

Monetary policy, parameter uncertainty and optimal learning

Volker Wieland*

Board of Governors of the Federal Reserve System, Washington, DC 20551, U.S.A.

Abstract

Since central banks have limited information concerning the transmission channel of monetary policy, they are faced with the difficult task of simultaneously controlling the policy target and estimating the impact of policy actions. A tradeoff between estimation and control arises because policy actions influence estimation and provide information which may improve future performance. I analyze this tradeoff in a simple model with parameter uncertainty and conduct dynamic simulations of the policymaker's decision problem in the presence of the type of uncertainties that arose in the wake of German reunification. A policy that separates learning from control may induce a persistent upward bias in money growth and inflation, just as observed after unification. In contrast, the optimal learning strategy which exploits the tradeoff between control and estimation significantly improves stabilization performance and reduces the likelihood of inflationary bias.

Key words: optimal control with unknown parameters, Bayesian learning, monetary policy, structural change, learning by doing, inflation targeting

JEL classification: E52, E40, D83, C44

*Correspondence: vwieland@frb.gov, Fax: (202) 452-2301, Phone: (202) 736-5620

This paper is a substantially revised version of chapter 2 of my Ph.D. dissertation at Stanford University. I would like to thank John B. Taylor, Kenneth Judd, John Leahy, Ronald McKinnon, Michael Horvath, Thomas Sargent, Andrew Levin, Athanasios Orphanides, Richard Porter, Thomas Cosimano, Ronald Balvers, Jürgen von Hagen, Manfred Neumann and David Wilcox for many helpful suggestions and Geoff Lanyon for assistance. I have also benefitted from discussions with seminar participants at Stanford University, the Federal Reserve Board, the University of Bonn and the Bundesbank. All remaining errors are my own. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting those of the Federal Reserve Board. Financial support at an early stage through a Bradley dissertation fellowship from the Center of Economic Policy Research at Stanford has been strongly appreciated. Furthermore, I would like to thank the Center for European Integration Studies at Bonn University for the hospitality extended while I was revising this paper.

1. Introduction

Monetary policy is conducted in an uncertain economic environment where policymakers have limited knowledge of the complex and unstable relationship between policy instruments such as the short-term nominal interest rate and policy targets such as the rate of inflation. In fact, macroeconomists can at best provide a rough estimate of the economy's response to a specific policy action. As a consequence, monetary policymakers are confronted with a complex simultaneous control and estimation problem. On the one hand, they attempt to control target variables such as inflation as best as possible based on current knowledge and information. On the other hand, they have to examine and perhaps revise their view of how the economy responds to policy actions as soon as new information becomes available and then adjust policy accordingly. This is particularly important at times when the economy is undergoing structural changes such as, for example, following German unification or the formation of the European monetary union.

Recent research on monetary policy and transition such as Bertocci and Spagat (1993) and Balvers and Cosimano (1994) has emphasized that policymakers face a tradeoff between control and estimation whenever they are uncertain about parameters that influence policy effectiveness. This tradeoff arises because policy actions may provide new information about the relationship between policy instrument and target. In particular, it may be optimal to choose a policy setting that worsens current outcomes but yields new information that will make it possible to improve policy performance in the future. Thus, how quickly policymakers will learn about relevant parameters following structural changes such as German unification will depend, among other factors, on their own actions. For example, if the policymakers' estimates are biased and policy is set without considering how it will affect these estimates, one might well expect to observe a persistent deterioration in stabilization performance. Alternatively, if policymakers take into account the informational effects of policy actions and learn optimally, long-lasting biases in target variables should be less likely. Such optimal policy may incorporate a significant degree of experimentation.

This paper investigates the likelihood of a persistent bias in monetary policy due to parameter uncertainty and explores to what extent optimal policy improves the speed of learning.

The analysis is based on a dynamic programming framework with Bayesian learning about unknown parameters. The asymptotic behavior of parameter estimates and policies in such a framework has been studied by Easley and Kiefer (1988) and Kiefer and Nyarko (1989). Wieland (1999) presents a numerical algorithm for computing the optimal policy in such a Bayesian learning model. Here, I use this methodology to study the quantitative implications of parameter uncertainty and learning for monetary policy.

The practical relevance of such learning considerations in the aftermath of structural changes is shown by studying monetary policy after German unification. Economic union with East-Germany generated substantial uncertainties concerning the transmission of monetary policy and subsequent years witnessed a significant increase in money growth and inflation. I examine under what conditions learning dynamics could lead to such a persistent increase in inflation using a simple model of monetary policymaking. This exercise has a positive and a normative dimension. Actual central bank behavior under parameter uncertainty may be better represented by a policy that treats control and estimation separately. Such an approach corresponds to the cautious, gradualist policy recommended originally by Brainard (1967) and widely discussed in the policy literature (see for example Blinder (1998)). It implies that the policy action which would be optimal given available estimates of the relationship between policy instrument and target, is adjusted to account for the degree of precision of these estimates. This approach involves learning albeit in a passive manner. Once new data becomes available the policy stance changes to account for changes in the parameter estimates and the associated degree of precision. This passive learning policy is myopic because it neglects the information effect of policy actions. Normative implications arise from a comparison to the optimal learning policy that exploits this information effect.

The paper is structured as follows. Section 2 reviews related literature. Section 3 formulates the basic control problem with unknown parameters in a dynamic programming framework with Bayesian learning. Convergence of parameter estimates and policy actions is shown by Easley and Kiefer (1988) and Kiefer and Nyarko (1989). They also show that the limit beliefs of the decision maker need not coincide with the true parameter values generating the data. Using numerical methods I compute policy functions that map the policymaker's beliefs about the unknown parameters into policy actions. I consider two unknown parameters and model

beliefs as a bivariate normal distribution. This specification is conducive to practical applications since it implies that the unknown parameters may be any real number from the perspective of the policymaker. The myopic, passive learning policy can be calculated analytically, while the optimal policy has to be approximated numerically, because Bayesian learning introduces a nonlinearity into the dynamic optimization problem.¹ Section 4 reviews the impact of German unification on money growth and inflation, the performance of the Bundesbank's monetary targeting strategy, and discusses the empirical evidence for structural changes in money demand and supply relationships. Section 5 provides an interpretation of the basic learning and control problem in terms of German monetary policy and describes the calibration of the learning model. Section 6 presents dynamic simulations of this model under passive and optimal learning policies. These simulations show that passive learning can result in a persistent upward bias in money growth and inflation due to mistaken beliefs about money demand and supply parameters. The optimal policy, which is approximated by numerical methods, substantially reduces the likelihood of such a policy bias, albeit at the expense of somewhat higher initial variability. I also provide a detailed sensitivity analysis. Section 7 concludes.

2. Literature Review

This paper is related to a large literature on optimal decision-making under parameter uncertainty. In a well-known paper Brainard (1967) using a model without learning showed that multiplicative parameter uncertainty provides a rationale for a cautious, gradualist approach to policymaking. Prescott (1972) studied a multi-period control problem with learning about one unknown parameter by means of numerical methods, yet found little difference between passive and optimal learning except under a high degree of uncertainty. The dynamic programming algorithm used in this paper is very similar in spirit to the approach taken by Prescott (1972). However, I consider a problem with two unknown parameters, which allows for more complex

¹ The numerical algorithm is described in the appendix. This specification results in a dynamic programming problem with five state variables (the two means, the two variances and the covariance of the normal distribution). Numerical dynamic programming problems of this scale are rarely dealt with because the so-called curse of dimensionality implies almost prohibitive computational costs. Approximations of the optimal policy functions used in this paper typically require several days computing time on a Sparc20 Sun workstation.

learning dynamics including asymptotic convergence to incorrect limit beliefs. Thus, the paper bridges a gap between the computational literature on optimal experimentation and the theoretical literature on incomplete learning (for example Easley and Kiefer (1988) and Kiefer and Nyarko (1989)).² An alternative computational approach that originated from engineering has been applied in the dual control literature (for example Tse and Bar-Shalom (1973), Kendrick (1981), (1982), Mizrach (1991), Amman and Kendrick (1995) and others). An advantage of the algorithm used here is that it approximates policy and value functions for this nonlinear dynamic programming problem over a wide range of the state space and is flexible to allow for non-differentiabilities in the value function and discontinuities in the optimal policy. A more detailed discussion of these properties and a comparison to the dual control literature can be found in Wieland (1999). The main drawback is that the algorithm used here quickly falls victim to the “curse of dimensionality” when applied to models which require more state variables, either to describe beliefs about additional unknown parameters or to account for other dynamics in addition to learning.

Following the theoretical literature on optimal learning researchers concerned with monetary policy such as Bertocci and Spagat (1993) and Kasa (1996) have argued that the optimal policy under parameter uncertainty is more activist than a cautious, Brainard-type policy.³ Bertocci and Spagat studied the characteristics of such an optimal policy under fairly stringent restrictions on the type of parameter uncertainty. Balvers and Cosimano (1994) used a more general learning model with two time-varying parameters to study the optimal transition from high to low inflation. They found the optimal policy to exhibit some degree of gradualism. Their analysis however depends on restrictions on the type uncertainty and the form of the policymaker’s objective function, which render the dynamic programming problem to be solved deterministic. Furthermore, none of these papers provides quantitative results regarding the likelihood of a policy bias due to mis-specified beliefs and incomplete learning. Finally, it should be noted that optimal policy under uncertainty also has a strategic dimension that is not considered by the above authors nor in this paper. Caplin and Leahy (1996) show that in a game

² Early research on asymptotic learning behavior includes Taylor (1974) and Anderson and Taylor (1976).

³ Earlier work on policy and parameter uncertainty includes Craine (1979) and Tinsley and von zur Muehlen (1981).

between a policymaker who attempts to stimulate the economy and is uncertain about the response of potential investors, a cautious policy move may be ineffectual, because it may lead investors to wait for further policy easings.⁴

Before moving on, it is useful to compare the learning and control problem in this paper to macroeconomic models with least-squares learning in the vein of Marcet and Sargent (1989). There is an analogy between the passive learning strategy considered here and bounded rationality as it is used in macroeconomics. In the first case, policymakers neglect to take into account the effect of current policy actions on their beliefs and through this channel on future policy performance. In the latter framework, market participants, when forming their expectations about future outcomes do not take into account that the process of expectation formation affects future outcomes. Sims (1988) and Chung (1990) have applied the bounded rationality approach to the central bank's decision problem under parameter uncertainty about the Phillips curve. They show that a naive learning rule may generate persistent econometric illusions about the nature of the Phillips curve. However, in their framework the policymaker does not take into account the effect of policy on the speed of learning.

3. A simple model of learning and control

The basic decision problem faced by the policymaker in this paper is to control a target variable (π), when the relationship between the policy instrument (i) and the target variable is a linear regression equation with two unknown parameters (β_0, β_1):

$$\pi_t = \beta_0 + \beta_1 i_t + \epsilon_t \quad (1)$$

At time t the policymaker chooses a value for the instrument based on current beliefs about the unknown parameters. Then a shock ϵ_t occurs and a new observation becomes available. The

⁴ An alternative approach to study this strategic dimension is taken by Basar and Salmon (1990).

shock term is assumed to be normally distributed with mean zero and known variance σ^2 .⁵ Before choosing next period's instrument setting, the policymaker updates the parameter estimates based on this new data. The policymaker's beliefs with respect to the vector of unknown parameters $\beta=(\beta_0,\beta_1)'$ are modeled as a bivariate normal distribution:

$$p(\beta | \cdot) = N\left(b_t, \Sigma_t\right) \quad \text{where} \quad \Sigma_t = \begin{pmatrix} v_0 & v_{01} \\ v_{01} & v_1 \end{pmatrix}_t \quad (2)$$

All information that is available at the beginning of period t is summarized in the vector of means $b=(b_{0,t},b_{1,t})$ and the variance-covariance matrix Σ (with variances $(v_{0,t},v_{1,t})$ and covariance v_{01}). At the end of the period the new data is used to update beliefs according to

$$\Sigma_{t+1} = \left[\Sigma_t^{-1} + X_t'X_t \right]^{-1} \quad b_{t+1} = \Sigma_{t+1} \left[X_t' \pi_t + \Sigma_t^{-1} b_t \right] \quad (3)$$

where $X_t=(1,i_t)$. A derivation of these updating equations from Bayes rule can be found in Zellner (1971).⁶ They represent a learning channel, through which the current policy action i_t affects next period's estimates of the unknown parameters and thus indirectly future realizations of i and π . Both, the instrument and the target variable may appear in the policymaker's loss function, which I denote here by $U(\pi, i)$. In section 5, U is defined to be standard quadratic loss function. Taking expectations with respect to β and ϵ , one obtains the expected one-period loss $L(i, b, \Sigma)$,

$$L(i, b, \Sigma) = \int_R \int_{R^2} U(\beta_0 + \beta_1 i + \epsilon, i) p(\beta | b, \Sigma) q(\epsilon) d\beta d\epsilon \quad (4)$$

where $q(\epsilon)$ represents the normal density function of the shocks. A policymaker who deals with control and estimation separately will proceed in two steps. In a first step, policy is set to minimize expected one-period loss based on current beliefs. In a second step, new observations are used to update beliefs and the policy setting is adjusted accordingly. Such a passive learning

⁵ This assumption is convenient for computational purposes and standard in the literature. If prior beliefs are represented by a normal distribution and the shocks are normally distributed with known variance, then the posterior belief is also a normal distribution.

⁶ They can also be derived by applying the Kalman filter to recursive least squares (see Harvey (1993)).

strategy is myopic, because it disregards the effect of policy actions on future beliefs.

A stationary policy is defined as a function $H(b, \Sigma)$, which selects an action i based on the current state (b, Σ) . Given a specification of the policymaker's loss function, the myopic policy $H^{my}(b, \Sigma)$ which minimizes $L(i, b, \Sigma)$ can typically be derived analytically. This is not the case for the optimal policy $H^{opt}(b, \Sigma)$ which minimizes the discounted sum of expected current and future losses:

$$\begin{aligned} \text{Min } E \left[\sum_{j=0}^{\infty} \delta^j L(i_{t+j}, b_{t+j}, \Sigma_{t+j}) \mid b_t, \Sigma_t, H \right] \\ [i_{t+j}]_{j=0}^{\infty} \end{aligned} \quad (5)$$

s.t. (3)

where δ denotes the discount factor. If β were known, this dynamic optimization problem would reduce to a sequence of static problems and the myopic policy would be equivalent to the optimal policy. However, with β unknown, beliefs change over time and form an explicit link between present and future periods. Estimation and control then cannot be separated because future beliefs $((b, \Sigma)_{t+j}, j=1, 2, \dots)$ depend on the whole history of policy actions $(i_{t+j-1}, j=1, 2, \dots)$ and thus on the policy function H . The effect of policy on future beliefs and the expectations operator is also apparent from the Bellman equation associated with this dynamic programming problem:

$$V(b, \Sigma) = \text{Min}_i \left[L(i, b, \Sigma) + \delta \int V(B(i, b, \Sigma, \beta_0 + \beta_1 i + \epsilon)) p(\beta \mid b, \Sigma) q(\epsilon) d\beta d\epsilon \right] \quad (6)$$

$V(b, \Sigma)$ denotes the value function and the two terms on the right-hand side characterize the tradeoff between control and estimation. The first term represents the expected current loss, while the second term is the expected continuation value in the next period which reflects the improvement in future performance due to improved information. Next period's beliefs have been replaced with the Bayes operator $B(\cdot)$ as a stand-in for the updating equations (3). They depend on the realization of the dependent variable $\pi = \beta_0 + \beta_1 i + \epsilon$. The conditional distribution of π is a

function of i , the beliefs $p(\cdot)$ and the distribution of the shocks $q(\cdot)$.

As shown by Easley and Kiefer (1988) and Kiefer and Nyarko (1989) a stationary optimal policy exists and the value function is continuous and satisfies the Bellman equation. Using standard dynamic programming methods, one can show that Blackwell's (1965) sufficiency conditions of monotonicity and discounting are satisfied and (6) has a fixed point in the space of continuous functions which is the value function $V(\cdot)$. Policy and value functions can be obtained by means of an iterative algorithm based on the Bellman equation starting with an initial guess about $V(\cdot)$. However, the integration in (6) can usually not be carried out analytically, because the updating equations are nonlinear. Numerical approximation involves high computation costs, because value and policy functions are five-dimensional functions and may exhibit important non-differentiabilities and discontinuities. These issues are discussed in more detail in Wieland (1999).

It remains to discuss the asymptotic properties of beliefs and policy actions which are the main focus of Kiefer and Nyarko (1989). Standard convergence results do not apply to this controlled regression. The reason is that along any sample path for which parameter estimates converge the sequence of actions may also be converging. But if actions converge rapidly they may not generate enough information for identifying the unknown parameters. Therefore it need not be the case that the limit distribution is centered on the true parameter values (i.e. limit beliefs do not necessarily coincide with $(\beta, [0])$). Kiefer and Nyarko (1989) prove convergence under fairly general conditions and characterize the set of possible limit beliefs and policies for the case of a simple regression. I will return to these issues in section 6 in the context of dynamic simulations of alternative policies.

4. Monetary policy, parameter uncertainty and learning after German unification

4.1. The impact of German unification

No other industrial country in recent history has experienced structural change comparable in magnitude to German unification. For West-Germany it implied an increase in population of 25 percent and an increase in GDP of about 10 percent. The introduction of the DM in the East German Länder in the fall of 1990 required a huge one-time expansion of the German money

supply as shown in the top left panel of figure 1. M3, for example, increased by almost 16 percent at that time.

Not surprisingly monetary policymakers in Germany were faced with a difficult control and learning problem following unification. Monetary union confronted the Bundesbank with uncertainties regarding the implementation and effects of monetary policy, which extended far beyond the determining the appropriate size of the adjustment in the money stock in 1990. For example, while the process of consolidating and integrating the rudimentary East German credit system into a modern two-tier banking system inevitably took some time, East German banks had to be immediately included into the Bundesbank's interest rate and liquidity management following the currency conversion. This transition created uncertainty concerning money multipliers and money supply management.⁷ Furthermore, unification introduced uncertainties regarding the demand for money. The Bundesbank initially had little information about the portfolio choices East Germans would make in the new financial environment with regard to money and other assets and the potential for velocity shifts and changes in medium- and long-run economic growth made it difficult to predict money demand with much precision. Policymakers also had to consider that unification might induce permanent changes in the output and inflation dynamics of the Germany economy.

4.2. The performance of monetary targeting

The Bundesbank has a long tradition of pursuing its long-run objective of price stability by means of a medium-term monetary targeting strategy. Since December 1974 it has set annual monetary targets, first for the growth rate of the central bank money stock, a measure that is

⁷ The Bundesbank subjected East-German banks to minimum reserve requirements as early as August 1990. However they had to adopt unconventional methods in providing central bank funds to the banks such as providing more generous discount quotas or giving permission to discount trade bills with lower-grade co-signers than what was required of West German banks. The instruments customary in West Germany such as securities repurchase agreements require that the banks put up assets as collateral such as high-quality trade bills and securities. East German banks did not have such assets in sufficient quantity and even a year after unification the provision of funds to East German banks through securities repurchase agreements only amounted to 23 percent of the total compared to 47 percent in West Germany. As East German banks were unfamiliar with the Bundesbank's money management methods, they initially tended to hold large amounts of liquid funds. Von Hagen (1993) discusses some of these problems in more detail and provides an econometric analysis of money supply and demand uncertainties associated with unification.

close to the monetary base, and since 1988 for the growth rate of M3. At times, the Bundesbank has announced a single target rate but more frequently an annual target range, which implied a widening cone for the growth rate of M3 throughout the year. Following unification M3 growth substantially exceeded the announced target range in 1991, 1992 and 1993, barely declined to the upper limit of the target range towards the end of 1994, then undershot the target range substantially in 1995, only to exceed the target range again in 1996. Only by 1997 M3 growth entered and remained within the announced target range.⁸ The increase in money growth in the early 1990s from about 6 to 9 percent, on average, was accompanied by an increase in inflation to about 4.5 percent from 1991 to 1994 – more than twice the pre-unification six-year average. The growth rates of M3, the CPI and GDP deflator before and after unification are shown in Figure 1 and the averages in Table 1.

(FIGURE 1 ABOUT HERE)

(TABLE 1 ABOUT HERE)

A more detailed look at these developments reveals the difficulties that arose in predicting and controlling money growth and inflation after German unification. In spite of the problems that may have been associated with engineering the initial level adjustment in the money stock in the second half of 1990, the Bundesbank was able to keep M3 growth within the announced target range of 3 to 5 percent for 1991 throughout the first half of that year. However, by the final quarter of 1991 money growth accelerated to an annualized rate of about 10 percent.⁹ This rapid increase led the Bundesbank to conclude in February 1992 that “current monetary growth is much faster than appears consistent in the somewhat longer term with the objective of monetary stability”.¹⁰ M3 continued to grow at an annual rate of 9.5 percent in 1992, more than twice the midpoint of the announced target range of 3.5 to 5.5 percent. This growth was attributed to a large extent to unforeseeable special factors and uncertainties in financial markets

⁸ A chart showing M3 growth relative to the target range is discussed in detail in the Bundesbank’s monthly report.

⁹ The annual growth rate from 90:Q4 to 91:Q4 was reported as 5.2 percent in the monthly report of February 1992.

¹⁰ Monthly Report of the Deutsche Bundesbank, Vol. 44, No. 2, February 1992, p. 13.

beyond the Bundesbank's control.¹¹

Following these developments the target range for 1993 was increased to a range of 4.5 to 6.5 percent. Nevertheless, M3 growth again exceeded the target and averaged 8.8 percent during 1993. M3 continued to increase very rapidly in the first half of 1994 at an average rate of over 10 percent. In the second half, however, money growth slowed dramatically and came in just below the upper edge of the 1994 target range of 4 to 6 percent by year's end. The Bundesbank then decided to stick with a target range of 4 to 6 percent for 1995 and reaffirmed its intention to persist with the concept of monetary targeting as well as the target variable M3, while noting that the definition of the relevant money stock had become more difficult partly due to the emergence of money market funds.¹² All through 1995, money growth remained distinctly below target, averaging 2 percent for the year. 1996 instead witnessed another bout of rapid growth in M3 reaching almost 8 percent – significantly above the widened target range of 4 to 7 percent for that year. Following this roller coaster ride the monthly report of January 1997 concluded that “monetary growth again found itself in troubled waters last year”, but emphasized that “the Bundesbank is abiding by its strategy of monetary targeting in spite of the persistent volatility of shorter-term monetary growth and the fact that the target was again not met”.¹³ By that time however, inflation had already declined to the pre-unification average of about 2 percent. Furthermore, by mid-1997 M3 reentered the target range and has remained within the range since then.

In summary, the period of 1991 to 1996 was characterized by four misses and two near misses of the monetary target range, all but one on the upper side of the range. Inflation more than doubled from about 2 percent in 1990 to about 5 percent in 1993 and only returned to pre-

¹¹As stated in the Monthly Report of the Deutsche Bundesbank, Vol. 44, No. 12, December 1992, p. 10, “This overshooting is due, as mentioned before, to some unforeseeable special factors beyond the Bundesbank's control - such as huge inflows of funds in the wake of the turmoil in the European Monetary System, the rise in cash holdings attributable to the new tax on interest income, the sizeable interest subsidies, the inverse interest rate pattern, and uncertainties in the financial markets which have tended to foster the growth of high-yielding short-term financial assets.”

¹²See the Monthly Report of the Deutsche Bundesbank, Vol. 47, No. 1, January 1995, pp. 25-26, for a discussion of these definition problems. As to the preceding growth in money, the Bundesbank emphasized at that time that “Transitional disruptive factors, as for instance the bloating of the money stock at the beginning of 1994 that was clearly attributable to special factors, therefore constituted no reason to abandon the strategy of monetary targeting”, (p. 24).

¹³Monthly Report of the Deutsche Bundesbank, Vol. 49, No. 1, January 1997, p. 18.

unification levels by 1995. Of course, deviations from the monetary target have occurred before and may be explained in different ways. The Bundesbank's monthly reports cited a variety of special factors as the source of unavoidable money growth subsequent unification. Here I suggest an explanation for persistent misses of the monetary and inflation targets that is based on uncertainty and learning by the policymaker about the economy's response to policy actions and is consistent with the concerns about special factors and unavoidable temporary deviations expressed in the monthly reports. I show that learning about changes in the economy's response to policy may take a very long time, even if all available information is used to update relevant parameter estimates. In updating these estimates, the policymaker attempts to distinguish between transitory shocks and changes in underlying parameters. The speed of learning in such situations depends on policy actions and I will show that a passive learning approach may therefore induce persistent deviations from policy targets.

4.3 Empirical evidence of structural change in money demand and supply

In order to develop an explanation of the sustained increase in money growth and inflation following unification based on a learning framework, I first need to document the uncertainties regarding monetary control at that time. Rather than conducting an empirical investigation myself, I simply summarize the diverse results available from the existing literature.

The prospect and then the implementation of German monetary union fueled the public debate on the stability of money demand and supply relationships and the viability of monetary targeting. As more and more data became available a literature on testing the stability of M3 money demand after unification developed with contributions by von Hagen (1993), Gerlach (1994), the OECD (1993) and (1994), Hansen and Kim (1995), Kole and Meade (1995), the Bundesbank¹⁴, Issing and Tödter (1995), Neumann (1996), Scharnagl (1998) and Wolters, Teräsvirta and Lütkepohl (1998) among others. The results from this research are rather mixed. The authors disagree whether or not M3 demand remained stable, and if not, whether a shift occurred in the long-run equilibrium relation between money and its determinants or in the short-

¹⁴ Monthly Report of the Deutsche Bundesbank, July 1995, Appendix: Empirical study of the stability of money demand in Germany, pp. 29-35, and more recently in the Monthly Report of the Deutsche Bundesbank, August 1997, Appendix: On the stability of money demand, pp. 27-32.

run dynamics, and whether it was a one-time level adjustment or a more permanent structural change. On the one hand, Issing and Tödter (1995), Scharnagl (1998) and the Bundesbank (1995, 1997) find little evidence that the long-run relation changed after unification. The Bundesbank studies also reject the hypothesis of a permanent shift in short-run dynamics and conclude that unification affected monetary growth only temporarily. On the other hand, von Hagen (1993) finds that velocity and multiplier forecast errors based on pre-unification estimates increased sharply after unification and detects instabilities in broad money demand and shifts in the money multiplier. Hansen and Kim (1995) find evidence of a structural break in the cointegration relationship of money demand at unification and Neumann (1996) considers a downward shift in the income elasticity of money demand. Useful comparisons of the results in the literature are provided by Kole and Meade (1995) and Wolters et al. (1998). Kole and Meade (1995) conclude that the equilibrium money demand relation appears to have been largely stable since unification despite some changes in short-run dynamics and a level shift in 1990. However, they caution that it may be too early to judge whether long-run money demand in Germany has truly remained stable. Wolters et al. (1998) find a clear structural break in 1990 but once this break is accounted for they judge the money demand relationship to be stable. They state that their results are at variance with the findings of the Bundesbank study on money demand but conclude that the Bundesbank's strategy of monetary targeting remains viable.

The proliferation of studies on M3 following unification and the substantial range of findings provide clear evidence that uncertainty about money demand and supply increased and remained high for a number of years. This uncertainty made it difficult for policymakers to distinguish between temporary and permanent changes and confronted them with a complex learning problem.

5. Interpreting the learning problem in terms of German monetary policy

5.1. The relationship between policy instrument and policy target

It remains to link the generic learning problem of section 3 to the uncertainties faced by monetary policymakers subsequent unification. Here, I provide a formal interpretation with an eye towards the Bundesbank's procedure for deriving its annual money growth target. Due to the

computational costs associated with numerically approximating the optimal policy several simplifying assumptions are inevitable. In particular, since the bivariate normal learning problem already requires five state variables for modeling learning dynamics, further state variables describing additional dynamics in output and inflation determination cannot be accommodated in the analysis. An advantage of these assumptions is that any finding of persistent deviations of inflation from target will be exclusively attributable to learning dynamics.

The Bundesbank's monetary targeting strategy as described in articles by Bundesbank economists (see Tödter and Reimers (1994) or Tödter and Ziebarth (1997)) and the Bundesbank itself¹⁵ is geared towards the long-term relationship between money and prices as reflected in the quantity theory.¹⁶ According to the quantity equation the product of the money stock, M , and the velocity of circulation, V , equals the product of the price level, P , and the real gross domestic product, Y . Written in logarithms (lower-case), this implies:

$$m_t + v_t = p_t + y_t \quad (7)$$

Following the so-called P-Star approach of Hallman, Porter and Small (1991), the above papers use equation (7) to define an equilibrium price level, P^* , which is the money stock M per unit of real production potential, Y^* , at the equilibrium velocity, V^* . Written in logarithms this implies $p^* = m - y + v^*$. In addition, they assume that there exists a stable long-run (cointegrating) money demand relationship such as

$$m_t - p_t = \alpha_0 + \alpha_1 y_t + \epsilon_{dt} \quad (8)$$

where ϵ_{dt} measures zero-mean short-run deviations from equilibrium. Using equations (7) and (8), equilibrium velocity v^* can be expressed as a function of the parameters (α_0, α_1) ¹⁷ and potential output, and one can solve for the equilibrium price level, p^* :

¹⁵ Monthly Report of the Deutsche Bundesbank, January 1992, pp. 27-28.

¹⁶ This strategy is also described by von Hagen (1993), Clarida and Gertler (1997) and Dornbusch (1997).

¹⁷ Issing and Tödter (1995) estimate the income elasticity of money demand, α_1 , in Germany at about 1.5. Given an average growth rate of potential output of about 2.5 percent, this would imply a trend decline in the velocity of circulation of 1.25 percent per year.

$$p_t^* = m_t - \alpha_0 - \alpha_1 y_t^* \quad (9)$$

Given its long-run inflation target, π^* , the Bundesbank derives an annual target for money growth, Δm^* , from (9), which takes due account of forecasts of the growth of production potential and the trend change in the velocity of circulation.

$$\Delta m_t^* = \pi^* + \alpha_1 \Delta y_t^* \quad (10)$$

If the policymaker succeeds in increasing the money stock at this rate, then the equilibrium price level p^* and, once inflation and output dynamics have expired, the actual price level p will grow at the target rate π^* . Of course, inflation and output dynamics may imply sluggish adjustment of actual inflation, π_t , to the growth rate of the equilibrium price level, Δp_t^* . Both, the Bundesbank (1992) and Tödter and Ziebarth (1997) posit an error correction equation for actual inflation such as

$$\pi_t = \lambda_0 \pi_{t-1} + (1 - \lambda_0) \Delta p_t^* + \lambda_1 (p_t^* - p_t) + \epsilon_{p,t} \quad (11)$$

where $0 < \lambda_0 < 1$ and $\epsilon_{p,t}$ is a zero-mean random price shock. The first two terms on the right-hand side of (11), lagged inflation and equilibrium inflation, are motivated by adaptive expectations of wage and price setters. The third term, the so-called price gap ($p^* - p$), can be decomposed into an output gap ($y - y^*$) and a liquidity gap ($v^* - v$). As argued by Hallman et al. (1989), under certain conditions this inflation equation is equivalent to a standard expectations-augmented supply curve or an accelerationist Phillips curve. As noted above, I will simplify the inflation equation (11) for the purposes of this paper by assuming that deviations from equilibrium inflation are i.i.d., that is, I assume $\lambda_0 = 0$, $y_t = y_t^* + \epsilon_{y,t}$ and $v_t = v_t^* + \epsilon_{v,t}$, where $\epsilon_{v,t}$ and $\epsilon_{y,t}$ are i.i.d. normally-distributed zero-mean random variables and the growth rate of potential output, Δy^* , is a constant.

To complete the monetary control problem it is necessary to recognize that monetary policy cannot exert direct control over money growth. The policy instrument that the Bundesbank influences on a daily basis is the interest rate on securities repurchase agreements, which is

closely matched by the day-to-day money market rate. I assume that the short- and medium-term relationship between the interest rate and money and inflation is governed by the standard liquidity effect, that is, lower interest rates imply higher money growth and inflation:

$$m_t = m_{t-1} + \gamma_0 - \gamma_1 i_t + \epsilon_{m,t} \quad (12)$$

The parameters γ_0 and γ_1 reflect short-run dynamics in money demand and money multipliers, while the random shock $\epsilon_{m,t}$ represents temporary special factors that cause fluctuations in money growth and inflation. Combining (9), (10), (11) and (12) one can derive the impact of policy on inflation and interpret the parameters and shocks in the simple learning model of section 3 in terms of some of the uncertainties that arose after unification:

$$\pi_t = \beta_0 + \beta_1 i_t + \epsilon_t \quad (13)$$

where $\beta_0 = \gamma_0 - \alpha_1 \Delta y^*$ $\beta_1 = -\gamma_1$ $\epsilon_t = \epsilon_{p,t} + \epsilon_{m,t} + \lambda_1 (\epsilon_{y,t} - \epsilon_{v,t})$

Equation (13) relates the parameters β_0 and β_1 to parameters that play a role in the monetary targeting strategy. The composite shock ϵ is a function of price, velocity, supply and monetary shocks. I will assume that it is normally distributed with zero mean and known variance σ^2 .

5.2. The policymaker's decision problem

The final ingredient that is needed to proceed with an analysis of optimal monetary policy is the policymaker's objective function. I use a standard quadratic loss function, which includes inflation deviations from target and the policy instrument with a relative weight ω . The latter term captures central banks' well-known tendency to smooth interest rates. Then, the interest rate is chosen to minimize the discounted sum of expected current and future losses

$$\text{Min}_{[i]_{t=0}^{\infty}} E \left[\sum_{t=0}^{\infty} \delta^t ((\pi_t - \pi^*)^2 + \omega i_t^2) \mid I_t \right] \quad (14)$$

where δ denotes the policymaker's discount factor. If all the parameters were known, the optimal

interest rate setting would be:

$$i_t = - \frac{\beta_1 (\beta_0 - \pi^*)}{\beta_1^2 + \omega} \quad (15)$$

If instead the parameters are uncertain, the policymaker faces the same joint control and estimation problem as in section 3. Her beliefs about β_0 and β_1 are modeled as a bivariate normal distribution $p(\beta_0, \beta_1 | I_t) = N(b_0, b_1, \Sigma)$, where b_0 and b_1 are the means and Σ the covariance matrix with variances v_0, v_1 and the covariance v_{01} . Bayes rule implies the updating equations

$$\Sigma_{t+1} = \left[\Sigma_t^{-1} + X_t' X_t \right]^{-1} \quad b_{t+1} = \Sigma_{t+1} \left[X_t' \pi_t + \Sigma_t^{-1} b_t \right] \quad \text{where } X_t = (1 \ i_t) \quad (16)$$

which correspond exactly to equation (3) in section 3.

One possible strategy for the policymaker is to treat control and estimation separately. The first step of such a passive learning strategy involves computing the optimal instrument setting taking the point estimates and the associated degree of uncertainty as fixed. The second step is to re-estimate the parameters based on the outcome of the policy decision. Because there are no links between current and future periods other than the updating of beliefs, this passive learning policy is equivalent to minimizing current expected loss given current beliefs on a period-by-period basis:

$$i_t = - \frac{v_{01,t} + b_{1,t} (b_{0,t} - \pi^*)}{v_{1,t} + b_{1,t}^2 + \omega} \quad (17)$$

Compared to the decision rule in (15), the parameters are replaced with the point estimates. Furthermore, the variance of the slope parameter $v_{1,t}$ and the covariance $v_{01,t}$ introduce an element of caution to policy as in Brainard (1967). The passive learning strategy is myopic, because it

neglects the information effect of policy actions. As discussed in section 3 the dynamically optimal policy would account for the effect of policy actions on future beliefs and future performance. Since this policy cannot be derived analytically due to the nonlinear updating equations, I use the numerical dynamic programming algorithm discussed in the appendix to approximate it numerically.

5.3. Calibration

Finally, to be able to simulate the learning dynamics under the passive and optimal strategies it is necessary to calibrate the model. In calibrating model parameters and initial beliefs, I take into account observed sample moments. In particular, I take into account that prior to unification, inflation, money and output growth roughly averaged 2, 6 and 2.5 percent respectively, while the call money rate was on average 5 percent (see the top part of Table 2). As discussed in section 4, inflation and money growth increased substantially after unification and only began to return to pre-unification levels after the Bundesbank had raised interest rates beyond 9 percent in 1992 and 1993.

The calibration is based on the following rationale. Relevant structural parameters may have changed with unification and as a consequence the policy stance may have been too easy. Thus, one can speculate that money growth and inflation could have been kept at pre-unification averages, if the policymaker had only raised interest rates more quickly to a higher level. Consistent with this rationale, I calibrate the values of the underlying “true” parameters after unification such that a short-term interest rate of about 10 percent is needed to keep inflation from rising beyond 2 percent and money growth beyond 6 percent respectively. Prior to the post-unification increase in money growth and inflation, short rates of about 5 percent may have seemed sufficient to keep inflation and money growth steady. Therefore, I calibrate the policymaker’s initial beliefs such that an interest rate of 5 percent is expected to be consistent with inflation remaining at 2 percent.

Table 2 provides a summary of the calibration. Unfortunately the restrictions implied by the above rationale are not sufficient to pin down all the relevant model coefficients. Regarding the true data-generating process I need to pick values for $(\beta_1, \sigma^2, \omega)$. Regarding the policymaker’s initial beliefs I choose values for (b_1, v_1, ρ) . Since these values are not based on a

specific rationale, I will investigate the sensitivity of the simulation results to alternative values later on. For the slope β_1 , I pick a value of -0.7. Thus, the interest rate has to be raised by about 1.5 percentage points to lower money growth and inflation by 1 percentage point. Given β_1 , the values of β_0 , Δy^* , α_1 , γ_0 and γ_1 are set to obtain the observed averages of money growth, output, inflation and interest rates as shown in the middle part of Table 2. Then, I choose values of the preference parameters, ω and π^* , that are consistent with the steady-state inflation and interest rate. I assume that the policymaker assigns a positive weight to interest variability, $\omega=0.14$, and a target rate π^* of zero which reflects the long-run goal of price stability. Finally, the variance of the shocks σ^2 is set to 1.

The rationale for setting initial beliefs is that the policymaker expects inflation to remain stable, if the interest rate remains close to the pre-unification level of 5 percent. Thus, given an estimate b_1 of the slope β_1 , the estimate b_0 of the intercept β_0 has to be consistent with an expected inflation rate of 2 percent. This feature is illustrated in Figure 2. The solid line represents the inflation impetus of policy under the true parameter values, while the dashed line is associated with a slope estimate b_1 of -0.8. It remains to specify the degree of uncertainty. As a benchmark I consider a scenario with fairly imprecise estimates. The variance of the slope estimate v_1 is set to 0.25. Together with an estimate b_1 of -0.8, this implies a t-statistic of 1.6. As shown in Table 2, the covariance v_{01} that would be consistent with the policymaker's expectation of stable inflation equals the product of v_1 and the perceived steady-state interest rate.¹⁸ Furthermore, for a given choice of the correlation coefficient ρ , one can derive the variance of the intercept v_0 that is consistent with v_{01} and v_1 .

In summary, once β_1 , σ^2 , ω , b_1 , v_1 and ρ are set, the values of the remaining model coefficients follow from the rationale outlined above. Later on I will investigate how sensitive the simulation results are to alternative values of these six coefficients.

(FIGURE 2 ABOUT HERE)

(TABLE 2 ABOUT HERE)

¹⁸ For a derivation of this result see Greene (1993) pp. 155-57.

6. Policy bias and stabilization performance

6.1 *Passive versus optimal learning*

Dynamic simulations of the calibrated model show that the passive learning policy often adjusts very slowly to a change in the unknown parameters and thereby may induce a persistent increase in money growth and inflation of the magnitude observed after German unification. An example of such a simulation is provided in Figure 3. The six different panels depict inflation, money growth and interest rate sequences as well as the evolution of the policymaker's beliefs about the unknown parameters for a given sequence of shocks. The policymaker's initial belief is as specified in Table 2 with the initial parameter estimates corresponding to the dashed line in Figure 2. In each panel the simulated sequences are shown with the steady state values to which they would converge, if the policymaker were to learn the true values of the unknown parameters β_0 and β_1 .

Turning first to inflation, the top left panel in Figure 3 shows the inflation rate rising quickly above 4 percent and remaining at this level for more than 30 periods before it returns to the originally intended steady-state of 2 percent. Two different inflation series are shown in this panel, the actual inflation rate observed by the policymaker as well as a "core" inflation rate that is obtained by subtracting the effect of temporary random shocks ϵ from the observed inflation rate. This measure directly reflects the inflationary bias that is induced by policy. As can be seen from the top left panel, the policymaker starts by raising interest rates somewhat during the first few periods, but falls substantially short of raising them to the level that would be required to keep inflation from rising. The required level is about 10 percent as determined in the calibration. Instead, interest rates remain around 7 percent for some time and as a consequence inflation and money growth pick up. Policy-induced money growth, that is the observed money growth series minus shocks, is shown to increase to about 8 percent in the top left panel.

(FIGURE 3 ABOUT HERE)

The reason why the policymaker does not immediately respond to the increase in inflation and money growth by raising interest rates further is related to the evolution of beliefs. The sequences of point estimates of the intercept and slope parameters b_0 and b_1 , are depicted in the middle left panel of Figure 3. For over 30 periods the policymaker believes that the intercept is

about 5.8, much smaller than the true value of 9, and the slope about -0.3, much flatter than the true slope of -0.7. This belief is reflected in the dashed curve in the bottom left panel which is based on the point estimates in the 20th period. As a result of these beliefs the central bank is pessimistic about its ability to avoid an increase in inflation. The flat slope estimate implies that interest rates much higher than 10 percent would be required to drive inflation immediately back down to 2 percent. The policymaker is reluctant to tighten policy as aggressively because of the interest rate smoothing term in his loss function and consequently regards some increase in inflation as unavoidable. The question is, why does the central bank adhere to this pessimistic belief for such a long time. As can be seen from the bottom left panel, the passive learning policy induces many observations on inflation at interest rate levels between 5 and 7.5 percent. These observations are not very informative and do not lead the policymaker to revise his incorrect beliefs. This is also reflected in the sequences of variances and covariances in the middle and bottom right panel which remain almost unchanged for this period of time. As a result, the passive learning policy reinforces incorrect beliefs and induces a persistent policy bias.

Figure 4 shows a dynamic simulation of the optimal learning policy with the same sequence of shocks and prior beliefs that underlie the passive learning simulation in Figure 3. The discount factor δ associated with this policy is 0.95. As can be seen from the top left and right panel of Figure 4, as soon as inflation rises, the optimizing policymaker experiments with a tighter monetary policy in order to find out whether the pessimistic belief regarding the relationship between policy and inflation is correct. The policymaker learns rather quickly that this belief is wrong, raises interest rates to about 10 percent and thereby reduces inflation.¹⁹ Thus, the optimal learning policy avoids the emergence of a sustained policy bias due to incorrect beliefs. The speed of learning is increased, although at the expense of somewhat higher initial variability. The evolution of point estimates and associated variances in the two middle panels of Figure 4 indicates that a few periods of active probing are sufficient to achieve a sharp

¹⁹ At this point it is appropriate to note again the medium-term nature of these simulations. Based on the Fisher effect, one would expect interest rates to decline again after inflation has returned to a 2 percent level. Incorporating this effect would have required a more general model that allows for further dynamics in addition to the learning dynamics. Because of the substantial increase in computing time associated with the optimal learning policy that would result from including further state variables to model such dynamics, I did not incorporate this effect in the model used here.

improvement in the precision of parameter estimates. As soon as the interest rate approaches the value that would be chosen under certainty, active probing stops because further improved knowledge about the unknown parameters does not translate into sizeable stabilization gains. This explains why the degree of uncertainty only declines very slowly from the 10th period on. The lower right panel shows that the first six observations are sufficient to obtain estimates that are close enough to the true parameter values to put policy on the right course.

(FIGURE 4 ABOUT HERE)

A remaining issue in interpreting these simulations is the meaning of a time period. What is the relevant frequency depends on data availability and the associated decision horizon. If the focus is on variables such as the GDP deflator as a measure of inflation, then a period may be a quarter. In the case of CPI and money growth, the relevant estimation and decision horizon may be a month.

6.2 *Incomplete learning and the emerging policy bias*

The policy bias that emerges under passive learning in the dynamic simulation in Figure 3 is closely related to the possibility of incomplete learning that has been investigated in the theoretical literature on Bayesian learning. The basic intuition behind this possibility is simply the following: if the policy instrument, which is the right-hand side variable in the regression equation (1), does not exhibit enough variation, it may not be possible to correctly identify the unknown parameters. Kiefer and Nyarko (1989) have shown that the policy action i_t and the posterior process of beliefs $(b, \Sigma)_t$ converge to a limit $(\bar{i}, (\bar{b}, \bar{\Sigma}))$ and that there are multiple limits which do not necessarily coincide with the true parameter values. For the case of a simple regression they have also shown that these limit belief and policy pairs exhibit three properties that can be used to characterize the set of possible limits: (1) the updating equations must have a fixed point at the limit belief and policy pair; (2) in the limit the mean of the target variable must be predicted correctly given the limit policy (which, however, does not mean that both parameters are correctly identified), and (3) the limit policy minimizes one-period loss given the limit belief. These properties imply a system of equations and inequality conditions with multiple solutions that is summarized in Table 3.

(TABLE 3 ABOUT HERE)

Each solution represents a belief that is self-reinforcing under the passive learning policy. All incorrect limit beliefs are associated with positive limit variances. Furthermore, the limit covariance v_{01} implies a limit correlation coefficient $\rho(=v_{01}/(v_0v_1)^{1/2})$ of +1 or -1. An example for such an incorrect limit belief and policy pair is:

$$\bar{i} = 7.2 \quad \bar{b} = \begin{pmatrix} 5.7927 \\ -0.2545 \end{pmatrix} \quad \bar{\Sigma} = \begin{pmatrix} 7.776 & -1.08 \\ -1.08 & 0.15 \end{pmatrix} \quad (18)$$

In fact, the policymaker's beliefs remain very close to this belief and policy pair during the period of biased policy shown in Figure 3. While the parameter estimates and policy instrument do not ultimately converge to this incorrect limit belief and policy pair, they vary around the values given in (18) for about 30 periods until they finally converge to the correct limit belief and policy. During this time the correlation coefficient ρ approaches -1 as shown in the lower right panel of Figure 3.

6.3 The likelihood of a policy bias

The learning dynamics that arise in the preceding two dynamic simulations depend of course on the specific sequence of shocks. Large-scale simulation exercises are necessary to assess the likelihood of a policy-induced inflation bias under passive versus optimal learning. For this purpose, I conduct 1000 simulations with identical prior beliefs but different sequences of shocks of a length of 100 periods that are drawn from an $N(0,1)$ distribution. The results of these simulations are reported in Table 4. First, I check how many of the simulations exhibit a policy-induced inflation bias that lasts at least as long as the bias that emerges in the dynamic simulation in Figure 3. More specifically, I count the number of simulations for which policy-induced inflation in period 30, that is inflation minus the effect of the temporary random shock ϵ is more than 1 percent above the steady-state inflation rate under certainty: $(\pi_{t=30} - \epsilon_{t=30} - 2) > 1$. As shown in the first row of the first column of Table 3, under passive learning about 30 percent of the simulations exhibit an inflation bias greater than 1 percent for at least 30 periods. The average inflation bias for these simulations during the first 30 periods is about 1.5 percent points. During the same period interest rates are on average set 2.3 percent points too low. The results in the third column of Table 4 show that the average inflation bias is about 0.75 percent

when averaging over all 1000 simulations. If the policymaker is learning optimally however, a policy-induced inflation bias is much less likely. As reported in the second column of Table 4, under optimal learning only about 6 percent of the simulations are characterized by a bias larger than 1 percent that persists for more than 30 periods.

TABLE 4 ABOUT HERE

6.4 Sensitivity analysis

The simulation results reported so far depend on the specific prior belief and also on the specific values of the underlying true parameters. The rationale for the calibration in section 5 does not uniquely pin down all of these values. As to the initial beliefs, the slope estimate b_1 , the associated variance, v_1 , and the correlation coefficient, ρ , were chosen freely. Similarly, I chose values for the variance of the shocks, σ^2 , the weight on interest rates in the policymaker's loss function, ω , and the true slope, β_1 . For this reason, I have conducted a large number of simulations to assess how the likelihood of a policy bias is affected by these assumptions.

Figure 5 summarizes the outcome of this sensitivity analysis. The six panels show the likelihood of a policy-induced inflation bias that is greater than 1 percentage point and lasts for more than 30 periods (the same measure as in Table 4) as a function of b_1 , v_1 , ρ , σ^2 , ω and β_1 respectively. The solid line always refers to outcomes under the passive learning policy with the outcome reported in Table 4 denoted by a circle. Outcomes under the optimal policy are represented by a dashed line with the outcome in Table 4 denoted by a cross. The top left panel of Figure 5 shows that the number of simulations that exhibit a bias increases with the degree of precision of the initial beliefs. I consider values of v_1 that would imply t-statistics for the slope estimate, b_1 , between 1 and 2.5. The intuition behind this result is the following. The more confident the policymaker is about the initial, incorrect parameter estimates, the less weight he gives to new data in revising his beliefs, and thus the more likely is a sustained policy bias due to incorrect beliefs. The top right panel of Figure 5 shows that a prior belief that implies a flatter slope estimate is more likely to lead to a sustained policy bias. The reason is that a flat slope represents a pessimistic view about how aggressive the policy tightening needs to be to keep inflation from rising. The flatter the slope, the more likely the policymaker will consider some increase in inflation as unavoidable.

The earlier discussion regarding incomplete learning suggested that the frequency with which a policy bias occurs will depend on how close in some sense the initial belief is to a self-reinforcing incorrect limit belief. Since all self-reinforcing incorrect beliefs are associated with a correlation coefficient ρ of -1 or 1, the initial value of ρ is a simple measure of closeness to such a self-reinforcing belief. The middle left panel confirms that the likelihood of a policy bias declines as the initial ρ is increased from -1 to 0.

Temporary random shocks prevent the policymaker from calculating the unknown parameters based on a few observations. One would expect that it becomes more difficult for the policymaker to detect changes in the unknown parameters, the larger the variance of these shocks. The evidence in the middle left panel of Figure 5 confirms that the likelihood of a policy bias due to incorrect beliefs increases with the variance of the shocks. The two lower panels show that a policy bias occurs more frequently, the higher the weight on interest rate variability in the policymaker's loss function, and less frequently, the further away the true value of the slope parameter is from the policymaker's estimate.²⁰ Note that changing any one of these six coefficients implies different values of some of the other coefficients which are derived in the same way as in Table 2.

(FIGURE 5 ABOUT HERE)

7. Conclusion

This paper shows that the tradeoff between control and estimation that arises in the presence parameter uncertainty can have quantitatively important implications for monetary policymaking, in particular following structural changes such as German unification. Focusing on the increase in money growth and inflation after German unification, a careful reading of the Bundesbank's monthly reports suggests that uncertainty and learning considerations regarding temporary versus permanent factors behind this increase were of special concern. Furthermore, the proliferation of studies on money demand following unification and the substantial range of findings indicate a high degree of uncertainty surrounding the determinants of money growth for

²⁰ I do not provide results under optimal learning in these cases, since computing numerical approximations of the optimal policy under alternative coefficient values would have involved substantial additional computation costs.

many years after unification. The paper shows how a basic model of Bayesian learning with two unknown parameters can be used to describe some of the decision problems faced by German monetary policymakers. Dynamic simulations of this learning model indicate that a passive learning strategy, which does not take into account fully the interaction between learning and control, may induce a persistent inflationary bias due to misspecified beliefs about the impact of policy on money growth and inflation. This bias may well be of similar magnitude and longer-lasting than the increase in inflation observed for several years after unification. A detailed sensitivity analysis confirms that the probability of such a bias is significant for a wide range of initial beliefs and values of the underlying parameters. A policy that optimally takes into account the learning channel, however, is found to experiment with tighter monetary policy, which substantially reduces the likelihood of a sustained policy bias due to incorrect beliefs. This finding shows that optimal learning can have important quantitative implications for stabilization policy. Of course, given the simplifying assumptions regarding the underlying model this finding should not be interpreted as a test of the optimality of German monetary policy. It should be noted, however, that the numerical dynamic programming algorithm developed for approximating the optimal policy allows for more complicated learning dynamics than previous studies of learning by central banks and could be used to study decision problems under parameter uncertainty in many other areas.

Several new questions arise from this research. First, the tradeoff between control and estimation may be relevant to monetary policymaking during other episodes of structural change. An example that immediately comes to mind is the European Central Bank's policy in the new European monetary union. As far as U.S. monetary policy is concerned one could think of the so-called "missing money" period in the mid 1970's, or more recently, of uncertainty about the natural unemployment rate (see Staiger, Stock and Watson (1997)) and the possibility of improvements in productivity trends. Second, the dynamic simulations in this paper show that parameter uncertainty and learning by the policymaker may induce nonstationary behavior in economic observables such as money growth and inflation. This effect is usually neglected in econometric studies of macroeconomic relationships even though it would call standard estimation

results into question. An exception is the study by Horvath (1991).²¹ Third, questions regarding central bank credibility and the interaction between the private sector's and the policymaker's learning behavior, especially when they do not share the same information set and have different beliefs, would seem to be of special interest.

²¹ Horvath uses the problem of a policymaker that maximizes the discounted sum of tax revenues subject to a Laffer curve with an unknown parameter to demonstrate that learning affects the exogeneity status of policy variables and proposes an extension of the standard econometric definition of exogeneity.

Appendix

The numerical dynamic programming algorithm

The algorithm for computing the value function and optimal policy relies on successive application of the functional operator T which is based on the Bellman equation and defines a contraction mapping:

$$Tw = \underset{i}{\text{Min}} \left(L(i, \theta) + \delta \int w(B(i, \theta, \beta_0 + \beta_1 i + \epsilon)) p(\beta | \theta) q(\epsilon) d\beta d\epsilon \right) \quad (\text{A.1})$$

$L(\cdot)$ denotes current expected loss as a function of the current policy action i and the current state θ which consists of the means and covariance matrix (b, Σ) of the bivariate normal distribution representing the policymaker's beliefs about the unknown parameters $\beta = (\beta_0, \beta_1)$. $w(\cdot)$ refers to a continuous bounded function that is defined on the relevant state space Θ and constitutes an approximation of the value function V . δ denotes the discount factor and the integral represents the expected value in the next period. Next period's beliefs have been substituted out using the relevant updating equations (3) which are denoted by $B(\cdot)$. $p(\beta | \theta)$ is the normal density function corresponding to the policymaker's beliefs. Finally, $q(\epsilon)$ denotes the normal density function of the $N(0,1)$ shock term.

We work with the space of continuous bounded functions mapping the state space Θ into the real line. This is a complete metric space in the sup metric

$$d(w^0, w^1) = \sup_{\Theta} |w^0(\theta) - w^1(\theta)| \quad (\text{A.2})$$

where w^0 and w^1 are continuous bounded functions. The operator T maps a continuous bounded function w into a continuous bounded function Tw . As shown by Kiefer and Nyarko (1989) Blackwell's (1965) sufficiency conditions, monotonicity and discounting, are satisfied and T is a contraction mapping such that

$$d(Tw^1, Tw^0) \leq \delta d(w^1, w^0) \quad (\text{A.3})$$

Thus, T has a unique fixed point, the value function V , which can be calculated by value iteration, meaning successive application of T . $T^n w$ converges to V uniformly as $n \rightarrow \infty$. A convenient starting value w^0 is the one-period loss function $L(\cdot)$.

The algorithm is implemented as follows: first, calculate starting values w^0 for a grid of points in the state space Θ and save them in a table; second, calculate w^1 by applying the operator T to w^0 and update said table. This step involves a minimization with respect to the control i for each grid point in the state space. This minimization in turn requires repeated evaluation of the following integral:

$$\int w^0(B(i, \theta, \beta_0 + \beta_1 i + \epsilon)) p(\beta | \theta) q(\epsilon) d\beta d\epsilon \quad (\text{A.4})$$

The updating equations $B(\cdot)$ and the two normal density functions are known. The values of w^0 at the grid points can be read from the table and the values in between grid points can be approximated by multilinear interpolation.²² The advantage of linear interpolation is that it preserves the shape of the function, positivity and monotonicity. Thus, even though the algorithm only remembers a discrete approximation of w^0 , when computing values of w^1 linear interpolation guarantees that Blackwell's sufficiency conditions are satisfied and the algorithm remains a contraction mapping. Based on the updating equations, the density functions and the table w^0 the above integral can be evaluated using Gaussian quadrature. The minimization step however, is nontrivial because there exist multiple local maxima. To ensure that the global minimum is found, I first conduct a rough grid search, save the minimum, and then conduct a golden section search to compute the minimum more precisely. This two-step procedure is slow but secure. For each grid point in Θ this minimum is used to update the table approximating the value function, $w^1(\cdot)$. The associated set of controls provides an approximation to the optimal policy, $i=h^1(\cdot)$. This procedure is repeated to obtain w^2 and so on until the difference between two successive approximations is sufficiently small. A more detailed discussion of numerical dynamic programming, as well as the optimization and quadrature techniques utilized here can be found in Judd (1998).

Precision and computation costs

Since the algorithm is a contraction mapping, it is straightforward to construct a bound on the approximation error. If $w^{n+1} = Tw^n$, then $d(w^{n+1}, w^n) \leq \delta d(w^n, w^{n-1})$ and after iterating k times $d(w^{n+1+k}, w^{n+k}) \leq \delta^{1+k} d(w^n, w^{n-1})$. This implies the following upper bound on the approximation error:

$$d(V, w^n) \leq \sum d(w^{n+1+k}, w^{n+k}) \leq \frac{\delta}{1-\delta} d(w^n, w^{n-1}) \quad (\text{A.5})$$

This bound only depends on the discount factor and the maximal distance between the approximations obtained in the last two iterations. After every iteration, the algorithm checks all the grid points for the maximum of $|w^n(\theta) - w^{n-1}(\theta)|$ and computes the bound on the approximation error. The value iterations are stopped once the discounted relative difference between two successive approximations is within 0.5%. This error bound neglects that the numerical maximization and integration procedures as well as the linear interpolation of w between grid points are additional sources of approximation error. For a detailed sensitivity study which shows that the error due to these sources is very small see Wieland (1999).

A drawback of the algorithm is that computational effort increases geometrically with the number of state variables. For example, if one uses N grid points for each dimension, the integration and optimization procedures described above have to be carried for each grid point i.e. N^5 times to complete one value iteration. The search for the optimum is especially time consuming, because of the existence of multiple local optima. An additional factor in terms of

²² Since we can use $L(\cdot)$ or another known function as starting point w^0 , the values in between grid points could be calculated exactly in the first iteration. However, in subsequent iterations, the functional form of w^n is not known and values in between grid points are obtained by interpolation.

computational effort is that the number of value iterations required to achieve convergence within a set maximal error bound increases rapidly with the discount factor. To reduce the number of value iterations required for convergence within a set maximal error bound, I introduce policy iterations. A policy iteration implies the application of the following operator:

$$T^P w = L(h(\theta), \theta) + \delta \int w(B(i, \theta, \beta_0 + \beta_1 h(\theta) + \epsilon)) p(\beta | \theta) q(\epsilon) d\beta d\epsilon \quad (\text{A.6})$$

Here $h(\theta)$ refers to the approximation of the policy function obtained from the preceding value iteration. Thus, following every value iteration, the operator T^P is repeatedly applied to the approximation of the value function w (for a given $h(\cdot)$) until it converges within a given bound. Such policy iterations can be carried out fairly quickly because they do not involve an optimization step. As a result, the number of value iterations as well as overall computing time is reduced. This is particularly useful, when considering high values of the discount factor.

The approximations underlying the simulations in Figures 4 and 5 and Table 3 are based on a five-dimensional grid with about 20 grid points for each dimension and use 15-point Gaussian quadrature. The grid points were chosen so that the distance between them is smaller in areas of high curvature. Convergence as defined by a 0.5% maximal error was typically achieved after about one week on a SPARC 20 SUN work station.

References

- Amman, H., and D. Kendrick, 1995, Nonconvexities in stochastic control models, *International Economic Review* 36, 455-475.
- Anderson, T., and J.B. Taylor, 1976, Some experimental results on the statistical properties of least squares estimates in control problems, *Econometrica* 44, 1289-1302.
- Balvers, R. and T. Cosimano, 1994, Inflation variability and gradualist monetary policy, *Review of Economic Studies*, 721-738.
- Basar, T. and M. Salmon, 1990, Credibility and the value of information transmission in a model of monetary policy and inflation, *Journal of Economic Dynamics and Control* 14, 97-116.
- Bertocci, G. and M. Spagat, 1993, Learning, experimentation and monetary policy, *Journal of Monetary Economics* 32, 169-178.
- Blackwell, D., 1965, Discounted dynamic programming, *Annals of Mathematical Statistics* 36, 226-335.
- Blinder, A., 1998, *Central banking in theory and practice*, (MIT Press, Cambridge).
- Brainard, W., 1967, Uncertainty and the effectiveness of policy, *American Economic Review* 57, 411-25.
- Caplin, A. and J. Leahy, 1996, Monetary policy as a process of search, *American Economic Review* 86, 689-702.
- Chung, H., 1990, Did policy makers really believe in the Phillips curve?: an econometric test, unpublished Ph.D. dissertation, University of Minnesota.
- Clarida, R. and M. Gertler, 1997, How the Bundesbank conducts monetary policy, in, Romer, C. and D. Romer, eds., *Reducing inflation: motivation and strategy*, (NBER and Chicago Press).
- Craine, R., 1979, Optimal monetary policy with uncertainty, *Journal of Economic Dynamics and Control* 1, 59-83.
- Deutsche Bundesbank, Monthly report, various issues, 1990 to 1998.
- Dornbusch, R., 1997, Comment on "How the Bundesbank conducts monetary policy," in, Romer, C. and D. Romer, eds., *Reducing inflation: motivation and strategy*, (NBER and Chicago Press).
- Easley, D. and N.M. Kiefer, 1988, Controlling a stochastic process with unknown parameters, *Econometrica* 56, 1045-1064.

- Gerlach, S., 1994, German unification and the demand for German M3, Bank for International Settlements Working Paper 21.
- Greene, W., 1993, *Econometric analysis*, (Macmillan, New York).
- Hallman, J., Porter, R., and D. Small, 1991, Is the price level tied to the M2 monetary aggregate in the long run?, *American Economic Review* 81, 841-858.
- Hansen, G. and J. Kim, 1995, The stability of German money demand: test of the cointegration relation, *Weltwirtschaftliches Archiv* 131, 286-301.
- Harvey, A., 1993, *Time series models*, 2nd edition, (MIT Press, Cambridge).
- Horvath, B., 1991, Are policy variables exogenous? The econometric implications of learning while maximizing, *Lecture Notes in Economics and Mathematical Systems* 364, (Springer).
- Issing, O. and K. Tödter, 1995, Geldmenge und Preise im vereinigten Deutschland, in D. Duwendag, ed., *Neuere Entwicklungen in der Geldtheorie und Währungspolitik*, Schriften des Vereins für Socialpolitik, 97-123.
- Judd, K., 1998, *Numerical methods in economics*, (MIT Press, Cambridge).
- Kasa, K., 1996, Will the fed ever learn?, unpublished manuscript, Federal Reserve Bank of San Francisco, Research Department.
- Kendrick, D., 1981, *Stochastic control for economic models*, Economic Handbook Series, McGraw Hill, New York.
- Kendrick, D., 1982, Caution and probing in a macroeconomic model, *Journal of Economic Dynamics and Control* 4, 149-170.
- Kiefer, N., and Y. Nyarko, 1989, Optimal control of an unknown linear process with learning, *International Economic Review* 30, 571-586.
- Kole, L. and E. Meade, 1995, German monetary targeting: a retrospective view, *Federal Reserve Bulletin* 81, 917-931.
- Marcet, A., and T. Sargent, 1989, Convergence of least squares learning mechanisms in self-referential linear stochastic models, *Journal of Economic Theory* 48, 337-68.
- Mizrach, B., 1991, Non-convexities in a stochastic control problem with learning, *Journal of Economic Dynamics and Control* 15, 515-538.
- Neumann, M., 1996, Monetary targeting in Germany, IMES Discussion Paper Series, Institute

- for Monetary and Economic Studies, Bank of Japan, Discussion Paper 96-E-15.
- OECD, 1993, 1992-93 Economic survey, Germany, (OECD, Paris).
- OECD, 1994, 1003-94, Economic survey, Germany, (OECD, Paris).
- Prescott, E., 1972, The multi-period control problem under uncertainty, *Econometrica* 40, 1043-1048.
- Scharnagl, M., 1998, The stability of German money demand: not just a myth, *Empirical Economics* 23.
- Sims, C. A., 1988, Projecting policy effects with statistical models, *Revista de Analisis Economico* 3, 3-20.
- Staiger, D., J. Stock and M. Watson, 1997, How precise are estimates of the natural rate of unemployment?, in, Romer, C. and D. Romer, eds., *Reducing inflation: motivation and strategy*, (NBER and Chicago Press).
- Taylor, J.B., 1974, Asymptotic properties of multiperiod control rules in the linear regression model, *International Economic Review* 15, 472-484.
- Tinsley, P. and P. von zur Muehlen, 1982, A Maximum Probability Approach to Short-run Policy, *Journal of Econometrics* 15, 31-48.
- Tödter, K. and H. Reimers, 1994, P-star as a link between money and prices in Germany, *Welwirtschaftliches Archiv* 130, 273-289.
- Tödter, K. and G. Ziebarth, 1997, Price stability versus low inflation in Germany: an analysis of costs and benefits, NBER Working Paper 6170.
- Tse, E. and Y. Bar-Shalom, 1973, An actively adaptive control for linear systems with random parameters, *IEEE Transactions on Automatic Control* 17, 38-52.
- von Hagen, J., 1993, Monetary union, money demand, and money supply: a review of the german monetary union, *European Economic Review*, 1993, 803-836.
- Wieland, V., 1999, Learning by doing and the value of optimal experimentation, forthcoming, *Journal of Economic Dynamics and Control*.
- Wolters, J., Teräsvirta, T. and H. Lütkepohl, 1998, Modeling the demand for M3 in the unified Germany, *Review of Economics and Statistics*, 399-409.
- Zellner, A., 1971, *Introduction to Bayesian inference in econometrics*, New York, Wiley, 1971.

Table 1. Average Money Growth and Inflation Rates Before and After Unification

	M1	M2	M3	GDP Deflator	CPI
1984:Q1 - 1990:Q1	6.73	6.94	5.86	2.23	1.53
1991:Q3 - 1994:Q2	8.80	9.66	8.67	4.02	4.35
1994:Q3 - 1998:Q3	6.94	0.57	4.20	1.21	1.53

Note: The growth rates are given in annualized percentage terms and correspond to the dashed lines in figure 1.

Table 2. Calibration of Model Parameters and Beliefs

Data:

inflation, money and output averages (84:1–90:1): $\bar{\pi}_{Defl.} = 2.2$ $\Delta \bar{m}_{M3} = 5.9$ $\Delta \bar{y}_{GDP} = 2.7$
 call-money rates, (low: 84:1–90:1), (high: 92–93): $\bar{i}_{low} = 5.0$ $\bar{i}_{high} = 8.5$

Calibration of the underlying (true) parameter values:

observed averages when parameters are known: $\bar{\pi} = 2$ $\Delta \bar{m} = 6$ $\Delta \bar{y} = 2.5$ $\bar{i} = 10$

choose values for the free parameters: $\beta_1 = -0.7$ $\sigma^2 = 1$ $\omega = 0.14$

the remaining parameter values are chosen

s.t. the observed averages match the data: $\beta_0 = \bar{\pi} - \beta_1 \bar{i} = 9$

$$\Delta y^* = 2.5$$

$$\alpha_1 = (\Delta \bar{m} - \bar{\pi}) / \Delta y^* = 1.6$$

$$\gamma_1 = -\beta_1 = 0.7$$

$$\gamma_0 = \Delta \bar{m} + \gamma_1 \bar{i} = 13$$

$$\pi^* = \beta_0 + (\omega / \beta_1 + \beta_1) \bar{i} = 0$$

Calibration of the policymaker's beliefs:

averages initially expected by the policymaker: $\bar{\pi} = 2$ $\bar{i} = 5$

free parameters regarding initial beliefs: $b_1 = -0.8$ $v_1 = 0.25$ $\rho = -0.85$

the remaining parameters are chosen to

obtain the averages expected initially: $b_0 = \bar{\pi} - \beta_1 \bar{i} = 6$

$$v_{01} = -v_1 \bar{i} = -1.54$$

$$v_0 = v_1 (v_{01} / \rho)^2 = 13.1$$

Note: Sensitivity studies regarding $(\beta_1, \sigma^2, \omega, b_1, v_1, \rho)$ are reported in section 6 and Figure 5.

Table 3. The Set of Possible Limit Beliefs and Actions

belief invariance	$\bar{v}_0 + \bar{v}_{01} \bar{i} = 0 \Rightarrow b_0, v_0 \text{ invariant}$ $\bar{v}_{01} + \bar{v}_1 \bar{i} = 0 \Rightarrow b_1, v_1, v_{01} \text{ invariant}$
prediction	$\beta_0 + \beta_1 \bar{i} = \bar{b}_0 + \bar{b}_1 \bar{i}$
optimality	$\bar{i} = -\frac{(\bar{v}_{01} + \bar{b}_0 \bar{b}_1)}{(\bar{v}_1 + \bar{b}_1^2 + \omega)}$
semi-positive-definiteness of Σ	$v_0 v_1 - v_{01}^2 \geq 0$
non-negativity of variances	$v_0, v_1 \geq 0$

Table 4. Average Policy Bias under Passive and Optimal Learning

	Only Selected Simulations		All Simulations	
	which exhibit a policy bias > 1 % for at least 30 periods ¹⁾		(1000)	
	Passive Learning	Optimal Learning ²⁾	Passive Learning	Optimal Learning ²⁾
Percentage of Simulations	29.5 %	6.3%	--	--
Average Inflation Bias	1.51	1.0	0.75	0.40
Average Money Growth Bias	1.51	1.0	0.75	0.40
Average Interest Rate Bias	-2.33	-1.76	-1.23	0.71

Notes:

1) That is, if $|\pi_{30} - \epsilon_{30} - 2| > 1$, which is equivalent to $|\beta_0 + \beta_1 i_{30} - 2| > 1$.

2) The discount factor δ for the optimal policy is set to 0.95.

Figure 1. Money and Inflation Before and After German Unification




Figure 2: Perceived and Actual Inflation Impetus of Policy




Figure 3. Passive Learning Policy - Dynamic Simulation






Figure 4. Optimal Learning Policy - Dynamic Simulation



Note: The discount factor β for the optimal policy is set to 0.95.

Figure 5
Sensitivity Analysis: Likelihood of a persistent inflation bias (magnitude > 1%, length > 30 periods)



Notes: 1. o and x correspond to the outcomes under passive and optimal learning in Table 4.

2. The discount factor δ for the optimal policy is set to 0.95

3. The outcomes for the optimal policy under alternative values of the variance of shocks is based on a policy function that assumes that the variance of shocks is equal to 1.