Long-run inflation expectations and monetary policy

Antulio N. Bomfim and Flint Brayton*

Introduction

Macroeconomic models are frequently used to simulate the transitional aggregate dynamics that are set into motion by a shift in monetary policy to alter the rate of inflation. A standard result is that the cost of lowering (or raising) the rate of inflation – the integral over time of the deviation of unemployment from its path in the absence of the policy change – varies with how quickly inflation expectations adjust. The more sluggish are expectations, the larger is the unemployment cost per unit of inflation change.

In the Federal Reserve Board's FRB/US macroeconomic model, expectations of long-run inflation play an important role in inflation dynamics. Several different simulation options for the formation of these expectations are available, and as described by Bomfim, Tetlow, von zur Muehlen and Williams (1997), the model's estimate of the unemployment "sacrifice ratio" associated with a change in inflation is affected significantly by the particular expectations mechanism selected. Up until now, however, there has been little empirical basis on which to decide how best to characterize the evolution of long-run inflation expectations. The purpose of this paper is to strengthen the empirical underpinning of this key part of the expectations mechanism in FRB/US by proposing and estimating simple learning rules for the determination of long-run inflation expectations.

Given that inflation in the long run is commonly regarded as a monetary phenomenon, it is natural to look for a connection between long-run inflation expectations and the conduct of monetary policy. Although one might search for evidence of revisions to expectations at times of announcements of policy changes, our prior is that participants in the economy are more likely to scrutinize policy actions more closely than announcements for evidence of a policy shift. Thus, we examine how well various models of learning empirically capture the speed with which long-run inflation expectations respond to a change in monetary policy.[1] The empirical results are then used to construct a version of the FRB/US model in which the expectations held by the private sector about monetary policy are specified as the outcome of learning in a stochastic environment.

A monetary policy regime is typically characterized as a policy reaction function whose structure and coefficients implicitly reflect long-term policy objectives and the speed with which deviations from targets are planned to be eliminated. Changes in policy objectives, including the speed of adjustment, alter the reaction function's coefficients. The models of learning that we examine – rolling regressions and Kalman filtering – yield "real-time" estimates of the coefficients of a posited reaction function. For each approach to learning, the time series of coefficient estimates provides a time series of perceived inflation targets.

---

* Helpful comments were provided by Franck Sedillot and participants of the Bank for International Settlements' 1998 Meeting of Central Bank Model Builders and Econometricians. We thank Steve Sumner for excellent research assistance. The opinions presented in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Board.

1 Our linking of long-run inflation expectations to perceptions of the conduct of monetary policy is not the only approach that has been used to characterize long-run inflation expectations. For alternative approaches, see Kozicki, Reifschneider and Tinsley (1996), who present a proxy based on a time-varying intercept in an estimated equation for the rate of inflation, and Kozicki and Tinsley (1996), who describe a measure derived from the term structure of interest rates under the assumption that the real rate of interest is stationary.

The first half of the paper compares available survey data on long-run inflation expectations with our constructed time series of real-time perceptions of the inflation objective of policy. Tentative regression evidence suggests the survey data on long-run inflation expectations is related more closely to our real-time learning constructions than to actual inflation. Nonetheless, the correspondence between the surveys and the constructions is modest, and the most important regressor is the lagged reading from the survey itself. The second half of the paper presents simulations of a version of FRB/US augmented to incorporate several of the real-time learning models. The simulations indicate that some of the learning approaches yield estimates of the unemployment sacrifice ratio that are in accord with the range of conventional estimates.

1. Survey measures of long-run inflation expectations

Survey data for the United States on long-run inflation expectations is sparse and available only since 1980. We use two series in our analysis:

- π^s_{mich} is the median inflation expectation over a 5 to 10-year horizon from the Survey Research Center at the University of Michigan. Survey respondents are a random sample of individuals. π^s_{mich} starts in 1980:Q1.[2]
- π^s_{h-p} is a measure of inflation expectations over a 10-year horizon spliced from two surveys. The first segment (1980 through mid-1991) is taken from Richard Hoey's "decision-makers" poll; subsequent observations are from the "Survey of Professional Forecasters" compiled by the Federal Reserve Bank of Philadelphia. π^s_{h-p} starts in 1980:Q3.[3]

Figure 1

Survey data on long-run inflation expectations

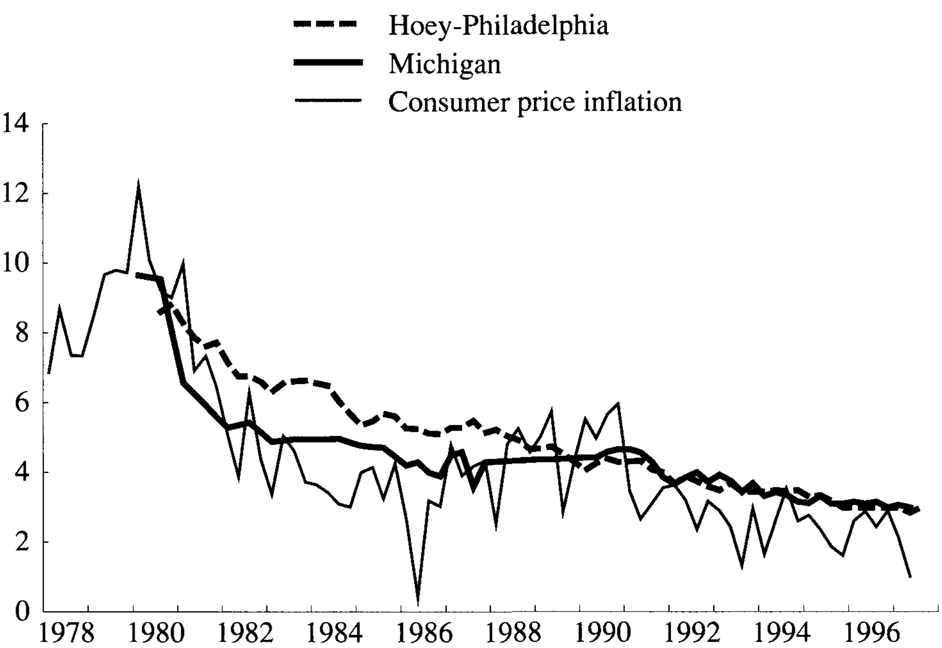

2 Only two observations per year are available from 1980 to 1985 and the series has a gap without observations from 1988:Q1 to 1990:Q1. Missing entries are interpolated linearly. Prior to 1980, a single observation exists for 1979:Q1.

3 Prior to 1980:Q3, the Hoey survey was also conducted in 1978:Q3 and 1979:Q1.

As shown in Figure 1, π^s_{mich} declines fairly rapidly in the early 1980s while the drop in π^s_{h-p} is more gradual. The two series converge by 1990 and subsequently edge down in tandem to 3% by 1996. The general consensus holds that monetary policy in the United States shifted in late-1979 to one aiming toward a substantial reduction in the rate of inflation. Neither survey shows a one-time drop in long-run inflation expectations in the immediate aftermath of the policy shift, although, admittedly, the fact that each survey only starts during 1980 makes this conclusion a bit tentative. A question we examine is whether the less-than-immediate response of the two expectations measures is better captured by a learning model in which policy changes become more apparent over time through observation of the changing relationship between the short-run policy instrument and macroeconomic conditions, or whether the survey expectations are simply adjusting to lower inflation as it emerges.

2. A simple model of monetary policy

We assume that historical US monetary policy can be (approximately) represented by an equation for the Federal funds rate in which the explanatory variables are lagged values of the Federal funds rate and current and lagged values of inflation and the deviation of the unemployment rate from an estimate of the natural rate. This specification is closely related to the policy rule proposed by Taylor (1993). While it may be that other macroeconomic or financial factors have influenced policy during certain periods, the posited relationship appears to capture much of the movement in the Federal funds rate since 1966, as long as some variation over time in its coefficients is permitted.

Our starting point is a general dynamic specification in which the Federal funds rate (i) depends on four quarterly lags of the funds rate and the current and first three lagged values of both consumer price inflation (π) and the unemployment gap ($\tilde{u}$):[4]

$$i_t = \alpha + \sum_{i=1}^{4} \beta_i i_{t-i} + \sum_{i=0}^{3} \gamma_i \pi_{t-i} + \sum_{i=0}^{3} \delta_i \tilde{u}_{t-i} \qquad (1)$$

Given a set of parameter estimates, the rate of inflation desired by policymakers (π^*) can be calculated as

$$\pi^* = \left(\alpha - \left(1 - \sum \beta\right) r^*\right) / \left(1 - \sum \beta - \sum \gamma\right) \qquad (2)$$

if it assumed that (i) the long-run real rate of interest (r^*) is a known constant and (ii) that the equilibrium nominal rate of interest moves one-for-one with equilibrium inflation. Note that the standard inflation stability condition associated with policy rules such as this – that the nominal funds rate change more than one-for-one with changes in inflation – is equivalent to the denominator of equation (2) being negative. If on the other hand, the denominator is zero and interest rates move one-for-one with inflation, monetary policy has no particular inflation target and accepts the current rate of inflation, whatever it is.

Not surprisingly, the coefficients of equation (1) are unstable over time. Figure 2 makes this point graphically. The series plotted are sequences of sub-sample tests for coefficient stability of a reaction function estimated from 1966:Q1 to 1996:Q4. The tests statistics, which are shown for every quarter in this span except the very beginning and end, are reported as ratios to the 5% critical value, and thus numbers greater than 1.0 represent rejections of stability at this significance level. Two test sequences are plotted, one for equation (1) and a second for a version of the reaction function whose dynamic structure has been simplified to eliminate insignificant regressors,

---

4 Simultaneity bias is not an issue in estimating this relationship if the reasonable assumption is made that inflation and unemployment are unaffected contemporaneously by the Federal funds rate.

$$i_t = \alpha + \sum_{i=1}^{3} \beta_i i_{t-i} + \gamma \sum_{i=0}^{3} (.25\pi_{t-i}) + \sum_{i=0}^{2} \delta_i \tilde{u}_{t-i} \tag{3}$$

From here on, equation (1) will be referred to as the "long-lags" reaction function and equation (3) as the "short-lags" variant.

Figure 2

Chow test sequences for coefficient stability

Ratio of test statistic to 5% critical value

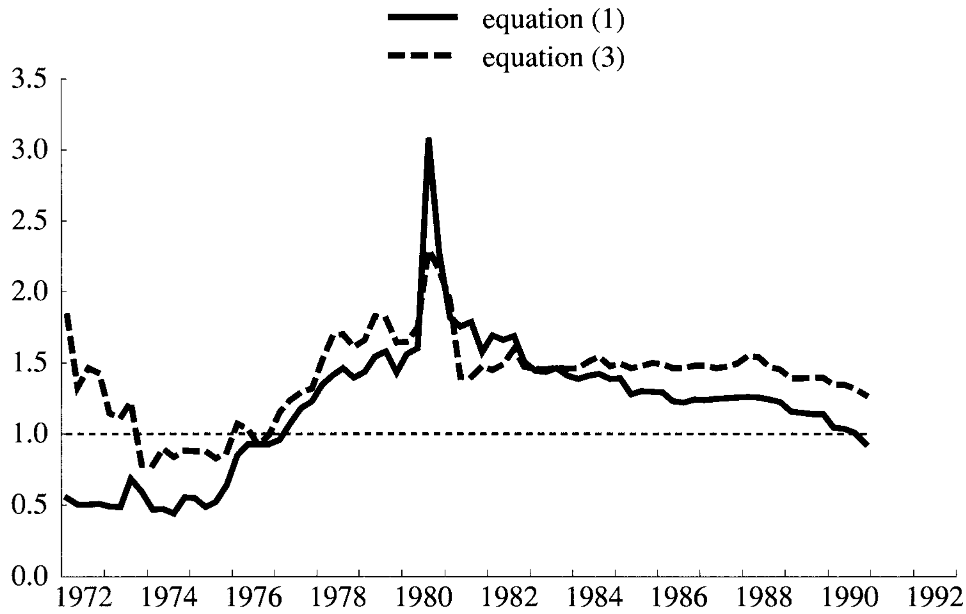

Although the figure seems to reveal that instability of the reaction function coefficients is pervasive, in fact simply permitting the constant to shift at the end of 1979 leads to a much more stable result.[5] The particular dating of the intercept shift was not chosen on any statistical grounds; rather, the selected switch-point conforms to the commonly held view that monetary policy changed at that time to one aiming to reduce the rate of inflation. Based on the formula given in (2) and the assumption that the real rate of interest is 2%, the "short-lags" specification with an intercept shift indicates that the target rate of inflation fell 4 percentage points from about 6½% in the period up

Table 1

Estimates of the target rate of inflation

Equation	1966-79			1980-96		
	value	95% range		value	95% range	
Long lags	6.44	5.37	7.54	2.75	1.42	3.81
Short lags	6.63	5.45	7.88	2.50	0.99	3.71

5 For the "long-lags" equation, the test statistic for structural change at 1980:Q1 has a p-value of 0.002 when the null hypothesis includes a constant intercept and a p-value of 0.10 when the null includes a shifting intercept. The corresponding p-values for the "short-lags" equation are 0.002 and 0.17.

through 1979 to 2½% since then.[6] The estimates from the "long-lags" version are similar. Confidence ranges around these values are fairly wide and, at the 95% level, encompass values more than 1 percentage point higher or lower than the point estimates.

3. Modeling long-run inflation expectations

We now turn to the question of what someone knowing the general form of the Federal funds rate reaction function could have deduced about policymakers' inflation objectives at different points of time. Hindsight enables the identification of a shift in the inflation target of policy at the end of 1979, but at the time sorting out exactly how policy was changing was undoubtedly difficult. For example, clear identification of the policy change as a lowering of the inflation objective, a more aggressive response to deviations of actual inflation from its target, or some combination of the two, was probably not possible immediately.

Three real-time approaches to estimating the policy reaction function are employed to construct time series of hypothetical perceptions of the policy objective for inflation. The first uses rolling regressions having an estimation interval (window) of fixed length, the second uses rolling regressions in which the estimation interval expands over time but data observations are given less weight as they recede from the end of each estimation period, and the third is the Kalman filter. In each case, the perceived inflation target for any particular quarter is calculated according to equation (2), using the real-time estimates of the reaction function coefficients for that date.

The Kalman filter is the optimal estimation approach when the reaction function coefficients are believed to vary over time as random walks. The first two are more ad hoc in design, though one can think of the optimal window length or decay rate for the rolling regressions as balancing the cost of slower identification of a policy shift as the window lengthens or the decay rate diminishes against the risk of falsely identifying a policy shift when the past is "forgotten" too quickly. The rolling regression approach with declining data importance weights shares one desirable feature with the Kalman filter: Each updates the reaction function coefficient vector in proportion to the gap between the observed value of the Federal funds rate and the value predicted on the basis of the prior estimate of the coefficients. No revision is made to the coefficient vector if there is no surprise to the funds rate.

3.1 Rolling regressions

Rolling regressions were estimated for a variety of window lengths and decay rates. Figure 3 shows the constructed perception of the inflation target derived from the rolling estimation of the "short lags" equation with a 15-year window. This window length yields a constructed series that matches the general pattern of the two inflation surveys somewhat better than do series based on other window lengths. The dotted lines in the figure represent a 95% confidence band around the rolling-regression estimate. Most of the observations from the surveys lie well inside these bands, with the only exceptions occurring at the beginning of the period when some of the high initial survey values lie above the confidence band. Initial values in 1980 of the constructed series, about 7% expected inflation, lie below the survey responses which range between 8 and 10%, and both surveys tend to fall more rapidly than does the constructed measure in the early 1980s. Note the confidence band is not shown for 1996 because it becomes very wide. As the high inflation years of the mid-1970s and

---

6 Our results are robust to variation in r^*. As can be seen in equation (2), the effect of r^* on π^* depends on the degree of interest-rate smoothing in the policy rule, which we measure as $\Sigma\beta$. In practice, we find that there is substantial smoothing – the sum of the estimated β_i is close to 1 – and the effect of r^* on π^* is rather small.

very early 1980s gradually fall out of the rolling regression sample, the estimates of the coefficients of the divisor in the formula for the constructed inflation target become less precise.

Figure 3

Perceived inflation objective: rolling regressions with 15-year window

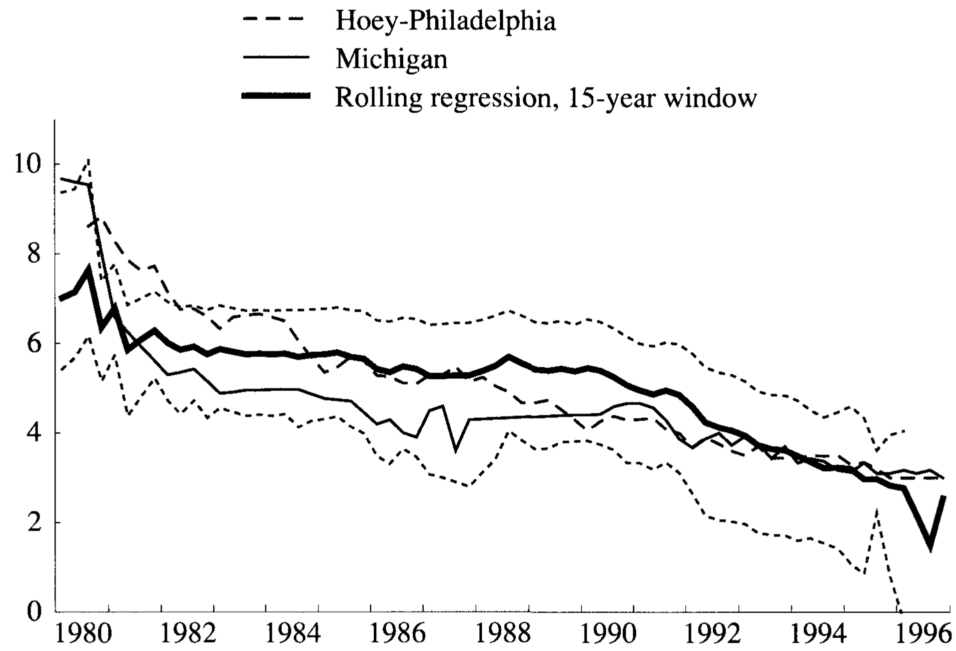

Figure 4

Perceived inflation objective: rolling regressions with 2% decay per quarter

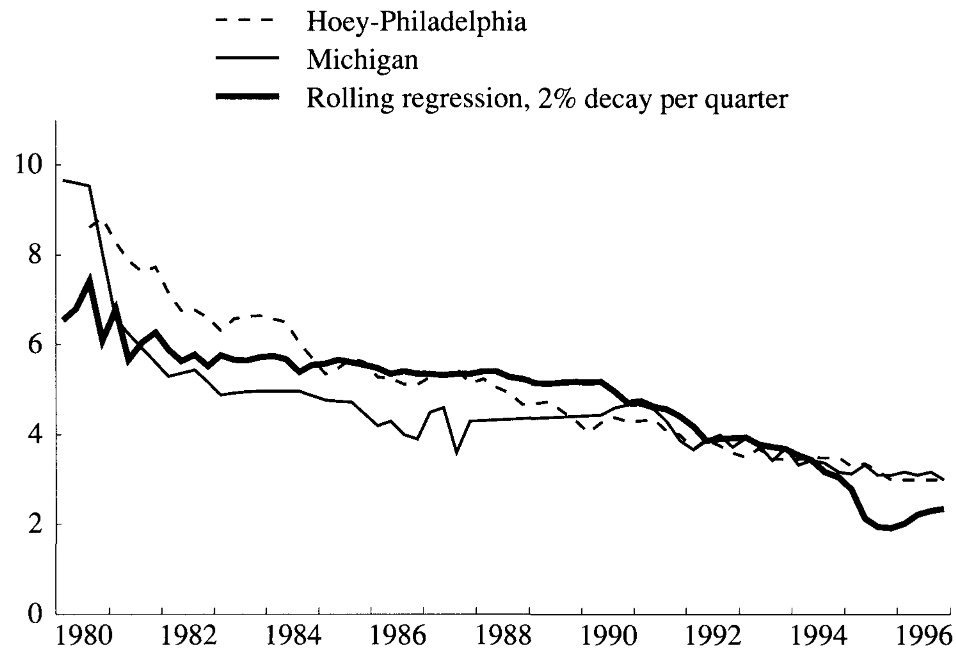

An alternative to the fixed-window regression approach is one in which the estimation sample expands over time but observations are given relatively less weight are they recede into the past. Figure 4 shows a constructed measure of perceived long-run inflation derived under the assumption that the data importance weights decline 2% per quarter. The properties of this series are generally similar to those of the series based on the rolling regression with the 15-year window.

3.2 Kalman filter

To illustrate the Kalman filter approach, we start with a simple Taylor-like policy function

$$i_t = \theta_1 i_{t-1} + \theta_2 \left(\bar{\pi}_t - \pi_t^*\right) + \theta_3 \tilde{u}_t + \theta_4 \tilde{u}_{t-1} + \theta_5 \left(r^* + \bar{\pi}_t\right) + e_t \tag{4}$$

where $\bar{\pi}_t$ is a 4-quarter moving average of inflation, e_t represents i.i.d. shocks to the reaction function, and $\theta_1 + \theta_5 = 1$.[7]

Consider now a framework where the private sector knows the functional form of the reaction function, but not its coefficients or the potentially time-varying inflation target. Agents use a recursive least-squares algorithm to estimate the θ_i parameters and assume that the unobserved inflation target follows a random walk.

$$\pi_t^* = \pi_{t-1}^* + \varepsilon_t \tag{5}$$

where we assume that ε_t is white-noise and uncorrelated with e_t. This specification of the reaction function differs from the one used for the rolling regression approaches in that only one parameter of the policy rule, the inflation target, is explicitly assumed to vary stochastically.

Figure 5

Perceived inflation objective: Kalman filter

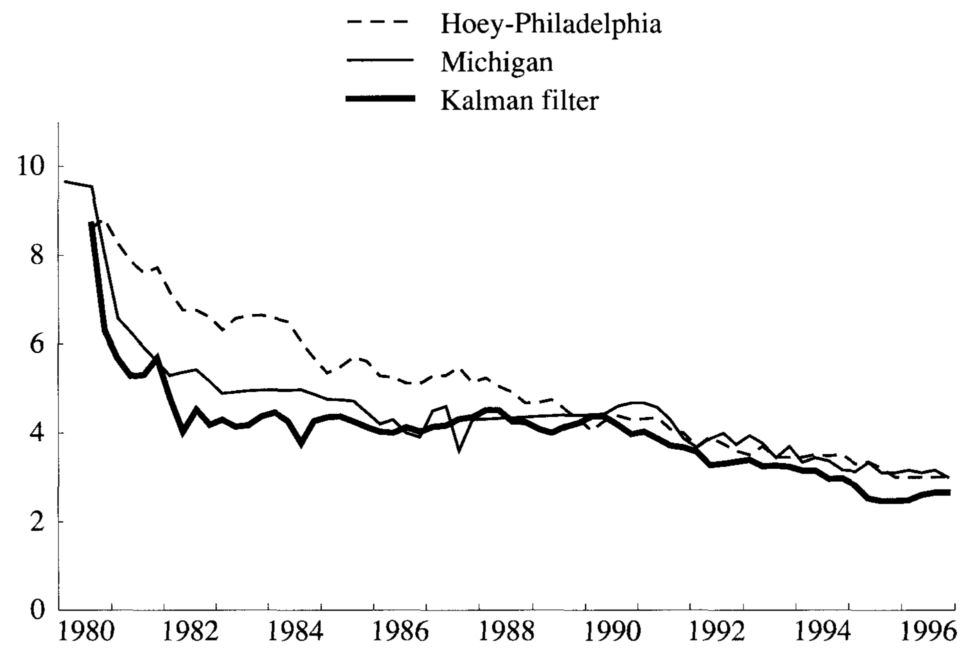

The private sector's learning problem can be summarized by the following state-space form:

[7] It is straightforward to see that equation (4) can be mapped into equation (3).

$$i_t = x_t' \Gamma_t + e_t \tag{6}$$

$$\Gamma_t = \Gamma_{t-1} + \eta_t \tag{7}$$

where $x_t \equiv [1, i_{t-1}, \bar{\pi}_t, \tilde{u}_t, \tilde{u}_{t-1}]'$, $\Gamma_t \equiv [\theta_5 r^* - \theta_2 \pi_t^*, 1-\theta_5, \theta_2 + \theta_5, \theta_3, \theta_4]'$, and η_t has zeroes everywhere, except for its first entry. Thus, given equations (6) and (7), agents update their estimates of the policy parameters (θ_i) and the inflation target (π_t^*) as each new quarter of data becomes available. The results are summarized in Figure 5. The thick solid line in the figure corresponds to the private sector's perceived inflation target under standard Kalman filter (KF) learning. The constructed series is broadly consistent with the survey data, tracking the Michigan series particularly well. In contrast to the learning mechanisms based on rolling regressions, the generated series drops quite rapidly in the early eighties, from about 9% in 1980 to near 4% in 1982 – this decline is comparable to the one registered by the Michigan survey, but faster than suggested by the Hoey-Philadelphia data. Turning to more recent readings, the standard KF-based learning algorithm places long-run inflation expectations at about 2½% in early 1997, about 50 basis points below both surveys.[8]

3.3 Are the models consistent with the surveys?

Simple regressions are used to characterize more formally the relationship between the inflation surveys and the constructed series,

$$\pi_t^s = \alpha_0 + \alpha_1 \pi_{t-1}^s + \alpha_2 \pi_t^c + \alpha_3 \pi_t , \tag{8}$$

in which one of the inflation surveys (π^s) is regressed on a constant, its own lag, a constructed target inflation perception (π^c), and actual inflation (π).

Table 2 reports a pair of regressions for each combination of the two inflation surveys and the three constructed inflation perceptions presented above. For each combination, the first regression restricts the intercept to zero and the sum of the other coefficients to be one, while the second regression is unrestricted. For the Michigan inflation survey (π_{mich}^s), coefficients on the perceived inflation target are uniformly larger and statistically more significant that are coefficients on actual inflation.[9] Indeed, the coefficients on the perceived targets tend to be highly significant while most coefficients on actual inflation are insignificantly different from zero. Nonetheless, the most important regressor is the lagged value of the survey, whose coefficient ranges between 0.7 and 0.8.

Qualitative aspects of the regressions for the Hoey-Philadelphia survey (π_{h-p}^s) are similar to those for the Michigan survey. The perceived inflation targets tend to be more significant than is actual inflation and the survey data are quite inertial. Quantitatively, for the Hoey-Philadelphia survey, the degree of difference in significance of the perceived target and actual inflation is reduced, and coefficients on the lagged survey observation are even higher.[10]

---

8 We have experimented with a version of the Kalman filter approach that allows for subsample variation in r^* and found that our results are little changed.

9 Estimation results in Table 2 are little affected if the inflation perceptions are entered with a lag rather than contemporaneously, if the first lag of actual inflation is used in place of its current value, or if lags 0 to 3 or lags 1 to 4 of actual inflation are entered. Furthermore, additional regressions indicate that Granger causality runs from the perceived targets to the Michigan survey but not vice versa.

10 Another difference is that the Hoey-Philadelphia survey and the perceived inflation targets each Granger causes the other.

Table 2

Regressions of survey inflation on constructed series

Survey inflation	Constructed inflation	Coefficient: cnst	π^s_{t-1}	π^c_t	π_t	Regression standard error
π^s_{mich}	w = 15	–	0.80	0.14	0.06	0.304
			(14.2)	(3.5)	(1.8)	
		0.39	0.72	0.14	0.02	0.263
		(2.7)	(13.9)	(3.3)	(0.7)	
	d = 0.02	–	0.79	0.16	0.05	0.303
			(14.5)	(3.9)	(1.6)	
		0.44	0.73	0.13	0.02	0.265
		(3.2)	(14.0)	(3.2)	(0.8)	
	Kalman	–	0.71	0.34	-0.04	0.247
			(17.2)	(8.0)	(1.5)	
		0.28	0.58	0.39	-0.00	0.238
		(2.1)	(9.7)	(5.2)	(0.1)	
π^s_{h-p}	w = 15	–	0.89	0.07	0.04	0.184
			(30.7)	(2.3)	(2.9)	
		0.09	0.91	0.05	0.01	0.180
		(0.9)	(30.5)	(1.3)	(0.5)	
	d = 0.02	–	0.87	0.10	0.053	0.178
			(32.2)	(3.3)	(2.2)	
		0.08	0.89	0.07	0.01	0.176
		(0.9)	(29.8)	(2.0)	(0.4)	
	Kalman	–	0.91	0.09	0.00	0.175
			(58.7)	(3.8)	(0.1)	
		-0.03	0.90	0.11	-0.00	0.178
		(0.2)	(29.8)	(1.7)	(0.1)	

Notes: t-statistics are shown in parentheses. Michigan (π^s_{mich}): sample period is 1980:Q3 – 1996:Q4. Hoey-Philadelphia (π^s_{h-p}): sample period is 1981:Q1 – 1996:Q4.

On the whole, the regression tests support the view that long-run inflation expectations move, in part at least, with changing perceptions about monetary policy objectives as derived from learning models. The adaptive expectations view has less support.

4. Simulation analysis

The historical analysis suggests that for FRB/US simulations which assume that the private sector has incomplete information about the nature of monetary policy, gradual adjustment of long-run expectations might be better specified as the outcome of learning about the parameters of the policy reaction function than as a simple partial adjustment to observed inflation. The remainder of this paper describes some exploratory FRB/US simulations in which the evolution of long-run inflation expectations is modeled as the outcome of learning. Because learning is a process of signal extraction, a realistic analysis of alternative learning mechanisms requires a stochastic rather than deterministic simulation environment.

4.1 Long-run inflation expectations in FRB/US

As mentioned in the introduction, FRB/US has several different simulation options for expectations formation. The option employed here bases expectations on the forecasts of a VAR system that at its core has a set of three equations for the Federal funds rate, inflation and the output gap. The VAR system is restricted so that as the planning horizon lengthens, period-by-period inflation expectations approach the long-run inflation expectation.[11] The long-run inflation expectation is an anchor or "endpoint" that at any point of time is predetermined in the calculation of expectations having a shorter horizon. In FRB/US simulations, the manner in which the inflation expectations endpoint moves over time has up until now been specified as either adaptive, in the sense of adjusting gradually toward actual inflation, or as embodying full knowledge of the true long-run policy objective for inflation. The simulations reported next instead use one of the regression-based learning algorithms.

4.2 Design of stochastic simulations

The FRB/US model consists of about 40 stochastic equations, numerous identities and about 100 exogenous variables. For stochastic simulations, equations are added for 10 key exogenous variables, such as the price of oil, so that they can be easily given random shocks. For the 50 stochastic equations in the augmented model, shocks are bootstrapped from historical residuals. In each period simulated, a historical quarter between 1966:Q1 and 1995:Q4 is randomly chosen and the vector of equation residuals associated with that quarter is drawn. Most FRB/US equations have residuals that are serially uncorrelated. For a few financial equations, however, residuals are serially correlated, and AR(1) error-propagation equations are added to the model in these instances. Monetary policy is characterized by the version of the "short-lags" Federal funds rate reaction function estimated from 1966:Q1 to 1996:Q4 that allows for a shift in its intercept at the end of 1979.

Several special issues arise in stochastic simulations that incorporate learning algorithms such as those discussed above. One is the need for initial conditions from which to start the algorithms. Although the last 15 or 20 years of US macroeconomic data could serve this purpose, with the stochastic simulations running from the present out into the future, other considerations make it easier to start from a deterministic baseline characterized by steady-state balanced growth. For this reason, the stochastic simulations reported in this paper have an initial 15-year period in which the long-run inflation expectation is exogenous. Then, with a long enough simulated "history" available, one of the learning algorithms is switched on for the remaining 35 years of each simulation.

A second issue concerns the shocks applied in the stochastic simulations to the Federal funds rate reaction function. The historical residuals of this equation are quite variable – the standard deviation is about 100 basis points – and include several outliers. An important question is to what degree these residuals, especially the outliers, represent actual surprises to participants in the economy and to what degree they reflect well-understood responses of the funds rate to special short-run factors that are not included in the reaction function. One example is the credit-control episode of 1980, which led to large short-run gyrations in GDP and interest rates as well as large residuals to the estimated equation for the Federal funds rate. Outside of a few episodes, however, it is more difficult to gauge the appropriate magnitude of "true" errors. In the stochastic simulations, two modifications are made to the funds rate shocks. First, to reduce the influence of outliers, the residuals are drawn from a normal distribution rather than from the historical set residuals. Second, to examine the sensitivity of simulation results to the magnitude of funds rate shocks, the standard deviation of the normally-distributed shocks is chosen alternatively as one-half or the same as the standard deviation of the historical residuals.

---

[11] The role of long-run expectations in the FRB/US model is discussed in more detail in Brayton, Mauskopf, Reifschneider, Tinsley and Williams (1997) and Bomfim and Rudebusch (1997).

Finally, a metric that will be used to evaluate the performance of the alternative learning procedures is the unemployment sacrifice ratio associated with a monetary policy shift that aims to reduce the rate of inflation one percentage point. For each particular learning procedure analyzed, two sets of stochastic simulations are run, one set in which the policy target for inflation is constant over time and a second set in which a one percentage point reduction in the inflation target is introduced in year 20. Because the same sequence of shocks is drawn in each stochastic set, pairwise comparisons of individual simulations can be made. Each simulation set consists of 50 replications.[12]

4.3 Simulation results

Table 3 summarizes the results of the FRB/US stochastic simulation experiments. Each row corresponds to a particular experiment whose design is described in the left pair of columns. The middle four columns report the standard deviations of key macroeconomic variables from the set of stochastic simulations in which the policy inflation target is held constant, and the three columns on the right present statistics based on comparing the disinflationary set of stochastic simulations with the set having the constant inflation target.

Table 3

Stochastic simulation results

Simulation design [1]		Statistics under constant π^{*} [2] (standard deviation)				Disinflation statistics [3]		
						Sacrifice ratio		Years for π^e
π^e	ε_i	π	π^e	$\tilde{x}$	i	median	s.e.	to fall 0.9
d=0.04	1.0	1.85	1.12	3.01	2.88	2.88	1.20	11.50
d=0.03	1.0	1.79	0.85	2.98	2.79	2.92	0.95	12.50
d=0.02	1.0	1.74	0.56	2.94	2.70	3.52	0.77	12.75
d=0.05	0.5	1.71	0.72	2.83	2.45	2.51	0.75	8.50
d=0.04	0.5	1.67	0.52	2.79	2.40	2.69	0.80	9.75
d=0.03	0.5	1.65	0.37	2.78	2.35	3.02	0.60	12.50
d=0.02	0.5	1.64	0.24	2.77	2.31	3.69	0.59	19.75
w=10	1.0	1.88	1.18	3.08	2.99	2.97	2.00	8.25
w=15	1.0	1.74	0.64	2.93	2.71	3.39	0.79	14.75
w=8	0.5	1.74	0.86	2.85	2.55	2.64	0.98	8.50
w=10	0.5	1.68	0.55	2.81	2.39	2.64	0.97	10.75
w=15	0.5	1.64	0.30	2.76	2.30	3.43	0.73	13.00
kf	1.0	1.79	0.70	2.96	2.79	3.09	1.03	12.00

[1] Rolling-regression learning is denoted by w (= years) for fixed window or d (= decay rate) for expanding sample with declining data weights. "kf" denotes the Kalman filter. ε_i is the scale factor applied to shocks to the Federal funds rate.

[2] Standard deviations for quarterly observations on inflation (π), the long-run expected inflation(π^e), deviation of output from potential ($\tilde{x}$) and the Federal funds rate (i). Standard deviations are calculated from years 20-50.

[3] Sacrifice ratio statistics are from year 30 of the simulations.

[12] Occasionally, the regression-based learning algorithms calculate a value for the long-run inflation expectation that is wildly high or low, probably when the learning algorithm estimates the reaction function coefficients imprecisely. To prevent rare occurrences of this sort from having a large impact on a particular simulation, upper and lower boundaries are placed on the permissible values of the long-run inflation expectation. The upper limit is 3 percentage points higher than the policy target for inflation in the constant inflation simulations and the lower limit is 3 percentage points lower than the policy target in the disinflation simulations.

Focusing first on the rows that correspond to the rolling regression approaches which performed best in the historical analysis (d=0.02 and w=15), the median unemployment sacrifice ratio is close to 3½, a value which is well above current "consensus" estimate of 2 or so – see, e.g. Ball (1994). This finding is not affected if the magnitude of the shocks to the funds rate equation is scaled down by one-half. These particular parameterizations of the rolling regressions also result in disinflations that are very slow: The average length of time it takes for the long-run inflation expectation to fall 0.90 percentage points, or 90% of its ultimate decline, is well over 10 years. Because the process of learning about the change in the policy objective for inflation occurs very gradually, and the degree of inertia in expectations is high, a sizeable increase in the unemployment rate is required to lower the rate of inflation.

The learning process is accelerated and the median sacrifice ratio reduced by adjusting the rolling regression parameters to speed the rate of decline of the data importance weights or shorten the estimation window. When the shocks to the Federal funds rate have their full historical variability, raising the decay rate to 4% per quarter or shortening the regression window to 10 years reduces the simulated sacrifice ratio to 3. If the funds rate shocks are reduced to one-half their historical variability, a decay rate of 5% or a regression window of 8 years results in a sacrifice ratio of a bit more than 2½.

Up to this point, the shortening of the actual or effective length of the rolling regressions has little deleterious effect on measures of macroeconomic volatility when the policy holds the target rate of inflation fixed. Increases in the standard deviations of inflation, output and the Federal funds rate are minor. Note, however, that the variability of the sacrifice ratio across individual simulations gets substantially larger in some cases. Further contraction of length of the rolling regressions leads to much higher volatility of expected long-run inflation and this spills over into higher macroeconomic variability.

The last two rows of Table 3 summarize the results of stochastic simulations based on Kalman filter learning. As shown in the middle columns, the measures of macroeconomic volatility reported in the table are little affected by allowing the FRB/US agents to use the Kalman filter to form long-run inflation expectations. The same is not true for the disinflation statistics: The median value of the sacrifice ratio is about 3, roughly ½ point lower than the sacrifice ratios implied by the rolling regressions that performed best in the historical analysis, but still on the high end of estimates reported in the empirical literature.

Concluding remarks

Long-run inflation expectations play an important role in the short-run macroeconomic dynamics of the FRB/US model. Yet, FRB/US has so far lacked an empirically based approach to modeling the evolution of long-run inflation expectations when the private sector is uncertain about the ultimate inflation goals of the policymaker. In its narrowest sense, we view this paper as an attempt to fill in this gap. The learning schemes estimated here come from estimated learning processes that are based on explicit beliefs about monetary policy and macroeconomic conditions. In a broader sense, this paper illustrates how real-time learning can be incorporated into a large-scale macroeconomic model in a way that attempts to be both data- and model-consistent.[13] We see this as an important contribution given the two most common modeling alternatives entail assuming either

[13] Other researchers who looked into similar issues include Hall and Garratt (1995), who examined similar issues in the context of the London Business School Model, and Fuhrer and Hooker (1993), who analyzed the economic implications of alternative learning schemes in a small-scale macro model.

that the agents have full and complete knowledge of the workings of the economy – the rational expectations hypothesis – or that agents are limited to passively responding to actual developments in lagged inflation.

We should also emphasize that the nature of our findings extends beyond the interest of large-scale macro modelers. In particular, our method allowed us to estimate survey-independent measures of market participants' long-run inflation expectations. This is of value in and of itself given that the available survey data often cover only a short span of time. More important, armed with our constructed time series, we plan to examine how different learning models conform to historical and perceived conditions in financial markets – e.g., ex-ante long-term real interest rates – and how well they anticipate future developments in inflation.

References

Ball, Laurence (1994): "What determines the sacrifice ratio?", in N. G. Mankiw (ed.), *Monetary Policy*, Chicago: University of Chicago Press, pp. 155-82.

Bomfim, Antulio N. and Glenn D. Rudebusch (1997): "Opportunistic and deliberate disinflation under imperfect credibility". Manuscript, Federal Reserve Board.

Bomfim, Antulio, Robert Tetlow, Peter von zur Muehlen and John Williams (1997): "Expectations, earning, and the costs of disinflation: experiments using the FRB/US model". Bank for International Settlements, *Conference Papers*, Vol. 4, pp. 222-43.

Brayton, Flint, Eileen Mauskopf, David Reifschneider, Peter Tinsley and John Williams (1997): "The role of expectations in the FRB/US macroeconomic model". *Federal Reserve Bulletin*, pp. 227-45, April.

Fuhrer, Jeffrey and Mark A. Hooker (1993): "Learning about monetary regime shifts in an overlapping wage contract model". *Journal of Economic Dynamics and Control*, 17, pp. 531-53.

Hall, Stephen and Anthony Garratt (1995): "Model consistent learning and regime switching in the London Business School model". *Economic Modelling*, 12, pp. 87-95.

Kozicki, Sharon, Dave Reifschneider and Peter Tinsley (1996): "The behaviour of long-term interest rates in the FRB/US model". Bank for International Settlements, *Conference Papers*, Vol. 2, pp. 215-50.

Kozicki, Sharon and Peter Tinsley (1996): "Moving endpoints and the internal consistency of agents' ex ante forecasts". Federal Reserve Board, *FEDS Working Paper*, No. 47.

Taylor, John (1993): "Discretion versus policy rules in practice", in *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, pp. 195-214.