

Stabilization Theory and Policy: 50 Years after the Phillips Curve

Stephen J. Turnovsky*
University of Washington

May 2008

*I am grateful to David Kendrick for discussions on these issues over the years, as well as his comments on this draft. Support was provided in part by the Castor Endowment at the University of Washington.

(Plenary address to be presented at ESAM08 Conference, Markets and Models: Policy Frontiers in the A.W.H. Phillips Tradition, July 9-11, 2008, Wellington, New Zealand)

1. Introduction

A key objective of macroeconomic policy is to maintain economic stability. Interest in the topic originated with Tinbergen (1952). Employing a static linear framework he proved the now classic proposition stating that under certainty the policymaker need use only as many policy instruments as there are independent target variables in order to achieve any desired values for these target variables.¹ Extra policy variables are redundant, while with insufficient instruments not all objectives can be achieved simultaneously. In practice, with the economy being inherently dynamic, it is clearly important to cast the theory of stabilization in a dynamic context, thereby enabling us to consider the stability of the economy as it evolves over time as well as in response to unforeseen stochastic disturbances that may occur at any point in time.²

Few would dispute the proposition that Bill Phillips was a pioneer in the development of dynamic stabilization policy. His contributions were manifest in two seemingly different, but as it turns out highly inter-related, areas. The first, and more direct, contribution was contained in a pair of papers published in the *Economic Journal* in the 1950's; Phillips (1954, 1957). These papers draw upon his background as an engineer and are the first papers to apply feedback control methods to the stabilization of a macroeconomy. Today that is a burgeoning field, and despite challenges stemming from the subsequent development of rational expectations, the application of control methods is now an integral part of the analysis of dynamic economic systems.³

Like many fundamental contributions, Phillips' initial contribution was simple. Previously, Samuelson (1939) and Hicks (1950) had shown how, if one combines the multiplier in consumption with the accelerator in investment, one can derive a dynamic equation determining the evolution of

¹ Brainard (1967) re-examined the Tinbergen proposition in a simple stochastic model and showed how it ceases to apply once multiplicative stochastic disturbances are introduced. Henderson and Turnovsky (1972) showed how adjustment costs associated with policy instruments also leads to its break down.

² Preston (1974) referred to Tinbergen's theorem as one of "static controllability" and using results from control engineering developed an analogous condition for the controllability of a linear dynamic system.

³ There are in fact several societies and research outlets specifically focused on these types of issues. In the late 1970s, the Society of Economic Dynamics and Control (now the Society of Economic Dynamics) and the *Journal of Economic Dynamics and Control*, were founded to foster the application of control methods to economics. More recently, the Society for Computational Economics and journals such as the *Review of Economic Dynamics*, *Macroeconomic Dynamics*, and *Computational Economics* have become established and are further testimony to the flourishing research activity in this general area.

national income (output). The precise nature of this relationship depends upon how the lags generating the dynamics are introduced and many ways to accomplish this exist. The dynamics can be expressed in discrete time, as employed by Samuelson and Hicks, or in continuous time, as for example illustrated by Allen (1956) and used by Phillips himself. Phillips took these simple aggregate models and showed how, if one introduces a government that instead of remaining passive follows some active policy intervention rule, then it will be able to influence the dynamic time path of the economy.

The second contribution relates to the celebrated Phillips curve (1958). While this was originally proposed as an empirical relationship between (wage) inflation and unemployment and has spawned generations of empirical research, its introduction into the macroeconomic system turns out to have potentially profound consequences for the efficacy of stabilization policy.

Our objective in this essay is to review and evaluate the impact of Phillips' seminal work of half a century ago on subsequent developments in macroeconomic stabilization policy. We begin by first establishing the importance of stabilization policy in the development of economics, and in so doing underscore the significance of Phillips' contributions. We shall then discuss various issues in more detail. Sections 3 and 4 briefly review the formulation of linear stabilization rules adopted by Phillips and show how these relate to the optimal stabilization rules that emerge from conventional linear-quadratic optimization problems. These originated in the engineering literature, but turned out to be most convenient for the formal analysis of optimal stabilization policy. Most of the early stabilization literature assumed fixed prices, or in any event was associated with "backward looking" or "sluggish" variables. However, the development of rational expectations in the 1960's and 1970's posed a challenge for stabilization policy and this is discussed in Section 5. This arose from the role of "forward-looking" inflationary expectations in the Phillips curve, and the effect this had in the design of optimal stabilization rules, through issues such as the "Lucas Critique" and the "time consistency" of policy.

Section 6 briefly comments on a long-standing debate, the merits of fixed policy rules versus discretionary or optimal policy. Section 7 discusses some of the more contemporary aspects of stabilization policy and will serve to illustrate the durability of Phillips' contributions. The

“expectations-augmented Phillips curve” and “New-classical Phillips curves” of the 1970s are now replaced by the “New-Keynesian Phillips Curve”. The methods of optimal linear-quadratic stabilization theory of the 1970s is now applied as an approximation to more general utility functions. In addition the linear feedback control rules initially proposed by Phillips have now been introduced into multi-agent dynamic games, while issues of learning are receiving increasing attention..

2. The Significance of Stabilization Policy

In the last 39 years, 11 of the Nobel prizes awarded in economics have been in the general area of macroeconomics.⁴ Of these, four include the area of stabilization policy and contain this term or some close substitute in the citation. The first of these was awarded to Milton Friedman in 1976, with the citation stating: “for his achievement in the fields of consumption analysis, monetary history and theory and for his demonstration of the complexity of **stabilization policy**”. In 1995 the prize was awarded to Robert Lucas “for having developed and applied the hypothesis of rational expectations, and hereby having transformed macroeconomic analysis and deepened our understanding of **economic policy**”. The 2004 recipients, Finn Kydland and Edward Prescott were cited “for their contributions to dynamic macroeconomics: the time consistency of **economic policy** and the driving forces behind the business cycles”. Most recently, the citation for the 2006 recipient, Edmund Phelps, included the statement “for his analysis of intertemporal tradeoffs in **macroeconomic policy**.” The more detailed statement provided by the Sveriges Riksbank refers explicitly to the “so-called Phillips curve” and to the “expectations augmented Phillips curve”, in describing Phelps’ contribution.

While the contributions for which these four prizes were awarded all extend beyond stabilization policy, particularly in the case of Friedman, who was in fact skeptical of active stabilization policy, they nevertheless share several common themes. These include concepts such as “economic policy”, “Phillips curves”, “price expectations”, “macroeconomic dynamics”. There

⁴ This is as categorized by Assar Lindbeck at http://nobelprize.org/nobel_prizes/economics/articles/lindbeck/index.html. Jan Tinbergen, Paul Samuelson, and John Hicks, who were also early recipients and made important contributions in this area, were cited for more general contributions.

are also substantive inter-relationships between the awards. While Phelps developed a formal technical derivation of the expectations-augmented Phillips curve, Friedman provided an informal version in his 1968 American Economic Association presidential address.⁵ This formulation is also closely related to the form of supply function that was central to Lucas' work (the "Lucas supply function"). In any event, it is evident that Bill Phillips' insights are reflected in these contributions and that had it not been for his untimely death in 1975, he, himself, surely would have been an early recipient of a Nobel prize.

3. Phillips' Policy Rules

3.1 A simple textbook macrodynamic model

The Phillips analysis was based on the dynamic multiplier-accelerator model, previously developed by Samuelson (1939) and Hicks (1950). There are numerous versions of this model and we shall introduce the simplest formulation employed by Phillips. He expressed it using continuous time, which is more convenient for the purpose of establishing the implied dynamic behavior, but essentially the same conclusions can be reached using discrete time as did Samuelson and Hicks.⁶

Aggregate demand of the economy at time t , $Z(t)$, is defined by

$$Z(t) = C(t) + I(t) + G(t) \tag{1}$$

where $C(t)$ denotes consumption, $I(t)$ denotes investment, and $G(t)$ denotes government expenditure. Dynamics can be introduced in various ways. Whereas Samuelson and Hicks did so by introducing lags into consumption and investment behavior, Phillips did so by assuming gradual product market clearance. This is specified by

$$\dot{Y}(t) = \alpha [Z(t) - Y(t)] \quad \alpha > 0 \tag{2}$$

⁵ Friedman (1968), Phelps (1968).

⁶ The contemporary literature on stabilization policy almost always employs discrete time; see e.g. Woodford (2003). Discrete time is in fact much more convenient for capturing some of the recent theoretical developments, which sometimes depend upon subtle issues of timing. For example, the difference between the "New-Classical" and "New-Keynesian" Phillips curve is one of timing, a difference that can be best captured using discrete time. In our exposition we shall introduce time in whichever way is more convenient.

where $Y(t)$ denotes aggregate supply at time t , and the dot denotes time derivative. If aggregate demand exceeds output, supply is increased at a rate proportional to excess demand and vice versa.

To complete the model, behavioral hypotheses must be introduced for consumption and investment. The simplest of these is to specify that consumption is proportional to current output

$$C(t) = cY(t) \qquad 0 < c < 1 \qquad (3)$$

and a constant rate of investment, $I(t) = I$. If, further, we assume that government spending is constant as well, then combining these equations we see that equilibrium output evolves in accordance with the simple equation, specifying the textbook dynamic multiplier model

$$\dot{Y}(t) = \alpha [(c-1)Y(t) + I + G] \qquad (4)$$

Phillips' contribution was to introduce various policy rules for $G(t)$. Much of this was developed and can be discussed in terms of this simple model. However, most of the subsequent literature, as well as much of Phillips' own contributions endogenized investment by employing some form of the accelerator theory. The effect of this is to increase the order of the equilibrium dynamics, thereby generating a richer array of time paths for output and other relevant variables. For present purposes, the simpler model suffices.

Before discussing the policy rules introduced by Phillips, we should briefly observe the behavior of the economy implied by equation (4). With I and G fixed, it is a linear first-order differential equation in Y , and provided $0 < c < 1$, the evolution is stable and will converge monotonically to the stationary equilibrium level

$$\bar{Y} = \frac{I + G}{1 - c} \qquad (5)$$

This will be immediately recognized as being the equilibrium level of income in the simplest static linear macroeconomic model.

3.2 Policy rules

Within this framework, Phillips introduced government expenditure as an active policy

variable that is continuously adjusted to meet certain specified objectives. In doing so he emphasized the lags associated with adjusting the policy instrument itself. These are often referred to as *policy lags*, and reflect delays in implementing decisions due to, for example, the political process. They are quite distinct from lags from the underlying economic structure, such as those embodied in the market disequilibrium relationship (2), which may be appropriately characterized as being *system lags*.⁷

Phillips assumed that the policy actually implemented at any point in time adjusts only gradually to past policy decisions. Thus, if $G^d(t)$ is the desired value of the policy variable chosen at time t (the policy decision), the actual value of the policy variable, $G(t)$, is adjusted in accordance with:⁸

$$\dot{G}(t) = \beta(G^d(t) - G(t)) \quad \beta > 0 \quad (6)$$

The desired value of the policy variable, $G^d(t)$, is related by some rule to the ultimate target objective that he took to be the stabilization of national income. Phillips proposed three such policy rules, which he called: (i) proportional policy, (ii) integral policy, and (iii) derivative policy, all of which influenced the dynamics of the economy in different ways, having both desirable and undesirable effects on its evolution. These terms did not originate with Phillips. Rather, they were part of the tradition of classical control, where engineers referred to them as PID feedback rules. We shall briefly discuss each in turn

3.1.1 Proportional policy

This was specified by Phillips to be

$$G^d(t) = -\gamma_p(Y - Y^*) \quad \gamma_p > 0 \quad (7a)$$

where Y^* is the (desired) target level of output. The parameter γ_p represents the intensity of the

⁷ These two kinds of lags are also sometimes referred to as being “inside lags” and “outside lags”, respectively.

⁸ Solving equation (6), the actual policy at time t is $G(t) = \beta \int_{-\infty}^t G^d(s) e^{-\beta(t-s)} ds$, which is seen to be an exponentially declining weighted average of past policy decisions. As $\beta \rightarrow \infty$, the desired policy is fully implemented immediately.

policy maker's desired policy response when output deviates from its target. According to the rule specified in (7a), if $Y(t) < Y^*$, then $G^d(t) > 0$; if $Y(t) > Y^*$, then $G^d(t) < 0$. Since the rule may require $G^d(t) < 0$, Phillips interprets it as “net fiscal stimulus” rather than pure government spending, which by its nature is non-negative. Thus (7a) asserts that the desired net fiscal stimulus is proportional, but opposite to, the deviation between current and desired level of output. Combining equations (4), (6) and (7a), the dynamic evolution of the economy is described by the following pair of equations

$$\begin{pmatrix} \dot{Y} \\ \dot{G} \end{pmatrix} = \begin{pmatrix} \alpha(c-1) & \alpha \\ -\beta\gamma_p & -\beta \end{pmatrix} \begin{pmatrix} Y \\ G \end{pmatrix} + \begin{pmatrix} \alpha I \\ \beta\gamma_p Y^* \end{pmatrix} \quad (8)$$

Three observations about this system are worth noting. First, the necessary and sufficient conditions for stability are: (i) $1 - c + \gamma_p > 0$, (ii) $1 - c + \beta/\alpha > 0$ so that it is clear that γ_p, β , which characterize the stabilization policy and implementation will influence the dynamics. While (i) and (ii) will certainly be met if $0 < c < 1$, it may be possible to stabilize the system even in the implausible event where $c > 1$. Second, the eigenvalues to (8) will be complex if and only if $4\alpha\beta\gamma_p > [\beta - \alpha(1-c)]^2$, implying that policy lags may induce cycles into the adjustment. This is presumably undesirable, but hardly surprising, since with policies taking time to implement, by the time $G^d(t)$ is yielding its desired effect, the conditions leading to that decision may have changed, causing the system to over-adjust during the transition.⁹

Third, output in (8) converges to the stationary level

$$\bar{Y} = \frac{1 + \gamma_p Y^*}{1 - c + \gamma_p} \neq Y^* \quad (9)$$

That is, in general, the level of output will fail to converge to its desired target level. This was viewed by Phillips as being an undesirable feature of the proportional policy rule, but in fact it can be regarded as reflecting an inadequate specification of the rule, as given in (7a). This formulation ignores the fact that given the behavior of the private sector as reflected by C, I , the government

⁹ The policy parameters β, γ_p also affect the *speed* of convergence. While this was not an aspect that concerned Phillips, speeds of convergence have assumed an important role in contemporary macrodynamics, particularly in the dynamics of growth.

must also choose an appropriate target level of expenditure, G^* , if it wishes to attain Y^* in the long run. This appropriate level is determined by the stationary relationship

$$G^* = (1-c)Y^* - I$$

Once this fact is recognized, it becomes natural to express (7a) in deviation form

$$G^d(t) - G^* = -\gamma_p(Y - Y^*) \quad (7a')$$

in which case the stability conditions remain unchanged, and if satisfied, ensure that output converges to its target, Y^* .

3.1.2 Integral policy

As an alternative policy, Phillips introduced the possibility that $G^d(t)$ is determined by the integral (sum) of past deviations in output from its target, rather than only just the current deviation. This is specified by

$$G^p(t) = -\gamma_i \int_0^t [Y(s) - Y^*] ds \quad \gamma_i > 0 \quad (7b)$$

Differentiating with respect to t enables the policy to be written in the equivalent form

$$\dot{G}^p(t) = -\gamma_i [Y(t) - Y^*] \quad (7b')$$

Expressed in this way, the rule asserts that the policy variable should be increased if output is above its target, and decreased otherwise. It is the form of policy adjustment rule specified by Mundell (1962) and others in their analysis of the assignment problem, relating the appropriate adjustment of policy instruments to targets.

Combining, equations (4), (6), and (7b') yields a system of three dynamic equations in $Y(t)$, $G^d(t)$, and $G(t)$:

$$\begin{pmatrix} \dot{Y} \\ \dot{G}^d \\ \dot{G} \end{pmatrix} = \begin{pmatrix} \alpha(c-1) & 0 & \alpha \\ -\gamma_i & 0 & 0 \\ 0 & \beta & -\beta \end{pmatrix} \begin{pmatrix} Y \\ G^d \\ G \end{pmatrix} + \begin{pmatrix} \alpha I \\ \gamma_i Y^* \\ 0 \end{pmatrix} \quad (10)$$

This yields several differences from the proportional rule. Assuming $0 < c < 1$,

$$[\beta + \alpha(1-c)](1-c) > \gamma_i$$

is a necessary and sufficient condition for (10) to be stable. This indicates a tradeoff between the intensity of the stabilization policy and lags in policy for stability to prevail. If the policy lags are sufficiently long (β small), it is possible for overly intensive policy to generate instability. Indeed, this was one of the concerns originally expressed by Friedman (1948). In the absence of policy lags ($\beta \rightarrow \infty$) the integral policy will always ensure stability, although it may be associated with cyclical adjustment if the adjustment is too intensive ($4\gamma > \alpha(1-c)^2$). In any event, if stable, Y will converge to Y^* , thereby avoiding one of the undesirable features associated with Phillips' specification of the proportionate rule.

3.1.3 Derivative Policy

The third policy rule introduced by Phillips, the derivative policy, is of the form

$$G^d(t) = -\gamma_d \dot{Y}(t) \quad \gamma_d > 0 \quad (7c)$$

That is, fiscal stimulus depends upon the current rate of change of output, behaving like a “negative accelerator”. For an appropriately chosen γ_d this can stabilize an otherwise unstable system, although it will not succeed in driving income to its target level.

Phillips also proposed combining these three policy rules, by postulating e.g.

$$G^d(t) = -\gamma_p (Y(t) - Y^*) - \gamma_i \int_{-\infty}^t [Y(s) - Y^*] ds - \gamma_d \dot{Y}(t) \quad (11)$$

showing how by the judicious choice of weights γ_j the policymaker can take advantage of the various desirable features of the individual policies, while reducing their unattractive aspects. For example, the presence of the integral component ensures that income converges to its target, while at the same time undesired cyclical adjustments associated with this policy can be reduced with the simultaneous use of the proportional and derivative policies. In this respect it is intriguing to observe that combining the policies as in (11) is a step in the direction of choosing the optimal

stabilization policy. Finally, we again emphasize that Phillips also introduced these policies into more complex models that include an accelerator determined investment demand, leading to higher order dynamic systems.

While Phillips developed these rules in the context of fiscal stabilization policy, early applications of stabilization theory also applied them to monetary stabilization issues; see e.g. Lovell and Prescott (1968), Sargent (1971). They also formed the basis for simulation studies involving both monetary and fiscal policies in larger macro models; see e.g. Cooper and Fischer (1974). Most contemporary research on stabilization policy has focused on monetary policy, with the structure of fiscal policy being directed more toward longer-run issues pertaining to economic growth and capital accumulation.¹⁰

4. Linear-Quadratic Optimal Stabilization

The policy rules introduced by Phillips were postulated on the grounds of their plausibility. They are not in general optimal, although, as we shall see, they appear as components of an optimal policy.

4.1 General approach

Beginning in the 1960's, interest developed in the question of optimal stabilization policy. The framework employed to address this issue was the linear-quadratic system, an adaptation of the "state-regulator problem" developed by control engineers; see e.g. Kalman (1960), Athans and Falb (1966), Bryson and Ho (1969). In general, this can be outlined as follows.

Consider an economy summarized by n state (target) variables, x , and m control (policy) variables, u . Assume that the structure of the economy can be expressed by the linear vector system:

$$\dot{x}(t) = A(t)x(t) + B(t)u(t) \tag{12a}$$

where $A(t)$ is an $n \times n$ matrix and $B(t)$ is an $n \times m$ matrix.¹¹ Assume further that the objective is to

¹⁰ For example, the "endogenous growth" literature pioneered by Romer (1986) and its extensions emphasizes the impact of tax rates and the role of public capital on growth. There is much less focus on monetary policy.

¹¹ A system of an arbitrary order can always be reduced to a first order system by redefining the higher-order derivatives as new state variables, so in this respect (12a) is a general representation of a linear system.

minimize the quadratic cost function

$$J \equiv \frac{1}{2} x'(T)Fx(T) + \frac{1}{2} \int_0^T [x'(t)M(t)x(t) + u'(t)N(t)u(t)]dt \quad (12b)$$

where F , $M(t)$ are positive semi-definite matrices, $N(t)$ is positive definite, T is the planning horizon, and primes denote the vector transpose operator.

In economic terms, the policymaker wishes to keep a set of target variables, and the corresponding values of the policy variables, as close as possible to their desired target values, with deviations from objectives being penalized by quadratic costs. The optimal (cost-minimizing) value of the control vector, $\hat{u}(t)$ is a linear feedback rule of the form

$$\hat{u}(t) = -N^{-1}(t)B'(t)P(t)x(t) \quad (13a)$$

where $P(t)$ is the unique positive semi-definite solution to the Riccati equation:

$$\dot{P}(t) = -M(t) - A'(t)P(t) - P(t)A(t) + P(t)B(t)N^{-1}(t)B'(t)P(t) \quad (13b)$$

and satisfies the boundary condition

$$P(T) = F \quad (13c)$$

The critical thing to note about this solution is that the optimal policy is a time varying linear feedback control law, in which, in general, all of the control variables are linear functions of all of the current state variables.

Several observations about this form of solution can be noted. First, by simple re-definition of variables it can be easily adapted to incorporate exponential time discounting, as economic applications typically employ. Second, while the quadratic function has the convenience of yielding linear optimal control laws it implies that positive and negative deviations from targets are weighted equally, which may or may not be appropriate. Third, if the time horizon is infinite and the matrices, A, B, M, N defining the optimal stabilization problems are constant over time, then the optimal policy summarized by (13) simplifies to the stationary rule

$$\hat{u}(t) = -N^{-1}B'Px(t) \quad (14a)$$

where P is the unique positive semi-definite solution to the matrix equation:

$$M + A'P + PA - PBN^{-1}B'P = 0 \quad (14b)$$

By substitution it then follows from (13a), that when policy is set optimally, the economy evolves in accordance with

$$\dot{x}(t) = [A - BN^{-1}B'P]x(t) \quad (15)$$

The policy rules (13a) [or (14a)] can be characterized as being a kind of “generalized proportional” policy of the type proposed by Phillips, in the sense that the current policy variables are related proportionately to the current state of the economy relative to its long-run target. Unlike Phillips, the elements of the feedback rule, as described by (13a) or (14a), are given by specific values, which may be required to follow specific time paths, depending upon whether or not the economic structure is constant or time-varying. By construction, (15) will ensure that the economy converges to its desired target value; problems of instability, which we saw could be associated with the inappropriate setting of the integral policy rule, do not arise.

One final point is that much of the interest in stabilization policy relates to stochastic systems. Among the earliest treatments was Howrey (1967) who extended an earlier discrete-time multiplier-accelerator model formulated by Baumol (1961), to allow for additive stochastic disturbances. But what if the parameters themselves in the basic structural equations such as (12a) are stochastic? Important work by Wonham (1963, 1968, 1969) showed that the optimal policy rules of the form summarized by (13) and (14) above extend to this case, where the feedback rules are shown to depend upon the variance-covariance matrix of the underlying stochastic parameters.¹² Early papers by Turnovsky (1973, 1976, 1977b) then applied these results to issues in macroeconomic stabilization policy, of the type pioneered by Phillips.

¹² There are, however, constraints on the variance-covariance matrix of the underlying stochastic parameters which in effect assert that control is possible only if the stochastic components are not too large. For a discussion of this stochastic stabilizability condition in the context of the conventional aggregate macroeconomic model, see Turnovsky (1977b).

Recently, Kendrick (2005) has provided an excellent overview of the applications of stochastic control methods to the class of linear-quadratic economic models outlined in this section. With the introduction of stochastic elements, the availability of information at the time a policy must be implemented becomes important. In particular, it becomes necessary to distinguish between open-loop control, when the entire time path for policy is solved at the outset, and feedback rules, when current policy is updated as new information becomes available. In the case of deterministic systems which are fully known, the two solutions coincide. There is no gain from feedback, although the open-loop solution can be expressed in feedback form as in (13), and as Phillips himself did. An important aspect of this distinction involves learning, an issue to which we return in Section 7.3 below.

4.2 Some special cases

To give a sense of how these early applications of optimal stabilization policy relate to Phillips' contributions, we consider several special cases.

First, suppose that the economy is purely one-dimensional, such as in equation (4) above, being expressed by

$$\dot{y}(t) = \alpha [-sy(t) + g(t)] \quad (16a)$$

where $y(t)$ and $g(t)$ denote output and government expenditure, both measured about their respective target values, and $s = 1 - c$ is the marginal propensity to save. Investment is constant, equal to its desired value and there are no policy lags. The policymaker's objective is to choose the fiscal instrument, $g(t)$, so as to minimize the quadratic cost function

$$J \equiv \frac{1}{2} \int_0^{\infty} [my(t)^2 + ng(t)^2] dt \quad (16b)$$

subject to the evolution (16a). The optimal policy for this problem is the linear feedback rule

$$g(t) = -\frac{\alpha}{n} py(t) \quad (17a)$$

where p is the positive solution to the quadratic equation

$$\frac{\alpha^2}{n} p^2 + 2\alpha s p - m = 0 \quad (17b)$$

The optimal policy summarized by (17) is a purely proportional one, as proposed by Phillips, though of the modified form (7a') form. In terms of his notation $\gamma_p > 0$, assuming the specific values implied by (17a), (17b), which in turn depend upon the underlying structural parameters.

Turnovsky (1973) extended this to the case where s or α were stochastic. In the former case, for example, (17b) is modified to

$$\frac{\alpha^2}{n} p^2 + (2\alpha\bar{s} + \alpha^2\sigma_s^2)p - m = 0 \quad (17b)$$

where \bar{s} , σ_s^2 denote the mean and variance of the stochastic marginal propensity to save. More stochastic variation in the savings propensity implies more intensive fiscal intervention.

More generally, suppose now that the economy is described by the following multiplier-accelerator model, with all variables expressed in deviation form about their steady-state values

$$\dot{y}(t) = \alpha[z(t) - y(t)] \quad (18a)$$

$$z(t) = cy(t) + i(t) + g(t) \quad (18b)$$

$$\frac{di(t)}{dt} = \zeta[v\dot{y}(t) - i(t)] \quad (18c)$$

where $z(t)$ is aggregate demand and now investment, $i(t)$, is expressed as a lagged accelerator, where the desired stock of capital is proportional to output.¹³ The system is now driven by two state variables, $y(t), i(t)$, and so optimal fiscal policy will be of the generic form

$$g(t) = \theta_1 y(t) + \theta_2 i(t) \quad (19)$$

where the feedback components, θ_1, θ_2 , are computed from (14a) and (14b) [see Turnovsky, 1973].

¹³ Equation (18c) may be derived as follows. Suppose the desired stock of capital is proportional to output, $k^d(t) = vy(t)$, and that the actual capital stock adjusts gradually to its desired value in accordance with $\dot{k}(t) = \zeta[k^d(t) - k(t)]$. Combining these two equations with the relationship $\dot{k}(t) = i(t)$ yields (18c).

Solving (18c), for $i(t)$ and substituting, we can express (19) in the form

$$g(t) = (\theta_1 + \theta_2 \zeta v) y(t) - \theta_2 \zeta^2 v \int_{-\infty}^t y(s) e^{-\zeta(t-s)} ds \quad (19')$$

This can be seen to be the sum of Phillips' proportional policy and a form of integral policy, where past outputs (or its deviations) have exponentially declining weights. Furthermore, if the dynamics can be represented by a second order differential equation (as can easily be done in some variants of the Samuelson-Hicks model of the business cycle), the optimal policy can be written as the sum of a purely proportional plus a derivative component; see Turnovsky (1977a)

As a final example, we go beyond Phillips' early work and introduce a form of the Phillips curve, augmented by "backward-looking expectations". Consider the simple monetary model:

$$\dot{y}(t) = \alpha[z(t) - y(t)] \quad (20a)$$

$$z(t) = d_1 y(t) - d_2 [r(t) - \pi(t)] \quad (20b)$$

$$p(t) = \kappa y(t) + \pi(t) \quad \kappa > 0 \quad (20c)$$

$$\dot{\pi}(t) = \rho[p(t) - \pi(t)] \quad \rho > 0 \quad (20d)$$

For this modified structure, (20b) is an IS curve where aggregate demand positively upon output and negatively upon the real interest rate, $r(t) - \pi(t)$. Equation (20c) is an expectations-augmented Phillips curve, where the current rate of inflation, $p(t)$, increases with output expected inflation, where the coefficient on expected inflation, $\pi(t)$, is unity.¹⁴ Inflationary expectations evolve in accordance with the backward-looking adaptive expectations scheme, (20d). By substitution, this economy can be reduced to the pair of dynamic equations

$$\dot{y}(t) = \alpha((d_1 - 1)y(t) - d_2[r(t) - \pi(t)]) \quad (21a)$$

$$\dot{\pi}(t) = \rho \kappa y(t) \quad (21b)$$

¹⁴ Much of the early empirical work on the Phillips curve was concerned with whether or not this coefficient is unity, an issue that has bearing on the existence or otherwise of a long-run unemployment-inflation tradeoff; see Turnovsky (1977a) for a discussion of this issue. Despite this early debate, (20c) is a consensus canonical specification of the expectations-augmented Phillips curve.

Assume that the policymaker sets the nominal interest rate, $r(t)$, to minimize quadratic costs associated with deviations of output and inflation from their respective target values. The optimal monetary policy will be of the generalized proportional form

$$r(t) = \varphi_1 y(t) + \varphi_2 \pi(t) \tag{22}$$

which is essentially a form of the widely-discussed Taylor (1993) rule.¹⁵

Other examples can also be found, but we have surely made the point that the form of policy rules proposed by Phillips (1954) played a central role in the early applications of optimal control theory to stabilization policy.

5. The Challenge of Rational Expectations

The dynamic system considered by Phillips, as well as the early applications of dynamic control theory that we have been discussing, [including the last example of the expectations-augmented Phillips curve] are of the classical type, in that all variables are assumed to evolve continuously from some given initial condition. In the jargon of contemporary macrodynamics, all variables are “backward-looking” or “sluggish”. This reflects the fact that economists were using the traditional techniques of differential equations as developed by applied mathematicians and control engineers, which of course was consistent with Phillips’ own academic background.

However, the development of rational expectations and its application to macrodynamics in the 1970’s introduced the notion that some economic variables, most notably financial variables, are “forward-looking”, incorporating agents’ expectations of the future. It is realistic to permit these variables to respond instantaneously to new information as it impinges on the economy, instead of forcing them to evolve gradually from the past. This was first illustrated in a simple monetary model by Sargent and Wallace (1973). They showed how, given the inherent instability of the underlying differential equation driving the dynamics in this model, plausible economic behavior requires that the forward-looking variable (in their case the price level) to jump so as to ensure that the economy

¹⁵ Taylor rules are feedback rules that tie the current interest rate to deviations in expected inflation and output, about their desired target levels. Taylor proposed the specific coefficients of 0.5 on the output variable and 1.5 on the inflation deviation. Turnovsky (1981) has analyzed in detail the optimal tradeoffs between unemployment and inflation in an expanded version of this model.

follows a bounded (stable) adjustment path. Most economic dynamic systems consist of a combination of sluggish variables, such as physical capital, which by their nature can be accumulated only gradually, and forward looking jump variables, such as exchange rates or financial variables, that are not so constrained. As a consequence, the standard dynamic macroeconomic system embodying rational expectations has a combination of stable and unstable dynamics, with the case of a unique convergent saddlepath arising when the number of unstable roots equals the number of jump variables; see Blanchard and Kahn (1980), Buiter (1984).

This represents a fundamentally different approach to macroeconomic dynamics from the earlier literature and the introduction of rational expectations has had a profound effect on the application of control methods to stabilization policy. Several issues arise and we shall discuss these in turn.

5.1 Computation of Optimal Policy Rules under Rational Expectations

The first issue is the task of solving for rational expectations equilibrium, even in the absence of any active stabilization policy. While the equilibrium economic structure may be conceptually straightforward, its solution is likely to be computationally challenging, depending upon the dating of the forward-looking expectations variables, their forecast horizons, and the dimensionality of the system. Solutions procedures have been proposed by several authors, including Blanchard and Kahn (1980), Fair and Taylor (1983), Buiter (1984), and more recently Sims (2001).

Currie and Levine (1985) provide a lucid description of the computation of optimal feedback rules for a continuous-time formulation containing both sluggish and forward-looking variables. They show how, in a system embodying rational expectations, one can partition the dynamic variables into predetermined (sluggish) variables and non-predetermined (forward-looking) variables, while taking account of the saddlepoint structure associated with the rational expectations equilibrium. For the usual quadratic loss function, the resulting optimal policy rule can be expressed in several alternative but equivalent forms. One form is as a linear function of both the predetermined and non-predetermined variables. Alternatively, it can be expressed as a linear function of the pre-determined state variables and the pre-determined co-state variables (those

associated with the non-predetermined variables). However, since the latter can be expressed as an integral of the vector of the underlying state variables, the optimal policy rule can be expressed as a generalized linear feedback rule on the state variables combined with an integral feedback rule on the state variables. To this extent the form of the Phillips' policies still prevail. Further details are provided in the Appendix.

5.2 The Lucas Critique

As we have been discussing, the objective of stabilization policy is to influence the time paths of a set of target variables, such as output, inflation, etc. Being forward-looking, a key feature of rational expectations, in contrast to the traditional adaptive expectations scheme such as (20d), is that it incorporates the agent's information regarding the structure of the economy. In particular, rational expectations will include the agent's perception of policy as part of the economic environment. Lucas (1976) made the profound observation that, in these circumstances, for policymakers to conduct policy under the assumption that the coefficients describing the evolution of the economy remain fixed and invariant with respect to its chosen policy is not rational. In the dynamic system, (8), for example, the behavioral parameters, α, c , will in general vary with the chosen policy parameter, γ_p . This dependence needs to be taken into account in determining the effects of policy rules, as well as for the determination of optimal stabilization policy.¹⁶

The Lucas Critique is a general proposition having far-ranging implications for analyzing economic policy. Its main message is that if we want to predict the effects of policy changes we must model the "underlying parameters" such as technology and preferences that govern individual behavior. To the extent that modern macroeconomics is based on intertemporal optimization of utility subject to production constraints, the macroeconomic equilibrium so derived is immune from the Lucas Critique in that it is conditional on government policy. We can then model policymaking as a game, whereby the government, acting as leader, makes its stabilization (policy) decisions

¹⁶ As a related observation, the Lucas Critique calls into question the practice of econometrically estimating the parameters of a reduced form equation such as (8). This is because as the policy varies, so do the structural parameter, and consequently the assumption that they remain fixed over a sample period is inappropriate. We should also note that the Lucas Critique does not apply to dynamic systems such as the original Phillips models, which are entirely backward-looking.

taking into account the reactions of the private sector. This approach is at the core of the voluminous optimal tax literature.

However, solving for optimal policy in this way may be difficult, particularly over time, and it furthermore, it may be unrealistic to assume that the policymaker knows precisely the private sector's response to its decisions. Amman and Kendrick (2003) propose an approximation based on the use of the Kalman filter. The idea is that the policymaker need not be able to predict exactly how private agents will respond to its policies. Rather, it can simply use the Kalman filter and update parameter estimates each period. While this means that the policymaker will always be one period behind, in his perception of the private sector's behavioral responses, they argue that this may be good enough for most applications of macroeconomic policy. Monte Carlo runs they run provide some support this view.

5.3 Policy Neutrality

One area where the Lucas Critique is particularly potent is in the role of the Phillips curve in determining the tradeoffs between inflation and unemployment. In this context the Lucas critique says that the nature of the tradeoff depends upon government policy. The issue of policy neutrality is an extreme form of this, and asserts that, because of the Lucas Critique, the tradeoff breaks down completely.

In an influential article, Sargent and Wallace (1976) provided an example to show that under rational expectations only unanticipated policy changes can have real effects, so that any feedback policy rule, such as the Phillips rules we have been discussing, will have no effect on output. In our example, the time path of output would become independent of the policy parameters such as γ_p , so that there is no tradeoff between output and inflation. It turns out that this policy neutrality proposition, as it is known, and which is potentially most damaging to the use of control theory as a tool of stabilization policy, is sensitive to model specification, and in particular to the timing of expectations. This can be illustrated by comparing two simple examples.

First, consider an economy represented by the pair of stochastic difference equations

$$y_t = -e[r_t - (E_{t-1}(P_{t+1}) - E_{t-1}(P_t))] + u_t \quad (23a)$$

$$P_t - P_{t-1} = \theta y_t + E_{t-1}(P_t) - P_{t-1} + v_t \quad (23b)$$

where y_t denotes output (in deviation form) in logarithms, P_t denotes the price level in logarithms, $E_{t-1}(\cdot)$ denotes expectations, formed at time $t-1$ and assumed to be rational, r_t is the nominal interest rate, and u_t, v_t are white noise random disturbances in demand and supply, respectively. In keeping with the contemporary literature we employ discrete time. Equation (23a) is a standard IS curve, relating output negatively to the real interest rate, where the expected rate of inflation over the period $(t, t+1)$ is based on information at time $t-1$. The second equation is the New-Classical Phillips curve.¹⁷ Assuming that the monetary authority treats the nominal interest rate r_t as the policy variable, the equilibrium value of output y_t can be shown to be

$$y_t = -e[r_t - E_{t-1}(r_t)] + u_t \quad (24)$$

The point about (24) is that (the deviation of) output depends only upon the unanticipated component of monetary policy. Any feedback policy rule based on past observed data that the monetary authority follows is fully incorporated into private agents' expectations and thus is fully negated in terms of its effects on current output.

Things change dramatically, however, if we now modify (23a) to

$$y_t = -e[r_t - (E_t(P_{t+1}) - P_t)] + u_t \quad (23a')$$

The only difference is that we modify expected inflation for period $(t, t+1)$ to be conditional on information at time t , when the actual price level is observed. The rational expectations solution for output under this seemingly modest reformulation is

$$y_t = -\frac{e}{(1+e\theta)}[r_t - E_{t-1}(r_t)] - \frac{e}{(1+e\theta)} \sum_{i=1}^{\infty} [E_t(r_{t+i}) - E_{t-1}(r_{t+i})] + \frac{(u_t - ev_t)}{(1+e\theta)} \quad (24')$$

In addition to the current unanticipated component of monetary policy, given by the first term of (24'), current output now depends upon the sum of all revisions to future monetary policy between

¹⁷ Lucas adopted what has become known as the "Lucas supply function", which replaces (23b) by $y(t) = \lambda[P_t - E_{t-1}(P)] + v'_t$, so output deviations depend upon unanticipated price movements.

time $(t-1)$ and t , which takes account of new information acquired at time t . By impacting the forecast of inflation a feedback policy rule will now exert an impact on current output.

As a simple example, suppose that the monetary authority sets the interest rate in accordance with the rule

$$r_t = \lambda \left(\frac{u_{t-1} - e v_{t-1}}{1 + e\theta} \right)$$

This is feedback rule whereby the monetary authority adjusts the interest rate in response to last period's stochastic shocks, which are known at time t . With u_t, v_t being white noise, taking expected values over relevant periods (24') reduces to

$$y_t = \left(1 - \frac{\lambda e}{1 + \theta e} \right) \frac{u_t - e v_t}{1 + e\theta}$$

which clearly is influenced by the policy rule.¹⁸ Indeed, setting $\lambda = \theta + 1/e$ succeeds in stabilizing output completely. The presence of policy neutrality thus depends critically upon the available information, an issue that is taken up at greater length by Canzoneri, Henderson, and Rogoff (1983).

5.4 Time Consistency of Optimal Policy

The fourth issue to arise when the system contains forward-looking jump variables is that of "time consistency". This concept was first introduced into the economics literature by Strotz (1955-56). It describes a situation where an agent's preferences change over time, so that what is preferred at one instant of time is no longer preferred at some other later point in time. This issue has ramifications for various aspects of economics and in particular for stabilization policy. In this context the issue is whether or not a future policy decision that forms part of an optimal plan formulated at some initial date remains optimal when considered at some later date, even though no relevant information has changed in the meantime. If it is not, the optimal plan is said to be time inconsistent.

¹⁸ We should point out that setting the nominal interest rate, as in this example, leads to an indeterminate price level, an issue that has generated some debate, particularly in the context of the so-called monetary instrument problem; see Poole (1970), Parkin (1978), Turnovsky (1980), and McCallum (1981). This aspect does not invalidate the point we wish to make, and in any event can be easily rectified by introducing real money balances into the aggregate demand function.

This problem was emphasized by Kydland and Prescott (1977) who argued that the problem of time inconsistency has grave implications for the application of control theory methods to problems of economic stabilization. In the abstract to their paper they write “...We conclude that there is no way control theory can be made applicable to economic planning when expectations are rational.” In the conclusions they argue “...active stabilization may very well be dangerous and it is best that it not be attempted. Reliance on policies such as a constant growth in the money supply and constant tax rates constitute a safer course of action.” These are strong statements and many people in the community of control theorists view this assessment of the role of control theory to stabilization policy as overly pessimistic.

But the question of time consistency (or inconsistency) is important, and attempts to resolve it have generated a lot of research. The pursuit of time inconsistent policies will eventually cause the government to lose credibility and issues such as commitment and reputational equilibria have been analyzed by a number of authors; see e.g. Barro and Gordon (1983), Backus and Driffill (1985).

One simple solution, very much within the spirit of the linear-quadratic framework, is the following. As noted, the attainment of a rational expectations equilibrium involves an initial jump in the forward-looking variable. These initial jumps presumably impose real dislocational costs on the economy, and these should be taken into account in the design of the optimal policy system. Stemp and Turnovsky (1987) show how if these initial costs are large enough that it may cease to be optimal for the policymaker to re-optimize along a transitional path.

To see this assume that the policymaker’s objective function is to

$$\text{Minimize} \quad \int_0^{\infty} [ay^2 + (1-a)p^2]e^{-\beta t} dt + k|P(0) - P_0|^q$$

This cost function now has two components. The first is the standard quadratic loss function, asserting that the policymaker’s target is to achieve a zero rate of inflation ($p=0$) at a full employment level of output ($y=0$). One objective is to minimize the discounted intertemporal deviations from these targets, with $0 \leq a \leq 1$ reflecting the relative importance assigned to these two objectives.

In a world of rational expectations, a change in monetary policy will cause an initial

unanticipated discontinuous jump in the price level, $P(0)$, from its previously inherited level, P_0 . Given sluggishness in the economy this causes jumps in real magnitudes, which impose structural adjustment costs [e.g. labor reallocation] on the economy and these should be taken into account in assessing the overall benefits of the optimal stabilization policy. These initial adjustment costs are not reflected in the conventional term, but are incorporated in the second term. Stemp and Turnovsky show that the time consistency or otherwise depends critically upon q , being time consistent if $q \leq 1$ and time inconsistent otherwise.

6. Rules versus Optimal Policy

Despite the fact that the generic form of the optimal policy rule is the generalized proportional policy as set out in (13), from a practical point of view the policy may turn out to be extremely complicated to compute, especially for a large system, and even more so if it includes forward-looking variables. This leads to the question of the gains from applying optimal control over using some simple, but reasonable, policy such as the three rules proposed by Phillips, or the Taylor rule, or perhaps even doing nothing at all.

This is an old question, predating Phillips, going back to Friedman (1948) and early discussions of policy rules versus discretionary policy. At that time Friedman advanced the proposition that due to the length and variability of lags in the effects of monetary policy, the monetary authority should abandon discretionary monetary management and simply should allow the money supply to grow at a fixed rate. Indeed, our discussion of the Phillips rules provides some support for this view. We have seen that the presence of policy lags can introduce unwanted cycles in the economy, and even instability that otherwise would not exist. But the debate of rules versus discretion raises several issues, most important of which relate to the information that the policymaker possesses. Here we briefly note some of them.

First, suppose that the world is deterministic and the policymaker knows the true structure. Then by its nature the optimal policy dominates and so the question is whether the gains in economic stability are sufficient to justify the effort involved in computing the optimal rule over something much simpler and mechanistic. Some years ago Feldstein and Stock (1994) addressed this question

in an analysis where the objective is to target nominal income. They reached the conclusion that there is little difference between a very simple adaptive rule and an optimal policy. If this kind of proposition is robust, then simple policy rules of the type originally proposed by Phillips will continue to play an important role in the stabilization of dynamic economic systems.

Second, what if there is uncertainty? This introduces different levels of complications. The optimal policy model introduced in Section 4 assumed that everything is known except for the fact that some of the parameters describing the economy are stochastic. As we have seen this will influence the optimal setting of the associated policy instrument, just as it did in Brainard's (1967) early analysis. But one of the important results obtained by Wonham (1969) is that feedback control in a system with stochastic parameters, whereby the effects of policy become stochastic, is feasible if and only if the noise is not too great. In this case, it is possible for the policy instrument used for stabilization to introduce too much noise into the system implying that the economy will actually be more stable if the policymaker does not intervene.

Third, and most fundamentally, all optimal policies we derive are specific to an assumed economic structure, rather than the true economic environment that policymakers do not and cannot know. What are the merits of employing the optimal policy to the wrong model, rather than some arbitrary alternative rule? This issue is addressed in detail by Brock, Durlauf, Nason, and Rondina (2007). Their approach is to construct a model space that includes all candidate models for the economy, evaluate the policies for each of the candidate models, and then determine how to draw policy inferences given the fact that the true model is unknown. In contrast to the usual robustness analysis that measures misspecification relative to some baseline model, they acknowledge the global nature of model uncertainty. They focus on the sensitivity of the rules to model uncertainty, rather than on the derivation of optimal rules in the presence of model uncertainty. The other issue they address concerns the tradeoffs of policy on variances of different frequencies; policies that may reduce the variance of high frequency fluctuations may come at the expense of enhanced longer-term fluctuations.

7. Recent Developments

In this section we briefly note some of the more recent aspects of stabilization policy that pertain to Phillips' contribution.

7.1 New-Keynesian Phillips Curve

The New-Keynesian Phillips curve is based on a model of optimal pricing in imperfect competition and a theory of price stickiness; see e.g. Roberts (1995), McCallum and Nelson (1999), and Woodford (2003). It is of the generic form

$$P_t - P_{t-1} = \theta y_t + \beta (E_t(P_{t+1}) - P_t) \quad 0 < \beta < 1 \quad (25)$$

It differs from the New-Classical Phillips curve in that the expected inflation to which the current inflation is reacting extends for the next period $(t, t+1)$, rather than for the previous period $(t-1, t)$. This has important consequences for stabilization policy. To see this, we shall combine (25) with (23a), for which the New-Classical Phillips curve yields policy neutrality.

The form of the rational expectations solution depends upon the magnitude of $\beta + \theta e$. We consider the case $\beta + \theta e < 1$ when the unique stable solution for y_t is

$$y_t = -er_t - e^2\theta \sum_{j=1}^{\infty} (\beta + \theta e)^{j-1} E_{t-1}(r_{t+j}) + u_t \quad (26)$$

It is clear that interest rate rules based on past information will influence current output. For example, if $r_t = \mu u_{t-1}$, then the solution to (26) is

$$y_t = -e\mu u_{t-1} + u_t$$

which is clearly dependent upon the policy parameter μ . The case $\beta + \theta e > 1$ is associated with non-uniqueness issues of the type identified in rational expectations models by Taylor (1977), but depending upon how the non-uniqueness is resolved, policy rules will have real effects.

7.2 Multi-agent Stabilization

Thus far we have focused on a single decision-maker, acting in isolation. In reality, many economic situations are characterized by multiple decisionmakers operating in an interactive

environment. The decisions made by one agent influence the other, and vice versa, giving rise to strategic behavior that we can analyze as a dynamic game. As is well known crucial factors determining the equilibrium outcome include (i) the availability of information at the time decisions are made, (ii) the sequencing of the decisions by the agents, and (iii) whether they behave non-cooperatively to maximize their own individual welfare, or cooperatively to maximize their joint well-being.

Insofar as stabilization policy is concerned there are two main areas where strategic interaction is particularly important. The first is the interaction between monetary and fiscal policy, allowing for the fact that the central bank may have different objectives from the Treasury. This gets into issues relating to credibility of policy, reputation, and political aspects that are somewhat removed from the approach to stabilization that we are discussing here; see e.g. Persson and Tabellini (1999).

The second application is in the area of international economic policy coordination, and in particular monetary and exchange rate policy. Miller and Salmon (1985) and Oudiz and Sachs (1985) have analyzed two country dynamic games of monetary policy, which are direct generalizations of the class of optimal policy model summarized in Section 4. To give a flavor of this suppose the policy maker of Country 1 wishes to solve the following dynamic optimization problem

$$\min \int_0^{\infty} w'(t)Q_1w(t)dt \quad \text{where } w'(t) \equiv [x(t) \quad u_1(t) \quad u_2(t)]' \quad (27a)$$

subject to the dynamic evolution

$$\dot{x}(t) = Ax(t) + B_1u_1(t) + B_2u_2(t) \quad (27b)$$

for given $x(0)$. Policy maker in Country 2 solves an analogous problem.

As in previous examples the objective is to minimize a quadratic loss function, which depends upon Country 1's controls, $u_1(t)$, which of course this policymaker sets, Country 2's controls, $u_2(t)$, which he may react to, and a common set of target variables, $x(t)$. The latter may be more relevant to one country more than the other and may also include non-predetermined

variables as well as pure state variables. Equation (27b) describes the evolution of the state variables, which depend in part upon the choices each policymaker makes.

For this setup Miller and Salmon discuss open loop and feedback Nash and Stackelberg solutions. For feedback Nash, for example, each policymaker sets his controls in accordance with the feedback rule

$$u_1(t) = R_1 x(t), u_2(t) = R_2 x(t) \quad (28)$$

taking the other's as given when making his decision. The components of the feedback are determined by a generalized Riccati type equation which involves the structural parameters of both economies. It will also depend upon the specific rule defining the dynamic game. But the point we wish to make is that (28) are generalized proportional policy rules directly analogous to (14a) discussed earlier and hence the relationship to Phillips' early work extends to this type of analysis.¹⁹

7.3 Utility Maximization

The optimal stabilization rules we have derived have been chosen so as to minimize quadratic costs involving the deviations of the state variables and the control (policy) variables about some stationary level. Many variants of this criterion can be found, varying in such aspects as to whether the deviations in output are measured relative to the full employment level, the frictionless level of output, etc.

Apart from the limitation noted earlier that the quadratic function is weighting positive and negative deviations equally, it suffers from the more serious criticism that it may or may not be an appropriate representation of welfare, which presumably is the issue of ultimate concern. Indeed, for almost three decades now, the "representative agent model" has been the standard macroeconomics paradigm, although it too has been the source of criticisms.²⁰ With macroeconomic equilibrium being derived through utility maximization this framework is much more oriented toward analyzing welfare issues and therefore addressing issues pertaining to optimal policy making.

¹⁹ For an overview of the literature on linear quadratic differential games, see Engwerda (2005).

²⁰ See Colander (2006).

Recently, several authors have sought to examine the relationship between utility maximization and the conventional stabilization criteria that we have been adopting; see in particular Woodford (2003) where this is discussed in great detail. There he establishes conditions under which the quadratic loss function, so widely employed in stabilization policy, can be viewed as a second-order approximation to the expected value of a more general utility function. Here we informally sketch the relationship in a simple example.

Suppose welfare is represented by a utility function of the form $U(c, g)$, where c denotes consumption and g denotes government expenditure (the control variable). Suppose further that through stabilization c and g are restricted to stochastic fluctuations about their respective mean levels (\tilde{c}, \tilde{g}) . Employing a second order approximation to $U(c, g)$ about (\tilde{c}, \tilde{g}) , and taking expected values, we may write

$$EU(c, g) \cong U(\tilde{c}, \tilde{g}) + \frac{1}{2}U_{cc}E(c - \tilde{c})^2 + U_{cg}E(c - \tilde{c})(g - \tilde{g}) + \frac{1}{2}U_{gg}E(g - \tilde{g})^2 \quad (29)$$

Assume that output is produced by the production function $y = f(k)$, where k denotes capital stock. If the agent maximizes intertemporal utility, it is well known that equilibrium consumption along an evolving stable adjustment path is of the form $c = c(k, g)$, which may be linearly approximated by

$$c - \tilde{c} \cong c_k(k - \tilde{k}) + c_g(g - \tilde{g})$$

Substituting this linear approximation into (29) yields a second-order approximation to expected utility of the form

$$EU(c, g) \cong U(\tilde{c}, \tilde{g}) + X$$

where X is a quadratic loss term involving the state variable k and the control variable g . For the simple production function, the state variable can be immediately transformed to y , as in the stabilization literature. In order for the quadratic loss function to give the correct welfare rankings of different stabilization policies it must be the case that $U(\tilde{c}, \tilde{g})$ is independent of policy, or at least is only weakly sensitive to it. One case where it is independent is if the utility function is of the form $U(c + g)$. With capital stock constant in steady state, product market equilibrium implies

$\tilde{c} + \tilde{g} = f(\tilde{k})$, where the steady-state stock of capital is fixed and determined by the marginal product condition $f'(\tilde{k}) = \rho$, the rate of time discount.

7.4 Learning

Throughout this discussion we have assumed that the policymaker has complete knowledge of the true underlying economic structure. In the case of deterministic systems all parameters are known, as is their evolution if they are time-varying. In the case of stochastic systems all characteristics of relevant probability distributions are also known; only the specific random outcomes are unknown until they occur. In reality, of course, policymakers do not know the true system. Even if they know the broad structure of the economy, such as the general qualitative relationship among the variables, they will at best have only some estimate of the relevant parameters, and worse still, they are unlikely to even know the general structure of the economy, as Brock et al. (2007) have emphasized. At best, policymakers and agents in general may learn about the structure of the economy as it evolves over time.

The qualitative information about the economic structure becomes particularly important in dynamic models involving rational expectations, the key characteristic of which is that they incorporate agents' perceived structure of the economy. As a result of this, their beliefs about the economy will influence its actual evolution. The fact that applications of rational expectations assume complete knowledge has been a source of its criticism. While this is a reasonable objection, we view the traditional rational expectations specification as a useful benchmark, with its underlying characteristic of forward-looking behavior providing significant insights into macroeconomic dynamics in general and stabilization policy in particular.

To incorporate learning is challenging and raises many issues. By its nature, learning is a gradual process that takes place over time. The interaction of the dynamics of this process with that of the system itself is important and not all learning processes need be stable. The most comprehensive general study of learning as an element of macroeconomic dynamics is Evans and Honkapohja (2001) which itself draws heavily from their past research.

They emphasize the method of expectational stability. The key element of this concept is

that it involves a mapping from the perceived law of motion (dynamic structure), which in general is incorrect, to the corresponding actual law of motion, which incorporates this incorrect information. If the system is expectationally stable, learning process for the unknown parameters will converge to the true values and the agent will ultimately learn the true economic process. It is possible however, for the learning process to diverge, and cause the dynamics of the overall system to diverge as well.

Several issues in this process arise and should be mentioned. First, the time period involved in updating information on parameters need not coincide with the time interval that characterizes the system dynamics. Second, it is possible for updating of information to involve nonlinear relationships, leading to a multiplicity of solutions, and for learning not converging to any of them; see Blanchard and Fischer (1989). Third, information and learning is almost certainly not uniform across the economy; different agents have different degrees of information and different capacities to learn. Evans and Honkapohja focus primarily on learning by private agents, but the same issues apply to policymakers engaged in optimal policy making. Fourth, learning may take different forms, the two most common being least squares learning and Bayesian learning.

The learning procedures we have been outlining can be characterized as being “passive”, in the sense that the agent learns about the relevant parameters over time as the system evolves and information is updated. Kendrick (2005) contrasts this with “active” learning sometimes referred to as “dual control”. In the stabilization policies we have been considering in previous sections, the policy instrument is used for a single purpose, namely to help move the economy toward its target. In contrast, in dual control the policy variable are used for two purposes. In addition to the usual stabilization objective, the second is to perturb the system so as to enhance learning of the relevant parameters and thereby improve control performance at later stages. This form of learning was introduced originally by Kendrick (1982) and later by Amman and Kendrick (1994), using techniques previously developed in the control literature by Tse and Bar-Shalom (1973).

8. Conclusions

It is evident that Bill Phillips through his initial contributions to dynamic stabilization policy in conjunction with the Phillips curve has had a profound impact on the theory of economic policy.

First, the policy rules he proposed frequently lie in the class of optimal policies and thus serve as a useful benchmark, thereby assisting in the interpretation of more complex optimal policy rules. Indeed the relationship of the Phillips policy rules to the optimal applies, not only to traditional optimal policy making based on sluggish backward looking systems, but it also applies to systems involving forward looking expectations, as well as multi-agent strategic policymaking problems.

The Phillips curve has been a remarkably resilient concept and has remained a key component of the output-inflation tradeoffs that characterize stabilization policy. Beginning with original negative inflation-unemployment relationship, through the (backward-looking) expectations-augmented Phillips curve of the 1960s, to the New-classical Phillips curve of the 1970s and most recently the New Keynesian Phillips curves of the 1990s it has been a central component of short-run macrodynamic models for half a century.

Indeed, the implementation of rational expectations presented a serious challenge to the use of control theory as an instrument of macroeconomic stabilization policy. But it is fair to say that macroeconomists have accepted the challenge and that the methods of control theory are being applied more successfully than ever to dynamic macro models involving rational expectations. The economics profession owes a great debt to Bill Phillips for introducing them to these analytical tools.

Appendix

Optimal stabilization rule when dynamics include both backward looking and forward-looking variables

In this Appendix we briefly provide the details in this case. Our summary draws on the lucid exposition of Currie and Levine (1985). Consider the optimal stabilization problem

$$\text{Min} \int_0^{\infty} [y' My + u' Nu] dt \tag{A.1a}$$

subject to

$$\dot{y} = Ay + Bw \tag{A.1b}$$

where w is a vector of control variables and y , the vector of state variables are partitioned into pre-determined variables, z , and non-predetermined variables, x , namely $y' = [z \ x]'$. There are p costate variables partitioned correspondingly, $p' = [p_1 \ p_2]'$. For the predetermined variables, $z(0) = z_0$, with $p_1(0)$ being free, while for the non-predetermined variables $p_2(0) = 0$

The optimal rule for the deterministic control problem is of the form

$$u(t) = -N^{-1}B'p(t) \quad (\text{A.2a})$$

where

$$\begin{pmatrix} \dot{y}(t) \\ \dot{p}(t) \end{pmatrix} = \begin{pmatrix} A & -J \\ -M & -A' \end{pmatrix} \begin{pmatrix} y(t) \\ p(t) \end{pmatrix} \equiv H \begin{pmatrix} y(t) \\ p(t) \end{pmatrix} \quad (\text{A.2b})$$

and $J = -BR^{-1}B'$, together with the transversality conditions $\lim_{t \rightarrow \infty} p(t)y(t) = 0$, which in this case implies $\lim_{t \rightarrow \infty} p(t) = \lim_{t \rightarrow \infty} y(t) = 0$. If H has $2n$ distinct eigenvalues, n of these associated with the predetermined variables $[z \ p_2]'$ will be stable and n associated with the non-predetermined variables $[x \ p_1]'$ will be unstable. Rearranging and partitioning (A.2b) accordingly, we may write

$$\begin{pmatrix} \dot{z} \\ \dot{p}_2 \\ \dot{p}_1 \\ \dot{x} \end{pmatrix} = \begin{pmatrix} A_{11} & -J_{12} & -J_{11} & A_{12} \\ -M_{21} & -A'_{22} & -A'_{21} & -M_{22} \\ -Q_{11} & -A'_{12} & -A'_{11} & -M_{12} \\ M_{21} & -J_{22} & -J_{21} & A_{22} \end{pmatrix} \begin{pmatrix} z \\ p_2 \\ p_1 \\ x \end{pmatrix} \equiv H \begin{pmatrix} z \\ p_2 \\ p_1 \\ x \end{pmatrix} \equiv \begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix} \begin{pmatrix} z \\ p_2 \\ p_1 \\ x \end{pmatrix} \quad (\text{A.3})$$

We form the matrix of left eigenvectors of H , Q say, and order it so that the first n rows are the eigenvectors associated with the n stable eigenvectors. We then partition it so that

$$Q = \begin{pmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{pmatrix}$$

Stability then imposes the relationship on the non-predetermined variables

$$\begin{pmatrix} p_1 \\ x \end{pmatrix} = -Q_{22}^{-1}Q_{21} \begin{pmatrix} z \\ p_2 \end{pmatrix} \equiv -\begin{pmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{pmatrix} \begin{pmatrix} z \\ p_2 \end{pmatrix} \quad (\text{A.4})$$

The optimal feedback rule (A.2a) can thus be written

$$u(t) = -N^{-1}B'[-R_{11}z(t) - R_{12}p_2(t), p_2(t)]' \equiv D \begin{pmatrix} z(t) \\ p_2(t) \end{pmatrix} \quad (\text{A.5})$$

where $D \equiv -N^{-1}[-B_1R_{11}, B_2 - B_1R_{12}]$ and $B' \equiv [B_1 \ B_2]$. Thus we may write

$$\begin{pmatrix} \dot{z} \\ \dot{p}_2 \end{pmatrix} = [H_{11} - H_{12}R] \begin{pmatrix} z \\ p_2 \end{pmatrix} \equiv C \begin{pmatrix} z \\ p_2 \end{pmatrix} \quad (\text{A.6})$$

the solution to which is

$$\begin{pmatrix} \dot{z} \\ \dot{p}_2 \end{pmatrix} = e^{Ct} \begin{pmatrix} z_0 \\ 0 \end{pmatrix} \quad (\text{A.7})$$

From (A.4) we see that

$$x = -R_{21}z - R_{22}p_2$$

so that

$$p_2 = -R_{22}^{-1}(x + R_{21}z)$$

Substituting into (A.5) we obtain

$$w(t) = -N^{-1}B'S \begin{pmatrix} z(t) \\ p_2(t) \end{pmatrix} \quad (\text{A.8})$$

where $S_{11} \equiv -R_{11} + R_{12}R_{22}^{-1}R_{21}$; $S_{21} \equiv -R_{22}^{-1}R_{21}$; $S_{12} \equiv R_{12}R_{22}^{-1}$; $S_{22} \equiv -R_{22}^{-1}$

References

- Allen, R.G.D., (1956), *Mathematical Economics*, McMillan, London.
- Amman, H. and D.A. Kendrick (1994), "Active Learning: Monte Carlo Results," *Journal of Economic Dynamics and Control* 18, 119-124.
- Amman, H. and D. A. Kendrick (2003), "Mitigation of the Lucas Critique with Stochastic Control Methods", *Journal of Economic Dynamics and Control* 27, 2035-2057.
- Athans, M. and P. Falb, (1966), *Optimal Control*, McGraw-Hill, New York.
- Backus, D.K. and J. Driffill, (1985), "Inflation and Reputation," *American Economic Review* 75, , 530-538.
- Barro, R.J. and D.B. Gordon, (1983), "Rules, Discretion, and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics* 12, 101-121.
- Baumol, W. (1961), "Pitfalls in Contracyclical Policies: Some Tools and Results," *Review of Economics and Statistics* 43, 545-556.
- Blanchard, O. J. and C. M. Kahn (1980), "The Solution of Linear Difference Models under Rational Expectations," *Econometrica* 48, 1305-1311.
- Brainard, W. (1967), "Uncertainty and the Effectiveness of Policy," *American Economic Review, Proceedings* 57, 411-425.
- Brock, W.A., S.N. Durlauf, J.M. Nason, and G. Rondina, (2007), "Simple versus Optimal rules as Guides to Policy," *Journal of Monetary Economics* 54, 1372-1396.
- Bryson, A.E. and Y.C. Ho (1966), *Applied Optimal Control*, Blaisdell, Waltham MA
- Buiter, W.H., (1984), "Saddlepoint Problems in Continuous Time Rational Expectations Models: A General Method and Some Macroeconomic Examples," *Econometrica* 52, 665-680.
- Canzoneri, M.B., D.W. Henderson, and K. Rogoff (1983), "The Information component of the Interest Rate and Optimal Monetary Policy," *Quarterly Journal of Economics* 98, 545-565.
- Colander, D. ed., (2006), *Post-Walrasian Macroeconomics*, Cambridge University Press, Cambridge U.K.
- Cooper, J.P. and S. Fischer, (1974), "Monetary and Fiscal Policy in the Fully Stochastic St Louis Econometric Model," *Journal of Money, Credit, and Banking* 6, 1-22.

- Currie, D. and P. Levine, (1985), "Macroeconomic Policy Design in an Interdependent World," in W.H. Buiter and R.C. Marston (eds.) *International Economic Policy Coordination*, Cambridge University Press, Cambridge, U.K.
- Engwerda, J.C. (2005), *LQ Dynamic Optimization and Differential Games*, Wiley, Chichester.
- Evans, G.W. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton University, Princeton N.J.
- Fair, R.C. and J. B. Taylor (1983), "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models," *Econometrica* 51, 1169-1185.
- Feldstein, M. and J.H. Stock, (1994), "The Use of Monetary Aggregate to Target Nominal GDP," in N.G. Mankiw (ed.), *Studies in Business Cycles vol. 29*, University of Chicago Press, Chicago.
- Friedman, M., (1948), "A Monetary and Fiscal Framework for Economic Stability," *American Economic Review* 38, 1961, 245-264.
- Friedman, M., (1968), "The Role of Monetary Policy," *American Economic Review* 58, 1-17.
- Henderson, D.W. and S.J. Turnovsky (1972), "Optimal Macroeconomic Adjustment Under Conditions of Risk," *Journal of Economic Theory* 4, 58-71.
- Hicks, J.R., *A Contribution to the Theory of the Trade Cycle*, Oxford University Press, Oxford, 1950.
- Howrey, P., (1967), "Stabilization Policy in Linear Stochastic Models," *Review of Economics and Statistics* 49, 404-411.
- Kalman, R.E. (1960), "New Approach to Linear Filtering and Prediction Problems," *Transactions of the ASME-Journal of Basic Engineering* 82 (Series D), 35-45.
- Kendrick, D.A., (1982), "Caution and Probing in a Macroeconomic Model," *Journal of Economic Dynamics and Control* 4, 149-170.
- Kendrick, D.A. (2005), "Stochastic Control for Economic Models: Past, Present, and the Paths Ahead," *Journal of Economic Dynamic and Control* 29, 3-30.
- Kydland, F.E. and E.C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 1977, 473-491.

- Lovell, M.C. and E. Prescott, "Money Multiplier Accelerator Interaction," *Southern Economic Journal*, 35, 1968, 60-72.
- Lucas, R.E., "Econometric Policy Evaluation: A Critique," in *The Phillips Curve and Labor Markets*, K. Brunner and A.H. Meltzer (eds.) Carnegie Rochester Conference Series on Public Policy, North-Holland Amsterdam, 1976.
- McCallum, B.T., (1981), "Price Level Determinacy with an Interest Rate Policy Rule and Rational Expectations," *Journal of Monetary Economics* 8, 319-329.
- McCallum, B.T. and E. Nelson, (1999), "An Optimizing IS-LM specification of Monetary Policy and Business Cycle Analysis," *Journal of Money, Credit, and Banking* 31, 296-316.
- Miller, M.H. and M. Salmon, (1985), "Policy Coordination and Dynamic Games," in W.H. Buiter and R.C. Marston (eds.) *International Economic Policy Coordination*, Cambridge University Press, Cambridge, U.K.
- Mundell, R.A., "The Appropriate Use of Monetary and Fiscal Policy for Internal and External Stability," *IMF Staff Papers*, 9, 1962, 70-79.
- Oudiz, G. and J. Sachs, (1985), "International Policy Coordination in Dynamic Macroeconomic Models," in W.H. Buiter and R.C. Marston (eds.) *International Economic Policy Coordination*, Cambridge University Press, Cambridge, U.K.
- Parkin, J.M., (1978), "A Comparison of Alternative Techniques of Monetary Control under Rational Expectations," *Manchester School* 46, 252-287.
- Persson, T. and G. Tabellini, (1999), "Political Economics and Macroeconomic Policy," in J.B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics* vol. 1C, North-Holland, Amsterdam.
- Phelps, E.S., (1968), "Money-Wage Dynamics and Labor-Market Equilibrium," *Journal of Political Economy* 76, 678-711.
- Phillips, A.W., (1954), "Stabilisation Policy in a Closed Economy," *Economic Journal* 64, 290-323.
- Phillips, A.W., (1957), "Stabilisation Policy and the Time Form of Lagged Responses," *Economic Journal* 67, 265-277.
- Phillips, A.W. (1958), "The Relation Between Unemployment and the Rate of Change of Money

- Wage Rates in the United Kingdom, 1861-1957," *Economica* 25, 283-299.
- Poole, W. (1970), "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model," *Quarterly Journal of Economics* 84, 197-216.
- Preston, A.J. (1974), "A Dynamic Generalization of Tinbergen's theory of Policy," *Review of Economic Studies* 41, 65-74.
- Roberts, J.M. (1995), "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit, and Banking* 27, 975-984.
- Romer, P.M., (1986), "Increasing Returns and Long-run Growth," *Journal of Political Economy* 94, 1002-1037.
- Samuelson, P.A., "Interaction Between the Multiplier Analysis and the Principle of Acceleration," *Review of Economic Statistics*, 21, 1939, 75-78.
- Sargent, T.J, (1971), "The Optimum Monetary Instrument Variable in a Linear Economic Model," *Canadian Journal of Economics* 4, 50-60.
- Sargent, T.J. and N. Wallace, "The Stability of Models of Money and Growth with Perfect Foresight," *Econometrica*, 41, 1973, 1043-1048.
- Sargent, T.J. and N. Wallace, "Rational Expectations and the Theory of Economic Policy," *Journal of Monetary Economics*, 2, 1976, 169-183.
- Sims, C.A. (2001), "Solving Linear Rational Expectations Models," *Computational Economics* 20, 1-20.
- Stemp, P.J. and S.J. Turnovsky, "Optimal Monetary Policy in an Open Economy," *European Economic Review*, 31, 1987, 1113-1135.
- Strotz, R.H. (1955-56) "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* 23, 165-180.
- Taylor, J.B. (1977), "Conditions for Unique Solutions in Stochastic Macroeconomic Models with Rational Expectations," *Econometrica* 45, 1377-1385.
- Taylor, J.B. (1993), "Discretion versus Policy Rules in Practice," *Carnegie Rochester Conference Series on Public Policy* 39, 195-214.
- Tinbergen, J., (1952), *On the Theory of Economic Policy*, North-Holland, Amsterdam.

- Tse, E. and Y. Bar-Shalom, (1973), "On Actively Adaptive Control for Linear Systems with Random Parameters," *IEEE Transactions on Automatic Control AC-18*, 98-108.
- Turnovsky, S.J., (1973), "Optimal Stabilization Policies for Deterministic and Stochastic Linear Economic Systems," *Review of Economic Studies* 40, 79-96.
- Turnovsky, S.J., (1976), "Optimal Stabilization Policies for Stochastic Linear Systems: The Case of Correlated Multiplicative and Additive Disturbances," *Review of Economic Studies* 43, 191-194.
- Turnovsky, S.J., (1977a), *Macroeconomic Analysis and Stabilization Policy*, Cambridge University Press, Cambridge, U.K..
- Turnovsky, S.J., (1977b), "Optimal Control of Linear Systems with Stochastic Coefficients and Additive Disturbances," in J. D. Pitchford and S. J. Turnovsky, (eds.), *Applications of Control Theory to Economic Analysis*, North-Holland, Amsterdam, 293-335.
- Turnovsky, S.J., (1980), "The Choice of Monetary Instrument under Alternative Forms of Price Expectations," *Manchester School* 48, 39-63.
- Turnovsky, S.J. (1981), "The Optimal Intertemporal Choice of Inflation and Unemployment: An Analysis of the Steady State and Transitional Dynamics," *Journal of Economic Dynamics and Control* 3, 357-384.
- Wonham, W.M., (1963), *Stochastic Problems in Optimal Control*, RIAS Technical Report 63-14.
- Wonham, W.M., (1968), "On a Matrix Riccati Equation of Stochastic control," *SIAM Journal of Control* 6, 681-697.
- Wonham, W.M., (1969), "Random Differential Equations in Control Theory," in A.T. Barucha-Reid (ed.), *Probabilistic Methods in Applied Mathematics II*, Academic Press, New York.
- Woodford, W.M. (2003), *Interest and Prices*, Princeton University Press, Princeton, N.J.