from equation (3) as well as σ and the cutoffs.

Table 3 displays the estimated parameters. The models all need a large σ to match the


data. The large σ is needed because at every point in time, even when the VAR expectation is significantly positive or negative, there are a substantial number of pessimists and optimists. The standard deviation is roughly four for the VAR and real time models and above seven for the learning model. It is about 1.5 for the delay model and 2.8 for the extrapolation model. In addition, as shown in Table 1A, since there are many people who answer "about the same" to the unemployment question the large standard deviation forces the upper and lower cutoffs to be fairly large. For example in the VAR model, an individual must expect unemployment to rise by more than 1.72 percentage points before they will answer "more unemployment" as opposed to about the same. This seems unreasonable; who would think that an increase in unemployment by 1.5 percentage points is unemployment being "about the same"? This discrepancy is more dramatic for the lower cutoffs since people are unconditionally much more likely to be pessimistic than optimistic. The lower cutoffs are: -4, -9, -6 for the VAR, LSL and real-time data models respectively. The upper cutoffs are more reasonable for the delay and extrapolation model: 0.57 and 1.1 respectively. The lower cutoffs for these models are -2 and -3 respectively.


The delay parameter, λ , is estimated to be 0.126 and the percent of extrapolators is estimated to be 0.38. We can rank models by the minimum of the objective (4) that is obtained. The delay model does best, followed by the extrapolation model then the VAR model, the LSL model and then finally the real time data model. Even though the extrapolation model does best in matching the facts highlighted in this paper, the delay model does better by this metric. This is because the delay model creates a large amount of heterogeneity. At each point in time there are 87 different forecasts, which differ based on the age of the information they use, and the model assumes a random distribution of individuals around each of these forecasts. This strategy creates an enormous amount of heterogeneity which helps the delay model best fit the overall data.'

Figures 6-8 redo the analysis of figure 2 with the simulated data. (Full regression results are available in the appendix: Tables 5A and 6A.) First, the simulated models come fairly



Note: Solid line is the actual level of pessimism after the recession, the dashed line is the actual level of pessimism at the beginning the recession. Along with these series, I plot the prediction of the VAR model (x), LSL Model (circle) and Real Time Model (triangle). For these models too, solid is after recession, dashed is at the beginning of the recession.





close to the actual percent of pessimists in the data. They do this however with a large estimate of sigma and therefore by making the percent pessimists and optimist fairly insensitive to the mean of the distribution, the expectational forecast. All models underestimate the level of pessimism at end of the recession. The extrapolation model comes the closest to matching the data. It underestimates the level of pessimism by about 8 percentage points, while the other models are off by 12 percentage points. One of the important facts in this paper was that the level of pessimism was higher after the recession than at the beginning of the recession. More people expect unemployment to rise at the end of a recession than do at the beginning of a recession. Only the extrapolation model matches this fact. All other models predict lower levels of pessimism at the end of a recession (solid line) than at the beginning of the recession (dashed line).

Table 4 regresses the simulated fraction of pessimists on the VAR forecast and the lag change in unemployment to see if the models can replicate the relationship outlined in table 2. In the data, for this regression, the coefficients on the VAR forecast and the lag change were almost the same. The extrapolation model comes closest to matching this fact: the coefficient on the VAR forecast is 1.5 times as large as the coefficient on lag change. For the delay model, the difference between the coefficients is a factor of 4. The rational models


fair poorly. The VAR based model and the LSL model predict that the coefficient on the VAR forecast should be 20 times the coefficient on the lag change. The real time data model predicts the coefficient on the lag change should be negative.

6 Univariate Forecasting

While it seems that extrapolative expectations helps understand these expectations data one might wonder if extrapolation is the correct structural model for the agents' expectations. It might be possible that extrapolation fits the data because it arises from a different rule that the agents follow instead of just blind extrapolation. Ball (2000) explores this possibility. He argues that the true structural model is one in which agents make univariate forecasts (the inflation forecast from a regression of a variable on only its lags). He shows his model implies that inflation expectations will look adaptive in the post-war period because a univariate inflation forecast using post-war inflation data puts a weight close to one on lagged inflation. However, the model's predictions would be much different for the pre-war data when inflation was not very persistent and consequently a univariate forecast does not look much like adaptive expectations at all. To see if a univariate forecasting rule can explain the patterns of pessimism I repeat the regressions (1) for before and after the recession but instead of using the VAR forecast I use the univariate forecast from regressing unemployment on four lags of unemployment.

Figure 9 contains the results. The univariate forecasts are generally smaller in magnitude than those of the VAR. However the univariate forecasts are of the same sign as the VAR forecasts. This model therefore cannot explain the respondent's erroneous prediction of the sign of the unemployment change.

Testing the Ball model also reveals how puzzling these pessimistic expectations really are. Even a naive and simple forecasting technique, ignoring all information in the economy except the past rates of unemployment is able, on average, to correctly predict the sign of



future unemployment changes. This achievement, however, is a task many households in the survey are unable to perform.

Finally, figure 9 also displays the exact univariate forecast. Recall that since GDP is only available quarterly, it was necessary to use the procedure described in figure 1A to assign monthly expectations from the quarterly VAR. Since unemployment is available monthly I am able to drop this assumption— which was used to calculate the univariate forecast—and use the exact unemployment expectation from the univariate unemployment forecast, which is used to calculate the univariate forecast (exact) line. As one can see the two univariate forecasts are similar and therefore indicate that the assignment procedure is fairly innocuous. In fact, the exact expectations are farther from zero, supporting my previous claim that the assignment procedure brings the VAR expectation closer to the data. Despite this aid, the VAR still misses important features of the expectations data.

7 Who Make These Errors and Does it Matter?

7.1 Characteristics Associated with Expectational Errors

The previous sections of the paper have argued that individual's unemployment expectations contain predictable patterns of errors across the business cycle, namely insufficient pessimism at the beginning of recessions and excessive pessimism during recoveries. Now, I use the survey data to examine characteristics associated with those who make errors and the impact unemployment expectations errors have on buying attitudes.

To examine who is more likely to make errors in their unemployment expectations I estimate the following regressions: $Y_{it} = F(\alpha + X_{it}\beta + Z_t\gamma) + \varepsilon_{it}$. Here Y_{it} = either $Error_{it}$ which equals one if the individual's response differs from the VAR forecast and is zero otherwise¹⁵ or $PessimisticError_{it}$ which equals one if the individual expects unemployment to rise and the VAR predicts it will not rise and it is equal to zero otherwise. X_{it} is a vector of individual characteristics and Z_t contains the lag unemployment rate and year fixed effects. I estimate the model both by OLS ($F(x) = x$) and probit ($F =$ standard normal cdf). While the data allow me to control for many individual characteristics, including the individual's own assessment of his finances, it is not a panel survey and I can not control for unobserved heterogeneity with individual fixed effects.

Table 5 contains the regression results. The first column contains the OLS results; the second column contains the marginal effects (evaluated at the means) from the probit model. I find that men are 2-3% less likely to make expectational errors. I find a non-linear effect of age. Though the best and worst by age differ only by about two percent. An additional grade of education reduces the probability you make an error by 0.3 percent and college graduates are 0.6 percent less likely to make an error. Taken together these estimates imply education does little to mitigate these errors. Comparing a college graduate (with 16 years of education) with a individual with a 6th grade education, the college graduate is only

¹⁵Here I use upper and lower cutoffs of .2 and -.2 to classify responses into about the same, more unemployment and less unemployment. Results were robust to other choices for the cutoffs.

3.6% less likely to make an error in his unemployment expectations. Similarly, income has only a small impact on the probability of making an expectational error. The coefficient on income is -.00016. Since income is measured in *thousands* of year 2000 dollars, an individual with an annual income that is 100,000 dollars greater is only 1.6% less likely to make an error¹⁶. Finally, an individual being optimistic about his own finances has a small effect on the probability of making an error as does the individual's race.

Second, still in table 5, we turn to the characteristics associated with making a pessimistic error, expecting unemployment to rise when the VAR predicts otherwise. The coefficient on being male has fallen in half. The age coefficients imply that the best and worst by age differ by about 4%. Being a college graduate reduces the probability of making a pessimistic error by 0.8 percent, and the difference in the probability of making a pessimistic error between a college graduate and an individual with a 6th grade education is only 4.8%. The effect of income on the probability of making a pessimistic error is the same as its effect on the probability of making a general error. Individuals optimistic about their own financial situation are 7% less likely to make a pessimistic error. Taken with the results on income, the results on education imply that these errors are not made solely by an economically unimportant fraction of the population but are made fairly evenly across the income and education distributions. Finally, Blacks, Hispanics and Native Americans are significantly more likely to make a pessimistic error relative to the omitted race group, Whites.

7.2 The Effect of Errors on Buying Attitudes

Next I examine the effect of unemployment expectations on buying attitudes. The survey asks the following three questions of the respondents. 1."Generally speaking, do you think now is a good time or a bad time to buy a house?" 2."Speaking now of the automobile market, do you think the next 12 months or so will be a good time or a bad time to buy a car?" and 3."About the big things people buy for their homes, such as furniture, a refrigerator stove,

¹⁶Using income in logs did not change the estimates of the income effect significantly.

television, and things like that. Generally speaking, do you think now is a good or a bad time for people to buy major household items?"

For each item I run the following two regressions:

$$BadBuy_{it} = F(\alpha + \beta Pessimist_{it} + \gamma X_{it} + \delta Z_t) + \varepsilon_{it} \quad (5)$$

$$BadBuy_{it} = F(\alpha + \beta PessError_{it} + \gamma X_{it} + \delta Z_t) + \varepsilon_{it} \quad (6)$$

where $BadBuy_{it}$ indicates the individual believes it to be a bad time to buy the item mentioned in the question above, $Pessimist_{it}$ equals one if the individual expects unemployment to rise and is zero otherwise, $PessError_{it}$ equals one if the individual expects unemployment to rise contrary to the VAR and is zero otherwise, X_{it} is a vector of individual characteristics as in the previous section and Z_t contains the lag unemployment rate and year fixed effects. The model is estimated by OLS ($F(x) = x$) and probit ($F =$ standard normal cdf). We would expect both β coefficients to be positive, being pessimistic about the future is correlated with thinking that it is a bad time to make a large purchase.

Table 6 shows the results of the regressions, suppressing the coefficients on the control variables. Both being pessimistic about future unemployment changes and making a pessimistic error has important effects on buying attitudes. Being pessimistic or making a pessimistic error results in being 8% less likely to think it is a good time to buy a house, 10% less likely to think it is a good time to buy a car, and 8% less likely to think it is a good time to make a major purchase of a durable good.

8 Conclusion

In this paper I have compared unemployment expectations of households, as measured by the Michigan Survey of Consumers, with the predictions of a four variable VAR containing GDP, the unemployment rate, the inflation rate, and the fed funds rate. Three important facts emerged. First, concerning the fall in unemployment at the end of the recession, there

are above average levels of pessimism with half to one-third of the population expecting unemployment to *rise* even though the VAR easily predicts the fall in unemployment. Second, concerning the rise in unemployment at the beginning of a recession, fewer people had expected unemployment to rise than at the end of a recession when unemployment is falling even though the VAR predicted these changes. Finally, when regressing the percent of the population that expects unemployment to rise on the VAR prediction of the future unemployment change and the lag unemployment change, I found the coefficient on the lag change was roughly the same magnitude as the coefficient on the VAR prediction. A model with a random expectations distribution around the VAR expectation underestimated the coefficient on the lag change by a factor of 15.

I then examined the ability of other models of expectations to match these facts. I showed that these puzzles were not due to the unavailability of revised data. Then I showed that while least squares learning could be useful for understanding why pessimism is low at the very start of the recession, it had an even harder time than the VAR explaining the pessimism at the end of the recession. Similarly, delayed updating of expectations helped explain why there are so few pessimists at the beginning of a recession but not why there are so many pessimists at the end of a recession. An extrapolative expectations model where agents partially form expectations by extrapolating current trends into the future, explained both insufficient pessimism at the beginning of the recession and excessive pessimism at the end of the recession. In its fifth section, the paper demonstrated that among data simulated from the different expectational models, only data from the extrapolation model could match the facts outlined previously and the sixth section showed that even a simple univariate forecast could forecast the sign of unemployment changes.

Finally the paper concluded by showing that while those with more education or a greater income are less likely to make expectational errors (defined as differing from the VAR) the difference is not substantial. The effect of income and education on the probability that an agent makes an error is almost negligible; therefore expectational errors are not confined to

an economically insignificant fraction of the population. The paper also showed that when an individual makes a pessimistic expectational error, that is expecting unemployment to rise when in fact the VAR predicts it will fall, they are 8%-10% less likely to think that it is a good time to make a major purchase relative to the rest of the population.

References

- Ball, Laurence. 2000. Near-Rationality and Inflation in Two Monetary Regimes. *The Johns Hopkins University, Department of Economics Economics Working Paper Archive*, Oct.
- Carlson, J.A., & Parkin, M. 1975. Inflation Expectations. *Economica*, **42**, 123–138.
- Carroll, Christopher D. 2003. Macroeconomic Expectations Of Households And Professional Forecasters. *The Quarterly Journal of Economics*, **118**(1), 269–298.
- Coibion, Olivier, & Gorodnichenko, Yuriy. 2008. What Can Survey Forecasts Tell Us About Information Rigidities? *NBER Working Paper 14586*, December.
- Croushore, Dean, & Stark, Tom. 1999. A real-time data set for macroeconomists. *Federal Reserve Bank of Philadelphia*. <http://www.phil.frb.org/econ/forecast/reaindex.html>.
- Curtin, Richard T. 2003. Unemployment Expectations: The Impact of Private Information on Income Uncertainty. *Review of Income and Wealth*, **49**(4), 539–554.
- Evans, George, & Honkapohja, Seppo. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Gabaix, Xavier, & Laibson, David. 2001. The 6D Bias and the Equity Premium Puzzle. *NBER Macroeconomics Annual*.
- Howrey, E. Phillip. 2001. The Predictive Power of the Index of Consumer Sentiment. *Brookings Papers of Economic Activity*, **2001**(1), 175–207.
- Mankiw, N. Gregory, & Reis, Ricardo. 2002. Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve. *The Quarterly Journal of Economics*, **117**(4), 1295–1328.
- Mankiw, N. Gregory, Reis, Ricardo, & Wolfers, Justin. 2003. Disagreement about Inflation Expectations. *National Bureau of Economic Research, Inc Working Paper*, June.
- Newey, Whitney K, & West, Kenneth D. 1987. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, **55**(3), 703–08.
- Orphanides, Athanasios. 2001. Monetary Policy Rules Based on Real-Time Data. *American Economic Review*, **91**(4), 964–985.
- Orphanides, Athanasios, & van Norden, Simon. 2002. The Unreliability Of Output-Gap Estimates In Real Time. *The Review of Economics and Statistics*, **84**(4), 569–583.
- Reis, Ricardo. 2004. Inattentive Consumers. *National Bureau of Economic Research Working Papers*, Nov.
- Sargent, Thomas. 2001. *The Conquest of American Inflation*. Princeton University Press.

Souleles, Nicholas S. 2004. Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys. *Journal of Money, Credit and Banking*, **36**(1), 39–72.

Table 1: Pessimism at the Beginning and After Recession

Months	Pessimist	RealUrateChange	ExpUrateChange	Pvalue Pessimist	Pvalue RealUrateChange	Pvalue ExpUrateChange
Before 0	0.336*** [0.054]	0.300*** [0.048]	0.404*** [0.149]	0.0218	<.001	0.004
1	0.324*** [0.033]	0.333*** [0.100]	0.404*** [0.149]	<.001	<.001	0.01
2	0.319*** [0.059]	0.467*** [0.073]	0.651*** [0.152]	0.0761	<.001	<.001
3	0.342*** [0.051]	0.767*** [0.139]	0.705*** [0.145]	0.23	<.001	<.001
4	0.352*** [0.053]	1.000** [0.393]	0.705*** [0.145]	0.8898	<.001	<.001
5	0.414*** [0.065]	1.100*** [0.347]	0.918*** [0.228]	0.4563	<.001	<.001
After 0	0.480*** [0.031]	-0.333*** [0.028]	-0.461* [0.256]			
1	0.440*** [0.010]	-0.300*** [0.048]	-0.461 [0.297]			
2	0.424*** [0.005]	-0.500*** [0.166]	-0.694*** [0.138]			
3	0.407*** [0.018]	-0.733* [0.373]	-0.754*** [0.097]			
4	0.360*** [0.005]	-1.167** [0.481]	-0.881*** [0.249]			
5	0.362*** [0.025]	-1.267** [0.518]	-1.181*** [0.330]			
Observations	186559	318	318			

The table contains the results of regressing, on a vector of variables indicating months since unemployment began to rise and months since it began to fall, if the individual expects unemployment to rise, the actual change in unemployment, and the VAR prediction of the change. The rightmost columns contain the p-values of the test that the before and after coefficients are equal (by month). Standard errors are corrected for within month correlation in the first column and autocorrelation in the others using Newey-West with 3 lags. * significant at 5%; ** significant at 1%

Table 2: Pessimism and Lag Unemployment Changes

	Pessimist	Pessimist
ExpectedUrateChange	0.057*** [0.013]	0.026 [0.017]
LagRealUrateChange	0.062*** [0.008]	0.058*** [0.008]
RealUrateChange	---	0.034*** [0.012]
Constant	0.348*** [0.009]	0.348*** [0.008]
Observations	314	314
R-squared	0.45	0.49

This table contains the results from regressing the percent of the population who expect unemployment to rise on a VAR prediction of the unemployment change and the lag unemployment change. The second column adds the actual unemployment change as a regressor. Standard errors are corrected for autocorrelation using a Newey-West procedure with 3 lags. * significant at 5%; ** significant at 1%

Table 3: Estimates of the Parameters of the Heterogenous Expectations Models

	VAR Model	LS Learning Model	Real Time Model	Delay Model	Extrapolation Model
Sigma	4.1 (0.05)	7.82 (0.12)	5.37 (0.09)	1.49 (0.05)	2.77 (0.03)
Upper Cutoff	1.72 (0.02)	2.96 (0.05)	2.4 (0.04)	0.55 (0.024)	1.14 (0.02)
Lower Cutoff	-4.32 (0.05)	-8.66 (0.13)	-5.52 (0.09)	-1.61 (0.04)	-3.03 (0.03)
Lambda	-- --	-- --	-- --	0.126 (0.006)	0.99 (0.08)
Percent Extrapolators	-- --	-- --	-- --	-- --	0.38 (0.007)
Function Value	12157	14551	15070	6508	9713

This table gives the estimates of the parameters of the expectational models. The distribution of expectations is assumed to be normally distributed with mean zero and variance sigma around the calculated expectation of the model, the individual is assigned to expect unemployment to rise if they expect a change above the upper cutoff, to expect unemployment to fall if they expect unemployment to fall more than the lower cutoff. lambda is the probability of updating expectations in the delay model and the extrapolation parameter in the extrapolation model. Finally the percent of extrapolators are the percent of individuals who form their expectation by extrapolating past unemployment rate changes. In parentheses are the square root of the appropriate diagonal entry of the matrix $[\partial (F1 \dots FN/\partial \beta)W(\partial (F1 \dots FN/\partial \beta))']^{-1}$. evaluated at the estimated parameters, where W is the (inverse variance) weighting matrix and (F1 ... FN) are the predictions of the model.

Table 4: Simulated Pessimism and Lag Unemployment Changes

	Data	VarPessimists	LSLPessimists	RealTimePessimists	Delay Pessimists	Extrap Pessimists
ExpectedUrateChange	0.057*** [0.013]	0.085*** [0.001]	0.060*** [0.005]	0.062*** [0.002]	0.100*** [0.002]	0.078*** [0.001]
LagRealUrateChange	0.062*** [0.008]	0.004*** [0.001]	0.003 [0.003]	-0.012*** [0.002]	0.026*** [0.003]	0.050*** [0.002]
Constant	0.348*** [0.009]	0.340*** [0.001]	0.337*** [0.003]	0.338*** [0.001]	0.348*** [0.002]	0.346*** [0.001]
Observations	314	314	314	314	314	314

This table repeats the first regression from table 2 but instead of using the actually number of pessimists it used the number simulated from the different expectational models.

Table 5: Predictor of Expectational Errors and Pessimist Errors

	Prob. Make Error		Prob. Make Pessimistic Error	
	OLS	Probit	OLS	Probit
Male	-0.023*** [0.002]	-0.025*** [0.002]	-0.012*** [0.002]	-0.012*** [0.002]
Age	-0.00182*** [0.00040]	-0.00180*** [0.00042]	0.00553*** [0.00040]	0.00561*** [0.00042]
Age Squared	0.000017*** [0.000004]	0.000017*** [0.000004]	-0.000057*** [0.000004]	-0.000058*** [0.000004]
Education (Highest Grade Attained)	-0.003*** [0.001]	-0.003*** [0.001]	-0.004*** [0.001]	-0.004*** [0.001]
College	-0.006* [0.004]	-0.007* [0.004]	0.008** [0.004]	0.008** [0.004]
Married	-0.00367 [0.00246]	-0.00376 [0.00259]	-0.00650*** [0.00249]	-0.00676*** [0.00254]
Income (in thousands of year 2000 Dollars)	-0.00016*** [0.0000]	-0.0002*** [0.0000]	-0.00016*** [0.0000]	-0.0002*** [0.0000]
Optimistic About Own Finances	-0.0048** [0.0024]	-0.0061** [0.0025]	-0.0691*** [0.0024]	-0.0698*** [0.0025]
Black	0.0076* [0.0041]	0.0081* [0.0043]	0.0891*** [0.0042]	0.0911*** [0.0047]
Hispanic	0.0131** [0.0058]	0.0121** [0.0061]	0.0193*** [0.0057]	0.0202*** [0.0060]
Native American	-0.0032 [0.0127]	-0.0033 [0.0134]	0.0495*** [0.0127]	0.0525*** [0.0141]
Asian	0.0160* [0.0095]	0.0144 [0.0099]	0.0084 [0.0093]	0.0079 [0.0098]
Lag Unemployment Rate			0.1016*** [0.0038]	0.1145*** [0.0040]
Constant	0.6745*** [0.0125]		-0.4548*** [0.0261]	
Observations	170781	170781	132388	132388
R-squared	0.07		0.062	

This table contains the results from regressing if the agent's response differs from the VAR (Error) and if the agent's response is "more unemployment" when the VAR predicts otherwise (Pessimistic Error) on individual characteristics. The first column, in each panel, gives OLS estimates the second column gives estimates of the marginal effects from a probit model. Standard errors in brackets. The pessimistic error regression contains fixed effects for years. Observations fall in these regressions because I exclude years in which it is impossible for there to be a pessimistic error because the VAR predicts unemployment will rise. *significant at 5%; ** significant at 1%


Table 6: Effects of Expectational Errors on Buying Attitudes

	Bad Time Buy Home		Bad Time Buy Car		Bad Time Buy Durables	
	OLS	Probit	OLS	Probit	OLS	Probit
Pessimist	0.082*** [0.002]	0.094*** [0.003]	0.096*** [0.002]	0.100*** [0.003]	0.082*** [0.002]	0.083*** [0.002]
Observations	166995	166995	160801	160801	161348	161348
R-squared	0.183		0.076		0.072	
Pessimistic Error	0.078*** [0.003]	0.097*** [0.003]	0.101*** [0.003]	0.110*** [0.003]	0.078*** [0.003]	0.083*** [0.003]
Observations	166995	166995	160801	160801	161348	161348
R-squared	0.18		0.073		0.069	

This table presents the results from regressing individual responses to questions about their buying attitudes on if they expect unemployment to rise (pessimist), if they expect it to rise when the VAR predicts otherwise (pessimistic error) and on individual characteristics. The first column, in each pair, gives OLS estimates and the second column gives the marginal effects from a probit model. All regression includes year fixed effects and all the control variables from the previous regressions in Table 13. Standard errors in brackets. * significant at 5%; ** significant at 1%.

A Supplementary Figures and Tables

Figure 1A: Assignment of Expectations



The letter t indexes months and the quarters are marked.

For months t and $t+1$ I assign the unemployment expectation $E_{Q0}[U_{Q4}] - U_{Q0}$

For month $t+2$ I assign the unemployment expectation $E_{Q0}[U_{Q5}] - U_{Q1}$

Appendix Table 1A: Summary Statistics

<u>Table 1a: Unemployment Expectations</u>	More Unemployment	About the Same	Less Unemployment	Does Not Know	Did not Answer
Number of Responses	65124	88913	29642	2223	657
Percent Responding	34.91	47.66	15.89	1.19	0.35

<u>Table 1b: Summary Statistics for Characteristics</u>	N	Mean	Standard Deviation	Percent Responding
Income (In Year 2000 Dollars)	174126	48510	4328	93.3
Age	184932	44.481	17.13	99.1
Highest Grade Completed (1 to 17)	183552	13.23	2.65	98.4
College Graduate (Yes = 1)	186559	0.31	0.461	100.0

<u>Table 1c: Percentiles for Income and Education</u>	Income	Highest Grade Completed
10	13,186	10
20	20,487	12
30	28,355	12
40	35,553	12
50	42,602	13
60	49,019	14
70	58,311	15
80	67,232	16
90	81,828	17

<u>Table 1d: Race</u>	Number	Percent
White	154,073	84.8
Black	16,006	8.8
Hispanic	7,579	4.2
Asian	1,453	0.8
Native American	2,640	1.5

<u>Table 1e: Responses to Buying Attitudes Questions</u>	Good	Pros and Cons	Bad	Doesn't Know	Did Not Answer
House	63.9	2.3	31.2	2.2	0.5
Car	59.4	3.3	30.6	5.8	0.8
Durable Good	67.6	5.0	21.3	5.1	1.0

This table gives summary statistics for the Survey of Consumers data. Table 1a tabulates responses to the 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?" and Table 1e tabulates responses to questions on whether or not they think it is a good time to buy a house, car or durable good.

Appendix Table 2A: Pessimism and Optimism After a Recession

Months After Unemp. Fall	Pessimist	Real Urate Change	VAR Urate Change	LSL Urate Change	Real Time Urate Change
0	0.480*** [0.031]	-0.333*** [0.028]	-0.461* [0.256]	-0.966*** [0.334]	-0.513*** [0.078]
1	0.440*** [0.010]	-0.300*** [0.048]	-0.461 [0.297]	-0.960* [0.494]	-0.571*** [0.176]
2	0.424*** [0.005]	-0.500*** [0.166]	-0.694*** [0.138]	-1.165*** [0.286]	-0.728*** [0.163]
3	0.407*** [0.018]	-0.733* [0.373]	-0.754*** [0.097]	-1.263*** [0.235]	-0.771*** [0.134]
4	0.360*** [0.005]	-1.167** [0.481]	-0.881*** [0.249]	-1.430*** [0.434]	-0.904*** [0.200]
5	0.362*** [0.025]	-1.267** [0.518]	-1.181*** [0.330]	-1.642*** [0.417]	-1.155*** [0.426]
6	0.332*** [0.019]	-1.300*** [0.458]	-1.167*** [0.103]	-1.440*** [0.210]	-1.125*** [0.225]
7	0.338*** [0.032]	-1.400*** [0.502]	-1.208*** [0.184]	-1.492*** [0.377]	-1.131*** [0.265]
8	0.292*** [0.051]	-1.233** [0.529]	-1.215*** [0.195]	-1.431*** [0.264]	-1.125*** [0.311]
9	0.240*** [0.050]	-1.333*** [0.486]	-1.307*** [0.135]	-1.555*** [0.197]	-1.155*** [0.236]
10	0.233*** [0.061]	-1.367** [0.563]	-1.350*** [0.102]	-1.636*** [0.187]	-1.087*** [0.271]
11	0.226*** [0.041]	-1.367** [0.640]	-1.381*** [0.157]	-1.640*** [0.145]	-1.111*** [0.309]
Observations	186559	318	318	318	318

The table contains the results of regressing, on a vector of variables indicating months since unemployment began to fall at the end of a recession, if the individual expects unemployment to rise, the actual change in unemployment, the VAR prediction, the least squares learning prediction, and the real time data prediction of the change. Standard errors are corrected for within month correlation in the first two columns and autocorrelation in the other columns using a Newey-West procedure with 3 lags. * significant at 5%; ** significant at 1%

Appendix Table 3A: Pessimism At Beginning of A Recession

Mos After Unemp. Rises	Pessimist	Real Urate Change	VAR Urate Change	LSL Urate Change	Real Time Urate Change
0	0.336*** [0.054]	0.300*** [0.048]	0.404*** [0.149]	0.173 [0.287]	0.841*** [0.076]
1	0.324*** [0.033]	0.333*** [0.100]	0.404*** [0.149]	0.173 [0.287]	0.841*** [0.076]
2	0.319*** [0.059]	0.467*** [0.073]	0.651*** [0.152]	0.39 [0.403]	1.071*** [0.095]
3	0.342*** [0.051]	0.767*** [0.139]	0.705*** [0.145]	0.459 [0.387]	0.854*** [0.107]
4	0.352*** [0.053]	1.000** [0.393]	0.705*** [0.145]	0.459 [0.387]	0.854*** [0.107]
5	0.414*** [0.065]	1.100*** [0.347]	0.918*** [0.228]	0.597 [0.519]	1.086*** [0.205]
6	0.477*** [0.101]	1.233*** [0.373]	0.806*** [0.228]	0.616 [0.520]	0.892*** [0.130]
7	0.483*** [0.088]	1.267*** [0.217]	0.806*** [0.228]	0.616 [0.520]	0.892*** [0.130]
8	0.425*** [0.091]	1.400*** [0.166]	0.928*** [0.247]	0.691 [0.622]	1.045*** [0.148]
9	0.450*** [0.075]	1.433*** [0.055]	0.871*** [0.245]	0.764 [0.649]	0.943*** [0.193]
10	0.454*** [0.069]	1.600*** [0.048]	0.871*** [0.245]	0.764 [0.649]	0.943*** [0.193]
11	0.480*** [0.072]	1.567*** [0.154]	1.066*** [0.271]	0.943 [0.747]	1.176*** [0.223]
Observations	186559	318	318	318	318

The table contains the results of regressing, on a vector of variables indicating months since unemployment began to rise at the beginning of a recession, if the individual expects unemployment to rise, the actual change in unemployment, the VAR prediction the least squares learning, and the real time data expectation of the change. Standard errors are corrected for within month correlation in the first column and autocorrelation in the other columns using a Newey-West procedure with 3 lags. * significant at 5%; ** significant at 1%

Appendix Table 4A: Why Delay Model Doesn't Match Pessimism at the End of the Recession

Lag/Months	0	1	2	3	4	5	6	7	8	9	10	11
0	-0.34	-0.37	-0.70	-0.77	-0.78	-1.00	-1.07	-1.15	-1.27	-1.27	-1.35	-1.38
1	-0.42	-0.39	-0.68	-0.66	-0.74	-1.01	-1.08	-1.00	-1.09	-1.22	-1.26	-1.37
2	-0.58	-0.60	-0.71	-0.59	-0.57	-0.78	-0.81	-0.94	-1.06	-1.04	-0.93	-1.00
3	-0.36	-0.21	-0.34	-0.65	-0.64	-0.67	-0.65	-0.64	-0.73	-0.77	-0.93	-0.99
4	-0.02	-0.23	-0.52	-0.37	-0.20	-0.31	-0.60	-0.59	-0.57	-0.58	-0.57	-0.65
5	0.06	-0.09	-0.19	-0.09	-0.24	-0.48	-0.34	-0.19	-0.27	-0.50	-0.48	-0.46
6	0.10	0.12	0.04	0.01	-0.11	-0.20	-0.13	-0.23	-0.41	-0.29	-0.16	-0.23
7	0.00	-0.07	-0.10	0.03	0.05	-0.01	-0.04	-0.12	-0.18	-0.15	-0.20	-0.33
8	-0.08	-0.08	-0.01	-0.02	-0.06	-0.09	-0.03	-0.01	-0.05	-0.07	-0.12	-0.16
9	0.00	-0.02	-0.05	-0.08	-0.08	-0.04	-0.04	-0.07	-0.09	-0.07	-0.05	-0.07
10	-0.01	-0.01	-0.04	-0.03	-0.04	-0.07	-0.08	-0.08	-0.06	-0.06	-0.07	-0.09
11	0.00	-0.01	-0.02	-0.03	-0.04	-0.06	-0.06	-0.07	-0.09	-0.09	-0.09	-0.08
12	-0.01	-0.01	-0.03	-0.03	-0.03	-0.04	-0.06	-0.06	-0.08	-0.08	-0.09	-0.10
13	-0.02	-0.02	-0.03	-0.02	-0.02	-0.04	-0.04	-0.04	-0.06	-0.07	-0.08	-0.09
14	-0.01	-0.01	-0.03	-0.03	-0.03	-0.04	-0.04	-0.04	-0.05	-0.06	-0.06	-0.07
15	-0.01	0.00	-0.01	-0.02	-0.02	-0.03	-0.03	-0.04	-0.05	-0.05	-0.05	-0.06
16	-0.02	-0.02	-0.02	-0.01	-0.01	-0.02	-0.02	-0.03	-0.04	-0.04	-0.04	-0.05
17	-0.02	-0.02	-0.01	-0.02	-0.02	-0.02	-0.02	-0.01	-0.02	-0.03	-0.03	-0.04
18	-0.02	-0.02	-0.02	-0.02	-0.02	-0.01	-0.02	-0.02	-0.02	-0.02	-0.02	-0.03
19	-0.03	-0.02	-0.03	-0.02	-0.02	-0.02	-0.02	-0.02	-0.01	-0.03	-0.03	-0.02

The average expected unemployment change at n months into the recession (top row) based on information k periods old (first column).

Appendix Table 5A: Simulated Pessimism and Optimism After a Recess

Mos After Fall	Pessimist	VarPess	LSLPess	RealTimePess	DelayPess	ExtrapPess
0	0.480*** [0.031]	0.309*** [0.017]	0.308*** [0.015]	0.294*** [0.005]	0.318*** [0.025]	0.371*** [0.050]
1	0.440*** [0.010]	0.306*** [0.020]	0.309*** [0.022]	0.290*** [0.011]	0.315*** [0.026]	0.363*** [0.049]
2	0.424*** [0.005]	0.278*** [0.011]	0.299*** [0.013]	0.280*** [0.010]	0.285*** [0.011]	0.348*** [0.030]
3	0.407*** [0.018]	0.272*** [0.007]	0.295*** [0.010]	0.277*** [0.008]	0.285*** [0.011]	0.341*** [0.033]
4	0.360*** [0.005]	0.272*** [0.013]	0.288*** [0.019]	0.269*** [0.012]	0.282*** [0.009]	0.338*** [0.035]
5	0.362*** [0.025]	0.255*** [0.017]	0.279*** [0.018]	0.256*** [0.025]	0.261*** [0.015]	0.320*** [0.010]
6	0.332*** [0.019]	0.249*** [0.014]	0.287*** [0.009]	0.256*** [0.013]	0.257*** [0.019]	0.309*** [0.013]
7	0.338*** [0.032]	0.243*** [0.012]	0.285*** [0.016]	0.256*** [0.016]	0.254*** [0.015]	0.300*** [0.013]
8	0.292*** [0.051]	0.234*** [0.017]	0.288*** [0.012]	0.257*** [0.018]	0.243*** [0.019]	0.289*** [0.011]
9	0.240*** [0.050]	0.234*** [0.012]	0.282*** [0.009]	0.254*** [0.014]	0.241*** [0.025]	0.277*** [0.004]
10	0.233*** [0.061]	0.227*** [0.007]	0.279*** [0.008]	0.259*** [0.016]	0.238*** [0.020]	0.264*** [0.004]
11	0.226*** [0.041]	0.226*** [0.011]	0.278*** [0.006]	0.258*** [0.018]	0.230*** [0.022]	0.252*** [0.002]
Observations	186559	314	314	314	314	314

This table repeats the analysis of table 2A, replacing the actual number of pessimists in the data with the number calculated from the simulated expectational models.

Appendix Table 6A: Simulated Pessimism At Beginning of A Recession

Mos After Unemp. Rises	Pessimist	VarPess	LSLPess	RealTimePess	DelayPess	ExtrapPess
0	0.336*** [0.054]	0.375*** [0.014]	0.361*** [0.014]	0.386*** [0.005]	0.401*** [0.005]	0.358*** [0.014]
1	0.324*** [0.033]	0.375*** [0.014]	0.361*** [0.014]	0.386*** [0.005]	0.401*** [0.005]	0.356*** [0.012]
2	0.319*** [0.059]	0.399*** [0.013]	0.372*** [0.020]	0.402*** [0.007]	0.413*** [0.010]	0.382*** [0.013]
3	0.342*** [0.051]	0.403*** [0.014]	0.375*** [0.019]	0.387*** [0.008]	0.429*** [0.013]	0.385*** [0.014]
4	0.352*** [0.053]	0.403*** [0.014]	0.375*** [0.019]	0.387*** [0.008]	0.429*** [0.013]	0.392*** [0.010]
5	0.414*** [0.065]	0.411*** [0.013]	0.382*** [0.025]	0.403*** [0.015]	0.432*** [0.014]	0.412*** [0.018]
6	0.477*** [0.101]	0.412*** [0.022]	0.383*** [0.026]	0.389*** [0.009]	0.449*** [0.019]	0.396*** [0.017]
7	0.483*** [0.088]	0.412*** [0.022]	0.383*** [0.026]	0.389*** [0.009]	0.449*** [0.019]	0.411*** [0.022]
8	0.425*** [0.091]	0.416*** [0.024]	0.387*** [0.031]	0.400*** [0.011]	0.443*** [0.021]	0.418*** [0.022]
9	0.450*** [0.075]	0.419*** [0.023]	0.391*** [0.032]	0.393*** [0.014]	0.460*** [0.026]	0.418*** [0.023]
10	0.454*** [0.069]	0.419*** [0.023]	0.391*** [0.032]	0.393*** [0.014]	0.460*** [0.026]	0.415*** [0.020]
11	0.480*** [0.072]	0.418*** [0.025]	0.400*** [0.037]	0.410*** [0.016]	0.447*** [0.027]	0.429*** [0.025]
Observations	186559	314	314	314	314	314

This table repeats the analysis of table 3A, replacing the actual number of pessimists in the data with the number calculated from the simulated expectational models.