

Carroll, Christopher D.

Working Paper

The epidemiology of macroeconomic expectations

Working Paper, No. 462

Provided in Cooperation with:

Department of Economics, The Johns Hopkins University

Suggested Citation: Carroll, Christopher D. (2001) : The epidemiology of macroeconomic expectations, Working Paper, No. 462, The Johns Hopkins University, Department of Economics, Baltimore, MD

This Version is available at:

<https://hdl.handle.net/10419/72015>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Epidemiology of Macroeconomic Expectations

Christopher D. Carroll[†]
ccarroll@jhu.edu

December 7, 2001

ABSTRACT

Since the foundational work of Keynes (1936), macroeconomists have emphasized the importance of agents' expectations in determining macroeconomic outcomes. Yet in recent decades macroeconomists have devoted almost no effort to modeling actual empirical expectations data, instead assuming all agents' expectations are 'rational.' This paper takes up the challenge of modeling empirical household expectations data, and shows that a simple, standard model from epidemiology does a remarkably good job of explaining the deviations of household inflation and unemployment expectations from the 'rational expectations' benchmark. Furthermore, a microfoundations or 'agent-based' version of the model may be able to explain, in a way that still permits aggregation, stark rejections of the pure rational expectations framework like Souleles's (2002) finding that members of different demographic groups have sharply different predictions for macroeconomic aggregates like the inflation rate.

Keywords: inflation, expectations, unemployment, monetary policy, agent-based modeling

JEL Classification Codes: E0, E3

I am grateful to Jason Harburger, Jirka Slacalek, and Johanna Francis for excellent research assistance, to Carl Christ, Bill Dickens, Daniel Leigh, Jennifer Manning, Serena Ng, John Roberts, and Martin Sommer, for valuable feedback, and to Richard Kwok for guidance to the relevant epidemiology literature.

The data, econometric programs, and simulation 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.

1 Introduction

Ever since the foundation of macroeconomic theory by John Maynard Keynes (1936), economists have understood that macroeconomic outcomes depend critically upon expectations about those outcomes. Keynes himself believed that in the short run, economies could experience booms and busts that reflected movements in the ‘animal spirits’ of business leaders (a view that has some appeal at the current moment of dot-com hangover), but the basis for most of today’s macro models was laid in the ‘rational expectations revolution’ of the 1970s. Led by Lucas, Sargent, Barro, and others, this approach made a set of assumptions that were much stronger than rationality alone. In particular, the framework assumes that all agents in the economy are not merely rational, but also share identical (correct) beliefs about the structure of the economy, and have instantaneous and costless access to all the latest economic data. Each agent combines these data with the true macroeconomic model to obtain a forecast for the future path of the economy, on the assumption that all other agents have identical beliefs and information (and therefore forecasts).

This set of assumptions has turned out to be a powerful vehicle for macroeconomic modeling, but has never been free from the criticism that it does not resemble the real world of conflicting opinions and forecasts, workers (and even some business leaders) who may not pay much attention to macroeconomic matters, and information that can sometimes be costly to obtain and process. Rational expectations models also have problems explaining some robust stylized facts, such as the apparent inexorability of the tradeoff between inflation and unemployment rate (see Ball (1994) or Mankiw (2001)). Despite these problems, the rational expectations framework remains the dominant approach in macro theory, partly because it tends to be mathematically more tractable than alternatives that relax one or another of the framework’s assumptions.

This paper proposes a tractable alternative framework for the formation of a typical person’s expectations. Rather than having their own macroeconomic model and constantly feeding it the latest statistics, typical people are assumed to obtain their views about the future path of the economy from the news media. Furthermore (and importantly), not every person pays close attention to all macroeconomic news; instead, people are assumed to absorb the economic content of news reports probabilistically, so that it may take quite some time for news of changed macroeconomic circumstances to penetrate to all agents in the economy.

Roberts (1998) and Mankiw and Reis (2001) have recently proposed aggregate expectations equations that are mathematically very similar to aggregate implications of the baseline model derived here. Roberts’s work was motivated by his separate findings in Roberts (1995, 1997) that empirical macro models perform better in a variety of dimensions when survey-based inflation expectations are used in place of constructed model-consistent rational expectations. Mankiw and Reis (2001) obtain

similar findings, and particularly emphasize the point that these models can explain the inexorability of an inflation-unemployment tradeoff much better than the standard model with rational expectations does.

However, neither Roberts (1998) nor Mankiw and Reis (2001) devoted much effort to explaining *why* the dynamics of aggregate expectations should evolve as they proposed (though Roberts does offhandedly suggest that his equation might result from the diffusion of press reports). Mankiw and Reis motivate their model loosely by suggesting that there are calculation and information-processing costs that must be paid every time an agent updates his macroeconomic forecast; however, they do not provide an explicit processing-costs microfoundation, and indeed it seems unlikely that any plausible assumption about information costs would lead to exactly the aggregate equation they use.

In contrast, I provide an explicit microfoundation for an aggregate expectations equation, grounded in mathematically precise models from theoretical epidemiology. Rather than tracking the spread of a disease through a population, the model will track the spread of a piece of information (specifically, the information corresponding to the latest rational forecast of inflation).

The model's foundation in epidemiology should prove fruitful, both because it has novel direct implications of its own, and because the rich preexisting literature on more complex models of the spread of disease may yield further important and testable insights.¹ For example, taking literally the proposition that exposure to news stories is analogous to contact with an infectious agent leads to the hypothesis that during periods when there is more news coverage of inflation, news should spread faster and household expectations should be closer to rational expectations. These predictions are tested and confirmed using an index of the intensity of news coverage about inflation.

The paper's final contribution is to relate the model to the burgeoning literature on agent-based modeling and social networks that has been pioneered, among other places, at the Santa Fe Institute and the Center for Social and Economic Dynamics (CSED) at the Brookings Institution. The agent-based modeling approach is useful here for two reasons. First, this approach permits tests of the robustness of the baseline model; using agent-based simulations it is straightforward to explore a variety of plausible alternative assumptions about the spread of information that do not yield the simple analytical equation for the dynamics of aggregate expectations obtained for the baseline model. The main finding in these explorations is that the equation implied by the baseline model is a good approximation to the dynamics of aggregate expectations under several alternative assumptions about the information transmission process. The second advantage of the agent-based approach is that it generates agent-level statistics that can be compared to the household-level survey data (heretofore mostly neglected) from which the aggregate expectations indexes are constructed. As a large literature

¹For another example of the application of epidemiological models to economics see Kremer (2000).

over the last decade has shown, the ability to test aspects of a macro model with micro data can enrich both micro and macroeconomics. Here, I show that the qualitative pattern of the changes in the cross-household standard deviation of expectations in the data roughly matches the predictions of the model (though the level of the standard deviation is greater in the data than in the baseline model).

2 The Epidemiology of Expectations

2.1 A Special Case of the Kermack-McKendrick Model

Epidemiologists have developed a rich set of models for the transmission of disease in a population. The baseline framework is called the Kermack-McKendrick (1927) model. This model consists of a set of assumptions about who is susceptible to the disease, who among the susceptible becomes infected, and whether and how individuals recover from the infection (leading to an alternate designation as the ‘SIR’ framework).

The standard assumption for the discrete-time Kermack-McKendrick/SIR model is that a susceptible individual who is exposed to the disease in a given period has a fixed probability p of catching the disease. Designating the set of newly infected individuals in period t as N_t and the set of susceptible individuals as S_t ,

$$N_t = pS_t. \tag{1}$$

The next step is to determine susceptibility. The usual assumption is that in order to be susceptible, a healthy individual must have contact with an already-infected person. In a population where each individual has an equal probability of encountering any other person in the population (a ‘well-mixed’ population), the growth rate of the disease will depend upon the fraction of the population already infected; if very few individuals are currently infected, the small population of diseased people can infect only a small absolute number of new victims.

However, there is a special case that is even simpler. This occurs when the disease is not spread person-to-person, but through contact with a ‘common source’ of infection. The classic example is Legionnaire’s disease, which was transmitted to a group of hotel guests via a contaminated air conditioning system (see Fraser et. al. (1977) for a description from the epidemiological literature). Another application is to illness caused by common exposure to an environmental factor such as air pollution. In these cases, the transmission model is extremely simple: Any healthy individual is simply assumed to have a constant probability per period of becoming infected from the common source. This is the case we will examine, since below we will assume that news reports represent a ‘common source’ of information available to all members of the population.

One further assumption is needed to complete the model: The probability that someone who is infected will recover from the disease. The simplest possible assumption (which we will use) is that infected individuals never recover.

Under this set of assumptions, the dynamics of the disease are as follows. In the first period, proportion p of the population catches the disease, leaving $(1 - p)$ uninfected. In period 2, proportion p of these people catch the disease, leading to a new infection rate of $p(1 - p)$ and to a fraction $p + p(1 - p)$ of the population being infected. Spinning this process out, it is easy to see that starting from period 0 at the beginning of which nobody is infected, the total proportion infected at the end of t periods is

$$\text{Fraction Ill} = p + p(1 - p) + p(1 - p)^2 \dots + p(1 - p)^t \quad (2)$$

$$= p \sum_{s=0}^t (1 - p)^s \quad (3)$$

whose limit as $t \rightarrow \infty$ is $p/p = 1$, implying that (since there is no recovery) everyone will eventually become infected. In the case where ‘infection’ is interpreted as reflecting an agent’s knowledge of a piece of information, this simply says that eventually everyone in the economy will learn a given piece of news.

The last section of the paper explores a richer epidemiological model in which disease is spread person-to-person as well as through contact with the news media. And further extensions are possible, and would probably prove interesting; as an illustration of some of the possibilities, see Kremer (2000).

2.2 The Epidemiology of Inflation Expectations

Now consider a world where most people form their expectations about future inflation by reading newspaper articles. Imagine for the moment that every newspaper 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. Finally, suppose that at any point in time t all newspaper articles print identical forecasts.

Assume that not everybody reads every newspaper article on inflation. Instead, reading an article on inflation is like becoming infected with a common-source disease: In any given period each individual faces a constant probability λ of becoming ‘infected’ with the latest forecast by reading an article. Individuals who do not encounter an inflation article simply continue to believe the last forecast they read about.²

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

$$\pi_{t+1} = \log(p_{t+1}) - \log(p_t), \quad (4)$$

²This is mathematically very similar to the Calvo (1983) model in which firms change their prices with probability p .

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 $N_t[\pi_s]$, by analogy with equation (2) we have that

$$M_t[\pi_{t+1}] = \lambda N_t[\pi_{t+1}] + (1 - \lambda) \{ \lambda N_{t-1}[\pi_{t+1}] + (1 - \lambda) (\lambda N_{t-2}[\pi_{t+1}] + \dots) \} \quad (5)$$

The derivation of this equation is as follows. In period t a fraction λ of the population will have been ‘infected’ with the current-period newspaper forecast of the inflation rate next quarter, $N_t[\pi_{t+1}]$. Fraction $(1 - \lambda)$ of the population retains the views that 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).³ Mankiw and Reis loosely motivate the equation by arguing that developing a full-blown inflation forecast is a costly activity, which people might therefore engage in only occasionally.

It is undoubtedly true that developing a reasonably rational quarter-by-quarter forecast of the inflation rate 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!). If this were really what people were doing, one might expect them to make forecasts only very rarely indeed.

However, reading a newspaper article about inflation, or hearing a news story on television or the radio, is not costly in either time or money. There is no reason to suppose that people need to make forecasts themselves if news reports provide such forecasts essentially for free. Thus the epidemiological derivation of this equation seems considerably more attractive than the loose calculation-costs motivation provided by Mankiw and Reis, both because this is a fully specified model and because it delivers further testable implications (for example, if there is empirical evidence that people with higher levels of education are more likely to pay attention to news, the model implies that their inflation forecasts will on average be closer to the rational forecast; see below for more discussion of possible variation in λ across population groups).

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 (5), and even if they did it is very unlikely that a typical person would be able to remember the

³It is also similar to a formulation estimated by Roberts (1997), except that Roberts uses past realizations of the inflation rate rather than past rational forecasts.

detailed pattern of inflation rates far into the future. In order to relax these unrealistic assumptions 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 people believe that future changes in the fundamental rate are unforecastable; that is, after the next period the fundamental rate follows a random walk. Furthermore, suppose the person believes 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 person believes that the inflation process is captured by

$$\pi_t = \pi_t^f + \epsilon_t \tag{6}$$

$$\pi_{t+1}^f = \pi_t^f + \eta_{t+1} \tag{7}$$

$$\pi_{t+2}^f = \pi_{t+1}^f + \eta_{t+2} \tag{8}$$

...

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 in period t . We further 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 fundamental inflation rate are unforecastable, and transitory shocks are expected to go away.⁴

Before proceeding it is worth considering whether this is a plausible view of the inflation process; we would not want to build a model on an assumption that people believe something patently absurd. Certainly, it would not be plausible to suppose that people always and everywhere believe that the inflation rate is characterized by (6)-(8); for example, Ball (2000) shows that in the US from 1879-1914 the inflation rate was not persistent in the US, while in other countries there have been episodes of hyperinflation (and rapid disinflation) in which views like (6)-(8) would have been nonsense.

However, the relevant question for the purposes of this paper is whether this view of the inflation process is plausible for the period for which I have inflation expectations data. Perhaps the best way to examine this is to ask whether the univariate statistical process for the inflation rate implied by (6) and (7) is strongly at odds with the actual univariate inflation process. In other words, after allowing for transitory shocks, does the inflation rate approximately follow a random walk?

The appropriate statistical test is an augmented Dickey-Fuller test. Table 1 presents the results from such a test. The second row shows that even with more than 160 quarters of data it is not possible to reject at a 5 percent significance level the proposition

⁴Note that we are allowing people to have some idea about how next quarter's fundamental rate may differ from the current quarter's rate, because we did not impose that consumers' expectations of η_{t+1} must equal zero.

Lags	Degrees of Freedom	ADF Test Statistic
0	166	3.59***
1	165	2.84*
2	164	2.28

This table presents results of standard Dickey-Fuller and Augmented Dickey-Fuller tests for the presence of a unit root in the core rate of inflation (results are similar for CPI inflation). The column labelled ‘Lags’ indicates how many lags of the change in the inflation rate are included in the regression. With zero lags, the test is the original Dickey-Fuller test; with multiple lags, the test is an Augmented Dickey Fuller test. In both cases a constant term is permitted in the regression equation. The sample is from 1959q3 to 2001q2 (quarterly data from my DRI database begin in 1959q1. In order to have the same sample for all three tests, the sample must be restricted to 1959q3 and after.) One, two, and three stars indicate rejections of a unit root at the 10 percent, 5 percent, and one percent thresholds. RATS code generating these and all other empirical results is available at the author’s website.

Table 1: Dickey-Fuller and Augmented Dickey-Fuller Tests for a Unit Root in Inflation

that the core inflation rate follows a random walk with a one-period transitory component - that is, it is not possible to reject the process defined by (6)-(8).⁵ When the transitory shock is allowed to have effects that last for two quarters rather than one, it is not possible to reject a random walk in the fundamental component even at the 10 percent level of significance (the last row in the table).

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).

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. This is not an implausible assumption; newspaper articles on inflation (which are usually published the day after inflation statistics are released) often contain interviews with expert forecasters, who are frequently asked

⁵The 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.

for their forecasts of future inflation rates. The most common such forecast is for the inflation rate over the next year. It is not implausible, therefore, to suppose that at least part of what readers take away from such stories is a sense of what to expect for the average 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 (5). 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

$$\pi_{t,t+1} = 4(\log p_{t+1} - \log p_t) \quad (9)$$

$$= 4\pi_{t+1} \quad (10)$$

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

$$\pi_{t,t+4} = \pi_{t+1} + \pi_{t+2} + \pi_{t+3} + \pi_{t+4} \quad (11)$$

$$= \pi_{t+1}^f + \epsilon_{t+1} + \pi_{t+2}^f + \epsilon_{t+2} + \pi_{t+3}^f + \epsilon_{t+3} + \pi_{t+4}^f + \epsilon_{t+4} \quad (12)$$

$$= \pi_{t+1}^f + \epsilon_{t+1} + \pi_{t+1}^f + \eta_{t+2} + \epsilon_{t+2} + \pi_{t+1}^f + \eta_{t+2} + \eta_{t+3} + \epsilon_{t+3} + \pi_{t+1}^f + \eta_{t+2} + \eta_{t+3} + \eta_{t+4} + \epsilon_{t+4}. \quad (13)$$

To proceed, it will be necessary to consider people's expectations about future inflation rates. 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 we made earlier about the stochastic processes for ϵ and η imply that $F_t[\epsilon_{t+n}] = F_t[\eta_{t+n+1}] = 0$ for all $n > 0$.

Applying the F_t operator to both sides of (13) implies 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:

$$F_t[\pi_{t,t+4}] = 4F_t[\pi_{t+1}^f] \quad (14)$$

$$= F_t[\pi_{t,t+1}^f] \quad (15)$$

Now for an important conclusion: If people believe that the forecasts printed in the newspaper also embody the same view of the inflation process embodied in (6)-(8), then an identical analysis leads to the conclusion that (defining the 'newspaper expectations' operator N_t similarly to the consumer's expectations operator):

$$N_t[\pi_{t,t+4}] = 4N_t[\pi_{t+1}^f] \quad (16)$$

$$= N_t[\pi_{t,t+1}^f] \quad (17)$$

Thus, the newspaper forecast contains only a single important piece of information about the future: a projection of the fundamental inflation rate over the next year, which the unit root theory 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:

$$F_t[\pi_{t,t+1}] = F_t[\pi_{t,t+4}] = F_t[\pi_{t,t+4}^f] = N_t[\pi_{t,t+4}^f] = N_t[\pi_{t,t+4}]. \quad (18)$$

The rightmost equality holds because the consumer assumes the newspaper has no information about ϵ_{t+n} or η_{t+n+1} , so for $n > 0$, $N_t[\epsilon_{t+n}] = N_t[\eta_{t+n+1}] = 0$. 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 $F_t[\epsilon_{t+n}] = F_t[\eta_{t+n+1}] = 0$.

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

$$F_{t-1}[\pi_{t-1,t+3}] = F_{t-1}[\pi_{t,t+4}] \quad (19)$$

$$F_{t-2}[\pi_{t-2,t+2}] = F_{t-2}[\pi_{t,t+4}] \quad (20)$$

$$\dots \quad (21)$$

Now note that an equation similar to (5) 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 (19)-(21) 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

$$\begin{aligned} M_t[\pi_{t,t+4}] &= \lambda F_t[\pi_{t,t+4}] + (1 - \lambda) \{ \lambda F_{t-1}[\pi_{t-1,t+3}] + (1 - \lambda) (\lambda F_{t-2}[\pi_{t-2,t+2}] + \dots) \} \\ &= \lambda F_t[\pi_{t,t+4}] + (1 - \lambda) M_{t-1}[\pi_{t-1,t+3}] \end{aligned} \quad (22)$$

$$= \lambda N_t[\pi_{t,t+4}] + (1 - \lambda) M_{t-1}[\pi_{t-1,t+3}]. \quad (23)$$

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 constructed.⁶

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 examining the underlying logic that might produce it).

Readers uncomfortable with the strong assumptions needed to derive (23) may be happier upon noting that the equation

$$M_t[\pi_{t,t+4}] = \lambda N_t[\pi_{t,t+4}] + (1 - \lambda)M_{t-1}[\pi_{t,t+4}] \quad (24)$$

can be derived without any assumptions on consumers' beliefs about the inflation process; the difference between (23) and (24) 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 projection of the difference between the inflation rates in quarters $t+1$ and $t+5$.

3 Estimation

Estimating equation (23) 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 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 (for details on the exact questions, and the controls for question validity, see Curtin (1996)). 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 median 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 instrument 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 5 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.

3.1 Do the Forecasts Forecast?

There is a substantial existing literature on the forecasting performance of various survey measures of inflation expectations including the Michigan Survey and the SPF. Early papers (Turnovsky (1970), Bryan and Gavin (1986)) claimed to find statistically significant biases in some survey measures of expected inflation, but a recent review paper by Croushore (1998) shows that some of those results were spurious (due to improper treatment of the data or econometric problems), and that none of the results hold up when the sample period is updated to include data for the last 10 or 15 years. Croushore specifically examines both the Michigan survey and the SPF, and finds no evidence of bias over the entire sample for either survey. In most respects he finds the SPF a better forecaster of inflation outcomes than the Michigan survey.⁷

For the purposes of this paper, there are two important ‘sniff-test’ questions that should be addressed before attempting formal estimation of the model (23). First, do both the Michigan and SPF survey measures of inflation expectations have statistically significant ability to predict future inflation? And, second, is the SPF forecast better in some sense than the Michigan forecast, as assumed in the model?

As a first step, consider the implications of the statistical test performed in Table 1, which was unable to reject the hypothesis that inflation has a unit root. The high serial correlation in the inflation rate discovered by the unit root test means that future levels of the inflation rate will be highly predictable based on the recent past history of inflation. The interesting question is therefore whether the survey forecasts have predictive power for the future inflation rate *beyond* what could be predicted based on past inflation data.

⁷Unfortunately, the well-known 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 the mean, but my model delivers predictions only for the mean and not for the median, so I neglect the Michigan median in my empirical work. See the appendix to Roberts (1997) for more evidence on the inefficiency of the Michigan survey.

Dependent Variable: $\pi_{t,t+4}$

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

Dependent Variable: $\pi_{t+4,t} - \pi_{t-2,t-1}$

Constant		$M_t[\pi_{t,t+4}] - \pi_{t-2,t-1}$	$S_t[\pi_{t,t+4}] - \pi_{t-2,t-1}$	DW Stat	R^2
-0.760 (0.158)***		1.005 (0.140)***		0.60	0.73
-0.263 (0.131)**			1.252 (0.140)***	0.45	0.77
-0.440 (0.189)**		0.360 (0.239)	0.854 (0.219)***	0.43	0.78

$M_t[\pi_{t,t+4}]$ 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. $\pi_{t-2,t-1}$ is the inflation rate between quarter $t-2$ and $t-1$, expressed at an annual rate. The column labelled DW Stat reports the Durbin-Watson statistic. All equations were estimated over the 1981q3 to 2000q2 period for which both Michigan and SPF inflation forecasts are available. 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 2: Forecasting Power Of Michigan and SPF Indexes

To answer this question, Table 2 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 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.⁸ Both measures of inflation expectations remain highly statistically significant even controlling for previously

⁸Actually, quarter $t-1$'s inflation rate is included among the regressors even though it is not known until about the middle of the first month of quarter t , so about a sixth of the people whose views are summarized in the Michigan survey cannot have known π_{t-1} . This makes it harder for

known inflation rates, and 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

An alternative, and more stringent, test of the predictive power of the surveys is whether they can predict the *change* in the inflation rate. Under the assumption that the fundamental inflation rate follows the unit root process outlined in equations (6) to (8) (which was consistent with the empirical tests from Table 1), the difference between the inflation rate over the coming year and the most recently published quarterly inflation statistic is (cf. (13))

$$\begin{aligned} \pi_{t,t+4} - \pi_{t-2,t-1} &= 4 F_t [\pi_t] + 4\eta_{t+1} + 3\eta_{t+2} + 2\eta_{t+3} + \eta_{t+4} + \epsilon_{t+1} + \epsilon_{t+2} + \epsilon_{t+3} + \epsilon_{t+4} \\ &\quad - 4(F_{t-1}[\pi_{t-1}] + \epsilon_{t-1}) \\ &= 4\eta_t + 4\eta_{t+1} + 3\eta_{t+2} + 2\eta_{t+3} + \eta_{t+4} \\ &\quad + \epsilon_{t+1} + \epsilon_{t+2} + \epsilon_{t+3} + \epsilon_{t+4} - 4\epsilon_t - 4\epsilon_{t-1} \end{aligned}$$

Earlier we assumed that people believe that professional forecasters have some knowledge of ϵ_t , η_t , and η_{t+1} , so even under the assumption that the inflation process has a unit root this equation tells us that there should be a substantial component of predictability in the change in the inflation rate. Furthermore, we have thus far made no assumptions about what the professional forecasters actually know about the future; for example, they may use the unemployment rate to forecast changes in inflation. Thus, there is even more scope for the SPF than for the Michigan survey to predict changes in the inflation rate.

Results from regressions of $\pi_{t,t+4} - \pi_{t-2,t-1}$ on the two survey forecasts of inflation are presented in the next three rows of table 2. Again, both survey measures have highly statistically significant explanatory power for the change in the inflation rate, and again the SPF forecast wins the horserace by a big margin.¹⁰

In sum, all these results are qualitatively what would be expected from the model as described above: the forecast from the survey of households has significant predictive

the Michigan survey to be statistically significant because the lagged inflation statistic has an unfair advantage since it incorporates information some surveyed people could not have known; the survey is significant despite this handicap. This issue does not arise with the SPF forecast because π_{t-1} has been reported by the time the SPF survey is conducted.

⁹Technical considerations suggest that it might be more appropriate to include several lags of quarterly inflation rates as the control variables for past inflation rates in the regressions reported in table 2. When this is done the results remain substantially the same. Results for the annual inflation rate are reported because they are less cumbersome to present. Quarterly results are available from the author, or can be generated by running the set of RATS programs available on the author’s web page that generated all empirical results in this paper.

¹⁰Similar results are obtained when last quarter’s inflation rate is replaced by the inflation rate over the year leading up to the previous quarter.

power for future inflation, but the SPF forecast has much more power.

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. It remains possible that inflation forecasts could have been constructed that were even better than those of the SPF. However, Croushore (1998) shows that it would have been difficult to improve upon the SPF forecasts using information that was available at the time the SPF forecasts were made.¹¹ Furthermore, from the standpoint of the model the appropriate measure of $N_t[\pi_{t,t+1}]$ is what people might actually be able to read about in a newspaper or hear about in a news report on television. Since the Survey of Professional Forecasters is designed to incorporate forecasts from the leading macroeconomic forecasters in each period, it is likely that the economists who are contacted by journalists writing about inflation will be the same people contributing their forecasts to the SPF. If this is true, the SPF forecast would be the right one to use for our purposes even if there were systematic biases in its forecasts.

3.2 Estimating the Stickiness of Inflation Expectations

We can now turn to the main question, which is whether the Michigan forecast can be reasonably well represented as a distributed lag of the SPF forecast; that is, whether (23) provides a good empirical description of the dynamics of the Michigan survey's measure of household inflation expectations.

¹¹Croushore points out an important flaw in previous studies of the rationality of forecasts; previous authors regressed inflation outcomes on information that was available at the time the SPF forecasts were made, and concluded that any finding of statistical significance indicated a failure of rationality. But such a test implicitly assumes that rationality means that forecasters should have known in advance not just that a data realization had occurred, but also the effects that that realization should have on future inflation rates. For example, it supposes that 'rational' forecasters in 1973 would have known the precise effects oil shocks would have on subsequent inflation rates. But no similar experience existed upon which forecasters could base their forecasts, so finding *ex post* that an oil shock variable is statistically significant does not imply that forecasters were irrational not to have predicted its effects *ex ante*.

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

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; P_t 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.

Table 3: Estimating and Testing the Baseline Model

To provide a baseline for comparison, the first line of Table 3 presents results for the simplest possible model: that the value of the Michigan index of inflation expectations $M_t[\pi_{t,t+4}]$ 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

$$M_t[\pi_{t,t+4}] = \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_t[\pi_{t-1,t+3}] + \nu_t, \quad (25)$$

where $S_t[\pi_{t,t+4}]$ is the corresponding SPF forecast. Comparing this to (23) provides the testable restriction that $\alpha_2 = 1 - \alpha_1$ or, equivalently,

$$\alpha_1 + \alpha_2 = 1. \tag{26}$$

Results from the estimation of (25) are presented as equation 1 in Table 3. The point estimates of $\alpha_1 = 0.36$ and $\alpha_2 = 0.66$ suggest that the restriction (26) is very close to holding true, and the last column presents formal statistical evidence on the question: It shows the statistical significance with which the proposition that $\alpha_1 + \alpha_2 = 1$ can be rejected. The p-value indicates 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 of the equation (a Q-test yields the same result), which is impressive because the individual series involved have very high degrees of serial correlation (cf. the unit root tests in table 1 and the Durbin-Watson statistics in table 2). 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) 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 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, this estimate 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. Since such a 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 further 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 considerably less precise.

Intuitively it might seem that if almost 70 percent of agents have inflation expectations that are of a vintage of a year or less, the behavior of the macroeconomy could

not be all that different from what would be expected if all expectations were completely up-to-date. The surprising message of Roberts (1997, 1995) and Mankiw and Reis (2001) is that this intuition is wrong. Mankiw and Reis show that an economy with $\lambda = 0.25$ behaves in ways that are sharply different from an economy with fully rational expectations ($\lambda = 1$), and argue that in each case where behavior is different the behavior of the $\lambda = 0.25$ economy corresponds better with empirical evidence (for example with respect to the effect of interest rate shocks on aggregate output).

Mankiw and Reis (2001) simply postulated $\lambda = 0.25$. What equation 2 of Table 3 indicates is that if the data are forced to choose a λ they are happy with that choice. However, we have not allowed the data to speak to the question of whether there is a better representation of inflation expectations than (25).

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 of table 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; the data do indeed want a constant in the regression equation. 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 standard deviation 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 believe in a positive inflation rate (of about 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 single stable point of the estimated equation. (See the final section of the paper for a demonstration that the presence of a significant constant term could reflect the presence of some social transmission of inflation expectations via conversations with neighbors, in addition to the news-media channel examined here).

A modification to the model that makes more sense than a constant term is to allow for the possibility that some people update their view of the fundamental inflation rate to the most recent *past* inflation rate rather than to the SPF forecast of the future inflation rate. Indeed, since most news coverage of inflation is prompted by the release of the most recent past inflation statistics (and since the new number is often in the headline of the news article) one might argue that it would be *more* likely for people

to update their expectations to the past inflation rate than to a forecast of the future rate.

We can examine this possibility by estimating an equation of the form

$$M_t[\pi_{t,t+4}] = \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_{t-1}[\pi_{t-1,t+3}] + \alpha_3 P_t[\pi_{t-5,t-1}]. \quad (27)$$

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

In the epidemiological sense, derivations similar to those provided above for the baseline model can be performed to show that this equation can be interpreted as representing a model in which there are two competing sources of ‘infection,’ the forecast of the future inflation rate and the actual past inflation rate. A given person’s view of the fundamental inflation rate may come from either of these two sources, but a given person can be infected from only one of these sources at a time.

Results from estimating this equation are presented in the next row of Table 3. 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 lagged value of the survey; furthermore, the Durbin-Watson suggests a substantial amount of negative serial correlation in the residuals of this equation.

In sum, it seems fair to say that the simple ‘sticky expectations’ equation (23) does a remarkably good job of capturing much of the predictable behavior of the Michigan inflation expectations index.

3.3 Timing Issues: Monthly vs Quarterly, Michigan vs SPF

However, there is a potential problem. If we want to take our model’s updating equation (23) seriously as reflecting the effect of news reports on people’s perceptions, it is logically necessary for the news reports to have been published before people’s expectations can move toward them. This obvious point raises some conceptual difficulties about the empirical results presented thus far, because the actual timing of the quarterly Michigan and SPF surveys does not match up with the implied time structure in the model we have presented and estimated.

Most newspaper stories on inflation appear the day after the monthly release of inflation statistics, which happens around mid-month every month. Between mid-month and the beginning of the next month, there are typically few economic statistics released that would have much impact on a rational forecast of the inflation rate. Given these facts, if the Michigan survey for month $t + 1$ were conducted entirely on the first day of month $t + 1$, there would be little reason to object to an updating equation of the form

$$M_{t+1}[\pi_{t+1,t+13}] = \lambda N_t[\pi_{t,t+12}] + (1 - \lambda)M_t[\pi_{t,t+12}], \quad (28)$$

where N_t designates the inflation forecast embodied in the typical newspaper article published in the middle of month t .

Now consider the circumstances under which it would be appropriate to proxy the newspaper forecasts N_t with the Survey of Professional Forecasters' forecast S_t . The SPF is conducted once a quarter, during the quarter's middle month. Forecasts are collected beginning right after the inflation data for the previous month are reported, which is exactly the time that most news articles on inflation are being written; although survey participants have a week or so to turn in their forecasts, typically little or no important economic data are released before the survey deadline, so there is little reason to expect the views reflected in the SPF to be much different from the views the news media will have reported in mid-month. Hence, if month t was an SPF survey month there would be little reason to object to proxying N_t with S_t in (28).

However, an analogous substitution would not be appropriate for month $t + 1$ or $t + 2$, because these months will have no new survey data and substituting an out-of-date SPF forecast for the newspaper forecast could lead to biases. Thus, equation (28) should be estimated using data only from the set of months for which an SPF was done in the prior month.

Appendix A discusses results of estimating the model using monthly data from months in which there was an SPF forecast in the previous month. The conclusion is that results obtained using monthly data are close to what would have been expected if there were no timing problems for the quarterly estimates. Thus there is no reason to worry that the baseline estimate of $\lambda = 0.25$ is biased by the timing difficulties laid out above. (The results are presented in an appendix because there are several technical and econometric issues that must be addressed that would interrupt the flow of the main argument of the paper without contributing anything of substantive prominence.)

3.4 Measuring Inflation News

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 on inflation people should be better informed.

This is testable. Appendix B describes the construction of an annual index of the number of news stories containing the word ‘inflation’ that began on the front pages of the New York Times or the Washington Post.¹² The index is plotted against the actual inflation rate in the top panel of Figure 1; 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.

The bottom panel of Figure 1 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

$$GAPSQ_t = \alpha_0 + \alpha_1 NEWS_t \quad (29)$$

Table 4 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 undetermined way, there seems to be little reason to exclude the first year of SPF data.

¹²The ‘began on the front page’ criterion is a standard way to control for the possibility that newspapers may shrink or grow substantially over time. Since the size of the front page has remained the same size over time, an increase in the number of first-page stories should be a genuine reflection of greater news importance.

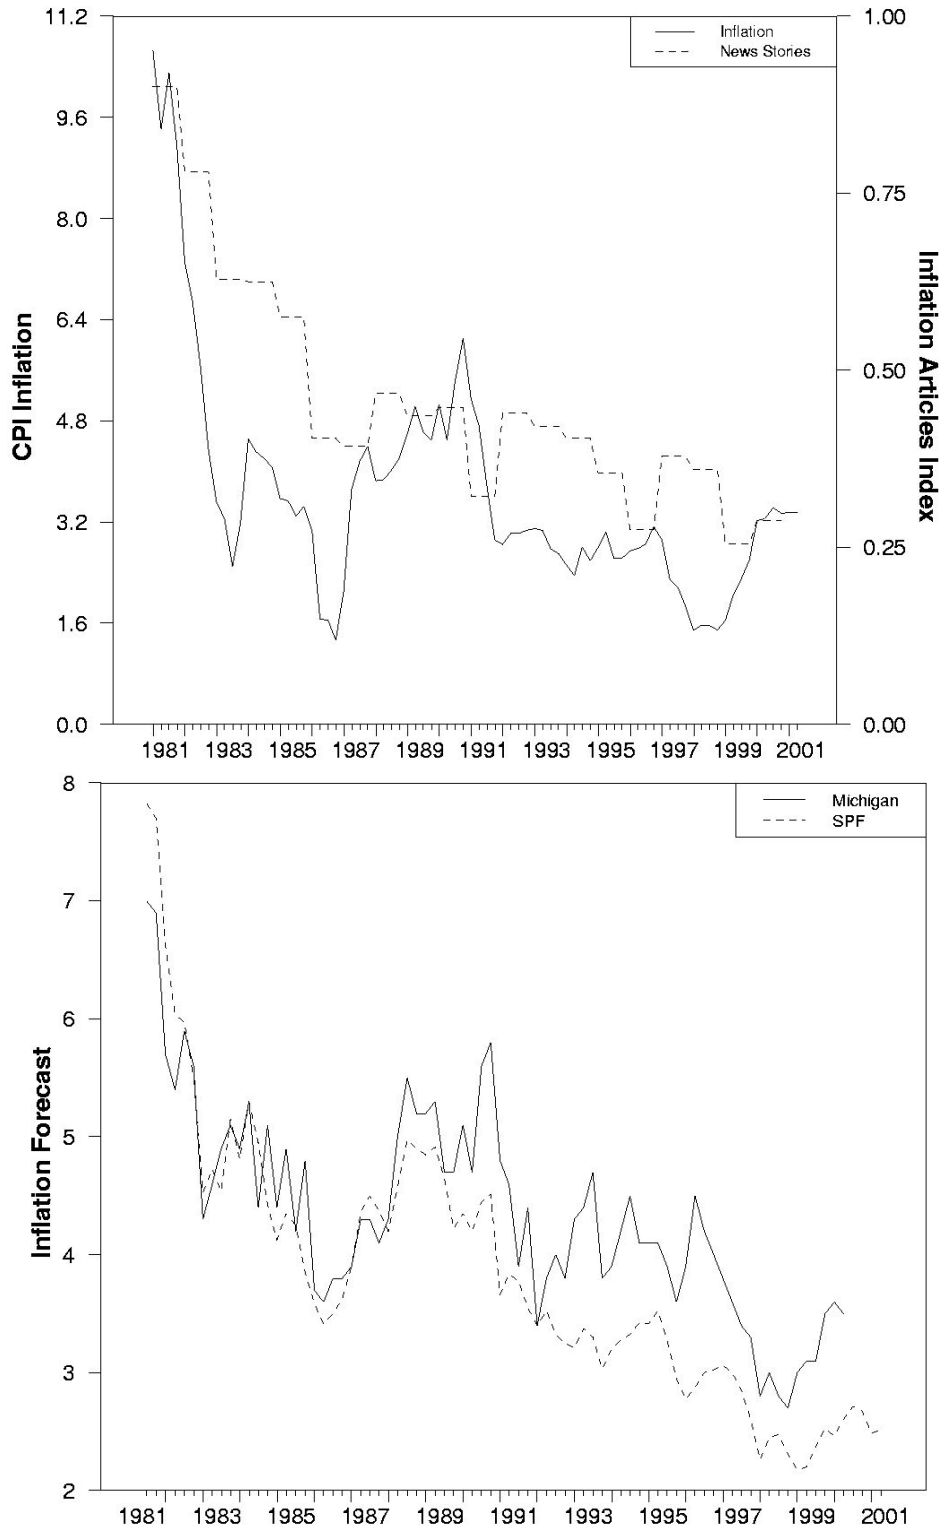


Figure 1: Inflation Versus News Stories, and Michigan Versus SPF Forecasts
21

$$\text{Estimating } \text{GAPSQ}_t = \alpha_0 + \alpha_1 \text{NEWS}_t$$

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

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 ‘infection’ is more likely when there is more coverage. The proposition that the infection rate is higher when there are more news stories can also be tested directly. Table 5 presents estimation results comparing the infection rate estimated during periods when there is more news coverage than average ($\text{NEWS}_t > \text{mean}(\text{NEWS})$) and less coverage than average ($\text{NEWS}_t < \text{mean}(\text{NEWS})$). The estimate of the infection rate 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).

Eqn	Sample	λ	Durbin-Watson	Q-Test p-value
1	All obs	0.273 (0.066)***	2.12	0.971
2	NEWS _t > mean(NEWS)	0.699 (0.176)***	1.57	0.216
3	NEWS _t < 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 5: Updating Speed Is Faster When There Is More News Coverage

There are two 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 (23), 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 epidemiological interpretation of the model proposed here.

4 Unemployment Expectations

If the epidemiological model of expectations 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

$$\bar{U}_{t,t+4} - \bar{U}_{t-4,t} = \gamma_0 + \gamma_1 M_t^U \quad (30)$$

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 inflation rate can be constructed from

$$\hat{U}_{t,t+4} = \hat{\gamma}_0 + \hat{\gamma}_1 M_t^U + \bar{U}_{t-4,t}. \quad (31)$$

When (30) is estimated, the coefficient on M_t^U is 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 6 presents a set of regression results for the household unemployment forecast that is essentially identical to the tests performed in Table 3 for inflation expectations.

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

Eqn	α_0	α_1	α_2	α_3	\bar{R}^2	Durbin-Watson	StdErr	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	$\alpha_1 + \alpha_2 = 1$ 0.109
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 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.

Table 6: Estimating and Testing the Baseline Model for Unemployment

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 epidemiological model performs even better in explaining unemployment expectations than 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.

Other extensions are also possible. The Michigan survey contains questions about future income, interest rates, and other important macro variables. In addition, another monthly survey of consumers conducted by the Conference Board has a wide variety of other forward-looking questions. It would be interesting to see whether the epidemiological model of expectations is broadly applicable, with a parameter value of around 0.25, for many of these variables.

5 Agent Based Models of Inflation Expectations

One of the most fruitful trends in empirical macroeconomics over the last fifteen years has been the growing effort to construct microfoundations that can be tested by examining microeconomic data. Broadly speaking, the goal is to find empirically sensible models for the behavior of the individual agents (people, firms, banks), which can then be aggregated to derive implications about macroeconomic dynamics. Separately, but in a similar spirit, researchers at the Santa Fe Institute, the CSED, and elsewhere have been exploring ‘agent-based’ models that examine the complex behavior that can sometimes emerge from the interactions between collections of simple agents.

One of the primary attractions of an agent-based or microfoundations approach to modeling household expectations is the prospect of being able to test the model using large microeconomic datasets. However, to my knowledge only two existing research papers have examined the raw household-level survey data underlying the University of Michigan’s aggregate expectations index, both by Nicholas Souleles (2002, 2000). For present purposes, the more interesting of these is Souleles (2002), which demonstrates (among other things) that there are highly statistically significant differences across demographic groups in forecasts of aggregate economic variables like the inflation rate. Clearly, in a world where everyone’s expectations were purely rational, there should be no demographic differences in expectations about the inflation rate.

An agent-based version of the framework proposed above could in principle account for such demographic differences. The simplest approach would be to assume that there are differences across demographic groups in the propensity to pay attention to economic news (different λ ’s); it is even conceivable that one could calibrate these differences using existing facts about the demographics of newspaper readership (or CNBC viewership).

Without access to the underlying micro data it is difficult to tell whether demographic heterogeneity in λ would be enough to explain Souleles’s findings about sys-

tematic demographic differences in macro expectations. Even without the raw micro data, however, an agent-based model has considerable utility. In particular, an agent-based approach permits us to examine the consequences of relaxing some of the model's assumptions to see how robust its predictions are. Given our hypothesis that Souleles's results on demographic differences in expectations might be due to differences in λ across groups, the most important application of the agent-based approach is to determining the consequences of heterogeneity in λ .

5.1 Heterogeneity in λ

Consider a model in which there are two categories of people, each of which makes up half the population, but with different newspaper-reading propensities, λ_1 and λ_2 .

For each group it will be possible to derive an equation like (23),

$$M_{i,t}[\pi_{t,t+4}] = \lambda_i N_t[\pi_{t,t+4}] + (1 - \lambda_i) M_{i,t-1}[\pi_{t-1,t+3}]. \quad (32)$$

But note that (dropping the π arguments for simplicity) aggregate expectations will just be the population-weighted sum of expectations for each group,

$$M_t = (M_{1,t} + M_{2,t})/2 \quad (33)$$

$$= \left(\frac{\lambda_1 + \lambda_2}{2} \right) N_t + ((1 - \lambda_1) M_{1,t-1} + (1 - \lambda_2) M_{2,t-1})/2 \quad (34)$$

Replace $M_{1,t-1}$ by $M_{t-1} + (M_{1,t-1} - M_{t-1})$ and similarly for $M_{2,t-1}$ to obtain

$$M_t = \left(\frac{\lambda_1 + \lambda_2}{2} \right) N_t + \left(1 - \left(\frac{\lambda_1 + \lambda_2}{2} \right) \right) M_{t-1} + \overbrace{\left(\frac{M_{1,t-1} + M_{2,t-1}}{2} \right)}^{=0} - M_{t-1} - \left(\frac{\lambda_1(M_{1,t-1} - M_{t-1}) + \lambda_2(M_{2,t-1} - M_{t-1})}{2} \right) \quad (35)$$

$$M_t = \hat{\lambda} N_t + (1 - \hat{\lambda}) M_{t-1} - \left(\frac{\lambda_1(M_{1,t-1} - M_{t-1}) + \lambda_2(M_{2,t-1} - M_{t-1})}{2} \right) \quad (36)$$

where $\hat{\lambda} = (\lambda_1 + \lambda_2)/2$.

Thus, the dynamics of aggregate inflation expectations with heterogeneity in λ have a component $\hat{\lambda} N_t + (1 - \hat{\lambda}) M_{t-1}$ that behaves just like a version of the model when everybody has the same λ equal to the average value in the population, plus a term (in big parentheses in (36)) that depends on the joint distribution of λ 's and the deviation by group of the difference between the previous period's rational forecast and the group's forecast.

Now consider estimating the baseline equation

$$M_t = \lambda N_t + (1 - \lambda)M_{t-1} \quad (37)$$

on a population with heterogeneous λ 's. The coefficient estimates will be biased in a way that depends on the correlations of N_t , M_{t-1} and M_t with the last term in equation (36), $\left(\frac{\lambda_1(M_{1,t-1}-M_{t-1})+\lambda_2(M_{2,t-1}-M_{t-1})}{2}\right)$. There is no analytical way to determine the magnitude or nature of the bias without making a specific assumption about the time series process for N_t , and even with such an assumption all that could be obtained is an expected asymptotic bias. The bias in any particular small sample would depend on the specific history of N_t in that sample.

The only sensible way to evaluate whether the bias problem is likely to be large given the actual history of inflation and inflation forecasts in the US is to simulate a model with households who have heterogeneous λ 's and to estimate the baseline equation on aggregate statistics generated by that sample.

Specifically, the experiment is as follows. A population of P agents is created, indexed by i ; each of them begins by drawing a value of λ_i from a uniform distribution on the interval $(\underline{\lambda}, \bar{\lambda})$. In an initial period 0, each agent is endowed with an initial value of $M_{i,0} = 2$ percent. Thus the population mean value $M_0 = (1/P) \sum_{i=1}^P M_{i,0} = 2$. For period 1, each agent draws a random variable distributed on the interval $[0, 1]$. If that draw is less than or equal to the agent's λ_i , the agent updates $M_{i,1} = N_1$ where N_1 is taken to be the 'Newspaper' forecast of the next year's inflation rate in period t ; if the random draw is less than λ_i the agent's $M_{i,1} = M_{i,0}$. The population-average value of M_1 is calculated, and the simulation then proceeds to the next period.

For the simulations, the 'news' series N_t is chosen as the concatenation of 1) the actual inflation rate from 1960q1 to 1981q2 and 2) the SPF forecast of inflation from 1981q3 to 2001q2. Then regression equations corresponding to (37) are estimated on the subsample corresponding to the empirical subsample, 1981q3 to 2001q2. Thus, the simulation results should indicate the dynamics of M_t that would have been observed if actual newspaper forecasts of inflation had been a random walk until 1981q2 and then had tracked the SPF once the SPF data began to be published.

The results of estimating (23) on the data generated by this simulation when the population is $P = 250,000$ are presented in Table 7. For comparison, and to verify that the simulation programs are working properly, equation (1) presents results when all agents' λ 's are exogenously set to 0.25. As expected, the simulation returns an estimate of $\lambda = 0.25$, and the equation fits so precisely that there are essentially no residuals.

The remaining rows of the table present the results in the case where λ values are heterogeneous in the population. The second row presents the most extreme example, $[\underline{\lambda}, \bar{\lambda}] = [0.00, 0.50]$. Fortunately, even in this extreme case the regression yields an estimate of the speed-of-adjustment parameter λ that, at around 0.26, is still quite

$$\text{Estimating } M_t[\pi_{t,t+4}] = \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_{t-1}[\pi_{t-1,t+3}] + \epsilon_t$$

λ Range	α_1	α_2	\bar{R}^2	Durbin- Watson	StdErr
[0.25,0.25]	0.250 (0.000)***	0.750 (0.000)***	1.000	2.58	0.002
[0.00,0.50]	0.261 (0.010)***	0.746 (0.009)***	0.999	0.10	0.039
[0.20,0.30]	0.249 (0.000)***	0.751 (0.000)***	1.000	2.20	0.002
[0.15,0.35]	0.245 (0.001)***	0.755 (0.001)***	1.000	0.50	0.006

$M_t[\pi_{t,t+4}]$ is mean inflation expectations in quarter t , $N_t[\pi_{t,t+4}]$ is the news signal corresponding to the SPF mean inflation forecast after 1981q3 and the previous year inflation rate before 1981q3. All equations are estimated over the period 1981q3 to 2001q2.

Table 7: Estimating the Baseline Model on Simulated Data with Heterogeneous λ s

close to the true average value 0.25 in the population. Interestingly, however, one consequence of the heterogeneity in λ is that there is now a very large amount of serial correlation in the residuals of the equation; the Durbin-Watson statistic indicates that this serial correlation is positive and a Q test shows it to be highly statistically significant.

Heterogeneous λ 's induce serial correlation primarily because the views of people with λ 's below $\bar{\lambda}$ are slow to change. For example, if the 'rational' forecast is highly serially correlated, an agent with a λ close to zero will be expected to make errors of the same size and direction for many periods in a row before finally updating.

The comparison of the high serial correlation that emerges from this simulation to the low serial correlation that emerged in the empirical estimation in Table 3 suggests that heterogeneity in λ is probably not as great as the assumed uniform distribution between 0.0 and 0.5. Results are therefore presented for a third experiment, in which λ 's are uniformly distributed between 0.2 and 0.3. Estimation on the simulated data from this experiment yields an estimate of λ very close to 0.25 and a Durbin-Watson statistic that indicates much less serial correlation than emerged with the broad $[0, 0.5]$ range of possible λ 's. Finally, the last row presents results when λ is uniformly distributed over the interval $[0.15, 0.35]$. This case is intermediate: the estimate of λ is still close to 0.25, but the Durbin-Watson statistic now begins to indicate substantial serial correlation.

On the whole, the simulation results suggest that the serial correlation properties of the empirical data are consistent with a moderate degree of heterogeneity in λ , but not with extreme heterogeneity. It is important to point out, however, that empirical data contain a degree of measurement and sampling error that is absent in the simulated data. To the extent that these sources can be thought of as white noise, they should bias the Durbin-Watson statistic up in comparison to the ‘true’ Durbin-Watson, so the scope for heterogeneity in λ is probably considerably larger than would be indicated by a simple comparison of the measured and simulated Durbin-Watson statistics. Thus the serial correlation results should not be taken as very serious evidence against substantial heterogeneity in λ .

A few last words on serial correlation. The important point in Mankiw and Reis (2001), as well as in work by Ball (2000) and others, is that the presence of some people whose expectations are not fully and instantaneously forward-looking profoundly changes the behavior of macro models. Thus, the discussion of serial correlation has an importance here beyond its usual econometric ramifications for standard errors and inference. If there are some consumers whose expectations are very slow to update, they may be primarily responsible for important deviations between the rational expectations model and macroeconomic reality.

5.2 Matching the Standard Deviation of Inflation Expectations

Thus far all our tests of the model have been based on its predictions for behavior of mean inflation expectations. Of course, the model also generates predictions for other statistics like the standard deviation of expectations across households at a point in time. Some households will have expectations that correspond to the most recent inflation forecast, while others will have expectations that are out of date by varying amounts. One prediction of the model is that (for a constant λ) if SPF inflation forecasts have remained stable for a long time, the standard deviation of expectations across households should be low, while if there have been substantial recent changes in the rational forecast of inflation we should expect to see more cross-section variability in households’ expectations.

This is testable. Curtin (1996) reports average values for the standard deviation for the Michigan survey’s inflation expectations over the period from 1978 to 1995; results are plotted as the solid line in figure 2. It is true that the empirical standard deviation was higher in the early 1980s, a time when inflation rates and SPF inflation expectations changed rapidly over the course of a few years, than later when the inflation rate was lower and more stable.

The short and long dashed loci in the figure depict the predictions of the homogeneous $\lambda = 0.25$ and heterogeneous $\lambda \in [0.0, 0.5]$ versions of the agent-based model.

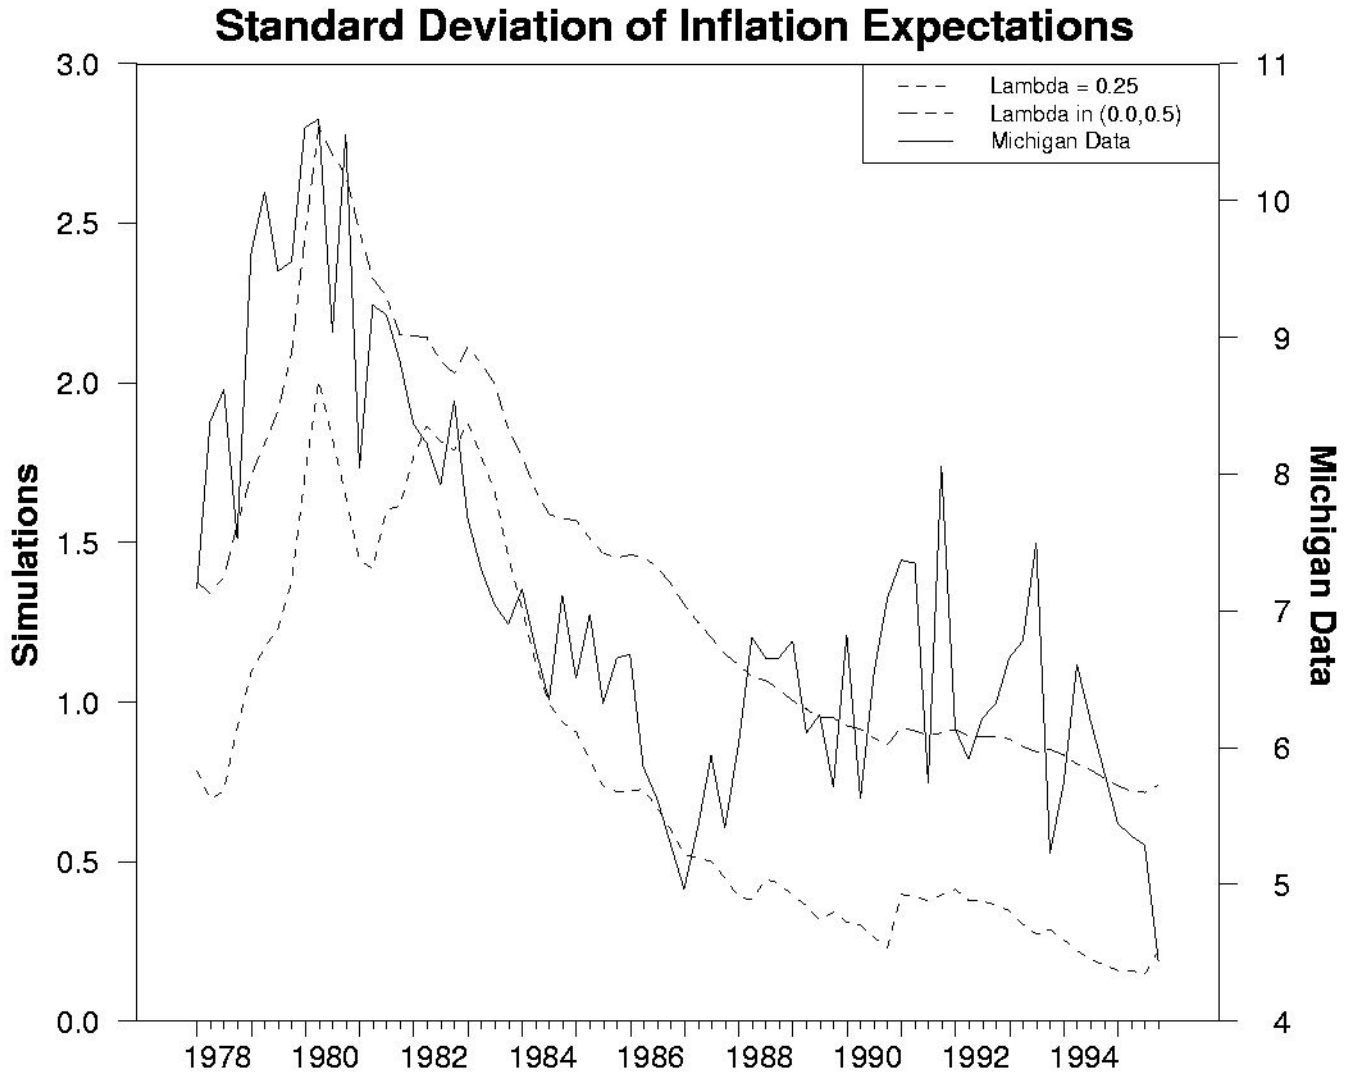


Figure 2: Standard Deviation of Inflation Expectations from Data and Simulations

There is considerable similarity between the time paths of the actual and simulated standard deviations: The standard deviation is greatest for both simulated and actual data in the late 1970s and early 1980s, because that is the period when the levels of both actual and expected inflation changed the most. In both simulated and real data the standard deviation falls gradually over time, but shows an uptick around the 1990 recession and recovery before returning to its downward path.

However, the *levels* of the standard deviations are very different between the simulations and the data; the scale for the Michigan data on the right axis ranges from 4 to 11, while the scale for the simulated standard deviations on the left axis ranges from 0 to 3. Over the entire sample period, the standard deviation of household inflation expectations is about 6.5 in the real data, compared to only about 0.5 in the simulated data.

Curtin (1996) analyzes the sources of the large standard deviation in inflation expectations across households. He finds that part of the extreme variability is attributable to small numbers of households with very extreme views of inflation. Curtin's interpretation is that these households are probably just ill-informed, and he proposes a variety of other ways to extract the data's central tendency that are intended to be robust to the presence of these extreme outlying households. However, even Curtin's preferred measure of dispersion in inflation expectations, the size of the range from the 25th to the 75th percentile in expectations, has an average span of almost 5 percentage points over the 81q3-95q4 period, much greater than would be produced by any of the simulation models considered above.¹³

The first observation to make about the extreme variability of household inflation expectations is that such variability calls into question almost all standard models of wage setting in which well-informed workers demand nominal wage increases in line with a rational expectation about the future inflation rate.¹⁴ If a large fraction of workers have views about the future inflation rate that are a long way from rational, it is hard to believe that those views have much impact on the wage-setting process. Perhaps it is possible to construct a model in which equilibrium is determined by average inflation expectations, with individual variations making little or no difference

¹³Curtin advocates use of the median rather than the mean as the summary statistic for 'typical' inflation expectations. However, the epidemiological model has simple analytical predictions for the mean but not the median of household expectations, so the empirical work in this paper uses the mean.

¹⁴The only prominent exception I am aware of is the two papers by Akerlof, Dickens, and Perry (1996, 2000) mentioned briefly above. In these models workers do not bother to learn about the inflation rate unless it is sufficiently high to make the research worthwhile. However such a model would presumably imply a modest upper bound to inflation expectation errors, since people who suspected the inflation rate was very high would have the incentive to learn the truth. In fact, Curtin (1996) finds that the most problematic feature of the empirical data is the small number of households with wildly implausibly high forecasts.

to individual wages. Constructing such a model is beyond the scope of this paper; but whether or not such a model is proposed, it seems likely that any thorough understanding of the relation between inflation expectations in the aggregate and actual inflation will need a model of how individuals' inflation expectations are determined.

The simplest method of generating extra individual variability in expectations is to assume that when people encounter a news report on inflation, the process of committing the associated inflation forecast to memory is error-prone.¹⁵

To be specific, suppose that whenever an agent encounters a news report and updates his expectations, the actual expectation stored in memory is given by the expectation printed in the news report times a mean-one lognormally distributed storage error. Since the errors average out in the population as a whole, this assumption generates dynamics of aggregate inflation expectations that are identical to those of the baseline model. Figure 3 plots the predictions for the standard deviation of inflation expectations across households of the baseline $\lambda = 0.25$ model with a lognormally distributed error with a standard error of 0.5. The figure shows that the change in the standard deviation of inflation residuals over time is very similar in the model and in the data, but the level of the standard deviation is still considerably smaller in the model. This could of course be rectified by including an additive error in addition to the multiplicative error. Such a proposed solution could be tested by examining more detailed information on the structure of expectations at the household level like that examined by Souleles (2002).

5.3 Social Transmission of Inflation Expectations

As noted above, the standard model of disease transmission is one in which illness is transmitted by person-to-person contact. Analogously, it is likely that some people's views about inflation are formed by conversations with others rather than by direct contact with news reports. For the purposes of this paper the most important question is whether the simple formula (23) would do a reasonably good job in capturing the dynamics of inflation expectations even when social transmission occurs.

Simulation of an agent-based model with both modes of transmission is straightforward. The extended model works as follows. In each period, every person has a probability λ of obtaining the latest forecast by reading a news story. Among the $(1 - \lambda)$ who do not encounter the news source, the algorithm is as follows. For each person i , there is some probability p that he will have a conversation about inflation with a randomly-selected other person j in the population. If j has an inflation forecast that is of more recent vintage than i 's forecast, then i adopts j 's forecast, and vice-versa.¹⁶

¹⁵Alternatively, one could assume that retrieval from memory is error-prone. The implications are very similar but not identical.

¹⁶This rules out the possibility that the less-recent forecast would be adopted by the person with

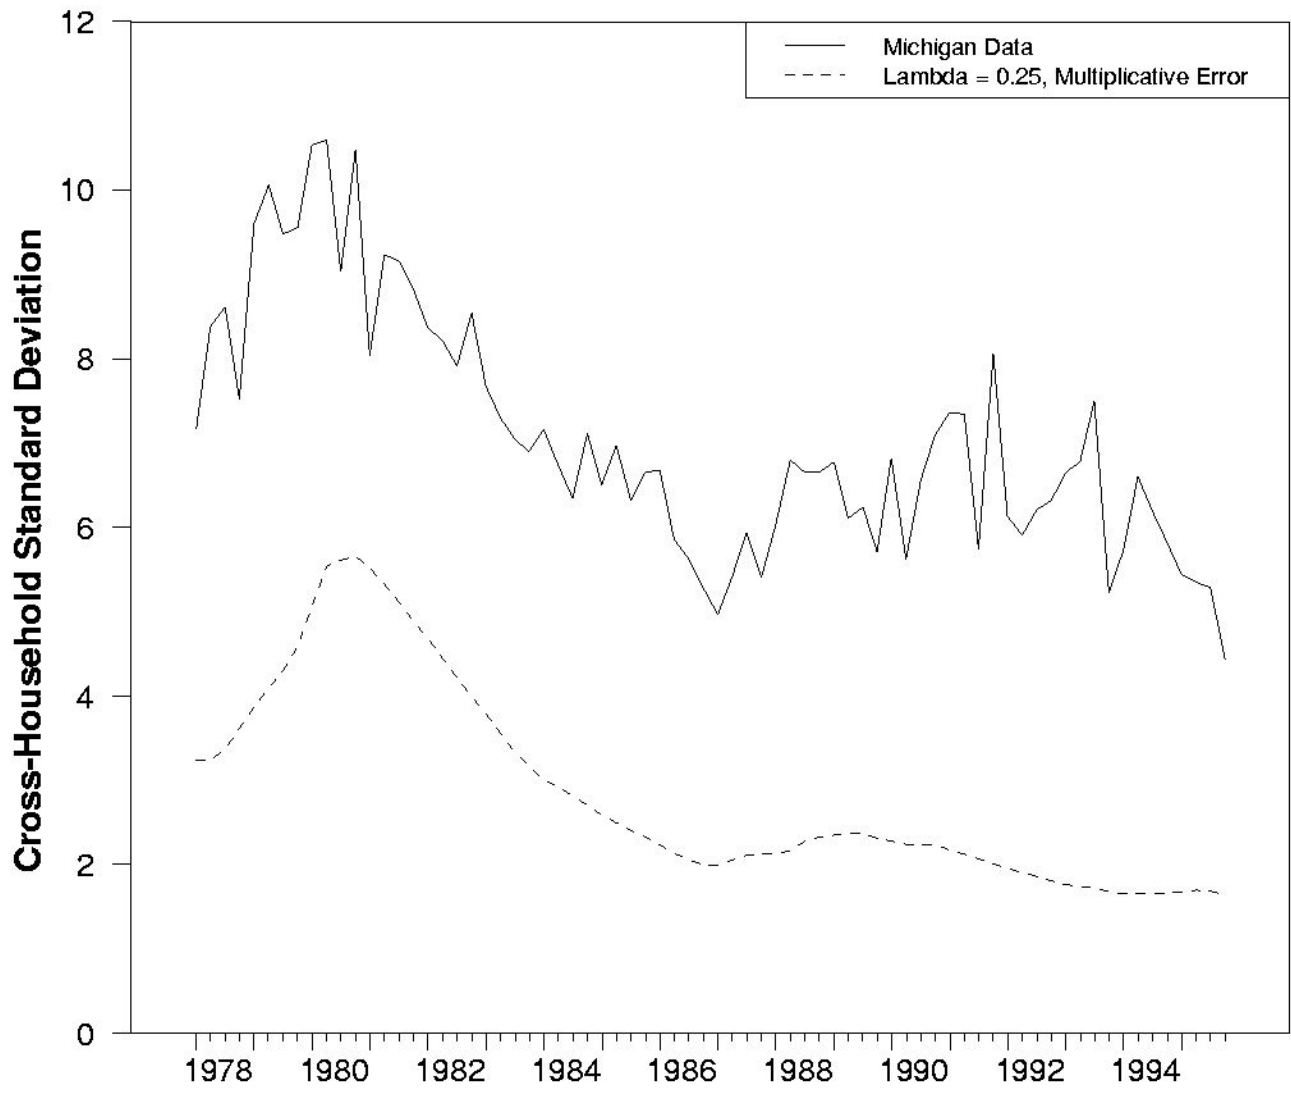


Figure 3: Standard Deviation of Inflation Expectations from Simulation with Memory Errors

$$\text{Estimating } M_t = \alpha_0 + \alpha_1 S_t + \alpha_2 M_{t-1} + \epsilon_t$$

Prob. of Social Exchange	α_0	α_1	α_2	\bar{R}^2	Durbin-Watson	StdErr	Test
							p-value
$p = 0.25$		0.311 (0.003) ^{***}	0.689 (0.003) ^{***}	1.000	2.19	0.020	$\alpha_1 + \alpha_2 = 1$ 0.1010
	0.008 (0.008)	0.302 (0.006) ^{***}	0.694 (0.006) ^{***}	1.000	2.07	0.020	$\alpha_0 = 0$ 0.3149
$p = 0.10$		0.276 (0.001) ^{***}	0.724 (0.001) ^{***}	1.000	2.05	0.009	$\alpha_1 + \alpha_2 = 1$ 0.1099
	0.002 (0.004)	0.273 (0.002) ^{***}	0.726 (0.002) ^{***}	1.000	1.95	0.008	$\alpha_0 = 0$ 0.5321

M_t is the mean value of inflation expectations across all agents in the simulated population; S_t is the actual annual inflation rate from 1960q1 to 1981q2, and the SPF inflation forecast from 1981q3 to 2000q2. Estimation is restricted to the simulation periods corresponding to 1981q3 to 2000q2 for which actual SPF data 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.

Table 8: Estimating Baseline Model on Random Mixing Simulations

Table 8 presents results of estimating equation (23) on the aggregate inflation expectations data that result from this agent-based simulation under a uniform fixed $\lambda = 0.25$ probability of news-reading. The first two rows present results when the probability of a social transmission event is $p = 0.25$. The primary effect of social transmission is to bias upward the estimated speed of adjustment term. The point estimate is about 0.31, or about 6 percentage points too high. However, the \bar{R}^2 of the equation is virtually 100 percent, indicating that even when there is social transmission of information, the common-source model does an excellent job of explaining the dynamics of aggregate expectations. The next row shows the results when the rate of social transmission is $p = 0.10$. Unsurprisingly, the size of the bias in the estimate of λ is substantially smaller in this case, and the model continues to perform well in an \bar{R}^2 sense.

A potential objection to these simulations is that they assume ‘random mixing.’ That is, every member of the population is equally likely to encounter any other member. Much of the literature on agent-based models has examined the behavior of populations that are distributed over a landscape in which most interactions occur

a more-recent information. The reason to rule this out is that if there were no directional bias (more recent forecasts push out older ones), the swapping of information would not change the distribution of forecasts in the population and therefore would not result in aggregate dynamics any different from those when no social communication is allowed.

between adjacent locations on the landscape. Often models with local but no global interaction yield quite different outcomes from ‘random mixing’ models.

To explore a model in which social communication occurs locally but not globally, I constructed a population distributed over a two dimensional lattice, of size 500x500, with one agent at each lattice point. I assumed that a fraction η of agents are ‘well informed’ - that is, as soon as a new inflation forecast is released, these agents learn the new forecast with zero lag. Other agents in the population obtain their views of inflation solely through interaction with neighbors.¹⁷ Thus, in this model, news travels out in concentric patterns (one step on the landscape per period) from its geographical origination points (the news agents, who are scattered randomly across the landscape). As in the random mixing model, I assume that new news drives out old news.

$$\text{Estimating } M_t = \alpha_0 + \alpha_1 S_t + \alpha_2 M_{t-1} + \epsilon_t$$

Up-to-date Agents	α_0	α_1	α_2	\bar{R}^2	Durbin-Watson	StdErr	Test
							p-value
$\eta = 0.25$		0.223 (0.025)***	0.708 (0.035)***	0.992	0.10	0.135	$\alpha_1 + \alpha_2 = 1$ 0.0000
	0.393 (0.003)***	0.306 (0.001)***	0.505 (0.001)***	1.000	0.91	0.009	$\alpha_0 = 0$ 0.0000
$\eta = 0.15$		0.098 (0.017)***	0.854 (0.027)***	0.988	0.13	0.116	$\alpha_1 + \alpha_2 = 1$ 0.0000
	0.473 (0.005)***	0.183 (0.001)***	0.589 (0.001)***	1.000	1.03	0.009	$\alpha_0 = 0$ 0.0000

M_t is the mean value of inflation expectations across all agents in the simulated population; S_t is the actual annual inflation rate from 1960q1 to 1981q2, and the SPF inflation forecast from 1981q3 to 2000q2. Estimation is restricted to the simulation periods corresponding to 1981q3 to 2000q2 for which actual SPF data 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.

Table 9: Estimating Baseline Model on Local Interactions Simulations

Results from estimating the baseline model on data produced by the ‘local interactions’ simulations are presented in table 9. For comparability with the baseline estimate of $\lambda = 0.25$ in the common-source model, I have assumed that proportion $\eta = 0.25$ of the agents in the new model are the ‘well-informed’ types whose inflationary expectations are always up to date. Interestingly, estimating the baseline model yields a

¹⁷For the purposes of the simulation, an agent’s neighbors are the agents in the eight cells surrounding him. For agents at the borders of the grid, neighborhoods are assumed to wrap around to the opposite side of the grid; implicitly this assumes the agents live on a torus.

coefficient of about $\alpha_1 = 0.22$ on the SPF forecast, even though 25 percent of agents always have expectations exactly equal to the SPF forecast. The coefficient on lagged expectations gets a value of about 0.71, and the last column indicates that a test of the proposition that $\alpha_1 + \alpha_2 = 1$ now rejects strongly. However, the regression still has an \bar{R}^2 of around 0.99, so the basic common-source model still does an excellent job of capturing the dynamics of aggregate inflation expectations.

The most interesting result, however, is shown in the next row: The estimation now finds a highly statistically significant role for a nonnegligible constant term. Recall that the only real empirical problem with the common-source model was that the estimation found a statistically significant role for a constant term.

Results in the next rows show what happens when the proportion of news agents is reduced to $\eta = 0.15$. As expected, the estimate of α_1 falls; indeed, the downward bias is now even more pronounced than with 25 percent well-informed. However, when a constant is allowed into the equation, the constant term itself is highly significant and the estimate of α_1 jumps to about 0.18, not far from the fraction of news agents in the population.

What these simulation results suggest is that the empirical constant term may somehow be reflecting the fact that some transmission of inflation expectations is through social exchange rather than directly through the news media. Furthermore, and happily, it is clear from the structure of the local interactions model that this population would eventually learn the true correct expectation of inflation if the SPF forecasts permanently settled down to a nonstochastic steady-state. Thus it is considerably more appealing to argue that the constant term reflects misspecification of the model (by leaving out social interactions) than to accept the presence of a true constant term (and its associated implication of permanent bias).

It is tempting to view the social learning simulations as a bit of a sideshow to the main thrust of this paper, which is about the surprisingly good fit of the common source epidemiological model. However, it is worth repeating the central lesson of Mankiw and Reis (2001) and others: The extent to which inflation can be reduced without increasing unemployment depends upon the speed with which a new view of inflation can be communicated to the *entire* population. It is not at all clear that the predictions about the medium-term inflation/unemployment tradeoff of a model with social transmission of expectations, or even of the common-source model with heterogeneous λ 's, are similar to the predictions of the homogeneous λ model examined by Mankiw and Reis (2001). Investigating these questions should be an interesting project for future research.

6 Conclusions

This paper has three main points.

The first is that it is high time, more than 25 years after the ‘rational expectations revolution’ and 65 years after Keynes’s emphasis on the centrality of expectations, that the examination of empirical data on expectations became a central part of macroeconomics. While there have been a few important prior contributions (particularly by Roberts (1997, 1998)), and a modest literature on consumer sentiment and consumption expenditures (see, e.g., Carroll, Fuhrer, and Wilcox (1994)), the bulk of the macroeconomics profession has ignored the rich empirical data available on actual household and business expectations in favor of the theoretical purity of rational expectations models.

The second point is that the abandonment of rational expectations need not lead macroeconomics into an atheoretical wilderness where the Lucas critique lurks behind every equation. There is an existing body of theory, in epidemiology as well as in the ‘small worlds’ literature, that can be directly applied to explain expectations data. This paper has shown that for inflation expectations and unemployment expectations, a remarkably simple epidemiological model does an excellent job in explaining the deviations of mean household expectations from a rational forecast.

The final point is that if we want to have macroeconomic models that are built on really secure microfoundations, and are therefore as immune as can reasonably be hoped from the Lucas critique, there is no substitute for building the model from the ground up, and comparing its predictions to whatever microeconomic data are available. Whether the enterprise is described as a search for the ‘microfoundations’ of macroeconomics or as ‘agent-based’ modeling, it seems reasonable to suppose that in the long run no macro model will be counted on as a reliable guide to macro behavior if its predictions for the expectations of individual actors bear little empirical resemblance to the observable expectations of actual individuals.

Appendix A: Timing Issues

This appendix explores implications of the very different sampling methodologies of the Michigan household survey and the Survey of Professional Forecasters.

As noted in the main text, timing issues suggest that when estimating the model with monthly data the timing problems are least serious when the sample is restricted to months t in which a Survey of Professional Forecasters was conducted and published. The appropriate estimating equation in this case is

$$M_{t+1}[\pi_{t+1,t+13}] = \lambda N_t[\pi_{t,t+12}] + (1 - \lambda)M_t[\pi_{t,t+12}]. \quad (38)$$

However, there is a problem in estimating this equation directly. The monthly Michigan data reflect surveys of only 500 households, and the estimated sampling error with such a small sample is no longer trivial enough to be ignored. Fortunately, Curtin (1996) has provided estimates of the magnitude of the sampling error. Curtin's results imply a typical sampling variance of about 0.09 per month, large enough to cause substantial bias in our estimate of λ if not corrected for. Fortunately, standard econometric formulas can be used to adjust parameter estimates when the variance of the sampling error is known.

The corrected point estimate of the speed-of-adjustment parameter is $1 - \lambda = 0.91$.¹⁸ Reassuringly, this estimate is right in the ballpark of what would have been expected if there were no data or reporting difficulties in the quarterly estimation procedure: The quarterly estimation of the baseline model produced a quarterly updating fraction of $(1 - \lambda) = 0.73$, which is statistically indistinguishable from the estimate of $0.92^3 \approx 0.78$ implied by the fact that a quarter contains three months. Thus, the procedure of estimating the model with the appropriate monthly data, being careful about measurement error, yields essentially the same answer as was obtained from quarterly data.

The preceding analysis glosses over a final problem: The Michigan survey data are collected on a continuous basis throughout a month, rather than all at once on the first day of the month. To see why this is potentially problematic, consider again a month t in which an SPF survey is conducted. Since inflation statistics are reported in mid-month, roughly half of the households surveyed in month t will have had the opportunity to read the news articles published upon the release of the statistics at mid-month, while the other half will not. In this case I have not been able to derive a 'clean' equation like (28) for the dynamics of inflation expectations.

¹⁸The programs, dataset, and econometric derivations that generate this result are included in the downloadable files associated with this paper on my website; the econometric theory, and derivation of the estimated sampling variance of 0.09, are laid out in `AppendixA_MonthlyExp.pdf`, the RATS program that estimates the model is `AppendixA_MonthlyExp.pgm`, and the documentation of the program is in `AppendixA_MonthlyPgm.pdf`.

To investigate the importance of this problem, I conducted the following simulation analysis. First, I specified a daily-frequency stochastic process for the ‘rational’ forecast of the next year’s inflation as a random walk with a daily innovation such that the quarterly innovations in the inflation forecast would match the standard deviation of the actual quarterly innovations in the SPF forecast. I assumed that households update to this true forecast using a λ parameter such that $(1 - \lambda)^n = 0.75$ where n is the number of days in a quarter. I then picked out the value of the ‘rational’ forecast at the midpoint of each quarter, and set a variable equal to that forecast. Finally, I aggregated the daily data to quarterly frequency, and performed a regression of the form

$$M_t = \lambda S_t + (1 - \lambda)M_{t-1}$$

where M_t now represents the quarterly average of the simulated daily values of the constructed household forecast and S_t is the value of the ‘true’ forecast at the midpoint of the quarter. This exercise produced an estimated $\lambda = 0.234$ as compared with the ‘correct’ $\lambda = 0.25$. Thus, it appears that the mismatch in timing between the Michigan and SPF surveys is unlikely to cause much of a problem in estimating the λ parameter.¹⁹

¹⁹These simulations are conducted by the program `AppendixA_TimingSims.pgm`.

Appendix B: Construction of Index of News Coverage

The index of news stories on inflation 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.

References

- AKERLOF, GEORGE, WILLIAM DICKENS, AND GEORGE PERRY (1996): “The Macroeconomics of Low Inflation,” *Brookings Papers in Economic Activity*, 1996:1, 1–76.
- (2000): “Near-Rational Wage and Price Setting and the Long-Run Phillips Curve,” *Brookings Papers in Economic Activity*, 2000:1, 1–60.
- BALL, LAURENCE (1994): “What Determines the Sacrifice Ratio?,” in *Monetary Policy*, ed. by N. Gregory Mankiw, chap. 5. University of Chicago Press, Chicago.
- (2000): “Near-Rationality and Inflation in Two Monetary Regimes,” *NBER Working Paper No. W7988*.
- BARSKY, ROBERT B. (1987): “The Fisher Hypothesis and the Forecastability and Persistence of Inflation,” *Journal of Monetary Economics*, 19, 3–24.
- 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> .
- CARROLL, CHRISTOPHER D., AND WENDY E. DUNN (1997): “Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets,” in *NBER Macroeconomics Annual, 1997*, ed. by Benjamin S. Bernanke, and Julio Rotemberg, pp. 165–229. MIT Press, Cambridge, MA, <http://www.econ.jhu.edu/people/ccarroll/macroann.pdf>; Methodological Appendix: <http://www.econ.jhu.edu/people/ccarroll/methods3.pdf>; Empirical Results and Simulation Programs: <http://www.econ.jhu.edu/people/ccarroll/cdfiles.html>; .
- CARROLL, CHRISTOPHER D., JEFFREY C. FUHRER, AND DAVID W. WILCOX (1994): “Does Consumer Sentiment Forecast Household Spending? If So, Why?,” *American Economic Review*, 84(5), 1397–1408.
- CROUSHORE, DEAN (1993): “Introducing: The Survey of Professional Forecasters,” *Federal Reserve Bank of Philadelphia Business Review*, pp. pages 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*.
- FRASER, DW ET. AL. (1977): “Legionnaires’ disease: description of an epidemic of pneumonia,” *New England Journal of Medicine*, 297(22), 1189–97.
- KERMACK, W. O., AND A. G. MCKENDRICK (1927): “Contributions to the Mathematical Theory of Epidemics,” *Proceedings of the Royal Academy of Sciences A*, 115, 700–721.
- KEYNES, JOHN MAYNARD (1936): *The General Theory of Employment, Interest, and Money*. Harcourt, Brace.
- KREMER, MICHAEL (2000): “An Epidemiological Model of Unions,” *Manuscript, Harvard University*.
- 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 Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *NBER Working Paper Number 8290*.
- NEWBY, WHITNEY K., AND KENNETH D. WEST (1987): “A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- 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*.
- SOULELES, NICHOLAS (2002): “Consumer Sentiment: Its Rationality and Usefulness in Forecasting Expenditure? Evidence from the Michigan Micro Data,” *Journal of Money, Credit, and Banking*.
- SOULELES, NICHOLAS S. (2000): “Household Securities Purchases, Transactions Costs, and Hedging Motives,” *Manuscript, University of Pennsylvania*.

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.