

MACROECONOMIC EXPECTATIONS OF HOUSEHOLDS AND PROFESSIONAL FORECASTERS*

Christopher D. Carroll
ccarroll@jhu.edu

Published in the Quarterly Journal of Economics, Volume 118, Number 1, February 2003

December 17, 2002

ABSTRACT

Economists have long emphasized the importance of expectations in determining macroeconomic outcomes. Yet there has been almost no recent effort to model actual empirical expectations data; instead, macroeconomists usually simply assume that expectations are ‘rational.’ This paper shows that while empirical household expectations are not rational in the usual sense, expectational dynamics are well captured by a model in which households’ views derive from news reports of the views of professional forecasters, which in turn may be rational. The model’s estimates imply that people only occasionally pay attention to news reports; this inattention generates ‘stickyness’ in aggregate expectations, with important macroeconomic consequences.

Keywords: inflation, expectations, unemployment, monetary policy

JEL Classification Codes: D84, E31

This paper is derived from a paper originally titled “The Epidemiology of Macroeconomic Expectations” written in connection with the conference “The Economy As an Evolving Complex System III” at the Santa Fe Institute in November 2001, honoring Kenneth Arrow’s contributions to the Santa Fe Institute and to economics. A paper with the original title will be published in the associated eponymous conference volume. That companion paper examines a variety of alternative epidemiological models of expectations transmission that generate results similar to those of the baseline model presented here.

I am grateful to Jason Harburger, Jennifer Manning, Jirka Slacalek, and Johanna Francis for excellent research assistance, to Robert Axtell, William Branch, Carl Christ, William Dickens, Michael Dotsey, Joshua Epstein, Marvin Goodfriend, Daniel Leigh, Bennett McCallum, Yash Mehra, Serena Ng, Athanasios Orphanides, Adam Posen, John Roberts, Martin Sommer, and Alexander Wolman for valuable feedback, and to Richard Kwok for guidance to the relevant epidemiology literature. Thanks also to seminar audiences at Johns Hopkins University, Harvard University, the Center on Social and Economic Dynamics at the Brookings Institution, the Federal Reserve Bank of Richmond, New York University, and the University of Cyprus for valuable feedback.

The data and econometric programs that generated all of the results in this paper are available on the author’s website, <http://www.econ.jhu.edu/people/carroll>.

* Department of Economics, The Johns Hopkins University, Baltimore MD 21218-2685.

I. Introduction

Ever since the traditional foundation of macroeconomics by John Maynard Keynes [1936], economists have understood that macroeconomic outcomes depend upon expectations. Keynes himself believed that economies could experience fluctuations that reflected movements in ‘animal spirits,’ but the basis for most of today’s macro models was laid in the rational expectations revolution of the 1970s. Early critics of the rational expectations approach complained that, in the words of Friedman [1979], such models lacked “a clear outline of the way in which economic agents derive the knowledge which they then use to formulate expectations,” but recent criticisms have focused on the difficulty rational expectations models have in reproducing various features of macroeconomic data like the high persistence of inflation (Fuhrer and Moore [1995]) and the apparent inexorability of the tradeoff between inflation and unemployment (Ball [1994]; Mankiw [2001]). The literature has consequently begun to explore the implications for macroeconomic dynamics of various alternative assumptions about expectations formation, most notably models of learning, see Sargent [1993] or Evans and Honkapohja [2001] for surveys.

Remarkably, however, there has been almost no work testing alternative models of expectations using actual empirical data on expectations. McCallum’s [2002] recent survey, for example, does not discuss results from a single paper that examines empirical expectations data. This is not for lack of data: The University of Michigan’s Survey Research Center has been collecting information on households’ expectations about inflation, unemployment, economic growth, interest rates, and other macroeconomic matters for almost 50 years, the Conference Board has conducted similar monthly surveys of households since the late 1970s, and the Survey of Professional Forecasters and its antecedents have collected data since the 1960s. While there has been some work testing (and usually rejecting) the rationality of these expectations,¹ aside from an insightful paper by Roberts [1998] and an impressive (and very recent) paper by Branch [2001] there appears to have been essentially no work proposing and testing positive alternative models for how empirical expectations are formed.²

This paper proposes and tests one such model. Rather than having full understanding of the ‘true’ macroeconomic model and constantly tracking the latest statistics to produce their own macroeconomic forecasts, typical people are assumed to obtain their macroeconomic views from the news media. Furthermore (and importantly), not every person pays close attention to all macroeconomic news; instead, individual people are assumed to absorb the economic content of news stories probabilistically, so that it may take quite some time for news of changed macroe-

¹See Croushore [1998], Thomas [1999], and Mehra [2002] for surveys.

²The only other even tangentially relevant papers I have found were by Fische and Idson [1990] who test a model of heterogeneous demand for information using two years’ worth of Michigan Survey data; a paper by Urich and Wachtel [1984] that tests rationality using survey data on money supply forecasts; and a paper by Dua and Ray [1992] that models SPF data using an ARIMA forecasting framework.

conomic circumstances to penetrate to all agents in the economy. Finally, the news media in turn are assumed to report the views of professional forecasters, who may themselves make rational forecasts. (The structure of the model was inspired by simple models of disease from the epidemiology literature; see the companion paper Carroll [forthcoming]) for more on the epidemiological foundations of the model and for a demonstration that more elaborate “epidemiological expectations” models generate results similar to those of the baseline model presented here.)

The baseline model provides a simple equation for the evolution of mean expectations that is very similar to an equation proposed in recent papers by Mankiw and Reis [2001, 2003]. Indeed, the model presented here could be viewed as providing microfoundations for the Mankiw and Reis equation. Another contribution is the derivation of a particularly simple specialization of that equation suitable for empirical work; this specialization turns out to yield an equation like one estimated by Roberts [1998], for which it can again be regarded as providing a microfoundation. Finally, the model’s explicit assumption that people derive their expectations from news reports (and the paper’s specific proposal for how to measure news coverage) respond to Friedman’s [1979] criticism of the unspecified nature of the expectations formation mechanism in rational expectations models.

The model is applied to estimate the evolution of inflation expectations and unemployment expectations from the Michigan Survey of Consumers. For inflation, the typical household is estimated to update expectations roughly once a year, while unemployment expectations appear to be updated slightly more frequently. Furthermore, in a horserace between a version of the model where people update their expectations either to the rational forward-looking forecast or to the most recently reported past statistics (the ‘adaptive expectations’ model), the data strongly prefer the forward-looking version of the model. Thus, the results can be interpreted as reflecting a plausible middle ground between fully rational expectations and adaptive expectations.

A final section briefly comments on the relationship between this model and some of the relevant existing empirical literature, with particular emphasis on the relationship of the model to sticky-price models and the model’s implications for the relationship between credibility and monetary policy. The implications of the model for macroeconomic dynamics are not addressed here, because the papers by Mankiw and Reis [2001, 2003] and Roberts [1995, 1997] address those questions and are directly applicable. Mankiw and Reis show that their model can explain many phenomena that are unexplained by fully rational models, including why disinflations are inevitably contractionary; why monetary policy affects the economy with considerable lags; why rapid economic growth leads to rising inflation; and why productivity slowdowns are associated with a rise in the natural rate of unemployment. The ability to solve all of these puzzles seems a large dividend in exchange for the small price of relaxing the assumption that all agents’ expectations are fully rational (in the sense required by typical rational expectations models) at every instant.

II. The Model

Consider a world where most people form their expectations about future inflation by reading newspaper articles. Imagine for the moment that every inflation article contains a complete forecast of the inflation rate for all future quarters, and suppose (again momentarily) that any person who reads such an article can subsequently recall the entire forecast.

Assume that not everybody reads every newspaper article on inflation. Instead, in any given period each individual faces a constant probability λ of encountering and absorbing the contents of an article on inflation. Individuals who do not encounter an article simply continue to believe the last forecast they read about. Thus, the framework is mathematically similar to the Calvo [1983] model of sticky prices in which firms change their prices with probability p .

Call π_{t+1} the inflation rate between quarter t and quarter $t + 1$,

$$\pi_{t+1} = \log(p_{t+1}) - \log(p_t),$$

where p_t is the aggregate price index in period t . If we define M_t as the operator that yields the population-mean value of inflation expectations at time t and denote the Newspaper forecast printed in quarter t for inflation in quarter $s \geq t$ as $N_t[\pi_s]$, we have that³

$$(1)M_t[\pi_{t+1}] = \lambda N_t[\pi_{t+1}] + (1 - \lambda) \{ \lambda N_{t-1}[\pi_{t+1}] + (1 - \lambda) (\lambda N_{t-2}[\pi_{t+1}] + \dots) \}.$$

The derivation of this equation is as follows. In period t a fraction λ of the population will have absorbed the current-period newspaper forecast for the next quarter, $N_t[\pi_{t+1}]$. Fraction $(1 - \lambda)$ of the population retains the views they held in period $t - 1$ of period $t + 1$'s inflation rate. Those period- $t - 1$ views in turn can be decomposed into a fraction λ of people who encountered an article in period $t - 1$ and obtained the newspaper forecast of period $t + 1$'s forecast, $N_{t-1}[\pi_{t+1}]$, and a fraction $(1 - \lambda)$ who retained their period- $t - 2$ views about the inflation forecast in period $t + 1$. Recursion leads to the remainder of the equation.

This expression for inflation expectations is identical to the one proposed by Mankiw and Reis [2001, 2003], except that in their framework updating agents compute their own forecasts under the usual assumptions of rational expectations. Mankiw and Reis motivate their assumption that forecasts are updated only occasionally by arguing that there may be costs to obtaining or processing information. It is undoubtedly true that developing a reasonably rational quarter-by-quarter forecast of inflation arbitrarily far into the future would be a very costly enterprise for a typical person (for example, it might require obtaining an economics Ph.D. first!). But Mankiw and Reis do not provide any formal model of information processing

³Here we are assuming that all newspapers report the same forecast for inflation; see Carroll [forthcoming] for a version that allows for the possibility that different newspapers print different forecasts; that paper shows that results in such a model are similar to those presented here.

costs that leads to their specification, and indeed it seems likely that a formal model of processing costs might imply an updating process quite different from the Poisson process Mankiw and Reis assume.⁴

The model proposed above can be regarded as a microfoundation for the Mankiw and Reis equation (1). Its value as a microfoundation is illustrated in the usual way: it provides additional testable implications that do not follow directly from the aggregate specification. In particular, the baseline model implies that in periods when there are more news stories on inflation, the speed of updating should be faster, an implication that is borne out in the empirical work below. It also provides implications for the analysis of the underlying micro data from the Michigan survey. In particular, Souleles [forthcoming] finds highly statistically significant differences across demographic groups in macroeconomic forecasts; the model suggests examining whether those differences can be explained by information on demographic differences in readership rates of newspapers, or more general data on differences in the extent to which different groups pay attention to economic matters.

Of course, real newspaper articles do not contain a quarter-by-quarter forecast of the inflation rate into the infinite future as assumed in the derivation of (1), and even if they did it is very unlikely that a typical person would be able to remember the detailed pattern of inflation rates far into the future. Furthermore, even if both of these assumptions were true, equation (1) would not be testable in its current form because the available survey data report only households' expectations about inflation rates over the next year. In order to derive implications from the model that are testable with these data, it turns out to be necessary to impose some structure on households' implicit views about the inflation process.

Suppose people believe that at any given time the economy has an underlying "fundamental" inflation rate. Furthermore, suppose they believe that future changes in the fundamental rate are unforecastable; that is, beyond the next period the fundamental rate follows a random walk. Finally, assume that people believe that the actual inflation rate in a given quarter is equal to that period's fundamental rate plus an error term ϵ_t which reflects unforecastable transitory inflation shocks (reflected in the 'special factors' that newspaper inflation stories often emphasize). Thus, the typical person believes that the inflation process is captured by

$$\begin{aligned}
 (2) \quad & \pi_t = \pi_t^f + \epsilon_t \\
 (3) \quad & \pi_{t+1}^f = \pi_t^f + \eta_{t+1}, \\
 & \vdots \qquad \qquad \qquad \vdots
 \end{aligned}$$

where ϵ_t is a transitory shock to the inflation rate in period t while η_t is the permanent innovation in the fundamental inflation rate π_t^f in period t . Now assume that consumers believe that values of η beyond period $t + 1$, and values of ϵ beyond period t , are unforecastable white noise variables; that is, future changes in the

⁴See Sims [2001] for a model grounded in information theory that provides a formal model of decisionmaking under information-processing constraints.

fundamental inflation rate are unforecastable, and transitory shocks are expected to go away.⁵

Before proceeding it is worth considering whether this is a plausible view of the inflation process; we would not want to assume that people believe something patently absurd. However, the near-unit-root feature of the inflation rate in the post-1959 period is well known to inflation researchers; some authors find that a unit root can be rejected for some measures of inflation over some time periods, but it seems fair to say that the conventional wisdom is that at least since the late 1950s inflation is ‘close’ to a unit root process. See Barsky [1987] for a more complete analysis, or Ball [2000] for a more recent treatment.

Note that the unit root (or near unit root) in inflation does not imply that future inflation rates are totally unpredictable, only that the history of inflation by itself is not very useful in forecasting future inflation *changes* (beyond the disappearance of the transitory component of the current period’s shock). This does not exclude the possibility that current and lagged values of other variables might have predictive power. Thus, this view of the inflation rate is not necessarily in conflict with the vast and venerable literature showing that other variables (most notably the unemployment rate) do have considerable predictive power for the inflation rate (see Staiger, Stock, and Watson [2001] for a recent treatment).

If we were to assume that households were rational and made their own inflation forecasts solely based on observed past inflation under the assumption of an inflation process like (2)-(3), then the rational forecast would be a geometrically declining weighted average of past inflation realizations; in this case rational expectations would be identical to adaptive expectations (Muth [1960]). However, we will assume that households believe that *experts* have some ability to directly estimate the past and present values of π through period t and η through period $t + 1$ (through deeper knowledge of how the economy works, or perhaps some private information); thus households can rationally believe that a forecast from a professional forecaster is more accurate than a simple adaptively rational forecast that they could construct themselves.

Suppose now that rather than containing a forecast for the entire quarter-by-quarter future history of the inflation rate, newspaper articles simply contain a forecast of the inflation rate over the next year. The next step is to figure out how such a one-year forecast for inflation can be integrated into some modified version of equation (1). To capture this, we must introduce a bit more notation. Define $\pi_{s,t}$ as the inflation rate between periods s and t , converted to an annual rate. Thus, for example, in quarterly data we can define the inflation rate for quarter $t + 1$ at an annual rate as

$$\begin{aligned}\pi_{t,t+1} &= 4(\log p_{t+1} - \log p_t) \\ &= 4\pi_{t+1},\end{aligned}$$

⁵Note that we are allowing people to have some idea about how next quarter’s fundamental rate may differ from the current quarter’s fundamental rate, because we did not impose that consumers’ expectations of η_{t+1} must equal zero.

where the factor of four is required to convert the quarterly price change to an annual rate.

Our hypothetical person's view is that the true *ex-post* inflation rate over the next year will be given by

$$\begin{aligned}
(4) \quad \pi_{t,t+4} &= \pi_{t+1} + \pi_{t+2} + \pi_{t+3} + \pi_{t+4} \\
&= \pi_{t+1}^f + \pi_{t+1} + \pi_{t+2}^f + \pi_{t+2} + \pi_{t+3}^f + \pi_{t+3} + \pi_{t+4}^f + \pi_{t+4} \\
&= \pi_{t+1}^f + \pi_{t+1} + \pi_{t+1}^f + \eta_{t+2} + \pi_{t+2} + \pi_{t+2}^f + \eta_{t+2} + \eta_{t+3} + \pi_{t+3} + \\
&\quad \pi_{t+3}^f + \eta_{t+2} + \eta_{t+3} + \eta_{t+4} + \pi_{t+4}.
\end{aligned}$$

Define $F_t[\bullet_s]$ as the agent's forecast (expectation) as of date t of \bullet_s , for an agent who updates his views from a news report in period t . Using this notation, the assumptions made earlier about the stochastic processes for π and η imply that

$$(5) \quad F_t[\pi_{t+n}] = F_t[\eta_{t+n+1}] = 0$$

for all $n > 0$. Applying the F_t operator to both sides of (4) reveals that the person's forecast of the inflation rate over the next year is simply equal to four times his forecast of the fundamental inflation rate for next quarter:

$$\begin{aligned}
F_t[\pi_{t,t+4}] &= 4F_t[\pi_{t+1}^f] \\
&= F_t[\pi_{t,t+1}^f].
\end{aligned}$$

Now for an important conclusion: If people believe that the forecasts printed in the newspaper embody the same view of the inflation process laid out in equations (2)-(3) and (5), then an identical analysis leads to the conclusion that (defining the 'newspaper expectations' operator N_t similarly to the consumer's expectations operator):

$$\begin{aligned}
N_t[\pi_{t,t+4}] &= 4N_t[\pi_{t+1}^f] \\
&= N_t[\pi_{t,t+1}^f].
\end{aligned}$$

Thus, from the consumer's point of view the newspaper forecast contains only a single important piece of information: a projection of the fundamental inflation rate over the next year, which the process (2)-(3) implies is the expected fundamental rate in all of the year's constituent quarters and all subsequent quarters as well. A consumer who reads the newspaper in period t , therefore, will update his expectations to equal the corresponding newspaper forecasts:

$$F_t[\pi_{t,t+1}] = F_t[\pi_{t,t+4}] = F_t[\pi_{t,t+4}^f] = N_t[\pi_{t,t+4}^f] = N_t[\pi_{t,t+4}].$$

The rightmost equality holds because the consumer assumes that for $n > 0$, newspaper has no information about π_{t+n} or η_{t+n+1} , so $N_t[\pi_{t+n}] = N_t[\eta_{t+n+1}] = 0$. The next equality to the left holds because we assume that when the consumer

reads the newspaper his views are updated to the views printed in the newspaper. The other two equalities similarly hold because $F_t[\pi_{t+n}] = F_t[\eta_{t+n+1}] = 0$.

Now note a crucial point: the assumption that changes in the inflation rate beyond period $t + 1$ are unforecastable means that

$$(6) \quad F_{t-1}[\pi_{t-1,t+3}] = F_{t-1}[\pi_{t,t+4}]$$

$$(7) \quad F_{t-2}[\pi_{t-2,t+2}] = F_{t-2}[\pi_{t,t+4}]$$

...

An equation similar to (1) can be written for projections of the inflation rate over the next year:

$$M_t[\pi_{t,t+4}] = \lambda F_t[\pi_{t,t+4}] + (1 - \lambda) \{ \lambda F_{t-1}[\pi_{t,t+4}] + (1 - \lambda) (\lambda F_{t-2}[\pi_{t,t+4}] + \dots) \},$$

and substituting (6)-(7) into this equation and replacing F_t with N_t on the assumption that the newspaper forecasts are the source of updating information, we obtain

$$\begin{aligned} M_t[\pi_{t,t+4}] &= \lambda F_t[\pi_{t,t+4}] + (1 - \lambda) \{ \lambda F_{t-1}[\pi_{t-1,t+3}] + (1 - \lambda) (\dots) \} \\ M_t[\pi_{t,t+4}] &= \lambda F_t[\pi_{t,t+4}] + (1 - \lambda) M_{t-1}[\pi_{t-1,t+3}] \\ (8) \quad M_t[\pi_{t,t+4}] &= \lambda N_t[\pi_{t,t+4}] + (1 - \lambda) M_{t-1}[\pi_{t-1,t+3}]. \end{aligned}$$

That is, mean measured inflation expectations for the next year should be a weighted average between the current ‘rational’ (or newspaper) forecast and last period’s mean measured inflation expectations. This equation is therefore directly estimable, assuming an appropriate proxy for newspaper expectations can be found.⁶

Readers uncomfortable with the strong assumptions needed to derive (8) may be happier upon noting that the equation

$$(9) \quad M_t[\pi_{t,t+4}] = \lambda N_t[\pi_{t,t+4}] + (1 - \lambda) M_{t-1}[\pi_{t,t+4}]$$

can be derived without any assumptions on consumers’ beliefs about the inflation process; the difference between (8) and (9) is only in the subscript on the π term inside the M_{t-1} operator. The assumptions made above were those necessary to rigorously obtain $M_{t-1}[\pi_{t,t+4}] = M_{t-1}[\pi_{t-1,t+3}]$. In practice, however, even a much more realistic view of the inflation process would likely imply a very high degree of correlation between the period- $t - 1$ projection of the inflation rate over the year beginning in quarter t and the period- $t - 1$ projection of the inflation rate over the year beginning in quarter $t + 1$. Indeed, three of out of the four quarters ($t + 2, t + 3$, and $t + 4$) are identical between the two projections; the only differences between the two measures would have to spring from the consumer’s period- $t - 1$ projection of the difference between the inflation rates in quarters $t + 1$ and $t + 5$.

⁶This equation is basically the same as equation (5) in Roberts [1998], except that Roberts proposes that the forecast toward which household expectations are moving is the ‘mathematically rational’ forecast (and he simply proposes the equation without constructing a microfoundation that might produce it).

III. Estimation

Estimating equation (8) requires us to identify data sources for population-mean inflation expectations and for ‘newspaper’ forecasts of inflation over the next year.

The University of Michigan’s Survey Research Center conducts a monthly survey of households that is intended to be representative of the population of the United States. One component of the survey asks households what they expect the inflation rate to be over the next year.⁷ I will directly use the mean inflation forecast from this survey as my proxy for $M_t[\pi_{t,t+4}]$.

Identifying the ‘newspaper’ forecast for next-quarter inflation might seem more problematic, but there is a surprisingly good candidate: The mean four-quarter inflation forecast from the Survey of Professional Forecasters (henceforth, SPF). The SPF, currently conducted by the Federal Reserve Bank of Philadelphia and previously a joint product of the National Bureau of Economic Research and the American Statistical Association, has collected and summarized forecasts from leading private forecasting firms since 1968. The survey questionnaire is distributed once a quarter, just after the middle of the second month of the quarter, and responses are due within a couple of weeks. The survey asks participants for quarter-by-quarter forecasts, spanning the current and next five quarters, for a wide variety of economic variables, including GDP growth, various measures of inflation including CPI inflation, and the unemployment rate. For more details on the SPF, see Croushore [1993].

As noted above, the typical newspaper article on inflation interviews some ‘experts’ on inflation. The obvious candidates for such experts are the set of people who forecast the economy for a living, so the pool of interviewees is likely to be approximately the same group of forecasters whose views are summarized by the SPF. Hence, it seems reasonable to identify N_t with the SPF inflation expectations data.

III.1. Do the Forecasts Forecast?

There is a substantial existing literature on the forecasting performance of various measures of inflation expectations including the Michigan Survey, the SPF, and an informal survey of economists known as the Livingston survey.⁸ Early papers (Turnovsky [1970], Bryan and Gavin [1986]) claimed to find statistically significant biases in some of the survey measures, but a recent review by Croushore [1998] shows

⁷Specifically, households are asked whether they think prices will go up, stay the same, or fall over the next year. Those who say ‘go up’ (the vast majority) are then asked ‘By about what percent do you expect prices to go up, on the average, during the next 12 months?’ For more details on the survey methodology, see Curtin [1996].

⁸Unfortunately, the recent survey paper by Thomas [1999] largely neglects the SPF, and focuses instead mainly on comparisons of the Michigan survey and the Livingston survey. Thomas finds the median of the Michigan survey to be a better forecaster than its mean, but my model delivers predictions only for the mean and not for the median, so I neglect the median forecast in my empirical work.

that some of those results were spurious (due to improper treatment of the data or econometric problems), and that the results claiming to reject rationality of the SPF fail to hold up when the sample period is updated to include data for the last ten or fifteen years. Croushore specifically examines the CPI forecasts of both the Michigan survey and the SPF, and in a ‘forecast improvement’ exercise finds evidence of systematic bias in the Michigan survey but not in the SPF. Roberts [1997] also finds that the Michigan survey’s inflation expectations measure fails standard rationality tests.

These results are suggestive, but are not precisely targeted on the question we are interested in: Whether the SPF forecast can be viewed as ‘more rational’ than the Michigan forecast, and whether there is evidence that information moves from the SPF forecasters to the Michigan households but not vice versa.

One of the simplest measures of forecast accuracy is the mean squared error. It is reassuring therefore that over the time period for which both SPF and Michigan forecasts are available, the *ex post* MSE of the SPF forecast is about 0.6 while the MSE for the Michigan survey is almost twice as large, about 1.1. (These are calculated by taking the square of the difference between the respective mean forecasts and the actual CPI core inflation over the corresponding time period.)

A natural next question is whether each of the two surveys has meaningful forecasting power for future inflation, and if so, whether the SPF forecast is better. As a first step, consider the implications of the near unit root in inflation. High serial correlation means that future levels of the inflation rate will be highly predictable based on the recent past history of inflation. Hence it is not very impressive to find that both surveys have highly significant predictive power for inflation (which they do), since this result could hold even if the forecasts were both mindless extrapolations of past inflation into the future. The interesting question is whether the survey forecasts have predictive power for the future inflation rate *beyond* what could be predicted based on past inflation data.

To answer this question, Table I presents a regression of the actual inflation rate over the next year on the Michigan and SPF measures of expected inflation, along with the most recent annual inflation statistic available at the time the SPF and Michigan forecasts were made. Both survey measures have highly statistically significant predictive power for future inflation even controlling for the inflation rate’s recent past history, but the SPF measure has substantially more predictive power. The ‘horserace’ regression results indicate that the Michigan survey measure contains no information that is not also included in the SPF measure, while the SPF forecast has highly statistically significant predictive power that is not contained in the Michigan survey.⁹ Note that this result implies that the Michigan forecast is *prima facie* irrational (using the usual definition in rational expectations mod-

⁹A more stringent test would be whether the surveys can predict the *change* in the inflation rate. See Carroll [2001] for this test, which again finds that both surveys have highly significant predictive power but the SPF has more power. A more extensive evaluation of the forecasting power of the indexes is provided in the archive of programs that generated all of the results in this paper, available on the author’s website.

els), since the information that forecasters possessed that allowed them to make a superior forecast was in principle also available to households. Thus, we can unambiguously conclude that the SPF forecast is ‘more rational’ than the Michigan forecast, and the difference is large in both statistical and economic terms.

A final preliminary check is suggested by the structure of the model, in which expectations are assumed to spread from forecasters to households. This suggests that the professional forecasts should Granger-cause the household forecasts, but not vice versa. Table II shows that there is indeed statistical evidence of Granger causality from the professional forecasts to household forecasts, but no Granger causality in the opposite direction.

Of course, a finding that the SPF forecast is better than the Michigan forecast does not necessarily imply that the SPF forecast is fully rational. However, Croushore [1998] reports results for a battery of optimality tests proposed in the *Handbook of Statistics* by Diebold and Lopez [1996]; the SPF fails only one of these tests, the DuFour test, which is actually partly a test of the symmetry of the forecast errors around zero. Since nothing in rational expectations theory requires errors to be symmetrically distributed, this test is arguably of less interest than the other tests. Finally, note that the question of the rationality of the SPF forecasts is logically separate from the enterprise here, which is to examine whether the Michigan forecasts can be well modeled as updating toward the SPF forecasts. Rationality of the SPF forecasts is interesting in and of itself, but is in principle an independent question that can be addressed separately (as in Croushore [1998]).

III.2. Estimating the Stickiness of Inflation Expectations

We now turn to the main question, which is whether the Michigan survey data can be reasonably well represented by the model (8).

To provide a baseline for comparison, the first line of Table III presents results for the simplest possible model: that the value of the Michigan index $M_t[\pi_{t,t+4}]$ is equal to a constant, α_0 . By definition the \bar{R}^2 is equal to zero; the standard error of the estimate is 0.88. The last column of the table is reserved for reporting the results of various tests that will be conducted as the analysis progresses. By way of example, the test performed for the benchmark expectations-constant model is whether the average value of the expectations index is zero, $\alpha_0 = 0$. Unsurprisingly, this nonsensical proposition can be rejected with an overwhelming degree of statistical confidence, as indicated by a p -value that says that the probability that the proposition is true is zero.

We begin to examine the baseline model’s ability to explain the Michigan data by estimating

$$(10) \quad M_t[\pi_{t,t+4}] = \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_t[\pi_{t-1,t+3}] + \epsilon_t,$$

where $S_t[\pi_{t,t+4}]$ is the corresponding SPF forecast. Comparing this to (8) provides

the testable restriction that $\alpha_2 = 1 - \alpha_1$ or, equivalently,

$$(11) \quad \alpha_1 + \alpha_2 = 1.$$

Results from the estimation of (10) are presented as equation 1. The point estimates of $\alpha_1 = 0.36$ and $\alpha_2 = 0.66$ suggest that the restriction (11) is very close to holding true, and the last column presents formal statistical evidence on the question: It shows that the statistical significance with which the proposition that $\alpha_1 + \alpha_2 = 1$ can be rejected is only about 0.18, so that the restriction is easily accommodated by the data at the conventional level of significance of 0.05 or greater. Estimation results when the restriction is imposed in estimation are presented in the next row of the table, which provides our first unambiguous estimate of the crucial coefficient: $\lambda = 0.27$. Note that the Durbin-Watson statistic indicates that there is no evidence of serial correlation in the residuals (a Q-test yields the same result), which is impressive because the individual series involved have very high degrees of serial correlation. This is evidence that the two variables are cointegrated, as would be expected if one were a distributed lag of the other.

The point estimate $\lambda = 0.27$ is remarkably close to the value of 0.25 assumed by Mankiw and Reis [2001, 2003] in their simulation experiments; unsurprisingly, the last column for equation 2 indicates that the proposition $\alpha_1 = \lambda = 0.25$ is easily accepted by the data. This estimate indicates that in each quarter, only about one fourth of households have a completely up-to-date forecast of the inflation rate over the coming year. On the other hand, it also indicates that only about 32 percent ($= (1 - 0.25)^4$) of households have inflation expectations that are more than a year out of date.

As noted above, Roberts [1998] estimated a similar equation, except that his proposal was that expectations move toward the mathematically rational forecast of inflation rather than toward the SPF measure. Since the ‘rational’ forecast is unobservable, he used the actual inflation rate and instrumented using a set of predetermined instruments, on the usual view that if the instruments are valid the estimation should yield an unbiased estimate of the coefficient on the true but unobservable rational forecast. However, this procedure is problematic if there was anything that the ‘rational’ forecaster did not know about the structure of the economy and had to learn from realizations over time; as Roberts acknowledges, it is also problematic if the structure of the economy changes over time. A final drawback to this approach is that instrumenting can cause a severe loss of efficiency. Since the theory proposed here is quite literally that household expectations move toward the SPF forecast, there is no reason to instrument. In the end, however, Roberts’s parameter estimates are similar to those obtained here, though with considerably larger standard errors.

What equation 2 of Table III indicates is that if the data are forced to choose an $\alpha_1 = 1 - \alpha_2$ they are happy with that restriction, and that a model that imposes the restriction has a highly statistically significant ability to fit the data. However, we have not allowed the data to speak to the question of whether there is a better representation of inflation expectations than (10).

The first avenue by which we might wish to let the data reject the specification is to allow a constant into the equation. Equation 3 presents the results. The last column indicates that the proposition that the constant term is zero can be rejected at a very high level of statistical significance; on the other hand, the improvement in fit that comes with a constant is rather modest: The standard error declines from about 0.43 without the constant to about 0.35 with it. Compared with a raw standard error for the dependent variable of about 0.88, this improvement in fit is not very impressive, even if it is statistically significant.

Furthermore, if the model is to be treated as a structural description of the true process by which inflation expectations are formed, the presence of a constant term does not make much sense. It implies, for example, that if both actual inflation and the rational forecast for inflation were to go to zero forever, people would continue to expect a positive inflation rate (of a bit under 2 percent) forever. It seems much more likely that under these circumstances people would eventually learn to expect an inflation rate of zero. This point can be generalized to show that if the actual inflation rate and the rational forecast were fixed forever at *any* constant value, people's expectations would never converge to the true, constant inflation rate, but instead would be perpetually biased (unless the true value happened to be exactly equal to the unique stable point of the estimated equation).

A more palatable explanation for the presence of a constant term is that the baseline model is not a perfectly accurate description of the process by which information is transmitted in the economy; in this case estimation of the misspecified model could result in a spuriously significant coefficient term. For example, Carroll [forthcoming] demonstrates that the presence of a significant constant term could reflect the presence of some social transmission of inflation expectations via conversations with neighbors (epidemiologically, the disease is locally communicable), in addition to the news-media channel examined here.

Another plausible modification to the model is to allow for the possibility that some people update their expectations to the most recent *past* inflation rate rather than to the SPF forecast of the future inflation rate. Since most news coverage of inflation is prompted by the release of past inflation statistics (and since the new past number is often in the headline of the news article) one might argue that it might seem *more* likely for people to update their expectations to the past inflation rate than to a forecast of the future rate. This corresponds to what is usually called a model of 'adaptive expectations.' As noted above, however, if people believe that the true inflation process is as described in (2)-(3), this adaptive expectations benchmark is also identical to the limited-information rational expectations forecast (again, remember that we are assuming that households believe professional forecasters know much more about the inflation process than is contained in its past history, so updating households could still believe that the SPF forecast is better than the adaptively rational forecast would be).

We can examine these possibilities by estimating an equation of the form

$$(12) \quad M_t[\pi_{t,t+4}] = \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_{t-1}[\pi_{t-1,t+3}] + \alpha_3 P_t[\pi_{t-5,t-1}],$$

where $P_t[\pi_{t-5,t-1}]$ represents the most recently published annual inflation rate as of time t .

Results from estimating this equation are presented in the next row of Table III. The past inflation rate is indeed highly statistically significant - but with a *negative* coefficient! The negative coefficient makes no sense, as it implies that a higher past inflation rate convinces people that the fundamental inflation rate is lower. The final row of the table, however, shows that when a constant is included in this regression, the past inflation rate is no longer statistically significant, while the forecast of the future inflation rate remains highly statistically significant. This seems to indicate that the significance of the past inflation rate is spurious, in the sense that the past inflation rate is just proxying for the missing constant term, which we have already acknowledged to be statistically significant. The last row in the table shows, surprisingly, that even when the SPF forecast is entirely absent, the lagged inflation rate has no explanatory power for the Michigan survey after controlling for the survey's own lagged value; furthermore, the Durbin-Watson statistic suggests a substantial amount of negative serial correlation in the residuals of this equation, in contrast with the baseline model.

In sum, it seems fair to say that the simple 'sticky expectations' equation (8) does a remarkably good job of capturing much of the predictable behavior of the Michigan inflation expectations index.¹⁰

III.3. Inflation News Coverage and Inflation Expectations

If we take literally the assumption that people derive their inflation expectations from news stories, we should expect that when there are more news stories people should be better informed. Figure I plots an index of the intensity of news coverage of inflation in the New York Times and the Washington Post against the actual inflation rate;¹¹ unsurprisingly, the intensity of news coverage of inflation was highest in the early 1980s when the actual inflation rate was very high, and inflation coverage has generally declined since then. Note, however, that the actual inflation

¹⁰One further robustness test is presented in Carroll [2001]: Estimation of the model on monthly rather than quarterly data. This gets around the timing problems caused by the fact that the Michigan index for a quarter reflects interviews continuously throughout the quarter while the SPF reflects forecasters' views at a point in time roughly halfway through the quarter. Estimates of the monthly λ are roughly a third the size of estimates of the quarterly rate, as would be expected if the quarterly estimates were unbiased.

¹¹The index was constructed as follows. For each newspaper $i \in \{\text{New York Times, Washington Post}\}$, for each year t since 1980 (when the Nexis index of both newspapers begins), a search was performed for stories that began on the front page of the newspaper and contained words beginning with the root 'inflation' (so that, for example, 'inflationary' or 'inflation-fighting' would be picked up). For each newspaper, the number of stories was converted to an index ranging between zero and 1 by dividing the number of stories in a given year by the maximum number of inflation stories in any year. Thus, the fact that the overall index falls to about 0.25 in the last part of the sample indicates that there were about a quarter as many front-page stories about inflation in this time period as there were at the maximum.

rate fell farther and faster than the news index; evidently, inflation remained an important story during the period when it was dropping rapidly.

The bottom panel of Figure I plots the SPF and Michigan forecasts since the third quarter of 1981 when the SPF first began to include CPI inflation. One striking feature of the figure is that during the high-news-coverage period of the early 1980s, the size of the gap between the SPF forecast and the Michigan forecast is distinctly smaller than the gap in the later period when there was less news coverage of inflation.

A formal statistical test of whether greater news coverage is associated with ‘more rational’ household forecasts (in the sense of forecasts that are closer to the SPF forecast) can be constructed as follows. Defining the square of the gap between the Michigan and SPF forecasts as $GAPSQ_t = (M_t - S_t)^2$, and defining the inflation index as $NEWS_t$, we can estimate the simple OLS regression equation

$$(13) \quad GAPSQ_t = \alpha_0 + \alpha_1 NEWS_t.$$

Table IV presents the results. Estimated over the entire sample from 1981q3 to 2000q2 the regression finds a negative relationship that is statistically significant at the 5 percent level after correcting for serial correlation. The second row shows that that if the first year of the SPF CPI forecasts is excluded the negative relationship is much stronger and statistically significant at better than the 1 percent level; however, aside from the possibility that the first few SPF CPI forecasts were problematic in some way, there seems to be little reason to exclude the first year of SPF data.

The finding that household inflation forecasts are better when there is more news coverage is an indirect implication of the model under the assumption that absorption of the SPF forecast is more likely when there is more inflation coverage. The proposition that the absorption rate is higher when there are more news stories can also be tested directly. Table V presents estimation results comparing the absorption rate estimated during periods when there is more news coverage than average ($NEWS_t > \text{mean}(NEWS)$) and less coverage than average ($NEWS_t < \text{mean}(NEWS)$). The estimate of λ is almost 0.7 during periods of intensive news coverage, but only about 0.2 during periods of less intense coverage; an F-test indicates that this difference in coefficients is statistically significant at the 5 percent level (and nearly at the 1 percent level).

There are several strands of the existing literature that deserve comment at this point. In two important recent papers, Akerlof, Dickens, and Perry [1996, 2000] have proposed a model in which workers do not bother to inform themselves about the inflation rate unless inflation gets high enough that ignorance would become costly. Since periods of high news coverage have coincided with periods of high inflation, this model is obviously consistent with the finding that mean inflation expectations are more rational during periods of high coverage. Indeed, in a way the ADP models are deeper than the one proposed here, because they provide an explanation for the intensity of news coverage which is taken as exogenous here:

The news media write more stories on inflation in periods when workers are more interested in the topic.

These results can also be viewed as somewhat similar to some findings by Roberts [1998], who estimates a model like (8), performs a sample split, and finds the speed of adjustment parameter much larger in the post-1976 period than in the pre-76 era. He interprets this as bad news for the model. However, the pre-76 era was one of much more stable inflation (until the last years) than the post-76 era, so the finding of a higher coefficient in the later years is very much in the spirit of the tests performed above, and is therefore consistent with the interpretation of the model proposed here.

IV. Unemployment Expectations

If the model of expectations proposed here is to be generally useful to macroeconomists, it will need to apply to other variables in addition to inflation. Another potential candidate is unemployment expectations; in previous work (Carroll [1992], Carroll and Dunn [1997]) I have found unemployment expectations to be a powerful predictor of household spending decisions, and since household spending accounts for two thirds of GDP, understanding the dynamics of unemployment expectations (and any deviations from rationality) should have considerable direct interest.

Unfortunately, however, the Michigan survey's question on unemployment does not ask households to name a specific figure for the future unemployment rate; instead, households are asked whether they expect the unemployment rate to rise, stay the same, or fall over the next year. Traditionally, the answers to these questions are converted into an index by subtracting the "fall" from the "rise" proportion. This diffusion index can then be converted into a forecast of the change in the unemployment rate by using the predicted value from a regression of the actual change in unemployment on the predicted change.

That is, the regression

$$(14) \quad \bar{U}_{t,t+4} - \bar{U}_{t-4,t} = \gamma_0 + \gamma_1 M_t^U$$

is estimated, where $\bar{U}_{t,t+4}$ is the average unemployment rate over the next year and $\bar{U}_{t-4,t}$ is the unemployment rate over the year to the present, and M_t^U is the Michigan index of unemployment expectations. With the estimated $\{\hat{\gamma}_0, \hat{\gamma}_1\}$ in hand a forecast of next year's unemployment rate can be constructed from

$$(15) \quad \hat{U}_{t,t+4} = \hat{\gamma}_0 + \hat{\gamma}_1 M_t^U + \bar{U}_{t-4,t}.$$

When (14) is estimated, the coefficient on M_t^U has a t-statistic of over 8, even after correcting for serial correlation. However, in a horserace regression of the actual change in unemployment on the Michigan diffusion index and the SPF forecast of the change in unemployment, the Michigan forecast has no predictive power. Thus, as with inflation, it appears that on average people have considerable

information about how the unemployment rate is likely to change, but forecasters know a lot more than households do.

Table VI presents a set of regression results for the household unemployment forecast that is essentially identical to the tests performed in Table III for inflation expectations.

The point estimate of the speed of adjustment parameter in row 3 is $\alpha_1 = 0.31$; the test reported in the last column of that row indicates that this is statistically indistinguishable from the estimate of $\lambda = 0.25$ obtained for inflation expectations. In most respects, in fact, the model performs even better in explaining unemployment expectations than in explaining inflation expectations. For example, row 3 indicates that the equation does not particularly want a constant term in it, while row 4 finds that the lagged level of the unemployment rate has no predictive power for current expectations even when a constant is excluded.

Nonetheless, this evidence should be considered with some caution. The process of constructing the forecast for the average future level of the unemployment rate, while apparently reasonable, may be econometrically and conceptually problematic. In particular, this method assumes that the *amount* by which unemployment is expected to change on average is related to the *proportion* of people who expect unemployment to rise or fall; in fact, there is no necessary linear relationship between these two quantities. Other econometric difficulties may come from the use of constructed variables on both the left and right hand sides of the equation. I view this model of unemployment expectations merely as secondary supporting evidence for the expectations modeling strategy pursued here, and therefore am not inclined to pursue these conceptual and econometric problems further, though they might be worth pursuing in later work.

V. Discussion and Relationship to Existing Literature

A potential criticism of this paper might be that expectations of households are unimportant for macroeconomic outcomes; instead, perhaps what matters are expectations of experts, which may be rational in the traditional sense. This is not a plausible criticism with respect to unemployment expectations, given the powerful influence that households' unemployment expectations have on household spending (Carroll and Dunn [1997]); households' consumption decisions surely depend on their own views rather than the views of others. In the case of inflation expectations, whether households' expectations are important presumably depends on whether sluggish household expectations are partly or largely responsible for the costly nature of disinflations. This seems a good guess, since credible preannounced disinflations should be costless (or nearly so) in an economy in which all agents have rational expectations (Ball [1994]). It seems likely that a good part of the influence of household expectations on inflation comes through a labor market channel. Standard models of unemployment, either in the search literature or the efficiency wage

literature, almost always rely upon an assumption that households' labor supply decisions are made by judging the appeal of available or expected real wage offers. It is an empirical fact that wage contracts are usually written in nominal terms, so these models of cyclical and structural unemployment entail an implicit assumption that households can translate nominal wage offers into real ones, which requires them to have expectations about inflation. Since wages are around two thirds of business costs, if households' inflation expectations affect nominal wage outcomes they must affect firms' pricing decisions through the usual 'wage-push' channel.

The potential importance of household expectations provides a new perspective on the ongoing debate about the importance of 'credibility' in monetary policy. Credibility has usually been thought of in terms of the beliefs of policy experts and private forecasters; for example, an extensive recent treatment in Bernanke et. al. [1999] judges credibility of inflation targeting regimes by examining how quickly professional forecasts converge to the stated target range. Indeed, some central banks (such as the Bank of England and the Bank of Israel) have begun to officially look at surveys of forecasters as well as yields on inflation indexed bonds as measures of such "expert" opinion. However, empirical tests of whether credibility matters for monetary policy have produced mixed results; see, e.g., Posen [1995, 1998] or Debelle and Fischer [1994], and for an insider's perspective and an excellent summary of the literature see Blinder [1998]. The results above, however, indicate that there are substantial gaps between beliefs of forecasters and of households, so credibility among experts may not be sufficient to achieve a desired inflationary outcome; the views of the experts need to be communicated effectively to the population to become effective.

In many cases the model is likely to yield similar, though not identical, conclusions to those obtained from models with price stickiness. Consider, for example, the finding of Ball, Mankiw, and Romer [1988] that the Phillips curve is steeper when inflation is higher, which they attribute to a reduction in price stickiness at high inflation rates (a further justification of such effects can be found in the recent paper by Dotsey, King, and Wolman [1999]). If the newsworthiness of inflation forecasts depends on the level of inflation and λ reflects the intensity of news coverage, then very similar implications could be derived in this model. However, the two interpretations could be distinguished if newsworthiness is related to the *change* in inflation (as figure I suggests).

A related question is why it appears to be easier to end high inflations than moderate ones (Sargent [1982, 1983]): Perhaps the attention of the population tends to be intensely focused on government inflation-fighting policies during hyperinflations (so $\lambda = 1$ and the model collapses to rational expectations), while attention may be focused on other matters during attempts to end moderate inflations. Ball [1994] shows that quicker disinflations seem to entail smaller sacrifices of output; again, this could reflect the fact that the policies needed to achieve a quick disinflation are more dramatic, and therefore more newsworthy, than in a gradualist approach. One way of testing these ideas for more recent episodes would be to construct indices of news coverage of inflation like those presented above for other countries during

disinflationary episodes.

Of course, the real world presumably combines some degree of price stickiness and a degree of expectational stickiness. The results in Ball [1995] showing strong interactions between credibility and price stickiness suggest that a model that combines sticky expectations and sticky prices might generate results different from the results obtainable with either feature alone; this would also be an interesting topic for future research.

VI. Conclusions

Given the consensus among economists that macroeconomic outcomes depend critically on agents' expectations, it is surprising that there has been very little effort to test positive models of expectations using the large body of empirical expectations data available from the Michigan Survey and the Conference Board.

This paper shows that a very simple model in which the typical household's expectations are updated probabilistically toward the views of professional forecasters does a good job of capturing much of the variation in the Michigan Survey's measures of inflation and unemployment expectations. In addition to fitting these data, the model can be interpreted as providing microfoundations for the aggregate expectations equation postulated recently by Mankiw and Reis [2001, 2003] and earlier by Roberts [1998]. As those papers show, macroeconomic dynamics are more plausible in a variety of dimensions, including the tradeoff between inflation and unemployment, the reaction of the economy to monetary shocks, and the relationship between productivity growth and unemployment, when expectations deviate in this way from the rational expectations benchmark.

There are many directions in which research could fruitfully proceed from here. First, the Michigan and Conference Board surveys contain many other expectational variables that could be studied to see whether the specification proposed here is widely applicable. Second, the model could be tested at the micro level using the raw household-level data from the surveys. (One approach would be to extend the model of Branch [2001], in which individuals choose among various alternative predictors, to include the SPF forecast among the competing predictors). Third, the implications of more sophisticated models of the spread of expectations in the population can be examined (see Carroll [forthcoming] for a start). And much more work remains to be done to investigate the empirical and theoretical properties of macroeconomic models in which expectations are formed in this way.

Finally, it is clear that in order for this framework to be a complete and general purpose tool, it will be necessary to develop a theory that explains the variations in the absorption parameter λ over time. For present purposes it was enough to show that λ is related to the intensity of news coverage, but that only pushes the problem one step back, to the need for a model of the extent of news coverage. Possibly the approach offered in a recent paper by Sims [2001] could help; Sims examines models of 'rational inattention' which imply that an agent with limited information

processing capacity should optimally ignore most macroeconomic data. The difficulty with applying the Sims framework to consumers directly is that solving the problem of deciding what to ignore is even harder than solving the full-information rational expectations model. However, if the news media were viewed as the agents who solve the information compression problem (since the information stream they can convey is obviously limited), the Sims framework might provide a useful formal model of how the news media go about deciding how much coverage to give to economic matters.

THE JOHNS HOPKINS UNIVERSITY

References

- Akerlof, George, William Dickens, and George Perry: “The Macroeconomics of Low Inflation,” *Brookings Papers in Economic Activity*, (1996:1), 1–76.
- Akerlof, George, William Dickens, and George Perry “Near-Rational Wage and Price Setting and the Long-Run Phillips Curve,” *Brookings Papers in Economic Activity*, (2000:1), 1–60.
- Ball, Laurence “What Determines the Sacrifice Ratio?,” in *Monetary Policy*, N. Gregory Mankiw, ed., chap. 5. (Chicago:, University of Chicago Press 1994) .
- “Disinflation with Imperfect Credibility,” *Journal of Monetary Economics*, XXXV (1995), 5–23.
- (2000): “Near-Rationality and Inflation in Two Monetary Regimes,” *NBER Working Paper No. W7988*.
- Ball, Lawrence, N. Gregory Mankiw, and David Romer (1988): “The New Keynesian Economics and the Output-Inflation Tradeoff,” *Brookings Papers on Economic Activity*, 1988(1), 1–65.
- Barsky, Robert B. (1987): “The Fisher Hypothesis and the Forecastability and Persistence of Inflation,” *Journal of Monetary Economics*, 19, 3–24.
- Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen (1999): *Inflation Targeting: Lessons from the International Experience*. Princeton University Press, Princeton, NJ.
- Blinder, Alan S. (1998): *Central Banking in Theory and Practice*. MIT Press, Cambridge, MA.
- Branch, William A. (2001): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Manuscript, Department of Economics, College of William and Mary*.
- Bryan, Michael F., and William T. Gavin (1986): “Models of Inflation Expectations Formation: A Comparison of Household and Economist Forecasts,” *Journal of Money, Credit, and Banking*, 18, 539–43.
- Calvo, Guillermo A. (1983): “Staggered Contracts in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–98.
- Carroll, Christopher D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 1992(2), 61–156, <http://www.econ.jhu.edu/people/ccarroll/BufferStockBPEA.pdf> .
- (2001): “The Epidemiology of Macroeconomic Expectations,” *NBER Working Paper Number 8695*.

- “The Epidemiology of Macroeconomic Expectations,” in *The Economy as an Evolving Complex System, III*, Larry Blume, and Steven Durlauf, ed. Oxford University Press forthcoming , <http://www.econ.jhu.edu/people/ccarroll/epidemiologySFI.pdf>.
- Carroll, Christopher D., and Wendy E. Dunn “Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets,” in *NBER Macroeconomics Annual, 1997*, Benjamin S. Bernanke, and Julio J. Rotemberg, ed., pp. 165–229. Cambridge, MA, MIT Press 1997 , <http://www.econ.jhu.edu/people/ccarroll/macroann.pdf>.
- Croushore, Dean (1993): “Introducing: The Survey of Professional Forecasters,” *Federal Reserve Bank of Philadelphia Business Review*, pp. 3–15.
- (1998): “Evaluating Inflation Forecasts,” *Federal Reserve Bank of Philadelphia Working Paper Number 98-14*.
- Curtin, Richard T. (1996): “Procedure to Estimate Price Expectations,” Manuscript, University of Michigan Survey Research Center.
- Debelle, Guy, and Stanley Fischer “How Independent Should a Central Bank Be?,” in *Goals, Guidelines, and Constraints Facing Monetary Policymakers*, Jeffrey C. Fuhrer, ed., pp. 195–221. Federal Reserve Bank of Boston 1994 , Federal Reserve Bank of Boston Conference Series 38.
- Diebold, Francis X., and Jose A. Lopez “Forecast Evaluation and Combination,” in *Handbook of Statistics*, G.S. Maddala, and C.R. Rao, ed., pp. 241–68. Amsterdam, North Holland 1996 .
- Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999): “State Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *Quarterly Journal of Economics*, 114, 655–690.
- Dua, Pami, and Subash C. Ray (1992): “ARIMA Models of the Price Level: An Assessment of the Multilevel Adaptive Learning Process in the USA,” *Journal of Forecasting*, 11, 507–16.
- Evans, George W., and Seppo Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- Fishe, Raymod P. H., and Todd L. Idson (1990): “Information-Induced Heteroscedasticity in Price Expectations Data,” *Review of Economics and Statistics*, 72(2), 304–312.
- Friedman, Benjamin M. (1979): “Optimal Expectations and the Extreme Information Assumptions of ‘Rational Expectations’ Macromodels,” *Journal of Monetary Economics*, 5, 23–41.

- Fuhrer, Jeffrey C., and George R. Moore (1995): “Inflation Persistence,” *Quarterly Journal of Economics*, 110(1), 127–159.
- Keynes, John Maynard (1936): *The General Theory of Employment, Interest, and Money*. Harcourt, Brace.
- Mankiw, N. Gregory (2001): “The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment,” *Economic Journal*, 111(471), C45–C61.
- Mankiw, N. Gregory, and Ricardo Reis (2001): “Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps,” *NBER Working Paper Number 8614*.
- (2003): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, Forthcoming.
- McCallum, Bennett T. (2002): “Recent Developments in Monetary Policy Analysis: The Roles of Theory and Evidence,” *Federal Reserve Bank of Richmond Economic Quarterly*, 88(1), 67–95.
- Mehra, Yash P. (2002): “Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality,” *Manuscript, Federal Reserve Bank of Richmond*.
- Muth, John F. (1960): “Optimal Properties of Exponentially Weighted Forecasts,” *Journal of the American Statistical Association*, 55(290), 299–306.
- Newey, Whitney K., and Kenneth D. West (1987): “A Simple Positive Semi-Definite, Heteroskedasticity and Autorcorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- Posen, Adam S. “Declarations Are Not Enough: Financial Sector Sources of Central Bank Independence,” in *NBER Macroeconomics Annual*, Ben S. Bernanke, and Julio J. Rotemberg, ed., pp. 258–74. Cambridge, MA, MIT Press 1995 .
- (1998): “Central Bank Independence and Disinflationary Credibility: A Missing Link?,” *Oxford Economic Papers*, 50(3), 335–59.
- Roberts, John M. (1995): “New Keynesian Economics and the Phillips Curve,” *Journal of Money, Credit, and Banking*, 27(4), 975–984.
- (1997): “Is Inflation Sticky?,” *Journal of Monetary Economics*, pp. 173–196.
- (1998): “Inflation Expectations and the Transmission of Monetary Policy,” *Federal Reserve Board FEDS working paper Number 1998-43*.
- Sargent, Thomas J. “The Ends of Four Big Inflations,” in *Inflation: Causes and Consequences*, Robert E. Hall, ed. Chicago, University of Chicago Press 1982 .

- “Stopping Moderate Inflation: The Methods of Poincare and Thatcher,” in *Inflation, Debt, and Indexation*, Rudiger Dornbusch, and M. Simonsen, ed., pp. 54–96. Cambridge, MA, MIT Press 1983 .
- (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.
- Sims, Christopher A. (2001): “Implications of Rational Inattention,” *Manuscript, Yale University*.
- Souleles, Nicholas (forthcoming): “Consumer Sentiment: Its Rationality and Usefulness in Forecasting Expenditure; Evidence from the Michigan Micro Data,” *Journal of Money, Credit, and Banking*.
- Staiger, Douglas, James H. Stock, and Mark W. Watson (2001): “Prices, Wages, and the U.S. NAIRU in the 1990s,” *NBER Working Paper Number 8320*.
- Thomas Jr., Lloyd B. (1999): “Survey Measures of Expected U.S. Inflation,” *Journal of Economic Perspectives*, 13(4), 125–144.
- Turnovsky, Stephen J. (1970): “Empirical Evidence on the Formation of Price Expectations,” *Journal of the American Statistical Association*, 65(December), 1441–54.
- Urich, Thomas, and Paul Wachtel (1984): “The Structure of Expectations of the Weekly Money Supply Announcement,” *Journal of Monetary Economics*, 13, 183–194.

TABLE I
Forecasting Power Of Michigan and SPF Indexes

Dependent Variable: $\pi_{t,t+4}$					
Constant	$\pi_{t-5,t-1}$	$M_t[\pi_{t,t+4}]$	$S_t[\pi_{t,t+4}]$	DW Stat	R^2
0.070	0.083	0.732		0.46	0.52
(0.526)	(0.145)	(0.204)***			
0.480	-0.220		1.036	0.52	0.64
(0.323)	(0.153)		(0.161)***		
0.437	-0.219	0.027	1.015	0.52	0.64
(0.545)	(0.152)	(0.261)	(0.241)***		

$M_t[\pi_{t,t+4}]$ is the period- t mean of the Michigan survey measure of household expectations for inflation over the next year. $S_t[\pi_{t,t+4}]$ is the period- t mean of the Survey of Professional Forecasters forecast of the inflation rate over the next year. $\pi_{t-5,t-1}$ is the inflation rate between quarter $t-5$ and $t-1$, the most recently available annual data available at time t . The column labelled DW Stat reports the Durbin-Watson statistic; similar results are obtained using the Box-Ljung Q statistic. All equations were estimated over the 1981q3 to 2000q2 period. Errors are corrected for heteroskedasticity and autocorrelation using a Newey-West [1987] procedure (a Bartlett modified kernel) with 4 lags. Results were not sensitive to alternative lag length choices. One, two, and three stars indicate, respectively, statistical significance at the 10, 5, and 1 percent levels.

TABLE II
Granger Causality Between Michigan and SPF Surveys

Eqn number	Dependent variable	Independent variables			Durbin-Watson	\bar{R}^2
		Constant	Sum of coefficients on			
			$S_{t-1}..S_{t-8}$	$M_{t-1}..M_{t-8}$		
1	S_t	0.53	1.13	-0.25	2.02	0.87
		(0.14)	(0.00)***	(0.32)		
2	M_t	1.27	0.58	0.18	1.96	0.71
		(0.01)***	(0.05)**	(0.01)***		

M_t is the Michigan household survey measure of mean inflation expectations in quarter t , S_t is the Survey of Professional Forecasters mean inflation forecast. p-values for exclusion tests are in parentheses below coefficient estimates. All equations are estimated over the period 1981q3 to 2000q2 for which both Michigan and SPF inflation forecasts are available. Box-Ljung Q-tests found no evidence of serial correlation in the residuals, so standard errors are not corrected for serial correlation; a serial correlation correction does not change results. One, two, and three stars indicate, respectively, statistical significance at the 10, 5, and 1 percent levels.

TABLE III
Estimating and Testing the Baseline Model

Estimating Equation $M_t[\pi_{t,t+4}] = \alpha_0 + \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_{t-1}[\pi_{t-1,t+3}] + \alpha_3 P_t[\pi_{t-5,t-1}] + t$

Eqn	α_0	α_1	α_2	α_3	\bar{R}^2	Durbin- Watson	StdErr	Test
								p-value
Memo:	4.34 (0.19) ^{***}				0.00	0.29	0.88	$\alpha_0 = 0$ 0.000
1		0.36 (0.09) ^{***}	0.66 (0.08) ^{***}		0.76	1.97	0.43	$\alpha_1 + \alpha_2 = 1$ 0.178
2		0.27 (0.07) ^{***}	0.73 (0.07) ^{***}		0.76	2.12	0.43	$\alpha_1 = 0.25$ 0.724
3	1.22 (0.20) ^{***}	0.51 (0.08) ^{***}	0.26 (0.09) ^{***}	0.84	1.74		0.35	$\alpha_0 = 0$ 0.000
4		0.49 (0.09) ^{***}	0.67 (0.08) ^{***}	-0.15 (0.05) ^{***}	0.79	2.26	0.40	$\alpha_1 + \alpha_2 + \alpha_3 = 1$ 0.199
5	1.26 (0.27) ^{***}	0.50 (0.08) ^{***}	0.25 (0.11) ^{**}	0.01 (0.05)	0.84	1.72	0.35	$\alpha_3 = 0$ 0.814
6			1.02 (0.04) ^{***}	-0.04 (0.05)	0.71	2.63	0.47	$\alpha_2 + \alpha_3 = 1$ 0.239

$M_t[\pi_{t,t+4}]$ is the Michigan household survey measure of mean inflation expectations in quarter t , $S_t[\pi_{t,t+4}]$ is the Survey of Professional Forecasters mean inflation forecast over the next year; $P_t[\pi_{t-5,t-1}]$ is the published inflation rate for the most recent one-year period. All equations are estimated over the period 1981q3 to 2000q2 for which both Michigan and SPF inflation forecasts are available. All standard errors are corrected for heteroskedasticity and serial correlation using a Newey-West procedure (a Bartlett kernel) with four lags. Results are not sensitive to the choice of lags. Box-Ljung tests found no evidence of serial correlation for equations 1-6; the Durbin-Watson is reported because it may be more familiar to most readers.

TABLE IV

Household Inflation Expectations Are More Accurate When There Is More News Coverage

Equation Estimated: $GAPSQ_t = \alpha_0 + \alpha_1 NEWS_t$

Sample	α_0	α_1	D-W Stat	R^2
1981q3-2000q2	0.94 (0.26) ^{***}	-1.03 (0.50) ^{**}	1.01	0.08
1982q3-2000q2	1.22 (0.25) ^{***}	-1.72 (0.46) ^{***}	1.08	0.14

GAPSQ is the square of the difference between the Michigan and SPF inflation forecasts. NEWS is an index of the intensity of news coverage of inflation in the New York Times and the Washington Post from 1981 to 2000. All standard errors are corrected for heteroskedasticity and serial correlation using a Newey-West [1987] procedure with four lags. Results are not sensitive to the choice of lags. {***, **, *} = {1 percent, 5 percent, 10 percent} significance.

TABLE V

Updating Speed Is Faster When There Is More News Coverage

Eqn	Sample	λ	Durbin-Watson	Q-Test p-value
1	All obs	0.273 (0.066) ^{***}	2.12	0.971
2	$NEWS_t > \text{mean}(NEWS)$	0.699 (0.176) ^{***}	1.57	0.216
3	$NEWS_t < \text{mean}(NEWS)$	0.210 (0.077) ^{**}	1.93	0.451

The equation is estimated in the form $M_t - M_{t-1} = \lambda(S_t - M_{t-1})$ which imposes the condition $\lambda + (1 - \lambda) = 1$. All standard errors are corrected for heteroskedasticity and serial correlation using a Newey-West [1987] procedure with four lags. Results are not sensitive to the choice of lags. {***, **, *} = {1 percent, 5 percent, 10 percent} significance.

TABLE VI

Estimating and Testing the Baseline Model for Unemployment

Estimating Equation $M_t[U_{t,t+4}] = \alpha_0 + \alpha_1 S_t[U_{t,t+4}] + \alpha_2 M_{t-1}[U_{t-1,t+3}] + \alpha_3 P_t[U_{t-5,t-1}] + \epsilon_t$

Eqn	α_0	α_1	α_2	α_3	\bar{R}^2	Durbin- Watson	Standard error	Test
								p-value
Memo:	6.38 (0.29)***				0.00	0.08	1.29	$\alpha_0 = 0$ 0.000
1		0.32 (0.07)***	0.68 (0.07)***		0.94	1.73	0.32	$\alpha_1 + \alpha_2 = 1$ 0.111
2		0.31 (0.07)***	0.69 (0.07)***		0.94	1.72	0.32	$\alpha_1 = 0.25$ 0.375
3	-0.03 (0.18)	0.32 (0.07)***	0.68 (0.07)***		0.94	1.74	0.32	$\alpha_0 = 0$ 0.847
4		0.32 (0.07)***	0.67 (0.09)***	0.01 (0.05)	0.94	1.73	0.33	$\alpha_1 + \alpha_2 + \alpha_3 = 1$ 0.112
5	-0.04 (0.18)	0.32 (0.07)***	0.67 (0.09)***	0.01 (0.06)	0.94	1.72	0.33	$\alpha_3 = 0$ 0.855

$M_t[U_{t,t+4}]$ is a forecast of the average unemployment rate over the next year in quarter t derived as described in the text from the Michigan survey measure of unemployment expectations; $S_t[U_{t,t+4}]$ is the mean of the SPF unemployment forecast over the next four quarters; $P_t[U_{t-5,t-1}]$ is the published unemployment rate for the most recent one-year period. All equations are estimated over the period 1978q1 to 2000q2 for which both Michigan and SPF unemployment forecasts are available. All standard errors are corrected for heteroskedasticity and serial correlation using a Newey-West procedure with four lags. Results are not sensitive to the choice of lags.

FIGURE I
Inflation Versus News Stories




FIGURE II
Michigan Versus SPF Forecasts

