5

Stabilization Policy — Should we?

In Chapters 3 and 4 we discussed the potential effectiveness (for output) of stabilization policy by which we mean feedback policy rather than 'automatic stabilizers' which will generally be effective for reasons we explained. We have shown that in general such policy can be effective. The question then arises of whether any such policy is optimal and, if so, how it should be designed. In practice we will focus on monetary policy, since fiscal policy is generally regarded as too slow in implementation and too uncertain in effect to be useful (besides the concern that varying spending programmes and tax rates may cause micro distortions). In this chapter we will discuss the issue of optimality: should we stabilize the economy? In the next we will discuss practical issues of design.

We will discuss two issues that bear on optimality: the welfare criterion and the question of credibility (or 'time inconsistency'). The first concerns whether it is not better to leave the economy alone, letting it react in accordance with the decisions of private agents; has the government reason to believe it can improve on these? The second issue concerns the problems raised by the government's natural desire to manipulate the economy in order to please its electorate; this may lead to the abandonment of inflation pledges in order to stimulate economic activity. Thus governments may find it hard to pursue a consistent anti-inflation policy and this will undermine its credibility.

THE WELFARE CRITERION

The yardstick of welfare that is most widely used by economists is the Pareto criterion. Pareto optimality is the condition in which no one can be made better off without someone else being made worse off; it can be illustrated for a two-person world by an Edgeworth box diagram (figure 5.1), in which the 'contract curve' joins the points of tangency of the two sets of indifference curves, these points being all Pareto-optimal. If it is then assumed that distributional considerations are absent, either because they do not matter or because the government achieves the socially desired distribution at all times, there can be a unique Pareto optimum for the specified distribution, illustrated by the circled dot in figure 5.1.

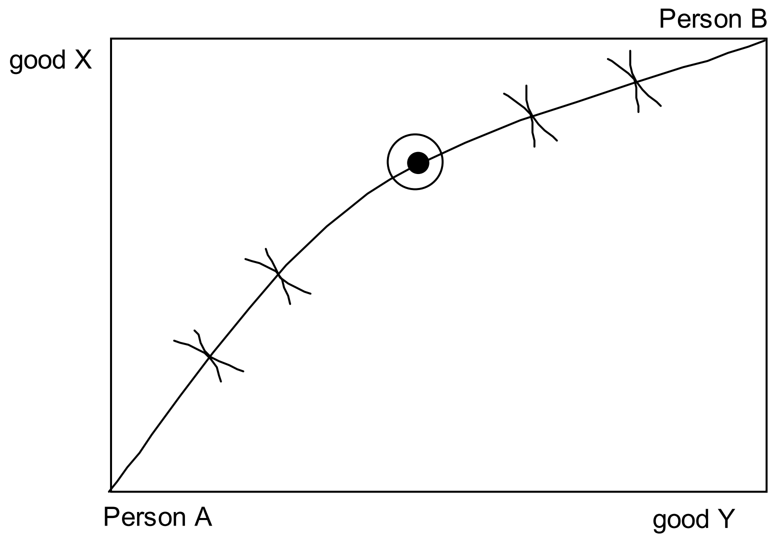

Figure 5.1: An Edgeworth box diagram

The proof that, for a given distribution function there exists a unique Pareto optimum has been established for an economy with well-behaved preferenes and technology in a Walrasian equilibrium when markets are competitive and complete and there are no distortions (discrepancies between private social costs due, for example, to externalities); the proof is due to Arrow, Debreu and Hahn (see Arrow and Hahn, 1971). In an economy of this type, the steady-state level of output at the Pareto optimum would be the y^* in our models.

Of course, in actual economies, with incomplete and some monopolistic markets and with distortions, y^*, the equilibrium level of output, will not be at a Pareto optimum. However, it is usual in macroeconomics to assume that these problems are the province of 'micro' policies, and that there should be no attempt by macroeconomic policy to push y systematically away from y^*; indeed we have already seen that any such attempt in an adaptive expectations model would cause ever-accelerating or ever-decelerating prices and in a rational expectations model would be completely frustrated. Instead y^* is taken, from the viewpoint of macroeconomic policy, to be the optimum output level in a steady state (it is, strictly, the 'second-best' optimum — but the best that can be achieved with macroeconomic policy instruments).

Supposing this to be so, the question arises of what is the optimum short-run output level. In Keynesian disequilibrium models with adaptive (or other backward-looking) expectations, if output is less (or more) than y^*, this is involuntary and suboptimal, the result of 'market failure'; hence the variation of output around y^* is a natural criterion for minimization.[1] Since output tends to y^* in the long run but y^* will typically be growing over time, we can measure this variation as the variance of output around its moving trend $\text{var}(y - y*)$. This is what we focused on in the discussion of effectiveness in chapter 4.

Sometimes in these models other measures have also been included in the minimand to represent the costs due to loss of consumer surplus not included in GDP — for example, the variance of inflation or interest rates as proxies for consumer and financial uncertainty. However, it is usually assumed in the context of stabilization policy that the variations in these costs across alternative policies are relatively small; this would also be our assumption.

In equilibrium models with adaptive expectations (such as the one implicit in Milton Friedman's, Friedman, 1968), output deviations from y^* arise because of expectations errors which could have been avoided by efficient use of available information. Again $\text{var}(y - y*)$ is a natural minimand because agents would wish they had made output decisions on the basis of good forecasts.

In equilibrium models with rational expectations, however, y only departs from y^* because of unavoidable expectations errors. Such models include the New Keynesian ones of chapter 4, because the nominal wage contracts in these models are voluntarily negotiated, in the full knowledge that when shocks occur the response to them will be constrained by the contract. (This leaves open the possibility that monetary policy can improve the private sector's feasible outcomes from its contract structure — we assume not here but return to the issue in chapter 6.)

Output is always at its 'desired' level in these models in the sense that, given available information, agents are maximizing their welfare subject to their private constraints, including the effects of shocks and associated expectations errors. Then, provided the level of distortion in

[1] If output, however, changes for supply reasons (e.g. crop failure), this would naturally be added into y^* and the variation computed around this adjusted figure.

the economy does not vary with $(y-y^*)$, government cannot improve and may reduce welfare by reducing $var(y - y^*)$ for the simple reason that it was already being maximized; an example of this is given by Sargent (1979, chapter 16) in a classical labour supply model where fiscal policy is effective because of intertemporal substitution, as discussed in chapter 4, and the same point has been stressed by Beenstock (1980).[2]

The proviso that the level of distortion does not vary with $(y-y^*)$will be violated in practice when unemployment benefits do not vary with wages; for when depressive shocks reduce output and employment, the gap between the private and social costs of unemployment will increase.

A simple way of evaluating the social gain to stabilization policy is shown in figure 5.2. Assume that the undistorted supply of labour would be at L^+; it is entirely inelastic. Assume also that if employed it could be put to work at the normal wage of w^*, this representing its social product because there is some long-term project (e.g. building a road) whose return is the same in slump as in boom. Then assume that policy can either stabilise employment (output) in face of shocks entirely at L^* or allow employment to vary to L_1 on the upside and L_2 on the downside. The extra employment has social utility, since it gets people closer to the optimum, L^+. We can measure the social loss as $w^*(L^+ - L)$.

We now notice from figure 5.2 that an upward shock has less absolute effect on employment than a downward shock because the supply curve gets progressively flatter as wages fall, getting closer and closer to the benefit level. This implies that as the economy fluctuates the average level of employment falls. It follows that the social loss is on average greater too. So plainly fluctuations lower welfare.

Policy makers would of course like to keep the economy at L^+, minimising fluctuations around this point. However, stabilization policies which attempt to boost L over L^* (y above y^*) systematically will be frustrated in the attempt and will create other problems for the economy (inflation in the case of monetary policy and high taxation in the case of fiscal policy). Hence var $(y - y^*)$ is the minimand subject to the constraint that stabilization policies are not on average biased away from y^*.

To sum up, in rational expectations models, there are prima facie reasons to believe that in general stabilization policy, even if effective, will not improve welfare, the main exception being where distortions are correlated with the cycle; the importance of this exception will vary with the nature of the benefit system.

---

[2] This situation is not one of full Pareto-optimality, as markets are incomplete and distortions assumed to exist; however, it is a situation of restricted Pareto-optimality (second-best), under the provisos given.

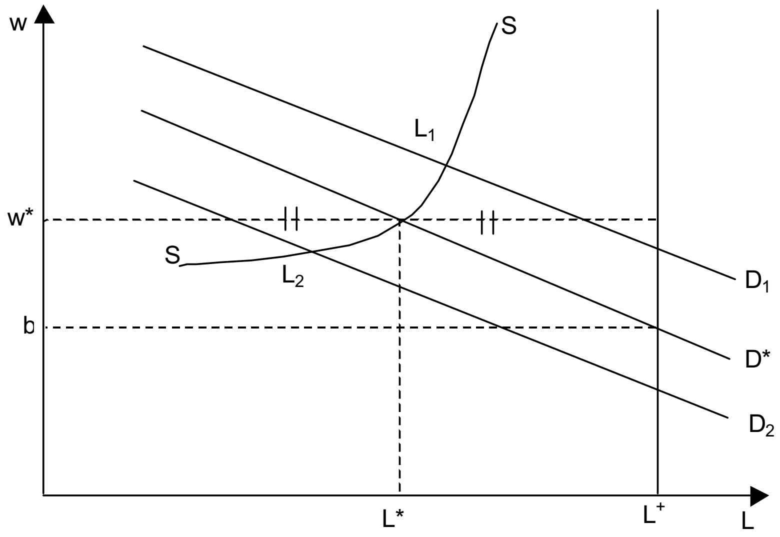

Figure 5.2: Benefit distortions and stabilization

This fundamental argument for stabilization policy is in principle also an argument for a change in the benefit system. Of course, if for some reason the benefit system cannot in practice be changed, then stabilization policy is justified on the 'second best' grounds we have described. But this does not imply that one should uncritically accept the system and the costs it incurs, even with stabilization policies to mitigate them.

OPTIMAL ECONOMIC POLICIES AND TIME INCONSISTENCY

A crucial feature of an economy with rational expectations is that current outcomes and the movement of the system's state depend in part upon anticipated future policy decisions.

The same feature will obviously not occur when expectations are formed in a backward-looking manner, since in these circumstances current outcomes and the movement of the system's state depend only upon current and past policy decisions and upon the current state: they need to allow for the full effects on private behaviour of people's expectations about the outcomes of the rules (we discuss this issue, often known as

Lucas' Critique, in chapter 6). However, having computed the rules appropriately, policy-makers face a further difficulty — or opportunity — if they are not bound in some way to stick to the rule they so compute but people expect them to stick to it nevertheless. For after one period of operating this rule, people have made certain commitments, expecting the rule to prevail; the policy maker now faces a new situation, in that he can exploit these commitments in making policy afresh. Such policy behaviour is known as 'time-inconsistency'.

Of course, matters will not rest there, since the people, thus tricked and exploited, will be less willing to believe that the fresh policy rule will prevail without yet further twists. There is also the question of whether people would have been willing to believe the original policy rule in the first place if the policy-makers could wriggle out of it. Finally, we may ask why policy-makers should wish to trick people in this way 'for the social good' — on the face of it, it is peculiar that social good can be achieved by trickery. However, we defer these developments of the argument until after an explanation of the mechanism behind time-inconsistency.

Kydland and Prescott (1977) showed how, in a dynamic economic system in which agents possess rational expectations, the optimal policy at time $t = 0$, which sets future paths for the control variables (taxes, subsidies, monetary growth), implies values of the control variables at some later time $t+i$, that will not be optimal when policy is re-examined at $t+i$, even though the preferences of agents are unchanged. This they called the time inconsistency of optimal plans.

Two examples (given by Kydland and Prescott, 1977; Fischer, 1980b) relate to examinations and patent policy. Optimal policy at the beginning of a course is to plan to have a mid-session exam. However, on the morning of the exam when all student preparation is complete, the optimal policy is to cancel the exam, saving the students the trouble of writing and the lecturer the trouble of grading. Optimal policy may analogously be to withdraw patent protection after resources have been allocated to successful inventive activity on the basis of continued patent protection.

The argument's formal structure can be seen in a simple two-period problem presented by Kydland and Prescott. Let $\pi = (\pi_1,\ \pi_2)$ be a sequence of policies for periods 1 and 2 and $x = (x_1,\ x_2)$ be the corresponding sequence for economic agents' decisions. The social objective function is given by

$$S(x_1, x_2, \pi_1, \pi_2) \tag{10}$$

Consider first the situation with backward-looking expectations, where no time inconsistency arises. For simplicity we will assume an absence of stochastic error terms, which does not affect the argument. In this case the model is:

$$x_1 = x_1(\pi_1) \tag{11}$$

$$x_2 = x_2(x_1, \pi_2) \tag{12}$$

The policy maker maximizes (1) subject to this model, meeting the first order conditions of a maximum (we assume the S function is well-behaved so that second-order conditions are met):

$$0 = \frac{\partial S}{\partial \pi_1} = \left(\frac{\partial S}{\partial x_1} + \frac{\partial S}{\partial x_2}.\frac{\partial x_2}{\partial x_1}\right)\frac{\partial x_1}{\partial \pi_1} + \frac{\partial S}{\partial \pi_1} \tag{13}$$

$0 = \frac{\partial S}{\partial \pi_2} = \frac{\partial S}{\partial x_2}.\frac{\partial x_2}{\partial \pi_2} + \frac{\partial S}{\partial \pi_2}$

The solution of (2) to (5) defines the optimal values of π_1 and π_2 (and the implied x_1 and x_2) which the policy maker chooses at the beginning of period 1. At the end of period 1, x_1 and π_1 have occurred, exactly as the model predicted and the policy maker intended, respectively. So at this point, were he to recompute his optimal second period policy π_2, he would simply solve (3) and (5) for π_2 and x_2. Since x_1 and π_1 are as implied by the complete four-equation solution earlier, the truncated solution for π_2 and x_2 conditional on them will also be the same as before (any reader doubtful of this should check this result with any set of four simultaneous equations). So the policy maker will automatically continue with his planned policy under the case of backward-looking expectations.

Now consider how rational expectations affect the situation, under the (artificial) assumption that people believe the policy-maker's announced plans. The model now changes because the outcome in period 1 is affected by expectations of period 2 policy

$$x_1 = x_1(\pi_1, \pi_2) \tag{14}$$

$$x_2 = x_2(x_1, \pi_2) \tag{15}$$

The policy maker's maximization at the start of period 1 is now given by

$$0 = \frac{\partial S}{\partial \pi_1} = \left(\frac{\partial S}{\partial x_1} + \frac{\partial S}{\partial x_2}.\frac{\partial x_2}{\partial x_1}\right)\frac{\partial x_1}{\partial \pi_1} + \frac{\partial S}{\partial \pi_1} \tag{16}$$

$$0 = \frac{\partial S}{\partial \pi_2} = \frac{\partial S}{\partial x_2}.\frac{\partial x_2}{\partial \pi_2} + \frac{\partial S}{\partial \pi_2} + \{\frac{\partial S}{\partial x_1} + \frac{\partial S}{\partial x_2}.\frac{\partial x_2}{\partial x_1}\}\frac{\partial x_1}{\partial \pi_2} \tag{17}$$

Equation (8) is the same as (4) under adaptive expectations, but (9) adds to (5) the last term, which reflects the public's rational anticipation of π_2 and its effect on period 1 events, x_1. The policy maker obviously wants that effect to be optimal.

At the end of period 1, when x_1 and π_1 have become history, the policy maker can re-optimize his setting of π_2 and, if he does so, his appropriate first-order condition is:

$$0 = \frac{\partial S}{\partial \pi_2} = \frac{\partial S}{\partial x_2} \cdot \frac{\partial x_2}{\partial \pi_2} + \frac{\partial S}{\partial \pi_2} \tag{18}$$

which reverts to (5). In other words, having influenced expectations in period 1, he now can neglect any effect of π_2 on the x_1 outcome (just as was the case throughout with adaptive expectations). He solves for π_2, x_2 from (7) and (10), given π_1, x_1 (and previously planned π_2, x_2) from (6) to (9).

Obviously the solution for π_2, x_2 will be different from the previous plan (it would be the same if and only if (10) were identical to (9)). Hence his optimal plan is necessarily time inconsistent.

In macroeconomic policy time inconsistency arises because of the desire among policy makers to exploit the Phillips curve short-run trade-off between output and inflation. If they merely carry forward the money supply plans everyone expects, output will be at the natural rate, y^* (as will unemployment). However, they would ideally prefer higher output (as stressed by Hillier and Malcolmson, 1984) because of the presence of distortions (such as unemployment benefits discussed earlier) which mean that private decisions generate too low an output and employment. So it is optimal to raise money supply growth (and inflation) above its expected level to stimulate output.

Take the following example, using a heavily simplified macro model; the example originally comes from Barro and Gordon (1983). Let the government set money supply growth, g_t, and assume that inflation is equal to this (one can think of this as the government using the money supply to target inflation directly). Output is determined by the Phillips curve:

$$y_t = y^* + \theta(g_t - g_t^e) \tag{19}$$

g_t^e is people's rational expectation of money supply growth formed before the t-period opens.

The government's (and also the social) utility function is:

$$U_t = -(y_t - ky^*)^2 - s(g_t - g^*)^2 \quad (k > 1) \tag{20}$$

$ky^*(> y^*)$ and g^* are the government's ideals for output and inflation, respectively.

Suppose the government announces its planned money growth in $t-1$; it is believed and therefore equals g_t^e. But then it is able to change its mind at the beginning of t, after g_t^e has been set, so that a different g_t can be delivered. The set-up is quite artificial, as we pointed out earlier and will shortly discuss.

Before the period the government maximizes U_t with respect to g_t, subject to the Phillips curve and $g_t = g_t^e$. In this case $y_t = y^*$ and the maximum is:

$$0 = \frac{\partial U_t}{\partial g_t} = -2s(g_t - g^*) \tag{21}$$

whence $g_t = g^*$.

However, once period t has begun, the government can re-optimize, maximizing U_t with respect to g_t again subject the Phillips curve, but now given $g_t^e = g^*$. Substituting for y_t from (12) in (11) gives:

$$U_t = -[(1-k)y^* + \theta(g_t - g^*)]^2 - s(g_t - g^*)^2 \tag{22}$$

Maximizing this we get:

$$0 = \frac{\partial U_t}{\partial g_t} = -2[(1-k)y^* + \theta(g_t - g^*)]\theta - 2s(g_t - g^*) \tag{23}$$

so that

$$g_t = g^* + \frac{\theta(k-1)y^*}{s+\theta^2} \tag{24}$$

Actual money supply growth is more stimulatory than expected as the government exploits the Phillips curve. Figure 5.3 illustrates this.

The concentric ellipses in figure 5.3 are the government's (society's) indifference curves, with B denoting its 'bliss point'. $P_L P_L$, $P_s P_s$ are respectively the long-run Phillips curve and the short-run, conditional on $g_t^e = g^*$ (we will return to $P_s' P_s'$ and point c). Point a is the optimal plan before period t opens. Point b is the optimal action after it opens.

As our heavy earlier hints have indicated, there is an artificiality in the assumptions which permit the government to fool rational private decision makers in this way. Even though it is for the general good, private individuals will try to avoid being fooled in this way, because as individuals that is their best decision. That is precisely what a distortion involves: the outcome of best private plans does not coincide with the social optimum.

Another way of stating this is to say that individuals will, in forming rational expectations, take account of the government's incentive not to carry through its announced plan. If people know the model, then they know the government's utility function and that it will maximize it.

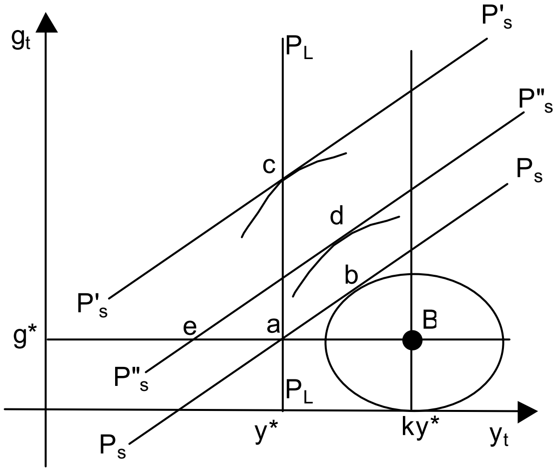

Figure 5.3: Time-inconsistent exploitation of the Phillips curve

They will be in a position to work out what the government will actually do as opposed to what it says it will do.

In our example, this has radical results. People form their expectations, g_t^e, by working back from what the government will do in period t given those expectations. The government will maximize U_t subject to the Phillips curve and g_t^e whatever it is. Substituting (11) into (12) gives:

$$U_t = -[(1-k)y^* + \theta(g_t - g_t^e)]^2 - s[g_t - g^*]^2 \tag{25}$$

and setting $\frac{\partial U_t}{\partial g_t} = 0$ solves for:

$$g_t = \frac{sg * + \theta^2 g_t^e + \theta(k-1)y^*}{s + \theta^2} \tag{26}$$

This is the rule people know the government will actually follow, regardless of its pious intentions. Therefore to form a rational expectation in period $t-1$ people take expectations of (18): in so doing they ensure that $g_t^e = E_{t-1}g_t$. Therefore they solve:

$$E_{t-1}g_t = \frac{sg * + \theta^2 E_{t-1}g_t + \theta(k-1)y^*}{s + \theta^2} \tag{27}$$

from which it emerges that:

$$g_t^e = E_{t-1}g_t = g^* + \frac{\theta(k-1)y^*}{s} \tag{28}$$

When the government comes to execute its optimal rule (18), whatever it may have announced in advance, it will find that it will generate money supply growth exactly as people expected, $g_t = g_t^e$ as in (20). This can be established by substituting (20) for g_t^e into (18) and solving for g_t.

In figure 5.3, this is point c, where $y_t = y^*$ on the long run Phillips curve. Yet the government has reached a tangency point of an indifference curve to the short run Phillips curve, $P_s'P_s'$, through that point.

What we see here is a combination of two factors. First, rational expectations is ensuring that in the absence of shocks people are not surprised so that their private output plans are not frustrated: $y_t = y^*$ on the long-run Phillips curve and $g_t = E_{t-1}g_t$. This is a restoration of Sargent and Wallace's principle of policy ineffectiveness in the context of a standard Phillips curve.

Secondly, the government is pursuing a particularly destabilizing rule for inflation. Instead of responding to past output, itself a rule that destabilizes money growth and inflation in the presence of shocks (but not here where shocks have been assumed away), it is responding to expectations of inflation in the attempt to stimulate output above its natural (inevitable) level. In other words, whatever money supply growth people may expect it to inject in an attempt to stimulate the economy, it goes one better in its attempt to reach tangency with the short-run Phillips curve. The process is only arrested by the escalating costs of inflation. It can be seen from figure 5.3, where the vertical distance between points b and c is given by $(k-1)\frac{\theta^2}{s(s+\theta^2)}$, that as the cost of inflation, s, tends to zero, point c tends to infinity.

Consequently, on top of any moderate destabilization of inflation that may occur if it also attempts to respond to past output, the government may produce a very large destabilization of inflation through this type of response. The solution forced by the conflict between the government's determination to reach tangency with the short-run Phillips curve and the people's rational expectations is the point c where both are finally satisfied.

In this case the 'time-inconsistency' of the government's optimal plans has been marginalized by people's rational expectations, because any 'plans' the government may form before period t opens are irrelevant. The only 'plan' of any interest is the computed outcome of the government's true rule. If we refer to this as the government's true plan, then we can say that it is time consistent, because it will be followed through in period t.

Clearly, this time-consistent result of the interaction between the government's rule and people's rational expectations delivers poor social utility. If somehow the government could achieve point a (the result of its original time-inconsistent optimal plan, had it been carried forward), that would be better, since it would achieve lower inflation together with $y_t = y^*$ as in c. Point b is better still but no longer available because of rational expectations.

In terms of the general two-period problem set out in the model of equations (1), (6) and (7), the equivalent of point a is the solution of equations (8) and (9) together with the model where the government determines its optimal plan assuming its own credibility. The equivalent of point b is the solution of (10) and (7) given x_1 and π_1 as determined by this original optimal plan and the full model — here it 'backslides'. The equivalent of point c is given by the joint solution of (10) and (18) with the model, where its future backsliding is fully taken into account in its decision for first-period policy and in the behaviour of private agents.

We have now exposed the artificiality of our original assumption that people automatically believe the plans the government announced at the start of the first period. In fact they will only believe them if the government commits itself in advance to carry them out, on pain of a penalty at least as great as any gain from backsliding. This penalty has to be administered by some agency outside the government's control: otherwise the government could abrogate the penalty. This incidentally creates a difficulty for a sovereign democratic government: if it is sovereign, may it not refuse to submit to any penalty, even one it agreed to? Clearly if it does not surrender its sovereignty over this matter then it cannot enjoy a point such as a in figure 5.3. Without an effective and sufficient penalty, the only possible point is c: not only b but also a is ruled out. (Point c, the 'time-consistent' policy case in this literature, would in the literature of contracts be called the 'self-enforcing' or 'incentive-compatible' case. In effect if no contract can be enforced against the interests of one party, then the only possible contract is one that this party will carry out anyway — Hart, 1983. The parallel here in policy is that no enforceable contract specifying a penalty can be drawn up with the policy maker.)

Since a is preferable to c, it is in the interest of policy-makers to find a mechanism for effective commitment. Various mechanisms have been suggested. The most widely canvassed is a constitutionally independent central bank such as the Bundesbank (set up after the Second World War) in Germany and the Federal Reserve Board in the USA; we shall explore this carefully below. Another is to join a currency board system of fixed exchange rates (as in the sterling area) whereby one's money supply is controlled by the dominant currency of the system: the European Monetary System evolved into the euro-area partly for this reason. In the case where the government itself controls its own money supply within a floating exchange rate system, the British government since the Second World War being one example, it may seek to create political embarrassment from backsliding.

The Thatcher government's Medium Term Financial Strategy — MTFS — was of this type. It set targets for the money supply and public sector deficit for a rolling four-year period: these targets and the associated commitments to reduce inflation were emphasized by the government as a litmus test of its fitness to govern. The implication was that the public should not re-elect it if it failed: clearly a dangerous commitment to fail on from a political viewpoint.

These mechanisms all have the key deficiency that governments are sovereign and cannot be bound by past commitments; therefore they will naturally re-optimise and whatever hurdles are put in their way they will contrive to leap over them — thus the time-inconsistency problem reappears (McCallum, 1995). Democratic governments for example have the power to change constitutions or to leave currency boards. They can plainly run public sector deficits, which ultimately may drive them to print money to avoid the rising cost of debt finance: central bank obstruction in these circumstances will be over-ruled, with the ultimate threat of removing the bank's independence. The Thatcher MTFS confronted the issue of deficits, and by making the public its judge also avoided the constitutional issue. The question still remained over how far it could get away with backsliding: for example, could it not have tried to persuade the public at an election that somewhat higher inflation then planned was for their good because it allowed lower unemployment (point *b* in figure 5.3)?

This deficiency implies that there is with each of these mechanisms some probability of effectiveness, some of backsliding; let us leave on one side for now where these probabilities might come from and how they might be assessed. The resulting solution will reflect these probabilities. Let us pursue this in the context of our inflation-unemployment trade-off example.

Suppose the government commits incompletely to a plan of $g_t = g^*$ so that there is a probability, π, of its pursing a backsliding policy g_t^b (that is, maximizing U_t given whatever expectations, g_t^e, have been formed on the basis of its 'commitment'). The consequential expectations, g_t^e, will be:

$$g_t^e = \pi g_t^b + (1 - \pi) g^* \tag{29}$$

To calculate g_t^b, we substitute for g_t^e in U_t obtaining

$$U_t = -[(1-k)y^* + \theta(g_t - \pi g_t^b - (1-\pi)g^*)]^2 - s(g_t - g^*)^2 \tag{30}$$

so that its backsliding strategy is given by setting $0 = \frac{\partial U_t}{\partial g_t}$ which yields:

$$g_t^b = g^* + \frac{\theta(k-1)y^*}{s+\theta^2(1-\pi)} \tag{31}$$

(Our procedure, it will be recalled, is to maximize with respect to g_t, given g_t^b, then to set $g_t^b = g_t$ in the resulting first-order condition, and solve for g_t^b.)

It follows that:

$$g_t^e = g^* + \frac{\pi\theta(k-1)y^*}{s+\theta^2(1-\pi)} \tag{32}$$

The resulting situation is illustrated in Figure 5.3 by the Phillips curve $P_s^{"}P_s^{"}$ and is worked out in Box 5.1. The outcome will either be d if the government backslides or e if it does not. Point d is clearly better than c (it may or may not be worse than a, depending on the size of s, the government's dislike of inflation and the extent of credibility in the uncertain case); backsliding when you have uncertain commitment is better than when you have none, but you might have been better off with totally certain commitment.

Box 5.1

THE VIRTUE OF (EVEN SOME) CREDIBILITY

When credibility is total, then

$$g_t = g_t^e = g^*$$

and

$$U_t = -[(1-k)y^*]^2 = U_a$$

i.e. utility in case a in Fig.5.3.
When there is no credibility (case c), we obtain

$$U_c = U_a(1 + \frac{\theta^2}{s})$$

After laborious substitutions we obtain for the backsliding case d under imperfect credibility

$$U_d = U_a(\frac{s(s+\theta^2)}{(s+\theta^2[1-\pi])^2})$$

and in case e (follow-through with g^*),

$$U_e = U_a(\frac{s+\theta^2}{s+\theta^2[1-\pi]})^2$$

Comparing U_e with U_c and U_d with U_a, we obtain in both cases an ambiguous ranking:

$$\frac{U_e}{U_c} = \frac{U_d}{U_a} = \frac{s(s+\theta^2)}{(s+\theta^2[1-\pi])^2}$$

Subtracting the denominator from the numerator yields:

$$\theta^2[s(2\pi-1) - \theta^4(1-\pi)^2]$$

so that $U_e \succ U_c$ and $U_d \succ U_a$ if for example π is high or if $\pi \succ 0.5$ and s is high or θ is low (it requires that $\pi \succ 0.5$).
The probability-weighted average of U_d and U_e is

$$U_w = \pi U_d + (1-\pi)U_e = U_a(\frac{s+\theta^2}{s+\theta^2[1-\pi]})$$

which is worse than U_a and better than U_e.

Point e, following through on uncertain commitment, though clearly worse than a and d, may or may not be greater than c. However the expected (probability-weighted) utility of points d and e is better than c and worse than a. It follows that some commitment, however uncertain, is better than no commitment at all, because it gets the government on to a lower Phillips curve, so a better trade-off: expectations of money supply growth, g_t^e, are reduced.

Uncertainty about whether the government will backslide for some reason is unavoidable given its sovereignty. The essential source of the uncertainty is whether the penalty triggered by backsliding (over-riding the central bank, being politically embarrassed by higher inflation, etc.) will be bigger than the gain from it (lower unemployment, higher output at the expense of somewhat higher inflation).

As we have seen, it pays the government to make a tough commitment to g^*. But no government can bind its successor, or even itself at a future date. If its preferences change (a new government or new ministers), or if circumstances change (an unexpected fall in y^*, for example), the penalty may not be sufficient.

Modelling this uncertainty is clearly difficult. Various attempts have been made. One, initiated by Backus and Driffill (1985 a,b), has focused on uncertainty about government preferences. The government may turn out to be 'wet' or 'dry' (respectively caring about both inflation and unemployment, and caring only about inflation): the probability, π, of backsliding (the effect of 'wetness') is constantly re-assessed in the light of policy actions.

The difficulty with this approach is to motivate, first, these differences in preferences and, second, the private sector's lack of knowledge about them. If we probe more deeply into the basis of politics — we do more on these lines in Chapter 8 below — we can see that parties wish to win votes, by appealing to popular preferences (those of the median voter), and they may also have 'ideology' (that is beliefs about what should be done allied to the belief that they can make this attractive to voters). Possibly we could define a wet government as one that wished to appeal to the median voter; a dry one as one that ideologically favoured zero inflation regardless of short-term popularity. Yet this raises the question: if voters truly reward reflationary time-inconsistency at the polls, would it be rational for a government to be dry, so guaranteeing loss of office?

Box 5.2

UNKNOWN PREFERENCES AND REPUTATION — A SIMPLE EXAMPLE

A government gets elected for a two-period term (a 'period' lasts two years). There are two sorts of politician: 'wet' with utility function (seen from the start of any term of government they might have)

$$U_w = -\frac{1}{2}\pi_t^2 + a(y_t - \overline{y}) + \delta[-\frac{1}{2}\pi_{t+1}^2 + a(y_{t+1} - \overline{y})]$$

and 'dry' with utility function

$$U_d = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\delta\pi_{t+1}^2$$

where in both cases $\delta \succ 1$ (because the second period is closest to the next election). The economy consists of a Phillips Curve

$$y_t = b(\pi_t - \pi_t^e)$$

and π_t^e is formed rationally with information at $t-1$ (the log of the natural rate, y^*, has been normalised at 0). The private sector is assumed to know nothing at all about the politicians' preferences,

so that the initial probability of the government being wet, $prob_w = 0.5$; it then looks for clues from government behaviour.
The dry-politician government would maximize its utility ($\frac{\partial U_d}{\partial \pi_t} = \frac{\partial U_d}{\partial \pi_{t+1}} = 0$) by setting $\pi_t = \pi_{t+1} = 0$. The wet-politician government faces a choice. If it chooses $\pi_t \succ 0$ in t, then the private sector will know in forming period $t+1$ expectations that the government is wet ($prob_w = 1$); if it chooses $\pi_t = 0$, the private sector will still be in the dark and $prob_w = 0.5$ still. in period $t+1$ it no longer matters whether the private sector knows because the period after the election $(t+2)$ does not count: politicians only care about winning the election. Hence maximizing U_w with respect to π_{t+1} yields simply

$$0 = \frac{\partial U_w}{\partial \pi_{t+1}} = -\pi_{t+1} + abor\pi_{t+1} = ab$$

Maximizing U_w with respect to π_t involves choosing the highest of two utilities, cases A and B:
(A) if $\pi_t = 0$: $U_w^A = -ab\pi_t^e - a\overline{y} + \delta[-\frac{1}{2}(ab)^2 + ab(ab - \pi_{t+1}^{e,A}) - a\overline{y}]$
(B) if $\pi_t \succ 0$, then the wet government will choose the π_t that maximizes u_w, that is $0 = \frac{\partial U_w}{\partial \pi_t} = -\pi_t + ab$ so $\pi_t = ab$ (remembering that π_{t+1}^e will be invariant to π_t if $\pi_t \succ 0$.) Its utility will then be:

$$U_w^B = -\frac{1}{2}(ab)^2 + ab(ab - \pi_t^e) - a\overline{y} + \delta[-\frac{1}{2}(ab)^2 + ab(ab - \pi_{t+1}^{e,B}) - a\overline{y}]$$

Subtracting U_w^A from U_w^B gives us:

$$U_w^B - U_w^A = -\frac{1}{2}(ab)^2 + (ab)^2 - \delta ab(\pi_{t+1}^{e,b} - \pi_{t+1}^{e,a})$$

which clearly depends on private expectations formed in t for $t+1$. In case A, people will still not know the government is wet and so $\pi_{t+1}^{e,A} = 0.5ab$; in case B people will know and so $\pi_{t+1}^{e,B} = ab$. So

$$U_w^B - U_w^A = \frac{1}{2}(ab)^2 - \delta ab(\frac{1}{2}ab) = \frac{1}{2}(ab)^2(1-\delta) \prec 0$$

Hence it pays a wet government to pretend it is dry in period t: $\pi_t = 0$, so 'building reputation' – before exploiting its reputation in the run-up to the election with a pre-election reflation. We can now incidentally work out π_t^e. Since people know that a wet government would choose this strategy, of zero inflation in the first period, therefore $\pi_t^e = 0$; both governments would therefore choose zero inflation, leaving people no way of distinguishing when forming expectations for the second period.

This is known as a situation of 'pooling equilibrium' in period t because both types of politician choose the same, 'pooled', outcome. However, the set-up assumes that the dry government is indifferent to being indistinguishable from a wet government in the first period. This is true above because the dry government's utility is unaffected by y_{t+1} and so by π^e_{t+1}.

Yet this is improbable since in winning the next election the dry government would be helped by lower unemployment and so by a lower π^e_{t+1}. We can allow for this by letting its preferences now be:

$$U_d = -\frac{1}{2}\pi_t^2 + \delta[-\frac{1}{2}\pi_{t+1}^2 + \alpha(y_{t+1} - \overline{y})]$$

where $\alpha \prec a$

In this, much more complicated, case the dry government also has a choice: it can choose $\pi_t = 0$ as before, knowing a wet government would match it or it can choose a π_t that is so *negative* in the first period that a wet government would find it too expensive to pretend and match it. If the dry government were to choose the latter strategy, we would have a 'separating equilibrium' where the wet governmentr does not pretend and hence cannot 'build reputation'; instead the private sector knows at once what sort of government it faces.

Consider the possibilities for the dry government:

(i) It chooses $\pi_t = 0$. By so doing it permits a wet government to pretend, as we know from the above steps it would.

(ii) The dry government chooses $\pi_t \prec 0$, such that the wet government would find it not worth pretending. To find out how negative π_t would have to be we check out the critical π_t just above which the wet government gives up pretending, call this π_t^*. π_t^* is such that $U_w^A = U_w^B$

So if $\pi_{wt} = \pi_t^*$, then $\pi_{t+1}^{e,A} = \frac{1}{2}ab$ and

$$U_w^A = -\frac{1}{2}(\pi_t^*)^2 + ab(\pi_t^* - \pi_t^e) - a\overline{y} + \delta[-\frac{1}{2}(ab)^2 + ab(ab - \frac{1}{2}ab) - a\overline{y}]$$

If $\pi_{wt} = ab$, then $\pi_{t+1}^{e,B} = ab$ and

$$U_w^B = -\frac{1}{2}(ab)^2 + ab(ab - \pi_t^e) - a\overline{y} + \delta[-\frac{1}{2}(ab)^2 - a\overline{y}]$$

Setting $U_w^A = U_w^B$ and cancelling terms gives us:

$$\pi_t^{*2} - 2ab\pi_t^* - (\delta - 1)(ab)^2 = 0$$

the unique negative solution of which is: $\pi_t^* = ab(1-\sqrt{\delta})$
We now ask whether it would pay the dry government to choose this, so forcing the wet government to stop pretending, or instead choose $\pi_t = 0$, so allowing it to pretend (for the dry to choose anything in between would be the worst of both worlds since it would suffer from negative inflation as well as permitting pretence). We compare U_d^A $(\pi_t = 0)$ with U_d^* $(\pi_t = \pi_t^*)$, so that $\pi_{t+1}^{e,0} = \frac{1}{2}ab$ and $\pi_{t+1}^{e,*} = 0$. We obtain:

$$U_d^* - U_d^A = \frac{1}{2}(ab)^2[\frac{\delta\alpha}{a} - (1-\sqrt{\delta})^2]$$

Considering this problem we can see at once that we require $\delta \succ 1$ and $\alpha \succ 0$ for a separating equilibrium to be possible. $\delta = 1$ implies $\pi_t^* = 0$; $\alpha = 0$ implies that $U_d^* \succ U_d^A$. Fig. 5.4 below shows the conditions (for $\delta \succ 1$) under which we get a separating equilibrium: where $\frac{\alpha}{a} \succ \frac{(1-\sqrt{\delta})^2}{\delta}$. For example, assuming $\delta = 2$ (the pre-election performance is twice as important for the election as the post-election one), then α needs to be only 9% of a; so the dry government needs to care about pre-election unemployment only a very little relative to a wet one.

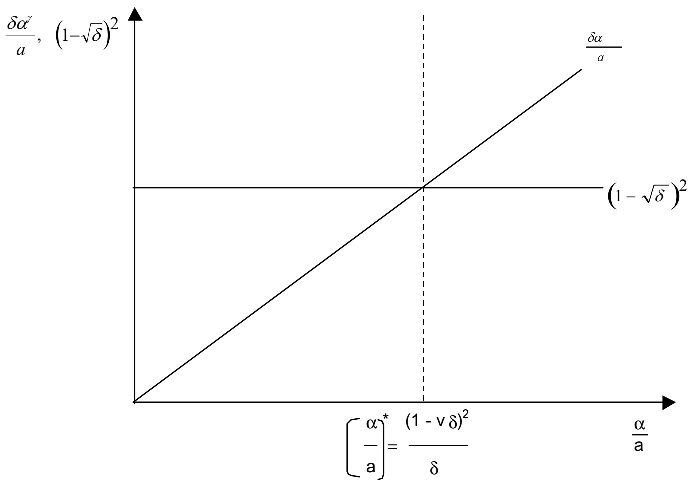

Figure 5.4: The Backus–Driffill model

As for revelation of preferences, Vickers(1986) showed that in many circumstances it would pay the dry government to signal its nature by some extreme act of dryness that a wet government would find too costly to imitate. In this case there would be a 'separating equilibrium' in which wage bargainers would know exactly what sort of government they faced. This question is explored in the example in Box 5.2.

Another approach, initiated by Barro and Gordon (1983), assumes that preferences are known and constant and that the penalty for backsliding is a 'loss of reputation', whereby the private sector (wage bargainers) disbelieve the government's announced plans for a certain period of time and assume the government will backslide during this period. The penalty is arbitrary in the Barro and Gordon model but we could reduce the arbitrariness by assuming for example that private disbelief continues until there is a change of membership in the government (e.g. an election cycle).

During the period of lost reputation, the government is considered to suffer the point c. This penalty is in general less than the gain to backsliding on g^* as an announced plan. People work out the announced plan ($g_t^p > g^*$) such that the penalty is just sufficient to offset any gain from backsliding. This is of course the plan that the goverment will also choose. It produces a point on the vertical Phillips curve between a and c. This is worked out in Box 5.3.

Box 5.3

REPUTATION THROUGH PRIVATE SECTOR WITHDRAWAL OF BELIEF

The government has an infinite-horizon utility function at t:

$$U_t = \sum_{i=0}^{\infty} \delta^i \{-\frac{1}{2}\pi_{t+i}^2 + a(y_{t+i} - \overline{y})\} \qquad \delta \prec 1$$

and faces a Phillips Curve

$$y_t = b(\pi_t - \pi_t^e)$$

π_t^e is formed rationally at $t-1$ for t according to the following agreed private sector one-period punishment strategy:

(1) if $\pi_{t-1} \succ \pi_{t-1}^e, \pi_t^e = ab$ (the government's discretionary chocie)

(2) if $\pi_{t-1} = \pi_{t-1}^e, \pi_t^e = \pi^a$, acredible government target announced in $t-1$ for t, t for $t+1$ etc (provisionally assumed constant).

Under (1) the government 'forfeits credibility', under (2) it 'achieves credibility'.
We start this game off by assuming that $\pi_{-1} = \pi^e_{-1}$.
Knowing the private strategy the government chooses a plan at -1 for $t \geq 0$. It has two basic strategies to compare:
(a) It can announce and deliver π^a throughout – a permanent follow-through strategy.
(b) It can cheat in even periods and suffer punishment in odd periods (cheating in odd ones must be worse than this because the gains are delayed). When it cheats it will choose $\pi = ab$ because this maximizes its (cheating) utility as we have seen.
Under (a) its utility is:

$$U^a_t = \frac{1}{1-\delta}[-\frac{1}{2}(\pi_a)^2 - a\overline{y}]$$

Under (b) it is:

$$U^b_t = \sum_{i(even)}^{\infty} \delta^i\{-\frac{1}{2}(ab)^2 + ab(ab-\pi_a) - a\overline{y}\} + \sum_{i(odd)}^{\infty} \delta^i\{-\frac{1}{2}(ab)^2 - a\overline{y}\}$$

Remembering that

$$\sum_{i(even)}^{\infty} \delta^i = 1 + \delta^2 + \delta^4 + ... = \sum_{i}^{\infty} \delta^{2i} = \frac{1}{1-\delta^2}$$

and analogously

$$\sum_{i(odd)}^{\infty} \delta^i = \frac{\delta}{1-\delta^2}$$

we obtain

$$U^b_t = \frac{1}{1-\delta^2}[\frac{1}{2}(ab)^2(1-\delta) - (1+\delta)a\overline{y} - ab\pi^a]$$

whence $\left(\text{remembering } \frac{1}{1-\delta^2} = \frac{1}{(1-\delta)(1+\delta)}\right)$

$$U^a_t - U^b_t = -\frac{1}{2(1-\delta)}\{(\pi^a)^2 - \frac{2ab}{1+\delta}\pi^a + \frac{(1-\delta)}{(1+\delta)}(ab)^2\}$$

The situation is illustrated by Figure 5.5 which shows U^a and U^b as functions of π^a, the credible target. $U^a = U^b$ at two points: $\pi^a = \frac{ab(1-\delta)}{1+\delta}$; ab (these are found as the solutions to the quadratic in

π^a that comes from setting $U^a = U^b$). Any point to the left of the first is not credible because it would pay the government to cheat. Any point between the first and second provides less utility than the first and so will not be chosen by the government. Therefore the first point, $\frac{ab(1-\delta)}{1+\delta}$, is the only feasible one: at it the government will never cheat, because the loss of utility from losing credibility next period just offsets the gain from cheating. Plainly π^a is a constant as assumed (as can be checked by seeing that this holds in every period).

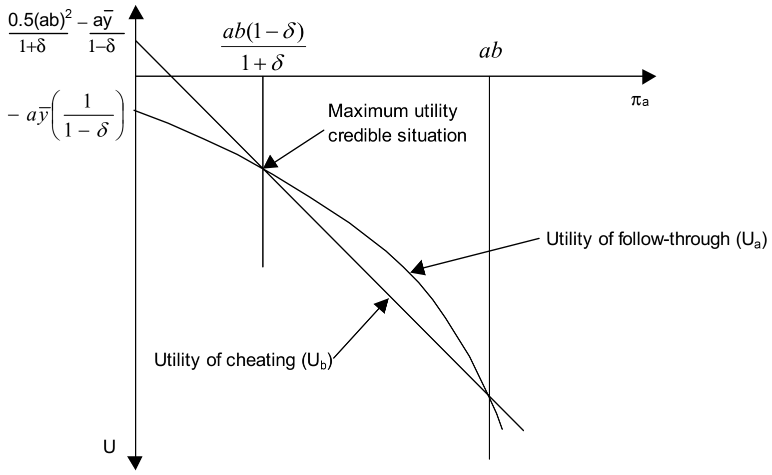

Figure 5.5: The Barro-Gordon model of credibility

This idea on its own simply works out what would happen if the penalty for backsliding on a plan were less than the penalty needed to prevent backsliding on g^*. The result is a steady $g_t^e = g_t > g^*$. The government would in fact be better off if it could somehow persuade people to disbelieve it for longer!

Uncertainty of a sort can be introduced by introducing shocks to the economy's natural rate or to the government's preferences which cause the government to choose a new g_t^e equilibrium, after initially reneging. As we have seen it will always pay the government to choose a permanent non-punishment equilibrium as soon as it can. If $g_t > g_t^e$, the government loses credibility: future g_t defaults to point c until credibility is restored. It then reverts to $g_t = g_t^e > g^*$ until the next shock producing $g_t > g_t^e$.

The difficulty with this approach is how to define and motivate a coordinated, non-arbitrary strategy on the part of private agents for 'triggering' belief or disbelief — their 'punishment strategy'. These agents, it must be remembered, are not voters here, but wage bargainers. Why should they collectively come together to agree a punishment strategy instead of getting on with their own private business? One can think of reasons if workers are represented by a single union or a well-integrated set of unions; and this case has been explored by Herrendorf and Neumann (1998). But in general the case is unconvincing.

We have therefore reviewed two approaches to uncertainty — one emphasizes the uncertainty of government preferences, the other that of the changing environment. Both are attempts to model the evolution of credibility, viewed as the possible selection by the government of either low inflation or reflationary policy. However both have arbitrary elements in them which make them unconvincing in general as an account of why governments choose to inflate or not.

More recently economists have returned to the original Kydland-Prescott model and attempted to embed it in a rational political process. There have been two main strands of thought.

The first is an institutional model of politics. In this approach it is assumed that institutions matter and do in practice constrain governments, including democratic ones. So it is accepted here that governments will be prone to time-inconsistency, implicitly because this will appeal to voters whose preferences are such as in the Barro-Gordon set-up. But there are difficulties for governments in overturning institutional barriers — they are for example seen as 'high-handed' or 'arbitrary' and therefore governments like to be seen as acting within the conventions; even if they could overthrow them, they will not do so lightly because it requires costly legislative time and ministerial energy.

These authors (the earliest of whom is Rogoff, 1985) suggest that the electorate may well accept some institutional reform, if it can be shown to have good long-term effects. Then the reform once made it will be respected by future governments. They suggest therefore establishing the institution of an independent central bank, with a remit (or 'institutionalized preferences' which could be summarised in the degree of 'inflation toughness' of the Governor) to give more attention to inflation than would be given by politicians reflecting popular preferences.

The earlier work gave no real role to stabilization since there were no shocks shifting the Phillips curve itself; only demand shocks introduced by monetary policy and inflation and causing output to move up and down the Curve. In this revival of the idea of Bank independence Phillips curve shocks (which could come from supply or from demand via other elements entering the curve, such as the exchange rate or real interest rates driving inter-temporal substitution) were introduced so that the government has another reason for varying inflation besides the desire to raise output above the natural rate: to stabilize the economy in the face of shocks.

The resulting set-up is illustrated in figure 5.6 (the same as figure 5.5 above, except that the Phillips curve is now stochastic). The voter, as represented by the government, has 'normal' preferences with a dislike of both inflation and unemployment. Unrestrained, these would produce the poor time-consistent outcome c on our original figure — high average inflation; however it would also stabilize shocks as shown at points f and g. Reflecting on their institutions however voters realize that this systematic result of the political process is not as good as could be achieved if institutional restraint was exercised through an independent central bank that was free to react to shocks and yet limited in its tolerance of inflation. (One might ask here why not look at other constitutional restraints? The answer would be that a purely constitutional law would lack the necessary flexibility Since we are setting the institution up for the long term we are concerned about its effect on expected (average) future social (i.e. the voters') utility. We vary a until this is maximized; it turns out (see Box 5.4) that this occurs somewhere between 0 (total intolerance of inflation) and popular preferences. Inflation would be zero at the first but stabilization would also be nil; at the second the inflation bias would be at its worst but stabilization would be at its most effective. A compromise is necessary.

Box: 5.4

THE INSTITUTIONAL APPROACH TO THE INDEPENDENT CENTRAL BANK

Suppose society has the utility function:

$$U_s = -\frac{1}{2}[\pi_t^2 + \alpha(y_t - \overline{y})^2]$$

and faces a Phillips Curve

$$y_t = c(\pi_t - \pi_t^e) + u_t$$

where u_t consists of shocks on the Phillips curve from supply and also demand (via real interest rates or exchange rates which enter the inflation/labour supply process) with variance σ^2. The institutional reformer wishes to choose a central bank set of 'preferences' (shorthand for the remit given in the central bank legislation) in order to mazimize society's expected utility (i.e. the average over the future – long – period during which the legislation will hold), that is

$$EU_s = -\frac{1}{2}[E\pi_t^2 + \alpha E(y_t - \overline{y})^2]$$

To motivate the set-up consider two extreme instructions to the central bank which we summarise as giving the central bank a set of preferences of the form:
(i) not to care at all about unemployment ($a = 0$)
(ii)to behave as if it had the same preferences as society ($a = \alpha$)
Using the central bank's first order conditions,
under (i) $\pi_t = 0 = \pi_t^e$ and so $EU_s = -\frac{1}{2}\alpha[\overline{y}^2 + \sigma^2]$
under (ii) $\pi_t = \frac{\alpha c^2 \pi_t^e + \alpha c\overline{y} - \alpha c u_t}{1+\alpha c^2}$. It follows that $\pi_t^e = ac\overline{y}$ so that $\pi_t = \alpha c\overline{y} - \frac{\alpha c}{1+\alpha c^2}u_t$ and $y_t = \frac{1}{1+\alpha c^2}u_t$ so $EU_s = -\frac{1}{2}\{\overline{y}^2(\alpha + \alpha^2 c^2) + (\frac{\alpha}{1+\alpha c^2})\sigma^2\}$
Compared with case (i), case (ii) offers a clear reduction in the effects of shocks due to the Bank's active stabilisation response, but a rise in the effect of $\overline{y}$ (the 'distortionary gap' between desired full employment output and its natural rate), due to the Bank's frustrated attempt to stimulate output beyond its natural rate; this latter being the effect on expected utility of the inflation bias.
We can now generalise the problem and find a compromise Bank preference for lower unemployment, letting $0 \prec a \prec \alpha$
This will produce analogously from the Bank's first order condition

$$\begin{aligned} \pi_t &= ac\overline{y} - \frac{ac}{1+ac^2}u_t \\ y_t &= \frac{1}{1+ac^2}u_t \end{aligned} \qquad (33)$$

and society will enjoy

$$EU_s = -\frac{1}{2}\left\{[\alpha + a^2c^2]\overline{y}^2 + \frac{\alpha + a^2c^2}{(1+ac^2)^2}\sigma^2\right\}$$

From this we can construct figure 5.7 where we show how EU_s varies with a. As $a \to \alpha$, the effect of the shocks falls as stabilisation rises but the inflation bias rises too. The two together reach a maximum at $\widehat{a}$. From $0 = \frac{\partial EU_s}{\partial a}$ we find the equation for $\widehat{a}$ is:

$$\frac{\widehat{a} - \alpha}{\widehat{a}(1+\widehat{a}c^2)^3} = \frac{\overline{y}^2}{\sigma^2}$$

which reveals that $\widehat{a}$ rises as σ^2 rises relative to $\overline{y}^2$. The reason is that the bigger the shocks are relative to the distortionary gap, the more the gain to stabilizing relative to the cost in higher inflation bias.

All this assumes the central bank's remit must be to maximize something of the same form as society's utility. But of course institutional design need not be restricted to this. Walsh (1995) pointed out that one could add a linear term to u_B; selected so that when this was maximized, inflation bias could be eliminated but optimal stabilization could be achieved. Let $U_B = -\frac{1}{2}(\pi_t^2 + \alpha[y_t - \overline{y}]^2) - \lambda\pi_t$ and set $\lambda = \alpha c\overline{y}$. The first order conditions now yield:

$$\begin{aligned} \pi_t &= -\frac{\alpha c}{1+\alpha c^2}u_t \\ y_t &= \frac{1}{1+\alpha c^2}u_t \\ EU_s &= -\frac{1}{2}[\alpha\overline{y}^2 + \frac{\alpha}{1+\alpha c^2}\sigma^2] \end{aligned}$$

which represents the first best result for society (shown as F on the diagram).

Some other ways of getting the same result are to adjust U_B by inserting an inflation target, π^*, that is artificially low or by setting a side-condition that inflation not be varied from its proper target except in response to shocks. These results can be extended to the more realistic case where output is persistent; in this case the optimal contract contains state-contingent restrictions — see the Appendix to chapter 11 on dynamic programming.

Recently it has been realized that institutional preferences[3] can be

[3] A number of authors have also considered objective functions that are non-quadratic in inflation and output deviations. For instance Orphanides and Wilcox (1996) consider the opportunistic approach of the central banker. Proponents of this approach hold that when inflation is moderate but still above the long-run objective, the Central Banker should not take deliberate anti-inflation action, but rather should wait for external circumstances – such as favorable supply shocks and unforeseen recessions – to deliver the desired reduction in inflation. While waiting for such circumstances to arise, the Central Banker should aggressively resist incipient increases in inflation. The policy maker who is endowed with these preferences gives higher weight to stabilizing output when inflation is low, but higher weight to inflation when inflation is above target. Such an approach is motivated by the views expressed by leading Central Bankers. For instance in testimony before the Senate committee that

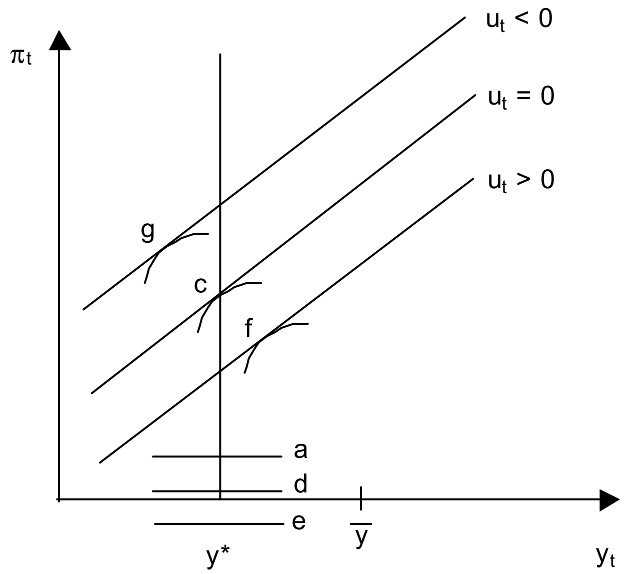

Figure 5.6: A tough central bank (a, d, e) and one with the same preferences as society (g, c, f)

rewritten so that the central bank delivers both zero inflation and maximum stability. In principle this is easily done (Walsh, 1995) by adding a

---

was meeting to consider his nomination to the Federal Reserve Board, former Vice Chairman Blinder summarized his views on this issue as follows:

'If monetary policy is used to cut our losses on the inflation front when luck runs against us, and pocket the gains when good fortune runs our way, we can continue to chip way at the already-low inflation rate.' [Blinder, 1994, p.4].

Nobay and Peel (1998) consider the case where Central Bankers have an asymmetric objective function. This is motivated in part by the European Central Bank's announced objectives according to which inflation should not be above 2 per cent, but is allowed to be below that level.

Specification of such asymmetric preferences leads to terms in the variance of inflation or output impacting on expected inflation. In particular there is a deflationary bias in the ECB example. Essentially the Central Bank has to shoot for a lower expected rate of inflation to allow for the conditional variance of shocks.

Other authors (see e.g. Beetsma and Jensen, 1997; Briault, Haldane and King, 1996) have explored the implications of assuming that the Central Banker's preferences, though quadratic, are uncertain. In particular the weight on say inflation has a stochastic component. Such specifications have implications for expected inflation and the efficacy of inflation targets as opposed to inflation taxes (Walsh, 1995) from a welfare perspective.

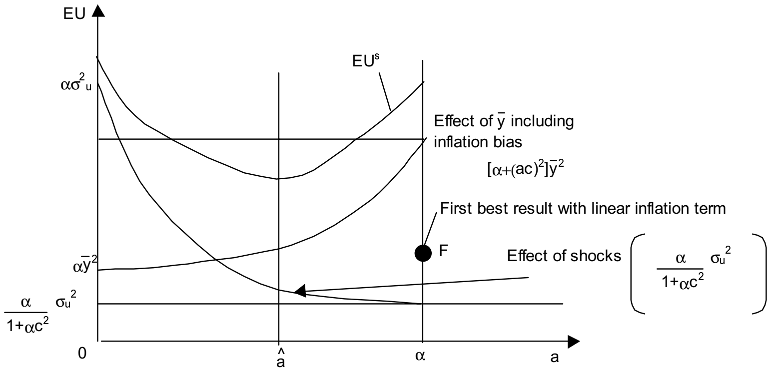

Figure 5.7: Choosing the optimal central bank preference parameter (a)

linear inflation target into the bank's preferences (it can be thought of as forcing an extra linear penalty on inflation); this can then be chosen to force the inflation bias to zero while the a parameter is left at the social preference value so guaranteeing that stabilization is maximized. There are other ways that this can be expressed such as telling the bank to stabilize the economy optimally subject to average inflation (i.e. inflation in the absence of shocks) being equal to target.

This is the institutionalist view. However it is challenged by those who argue that it only artificially 'solves' the time-inconsistency problem (McCallum, 1995). It relies on institutional restraint to prevent governments doing what suits them. Yet suppose that a government is restrained on one occasion by the cost of overriding the constitution; will it not find a way of doing it somehow at less cost in the future? Popular opinion will be forever straining at the leash demanding that its politicians deliver lower unemployment by a 'little bit of inflation'. When politicians are rewarded by popularity will they not in the end do whatever is necessary to gain it?

The institutionalists reply that popular opinion is also in favour of the reform when 'in institutional mode'. Yet this is inconsistent; popular opinion cannot be one thing in one mode and another thing when voting in elections. Either it wants restraint or it does not; it cannot do both at the same time.

One is driven to reject institutionalism as superficial, and failing to address the problem at its roots. The Institution is like the *deus ex machina* with which the Greek dramatists often ended their plays, just as a convenient device to end them but without inherent justification.

Greek audiences accepted it because they understood what it stood for: viz a way of giving the audience the moral of the story (delivered by the god). We have to probe deeper to see if a similar basis can be found for the institutional story.

To find one, we go to political economy and ask whether it could provide a solution. This second strand of thought is a fairly recent reaction to the institutionalist approach (Minford, 1995; Romer and Romer, 1997). Institutionalism has produced some empirical evidence in its favour (for example Cukierman, 1992), showing a correlation between the degree of Bank independence and the rate of inflation. This can be criticized on two levels. First, the definitions of independence are themselves influenced (implicitly if not explicitly) by success in keeping inflation down; the degree of 'toughness' (independence) of a central bank is necessarily evidenced by inflation itself. The example can be given here of the Bundesbank which always scores the most highly on measures of independence; however in strict terms it has no real independence as it could be nationalised by a simple majority of the lower house of parliament (the Bundestag). Its power comes, as board members will frequently say, from its support in public opinion. Second, correlation cannot prove causation here; it may be that as public opinion becomes disenchanted with inflationary policies, it decides to 'institutionalize' this by setting up an independent central bank.

The revisionist view is precisely this: that the electorate begins to understand the time-inconsistency problem as a potential source of inflation, much as an alcoholic understands that his short-term giving-in to the craving for the bottle is the cause of his disease. Wishing to ensure that inflation is kept low on average in spite of the desirability of low unemployment and of stabilizing the economy, it redesigns its own voting reflexes to reward politicians appropriately. It will henceforth penalize 'reflation' whose aim is to stimulate the economy beyond its natural rate; but it will reward stabilizing monetary policy, defined as that which, while reacting to shocks in a stabilising manner, does not on average drive output beyond the natural rate. The electorate does this because it recognises that its 'preferences' (as expressed in these rewards and penalties) are the root cause of political behaviour.

This view requires that the electorate is coordinated in some way so that its punishment strategy can be formulated and exercised. But this is not difficult to understand. We speak of the anti-inflationary 'climate of opinion' in Germany, or in the USA and Britain since the 1970s; what we mean is the popular reflexes that, no doubt heavily influenced by 'informed opinion', underlie voting behaviour.

Furthermore this view leads us to the practical case for an independent central bank: that it is an efficient low-cost way of organising the execution of monetary policy in line with such a public opinion of the problem — just as having independent courts is for the law, or privately-run (privatised) companies is for business. We can say that it does 'institutionalize' the solution effectively. In other words it tells us how we can make sense of the institutionalist case which in itself solves nothing. The solution of the time-inconsistency problem is at the basic level of public understanding, just as that of alcoholism is at the level of personal understanding. Without understanding no solution is possible. With it any set-up for monetary policy would in principle work (just as an alcoholic has many means in principle for stopping drinking to destruction); but a central bank is a convenient method in many ways (just as joining a group like Alcoholics Anonymous is for alcoholism). Empirically it makes sense of the wide variety of institutions under which countries have had both low and high inflation (Romer and Romer, 1997); an independent central bank is neither necessary nor sufficient for low inflation. It merely happens to be convenient for it in a democracy.

CONCLUSIONS

We began this chapter by asking whether there was any case in principle for the government to intervene in the business cycle via variations in interest rates or money supply growth. After all private agents have access to the same information as the government and their decisions would be expected, in the absence of micro distortions, to reflect social as well as private costs and benefits. However we identified the crucial distortions in the labour market (notably the presence of unemployment benefit) that made unemployment non-reflective of social costs; since the business cycle will under these conditions give more unemployment on the downside than less unemployment on the upside, greater cycle variability will produce a higher average distortion. Hence successful stabilization policy would reduce damaging unemployment on average.

We then set out the problem of time inconsistency as a cause of excessive inflation; the point being that governments, having started out with a sensible aim for low inflation, will have an incentive to stimulate the economy further than this in order to bring unemployment down closer to full employment. People, knowing this incentive, then react by raising their inflation expectations to the point where governments just have no further incentive to inflate; this of course means that inflation is higher than the target while at the same time the economy does no better on unemployment — 'stagflation'. Various models of inflation 'credibility' have been built on the idea that private agents may react to past behaviour in forming views of whether a government will in the future succumb to the temptation to 'reflate'; these models are attempts to see whether the natural operation of private markets and reputation can reduce this governmental tendency. However these models are unsatisfactory in a variety of ways and economists have turned to other ways of resolving the problem. Many economists and policy-makers called for the power over monetary policy to be either withdrawn totally from government (for example by a constitutional clause or by fixing the exchange rate to an external monetary anchor) or delegated to an independent central bank with a strictly limited mandate for stabilization and an overriding obligation to control inflation. This of course means that the power of monetary policy to stabilize is largely jettisoned. Some institutionalists have tried to deal with this problem by arguing that politics will generate institutions, most naturally an independent central bank, to restrain politicians' reflation instincts (as encouraged by popular opinion) while still allowing the flexibility to deal with shocks. The Bank's charter can be tuned to deliver the first-best combination of low inflation and output stability. However this is only a solution of the time-inconsistency problem in a superficial sense because it can only work if public opinion tolerates it and does not clamour for an override when reflation is attractive. We argued that the true solution of the problem lies in public opinion itself: understanding how it must reward and penalise political behaviour over monetary policy. Given such understanding, a convenient but in no way essential way to organise monetary policy is then via an independent central bank accountable to public opinion for generating low inflation and output stability.