

The Epidemiology of Macroeconomic Expectations

Christopher D. Carroll[†]

ccarroll@jhu.edu

April 15, 2003

ABSTRACT

Macroeconomists have long emphasized the importance of expectations in determining macroeconomic outcomes, and an enormous theoretical literature has developed examining many models of expectations formation. This paper proposes a new approach, based on epidemiological models, in which only a small set of agents (professional forecasters) formulate their own expectations, which then spread through the population via the news media in a manner analogous to the spread of a disease. The paper shows that the very simplest epidemiological model, called the ‘common source’ model, does a good job of explaining the dynamics of inflation and unemployment expectations, and more complicated epidemiological models produce dynamics similar to those that emerge from the common source model.

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

JEL Classification Codes: D84, E31

This paper was written in connection with the conference “The Economy As an Evolving Complex System III” at the Santa Fe Institute in November of 2001, in honor of Kenneth Arrow’s contributions to the Santa Fe Institute and to economics. A companion paper, “Macroeconomic Expectations of Households and Professional Forecasters,” presents empirical work that was part of an earlier draft of this paper; that paper was published in the *Quarterly Journal of Economics*, Vol. 118, Number 1 (February 2003), pp. 269–298.

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

The data, 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, Johns Hopkins University, Baltimore MD 21218-2685.

1 Introduction

Economists have long understood that macroeconomic outcomes depend critically upon agents' expectations. Keynes (1936) believed that economies could experience booms and busts that reflected movements in 'animal spirits' (a view that has some appeal at the current moment of dot-com hangover), but the basis for most of today's macro models is the rational expectations approach pioneered in the 1970s by Lucas, Sargent, Barro, and others. This approach makes a set of assumptions that are 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)). Partly in response to these problems, an emerging literature has been exploring models based on an assumption that agents engage in some form of data-based learning process to form their expectations; for surveys, see Sargent (1993) or Evans and Honkapohja (2001). But the rational expectations framework remains the dominant approach, partly because it tends to be mathematically more tractable than many proposed alternatives.

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, directly or indirectly. 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, in a way that resembles the spread of disease in a population, 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, 2002) have recently proposed aggregate expectations equations that are mathematically very similar to aggregate implications of the baseline epidemiological model of expectations proposed 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, 2002) 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, 2002) devote much effort to explaining *why* the dynamics of aggregate expectations should evolve as they proposed (though Roberts does suggest that his equation might result from the diffusion of press reports). Mankiw and Reis motivate their model by suggesting that there are costs either of obtaining or of processing inflation every time an agent updates; however, they do not provide an explicit information-costs or processing-costs microfoundation.

This paper provides an explicit microfoundation for a simple aggregate expectations equation, based on simple models of the spread of disease. Rather than tracking the spread of a disease through a population, the model tracks the spread of a piece of information (specifically, the latest rational forecast of inflation).

A companion paper, Carroll (2003), estimates the baseline model and finds that it does a good job at capturing the dynamics of household survey expectations about both inflation and unemployment. Furthermore, that paper shows that household inflation expectations are closer to the expectations of professional forecasters during periods when there is more news coverage of inflation, and that the speed with which household expectations adjust to professional expectations is faster when there is more news coverage.

After providing the epidemiological foundation for the model estimated in Carroll (2003) and summarizing the basic empirical results from that paper, this paper explores the implications of several extensions to the model, using the kinds of ‘agent-based’ simulation techniques pioneered at the Santa Fe Institute and at the Center for Social and Economic Dynamics (CSED) at the Brookings Institution.

The first extension is to allow heterogeneity in the extent to which different households pay attention to macroeconomic news. This version of the model is capable of generating demographic differences in macroeconomic expectations like those documented in Souleles (forthcoming), which are very hard to rationalize in a rational expectations model. Both this extension and the baseline version of the model are then simulated in order to derive implications for the standard deviation of inflation expectations across agents. When the simulated data are compared to the empirical data, the results are mixed. On one hand, the patterns over time of the empirical and the simulated standard deviations bear a strong resemblance, rising sharply in the late 1970s and early 1980s and then gradually falling off again. On the other hand, the *level* of the standard deviation in the empirical data is much higher than in the simulated data; however, I show that adding simple forms of memory error can make the model and the data match up reasonably well.

The next extension examines what happens when the model is generalized to a more standard epidemiological context: A ‘random mixing’ framework in which people can be ‘infected’ with updated inflation expectations by conversations with random other individuals in the population. It turns out that when the baseline framework that assumes infection only from news sources is estimated on simulated

data from the random mixing model, the baseline framework does an excellent job of capturing the dynamics of mean inflation expectations; this suggests that the ‘common source’ simplification is probably not too problematic.

The final extension is to a context in which people communicate only with near ‘neighbors’ in some social sense, rather than with random other individuals in the population. Simulation and estimation of the baseline model on this population finds that the baseline model again does a good job of capturing expectational dynamics; however, in one particular respect the results in this framework are a better match for the empirical results than is the baseline model.

The paper concludes with some general lessons and ideas for future research.

2 The Epidemiology of Expectations

2.1 The SIR Model

Epidemiologists have developed a rich set of models for the transmission of disease in a population. The general framework 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 the framework’s designation as the ‘SIR’ model).

The standard assumption 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.

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.¹

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 $t + 1$,

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

¹This subsection is largely drawn from Carroll (2003); however, that paper does not discuss the epidemiological interpretation of the derivations, as this derivation does.

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, 2002), except that in their model the updating agents construct their own rational forecast of the future course of the macroeconomy rather than learning about the experts’ forecast from the news media. The equation is also similar to a formulation estimated by Roberts (1997), except that Roberts uses past realizations of the inflation rate on the right hand side rather than past forecasts.

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. 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 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. Finally, 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

$$\begin{aligned}\pi_t &= \pi_t^f + \epsilon_t & (6) \\ \pi_{t+1}^f &= \pi_t^f + \eta_{t+1}, & (7) \\ \vdots & & \vdots\end{aligned}$$

where ϵ_t is a transitory shock to the inflation rate in period t while η_t is the permanent innovation in the fundamental inflation rate 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.³

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)-(7); 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)-(7) 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 that the core inflation rate follows a random walk with a one-period

³Note 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

transitory component - that is, it is not possible to reject the process defined by (6)-(7).⁴ 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. 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

⁴The 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.

an annual rate as

$$\pi_{t,t+1} = 4(\log p_{t+1} - \log p_t) \quad (8)$$

$$= 4\pi_{t+1} \quad (9)$$

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

Under this set of assumptions, Carroll (2003) shows that the process for inflation expectations can be rewritten as

$$M_t[\pi_{t,t+4}] = \lambda N_t[\pi_{t,t+4}] + (1 - \lambda)M_{t-1}[\pi_{t-1,t+3}]. \quad (10)$$

That is, mean measured inflation expectations for the next year should be a weighted average between the current 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.⁵

2.3 Estimates

Estimating (10) empirically requires the identification of empirical counterparts for household-level inflation expectations and newspaper inflation forecasts. Conveniently, the University of Michigan’s Survey Research Center has been asking households about their inflation expectations for well over thirty years. To be precise, households are first asked

“During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are right now?”

and then those who say “go up” (the vast majority) are asked

“By about what percent do you expect prices to go up, on the average, during the next 12 months?”

The Survey Research Center uses the answers to these questions to construct an index of mean inflation expectations, which is an almost exact counterpart to the object required by the theory. (For details of index construction, see Curtin (1996)).

Measuring the forecasts that people are assumed to encounter in the news media is a thornier problem. But typical newspaper articles on inflation tend to quote professional forecasters, and so it seems reasonable to use the mean forecast from the Survey of Professional Forecasters (SPF) as a proxy for what the news media are reporting.

⁵This 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).

Carroll (2003) estimates equation (10) using the Michigan survey index for M_t and the SPF for N_t . Results are reproduced in the upper panel of table 2 (where N_t changes to S_t to indicate the use of the SPF).

Table 2:
Estimating and Testing the Baseline Model

$$\text{Estimating Equation } M_t[\bullet_{t,t+4}] = \alpha_0 + \alpha_1 S_t[\bullet_{t,t+4}] + \alpha_2 M_{t-1}[\bullet_{t-1,t+3}] + \epsilon_t$$

Eqn	α_0	α_1	α_2	\bar{R}^2	Durbin- Watson	StdErr	Test
							p-value
Estimating Baseline Model on Inflation Expectations							
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
Estimating Baseline Model on Unemployment Expectations							
Memo:	6.27 (0.31)***			0.00	0.07	1.25	$\alpha_0 = 0$ 0.000
1		0.30 (0.07)***	0.69 (0.07)***	0.95	1.59	0.28	$\alpha_1 + \alpha_2 = 1$ 0.036
2		0.30 (0.07)***	0.70 (0.07)***	0.94	1.50	0.29	$\alpha_1 = 0.25$ 0.476
3	-0.03 (0.19)	0.30 (0.08)***	0.69 (0.07)***	0.95	1.60	0.29	$\alpha_0 = 0$ 0.890

$M_t[\bullet_{t,t+4}]$ is the Michigan household survey measure of mean inflation expectations or projected unemployment expectations in quarter t , $S_t[\bullet_{t,t+4}]$ is the Survey of Professional Forecasters mean inflation or unemployment forecast over the next year. Inflation equations are estimated over the period 1981q3 to 2000q2 for which both Michigan and SPF inflation forecasts are available; Unemployment equations are estimated over the period 1978q1 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.

The first line of the table ('Memo:') 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 is reserved for reporting the results of various tests; for example, the test performed in the 'Memo:' equation is for whether the

average value of the expectations index is zero, $\alpha_0 = 0$, which is rejected at a very high level of statistical significance, as indicated by a p-value of zero.

Equation 1 in the table reflects the estimation of an equation of the form

$$M_t[\pi_{t,t+4}] = \alpha_1 S_t[\pi_{t,t+4}] + \alpha_2 M_t[\pi_{t-1,t+3}] + \nu_t. \quad (11)$$

Comparing this to (10) provides the testable restriction that $\alpha_2 = 1 - \alpha_1$ or, equivalently,

$$\alpha_1 + \alpha_2 = 1. \quad (12)$$

The point estimates in equation 1 of $\alpha_1 = 0.36$ and $\alpha_2 = 0.66$ suggest that the restriction (12) is very close to holding true, and the last column confirms that the proposition is easily accepted by the data (the p -value is far above the usual critical level of 0.05 which would signal a rejection).

Estimation results when the restriction (12) is imposed in estimation are presented in the next row of the table, yielding our central estimate of the model's main parameter: $\lambda = 0.27$. This point estimate is remarkably close to the value of 0.25 assumed by Mankiw and Reis (2001, 2002) in their simulation experiments; unsurprisingly, the test reported in the last column for equation 2 indicates that the proposition $\alpha_1 = \lambda = 0.25$ is easily accepted by the data. Thus, the model implies 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.

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 (1995, 1997) and Mankiw and Reis (2001, 2002) 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.

Equation 3 in the table reports some bad news for the model: When a constant term is permitted in the regression equation, it turns out to be highly statistically significant. The model (10) did not imply the presence of a constant term, so this is somewhat disappointing. On the other hand, despite being statistically significant the constant term does not improve the fit of the equation much: The standard error of the estimate only declines from 0.43 to 0.35. Furthermore, a version of the model with a constant term cannot be plausibly interpreted as a true 'structural' model of the expectations process, since it implies that even if the inflation rate were to go to zero forever, and all forecasters were to begin forecasting zero inflation forever, households would never catch on. A more plausible interpretation of the positive constant term is that it may reflect some form of misspecification of the model. Below, I will present simulation results showing that if expectations are transmitted

from person to person in addition to through the news media, and an equation of the form of (10) is estimated on the data generated by the modified model, the regression equation returns a positive and statistically significant constant term; thus, the presence of the constant term can be interpreted as evidence that the simple ‘common-source’ epidemiological model postulated here may be a bit too simple.

The bottom panel of the table presents results for estimating an equation for unemployment expectations that is parallel to the estimate for inflation expectations.⁶ Equation 2 of this panel presents the version of the model that restricts the coefficients to sum to 1; the point estimate of the fraction of updaters is $\lambda = 0.31$, but this estimate is not significantly different from the estimate of $\lambda = 0.27$ obtained for inflation expectations or from the $\lambda = 0.25$ postulated by Mankiw and Reis (2001, 2002). The last equation shows that the unemployment expectations equation does not particularly want an intercept term, so the model actually fits better for unemployment expectations than for inflation expectations.

In sum, it seems fair to say that the simple ‘common-source’ epidemiological equation (10) does a remarkably good job of capturing much of the predictable behavior of both the inflation and the unemployment expectations indexes.

3 Agent Based Models of Inflation Expectations

One of the most fruitful trends in empirical macroeconomics over the last fifteen years has been the effort to construct rigorous microfoundations for macroeconomic models. 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 microfounded approach to modeling macroeconomic behavior is the prospect of being able to test a model using large microeconomic datasets. This is an opportunity that has largely been neglected so far in the area of expectations formation; I have found only three existing research papers that have examined the raw household-level expectations data from the Michigan survey. Two are by Nicholas Souleles (2000, forthcoming). For present purposes, the more interesting of these is Souleles (forthcoming), which demonstrates (among other things) that there are highly statistically significant differences across demographic groups in forecasts of several macroeconomic variables. Clearly, in a world where everyone’s expectations were purely rational, there should

⁶Some data construction was necessary to do this, because the Michigan survey does not ask people directly what their expectations are for the level of the unemployment rate, but instead asks whether they think the unemployment rate will rise or fall over the next year. See Carroll (2003) for details about how the unemployment expectations data are constructed.

be no demographic differences in such expectations.

An agent-based version of the epidemiological model 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 systematic 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 .⁷

3.1 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 (10),

$$M_{i,t}[\pi_{t,t+4}] = \lambda_i N_t[\pi_{t,t+4}] + (1 - \lambda_i) M_{i,t-1}[\pi_{t-1,t+3}]. \quad (13)$$

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

$$M_t = (M_{1,t} + M_{2,t})/2 \quad (14)$$

$$= \left(\frac{\lambda_1 + \lambda_2}{2} \right) N_t + ((1 - \lambda_1) M_{1,t-1} + (1 - \lambda_2) M_{2,t-1})/2 \quad (15)$$

⁷The final paper I know of that examines the micro data underlying the Michigan inflation expectations index is by Branch (2001), who proposes an interesting model in which individual consumers dynamically choose between competing models for predicting inflation, but are subject to idiosyncratic errors. He finds evidence that people tend to switch toward whichever model has recently produced the lowest mean squared error in its forecasts. This interesting approach deserves further study.

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

$$M_t = \left(\frac{\lambda_1 + \lambda_2}{2}\right) N_t + \left(1 - \left(\frac{\lambda_1 + \lambda_2}{2}\right)\right) M_{t-1} + \overbrace{\left(\frac{M_{1,t-1} + M_{2,t-1}}{2}\right)}^{=0} - M_{t-1} - \left(\frac{\lambda_1(M_{1,t-1} - M_{t-1}) + \lambda_2(M_{2,t-1} - M_{t-1})}{2}\right) \quad (16)$$

$$M_t = \hat{\lambda} N_t + (1 - \hat{\lambda}) M_{t-1} - \left(\frac{\lambda_1(M_{1,t-1} - M_{t-1}) + \lambda_2(M_{2,t-1} - M_{t-1})}{2}\right) \quad (17)$$

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 (17)) 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

$$M_t = \lambda N_t + (1 - \lambda) M_{t-1} \quad (18)$$

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 (17), $\left(\frac{\lambda_1(M_{1,t-1} - M_{t-1}) + \lambda_2(M_{2,t-1} - M_{t-1})}{2}\right)$. 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

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

λ Range	α_1	α_2	\bar{R}^2	Durbin- Watson	StdErr
[0.25,0.25]	0.250 (0.001)***	0.750 (0.001)***	1.000	1.94	0.006
[0.00,0.50]	0.265 (0.010)***	0.743 (0.009)***	0.999	0.11	0.039
[0.20,0.30]	0.249 (0.001)***	0.751 (0.001)***	1.000	2.03	0.005
[0.15,0.35]	0.244 (0.002)***	0.756 (0.002)***	1.000	0.64	0.009

$M_t[\pi_{t,t+4}]$ 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 3: Estimating the Baseline Model on Simulated Data with Heterogeneous λ s

1981q3 to 2001q2. Then regression equations corresponding to (18) 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 (10) on the data generated by this simulation when the population is $P = 250,000$ are presented in Table 3. 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 case the regression yields an estimate of the speed-of-adjustment parameter λ that, at around 0.26, is still quite 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 after a shock to the fundamental inflation rate, until 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 2 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. Furthermore, since the error term is very tiny (the standard errors are less than 4/100 of one percentage point), so any serial correlation properties it has cannot be of much econometric consequence. 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, 2002), 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 possibility of heterogeneity in λ , and the resulting serial correlation in errors, 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.

3.2 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




Figure 1: Standard Deviation of Inflation Expectations from Data and Simulations

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 1. 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. 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 high 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 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.⁸

The first observation to make about the excessive cross-section 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.⁹ 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 to individual wages. Constructing such

⁸Curtin 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.

⁹The 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.

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.¹⁰

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 2 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 (forthcoming).

3.3 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 (10) 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.¹¹

¹⁰Alternatively, one could assume that retrieval from memory is error-prone. The implications are very similar but not identical.

¹¹This rules out the possibility that the less-recent forecast would be adopted by the person with 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




Figure 2: 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	2.26	0.020	$\alpha_1 + \alpha_2 = 1$ 0.1357
	0.009 (0.009)	0.303 (0.006) ^{***}	0.694 (0.006) ^{***}	1.000	2.15	0.020	$\alpha_0 = 0$ 0.2939
$p = 0.10$		0.276 (0.001) ^{***}	0.724 (0.001) ^{***}	1.000	1.69	0.009	$\alpha_1 + \alpha_2 = 1$ 0.2172
	0.000 (0.004)	0.274 (0.003) ^{***}	0.725 (0.003) ^{***}	1.000	1.66	0.009	$\alpha_0 = 0$ 0.9006

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 4: Estimating Baseline Model on Random Mixing Simulations

Table 4 presents results of estimating equation (10) 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 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

distribution of forecasts in the population and therefore would not result in aggregate dynamics any different from those when no social communication is allowed.

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.¹² 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
$\eta = 0.25$		0.234 (0.025)***	0.696 (0.034)***	0.992	0.10	0.135	$\alpha_1 + \alpha_2 = 1$ 0.0000
	0.386 (0.003)***	0.315 (0.001)***	0.499 (0.001)***	1.000	1.03	0.007	$\alpha_0 = 0$ 0.0000
$\eta = 0.15$		0.098 (0.017)***	0.853 (0.027)***	0.988	0.13	0.117	$\alpha_1 + \alpha_2 = 1$ 0.0000
	0.475 (0.004)***	0.185 (0.001)***	0.587 (0.001)***	1.000	1.15	0.008	$\alpha_0 = 0$ 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 5: 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 5. 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 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.

¹²For 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.

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 always-up-to-date 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).

A final caveat is in order. The central lesson of Mankiw and Reis (2001, 2002) and others is that 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, 2002). Investigating this question should be an interesting project for future research.

4 Conclusions

This paper was written to provide a specific example of a more general proposition: That many of the puzzles confronting standard macroeconomic models today could be resolved by abandoning the mathematically elegant but patently false assumptions of rational expectations models and replacing them with more realistic and explicit models of how people obtain their ideas about economic topics, involving some form of learning or social transmission of knowledge and information. While the paper confines itself to presenting results from agent-based simulations of such a model of inflation expectaitons, the closely related work by Mankiw and Reis (2001, 2002) shows that macroeconomic dynamics are much more plausible when expectations are governed by a model like the ones explored here.

Other puzzles that might yield to such an approach are legion. For example, excess smoothness in aggregate consumption (relative to the rational expectations

benchmark) may reflect precisely the same kind of inattention posited for inflation expectations here (I am actively pursuing this possibility in ongoing work). A plausible explanation for the equity premium puzzle might be to suppose that it has taken a long time for news of the favorable risk/return tradeoff of stocks to spread from experts like Mehra and Prescott (1985) to the general population. The strong systematic relationship of productivity growth and the natural rate of unemployment documented, for example, by Staiger, Stock, and Watson (2001) and Ball and Moffitt (2001) may reflect workers' imperfect knowledge about productivity growth (and the slow social transmission of such information). The detailed dynamics of productivity itself can surely be captured better by models in which new technologies spread gradually in a population than by models in which new technologies instantaneously boost productivity upon the date of invention (which is the conventional 'technology shock' approach in rational expectations models). And a substantial literature now exists arguing that social transmission of information in a population of investing agents may be able to explain the excess volatility of asset prices compared to the rational expectations benchmark (see LeBaron (forthcoming) for a summary).

References

- Akerlof, George, William Dickens, and George Perry: “The Macroeconomics of Low Inflation,” *Brookings Papers in Economic Activity*, (1996:1), 1–76.
- Akerlof, George, William Dickens, and George Perry “Near-Rational Wage and Price Setting and the Long-Run Phillips Curve,” *Brookings Papers in Economic Activity*, (2000:1), 1–60.
- Ball, Laurence “What Determines the Sacrifice Ratio?,” in *Monetary Policy*, N. Gregory Mankiw, ed., chap. 5. (Chicago:, University of Chicago Press 1994) .
- (2000): “Near-Rationality and Inflation in Two Monetary Regimes,” *NBER Working Paper No. W7988*.
- Ball, Laurence, and Robert Moffitt (2001): “Productivity Growth and the Phillips Curve,” *Johns Hopkins University Department of Economics Working Paper Number 450*.
- Barsky, Robert B. (1987): “The Fisher Hypothesis and the Forecastability and Persistence of Inflation,” *Journal of Monetary Economics*, 19, 3–24.
- Branch, William A. (2001): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Manuscript, Department of Economics, College of William and Mary*.
- Calvo, Guillermo A. (1983): “Staggered Contracts in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–98.
- Carroll, Christopher D. (2003): “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, 118(1), 269–298, <http://www.econ.jhu.edu/people/ccarroll/epidemiologyQJE.pdf>.
- Curtin, Richard T. (1996): “Procedure to Estimate Price Expectations,” *Manuscript, University of Michigan Survey Research Center*.
- Evans, George W., and Seppo Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- Fraser, DW et. al. (1977): “Legionnaires’ disease: description of an epidemic of pneumonia,” *New England Journal of Medicine*, 297(22), 1189–97.
- Keynes, John Maynard (1936): *The General Theory of Employment, Interest, and Money*. Harcourt, Brace.
- LeBaron, Blake (forthcoming): “Agent Based Computational Finance: Suggested Readings and Early Research,” *Journal of Economic Dynamics and Control*.

- Mankiw, N. Gregory (2001): “The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment,” *Economic Journal*, 111(471), C45–C61.
- Mankiw, N. Gregory, and Ricardo Reis (2001): “Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps,” *NBER Working Paper Number 8614*.
- (2002): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117(4), 1295–1328.
- Mehra, Rajnish, and Edward C. Prescott (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15, 145–61.
- Roberts, John M. (1995): “New Keynesian Economics and the Phillips Curve,” *Journal of Money, Credit, and Banking*, 27(4), 975–984.
- (1997): “Is Inflation Sticky?,” *Journal of Monetary Economics*, pp. 173–196.
- (1998): “Inflation Expectations and the Transmission of Monetary Policy,” *Federal Reserve Board FEDS working paper Number 1998-43*.
- Sargent, Thomas J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.
- Souleles, Nicholas (forthcoming): “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*.