WP/88/114

INTERNATIONAL MONETARY FUND

Research Department

Model Uncertainty, Learning, and the Gains from Coordination *

Prepared by Atish R. Ghosh and Paul R. Masson

December 27, 1988

Abstract

The paper considers gains from international economic policy coordination when there is uncertainty concerning the functioning of the world economy, but also learning about the "true" model on the part of policymakers. The paper reports estimates of plausible alternative versions of a standard, two-country model. Activist policy (either coordinated or uncoordinated) may produce large welfare losses in the absence of learning, if policymakers believe in the wrong model; hence exogenous money targets and freely flexible exchange rates may be best. However, model learning (from observations on macroeconomic variables) causes coordinated policies to dominate activist uncoordinated policies or exogenous money targets.

JEL Classification Numbers:
0260, 4200, 4310, 4320

* Prepared for the sixth IFAC/SEDC/IFORS/IFIP symposium on "Dynamic Modeling and Control of National Economics," Edinburgh, Scotland, June 27-29, 1989; Mr. Ghosh, who is currently at Harvard University, was a consultant at the Fund when this paper was written. We would like to thank Claudia Lee for her encouragement and assistance. We have benefited from comments of participants at a Research Department seminar, and from conversations with Jeffrey Frankel and Dale Henderson. None of them should be held responsible for the views expressed in the paper, which are those of the authors alone, and in particular, do not represent those of the International Monetary Fund.

Contents

I. Introduction

In recent years there has been a marked increase in interest in international economic policy coordination, as evidenced by the proliferation of meetings of officials and of academic publications. Though the presumption is that policy coordination among the major industrial countries is a good thing, there exist valid doubts concerning the possibility of designing appropriate intervention and coordination rules when the effects of policies are uncertain. Speaking at the American Economic Association meetings in December, 1987, Martin Feldstein noted (Feldstein [1988, p. 10]):

> Uncertainties about the actual state of the international economy and uncertainties about the effects of one country's policies on the economies of other countries make it impossible to be confident that coordinated policy shifts would actually be beneficial.

Though coordinated policies may, ex post, turn out to be have been ill advised, the relevant question is whether they are likely to result in higher welfare on average than uncoordinated policies, despite the presence of such uncertainties.

The issue has been discussed in several recent papers, though a consensus on the implications of model uncertainty for the desirability of coordination has yet to be achieved. Frankel, in a series of articles (Frankel [1987],[1988], Frankel and Rockett [1988]) argues that model uncertainty makes coordination too risky and that, on average, countries are as well off pursuing non-cooperative policies as they are under coordination. Holtham and Hughes Hallett [1987], in contrast, show that criteria can be applied that diminish the likelihood that coordination will be bad. Ghosh and Ghosh [1986] and Ghosh and Masson [1988] have argued that the existence of model uncertainty does not necessarily preclude a beneficial role for the coordination of macroeconomic policies; indeed, it may in fact provide an additional incentive to coordinate policies internationally.

The differences in conclusions stem from two essential differences in approach. First is the question of whether to evaluate gains ex ante or ex post. Ghosh and Masson focus on the ex ante expected gains from coordination while Frankel considers the ex post actual gains after arbitrarily specifying which is the true model. The second issue concerns the nature of expectations formation. In Ghosh and Masson, policymakers are assumed to have rational expectations across the set of possible models and to take due account of the presence of model uncertainty in formulating policies whereas in Frankel's approach policymakers are assumed to have different subjective priors with respect to the different models, but each believes (wrongly) that he/she knows the correct model. 1/

In this paper we re-examine the issue of model uncertainty and policy coordination, highlighting the effects of the differences in approach described above. We use variants of a simple consensus open-economy model presented in Oudiz and Sachs [1985], which is based on the Dornbusch [1976] extension to the Mundell-Fleming model. 2/ If the subjective priors of policymakers (and of private agents) equal the objective probabilities, the average of welfare values achieved ex post--presumably the relevant criterion--will equal the ex ante expected welfare value. This was the case in our earlier paper; here we relax that assumption, and allow subjective priors and objective probabilities to differ, as does Frankel. We estimate alternative versions of the Oudiz-Sachs model, and perform stochastic simulations on the assumption that one or the other version is the true one, and agents assign non-zero probabilities to each of the models; we use the ex post welfare criterion in evaluating the gains from coordination. Our work differs from most of Frankel's work, however, in that we continue to maintain that policymakers take account of the model uncertainty, and maximize expected utility over the range of models.

This paper also makes a significant departure from all of the earlier coordination literature in that it abandons the purely static view of model uncertainty which has hitherto been adopted. Instead, it is assumed that agents update their priors over the set of possible models in a Bayesian fashion. 3/ There is an extensive literature on "learning rational expectations" (Brock [1972], Cyert and de Groot [1974], De Canio [1979], Bray [1982], Blume, Bray and Easley [1982], Bray and Savin [1986], and

---

1/ In a section entitled "Extensions with Uncertainty," Frankel and Rockett [1988] consider cases where models are not believed to be correct with certainty; if each policymaker averages over the possible models when coordinating policies, gains from coordination relative to non-cooperation are higher than when coordination is based on individual models.

2/ The multi-country models surveyed by Frankel are generally elaborations of this simple model structure.

3/ Frankel and Rockett [1988, p. 318] refer to this possibility in justifying their assumption of disagreement among policymakers, but do not treat it formally: "If one wishes to think of actors as perpetually processing new information in a Bayesian manner, so that their models over time would converge on any given reality in the limit, then one must admit that the speed of convergence is sufficiently slow, or else that reality is changing sufficiently rapidly, that policymakers have not been able to reach agreement on the true model."

Marcet and Sargent [1988]). The conclusion of this literature is that least squares learning seems to converge to rational expectations in a wide variety of circumstances, but that it may also converge to an incorrect model. In particular, Blume and Easley [1982] is closest to the setup of our paper, as the authors consider Bayesian updating of prior probabilities applied to a finite set of models. They construct examples of cases where both the true model and the wrong model are locally stable, i.e. where it is possible that agents do not converge to rational expectations. It is thus of interest to analyze in an empirical macroeconomic model whether convergence to rational expectations occurs, both in a cooperative and in a Nash equilibrium. Furthermore, since the process of learning affects both policy setting and the private sector's expectation formation, it will change the economy's equilibrium and hence the estimated gains from coordination.

We evaluate three regimes: a cooperative equilibrium in which policies are jointly chosen to maximize a weighted average of the two countries' utilities; a non-cooperative, or Nash, equilibrium in which each country maximizes its own utility, taking as given the actions of the foreign government; and finally a non-interventionist "pure float" exchange rate policy, where each country keeps the money supply at its exogenous target level. The models we use are based on an empirically-estimated, two country (U.S. and an aggregate rest-of-OECD) model with a number of structural variants. Although these structural differences are seemingly minor and innocuous, the differences in the reduced-form multipliers of the models are substantial, with the degree of model uncertainty similar to that in Frankel's study. 1/ Policy conflicts between the countries arise from structural shocks to the world economy--including money demand shocks, aggregate demand shocks and inflation shocks. In our simulation analysis we use drawings from a joint distribution describing these shocks that is based on the empirically-estimated covariance matrix.

Our main conclusions may be briefly summarized. First, relatively minor structural differences in the models translate into substantial differences in the reduced-form multipliers. Second, if optimal policies--whether cooperative or non-cooperative--assign sufficiently little weight to the true model, and there is no learning, then the economy can become dynamically unstable with potentially huge gains or losses from coordination. In contrast, uncoordinated policies involving freely-floating exchange rates and exogenous money targets never result in dynamic instabilities in our estimated version of the Oudiz-Sachs model. Therefore, in the absence of model learning, floating exchange rates may be optimal because they are more robust to model misspecification. We have argued

---

1/ The degree of model uncertainty for a particular reduced-form multiplier is defined below as the ratio of its average across models, squared, to the sum of its average squared and its variance.

previously (in Ghosh and Masson [1988]) that the lesson from model uncertainty is not that coordination is bad, but that policy--whether it is coordinated or not--should not be too activist. Our results here support that conclusion. Third, once Bayesian learning is allowed, optimal policies never result in dynamic instability and, even when very little initial probability is assigned to the true model, the ex post gains from coordination (when discounted back to the present) are always positive in our simulation exercises.

The reason for the latter result is simple. If policies are set to maximize ex ante expected utility, coordination only results in welfare deterioration ex post if the models are very different; in that case, however, it becomes very easy to distinguish between the models and learn which is the true model. The subjective priors therefore quickly converge to the true model and coordination is welfare improving. We would not want to exaggerate the relevance of this result to real-world policy choice; instead, it highlights the inadequacy of the assumption that one model is the "true" one. 1/ Nevertheless, experience of the past quarter century does provide evidence that policymakers abandon views of the world that can be seen to be wrong--such as the view that there exists a long-run tradeoff between unemployment and inflation that can be exploited by aggregate demand policy--and thereby avoid the more disastrous consequences of their actions.

The paper is organized as follows. Section II presents some of the theoretical arguments concerning the effects of learning on the gains from coordination. Section III describes the empirical model used for the simulations. Section IV describes the procedure used to evaluate the gains from coordination under model uncertainty and Bayesian learning. Section V presents the results of the simulation analysis. Section VI discusses some caveats to the results, while section VII offers some brief concluding remarks.

II. Gains From Policy Coordination: Ex ante and Ex post

In discussing the benefits of policy coordination it is useful to distinguish two concepts of the welfare gains. The first criterion is the ex ante expected gain from coordination--that is, the difference between the expected utility under the cooperative regime and the expected utility under the non-cooperative regime. The second is the ex post actual gain from coordination--the difference in actual welfare level after realizations of all random variables are known. In most of the literature, where the gains

---

1/ In the context of the quote in footnote 1 on page 1, we would argue that the speed of learning is unlikely to be the source of the problem, but rather the fact that reality is much more complex than the models, and is changing too rapidly.

from coordination arise from policy interdependence and the inefficiency of the Nash equilibrium (see Cooper [1968], Hamada [1974], Oudiz and Sachs [1984], *Canzoneri and Gray [1983]), the distinction is immaterial.* As long as the models are deterministic, or at least certainty equivalence obtains, both criteria give the same ranking of strategic regimes. 1/ The main conclusion of this literature is that the gains from coordination are usually positive though not necessarily so (Rogoff [1985]).

There is some recent literature on coordination under model uncertainty, in which it is assumed that policymakers are uncertain about the effects of their instruments on the target variables. In this context certainty equivalence does not apply, and the use of the ex ante welfare criterion can give very different results from the ex post measure. We will term "ex ante expected utility approach" the use of subjective priors to evaluate welfare gains; examples of this approach are Ghosh [1986], Ghosh and Ghosh [1987], Ghosh and Masson [1988], all of which assume that optimal policies are set on the basis of expected utility maximization (as in Brainard [1967]) using their subjective probability priors. The use of subjective priors to compute actual average level of welfare is legitimate as long as the subjective priors also correspond to the objective probabilities that the models are true. In this approach, the structure of the economy is viewed as changing over time and at any instant is a random realization from the set of possible models. Aside from the possibility that other distortions in the economy negate the gains from coordination (as in Rogoff's example mentioned above) the ex ante level of welfare can never be lower with coordination than without coordination.

In contrast, what we will call the "ex post welfare criterion" evaluates welfare using the "true" model: the economy is assumed to be correctly characterized by a single model, which policymakers however have not yet discovered. In this context, coordination may actually reduce welfare. Frankel and Rockett [1988], for instance, assume in the main part of their paper that policymakers ignore the presence of model uncertainty when choosing policies but assume (possibly wrongly) that they know the true model. Since the subjective priors of policymakers (which are unity for the model they believe and zero for all others) do not in general coincide with the objective probabilities (which are unity for the true model and zero otherwise) the average ex post level of welfare differs from the ex ante expected utility and coordination may decrease welfare.

1/ If the model is linear and the objective function is quadratic, then unobserved additive errors in the model's equations will affect welfare <u>ex post</u> through the variances of those errors, in a way that is independent of the actions of the policy makers. In contrast, if the model's parameters are stochastic, then <u>ex post</u> welfare will be affected in a way that depends on the policies chosen.

In order to illustrate the issues involved, it is useful to consider the following, static, reduced-form theoretical model (the discussion is based on Ghosh and Masson [1988]). Let the equations describing inflation (p) and output (y) of the home economy be:

$$p_t = \lambda M_t + \Theta(M_t - M_t^*) + \varphi + \epsilon_t^p$$

and

$$y_t = \gamma M_t + \nu M_t^* + \epsilon_t^y$$

where φ is a positive shock to inflation which is observed at time t, M and M^* are the domestic and foreign money supplies respectively, ϵ_t^p and ϵ_t^y are i.i.d. zero mean shocks which are not directly observable, and λ, Θ, γ, ν are all positive constants. Policymakers are assumed to have a quadratic objective function V defined over output and inflation, with relative weight ϕ, and discounted using factor β:

$$V = \underset{\text{wrt } M}{\text{Max}} \ -(1/2) \sum_{t=0}^{\infty} E\, \beta^t (p_t^2 + \phi y_t^2)$$

A similar objective function and reduced-form model applies to the fully symmetric foreign country.

In an uncoordinated equilibrium each country maximizes its utility taking as given the policy choice of the foreign country. In a cooperative equilibrium, a global social planner maximizes a weighted average of the two countries' objective functions:

$$V^c = \underset{\text{wrt } M,M^*}{\text{Max}} \ -(1/2)\{ \omega \sum_{t=0}^{\infty} E\, \beta^t (p_t^2 + \phi y_t^2) + (1-\omega) \sum_{t=0}^{\infty} E\, \beta^t (p_t^{*2} + \phi y_t^{*2})\}$$

In what follows we will drop the time subscript and only consider optimal policies in period t. Moreover, since the countries' parameters are the same and the shocks are symmetric, the two countries' policies are identical. We will report only the home country's values. The Nash, or uncoordinated equilibrium of this model is given by:

$$M^N = -(\lambda + \Theta)\varphi/[\lambda(\lambda + \Theta) + \phi\gamma(\gamma + \nu)]$$

Assuming a symmetric cooperative equilibrium (ω=0.5), the corresponding cooperative policy is:

$$M^C = -\lambda\varphi/[\lambda^2 + \phi(\gamma + \nu)^2]$$

The assumptions on the parameter values are sufficient (though not necessary) to ensure that the Nash equilibrium is too contractionary in response to an inflation shock:

$$M^N = M^{N*} < M^C = M^{C*}$$

The intuition is straightforward: in the uncoordinated regime, each country assumes that tighter monetary policy relative to the foreign country will result in an exchange rate appreciation and therefore a fall in the inflation rate. Since both countries attempt this simultaneously, their efforts are vitiated and the world simply inherits a contractionary bias (see Sachs [1983]). The cooperative planner recognizes the futility of competitive appreciations and therefore avoids this contractionary bias.

We now introduce model uncertainty into the framework by assuming that the parameters $(\theta,\gamma,\lambda,\nu)$ are uncertain, with means $(\mu_\theta,\mu_\gamma,\mu_\lambda,\mu_\nu)$ and variances $(\sigma^2_\theta,\sigma^2_\gamma,\sigma^2_\lambda,\sigma^2_\nu)$; covariances between pairs of parameters are assumed to be zero. Optimal policies under model uncertainty in each regime are given by:

$$M^N = -(\mu_\lambda+\mu_\theta)\varphi/[(\mu^2_\lambda + \sigma^2_\lambda + \mu_\theta\mu_\lambda) + \phi(\mu^2_\gamma + \sigma^2_\gamma + \mu_\gamma\mu_\nu)]$$

$$M^C = -\mu_\lambda\varphi/[(\mu^2_\lambda + \sigma^2_\lambda) + \phi(\mu^2_\gamma + \sigma^2_\gamma + \mu^2_\nu + \sigma^2_\nu)]$$

Although a "small" degree of model uncertainty about domestic multipliers raises the incentive to coordinate, as the degree of domestic multiplier uncertainty increases (σ^2_λ or $\sigma^2_\gamma \Rightarrow \infty$), the deflationary bias of the Nash regime is eliminated and the Nash and cooperative strategies converge (Ghosh and Masson [1988]). By contrast, greater uncertainty about the transmission effects ($\sigma^2_\nu \Rightarrow \infty$) 1/ causes a larger divergence between the Nash and cooperative strategies. These results indicate how model uncertainty affects the ex ante expected gains from coordination; since these policies are chosen to maximize the ex ante expected utility, by construction expected welfare cannot be lower under cooperation than under non-cooperation.

However, for any given realization from the set of possible models, coordination may result in lower welfare than non-cooperative behavior;

1/ σ^2_θ does not appear in these expressions since it embodies both transmission and domestic multiplier effects.

therefore, ex post welfare may be lower if governments have cooperated. For instance, in Frankel [1987], since there is one true model of the economy, but policymakers do not know which of several alternatives is the correct one, then, ex post, the cooperative regime may yield lower welfare than the non-cooperative regime. This is all the more likely if policymakers completely ignore model uncertainty in choosing policies, and hence are likely to be more activist than is optimal.

To simplify the discussion assume that the uncertainty about the parameters arises because there are a two possible models of the economy:

$$p^i = \alpha^i M - \theta^i M^* + \psi \quad i=1,2$$

where $\alpha=\lambda+\theta$, and that policymakers have a single objective, inflation ($\phi=0$). Agents assign prior probabilities of Π and $(1-\Pi)$ to each model. The non-cooperative strategy is:

$$M^N = M^{N*} = -\psi[\Pi\alpha^1 + (1-\Pi)\alpha^2]/[\Pi\alpha^1(\alpha^1-\theta^1) + (1-\Pi)\alpha^2(\alpha^2-\theta^2)]$$

and the cooperative strategy is:

$$M^C = M^{C*} = -\psi[\Pi(\alpha^1-\theta^1) + (1-\Pi)(\alpha^2-\theta^2)]/[\Pi(\alpha^1-\theta^1)^2 + (1-\Pi)(\alpha^2-\theta^2)^2]$$

A third possible policy is to refrain from intervention entirely, or rather to maintain the money supply at the level prevailing before the shock:

$$M^F = M^{F*} = 0$$

Such a policy corresponds to the combination of freely flexible exchange rates and a fixed target for the money supply.

Suppose model 1 is the true model; then the value function associated with each regime is given by:

$$V^N = -(1/2)\psi^2\{(1-\Pi)^2[\alpha^2(\alpha^1-\theta^1)-\alpha^1(\alpha^2-\theta^2)]^2/[\Pi\alpha^1(\alpha^1-\theta^1) + (1-\Pi)\alpha^2(\alpha^2-\theta^2)]\}$$

$$V^C = -(1/2)\psi^2\{(1-\Pi)^2[(\alpha^2-\theta^2)((\alpha^1-\theta^1)-(\alpha^2-\theta^2)]^2/[\Pi(\alpha^1-\theta^1)^2 + (1-\Pi)(\alpha^2-\theta^2)^2]\}$$

$$V^F = -(1/2)\psi^2$$

The level of welfare under cooperation is necessarily an increasing function of Π, the probability associated with the true model; whereas the second-best nature of the Nash equilibrium implies that welfare may be an increasing or a decreasing function of the probability assigned to the true model. Since the pure float regime does not require any policy intervention, the actual level of welfare is independent of policymakers' subjective priors. From the expressions for the value functions it is simple to find specific parameter values such that coordination is welfare inferior to non-cooperative policies as long as $\Pi \neq 1$. However, it is also possible to find parameter values such that the floating exchange regime is superior to both cooperative and non-cooperative policies, indicating that the use of any activist policy in the face of uncertainty may make matters worse. As argued in Ghosh and Masson [1988], however, this shows that caution is needed in setting policy in these conditions, and is not an indictment of coordination per se.

It is to be expected from the discussion above that the effects of learning on the ex ante gains from coordination will depend upon whether initially there is a large degree of domestic multiplier or transmission multiplier uncertainty. If uncertainty is primarily related to transmission effects then as agents learn about the model the gains from coordination should decrease, and conversely for uncertainty about the domestic effects of policies. The effects on ex post gains from coordination are more difficult to characterize: since welfare under cooperation must rise as agents learn about the model, but may either increase or decrease in the non-cooperative equilibrium, the effects on the gains from coordination are ambiguous.

It is an empirical question, to which we now turn, as to whether coordination is likely to reduce welfare either relative to a Nash strategy or to a pure floating exchange rate. In the absence of model learning, coordination may certainly reduce welfare ex post, and, in a dynamic model, may lead to dynamic instability of the economy if the wrong model is used to set policies. On the other hand, if learning is sufficiently fast, and ex ante optimal policies are sufficiently robust to model errors, coordination could still be welfare superior to non-cooperative policies, both ex ante and ex post.

III. The Empirical Model

The empirical model we adopt is a general, two-country Mundell-Fleming model with forward-looking exchange rate expectations (see Dornbusch [1976], Oudiz and Sachs [1985]). The model was estimated using the data that are the basis for the IMF's MULTIMOD model (Masson and others [1988]). Results are presented in Table 1, using the same notation as in Oudiz and Sachs [1985]. Our two countries are the United States and the rest of the world; data for the latter region resulted from aggregation of the remaining

Table 1. Parameter Estimates for Oudiz-Sachs Model, 1966-86.
(standard errors in parentheses)

Aggregate demand

$$q_t = -\underset{(.051)}{.110}\,(p_t - e_t - p_t^*) + \underset{(.206)}{.037}\,q_t^* - \underset{(.301)}{.178}\,(i_t - p_{t+1} + p_t)$$

$$+\underset{(.009)}{.021}\,t + \underset{(.167)}{.288}\,g_t + \underset{(2.302)}{5.391} \qquad \bar{R}^2 = .984 \quad \sigma = .021 \quad DW=1.78$$

$$q_t^* = \underset{(.096)}{.171}\,(p_t - e_t - p_t^*) + \underset{(.496)}{.514}\,q_t - \underset{(.207)}{.384}\,(i_t^* - p_{t+1}^* + p_t^*)$$

$$-\underset{(.025)}{.030}\,t + \underset{(.468)}{1.599}\,g_t^* - \underset{(6.068)}{5.726} \qquad \bar{R}^2 = .990 \quad \sigma = .020 \quad DW=1.15$$

Money demand

$$m_t - p_t = \underset{(.066)}{.225}\,q_t - \underset{(.413)}{1.419}\,i_t + \underset{(.503)}{3.453} \qquad \bar{R}^2 = .382 \quad \sigma = .038 \quad DW=1.02$$

$$m_t^* - p_t^* = \underset{(.078)}{.700}\,q_t^* - \underset{(.609)}{1.077}\,i_t^* + \underset{(.616)}{.100} \qquad \bar{R}^2 = .853 \quad \sigma = .051 \quad DW=0.67$$

Consumer price index

$$p_t^c = .899\,p_t + (1-.899)\,(p_t^* + e_t)$$

$$p_t^{c*} = .758\,p_t^* + (1-.758)\,(p_t - e_t)$$

Output price change

$$p_t - p_{t-1} = (p_{t-1}^c - p_{t-2}^c) + \underset{(.119)}{.095}\,[q_{t-1} - .027\,(t-1) - 7.373]$$

$$+\underset{(.108)}{.309}\,(q_{t-1} - q_{t-2}) - \underset{(.004)}{.009} \qquad \bar{R}^2 = .351 \quad \sigma = .011 \quad DW=2.33$$

$$p_t^* - p_{t-1}^* = (p_{t-1}^{c*} - p_{t-2}^{c*}) + \underset{(.088)}{.040}\,[q_{t-1}^* - .033\,(t-1) - 7.803]$$

$$+\underset{(.204)}{.880}\,(q_{t-1}^* - q_{t-2}^*) - \underset{(.008)}{.031} \qquad \bar{R}^2 = .458 \quad \sigma = .017 \quad DW=1.74$$

Table 1 (continued)

Exchange rate

$$e_{t+1} = e_t + i_t - i^*_t + \epsilon \qquad \sigma_\epsilon = .026$$

Variable definitions
(U.S. variables are unstarred, non-U.S. variables are starred)

---

e = nominal exchange rate (dollars per foreign currency)

g = real government spending on goods and services

i = *nominal short-term interest rate*

m = monetary base

p = GDP deflator

p^c = deflator of domestic absorption

q = real GDP

t = time trend

---

industrial countries in MULTIMOD. 1/ The estimation period, using annual data, was 1966 to 1986. Aggregate demand equations were estimated using instrumental variables: the instruments included money stocks, government spending, time, and lagged prices and output.

There are several things to note in Table 1, which can be compared to Table 7.2 in Oudiz and Sachs [1985, p. 297]. First, the aggregate demand equation contains additional exogenous terms: the log of government spending (labeled g and g^*) and a time trend. The former is intended to capture other influences on demand, and, in particular, the effect of fiscal policy (which is not, however, considered in the policy coordination experiments discussed below). The time trend is present because aggregate demand variations should be relative to potential output, which is here exogenous and assumed to grow smoothly over time.

Second, since in the Oudiz-Sachs model domestic output prices depend solely on wages, the latter are substituted out of the model; the same is done for the rate of inflation. The resulting price change equation also has a time trend, so that the level effect of output is relative to potential output, which was generated from a separate equation where the log of output was regressed on a constant and a time trend. Finally, the equation for the exchange rate, based on uncovered interest parity, has an error term. Even though the equation is not estimated, an allowance is made in the stochastic simulations below for the fact that the equation does not hold exactly. The interpretation given to errors in this equation is that of shifts in portfolio preferences. Given a proxy for the expected exchange rate--in particular, assuming that in the historical data expectations of the exchange rate were based on a random walk model--the standard deviation of that error is seen to be sizable.

The coefficient estimates are all of the right signs, and generally fairly well determined--in particular the effects of the real exchange rate on U.S. aggregate demand and of U.S. GDP on rest-of-world demand, the money demand parameters, and the change in GDP effects on the two regions' output price changes. Despite its simplicity, the model seems to fit the data fairly well. Nevertheless, residual serial correlation is evident in the money demand equations, and the specification embodies arbitrary constraints that are open to discussion. We proceed to relax some of these restrictions below, and estimate the resulting alternative models.

1/ In general, variables were aggregated by converting to a common currency and summing. GDP weights were used to aggregate interest rates. The exchange rate was taken to be the reciprocal of the MERM-weighted effective rate of the U.S. dollar (in index form, 1980=1). The rest of world price level and money supply are expressed in this "currency."

The computational burden of calculating optimal policies with model learning severely restricted the number of alternative models we could introduce; in this paper we consider three possible models. The two alternative models, models II and III, differ from the baseline model (model I) in two respects: in model II money balances are deflated by the consumer price index rather than the GDP deflator, while in model III there is a lagged dependent variable in the money demand equation and a non-vertical Phillips curve.

A first arbitrary feature of model I relates to the proper deflator for real money balances. As Branson and Buiter [1983] point out in their discussion of the Mundell-Fleming model, the effects of monetary and fiscal policies can be importantly different depending on whether the domestic output price or a broader index that includes foreign goods is used. The Oudiz-Sachs specification, excluding as it does the effects of terms of trade, conforms in this respect to the original Mundell-Fleming model. Deflating money balances not by p but by p^c yields the following equations:

$$m_t - p_t^c = \underset{(.070)}{.214}\, q_t - \underset{(.436)}{1.651}\, i_t + \underset{(.530)}{3.574} \qquad \bar{R}^2 = .397 \quad \sigma = .041 \quad DW=1.02$$

$$m_t^* - p_t^{c*} = \underset{(.087)}{.650}\, q_t^* - \underset{(.683)}{1.190}\, i_t^* + \underset{(.691)}{.541} \qquad \bar{R}^2 = .790 \quad \sigma = .057 \quad DW=0.66$$

The fit is approximately the same as the equations in Table 1 (the dependent variable is of course not the same so that they cannot be compared directly), and the coefficients are almost identical and equally well determined. Given our short sample period and high collinearity between the two price series, tests of one model against the other (using either nested or non-nested tests) are unlikely to distinguish between them. Instead, we take the equations with money balances deflated by p^c as an alternative model: this specification is included in model II.

A second major area of arbitrariness in the Oudiz-Sachs model presented in Table 1 is the dynamic specification. The money demand equation, whether specified with the output or consumption deflator, shows evidence of residual serial correlation. Most studies have allowed for the possibility that money balances adjust with a lag, and have included a lagged dependent variable (see for instance Goldfeld [1973]). If we do so, then the following equations result, which are included in model III:

$$m_t - p_t^c = \underset{(.051)}{.184}\, q_t - \underset{(.318)}{1.387}\, i_t + \underset{(.160)}{.680}\, (m_{t-1} - p_{t-1}^c) + \underset{(.857)}{.302}$$

$$\bar{R}^2 = .691 \quad \sigma = .029 \quad DW=1.63$$

$$m_t^* - p_t^{c*} = .277\, q_t^* - 1.579\, i_t^* + .702\, (m_{t-1}^* - p_{t-1}^{c*}) - .421$$
$$\quad (.094) \qquad (.456) \qquad (.141) \qquad\qquad (.494)$$

$$\bar{R}^2 = .909 \quad \sigma = .038 \quad DW=1.65$$

Model III also relaxes the unitary coefficient on the lagged rate of inflation in the domestic price change equation. Relaxing the restriction in the inflation equation lowers the standard errors of estimate in both the U.S. and rest-of-world equations, and the standard errors for the inflation coefficients imply rejection of the unitary coefficient for the United States (but not for the rest of world). Furthermore, the change in output becomes insignificant, so we dropped this variable and reestimated, yielding:

$$p_t - p_{t-1} = .611\, (p_{t-1}^c - p_{t-2}^c) + .364\, [q_{t-1} - .027\, (t-1) - 7.373] + .022$$
$$\quad (.098) \qquad\qquad (.101) \qquad\qquad\qquad (.006)$$

$$\bar{R}^2 = .783 \quad \sigma = .010 \quad DW=1.95$$

$$p_t^* - p_{t-1}^* = .238\, (p_{t-1}^{c*} - p_{t-2}^{c*}) + .491\, [q_{t-1}^* - .033\, (t-1) - 7.803] + .051$$
$$\quad (.146) \qquad\qquad (.116) \qquad\qquad\qquad (.010)$$

$$\bar{R}^2 = .761 \quad \sigma = .015 \quad DW=1.44$$

Though the fit of the equations is only marginally superior, the two specifications have quite different long-run properties. The Oudiz-Sachs model exhibits no long-run tradeoff between output and inflation, as in steady state both output prices and consumer prices grow at the same, steady rate. In contrast, with a non-unitary coefficient, different rates of inflation are associated with different rates of output growth; this possibility would be disputed by many economists, however (see Friedman [1968], for instance). On the basis of the estimation results, it may be reasonable to attribute some non-zero probability to the existence of such a tradeoff.

Models II and III therefore differ from model I in several respects: model II deflates money balances by the consumption deflator, while model III allows for lags in money demand and for a non-vertical Phillips curve. Although these structural differences seem minor, the reduced-form multipliers differ considerably across the models, as do the dynamic properties (see Table 2 for the eigenvalues of the three models). There are

Table 2. Eigenvalues of Models I-III

Model I	Model II	Model III
1.122	1.415	1.428
.331±.850i	.412±.875i	-.183±.554i
.404	.022	-.104±.188i
-.773	-.285	

Table 3. Measures of Reduced-Form Model Uncertainty
(Zeta Values in percent)

Effects of U.S. Money Supply on:				Effects of ROW Money Supply on:			
US Y	ROW Y	US P	ROW P	US Y	ROW Y	US P	ROW P
Models I-III 1/							
28	46	57	44	23	48	51	45
Frankel [1988] 2/							
68	--	45	52	19	72	24	25

1/ Average of dynamic multipliers over the first 10 periods.

2/ Second-year effects.

5 eigenvalues in each model, with one unstable eigenvalue in each case, corresponding to the one non-predetermined variable, the exchange rate.

A simple summary statistic for the degree of model uncertainty is given by the ratio:

$$\zeta = \mu^2/(\mu^2+\sigma^2)$$

where μ is the mean multiplier of an instrument on a particular target and σ^2 its variance (across the three models). This gives a dimensionless statistic which equals unity if there is no model uncertainty and approaches zero as the effectiveness of the policy instrument deteriorates, in the sense of Brainard [1967]. In a model with forward-looking variables the multipliers depend upon the anticipation of future policies as well as current policies and therefore are not independent of the regime under consideration. The ζ values reported in Table 3 are for an exogenous, permanent increase in the money supply of 4 percent. For the sake of comparison, we also report the ζ values for the models in Frankel's study, also calculated for such a money supply shock. As can be seen, his study of 12 different models incorporates about the same degree of uncertainty about the reduced-form multipliers as our three alternative models. 1/

IV. Optimal Policies under Model Uncertainty and Model Learning

In order to calculate the average ex post gains from coordination with model learning we adapt the algorithm developed by Ghosh and Masson [1988]. The logic of the model is as follows. In period t, the state vector x_t and a vector of subjective priors Π_t (with elements π_t^i) are inherited. Policymakers choose a vector of controls (i.e. policy instruments) u_t in order to influence their target vector τ_t. They do not attempt active learning, i.e. performing policy experiments in order to discover which of the models is correct. At the end of period t a vector of endogenous variables ω_t is observed which allows agents to update their priors, yielding Π_{t+1}.

1/ A footnote in the Brookings volume observes that the multipliers reported may not be strictly comparable since the authors' calculations of the multipliers did not follow identical procedures.

The dynamics of the world model are assumed to be given by:

$$\begin{bmatrix} x^i_{t+1} \\ e^i_{t+1} \end{bmatrix} = \begin{bmatrix} A^i & B^i \\ D^i & F^i \end{bmatrix} \begin{bmatrix} x^i_t \\ e^i_t \end{bmatrix} + \begin{bmatrix} C^i \\ G^i \end{bmatrix} u_t + \begin{bmatrix} \Theta^i_x \\ \Theta^i_e \end{bmatrix} \epsilon_t \quad i=1,\ldots,k$$

where x is a vector of state variables, e is a vector of jumping variables, u is the vector of controls, and ϵ is an unobserved vector white noise shock, distributed $N(0,\Sigma)$, where Σ is the variance-covariance matrix; $\{A^i, B^i, C^i, D^i, F^i, G^i, \Theta^i_x, \Theta^i_e\}$ are constant matrices associated with model i.

In addition, structural equations of the models map the state variables and the forward-looking variables into a vector of targets τ:

$$\tau^i_t = L^i x_t + M^i e_t + N^i u_t + \Theta^i_\tau \epsilon_t$$

Policymakers in each country are assumed to have preferences over the target vector which are represented by:

$$v = \max -(1/2) \sum_{t=0}^{\infty} \beta^t E\{\tau'_t \Omega \tau_t\}$$

and

$$v^* = \max -(1/2) \sum_{t=0}^{\infty} \beta^{*t} E\{\tau_t' \Omega^* \tau_t\}$$

respectively, where the expectation is taken with respect to both the uncertainty about parameter values and the current realization of the shock ϵ_t. We assume that although agents do not know the true model of the world economy, they do know the variance-covariance matrix of the additive shocks. 1/ Moreover, we do not allow heterogeneity of agents; both the private sectors and the governments start with the same priors across models and update them in the same fashion.

Following Oudiz and Sachs [1985], we derive the optimal decision rules by first calculating the dynamic programming solution for a (finite) T

1/ This is a somewhat heroic assumption, but simplifies the analysis considerably. Furthermore, our main conclusions do not appear to be excessively sensitive to the assumed variance-covariance matrix.

period horizon, and then letting $T \Rightarrow \infty$. Assume that by period T the exchange rate has stabilized and the agents in each country have the vector of priors Π_T and Π^*_T over the k possible models. Since the exchange rate has stabilized, the matrix equation given above for the global model implies:

$$e_T = (I-F^i)^{-1}\{D^i x^i_T + G^i u_T + \theta^i_e \epsilon_T\}$$

Therefore:

$$\begin{aligned} r^i_T &= \{L^i + (I-F^i)^{-1}D^i\}\ x_T + \{N^i + (I-F^i)^{-1}G^i\}u_T + \{\theta^i_r + (I-F^i)^{-1}\theta^i_e\}\epsilon_T \\ &\equiv \Phi^i_x x_T + \Phi^i_u u_T + \Phi^i_\epsilon\ \epsilon_T \end{aligned}$$

The home country's objective function is therefore:

$$v = \text{Max}\ -(1/2) \sum_{t=0}^{\infty} \beta^t \sum_{i=1}^{k} \pi^i_T \{\Phi^i_x x_T + \Phi^i_u u_T + \Phi^i_\epsilon\ \epsilon_T\}'\ \ \Omega\ \{\Phi^i_x x_T + \Phi^i_u u_T + \Phi^i_\epsilon\ \epsilon_T\}$$

and the foreign country's is given analogously. As shown in Ghosh and Masson [1988], the solutions under Nash and cooperative behavior may be written:

$$u^N_T = \Gamma^N_T(\Pi_T, \Pi^*_T)\ x_T \quad \text{and} \quad u^C_T = \Gamma^C_T(\Pi_T, \Pi^*_T)\ x_T$$

where $\Gamma^N(\cdot)$ and $\Gamma^C(\cdot)$ are non-linear functions of Π and Π^* and where the assumption of a zero mean for ϵ and certainty equivalence have been exploited.

Substituting back into the objective functions gives, for either policy regime, expressions of the form:

$$v_T = x_T'\ S_T(\Pi_T, \Pi^*_T)\ x_T \quad \text{and} \quad v^*_T = x_T'\ S^*_T(\Pi_T, \Pi^*_T)\ x_T$$

and the forward-looking variables are given by:

$$\begin{aligned} e_T &= (I-F^i)^{-1}\{(D^i + G^i\Gamma_T(\Pi_T, \Pi^*_T))x^i_T + \theta^i_e \epsilon_T\} \\ &\equiv \Lambda^i_T(\Pi_T, \Pi^*_T)\ x_T \end{aligned}$$

so that the expected value of the exchange rate is:

$$E(e_T) = \sum_{i=1}^{k} \pi^i \Lambda_T^i(\Pi_T, \Pi_T^*) \quad x_T$$

Now consider period T-1, in which the state vector x_{T-1} and the priors Π_{T-1} and Π^*_{T-1} have been inherited. Assuming that the model is not expected to change between periods, the first-order conditions may be solved to yield optimal policies:

$$u_{T-1} = \Gamma_{T-1}(\Pi_{T-1}, \Pi^*_{T-1}) \ x_{T-1} \quad \text{and} \quad u^*_{T-1} = \Gamma^*_{T-1}(\Pi_{T-1}, \Pi^*_{T-1}) \ x_{T-1}$$

and value functions:

$$v_{T-1} = x_{T-1}' \ S_{T-1}(\Pi_{T-1}, \Pi^*_{T-1}) \ x_{T-1} \text{ and } v^*_{T-1} = x'_{T-1} \ S^*_{T-1}(\Pi_{T-1}, \Pi^*_{T-1}) \ x_{T-1}$$

The procedure is repeated until stationary policy and value functions are obtained:

$$u_t = \Gamma(\Pi_t, \Pi_t^*) \ x_t \qquad \text{and} \qquad u_t^* = \Gamma^*(\Pi_t, \Pi_t^*) \ x_t$$

$$v_t(x_t, \Pi_t, \Pi_t^*) = x_t' \ S(\Pi_t, \Pi_t^*) \ x_t \quad \text{and} \quad v_t^*(x_t, \Pi_t, \Pi_t^*) = x_t' \ S^*(\Pi_t, \Pi_t^*) \ x_t$$

Having obtained the optimal stationary policy rules, the model is simulated forward, and policymakers' priors are updated using Bayesian inference. The forward simulation is conditional on a particular model, say model j, being true. Suppose that in period t, the world economy has inherited the state x_t and policymakers' priors over the models are given by Π_t and Π_t^*. The optimal policy in period t for the home country is given by:

$$u_t = \Gamma(\Pi_t, \Pi_t^*) \ x_t$$

A drawing from the shocks ϵ_t is made, and the state vector in t+1 is therefore given by

$$x_{t+1} = A^j x_t + B^j \Lambda^j(\Pi_t, \Pi_t^*) x_t + C^j u_t + \Theta_x^j \epsilon_t$$

At the beginning of period t+1, policymakers observe a vector of variables ω_{t+1}. Each of the k possible models implies a structural relationship for the observation vector ω_{t+1}^i:

$$\omega^i_{t+1} = V^i_x x_t + V^i_e e_t + V^i_u u_t + V^i_\theta \epsilon_t$$

where we assume that V^i_θ is invertible. 1/

Let $E(\omega^i_{t+1})$ be the expected value of ω^i_{t+1} (evaluated at $E(\epsilon_t) = 0$). The value of the shock implied by each model is therefore:

$$\epsilon^i_t = (V^i_\theta)^{-1}[\omega^i_{t+1} - E(\omega^i_{t+1})]$$

The new Bayesian priors are then given by:

$$\Pi^i_{t+1} = \frac{Pr(\epsilon^i|\Sigma)\ \Pi^i_t}{\left[\sum_{i=1}^{k} Pr(\epsilon^i|\Sigma)\ \Pi^i_t\right]}$$

where $Pr(\cdot)$ is the probability that a vector shock, distributed $N(0,\Sigma)$, takes the value ϵ^i. The state variables in period t+2 are then generated with a drawing for ϵ_{t+1}, and the whole process is repeated.

V. Simulation Results

In addition to the model parameters, the simulation analysis requires specification of the policymakers' discount factors--chosen to be 0.95, that is a discount rate of 5 percent per annum; the utility weights on each target; and the relative weight each country receives in the social planner's objective function. The utility weights were taken from the revealed preference estimates of Oudiz and Sachs [1984] with policymakers targeting inflation and output; in the cooperative solution, each country is given equal weight in the social planner's objective function. The observation vector includes output, interest rates, producer prices and the nominal exchange rate in each country. The optimal policies and value functions depend non-linearly on the probability assigned to each of the three models. The recursive optimization had to be done on a two

---

1/ If policy-makers observe fewer variables than the number of shocks in ϵ then they face a signal extraction problem concerning the shocks as well as the models.

dimensional grid; that is, for each possible combination, $(\Pi^1, \Pi^2, (1-\Pi^1-\Pi^2))$, the recursive algorithm outlined above must be solved. 1/

In the first set of simulations all agents assign an initial probability to each model and do not update their priors (to repeat, we assume throughout that all agents in the model--both private and public--share a common set of subjective priors). The results are reported conditional on each of the three models being the true model. The figures represent the total present value of disutility for the world economy and are expressed in terms of the GDP equivalent utility. As is evident below, some of the simulations exhibit explosive behavior and the present value of disutility may be undefined; in these cases we have simply marked the entry "explosive" and indicate whether coordination is welfare improving or deteriorating. Since there is an additive random shock, ϵ_t, all figures depend upon the specific realizations (which are drawn from a Normal distribution with zero mean and the estimated variance-covariance matrix of the "true" model 2/), and hence we have taken the average of ten stochastic simulations (the drawings are the same for all of the models). The optimal policies are designed to stabilize output and inflation against the specific random structural shocks applied to the model; as such the optimal policies under both cooperative and non-cooperative behavior are not easily interpretable and are therefore not reported.

There are a number of noteworthy points about Tables 4-6, in particular concerning the results in the top panel. First, on average the gains from coordination are not spectacular but are certainly measurable, amounting to a permanent increase of perhaps one or two percent of GDP per year; an estimate which is in line with those of previous studies, e.g. Oudiz and Sachs [1984]. Both the Nash and the cooperative equilibria are significantly better than a pure float regime for model I; this is not true of model III, however, where floating is intermediate between the other two regimes. Second, welfare under the cooperative equilibrium is an increasing function of the probability assigned to the true model. By contrast, in the non-cooperative regime, the relationship is not monotonic with welfare generally rising with the probability of the true model in simulations I and II but decreasing when model III is true. Welfare in the regime of freely floating exchange rates with exogenous money targets should not depend upon the prior weights associated with each model since there is no active policy intervention at all. This would hold exactly in a static model (or a dynamic model in which there were no forward-looking variables) but not in a

---

1/ We used a grid of eleven equally spaced intervals (0.0-1.0 inclusive) giving a total of 121 optimization problems to solve.

2/ None of the estimated correlations between shocks was significant at the 5 percent level, so a diagonal variance-covariance matrix was used to generate the shocks.

Table 4. Disutility Levels, No Model Learning
(Model I is the true model)

Probability assigned to models			Regime		
I	II	III	Cooperative	Nash	Float
0.1	0.45	0.45	41.5	65.8	80.7
0.2	0.40	0.40	41.1	60.3	81.2
0.3	0.35	0.35	40.8	57.3	81.7
0.4	0.30	0.30	38.5	56.7	81.1
0.5	0.25	0.25	38.5	55.0	81.6
0.6	0.20	0.20	37.3	54.8	80.6
0.7	0.15	0.15	37.3	53.4	80.8
0.8	0.10	0.10	36.6	53.6	78.6
0.9	0.05	0.05	36.6	55.7	78.4
1.0	0.00	0.00	36.5	55.7	80.9
0.1	0.9	0.0	393.6	104.5	78.1
0.2	0.8	0.0	58.5	103.0	78.2
0.1	0.0	0.9	Explosive	5200	88.0
0.2	0.0	0.8	4567	66.0	87.0

Table 5. Disutility Levels, No Model Learning
(Model II is the true model)

Probability assigned to models			Regime		
I	II	III	Cooperative	Nash	Float
0.45	0.1	0.45	39.2	37.5	52.4
0.40	0.2	0.40	37.9	37.7	52.3
0.35	0.3	0.35	34.0	38.1	52.1
0.30	0.4	0.30	33.7	38.2	51.9
0.25	0.5	0.25	31.5	38.7	51.7
0.20	0.6	0.20	31.3	38.9	51.6
0.15	0.7	0.15	29.8	39.4	51.4
0.10	0.8	0.10	29.5	39.0	51.3
0.05	0.9	0.05	28.4	40.0	51.2
0.00	1.0	0.00	28.4	42.0	51.1
0.9	0.1	0.0	30.5	40.2	56.2
0.8	0.2	0.0	30.2	40.3	54.5
0.0	0.1	0.9	2.27×10^6	6.8×10^5	53.0
0.0	0.2	0.8	12975	1002	52.7

Table 6. Disutility Levels, No Model Learning
(Model III is the true model)

Probability assigned to models			Regime		
I	II	III	Cooperative	Nash	Float
0.45	0.45	0.1	30.0	63.8	41.3
0.40	0.40	0.2	28.1	60.5	41.2
0.35	0.35	0.3	27.1	57.5	41.3
0.30	0.30	0.4	26.4	58.3	41.2
0.25	0.25	0.5	25.7	52.9	41.2
0.20	0.20	0.6	25.2	50.6	41.2
0.15	0.15	0.7	24.6	46.8	41.2
0.10	0.10	0.8	24.1	43.8	41.2
0.05	0.05	0.9	23.7	39.9	41.0
0.00	0.00	1.0	23.4	41.8	40.9
0.9	0.0	0.1	29.3	58.5	41.4
0.8	0.0	0.2	28.6	56.9	41.3
0.0	0.9	0.1	35.6	915.3	41.2
0.0	0.8	0.2	30.8	99.9	41.5

rational expectations model since the endogenous variables depend upon the expected future path of the nominal exchange rate, which in turn is a function of the subjective priors over the models held by the private sector.

Now consider the results reported below the line in Tables 4-6. Here we have assigned a low probability to the true model and a high probability to one of the competing models. Suppose, for example, that model I is true but agents assign a prior weight of only 10 percent to model I and 90 percent to model II. In that case, the cooperative regime actually yields substantially lower ex post welfare (Table 4). The situation is more critical when a large weight is attached to model III (and model I is correct); here the cooperative solution results in (eventually) explosive behavior of the economy or in large disutility. The Nash equilibrium also implies large disutility when the weights are 90 percent model III and 10 percent model I. In fact, the most robust policy to follow is one of complete non-intervention. A similar outcome results when model II is the true model, but a large weight is put on model III (Table 5). The large losses and possible instability no doubt result from the fact that model III implies a long-run tradeoff between output and inflation, which activist policies (either Nash or cooperative) try to exploit--unsuccessfully, it turns out, because the true model, either model I or model II, is in fact a natural rate model. In either case, coordination severely reduces welfare to the extent of making the sustainability of a coordinated regime infeasible. However, it is important to note that it is not only the cooperative regime that may result in low welfare when policymakers use the wrong model. For example, if model III is the true model, a low prior on the true model results in the Nash equilibrium being highly inefficient, so that there are very large gains from coordination (Table 6).

The conclusion that emerges from the above results is therefore that a policy rule that does not require active intervention--corresponding to exogenous money supplies and floating exchange rates--is the safest policy when policymakers are uncertain about the model. Our results therefore provide some support to the advocacy of fixed rules in preference to activist policies, which Friedman has long argued may be destabilizing (see, for instance, Friedman [1948]). 1/ However, when there is no uncertainty about the model the level of welfare under cooperation is considerably larger than under non-intervention.

The results suggest that when policymakers have the wrong priors over the models, and do not undertake updating, any policy intervention --coordinated or uncoordinated--can be dangerous. However, it is implausible that policymakers would not update their subjective priors when they find

---

1/ There are other considerations that are relevant to the choice between fixed rules and more activist policies; see for instance Fischer [1988].

that their expectations about the effects of their policies are consistently invalidated. In a second set of simulations, therefore, we assume that agents update their priors over the true model in the Bayesian fashion described above. Tables 7-9 report the discounted present value of disutility under cooperative and non-cooperative behavior, as well as the floating exchange rate regime with exogenous money targets, for simulations where initial probability priors are as given in the tables.

The important conclusion that emerges from these tables is that policies no longer become unstable even when the initial priors attributed to the true model are very low. Furthermore, we find that coordination is always welfare improving. Consider, for example, the case in which agents assign a prior probability of 10 percent to model I and 90 percent to model III, when model I is the correct model. In the absence of Bayesian updating, coordination reduces welfare substantially (Table 4) since the coordinated regime is dynamically unstable. From Table 7, however, Bayesian learning, combined with the same initial priors, results in a welfare gain equivalent to 20 percent of GDP (a disutility of 38.2, compared to 58.4 under Nash, expressed in GDP-equivalent percentages). 1/ Similarly, when model II is true, but policy makers assign a probability weight of 80 percent to model III, the large loss from policy coordination in the absence of learning (Table 5) becomes a modest gain (Table 8) once learning is allowed. With endogenous model learning, the gains from coordination range from the utility equivalents of about 13 percent of GDP to about 25 percent of GDP. Again, welfare in the cooperative equilibrium is an increasing function of the initial probability assigned to the true model, while in the Nash equilibrium the relationship is non-monotonic. The non-interventionist policy regime now performs consistently worse than the cooperative regime but is sometimes better than the Nash equilibrium.

The examples where policymakers assign a high initial weight to model III, but where in reality one of the other two models is correct, has some relevance to the history of demand management in the postwar period. Early models of the "Phillips curve" (see, for instance, Phillips [1958] and Lipsey [1960]) implied that there was a tradeoff between the rate of change of wages and prices and output or the unemployment rate. These models no doubt helped induce central banks and treasuries to engage in demand expansion, in an attempt to buy more output growth at what was judged to be an acceptable inflation cost. The experience of accelerating inflation beginning in the late sixties forced economists and policymakers to reconsider those models, and there has been a profound shift in policy away from short-term fine-tuning and to a concern for the medium-term inflation

---

1/ Nevertheless, a superior policy ex post would have been to follow a non-interventionist policy until policymakers had reduced their probability weight on the incorrect model (model III) sufficiently, and then had cooperated.

Table 7. Disutility Levels, Bayesian Model Learning
(Model I is the true model)

Probability assigned to models			Regime		
I	II	III	Cooperative	Nash	Float
0.1	0.45	0.45	36.8	56.1	78.6
0.2	0.40	0.40	36.7	60.3	79.0
0.3	0.35	0.35	36.7	56.3	79.3
0.4	0.30	0.30	36.6	56.3	80.0
0.5	0.25	0.25	36.6	56.3	80.6
0.6	0.20	0.20	36.6	56.4	80.7
0.7	0.15	0.15	36.5	56.4	80.7
0.8	0.10	0.10	36.5	55.9	80.9
0.9	0.05	0.05	36.5	56.0	81.1
1.0	0.00	0.00	36.5	55.7	80.9
0.1	0.9	0.0	38.3	60.6	79.1
0.2	0.8	0.0	38.0	60.1	79.9
0.1	0.0	0.9	38.2	58.4	78.5
0.2	0.0	0.8	38.1	59.0	79.3

Table 8. Disutility Levels, Bayesian Model Learning
(Model II is the true model)

Probability assigned to models			Regime		
I	II	III	Cooperative	Nash	Float
0.45	0.1	0.45	29.2	40.5	56.6
0.40	0.2	0.40	29.2	40.7	55.9
0.35	0.3	0.35	29.1	40.8	55.3
0.30	0.4	0.30	29.1	40.9	55.2
0.25	0.5	0.25	29.2	40.8	54.3
0.20	0.6	0.20	29.1	41.2	54.1
0.15	0.7	0.15	29.1	41.3	53.3
0.10	0.8	0.10	29.0	41.4	53.3
0.05	0.9	0.05	28.8	41.6	51.0
0.00	1.0	0.00	28.4	42.0	51.1
0.9	0.1	0.0	29.3	42.4	57.4
0.8	0.2	0.0	29.2	42.8	56.9
0.0	0.1	0.9	28.8	42.3	51.0
0.0	0.2	0.8	28.5	42.6	51.0

Table 9. Disutility Levels, Bayesian Model Learning
(Model III is the true model)

Probability assigned to models			Regime		
I	II	III	Cooperative	Nash	Float
0.45	0.45	0.1	24.4	45.6	43.3
0.40	0.40	0.2	24.3	45.1	43.1
0.35	0.35	0.3	24.1	44.9	42.8
0.30	0.30	0.4	23.9	44.5	42.6
0.25	0.25	0.5	23.9	44.7	42.3
0.20	0.20	0.6	23.7	44.1	42.0
0.15	0.15	0.7	23.7	44.2	41.8
0.10	0.10	0.8	23.6	43.5	41.5
0.05	0.05	0.9	23.5	43.1	42.5
0.00	0.00	1.0	23.4	41.8	40.9
0.9	0.0	0.1	25.9	51.1	45.4
0.8	0.0	0.2	25.8	50.1	44.2
0.0	0.9	0.1	25.4	49.2	42.5
0.0	0.8	0.2	25.1	48.1	42.5

consequences of policy. Moreover, the rationale for the policy changes has been acceptance of natural rate models which do not allow for monetary stimulus to have permanent positive effects on the level of activity. Friedman [1977, p. 470] commented on the change in policy in the following terms in his Nobel lecture:

> Government policy about inflation and unemployment has been at the center of political controversy. Idealogical war has raged over these matters. Yet the drastic change that has occurred in economic theory has not been a result of ideological warfare. It has not resulted from divergent political beliefs or aims. It has responded almost entirely to the force of events: brute experience proved far more potent than the strongest of political or ideological preferences.

VI. Caveats and Discussion

Having found that in the presence of model learning that the cooperative regime dominates non-cooperation and non-intervention in the context of our model, we tried to gauge the sensitivity of the results to our assumptions. The ex post performance of the coordinated regime depends on two factors: how robust is the optimal policy to model errors; and how fast is the rate of learning of the model. Clearly, the more diverse the models the more likely is the optimal policy based on the "wrong" priors to result in lower welfare. We also thought that we might have assumed too little model uncertainty in choosing the structurally similar models of section III. However, in further simulations with more diverse models we found that as the degree of model uncertainty increased (in the sense of lower ζ ratios) the rate of model learning increases. The intuition is straightforward: if the models are very different then the implied observation vector of each model is also very different and the updated priors assigned to the false models will be correspondingly low. Thus, although greater diversity between the models makes more likely the possibility of welfare-deteriorating policies, it also serves to reduce the model uncertainty (at least when there is a finite number of alternative models, as is the case here).

It is possible to decrease the rate of learning by increasing the variance of the additive shocks ϵ_t. The greater the variance of ϵ_t the greater the noise in the updating observations and therefore the slower the rate of learning. We checked the sensitivity of our results by using a variance matrix with ten times the estimated standard errors. Despite this large increase in error variances, however, we were unable to reverse our conclusions about the ranking of the coordinated and uncoordinated regimes in the presence of learning.

It must be emphasized however that our experiments assume knowledge of the variance-covariance matrix describing the shocks. Forcing agents to estimate it would introduce considerably more uncertainty, and make it more difficult to infer which model was correct. So would a combination of temporary and permanent shocks.

Another caveat is that we have simplified the problem to one where policymakers must just discover the unique, unchanging model describing the economy. An alternative assumption is that agents can never perfectly anticipate the true model when setting policies because the true model is stochastic. If the vector of subjective priors converges to the true probability that the model is realized then the average ex post welfare gain will--at least in large samples--equal the ex ante expected welfare gain from coordination. In that case, we return to the ex ante expected welfare criterion in which coordination is necessarily welfare improving. 1/ More realistic is the case where the true model is changing in a non-random way as a result of structural shifts, and where these shifts occur frequently enough so that the distribution describing the models is never completely learned. We have yet to explore this case.

VII. Conclusions

In this paper we have discussed whether coordination is likely to reduce the actual ex post level of welfare when policy makers are uncertain about the effects of policies. We have found no evidence in our simulations that policy coordination is likely to reduce welfare vis a vis non-cooperative policies; however, a simple non-intervention regime such as a pure floating exchange rate may be the most robust policy in the presence of model uncertainty. These conclusions are of course specific to a particular model, and are based here on an estimated version of the Mundell-Fleming model with sticky prices and rational exchange rate expectations. More experimentation with other models is no doubt necessary in order to gauge whether the conclusions can be generalized.

Once we introduced endogenous model learning we found that coordination always results in higher welfare: though the gains from coordination are not spectacular they appear to be at least positive. We were unable to generate losses from coordination even by increasing the variance of the additive noise. No doubt, the representation of the economy is much too simple. What the results suggest, however, is that the conclusion that coordination is as likely to decrease welfare as to increase it (see Frankel and Rockett

---

1/ This was verified in simulations in which the true model was drawn stochastically from the three possible models, each of which had equal probability, and agents undertook Bayesian learning as above.

[1988]) is not consistent with the joint assumptions that there is a single "true" model and that agents learn about that model in a Bayesian fashion.

Clearly, an important contribution of future research would be to characterize the degree and nature of model uncertainty in medium scale models of the world economy, and to gauge whether that is an adequate measure of real-world uncertainty. To this end, we intend in our further work to evaluate existing models using an explicitly Bayesian framework and to derive a stable distribution of possible models from which the economy's structure in any period may be viewed as a drawing.

References

Blume, Lawrence E., and David Easley, "Learning to Be Rational," Journal of Economic Theory, vol. 26 (April 1982), pp. 340-51.

_____, David Easley and Margaret Bray, "Introduction to the Stability of Rational Expectations Equilibrium," Journal of Economic Theory, vol. 26 (April 1982), pp. 313-17.

Brainard, William C., "Uncertainty and the Effectiveness of Policy," American Economic Review, vol. 57 (May 1967), pp. 411-25.

Branson, William H. and Willem H. Buiter, "Monetary and Fiscal Policy with Flexible Exchange Rates," in J. Bhandari and B. Putnam, eds. Economic Interdependence and Flexible Exchange Rates (Cambridge, Mass.: MIT Press, 1983), pp. 251-85.

Bray, Margaret, "Learning, Estimation, and the Stability of Rational Expectations," Journal of Economic Theory, vol. 26 (April 1982), pp. 318-39.

_____ and N.E. Savin, "Rational Expectations Equilibria, Learning, and Model Specification," Econometrica, vol. 54 (September 1986), pp. 1129-60.

Canzoneri, Matt and Joanna Gray, "Two Essays on Monetary Policy in an Interdependent World," International Finance Discussion Paper No. 219 (Washington, D.C.: Federal Reserve Board, 1983).

Cooper, Richard, The Economics of Interdependence (New York: McGraw-Hill, 1968).

Cyert, R.M. and M.H. De Groot, "Rational Expectations and Bayesian Analysis," Journal of Political Economy, vol. 82 (1974), pp. 521-36.

De Canio, Stephen J., "Rational Expectations and Learning from Experience," Quarterly Journal of Economics, vol. 93 (February 1979), pp. 47-57.

Dornbusch, Rudiger, "Expectations and Exchange Rate Dynamics," Journal of Political Economy, vol. 84 (December 1976), pp. 1161-76.

Feldstein, Martin, "Distinguished Lecture on Economics in Government: Thinking about International Economic Coordination," Journal of Economic Perspectives, vol. 2 (Spring 1988), pp. 3-13.

Fischer, Stanley, "Rules versus Discretion in Monetary Policy," NBER Working Paper No. 2518, February 1988.

Frankel, Jeffrey, "Ambiguous Policy Multipliers in Theory and in Empirical Models," in R. Bryant and others, eds. Empirical Macroeconomics for Interdependent Economies (Washington, D.C.: the Brookings Institution, 1988), pp. 17-26.

_____, "Obstacles to International Macroeconomic Policy Coordination," International Monetary Fund WP/87/29, April 1987.

_____ and Katharine Rockett, "International Macroeconomic Policy Coordination when Policy-Makers Do Not Agree on the True Model," American Economic Review, vol. 78 (June 1988), pp. 318-40.

Friedman, Milton, "A Monetary and Fiscal Framework for Economic Stability," [1948], reprinted in Essays in Positive Economics (Chicago: University of Chicago Press, 1953).

_____, "Inflation and Unemployment," Journal of Political Economy, vol. 85 (June 1977), pp. 451-72.

_____, "The Role of Monetary Policy," American Economic Review, vol. 58 (March 1968), pp. 1-17.

Ghosh, Atish, "International Policy Coordination in an Uncertain World," Economics Letters, vol. 21, no. 3, (1986), pp. 271-276.

_____ and Paul R. Masson, "International Policy Coordination in a World with Model Uncertainty," International Monetary Fund Staff Papers, vol. 35, (June 1988), pp. 230-58.

Ghosh, Atish and Swati Ghosh, "International Policy Coordination When the Model is Unknown," mimeo, Geneva, Switzerland, December 1986.

Goldfeld, Stephen, "The Demand for Money Revisited," Brookings Papers on Economic Activity, 3:1973, pp. 577-637.

Hamada, Koichi, "Alternative Exchange Rate Systems and the Interdependence of Monetary Policies," in R.Z. Aliber, National Monetary Policies and the International Financial System (Chicago: University of Chicago Press, 1974), pp. 13-33.

Holtham, Gerald and Andrew Hughes Hallett, "International Policy Coordination and Model Uncertainty," in R. Bryant and R. Portes, eds., Global Macroeconomics: Policy Conflict and Cooperation, (London, MacMillan,1987).

Lipsey, R.G., "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1862-1957: A Further Analysis," Economica, vol. 27 (February 1960).

Marcet, Albert and Thomas J. Sargent, "The Fate of Systems with 'Adaptive' Expectations," American Economic Review, vol. 78 (May 1988), pp. 168-72.

Masson, Paul R., and others, "MULTIMOD: A Multi-Region Econometric Model," Staff Studies for the World Economic Outlook (Washington, D.C.: International Monetary Fund, July 1988), pp. 50-104.

McKibbin, Warwick and Jeffrey Sachs, "Coordination of Monetary and Fiscal Policies in the OECD," in J. Frenkel, ed. International Aspects of Fiscal Policy, (Chicago, University of Chicago Press, 1988).

Oudiz, Gilles and Jeffrey Sachs, "Macroeconomic Policy Coordination among Industrial Economies," Brookings Papers on Economic Activity, 1:1984, pp. 1-75.

______, "International Policy Coordination in Dynamic Macroeconomic Models," in W.H. Buiter and R.C. Marston, eds. International Economic Policy Coordination, (Cambridge, Cambridge University Press, 1985).

Phillips, A.W., "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," Economica, vol. 25 (November 1958), pp. 283-99.

Rogoff, Kenneth, "Can International Monetary Policy Cooperation be Counterproductive?" Journal of International Economics vol. 18, (May 1985), pp. 199-217.

Sachs, Jeffrey, "International Policy Coordination in a Dynamic Macroeconomic Model," NBER Working Paper No. 1166, July 1983.