Optimal Nonlinear Policy: Signal Extraction with a Non-Normal Prior

Eric T. Swanson
Board of Governors of the Federal Reserve System
eswanson@frb.gov

Abstract

The literature on optimal monetary policy under uncertainty typically makes three major assumptions: 1) policymakers' preferences are quadratic, 2) the economy is linear, and 3) stochastic errors and policymakers' prior beliefs about unobserved variables are normally distributed. This paper relaxes the third assumption and explores its implications for optimal policy. The separation principle and certainty equivalence continue to hold in this framework, but policymakers' beliefs are no longer updated in a linear fashion. I consider in particular a class of models in which policymakers' priors about the natural rate of unemployment are diffuse in a region around the mean. When this is the case, it is optimal for policy to respond cautiously to small surprises in the observed unemployment rate, but become increasingly more aggressive at the margin. This model appears to match well statements by Federal Reserve officials, and the historical behavior of the Fed, in the late 1990's.

JEL Classification: E52

Version 1.2.1
April 11, 2001

I thank Brian Sack, Dave Small, Volker Wieland, and two anonymous referees for helpful discussions, comments, and suggestions. The views expressed in this paper, and all errors and omissions, should be regarded as those solely of the author, and are not necessarily those of the individuals listed above, the Federal Reserve System, or its Board of Governors.

1. Introduction

The literature on optimal monetary policy under uncertainty typically makes three major assumptions: 1) policymakers' preferences are well-approximated by a quadratic function, 2) the economy is well-approximated by a linear system of equations, and 3) stochastic errors and policymakers' priors about unobserved variables are normally distributed. While these assumptions contribute much in the way of tractability and simplicity to the models and their solutions, it is important to understand the effects of relaxing these constraints on the prescriptions for optimal policy. In particular, I demonstrate that relaxing the assumption of normality, in favor of a prior that is diffuse in a region around the mean, provides a very simple and intuitive way of modeling statements by Federal Reserve officials, and the Fed's behavior, in the late 1990's.

Consider, for example, statements by Governor Laurence Meyer of the Federal Reserve Board, in a speech made September 8, 1999 to the Philadelphia Council for Business Economics:

> There are, however, a couple of constructive policy responses *in light of prevailing uncertainties* about the level of excess demand and the forecast. First, *policymakers could update their estimates of the NAIRU and the output gap* (assuming, in the first place, that they find these concepts useful, as I do) in light of realizations of unemployment, output, inflation, and other variables... Second, policymakers could attenuate the response of the real federal funds rate to declines in the unemployment rate in a region around their estimate of the NAIRU. *But once the unemployment rate gets far enough below (or above) the estimated NAIRU so that confidence returns that the labor market is experiencing excess demand (or supply), then the more normal response of real interest rates to incremental declines in the unemployment rate would again become appropriate.*
>
> (Meyer (1999), emphasis added)

I demonstrate below that a diffuse-middled prior on the natural rate of unemployment, or NAIRU, leads very naturally to the statements in this quotation.

The paper proceeds as follows. Section two presents an illustrative model of the economy with which the policy problem can be studied. Section three works through the univariate signal extraction problem when policymakers' priors on the NAIRU are thin-tailed. Section four works through the multivariate signal extraction problem that results when the model economy is slightly richer and considers an example calibrated to U. S. data in the late 1990's. Section five discusses the results and concludes.

2. A Simple Model

The main points of this paper are independent of the exact model under consideration. For the purposes of clarity and illustration, it is therefore advantageous to work with a model that is as simple as possible, while still conveying all of the relevant intuition. I thus work with the following simple, two-equation, backward-looking model as a baseline example:

$$(u_t - u^*) = \theta(u_{t-1} - u^*) + \alpha(r_{t-1} - r^*) + \varepsilon_t \tag{1a}$$

$$\pi_t = \pi_{t-1} - \beta(u_{t-1} - u^*) + \nu_t \tag{1b}$$

where u_t is the unemployment rate at time t, r_t the real interest rate, π_t the inflation rate, u^* the "natural" rate of unemployment consistent with long-run equilibirium ($|\theta| < 1$), and r^* is the "natural" rate of interest.[1] It is assumed for simplicity that policymakers have direct control over the real interest rate r_t. The random variables ε_t and ν_t denote Gaussian mean-zero stochastic shocks to the system, and are orthogonal to all variables dated $t-1$ or earlier.[2]

Policymakers set interest rates so as to minimize a discounted sum of squared deviations of unemployment and inflation from long-run target values:

$$\min\ (1-\delta)\ E_t \sum_{s=t}^{\infty} \delta^{s-t} \left[(\pi_s - \pi^*)^2 + \gamma(u_s - u^*)^2\right] \tag{2}$$

where π^* denotes policymakers' long-run target for inflation, δ is policymakers' discount factor, and γ is the relative weight on unemployment stabilization. It is assumed that policymakers' target for unemployment is the long-run natural rate u^*.

The model is complicated somewhat by the fact that the natural rate of unemployment u^* is never observed. Policymakers only observe their past choices for the policy

[1] One might describe equation (1a) as an "IS curve" and equation (1b) as a "Phillips curve."

[2] One can also consider the effects of non-normally distributed error terms ε and ν. I maintain the assumption of normality here so as to isolate the effects of non-normally distributed priors about u^*. In fact, the results demonstrated below are all magnified if we consider error terms (ε and ν) that are fat-tailed rather than Gaussian.

instrument r, current and past values of the actual unemployment rate u, and current and past values of the realized inflation rate π, which they can use to help them infer the true value of u^*. This they do by Bayesian updating, taking into account observations of unemployment and inflation as they come in. It is assumed for simplicity that the natural rate of interest r^* is known.

Policymakers' problem is thus a straightforward discrete-time dynamic programming problem with quadratic objective, linear constraints, and an unobserved state variable, the solution of which is well known (Bertsekas (1987)):[3]

$$r_t - r^* = a\, E_t(u_t - u^*) + b\,(\pi_t - \pi^*) \tag{3}$$

where E_t denotes the mathematical expectation conditional on all information $\mathcal{I}_t$ available as of time t:

$$\mathcal{I}_t \equiv \{\theta, \alpha, \beta, \gamma, \delta, \sigma_\varepsilon^2, \sigma_\nu^2, F_{u^*|0}(\cdot), \pi^*, r^*, \pi_t, u_t, \pi_s, r_s, u_s \mid s < t\} \tag{4}$$

The constants a and b in (3) are determined by the parameters of the model and, along with the form of equation (3), are independent of the variances of ε and ν—this is the well-known property of certainty equivalence in linear-quadratic models. $F_{u^*|0}(\cdot)$ in (4) denotes policymakers' prior distribution on u^* at time 0. Note that the parameters of the model (θ, α, β, etc.) are assumed to be known with certainty.

Policymakers enter period t with prior beliefs about u^*, the expected value of which is $E_{t-1}u^*$. Based on this prior, and observations of lagged variables, policymakers have prior forecasts for the variables u_t and π_t (derived from equations (1)), namely:

$$E_{t-1}u_t = (1-\theta)E_{t-1}u^* + \theta u_{t-1} + \alpha(r_{t-1} - r^*) \tag{4a}$$

$$E_{t-1}\pi_t = \pi_{t-1} - \beta(u_{t-1} - E_{t-1}u^*) \tag{4b}$$

If the realized values of u_t and π_t come in close to these prior forecasts, then policymakers will see little reason to revise their beliefs, and $E_t u^*$ will be very close to $E_{t-1}u^*$. In contrast, if u_t and π_t contain substantial "surprises" relative to policymakers' prior expectations of these variables, then policymakers may be prompted to revise their beliefs about u^* more substantially. Sections 3 and 4 below consider this problem in detail.

[3] Note that this solution does not require normality of any of ε, ν, or u^*.

2.1 Separation of Estimation and Control

Policymakers' problem (minimizing (2) subject to (1)) has the well-known property of separability of estimation and control. Policymakers' problem thus can be separated into two stages: first, the unemployment gap $(u_t - u^*)$ is estimated on the basis of all available information in period t, and second, the interest rate r_t is set based on this estimate, as in (3). Note that this is a general property of the linearity in model (1), and holds for priors and error terms with completely general distributions.[4]

The division of policymakers' problem into estimation and control stages can be thought of in terms of the following figure:

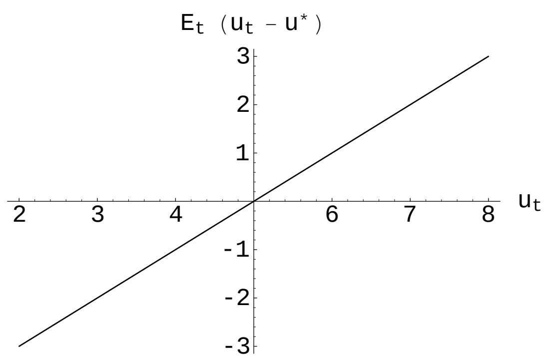

Figure 2.1a: Estimation

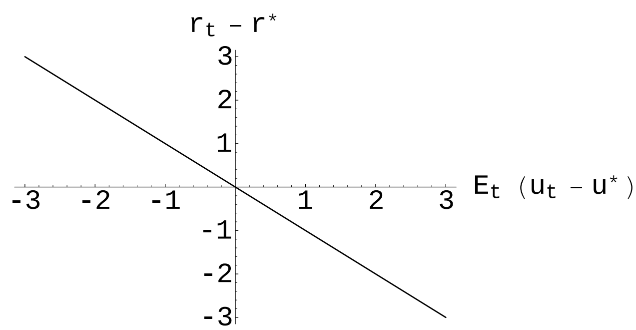

Figure 2.1b: Control

Note that the figure abstracts away from inflation equation (1b) for graphical simplicity, so that policymakers' estimation and control stages are each functions of only one variable, the unemployment rate in Figure 2.1a, and the estimated unemployment gap in Figure 2.1b.

Figure 2.1a depicts the estimated unemployment gap as a linear function of the observed value of u_t. This is the case, for example, when policymakers' priors about u^* and ε_t are independently normally distributed, resulting in the standard updating equation $E_t(u_t - u^*) = E_{t-1}(u_t - u^*) + (\mathrm{Var}_{t-1}\varepsilon_t/\mathrm{Var}_{t-1}u_t)(u_t - E_{t-1}u_t)$, where Var_{t-1} denotes policymakers' prior variance on the given variable conditional on information available at time $t-1$.

Figure 2.1b depicts the policy instrument as a linear function of the estimated unemployment gap, $E_t(u_t - u^*)$. This is the case for linear-quadratic models in general, and

[4] Doubting readers are referred to Bertsekas (1987), pp. 102–6, 292–3.

holds specifically for the optimal policy given in equation (3).[5] Combining the left and right panels of Figure 2.1 yields policymakers' reduced-form response to the unemployment rate u_t. Under the usual assumption that both panels of Figure 2.1 are linear, this reduced-form reaction function is also linear in u_t.

2.2 Two Motivations for Nonlinear Policy

One can go beyond the linear reduced-form response implied by Figure 2.1 by introducing nonlinearities into either the left-hand or right-hand panel of that figure. For example, Orphanides and Wieland (2000) do the latter, considering a model with a nonlinear Phillips curve, with a concave-to-convex shape (*i.e.,* a shape similar to $y = x^3$) that is like the one estimated by Filardo (1998). This approach yields an optimal nonlinear policy response to unemployment even in a world of perfect certainty, as well as a world in which disturbances and policymakers' priors are all normally distributed.[6]

While this approach has its merits, the quotation by Governor Meyer given earlier seems to indicate policymakers in the 1990's were concerned primarily with uncertainty about their estimates of NAIRU or full capacity, rather than important nonlinearities in the economy or non-quadratic preferences. A modification of the *left*-hand panel of Figure 2.1 (the estimation stage of policymakers' problem) thus seems more appropriate for modeling recent monetary policy behavior in the U. S.

3. Univariate Signal Extraction with a Non-Normal Prior

Policymakers' problem of estimating u^* is inherently one of signal extraction. To see this, note that policymakers never observe the true value of u^*, but only receive noisy observations of u^* through realizations of u_t and π_t.[7] Policymakers use these observations to update their beliefs about the true value of u^*.

[5] Of course, equation (3) is a function of both unemployment and inflation, the latter of which has been abstracted away from in Figure 2.1.

[6] Orphanides and Wieland (2000) also consider a case where policymakers' preferences are non-quadratic, in particular by introducing a "zone of indifference" for inflation rates between 0 and 2 percent (motivated, for example, by many inflation-targeting central banks' official charters). This leads to a very similar nonlinearity in the right-hand panel of Figure 2.1.

[7] To emphasize the signal extraction properties of equations (1), u^* could be shifted from the left-hand side to the right-hand side of equation (1a).

I begin by illustrating this process in detail when policymakers only have to perform univariate signal extraction. While multivariate signal extraction is not unduly difficult, the figures are much simpler for the univariate case, and little intuition is lost by the simplification. The more general multivariate signal extraction problem will be considered in the following section. Until then, I abstract away from the inflation equation (1b).

When policymakers' priors about u^* and ε_t are both normally distributed, one gets the standard result that the optimal estimate $E_t(u_t - u^*)$ is a linear function of the observed variable u_t, as discussed in the previous section. When policymakers' priors about u^* are not normally distributed, however, this linearity fails to obtain. For example, suppose policymakers' priors on u^* were distributed uniformly over the interval $[4, 6]$, as in Figure 3.1a below:

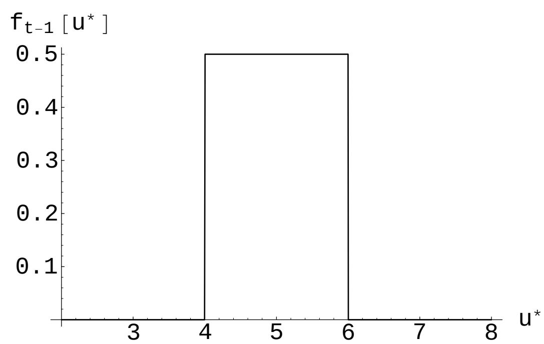

Figure 3.1a

Intuitively, in this figure policymakers' point estimate of the NAIRU is equal to 5, but this estimate is uncertain enough that any point in the interval $[4, 6]$ is deemed equally likely to be the true value. Policymakers' uncertainty drops off rapidly as one moves outside of the interval $[4, 6]$—in this idealized example, policymakers are certain that the NAIRU lies neither above 6 nor below 4.

This nonnormality of policymakers' beliefs will be reflected in their posterior estimates of u^* and the unemployment gap $\tilde{u}_t \equiv u_t - u^*$, after observing u_t. These estimates, which would have been linear in u_t had we assumed policymakers' priors to be normally distributed, now have the following functional forms with respect to u_t:[8]

[8] It is assumed in these figures that u_{t-1} and r_{t-1} are such that policymakers' prior forecast $E_{t-1}u_t$ equals 5 (and the expected unemployment gap $E_{t-1}\tilde{u}_t$ is zero). The graphs look exactly the same (up to a translation) for different prior forecasts of u_t, since it is only the unemployment surprise that enters into policymakers' updating equation. The standard deviation of the shock to the unemployment gap (ε_t) in this figure and those that follow is set equal to 0.4. This number comes from the calibration (to annual data) in Section 4, below.

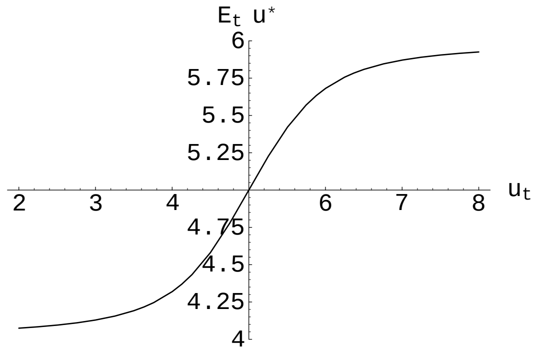

Figure 3.1b

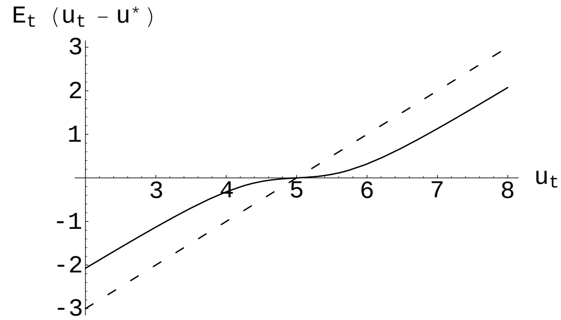

Figure 3.1c

Intuitively, because policymakers are so uncertain about the NAIRU within the interval $[4, 6]$, they are very willing to revise their estimate of u^* for observations of u_t that are well inside this interval, as is evident in Figure 3.1b. As observed unemployment moves farther away from policymakers' prior point estimate of 5, however, policymakers assign an increasingly smaller fraction of each increment of unemployment to u^*, this fraction approaching zero as u_t becomes more and more extreme relative to policymakers' priors. Correspondingly, the fraction of each increment of unemployment assigned to the unemployment gap, $\tilde{u}_t$, is close to 0 near the middle of Figure 3.1c, and approaches 1 as u_t moves out toward the edges of that figure.

Policymakers are thus very willing to revise their estimate of u^* (and set r_t cautiously) for small surprises in the observed unemployment rate. Note that this result accentuates the signal extraction-based motivation for policy attenuation discussed in Swanson (2000) and Svensson and Woodford (2000).[9] The optimal marginal response of policy increases substantially, however, as the surprise in unemployment becomes larger, approaching the marginal certainty-equivalent response in the limit. This matches very well the monetary policy prescription given by Governor Meyer in the quotation presented earlier.

3.1 Other Diffuse-Middled Distributions on the NAIRU

The increasingly aggressive marginal response of policy to observed unemployment u_t above is not specific to the uniform distribution in Figure 3.1a, but rather holds more generally

[9] Sack and Wieland (2000) survey the literature on optimal policy attenuation. The motivation given by Swanson (2000) and Svensson and Woodford (2000) is based on signal extraction, as in the present paper. The diffuse-middled priors in this paper accentuate the attenuation effect by increasing the variance of policymakers' priors without making the tails implausibly large. Note also that, unlike the present paper, Swanson (2000) takes the signal extraction problem one level deeper by assuming the *estimated* unemployment gap $E_t(u_t - u^*)$ is *itself* an indicator for an underlying variable of interest, such as "excess demand" or "inflationary pressures."

for any prior distribution on u^* that is diffuse in a region around the mean. Intuitively, the diffuse center of the distribution makes policymakers very willing to revise their beliefs about u^* for small surprises in observed unemployment u_t. This flexibility fades as the surprises move out toward the tails of policymakers' prior distribution.[10] A more rigorous treatment of this point is provided in the Appendix.

For example, policymakers' priors might be distributed less like the uniform distribution and more along the lines of $e^{-\frac{1}{2}(x-5)^4}$, as in the solid line in Figure 3.2a below (dashed lines depict uniform and normal distributions for comparison):[11]

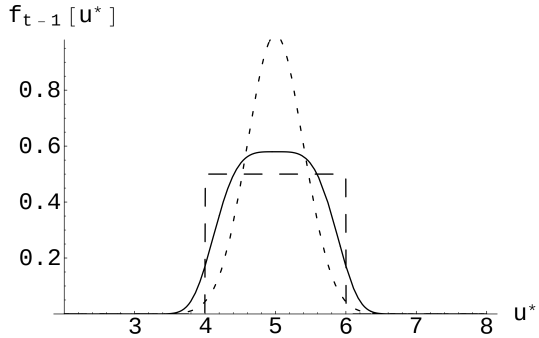

Figure 3.2a

Bayesian updating for this intermediate distribution leads to the following estimates of u^* and $\tilde{u}_t \equiv u_t - u^*$ as functions of observed unemployment u_t:

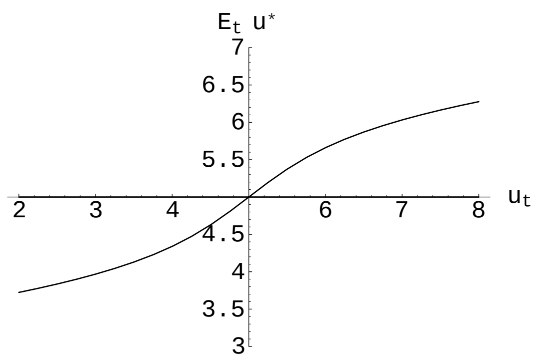

Figure 3.2b

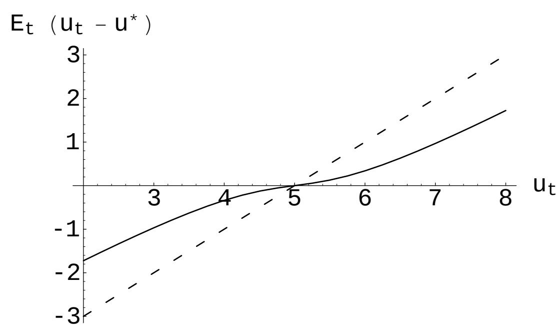

Figure 3.2c

[10] In addition, the tails of policymakers' prior influence the optimal marginal policy response for large surprises. If the distribution is thin-tailed (*i.e.*, $f(x)e^{\frac{1}{2}x^2}$ vanishes as $|x| \to \infty$), then policymakers' optimal marginal response approaches the certainty-equivalent marginal response for large surprises. If the tails are asymptotically Gaussian, then the optimal marginal response approaches a line with slope $\mathrm{Var}(\tilde{u}_t)/\mathrm{Var}(u_t)$, the signal-to-noise ratio associated with the sum of independent normal random variables.

[11] The exact formula for the solid-lined density is $ke^{-\frac{1}{2}((x-5)/.8)^4}$, where the normalization constant $k = 1/((.8)\,2^{5/4}\,\Gamma(5/4))$. The (short-dashed line) normal density is distributed $N(5, .16)$. Although the non-uniform densities in Figure 3.2a appear to have compact support, this is an artifact of the graph's resolution—the densities are in fact strictly positive over the entire real line. The integration required in going from the solid line in Figure 3.2a to Figures 3.2b and 3.2c is performed over the entire real line.

As in the case of the uniform distribution, policymakers are very willing to revise their estimate of u^* for small surprises in u_t. Optimal policy calls for an increasingly aggressive response at the margin, however, as u_t moves farther away from policymakers' prior forecast, again approaching the certainty-equivalent marginal response in the limit.[12]

3.2 Structural Change as a Motivation for Diffuse-Middled Priors

The question naturally arises as to how policymakers might have arrived at such a diffuse-middled prior on u^* to begin with. The examples above are essentially static in nature, and do not address this issue.

A reasonable answer to this question is provided by a model of structural change. Suppose that the usual state of the world is one in which policymakers' priors about u^* are close to normally distributed, but that policymakers know there is a very small probability p each period of a structural shift occurring—*i.e.,* a very significant shock to u^*. Policymakers do not observe such a shock to u^* directly, but would begin to infer it as observations of the data roll in.

Some evidence that this story may have been important in the Fed's thinking can be found in remarks by Federal Reserve Board Vice Chairman Roger Ferguson in the June 2000 issue of *The Region,* published by the Federal Reserve Bank of Minneapolis:

> I am comfortable with the idea that demand and supply imbalances affect prices and inflation in the short run. What we may have seen recently is that the point at which that trade-off starts to come into play might have moved from an unemployment rate that was up to 6 percent to an unemployment rate that is maybe somewhat lower. Some people argue that the so-called NAIRU, which is an element of the Phillips curve, may have moved down to 5 percent or even lower.
>
> I must say that anyone who has focused on the NAIRU recognizes that you can't get a point estimate that is immutable over time because the nature of the economy does change. Even now as we talk about imbalances, there is an implicit short-run Phillips curve concept embedded in the discussion without necessarily saying that the unemployment rate at which inflation starts to pick

[12] As noted briefly in footnote 2, one can also consider the effects of fat-tailed shocks ε on the results. Extending the model in this direction strengthens the nonlinearities depicted in the main body of the paper. Intuitively, for small surprises in u_t, policymakers divide the surprise relatively evenly between ε and u^*, while for large surprises in u_t, policymakers assign the preponderance of the surprise at the margin to ε, because it is fat-tailed.

up is exactly 4.5 percent, or 4.9 percent, or 5.2 percent or 5.5 percent. So I'm a bit of, how can I describe myself, in the middle in believing that there is a short-term trade-off between resource utilization and inflation, but not necessarily being wed to a specific point estimate on the short-run Phillips curve where inflation is likely to accelerate.

(Ferguson (2000))

These statements seem to imply that structural change and diffuse-middled distributions of the type considered earlier are indeed very relevant for thinking about the recent behavior of U. S. monetary policy.

We can incorporate structural change more concretely into our earlier framework as follows. Suppose that with large probability $1-p$ the NAIRU u^* is constant from one period to the next,[13] and that policymakers' priors about the true value of u^*, conditional on no structural change, are normally distributed, with a mean of 6 and standard deviation 0.25. In addition, suppose that, conditonal on the occurrence of structural change, policymakers' priors are normally distributed with the same mean 6, but with a standard deviation of 1.0.[14] Figure 3.3a below depicts this prior distribution for $p = .02$:

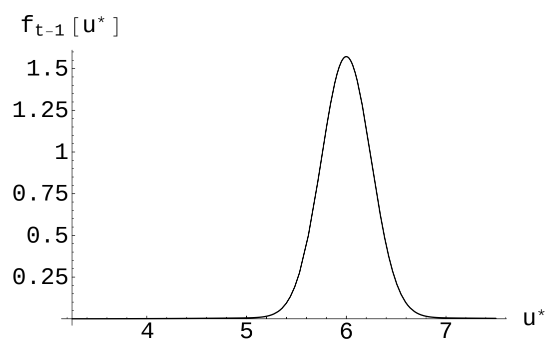

Figure 3.3a

Note that it is very nearly Gaussian, as the probability of structural change is fairly small.

Starting from this prior, an anomalous value for u_t causes policymakers to seriously consider the possibility that a structural break has in fact occurred. For example, starting

[13] One can allow u^* to drift from period to period by some small disturbance ξ_t without altering the story being told here.

[14] The implicit assumption here is that policymakers have no prior reason to believe any potential shift in u^* would be downward rather than upward. In fact, one might argue that observed technological innovations in computers and communications equipment did give policymakers a reason to believe u^* might move down rather than up. The model in the main body of the paper could easily be adjusted to describe such an asymmetric prior (e.g., by letting the mean, conditional on structural change, equal 5).

from the prior above, an observed value for u_t that comes in substantially below policymakers' prior expectation of 6,[15] leads to a diffuse-middled or even "two-regime" posterior distribution on u^*, as in Figure 3.3b below, which depicts policymakers' posterior after an observed unemployment rate of 4.3:

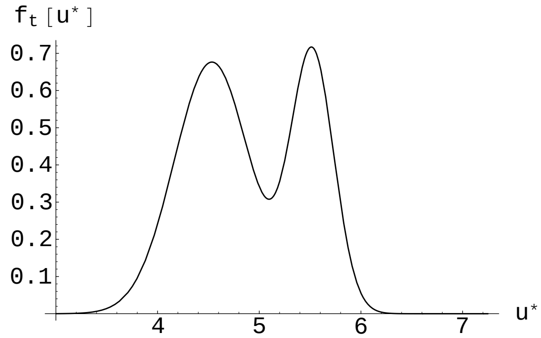

Figure 3.3b

This posterior clearly exhibits "New Economy" and "Old Economy" regimes: the former is centered around a value for u^* of about 4.5, and the latter around a value of about 5.5.[16] If a structural break has in fact occurred, the true value of u^* is likely to lie around 4.5. Alternatively, if a structural break has not occurred, the true u^* is likely to lie around 5.5. Policymakers are implicitly taking both of these possibilities into consideration in Figure 3.3b.

The two-regime posterior in Figure 3.3b possesses the same "diffuse-middled" property as the prior distributions considered earlier (see the Appendix). Subsequent observations of unemployment u_{t+1} that lie between the point estimate of u^* for each regime considered separately (here 4.5 and 5.5) will cause policymakers to significantly revise their beliefs about which regime they are in; policymakers' best estimate of u^* is thus highly responsive to changes in the observed unemployment rate in this region. Outside of this range, policymakers' marginal inference and policy responses return quickly to the usual linear function of observed unemployment. This is apparent in Figures 3.3c and d below, which depict the optimal subsequent estimates of u^* and the unemployment gap as

[15] In this simple one-equation version of the model, it is easily seen that optimizing policymakers will always have set r_{t-1} so that their prior expection $E_{t-1}u_t = 6$.

[16] This clear division into two distinct regimes results from the idealized structure and simplicity of the model. More generally, the process of structural change might be modeled as taking place more gradually, with the model itself having more parameters and dimensions of uncertainty, leading to a less sharp distinction between regimes, as well as a more gradual updating process.

functions of observed next-period unemployment u_{t+1}:[17]

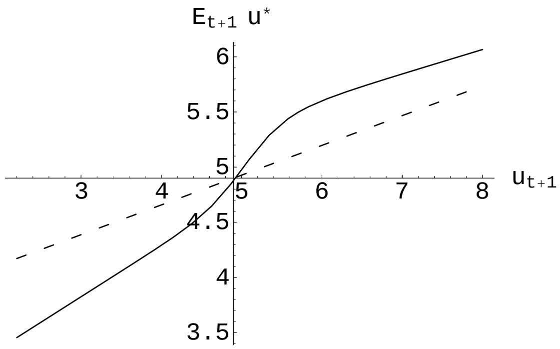

Figure 3.3c

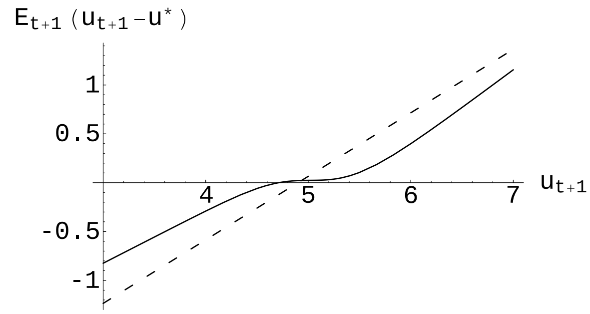

Figure 3.3d

Note that there is an asymmetry in these figures resulting from the asymmetric distribution in Figure 3.3b.

4. Multivariate Signal Extraction with a Non-Normal Prior

The above discussion highlights the basic point that diffuse-middled priors on the NAIRU make it optimal for policymakers to set interest rates cautiously in response to small unemployment surprises, but to become increasingly more aggressive on the margin as the size of the unemployment surprise increases. This observation continues to hold in more realistic, multi-equation models of the economy, though it is complicated somewhat by the need to perform multivariate signal extraction.

In this section, I reinstate the inflation equation of the model (1b), and let policymakers update their beliefs about the NAIRU in response to observations of both unemployment and inflation. For this more general case, policymakers' optimal inference and response functions are no longer functions simply of u_t, but now also depend on the inflation rate π_t. This results in graphs that look something like the following figures:

[17] The dashed lines in Figures 3.3c and d have slopes that are less than unity, corresponding to the signal-to-noise ratios associated with the sum of independent normal random variables. Although u^* is not normally distributed, the two tails of its density in Figure 3.3b closely approximate those of a normal random variable asymptotically, leading policymakers' optimal estimate of u^* and the unemployment gap $(u_{t+1} - u^*)$ to approach a line with slope corresponding to the signal-to-noise ratio. Were the tails to converge to zero more quickly than normal, as in the figures considered earlier, the asymptotic slope in Figure 3.3d would be 1 (and that in Figure 3.3c would be 0) as in those earlier figures.

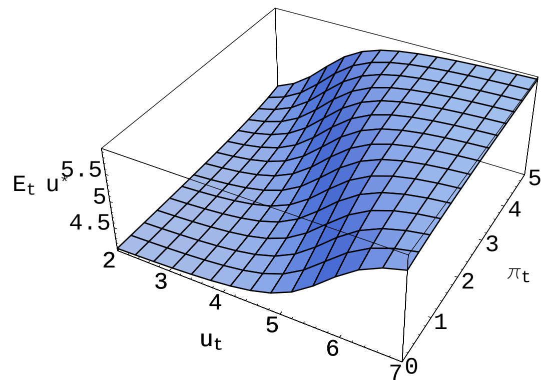

Figure 4.1a

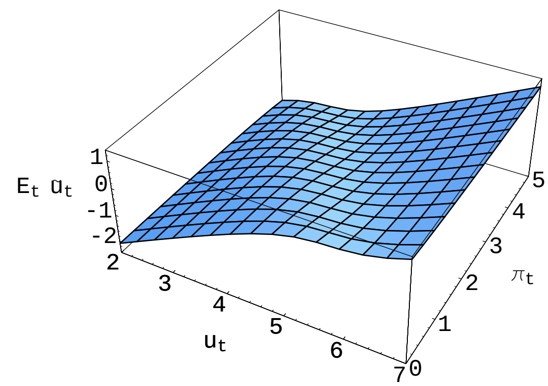

Figure 4.1b

where $\tilde{u}_t$ denotes the unemployment gap $(u_t - u^*)$. Cross-sections of these surfaces for a given value of π_t generate graphs that are analogous to the figures in section 3.[18] In the presentation of the results below, I will tend to focus on representative cross-sections rather than the full three-dimensional graphs, as the former present the results somewhat more clearly. However, the three-dimensional nature of the underlying optimal response graphs should be kept in mind.

The results of this section are presented in terms of a calibrated example, describing roughly the situation facing Fed policymakers in the late 1990's.

4.1 A Calibrated Example: The U.S. in the Late 1990's

The parameters of model (1) can be calibrated using annual U. S. data from 1960 to 1998. This yields the following values: $\theta = 0.75$, $\alpha = 0.15$, $\beta = 0.4$, and $r^* = 2.5$. In addition, the standard deviations of ε_t and ν_t are estimated to be 0.8 and 1.2, respectively.[19] These values correspond closely to those used by Reifschneider and Williams (1999) for a very similar model.

I have chosen to calibrate the model to annual rather than quarterly data for the following reasons. First, the story being told here is one of policymakers' beliefs, surprises in the data, and Bayesian inference. If the data are serially correlated, policymakers' revision

[18] One could also present cross-sections along the other dimension of Figures 4.1a and b—*i.e.,* for a given realization of the unemployment rate. This would yield graphs analogous to the figures of section 3 that present policymakers' optimal responses along the inflation dimension. I will present some graphs along these lines below.

[19] I assume for the time being (counterfactually) that ε_t and ν_t are uncorrelated, and relax this assumption later. Note also that the estimated standard deviations differ substantially across time periods. I will discuss some implications of this, and plausible values for these variance-covariance parameters, below.

process may be very gradual, and may be more naturally modeled as an annual rather than a quarterly frequency process. Second, policymakers update their priors based on "surprises" in the observed data. If the signal-to-noise ratio is substantially higher in annual than in quarterly data, then policymakers should be more willing to revise their beliefs in the face of surprises in annual data than they are to surprises in quarterly data. Thus, the nonlinearities demonstrated in the previous section may be more naturally thought of in terms of an annual model. Conceptually, however, the same exercise performed in this section could be undertaken using quarterly data.

For simplicity, I assume that policymakers' prior distribution on u^* is uniform, namely:

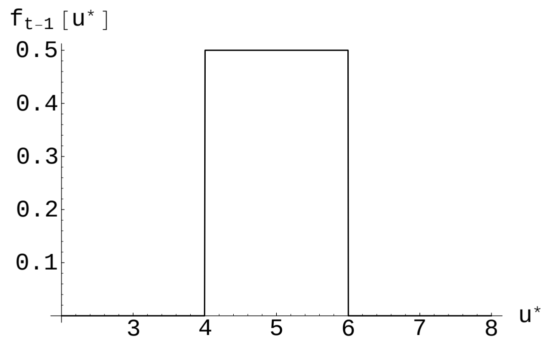

Figure 4.2

I choose to work with the uniform prior for a number of reasons. First, Bayesian updating is much easier computationally for a distribution with compact support, because integration does not need to be performed over the region where policymakers' priors have zero density. Second, it was shown in the previous section that diffuse-middled distributions other than the uniform yielded results that were qualitatively very similar to those for the uniform distribution, so the results here should be representative of those for diffuse-middled distributions more generally. Finally, the stylized appearance of the uniform prior helps to emphasize the fact that these Bayesian updating examples are by their very nature illustrative rather than definitive, and thus should not be taken too literally.

The U. S. economy in 1999 consisted of an inflation rate (PCE inflation) that was about 2.0 percent, an unemployment rate of about 4.2 percent, and a real federal funds rate of about 3.0 percent (these are annual averages). Taking the simple model (equations (1)) literally, this implies prior forecasts for unemployment u_t in 2000 of $(.25)(5)+(.75)(4.2)+(.15)(3.0-2.5)\approx 4.5$ and 2000 inflation (π_t) of $(2.0)-(0.4)(-0.8)\approx 2.3$ percent.

When data on unemployment and inflation for all of 2000 come in, policymakers will want to revise their beliefs on the basis of these observations. For example, if PCE inflation comes in at 2.3 percent as expected, but unemployment comes in at a value different from their prior forecast of 4.5, policymakers will revise their expectation of u^* and the unemployment gap $\tilde{u}_t$ accordingly. These revisions are depicted in the following graphs as functions of the realized value of the unemployment rate, holding realized inflation constant at 2.3 percent:[20]

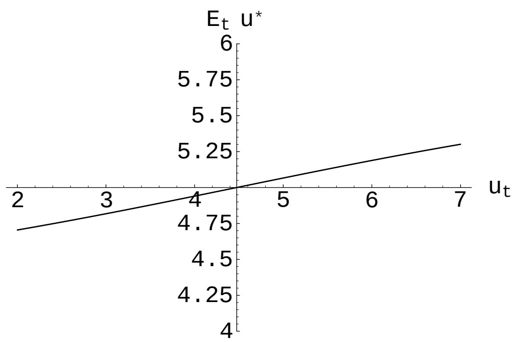

Figure 4.3a

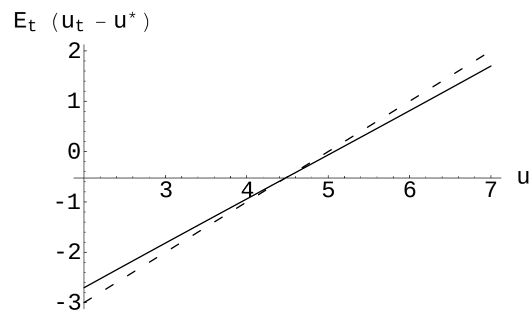

Figure 4.3b

Note that these graphs are essentially linear, so that a nonlinear response of interest rates would not be warranted.[21] This result hinges crucially on the fact that the estimated variances of the shocks to the simple model are quite large over the period 1960 to 1998 (standard deviations of 0.8 and 1.2 for surprises to unemployment and inflation, respectively). Because policymakers in the model know that this variance is so large, an unemployment realization of, say, 4.0 percent is not enough for them to significantly revise their beliefs (indeed, it is within one standard deviation of their prior forecast), and thus they revise their NAIRU estimate downward only very slightly, from 5.0 percent to 4.9. The effects on policymakers' beliefs are so small quantitatively that there is little reason for the types of nonlinear policy responses described earlier.

The variances estimated above, however, may not be the most appropriate values for use in the model. For example, the estimated shocks to this simple model in the 1970's

[20] The line in Figure 4.3a intersects the horizontal axis at the point $(4.475, 5.0)$. This intersection point states that, at realized inflation of 2.32 and unemployment of 4.475 (which were exactly policymakers' prior forecasts for these variables), policymakers' posterior expectation of u^* is exactly equal to their prior of 5.0. The line in Figure 4.3b intersects the horizontal axis at the point $(4.475, -0.525)$. This states that, at realized inflation of 2.32 and unemployment of 4.475, policymakers' posterior expectation of the unemployment gap is $4.475 - 5.0 = -0.525$.

[21] As in section 3, the dashed line in Figure 4.3b is a 45-degree line corresponding to the certainty-equivalent case where policymakers know, or act as if they know, with certainty that $u^* = 5.0$.

and early 1980's are very large; re-estimating the shock variances for the more recent period from 1983 to 1997 yields standard deviations of 0.6 and 0.7 for the unemployment and inflation surprises, respectively. Furthermore, shocks to the equations have tended to be substantially greater coming in and out of recessions than during normal times; given that we are not currently in a recession or coming out of one, the appropriate standard deviations for the surprises may be smaller still—0.4 and 0.6 for unemployment and inflation, respectively. If we take these last numbers as the appropriate values, then policymakers take a low realized value of unemployment much more seriously than they did in the previous example, and will tend to revise their beliefs about u^* more substantially, from 5.0 to 4.8 in Figure 4.4a below:[22]

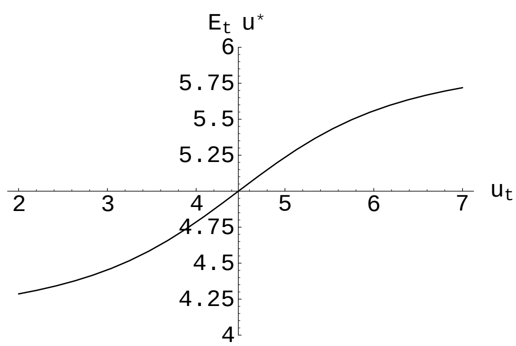

Figure 4.4a

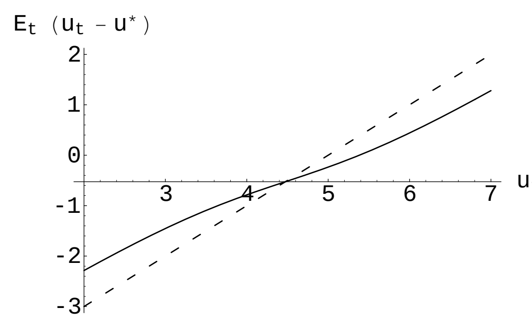

Figure 4.4b

Just as in the univariate examples of the previous section, policymakers here are very willing to revise their estimate of u^* for small surprises in u_t. As the size of the surprise gets larger, however, policymakers become increasingly more reluctant to make incremental revisions to this estimate, and the optimal marginal response of policy approaches the certainty-equivalent marginal response.

One can see similar effects looking along the inflation dimension. For example, if the unemployment rate in 2000 were to come in at the prior forecast of 4.5 percent, but inflation were to come in at a value different from policymakers' prior forecast of 2.3 percent, then

[22] Inflation is again assumed to come in at the expected value of 2.32 in these figures. Note that, while the standard deviation of the unemployment surprises in Figure 4.4b is the same as that in the figures of section 3 (namely, 0.4), the graph does not have the same degree of curvature that was present in those earlier figures. This is because Figures 4.4 assume a zero inflation surprise, so that the "joint surprise" in inflation and unemployment in Figures 4.4 is somewhat less than the univariate surprises that correspond to the horizontal axis in the figures of section 3.

policymakers would want to revise their estimates of u^* and the unemployment gap $\tilde{u}_t$ as in the following figure:[23]

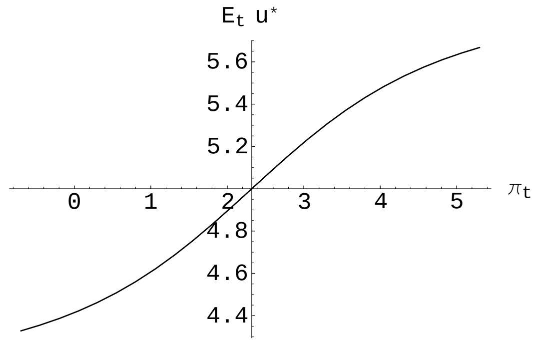

Figure 4.4c

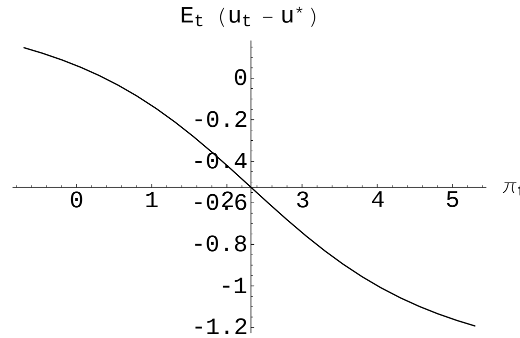

Figure 4.4d

Note that the estimated unemployment gap in Figure 4.4d is downward-sloping and does not grow without bound as in Figure 4.4b.[24] This follows from the fact that in Figure 4.4d, the observed unemployment rate is fixed at 4.5 percent and only the estimate of u^* is changing, in contrast to Figure 4.4b, in which both the observed unemployment rate and policymakers' estimate of u^* are changing.

It should be kept in mind that, as noted earlier, these graphs are intended to serve as representative cross-sections of the underlying three-dimensonal surfaces. Thus, policymakers' optimal estimates of u^* and $\tilde{u}_t$ in the previous two figures can be represented more generally as functions of unemployment and inflation both:

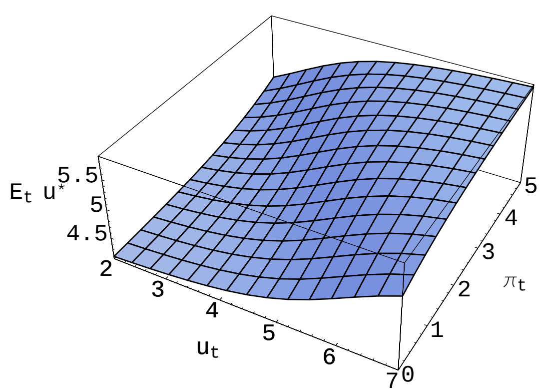

Figure 4.4e

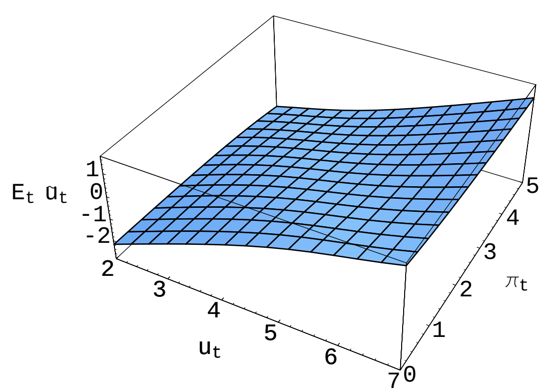

Figure 4.4f

[23] The points of intersection of the graphs with the horizontal axes are $(2.32, 5.0)$ in Figure 4.4c and $(2.32, -0.525)$ in Figure 4.4d. The reasons are the same as in footnote 20.

[24] Instead, it approaches horizontal asymptotes at 0.475 and -1.525, respectively.

In Figures 4.4a through d, we considered a surprise in only one variable in each figure. In Figures 4.4e and f, we can consider simultaneous surprises in unemployment and inflation of any magnitudes.

4.2 Correlation Between Surprises in Unemployment and Inflation

It was assumed for simplicity and clarity in the examples above that shocks to the unemployment rate and inflation (ε_t and ν_t) were uncorrelated. In fact, this assumption is both theoretically and empirically suspect. Theoretically, if the simple model in (1) omits important variables from the right-hand side that cause the error terms for unemployment (ε_t) and inflation (ν_t) to move in related ways, then policymakers will expect the surprises in the model to be correspondingly correlated. In addition, empirically, the estimated correlation of the error terms in model (1) is -0.4 over the full sample (1960 to 1998), and -0.6 over the period 1983 to 1997.

Taken at face value, these correlations imply that the situation facing policymakers in the late 1990's of low unemployment and low inflation *both* was even more surprising than the examples above would indicate, because historically unemployment and inflation have tended to surprise in opposite directions. This increase in the extent to which policymakers are surprised should lead them to revise their beliefs about u^* more substantially, and also lead them to become more averse to making incremental revisions in their estimate of u^*, so that the curvature that was characteristic of the figures above becomes more pronounced. This is demonstrated in Figures 4.5a and b, below:

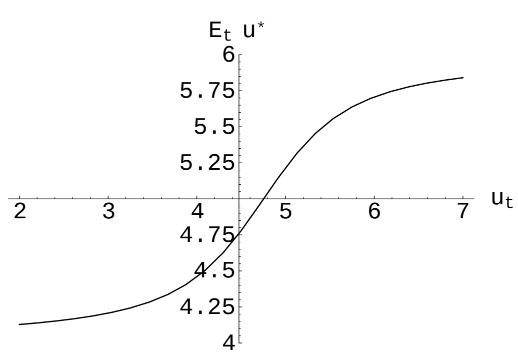

Figure 4.5a

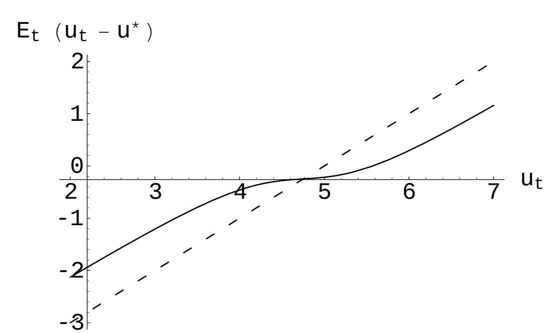

Figure 4.5b

Note that these figures assume a realized inflation rate of 1.8 percent, in line with the situation faced by policymakers in 1999, and a correlation between the error terms of -0.6.

This downward surprise in inflation causes policymakers' estimate $E_t u^*$ to be shifted downward for every value of u_t relative to Figure 4.4a—in fact, it takes an upward surprise in unemployment of $4.8 - 4.5 = 0.3$ to offset the downward surprise in inflation of 0.5, and lead to no net change in policymakers' estimate of u^*.

The full, three-dimensional plots of policymakers' estimates of u^* and $\tilde{u}_t$, as functions of unemployment and inflation both, are as follows:

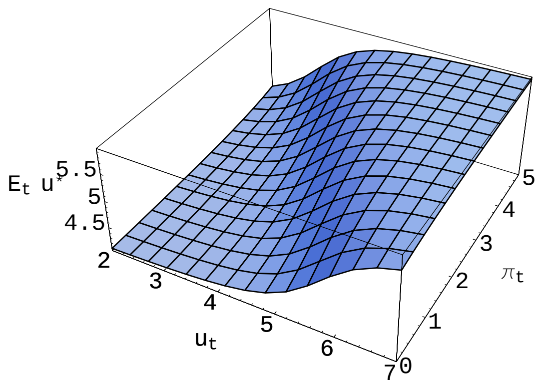

Figure 4.5c

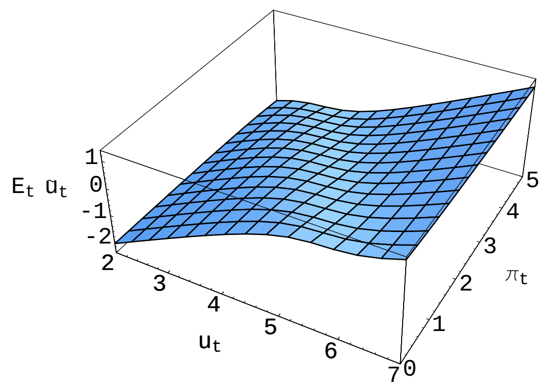

Figure 4.5d

It can be seen in the figures that, for a given downward surprise in unemployment, it takes a greater upward surprise in inflation for policymakers' estimate of u^* to remain unchanged. This is because some degree of negative correlation between the two surprises is already expected by policymakers *a priori*.

5. Discussion

The analysis and examples above clearly demonstrate a case for optimal nonlinear policy, as a function of observable variables, when policymakers' priors on unobserved variables of the system are non-normally distributed. In particular, if policymakers' priors are diffuse in a region around the mean, then it is optimal for policy to respond cautiously to small surprises in the observed unemployment rate, but become increasingly more aggressive at the margin as the size of the surprise gets larger. This matches very well the intuition presented in statements by Federal Reserve officials in 1999 and 2000.

Does it match the actual behavior of the Fed in the late 1990's? A strong case can be made that it does. For example, in a recent speech before the Joint Conference of the Federal Reserve Bank of San Francisco and the Stanford Institute for Economic Policy Research, Governor Meyer described the behavior of recent U. S. monetary policy as follows:

> In the fall of 1998, monetary policy responded both to the financial market distress and to the abrupt change in the forecast for growth... Once it became apparent that the U.S. economy was maintaining its momentum despite weaker foreign growth and that financial markets had returned toward normal, *the growing uncertainty about the output gap—reflecting the continuing contradiction of declining inflation and rising output gaps—made monetary policymakers cautious...*
>
> Beginning in mid-1999, with the estimated output gap increasing further and growth shifting to a still-higher gear, policymakers became more concerned about the possibility of overheating and, hence, the risks of higher inflation [and subsequently tightened policy significantly]...
>
> *Why did policymakers tolerate for a while further increases in the output gap, and why did they subsequently become more concerned about the inflation risks from further increases in the output gap? I think the change can be rationalized in terms of my discussion of the case for a nonlinear policy response under uncertainty.* As the unemployment rate fell farther below the best estimates of the NAIRU and the risk of overheating increased, policymakers became less tolerant of continued above-trend growth.
>
> (Meyer (2000), emphasis added)

The models of this paper thus may help us to better understand the motivations and performance of monetary policy, especially in light of the recent uncertainties about trend productivity growth and sustainable rates of labor and capacity utilization that U. S. monetary policymakers have faced.

5.1 More General Models

The results of this paper also apply more broadly than to just the simple illustrative model (1) considered throughout the main text. It is clear that the same ideas of signal extraction and updating with a non-normal prior apply to models with a greater number of indicator variables, with results that will have the same qualitative features as those discussed in the text—in this context, the "size of the surprise" referred to throughout the paper must be redefined as the overall size of the multivariate surprise, as in section 4.

It is also true that the results of this paper generalize very easily to any model that exhibits separation of estimation and control, including forward-looking models that possess this property.[25] The story being told in sections 3 and 4 of this paper is essentially one of estimation, so that as long as the control stage of the problem is separate from that of estimation, the same types of diagrams presented earlier will remain very relevant for thinking about the problem, and the same intuition applies.

5.2 Some Caveats

A few final caveats about the implementation and importance of such a nonlinear policy prescription are in order.

The examples above clearly illustrate the point that a quantity of primary importance for policymakers is the size of the surprise in observed variables that they face, relative to their prior standard deviation on that surprise. Correlations between the shocks in the model are also important in determining the magnitude of the overall surprise. Small surprises will tend to lead policymakers to a cautious response to observed unemployment in this framework, while larger surprises will tend to lead to an increasingly aggressive response of policy to unemployment at the margin.

This naturally raises the question as to what exactly constitutes a "small" or a "large" surprise. On the basis of the examples considered above, it appears that if the surprise is one of roughly 1 standard deviation, then it is not "large," yielding revisions and interest rate responses that are very well approximated by a linear (though possibly somewhat attenuated) function. However, if the surprise is instead taken to be one of roughly 2 standard deviations or more, then the nonlinearities in policymakers' optimal response become more pronounced. Thus, a guideline of roughly 2 standard deviations may serve as a starting rule of thumb for determining whether a surprise is "large" or "small," though it obviously should be taken with a substantial grain of salt.

In addition, it is clear that the baseline model (1) used in the examples above is extremely simple, and not representative of the full range of variables with which policymakers are actually concerned. A model more representative of that actually used by

[25] Pearlman, Currie, and Levine (1986) and Svensson and Woodford (2000) show that the forward-looking linear-quadratic model possesses this property.

policymakers (either formally or informally) would likely involve additional variables such as the prices of oil, commodities, exchange rates, assets, and a term structure of interest rates, to name just a few. Incorporating additional variables such as these into the model would explain correspondingly more variation in unemployment and inflation, and lead to a very different set of surprises in these equations than result from the very simple model specified in equations (1), plus lead to a whole new set of surprises in these other variables under consideration. Thus, again, it becomes very unclear what variances should be used when evaluating the relative magnitude of the surprise facing policymakers in practice.

6. Conclusions

Can optimal nonlinear policy describe the behavior of the Federal Reserve in the 1990s? The answer appears to be "yes," for a class of models in which policymakers' priors about the natural rate of unemployment are diffuse in a region around the mean. Optimal policy in this case is characterized by a cautious response to small surprises in observed unemployment and inflation, but an increasingly aggressive response to surprises in these variables at the margin. These features appear to match very well statements by Federal Reserve officials, and the historical behavior of the Fed, in the late 1990s.

Whether this nonlinearity is quantitatively significant in practice depends on the size of the overall "surprise" contained in the realized values of the observable variables, with two standard deviations not being an unreasonable benchmark for quantitative significance. Note that a "large" surprise in the examples above warranted both a large revision in policymakers' beliefs, *and* a more aggressive policy response at the margin.

Finally, it should be emphasized that the results demonstrated in this paper are not specific to models incorporating a Phillips curve type of relationship, but rather apply more generally to any signal extraction problem with non-normal priors on a key unobserved variable of the model. Since much of policymakers' current uncertainty about the state of the economy involves updating priors about unobserved variables through signal extraction, the effects in this paper should be regarded as being applicable to a fairly wide variety of situations.

Appendix: Mathematical Derivations

Assume policymakers' prior (time $t-1$) distribution on u^* is given by the density function $f(u^*)$. The mean of this distribution is assumed to exist, and without loss of generality, equals 0. The variance is assumed to exist and equals σ^{*2}. Given u^*, the (observable) unemployment rate u_t is distributed $N(u^*, \sigma^2)$, with density denoted by $\phi(u_t - u^*)$. I consider in this appendix what qualities of the density f lead to the prescriptions for policy given in the text.

I restrict attention to the case of one observable variable, u_t, for clarity. The analysis for multiple observable variables is essentially identical—one need only calculate the size of the joint surprise to determine the extent of updating of policymakers' priors that takes place.

The formula for policymakers' posterior on u^*, $E_t u^* \equiv E[u^*|u_t]$, is the usual:

$$E[u^*|u_t] = \frac{\int u^* \phi(u_t - u^*) f(u^*) du^*}{\int \phi(u_t - u^*) f(u^*) du^*} \tag{A1}$$

Policymakers set interest rates r_t based on the estimated unemployment gap, $u_t - E[u^*|u_t]$. The marginal responsiveness of policy to observed unemployment u_t is therefore given (up to a minus sign) by:

$$1 - \frac{\partial E[u^*|u_t]}{\partial u_t} = 1 - \frac{1}{\sigma^2}\mathrm{Var}[u^*|u_t] \tag{A2}$$

where the equality in (A2) follows from (A1) by differentiation through the integral. Note that for f normally distributed, $\mathrm{Var}[u^*|u_t]$ is independent of u_t, so that the marginal responsivenss of policy is constant.

The concavity or convexity of the optimal policy response is given by:

$$-\frac{\partial^2 E[u^*|u_t]}{\partial u_t^2} = -\frac{1}{\sigma^4}\mathrm{Skew}[u^*|u_t] \tag{A3}$$

The text emphasizes two features of optimal nonlinear policy that are thought to be particularly relevant: 1) the responsiveness of policy is attenuated (*i.e.,* cautious), relative to normal, for small surprises ($u_t \approx 0$), and 2) the marginal responsiveness of policy increases in the size of the surprise (u_t).

By equation (A2), property 1) above is equivalent to $\mathrm{Var}[u^*|u_t]$ being larger than normal when evaluated at $u_t = 0$.[26] I begin with the assumption that f is Gaussian, and consider what perturbations increase the conditional variance in (A2). This is essentially a calculus of variations problem. The following diagram serves as a useful illustration:

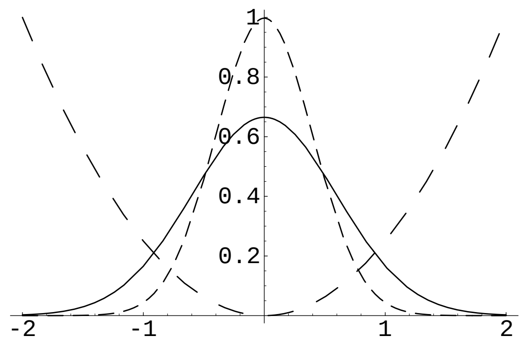

Figure A1

where the solid line in the figure depicts the density f, the short-dashed line ϕ, and the long-dashed line the function $y = x^2$. The quantity $\mathrm{Var}[u^*|u_t = 0]$ is (up to a constant factor) the integral of the product of these three functions:

$$\mathrm{Var}[u^*|u_t\!=\!0] = \frac{\int u^{*2}\phi(-u^*)f(u^*)du^*}{\int \phi(-u^*)f(u^*)du^*} \tag{A4}$$

Holding the denominator constant, we can think of increasing the value of this quantity by reducing f where the product $u^{*2}\phi(u^*)$ is small, and increasing f where the same product is large. Thus, we reduce f near 0, and increase f for intermediate values of u^*. Such an f will satisfy property 1) (attenuation near 0), as desired.

Analytically, we can prove these assertions via a variational argument on f. Consider perturbing the function f by an amount δ over the interval $[\xi\!-\!h, \xi\!+\!h]$, and call this new function f_δ. Note that we need not renormalize f_δ to have unit mass, as the denominator of (A4) makes this renormalization irrelevant. We have:

$$\begin{aligned} \frac{\partial}{\partial\delta} \int u^{*2}\phi(u^*)f_\delta(u^*)du^* \; &= \; \frac{1}{\delta}\int_{\xi-h}^{\xi+h} u^{*2}\phi(u^*)\delta du^* \\ &\approx \; 2h\xi^2\phi(\xi) \end{aligned} \tag{A5}$$

[26] There is actually nothing that prevents (A2) from being negative (as in the two-regime model in the text). In this case, I am interpreting a further decrease in (A2) (toward minus infinity) as the desired policy attenuation.

Similarly,

$$\frac{\partial}{\partial \delta} \int \phi(u^*) f_\delta(u^*) du^* \;=\; \frac{1}{\delta} \int_{\xi-h}^{\xi+h} \phi(u^*) \delta du^* \qquad \text{(A6)}$$
$$\approx\; 2h\, \phi(\xi)$$

The derivative of (A4) with respect to the δ-perturbation is then:

$$\frac{2h\xi^2\phi(\xi) \int \phi(u^*) f(u^*) du^* \;-\; 2h\,\phi(\xi) \int u^{*2} \phi(u^*) f(u^*) du^*}{\left(\int \phi(u^*) f(u^*) du^* \right)^2} \qquad \text{(A7)}$$

This quantity is less than 0 for $\xi = 0$. It is actually *greater* than 0 as $|\xi|$ tends toward infinity, though the effect is very small (because $\phi(\xi)$ is so small). More specifically, it is negative for $\xi^2 < E[u^{*2}|u_t = 0]$ and positive for $\xi^2 > E[u^{*2}|u_t = 0]$. Thus, perturbing f downward at 0, and pushing it upward for intermediate and large values of $|\xi|$, leads optimal policy to have property 1) emphasized in the text, as was to be shown.[27]

We can use the same technique to determine what distributions f possess property 2) (an increasingly aggressive response to unemployment at the margin). We know from equation (A3) that this is equivalent to Skew$[u^*|u_t]$ having sign opposite that of u_t—*i.e.*, policymakers' posterior distribution on u^* must be skewed toward 0.

Begin again with a density f that is Gaussian with variance σ^{*2}, and consider the perturbation f_δ, which equals f except over the interval $[\xi - h, \xi + h]$, where it is greater by an amount δ. A straightforward calculation yields:

$$\frac{\partial}{\partial \delta} \text{Skew}[u^*|u_t] \approx$$
$$\frac{2h\,\phi(u_t - \xi)}{\int \phi(u^*) f(u^*) du^*} \left[\xi^3 - 3\xi^2\mu - 3\xi E[u^{*2}|u_t] + 3\mu \text{Var}[u^*|u_t] + 6\mu^2 - \mu^3 \right] \qquad \text{(A8)}$$

where $\mu \equiv E[u^*|u_t]$. Note that the operators E and Var in (A8) are taken with respect to the base Gaussian density f. Since Skew$[u^*|u_t] = 0$ for all u_t when $\delta = 0$, we wish to know for what values of ξ the quantity (A8) is negative for all values of u_t in which we are interested. This requires an analysis of the cubic polynomial in ξ on the right-hand side of (A8).

[27] Note that equation (A7) holds for f non-Gaussian as well, so that the variational argument is valid for f significantly different from a normal distribution, as well as locally.

Since we are starting from a Gaussian base density f, we know that $\mu = \dfrac{\sigma^{*2}}{\sigma^2 + \sigma^{*2}} u_t$ and $\mathrm{Var}[u^*|u_t] = \dfrac{\sigma^2\sigma^{*2}}{\sigma^2 + \sigma^{*2}}$. For values of the surprise u_t that are not exceedingly large (much greater than 6, relative to its mean of 0), the constant (in ξ) term $3\mu\mathrm{Var}[u^*|u_t] + 6\mu^2 - \mu^3$ is necessarily positive. This implies that (A8) evaluated at 0 is postiive.

The limit of (A8) as ξ tends to infinity is also positive. Examination of signs of the coefficients of the ξ polynomial reveals that there are exactly two positive roots. Thus the ξ polynomial begins above 0, trends smoothly down into the negatives, and then trends smoothly back up into the positives, where it asymptotes to ξ^3.

To attain $\mathrm{Skew}[u^*|u_t] < 0$, it follows from this analysis that we want to perturb f downward near 0, downward at the tails, and upward for intermediate values of ξ. These are again exactly the diffuse-middled types of distributions described in the text. Note that in contrast to property 1), for property 2) the presence of thin tails in policymakers' prior distribution helps to generate the desired effect.

References

Bertsekas, Dimitri. *Dynamic Programming* (Prentice-Hall: Englewood Cliffs, NJ, 1987).

Ferguson, Roger. "Interview with Roger W. Ferguson, Jr." *The Region* 14 (June 2000), Federal Reserve Bank of Minneapolis, article by David Levy.

Filardo, Andrew. "New Evidence on the Output Cost of Fighting Inflation," *Federal Reserve Bank of Kansas City Economic Review* 83 (1998), pp. 33–61.

Meyer, Laurence. "Q & A on the Economic Outlook and Challenges Facing Monetary Policy," speech before the Philadelphia Council for Business Economics, Federal Bank of Philadelphia, September 8, 1999; transcript provided on the Federal Reserve Board's public web site.

Meyer, Laurence. "Structural Change and Monetary Policy," speech before the Joint Conference of the Federal Reserve Bank of San Francisco and the Stanford Institute for Economic Policy Research, Federal Reserve Bank of San Francisco, March 3, 2000; transcript provided on the Federal Reserve Board's public web site.

Orphanides, Athanasios and Volker Wieland. "Inflation Zone Targeting," *European Economic Review* 44 (2000), pp. 1351–87.

Pearlman, Joseph, David Currie, and Paul Levine. "Rational Expectations Models with Partial Information," *Economic Modelling* 3 (1986), pp. 90-105.

Reifschneider, David and John Williams. "Three Lessons for Monetary Policy in a Low Inflation Era," *Federal Reserve Board Finance and Economics Discussion Series* 1999–44 (1999).

Sack, Brian and Volker Wieland. "Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence," *Journal of Economics and Business* 53 (2000), pp. 205-28.

Swanson, Eric. "On Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules," *Federal Reserve Board Finance and Economics Discussion Series* 2000–32 (2000).

Svensson, Lars and Michael Woodford. "Indicator Variables for Optimal Policy," *NBER Working Paper* 7953 (2000).