Dynamic Policy Games in Economics
F. van der Ploeg and A.J. de Zeeuw, (Editors)

MONETARY POLICY IN MODELS WITH CAPITAL*

Finn E. Kydland

Graduate School of Industrial Administration
Carnegie-Mellon University
Pittsburgh, PA 15213, USA

Introduction

Over the past ten or twenty years there has been a clear trend in the direction of modelling aggregate phenomena, whether in a growth or a business-cycle context, as environments in which households and firms make optimal decisions. Such abstractions, especially when assuming competitive equilibria, imply considerable discipline in the sense that it is possible to construct models of quite complex dynamic phenomena with very few free parameters. If a model can be quantitatively restricted by empirical relations that are determined separately from the phenomena being modelled, it generally stands a better chance of giving a useful answer to the question being addressed.

In terms of thinking about policy, the introduction of optimising government initially seemed to offer the possibility of similar scientific benefits. So far, however, our models have raised more questions than they have answered. From a normative point of view, with a government objective, one can certainly determine the optimal policy within a given model. The problem lies in the implementation. As shown by Kydland and Prescott (1977), this government plan, in a dynamic economic environment without commitment, is intertemporally inconsistent. If the plan is implemented today, the government will generally have an incentive to change it in the future. Alternatively, if the time-consistent plan is used, then the results may be very inferior relative to the optimal plan with commitment.

From a positive standpoint, unless one can demonstrate that a commitment mechanism exists, the no-commitment time-consistent equilibrium would appear to be the candidate for a framework for understanding intertemporal policy-making. But this outcome, in some models, can be almost unbelievably bad for society, and it is hard to imagine that there would not be attempts to set up institutional arrangements, or pass laws that are difficult to change, in order that superior policies may be implemented.

To study the nature of monetary policy with or without commitment, including the possibility of commitment arising endogenously, basically two types of models have been used in the literature. One emphasises the role of money for the purpose of stabilisation policy due to a tradeoff between inflation and unemployment[1]. The other takes a public-finance approach[2]. In Section 1, I discuss why the former model is not in the spirit described above. The public-finance view, on the other hand, is a serious framework to build upon and is firmly grounded in the economic tradition described at the beginning of this introduction.

In order to evaluate policies, including both the optimal and the time-consistent ones, it is necessary to be able to determine the competitive equilibrium, given a policy rule that, generally, is a function of the aggregate state of the economy. In Section 2, a direct method which emphasises computability is illustrated for a particularly simple dynamic tax example, but in the context of a competitive industry. In Section 3, the aggregate general equilibrium case is outlined, and the example of a monetary policy rule is described. Section 4 discusses equilibrium definitions and other issues for the case in which the policy rule is chosen by an optimising government.

A major theme of this paper is that capital-theoretic issues are central to time inconsistency. This emphasis is maintained throughout the paper. At the end of each of Sections 3 and 4, I discuss models of money as a medium of exchange in which the demand for money is clearly dependent on the state of the aggregate economy. In Section 4, in particular, I informally outline a theory which would represent an attempt to account for high-frequency movements of the rental price of liquidity, but which also offers a temptation for policy-makers to excessively tax a form of accumulated capital, in this case in the household.

1. Time Inconsistency and Monetary Policy

In this section, I discuss two approaches that have been used for analysing time-consistency issues in the context of monetary policy. One, which is based on an expectational Phillips curve, suffers from at least two drawbacks. The first is that the level of the parameters is such that there is no economic reason why they would be even approximately policy invariant, especially under repeated play. Another related problem is that the model is not formulated in terms of parameters that economists are confident about measuring and restricting so as to get an idea of the quantitative importance of time inconsistency. The public-finance approach, on the other hand, is explicit, the models have parameters that can be measured, and monetary policy changes can imply real changes through any one of several channels. I predict that this, so far largely ignored, area will be an important topic for future research.

1.1 The Inflation-Unemployment Example

An example used in Kydland and Prescott (1977) is the following. The monetary authority maximises

$$-\pi_t^2 - \omega(u_t-u^*+u^o)^2, \qquad \omega, u^o > 0,$$

where u^* is the natural unemployment rate, π_t is the inflation rate, ω expresses the relative weights on the two terms, and the presence of the parameter u^o suggests that, for any given π_t, the preferred value of u is less than u^* (see our Figure 1 on p. 479). The Phillips curve is assumed to be linear:

$$u_t = -\lambda(\pi_t-\pi_t^e) + u^*,$$

where π_t^e is the expected inflation rate as of the beginning of period t. Substituting for u_t, the monetary objective can be written as

$$-\pi_t^2 + \omega[\lambda(\pi_t-\pi_t^e) - u^o]^2.$$

Some variations on this formulation have been used, for example making the natural unemployment rate stochastic as in Barro and Gordon (1983b), or making the second term linear, or letting both terms take a more general functional form; letting the optimal inflation rate be zero is without loss of generality.

This example demonstrates that, if people make their decisions first on the basis of expectations π_t^e, then the government will have an incentive to fool them by choosing $\pi_t > \pi_t^e$, thereby temporarily lowering the unemployment rate. In the long run, however, such a policy would lead to above-optimal inflation.

Barro and Gordon (1983a) made the important observation that the government can, in some cases, support lower inflation through reputation. Others have assumed that policy-maker type, whether high or low-inflation, is uncertain to the public, and various other forms of uncertainty have been introduced.

While this example serves to demonstrate the possible severity of time inconsistency within a framework that most economists are familiar with, it also has obvious weaknesses. For example, results have been derived for repeated play that depend on the parameters of this model. There is, however, no basis for arguing that the parameters are invariant to government policy, such as high versus low-inflation policy. The slope of the Phillips curve surely will depend crucially on the policy that individuals expect the government to follow. An example of a story behind it is an island model. This is not a likely basis for an invariant expectational Phillips curve. Justifications for why a positive value of u^o could be in the interest of the public have also been proposed. For example, it has been suggested that other tax distortions may become less severe as a result of unanticipated inflation. This is not unreasonable, but there is no economic foundation for arguing that the nature of these distortions is the same for high as for low-inflation policy.

Another strike against the model and its more elaborate variants is that, at least for the US over the last thirty years, it is not empirically plausible. If one detrends real GNP and the aggregate price level, whether measured by the GNP deflator or the Consumer Price Index, the price level clearly has been countercyclical[3]. This fact is illustrated in Figure 1. Using price changes (not in logs) rather than the price level, the measured cyclical inflation rate is slightly positively correlated with cyclical real GNP, but with a lag of three or four quarters. These empirical features suggest that monetary authorities have not played a significant role in generating output or employment movements in a way suggested by the model under consideration.

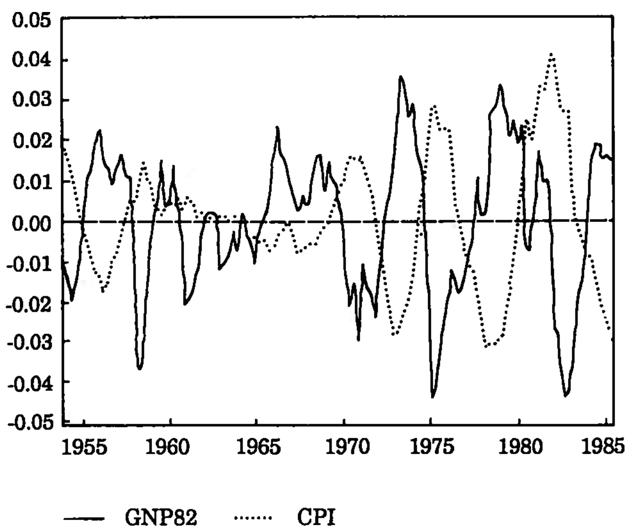

Figure 1: Per cent deviation, real GNP & CPI

There is also a recent literature which considers cooperative arrangements among countries. Some of these models ignore the optimising behaviour by private individuals and focus instead on the game among countries. This approach makes questions of commitment to cooperative arrangements (for example, through trigger-like strategies) quite similar to some of the literature in industrial organisation. Ignoring whether or not one can commit versus the atomistic inhabitants of the countries could make cooperation seem advantageous when it really is not. In contrast, Kehoe (1987) shows in a public-finance model with utility-maximising consumers and benevolent governments that, without commitment versus private agents, the cooperative solution may be inferior for reasons similar to those explaining why the time-consistent solution can be very suboptimal within a given country[4].

1.2 A Public-Finance Perspective

In contrast with this type of model of monetary policy-making, models of fiscal policy have been much more explicit in their specification of what decisions people make and what problems they solve. Our own main example (1977), involving specifically the investment tax credit, was of that type, as was our 1980 public-finance model (and Fischer's illuminating two-period version of it) in which the government can use different tax rates on labour and capital income in order to finance varying government expenditures. The result is that the government, without any commitment mechanism, would tax capital too heavily, as it is inelastically supplied in the short run, and thus would make future saving too low. More recently, the possibility that sustainable equilibria could arise that are not the limit of finite-horizon equilibria has been studied[5].

The omission of productive capital from many models removes a feature that was instrumental in producing large differences between time-consistent and optimal policy rules in our 1977 paper. Intuitively, the durability of capital must be an important reason. Investors would like to know tax rates on capital income many years into the future. An unexpected increase in the tax rate has little effect on the quantity of capital services supplied in the short run, but would affect savings and therefore growth in the longer run and have profound effects on the welfare of the society.

Does money play a significant role in promoting or discouraging growth? The usual examples of such a role are the inflation tax, which most people probably regard as quantitatively small under usual circumstances[6], and the tax on bondholders resulting from an unexpected increase in inflation.

Slightly more subtle, perhaps, but potentially quite important, is the possible increase in tax burden on capital due to interaction with the fiscal tax system. Examples are the value of depreciation write-offs for tax purposes or changes in the progressiveness of labour-income tax, assuming they are non-indexed. An increase in the progressiveness can be regarded as an increase in the tax on human capital[7]. Obviously, the net effect depends on what other simultaneous fiscal changes are made. An interesting and important empirical question is whether higher inflation tends to be associated with a significantly higher relative tax burden on physical or human capital and therefore tends to reduce growth.

2. Industry Equilibrium with Tax Policy

A prerequisite for any evaluation of policies is a method for determining competitive equilibrium, given a policy rule. Especially at the aggregate level, the assumption that households and firms are atomistic and do not behave strategically is essential. Any other assumption is probably unrealistic and leaves room for too many possibilities that are hard to evaluate and that leave too many degrees of freedom.

A method for computing competitive equilibrium enables one to compare alternative policy rules, such as the optimal policy rule (if it could credibly be committed to) as opposed to the time-consistent one. One could perhaps also evaluate simple rules which may be easier to implement while being not much inferior to the usually more complex optimal commitment policy.

In this section, I illustrate the determination of competitive equilibrium with tax policy in a particularly simple dynamic environment that we can think of as a competitive industry. This is a slightly simplified version of the main example used in Kydland and Prescott (1977). I compare the direct solution procedure with solving a social planner's problem, and point out that this latter method cannot be used when there are government policy rules that depend upon the state of the economy. In the next section, then, the more interesting case of an aggregate general equilibrium is considered.

In this model, investment planned in period t, x_t, and carried out in that period and the next, does not become productive until period t + 2. Thus, the law of motion for the firm-specific productive capital is

$$k_{t+1} = (1-\delta)k_t + x_{t-1}, \tag{2.1}$$

and in the aggregate

$$K_{t+1} = (1-\delta)K_t + X_{t-1}, \tag{2.2}$$

where δ is the physical rate of depreciation. The investment rate in period t is then

$$i_t = \varphi x_t + (1-\varphi)x_{t-1}, \tag{2.3}$$

where ψ is the fraction of the investment effort induced by plan x_t in the same period. Investment expenditures in period t are $qi_t + \gamma(i_t-\delta k_t)^2$, with a tax offset of $Z_t i_t$, where Z is the investment tax credit. Thus, the model combines cost of adjustment with the feature that it takes multiple periods before planned additions to the capital stock become productive[8]. These dynamic elements, and especially the time-to-build feature, turned out to have significant implications for the nature of time inconsistency in the model.

The cash-flow function of the typical firm is

$$P_t k_t - (q-Z_t)i_t - \gamma(i_t-\delta k_t)^2,$$

where units are chosen such that one unit of capital produces one unit of output[9]. The inverse aggregate demand function is assumed to be linear:

$$P_t = A_t - bK_t, \tag{2.4}$$

where b is positive, and A_t is a stochastic shift parameter that is subject to the first-order autoregressive process

$$A_{t+1} = \rho A_t + \epsilon_t, \qquad -1 < \rho < 1, \tag{2.5}$$

where ϵ is a positive independent random variable with mean μ and constant variance. Finally, we assume that the investment tax credit is chosen according to the rule

$$Z_t = z_0 + z_1 K_t + z_2 X_{t-1} + z_3 A_t. \tag{2.6}$$

Defining $R(k_t,x_{t-1},x_t,K_t,A_t,z_t)$ to be the objective function after we have substituted for i_t and P_t from (2.3) and (2.4), the typical firm's objective is to maximise

$$E \sum_{t=0}^{T} \beta^t R(k_t,x_{t-1},x_t,K_t,A_t,Z_t), \qquad 0 < \beta < 1,$$

where T could be infinity. Let $v_t(k_t,x_{t-1},K_t,X_{t-1},A_t)$ be the value to the firm of pursuing optimal decision rules for the remainder of the horizon, given the state of the economy at the beginning of period t. These functions satisfy the recursive relationship

$$v_t(k_t, x_{t-1}, K_t, X_{t-1}, A_t)$$

$$= \max_{x_t} E[R(k_t, x_{t-1}, x_t, K_t, A_t, Z_t) \tag{2.7}$$

$$+ \beta v_{t+1}(k_{t+1}, x_t, K_{t+1}, X_t, A_{t+1})],$$

for all t, given relations (2.1), (2.2), (2.5) and (2.6), along with expectation of future prices. These expectations can be expressed in terms of the expected laws of motion of industry-wide investment:

$$X_t = D_t^e(K_t, X_{t-1}, A_t, \pi_t),$$

which map uniquely into expectations of the process determining future prices through equations (2.2), (2.4) and (2.5).

Assuming now that the quadratic v_{t+1} has already been determined, one obtains a first-order condition at time t which implies a decision rule of the form

$$x_t = d_{0t} + d_{1t}k_t + d_{2t}x_{t-1} + d_{3t}K_t + d_{4t}X_{t-1} + d_{5t}A_t + d_{6t}X_t + d_{7t}Z_t, \tag{2.8}$$

where d_{it}, $i = 0,\ldots,7$, are the coefficients that are obtained for time period t. Note that this individual decision rule is made a function of the aggregate decision, X_t, in the industry. In order to get the aggregate or per-firm behavioural rule, which will also be the basis for the firms' expectations, we impose the conditions that $x_t = X_t$, $k_t = K_t$, and $x_{t-1} = X_{t-1}$, yielding

$$X_t = \frac{d_0}{1 - d_{6t}} + \frac{d_{1t} + d_{3t}}{1 - d_{6t}} K_t + \frac{d_{2t} + d_{4t}}{1 - d_{6t}} X_{t-1} + \frac{d_{5t}}{1 - d_{6t}} A_t + \frac{d_{7t}}{1 - d_{6t}} Z_t$$

$$\equiv D_t(K_t, X_{t-1}, A_t, Z_t). \tag{2.9}$$

The linearity of the decision rules allows us to aggregate in this way, so that the aggregate or per-firm values of the state variables contain the necessary information.

If $D_t \neq D_t^e$, aggregate behaviour will give rise to price distributions that are different from the expected distributions on the basis of which the individual decisions were made. In equilibrium, expectations are rational, and we therefore close the model by requiring that $D_t = D_t^e$ for all t. From a computational point of view, this can be handled by requiring that this condition hold at each iteration of the recursive procedure outlined above. To complete the computations for period t, then, we substitute for the linear relations (2.1), (2.2), (2.5) and (2.6), and the linear individual and aggregate behavioural rules, (2.8) and (2.9), in the right-hand side of the functional equation (2.7). The resulting function is quadratic and depends only on the individual and aggregate state of the economy at time t and represents the value function to be used in determining the (equilibrium) behaviour in period t - 1. Thus, this procedure incorporates the assumption that agents' expectations of future price distributions are rational.

For the infinite-horizon case, we can think of the above description as outlining one step in the successive approximations, the limit of which is generally a stationary decision rule. For this framework outlined, the aggregate behaviour of economic agents is then given by the relation

$$X_t = D(K_t, X_{t-1}, A_t, Z_t),$$

the coefficients of which are the same in every period.

In view of the results of Lucas and Prescott (1971), this procedure may appear unnecessarily complicated. Assuming for the moment that $Z_t = z^* =$ constant in every period, they showed that a competitive industry such as the one above behaves in equilibrium as if maximising a certain consumer surplus. Thus, we might consider solving this stand-in problem, which is to maximise

$$E \sum_{t=0}^{T} \beta^t S_t,$$

where

$$S_t = \int_0^{K_t} (A_t - bu)du - (q-z^*)I_t - \gamma(I_t - \delta K_t)^2,$$

subject to

$$K_{t+1} = (1-\delta)K_t + X_{t-1},$$

$$A_{t+1} = \rho A_t + \epsilon_t,$$

and where

$$I_t \equiv \psi X_t + (1-\psi)X_{t-1}.$$

This is a problem of smaller dimension the solution of which can be obtained by standard recursive methods. When the policy variable depends on aggregate industry behaviour, however, this simpler method is not immediately applicable. Since all firms are small, each firm assumes that it has no effect on future policies. Of course, if all firms invest less now, future capital stock will be lower, thus most likely future tax credits will increase. In the stand-in problem, this effect of current decisions on future taxes is recognised, and the competitive equilibrium therefore does not correspond to the solution of that maximisation problem, given the policy rule.

From an empirical standpoint, as pointed out by Lucas (1976), attempts at estimating equations like (2.9) and using them for policy evaluation will fail because the coefficients will not be invariant to changes in the policy rule (2.6). Furthermore, it is unlikely that this is a quantitatively unimportant phenomenon, especially in environments with structural dynamics such as the one considered here.

3. Dynamic Competitive Equilibrium with Government Policy

In this section, we outline a more general framework. It is capable of dealing with a variety of general equilibrium models with optimising households and firms facing government policy which depends on their aggregate behaviour. We concentrate on consumers' maximisation, although implicit is also firms' maximisation. The consumers are thought of as renting the input factors, including capital, to the firms.

In the spirit of thinking about computable models the quantitative, and not just qualitative, properties of which can be studied, we assume that the structure of the model is such that maximisation and equilibrium result in linear decision rules. This will typically require that a quadratic approximation be made for the utility function around its steady state which can usually be determined analytically. The utility function may be indirect in the sense that a budget or resource constraint has been substituted. In examples without policy in which this approach has been compared with exact methods, the properties have been very similar. A factor is that most aggregate fluctuations around their growth paths are small in terms of percentage. This formulation permits a great deal of quantitative discipline in that prior knowledge from outside the model can be used to restrict the parameters. Examples are capital depreciation rates, long-run ratios of key variables, including factor shares, elasticity of long-run labour supply, and so on.

Let x_t be the vector of decision variables for the representative consumer at time t, and y_t be the individual-specific vector of endogenous state variables summarising all the information needed for making decisions at time t. Typical examples are capital stocks or money holdings at the beginning of the period, although in some cases variables dated before t may have to be included. Also, let X_t and Y_t be the corresponding aggregate variables. For example, if y_t includes the agent's capital stock, then the aggregate capital stock is an element of Y_t. The laws of motion for the state variables are given by the linear equations

$$y_{t+1} = f(y_t, x_t) \tag{3.1}$$

$$Y_{t+1} = F(Y_t, X_t). \tag{3.2}$$

Generally, the function f will be the same as F. In addition, there may be state variables the paths of which are determined by autonomous processes. These variables are assumed to follow an autoregressive process

$$W_{t+1} = \Omega W_t + \eta_t, \tag{3.3}$$

where Ω is a matrix of fixed coefficients, and η is a random vector with fixed variance. Examples are technology shocks with serial correlations. Finally, Z_t is a vector of policy variables assumed to be determined by a linear policy rule

$$Z_t = G(Y_t, W_t), \tag{3.4}$$

which incorporates the government budget constraint and is correctly anticipated. Examples of variables in Z are tax rates, government purchases, or the change in the nominal money stock.

The consumer is assumed to maximise the indirect utility function

$$E \sum_{t=0}^{T} \beta^t u(x_t, y_t, P_t, Z_t),$$

which is obtained after having substituted for the private budget constraint. The price vector, P_t, is included because of this substitution. Relations determining the dynamic motion of prices will be endogenously determined.

The value function, $v_t(y_t, Y_t, W_t)$, denotes the (equilibrium) expected discounted value or utility at time t for a consumer in initial state y_t when pursuing optimal decision rules, and the initial aggregate state is (Y_t, W_t). Implicit in this value function is the assumption of rational expectations about future prices in the sense that their expected distributions are those generated by the equilibrium laws of motion of the economy. The decision rules and corresponding value functions can be determined recursively from the relation

$$v_t(y_t, Y_t, W_t) = \max_{x_t} E[u(x_t, y_t, P_t, Z_t) + \beta v_{t+1}(y_{t+1}, Y_{t+1}, W_{t+1})],$$

$$t = 0, 1, \ldots, T, \tag{3.5}$$

subject to constraints (3.1) - (3.4).

Assuming now that v_{t+1} has already been determined, and taking account of equations (3.1) - (3.4), the first-order conditions for a maximum at time t determine linear decision rules of the form

$$x_t = d_t(y_t,Y_t,W_t,X_t,P_t,Z_t). \tag{3.6}$$

As in Section 2, we note that the aggregate X_t is included in the right-hand side. These individual decision rules can be aggregated. It is convenient to think of X_t and Y_t in per-consumer terms. Thus, in the aggregate, we have $x_t = X_t$ and $y_t = Y_t$. Therefore

$$X_t = d_t(Y_t,Y_t,W_t,X_t,P_t,Z_t),$$

which can be written as

$$X_t = D_t(Y_t,W_t,P_t,Z_t). \tag{3.7}$$

The prices must be such that markets are cleared. These prices will depend on the aggregate state, i.e.,

$$P_t = P_t(Y_t,W_t,Z_t). \tag{3.8}$$

For example, if borrowing and lending are possible, then the price of these loans (which implicitly defines the interest rate) must be such that aggregate net loans are zero in every period. Or, as a second example, if consumers hold money from one period to the next, then the equilibrium price of money in terms of goods must be such that the aggregate amount of money that individuals wish to hold is equal to the amount supplied. The supply of money can in general depend on the state of the economy.

From a computational point of view, as in the preceding section, the equilibrium condition is required to hold at each iteration of the recursive procedure outlined above. To complete the computations for period t, therefore, we substitute for the linear laws of motion, (3.1) - (3.4), and the linear individual and aggregate behavioural rules, (3.6) and (3.7), in the right-hand side of the functional equation (3.5). The resulting function is quadratic and depends only on the individual and aggregate states (y_t,Y_t,W_t), and represents the value function to be used in determining the (equilibrium) behaviour in period t - 1. Thus, this procedure incorporates the assumption that agents' expectations of future price distributions are rational.

If v_{t+1} is quadratic, a quadratic function is maximised at time t subject to linear constraints. The optimal decision rules are then linear, and the new value function, v_t, is quadratic. By definition, v_{T+1} is zero and therefore trivially quadratic, and so all v_t are quadratic by backward induction.

We summarise this discussion in the following

DEFINITION. An equilibrium is a sequence of decision rules (3.6) and a sequence of price functions, $\{P_t(Y_t,W_t,Z_t)\}_{t=0}^{T}$, such that, for each individual, the decision rules solve recursively the functional equations (3.5) subject to laws of motion (3.1) - (3.4), the expected behaviour of the aggregate variables in each period, $D_t(Y_t,W_t,P_t,Z_t)$, which must be consistent with aggregation (or average) of the individual decision rules, and such that prices determined by the functions $P_t(Y_t,W_t,Z_t)$ clear markets for all t.

In this paper, we assume stationary models in the sense that utility functions and laws of motion are the same in every period. For infinite-horizon models, such structures generally yield stationary decision rules which we obtain by determining the limit as the horizon becomes large. For the framework outlined above, the aggregate behaviour of economic agents is then given by the relation

$$X_t = D(Y_t,W_t,P_t,Z_t),$$

the coefficients of which are the same in every period, and similarly for the price relations.

An example is the business-cycle model with money used in Kydland (1987). If n_t is hours of work in period t and T is the total time allocation per period, then net leisure in period t is $T - n_t + \ell(P_t m_t)$, where $\ell' > 0$ and $\ell'' < 0$, and m_t is the nominal quantity of money held by the typical household at the beginning of period t. The price level, P_t, is the inverse of the conventional price level. Thus, there is a tradeoff in the household between real money and leisure.

With this money-holding motive embedded in a real business cycle model, the resulting demand for money can be written as

$$m_t^d = m^d(y_t,Y_t,M_t,P_t),$$

which in the aggregate becomes

$$M_t^d = M^d(Y_t, P_t).$$

For this market to clear, the price level must satisfy

$$M^d(Y_t, P_t) = M^s(Y_t),$$

where M^s is an unchanging monetary policy rule. The resulting price function, $P(Y_t)$, is used by rational consumers in forming expectations about future price distributions. If M^d and M^s are both linear, then $P(Y_t)$ is also linear.

By keeping the model computable, it is possible to compare its covariance properties to those of the data. The aggregate behavioural relations, the laws of motion for the state variables, the equilibrium price relations, and the policy rules together form a system of difference equations. From this set of stochastic equations, repeated draws can be made of the same length as the data series available, and similar statistics can be computed for the artificial economy and the data. This, indeed, was the purpose for which the above-mentioned model of money as a medium of exchange was used. In particular, it is consistent with the observed countercyclical price level (conventionally defined) illustrated in Figure 1.

4. Optimising Government

So far, we have described the determination of equilibrium, given an exogenous policy function, without describing where the policy came from. This is a prerequisite for dealing with the case in which policy is determined from optimisation by the government. That is, the government wishes to maximise a social objective function

$$E \sum_{t=0}^{T} \beta_z^t S(X_t, Y_t, Z_t), \qquad 0 < \beta_z < 1,$$

subject to its budget constraint. For simplicity of notation, we have included the exogenous shocks in Y_t. If distributional considerations are not an issue, then a common abstraction is one in which all consumers are alike. In that case, a natural candidate for S is the equal-weighted or average utility function with $\beta_z = \beta$. It is well known that, in a dynamic environment such as one with productive capital, the optimal plan without commitment is still time inconsistent[10]. This is the case so long as, for example, tax rates are chosen by the policy-maker and allocations are selected by consumers.

The extension of our definition of equilibrium to include a time-consistent government continues to be recursive. Let $R_t(Y_t) = \Sigma_{\tau=t}^{T} \beta_z^{\tau-t} S(X_\tau, Y_\tau, Z_\tau)$ given that policy is selected consistently from time t on and the economy is competitive. Also, for any given Z_t, let $X_t = D_t(Y_t, P_t, Z_t)$ and $P_t = P_t(Y_t, Z_t)$ represent the competitive equilibrium resulting from solving the functional equation (3.5) subject to constraints (3.1) - (3.3) and after aggregation as described in going from equation (3.6) to (3.7). Then

$$R_t(Y_t) = \max_{Z_t} E[S(X_t, Y_t, Z_t) + \beta_z R_{t+1}(Y_{t+1})],$$

subject to $X_t = D_t(Y_t, P_t, Z_t)$ and the government budget constraint. The resulting solution is of the form $Z_t = G_t(Y_t)$.

The infinite-horizon case is now trickier. If there is a time-consistent equilibrium, one can be found by taking the limit of the above procedure. There may, however, be other equilibria that depend on how the policy-maker has behaved up until that time. A definition of a time-consistent stationary policy rule for the infinite horizon is given in Kydland and Prescott (1977, p. 481).

Note the contrast with the standard inflation-unemployment model. In that model without structural dynamics in the form of, for example, capital accumulation, time inconsistency arises because private agents make their decisions before the government does. It is more reasonable to consider environments in which, on day one of each time period, the government chooses tax rates and other policy variables, and, given these choices by the government, households and firms subsequently make their decisions for that time period. With capital accumulation of some form or another being a part of the model, time consistency is still an important consideration. This is because some private decisions in period t affect the state of the economy in the future and in the aggregate determine these tax rates through the tax policies. At the same time, period-t decisions are affected in an important way by these (expected) future tax rates.

A typical example is where government expenditures can be financed by taxing labour and capital income. The standard result is that, unless one can commit to the optimal plan (sometimes called the Ramsey plan), capital will be taxed too heavily. It is important to realise, however, that it is not necessary that capital be taxed directly for time consistency to be an issue. It will play a role so long as the quantity of what is being taxed is chosen as a function of some form of capital. For example, labour supply may depend on already accumulated capital. If so, that dependence gives rise to a time-inconsistency problem even if only the labour income is taxed. Another example is the case of intertemporally non-separable utility in leisure, which can be interpreted as a stand-in for households' allocation of part of their non-market time to, perhaps unobserved, capital in the households, such as quality of children, health, and so on.

The model of money as a medium of exchange outlined in the preceding section has the property that the demand for money is highly dependent on aggregate state variables, including physical capital and perhaps also unobserved capital in the household. It is possible to construct reasonable models of money that have a direct capital-theoretic element. Assume that there are at least two distinct ways of carrying out transactions. In the language of household production theory, one can think of the inputs as being real money, the allocation of time in that period, and the input of a form of household capital that results from previous uses of time. Whenever a significant change is made in the method households use for payment, such as one which significantly economises on the use of cash, there is probably some degree of learning taking place over a period of time as people gain experience with the new method.

Given that it takes time to accumulate this experience, one would expect considerably more movement in the rental price of liquidity services than otherwise. This would be consistent with the empirical puzzle of excess volatility of short-term interest rates relative to the standard demand for money function as demonstrated clearly in Lucas (1987), especially in his Figure 5.

Whenever a capital stock is being accumulated, an insight of the time-inconsistency literature is that it will be tempting for the government to tax it excessively, in this case indirectly. Whether this phenomenon can be of quantitatively significant magnitude in such a context is an interesting question for future research.

5. Conclusion

This paper has been concerned with monetary policy in environments in which dynamics are important due to capital-theoretic elements. These are also environments in which there is considerable scope for the phenomenon commonly referred to as the time inconsistency of optimal policy. In this context, three themes were elaborated on. The first is the importance of computable models that can be quantitatively restricted and in which the quantitative importance of the phenomena under study can be addressed. Such models could then be used, for example, as a basis for an assessment of the value of having a commitment mechanism.

The second theme is that two-player games, or games in which the decision problems of individuals in the economy are not explicitly formulated, are likely to be of little value in understanding aggregate government policy. An appropriate framework is one in which atomistic people's optimisation problems are explicitly formulated. In such models, unlike representative agent models the equilibrium of which can be determined by solving a stand-in planner's problem, the distinction between individual decision and state variables on the one hand and their aggregate counterparts on the other becomes important. Equilibria with optimising governments can then be used to obtain insights into many issues related to the operation of government policy.

The third theme is that incentives for capital accumulation are at the heart of time inconsistency. In particular, while most people may accept that view in the context of fiscal policy, it is less common in discussions of monetary policy. I argue that, in this sense, there is not really an important distinction between monetary and fiscal policy. Inflation is likely to affect capital-accumulation decisions, both physical and human. I also gave other examples of model features in the context of money which have capital-theoretic elements. The empirical relevance of such features is an important topic for future research.

Notes

* I have benefited from comments by Patrick Kehoe and Torsten Persson. The paper was written while the author was visiting the Federal Reserve Bank of Minneapolis. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. For some surveys, see Barro (1985), Blackburn and Christensen (1987), Chari, Kehoe and Prescott (1988), and Rogoff (1987).

2. See Calvo (1978), Lucas and Stokey (1983), and Persson, Persson, and Svensson (1987).

3. These results are from Kydland (1987).

4. See also van der Ploeg (1987).

5. See Chari and Kehoe (1988a,b).

6. This argument is made in Kydland (1983).

7. See, however, Cooley and Hansen (1987).

8. The cost-of-adjustment assumption is often used in industry or firm models so as to make investment expenditures smooth over time. This is probably inappropriate at the aggregate level, and certainly not necessary for the purpose mentioned above. In aggregate equilibrium, the real interest rate will adjust and affect the willingness to substitute between present and future consumption.

9. In the original model, labour was an input as well. This input had no intertemporal features, however, and therefore did not play a role for the issue of time inconsistency.

10. See Kydland and Prescott (1980).

References

Barro, R.J., 'Recent developments in the theory of rules versus discretion', Working Paper 12, University of Rochester, 1985.

Barro, R.J. and D.B. Gordon, 'Rules, discretion and reputation in a model of monetary policy', Journal of Monetary Economics, 12, 1983a, pp. 101-121.

Barro, R.J. and D.B. Gordon, 'A positive theory of monetary policy in a natural rate model', Journal of Political Economy, 91, 1983b, pp. 589-610.

Blackburn, K. and M. Christensen, 'Monetary policy and policy credibility: theories and evidence', Working Paper, 1987.

Calvo, G.A., 'On time consistency of optimal policy in a monetary economy', Econometrica, 46, 1978, pp. 1411-1428.

Chari, V.V. and P.J. Kehoe, 'Sustainable plans', Research Department Working Paper 377, Federal Reserve Bank of Minneapolis, 1988a.

Chari, V.V. and P.J. Kehoe, 'Sustainable plans and debt', Research Department Working Paper 399, Federal Reserve Bank of Minneapolis, 1988b.

Chari, V.V., P.J. Kehoe and E.C. Prescott, 'Time consistency and policy', Research Department Staff Report, Federal Reserve Bank of Minneapolis, 1988.

Cooley, T.F. and G.D. Hansen, 'The inflation tax and the business cycle', Working Paper, 1987.

Fischer, S., 'Dynamic inconsistency, cooperation and the benevolent dissembling government', Journal of Economic Dynamics and Control, 2, 1980, pp. 93-107.

Kehoe, P.J., 'International policy cooperation may be undesirable', Research Department Staff Report 103, Federal Reserve Bank of Minneapolis, 1986.

Kydland, F.E., 'Implications of dynamic optimal taxation for the evolution of tax structures: a comment', Public Choice, 41, 1983, pp. 229-235.

Kydland, F.E., 'The role of money in a business cycle model', Working Paper, Carnegie-Mellon University, 1987.

Kydland, F.E. and E.C. Prescott, 'Rules rather than discretion: the inconsistency of optimal plans', Journal of Political Economy, 85, 1977, pp. 473-491.

Kydland, F.E. and E.C. Prescott, 'Dynamic optimal taxation, rational expectations and optimal control', Journal of Economic Dynamics and Control, 2, 1980, pp. 79-91.

Notes

* I have benefited from comments by Patrick Kehoe and Torsten Persson. The paper was written while the author was visiting the Federal Reserve Bank of Minneapolis. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. For some surveys, see Barro (1985), Blackburn and Christensen (1987), Chari, Kehoe and Prescott (1988), and Rogoff (1987).

2. See Calvo (1978), Lucas and Stokey (1983), and Persson, Persson, and Svensson (1987).

3. These results are from Kydland (1987).

4. See also van der Ploeg (1987).

5. See Chari and Kehoe (1988a,b).

6. This argument is made in Kydland (1983).

7. See, however, Cooley and Hansen (1987).

8. The cost-of-adjustment assumption is often used in industry or firm models so as to make investment expenditures smooth over time. This is probably inappropriate at the aggregate level, and certainly not necessary for the purpose mentioned above. In aggregate equilibrium, the real interest rate will adjust and affect the willingness to substitute between present and future consumption.

9. In the original model, labour was an input as well. This input had no intertemporal features, however, and therefore did not play a role for the issue of time inconsistency.

10. See Kydland and Prescott (1980).