

Issues in the Design of Fiscal Policy Rules

Jonathan Portes* and Simon Wren-Lewis**

Abstract

Theory suggests that government should as far as possible smooth taxes and its recurrent consumption spending, which means that government debt should act as a shock absorber, and any planned adjustments in debt should be gradual. This suggests that operational targets for governments (e.g. for 5 years ahead) should involve deficits rather than debt, because such rules will be more robust to shocks. Beyond that, fiscal rules need to reflect the constraints on monetary policy, and the extent to which governments are subject to deficit bias. Fiscal rules for countries in a monetary union or fixed exchange rate regime need to include a strong countercyclical element. Fiscal rules should also contain a 'knock out' if interest rates hit the zero lower bound: in that case the fiscal and monetary authorities should cooperate to formulate a fiscal expansion package that allows interest rates to rise above this bound. In more normal times, the design of fiscal policy rules is likely to depend on the extent to which governments are subject to deficit bias, and the effectiveness of any national fiscal council. For example, governments that had not shown a history of deficit bias could aim to target deficits five years ahead (rolling targets), and these would not require cyclical adjustment. In contrast, governments that were more prone to bias could target a cyclically adjusted deficit at the end of their expected period of office. In both cases fiscal councils would have an important role to play, in ensuring plans were implemented in the first case and allowing for departures from target when external shocks occurred in the second.

* Director of the National Institute of Economic and Social Research, London
(j.portes@niesr.ac.uk)

** Professor of Economics and Fellow at Merton College, University of Oxford
(simon.wren-lewis@economics.ox.ac.uk)

Section 1 Introduction

Coen Teulings, the director of the CPB (the Netherlands Bureau for Economic Policy Analysis) – the world’s first fiscal council – observed¹ that there appeared to be a divergence in the views of economists compared to politicians about the role of fiscal policy following the financial crisis. One of the reasons he gave for this was there was no equivalent for fiscal policy of the Taylor rule for monetary policy: no simple rule to guide fiscal policymakers.

This paper is about the search for such a rule. However it is also why one single simple rule to guide fiscal policy may never be found. We highlight two reasons for this. First, basic theory suggests that fiscal policy actions should be very different when monetary policy is constrained in a fundamental way, while the reverse is not in general the case. There are two major examples of where this will be true. The first is when interest rates are at the zero lower bound. We discuss reasons why such a situation cannot be treated as a short-term deviation from an invariant fiscal rule. The second is where a country is part of a monetary union or a fixed exchange rate regime. The importance of aggressive countercyclical fiscal policy in a boom, highlighted by events in the Euro area before 2007, suggests that fiscal rules will be rather different within a monetary union, or with fixed rather than floating exchange rates.

The second reason why a fiscal equivalent of a Taylor rule may be elusive also reflects national differences, but in this case differences in political structure. While it is natural for economists to focus on fiscal rules that approximate optimal policy, the reason we need explicit rules is often to prevent non-benevolent behaviour by politicians. There are many potential reasons for what has been labelled ‘deficit bias’, which may vary depending on the institutional structure. The issue usually highlighted in discussion of monetary policy – time inconsistency – is just one of the factors that might generate deficit bias in the operation of fiscal policy. We discuss a number of areas where there may be a trade-off between rules that mimic optimal policy, and rules that are effective in countering deficit bias.

We begin in section 2 by looking at factors that are likely to be common across all monetary and political regimes. A clear implication from theory is that the government should as far as possible smooth taxes and recurrent consumption spending, which means that government debt should act as a shock absorber, and any planned adjustments in debt should be gradual. However the implications for the long run debt target are less clear. Section 3 looks at alternative reasons for deficit bias, and the extent to which there is a conflict between designing rules that encourage optimal policy, and rules that effectively discourage deficit bias.

Section 4 looks at specific types of rule that have been suggested in the literature or implemented by governments in the light of the considerations examined in previous sections. Should rules relate to government debt or government deficits? If rules are based on deficits, should

¹Teulings, C (2012) Why politicians prefer austerity to long-term fiscal reform:

http://www.mckinsey.com/features/government_designed_for_new_times/why_politicians_prefer_austerity_to_long-term_fiscal_reform

these involve cyclical adjustment? Is it appropriate for rules to be forward looking or date specific? Does it make sense to distinguish between consumption and investment spending? Section 5 identifies the reasons why optimal fiscal policy in the shorter term is crucially dependent on whether there are constraints on monetary stabilisation policy.

In many countries fiscal rules operate alongside fiscal councils: independent government sponsored institutions with some role in the budget planning process. Section 6 looks at whether these institutions might have some influence on the form of fiscal rule adopted. Section 7 concludes by suggesting what form actual rules might take in different types of economy.

Section 2. Optimal debt policy

Government debt, which is a stock concept, does not directly influence social welfare. Instead it is the things that influence debt that matter, like taxes and government spending. This leads to perhaps the most fundamental and most well known principle behind optimal debt policy, which is sometimes called tax smoothing. We can illustrate it in a somewhat stylised way as follows. Suppose social welfare declines as taxes rise, because taxes are distortionary. Suppose the social costs of a unit increase in taxes increases with the level of tax rates: raising taxes by 1% when taxes are already at 50% is worse than if taxes are 20%. To simplify the maths, we can capture this by making the loss in social welfare a quadratic term in the level of the tax rate (τ). The government would like to minimise these costs, but they need taxes to pay for government spending (g) and any interest on debt (d), where both are measured as a ratio to GDP.

We can represent the government's problem as

Minimise

$$\sum_{t=0}^{\infty} \beta^t [\tau_t^2 + 2\mu_t(d_t - (1+r)d_{t-1} - g + \tau_t)]$$

where β is a discount factor and r is the growth corrected interest rate (i.e. the real interest rate less the growth rate of real GDP). μ is a Lagrange multiplier capturing the assumption that the government must satisfy its budget constraint (there is no default), and we ignore financing through printing money. For simplicity we assume that government spending is not a choice variable, but the points discussed below easily generalise to when it is, as long as the costs of deviating from the optimal provision of public goods are also convex (e.g. quadratic).

The first order conditions for this optimisation problem are

$$\tau_t = \mu_t$$

$$\mu_t - \beta(1+r)\mu_{t+1} = 0$$

which can be combined as

$$\tau_t = \beta(1+r)\tau_{t+1}$$

together with an initial condition which is the inherited level of debt. The term combining the discount factor and the rate of interest, $\beta(1+r)$, is likely to be close to one. That in turn implies that, even if government spending changes substantially, then as long as such a change is temporary, tax rates adjust slowly over time, which is the tax smoothing idea. So, for example, if for a period government spending has to be unusually high (classically a war, but also perhaps because of a recession or natural disaster), it would be wrong to try and match this higher spending with higher tax rates. Instead taxes should only be raised by a small amount, with debt increasing instead, but

taxes should stay high after government spending has come back down, to at least pay the interest on the extra debt and perhaps also to bring debt back down again. Similarly, of course, a 'windfall gain', in the form of a temporary reduction in spending or rise in revenue, should be smoothed via a permanent but small reduction in tax rates.

Note that this smoothing principle has not required any assumptions about the macroeconomic impact of fiscal policy on demand, but simply the idea that the costs of taxation are increasing in the level of taxes. In addition, there is no distinction between government spending on consumption and investment here. So the Keynesian argument that debt and deficits should be shock absorbers is in addition to these tax smoothing arguments. Shocks to aggregate demand, for given levels of government spending and tax rates, are likely to lead to movements in government deficits which help dampen the impact of shocks: this is the automatic stabiliser. Attempts to target the deficit will amplify the impact of the shocks by switching off these automatic stabilisers. It is worth noting that the more consumers are able to smooth income shocks themselves, the less important these automatic stabilisers are likely to be.

The first implication of smoothing is that government debt should be a 'shock absorber', and that any adjustment in debt should be slow, to avoid sharp movements in taxes or government spending (see Kopits and Symansky, 1998). But what level should government debt adjust towards, albeit slowly? Here the exact size of the term $\beta(1+r)$ appears crucial. If $\beta(1+r) < 1$, this means that the rate at which we discount the cost of future taxes is high relative to the rate of interest, so there is no incentive for the government to control debt at all. The optimal thing for governments to do is to allow debt to gradually explode. This is not a case that we will consider further, except to note that it illustrates why an impatient (relative to society) government might cause difficulties.

If $\beta(1+r) = 1$, which is the case in many benchmark macroeconomic models (where agents are assumed to care about their children such that they act as if they live forever), then taxes stay constant over the whole planning horizon. The only way this can happen is if debt also stays constant at its inherited level. There is no long run debt target. The intuition is as follows. If we raised taxes today to reduce debt, this would have short run costs, but long run benefits. There is a trade-off. Barro's tax smoothing hypothesis tells us that when costs are increasing at the margin, the best thing to do is smooth costs over time. So the cost of raising taxes today will always exceed the discounted benefits of lower debt and taxes tomorrow, even though those benefits are permanent. This is sometimes called the "Random Walk Steady State Debt" result, because it implies that debt should accommodate shocks to it, rather than revert to some target level.

If $\beta(1+r) > 1$, then we get a very different and equally surprising result. Taxes gradually fall over time, until they eventually decline to zero. How can this happen, given that the government has spending to finance? The answer is that debt gradually declines to zero, and then the government starts to build up assets. Eventually it has enough assets that it can finance all its spending from the interest on those assets, and so taxes can be completely eliminated.

Both these conclusions about the long run debt target appear rather extreme, and very different from each other. Actually the difference between them is more apparent than real. The $\beta(1+r) = 1$ case is really an extreme case of tax smoothing, where the adjustment of debt is

infinitely slow, so we never adjust to any target. If $\beta(1+r) - 1$ was positive but extremely small, we would adjust debt downwards but at a very slow rate, and this might appear in the short term to be little different to the random walk steady state case. This emphasises that the key, and robust, result of this analysis is that debt adjustment should be slow.

In the analysis above we simply postulated quadratic costs due to the distortionary effects of taxes. More elaborate models have 'microfounded' these costs, most recently within the context of Keynesian models where prices are sticky, but labour supply is endogenous. Early examples are Benigno and Woodford (2004) and Schmitt-Grohe and Uribe (2004), and there is a more general discussion in Kirsanova, Leith and Wren-Lewis (2009). See also Marcet and Scott (2008). The tax smoothing hypothesis is originally due to Barro (1979).

How seriously should we take the case where $\beta(1+r) = 1$, implying there should be no debt target? There are a number of reasons for thinking that in most economies most of the time the rate of interest exceeds the discount rate. Equality comes from the workhorse intertemporal macro model, where agents care about their children as much as themselves, and so in effect live forever. If agents are less altruistic, or inheritance tax prevents them from behaving in this way, then the economy will generate a real interest rate which exceeds agent's impatience, so $(1+r)\beta > 1$. Equally we might have ethical objections to discounting the utility of future generations in this way, which the Stern review on climate change discussed in detail (Stern, N. (2006)), so $\beta = 1$.

Another reason why in practice $(1+r)\beta > 1$ is that there might be some element of default or inflation risk in owning government debt. Even if there is no default risk at present, the steady state random walk idea in principle allows debt to steadily rise if the economy is hit by a series of bad shocks, so at some point default risk may become relevant; alternatively, if agents assume that the government will choose to reduce the real value of the debt stock via inflation, this will raise nominal interest rates, and the associated uncertainty may raise r .

Even if we stay with the assumption that $\beta(1+r) = 1$, there are two reasons for believing that policy should aim to steadily reduce debt in normal times. First, shocks may be asymmetric, with the economy occasionally being hit by a large negative shock, like a financial crisis, while output never exceeds potential by more than a relatively small margin, as in Milton Friedman's "plucking" model of business cycle fluctuations (Friedman, 1993). This certainly seems more consistent with the stylised business cycle facts of recent decades in developed economies than the "standard" model with symmetric deviations from trend. In this case it is straightforward to show that tax smoothing actually implies a steady reduction in debt in between these shocks (Mash, 2010). Second, shocks may be symmetric, but large negative shocks like a financial crisis might mean that we enter a liquidity trap, so that fiscal expansion is required to assist monetary policy, while large positive shocks could be dealt with by monetary rather than fiscal contraction. There is no equivalent upper bound for interest rates, so prudent policy would reduce debt in normal times to make room for the liquidity trap possibility (Wren-Lewis, 2010). We discuss how a fiscal rule might operate during a liquidity trap in a later section. Recent developments suggest that this is more than just a theoretical possibility. The case for assuming that any rule should lead to gradual, if slow, debt reduction in "normal times" seems to be both strong in principle and to have been significantly bolstered by recent developments.

If $(1 + r)\beta > 1$, the implication that the government should eventually target a negative level of debt appears both extreme and impractical. However this implication has to be taken together with tax smoothing, which implies any adjustment towards the target should be very slow. As Leith et al (2011) show, the combination of a very ambitious long run target for debt and a very long adjustment period to achieve it might lead to quite practical goals for debt in the short to medium term. However it should be clear that this combination does raise issues when it comes to the design of a fiscal rule, which we discuss further below.

If $(1 + r)\beta > 1$, there is an additional reason why it might be desirable to eliminate government debt completely, and that is because it crowds out productive capital. In simple overlapping generation models, agents save to fund their retirement, and this determines the size of the capital stock. If agents have an alternative means of saving, which is to invest in government debt, then this debt displaces productive capital. Unless we think that current levels of productive capital in the economy are already too high (which is unlikely if $(1 + r)\beta > 1$), then reducing government debt will increase output and welfare.

This argument assumes that government debt and assets financing productive private capital are close substitutes. However the risk characteristics of the two assets is likely to be rather different, and this has given rise to a recent literature on the need for a large stock of safe assets to ensure financial efficiency and stability (see, in particular, Caballero and Farhi (2013)). What is less clear from this literature is whether this requires a large net stock of government debt, or just a large gross stock (which could be matched by financial assets held by the government).²

How does the distinction between government consumption and public investment change the argument? If individuals really did adjust their own consumption, as in the idealised Ricardian model, to look after the welfare of future generations, then it doesn't: the tax smoothing principle still applies. The same would be true if generations were selfish, but the policymaker acted as a benevolent social planner that treated generations equally apart from any discounting. However in this case issues of intergenerational equity do arise, which we discuss further below. If the public investment generated revenues in a non-distortionary manner, then this would reduce the overall burden of distortionary taxation the government would require. You would still want to smooth distortionary taxation, but the fact that the investment project partly 'paid for itself' in a non-distortionary way would add to its attraction.

Any policy to steadily reduce government debt will itself have strong implications for intergenerational equity. Future generations will enjoy a world with lower distortionary taxes, while the current generation will bear the cost of achieving that goal. If generations are not completely altruistic in the Barro sense, then this is clearly not a Pareto improvement. However requiring that policies designed to reduce government debt generate Pareto improvements is actually a recipe for deficit bias.

We can perhaps take away two conclusions from this discussion. The first is that the analysis of the optimum long run target for government debt is undeveloped, and is an area where

² This is discussed further in <http://mainlymacro.blogspot.co.uk/2013/01/safe-assets-and-government-debt.html>.

additional work could be usefully done. However, the case for aiming for a gradual reduction in debt levels seems to be reasonably strong in practice, particularly given the currently high levels of debt in most countries. The second, in contrast, is that tax smoothing implies that adjustment towards this long run target should be gradual and that shocks to debt should be accommodated in the short run rather than reversed quickly via policy. Fortunately, this means that we may have time to refine our ideas on what long run debt targets should be, without making too great a mistake in the meantime, as long as we adjust debt levels slowly.³

³ For this reason, choosing debt targets that are not too far away from recent levels, as the UK government did in 1998 (see Wren-Lewis, 2013) may be a reasonable strategy.

Section 3. Deficit bias, and the case for fiscal rules

It follows from the previous section that a welfare-maximising government would in general be expected to follow fiscal policies which broadly satisfied the following conditions:

- a gently declining path of debt over the medium term, but with blips in response to shocks
- broadly stable tax rates and recurrent government consumption.

Observed behaviour in the 1950s and 1960s was not inconsistent with this. However, as Calmfors and Wren-Lewis (2011), among many others, document, government debt as a share of GDP in the OECD as a whole has steadily increased since the 1970s. Between the mid 1970s and the mid 1990s the gross debt ratio nearly doubled in the OECD, rising from around 40 per cent of GDP to about 75 per cent. By the time of the 2008 crisis it had reached about 80%, despite the so-called "Great Moderation". This phenomenon, which runs contrary to the implications of the theory set out in section 2, is generally called 'deficit bias'.

Calmfors and Wren-Lewis survey a number of potential reasons for "deficit bias" (relative to what we might expect if governments were maximising social welfare):

(i) informational problems: this can include both over optimism by policymakers about future tax receipts, or asymmetric information where a government exploits a lack of public information to offer tax cuts or spending increases without informing the public that these will have to be reversed at some point. There may be an important intertemporal element here. While information on current deficits and problems may be good, information about future trends is inevitably less firm. As a result, a government may be tempted to replace current borrowing with leasing arrangements that imply future financial commitments (e.g. PFI in the UK).

(ii) impatience: this can be confined to the government, or shared by the government and the public (e.g. hyperbolic preferences).

(iii) exploiting future generations: increasing debt can transfer resources from future to current generations. If individuals or governments are altruistic towards future generations, this is a form of impatience.

(iv) electoral competition: politicians and the public may not be unduly impatient, but the prospect of electoral defeat can mean it is optimal for the governing party to raise debt to constrain what future governments can do.

(v) common-pool theory: groups lobby for particular fiscal actions, but because the costs of any single measure are relatively small, this cost is not internalised by the decision making process. This is clearest in the case of spending decisions (e.g. building a new bridge), but may also work on the tax side in terms of allowances or exemptions.

(vi) time inconsistency: as any fiscal decision that involves changing debt contains implications for future fiscal decisions, time consistency issues can easily arise, particularly when combined with some of the other factors leading to deficit bias.

This list is probably not exhaustive. For example, although a government may not be overoptimistic, it may be insufficiently risk averse relative to the public, and therefore fail to plan for large negative shocks along the lines suggested in the previous section. It also seems likely that different possible explanations for deficit bias may be relevant at different times or in different contexts, so it would be unwise to design rules that would address one or two of these potential factors in isolation.

Recognition that factors like these (particularly time-inconsistency) can lead to sub-optimal decision making has led to a considerable shift in the way monetary policy is implemented. In almost all major developed economies (and many developing ones) policymakers have sought to address these potential issues by delegating monetary policy to independent central banks; these central banks, in turn, are generally given targets for inflation, which range from the specific (as in the UK) to the general (as in the US.)

Given the perceived success of the independence/target model for monetary policy - at least up to the Great Recession - there has been increased interest among policymakers in applying this model, at least in part, to fiscal policy, while recognising that the issues in so doing are considerably more complex. (The differences between monetary and fiscal policy in this respect are discussed in Wren-Lewis, 2013c.) The two remedies which are generally suggested are fiscal councils and fiscal rules. We discuss fiscal councils in a later section, but here we do the same for fiscal rules.

In some cases rules may not be the best way of combating bias: for example if bias comes from overoptimistic forecasts, it may be sufficient to delegate the forecasting process to an independent body (like a fiscal council) that is not subject to over optimism. In some other cases it is difficult to see how rules can avoid deficit bias. If both the government and the private sector knowingly want to exploit future generations, then it is difficult to see why they would enact (or, having enacted, stick to) rules that prevented them doing so.

In some cases rules can be a useful device for overcoming bias that stems from the system of decision making rather than any non-benevolence in the actors concerned. Common pool problems could be a case in point, or difficulties caused by time inconsistency or democracy itself. In these cases, there would be less reason to worry about the complexity or robustness of rules, as there is no major incentive for the government to circumvent them.

However, some of the causes of deficit bias involve a non-benevolent government attempting to take advantage of asymmetric information. Here a rule is clearly meant to be a disciplining device on government, and it will be a device that at least some governments will seek to avoid. In this case there may be a clear conflict between the optimality of a rule and its effectiveness. One example is where it may be optimal to have targets for cyclically adjusted rather than actual deficits, but any cyclical adjustment mechanism will be complex and potentially open to manipulation, and so rules based on cyclically adjustment may be less effective at constraining non-benevolent policymakers. Another is when it may be desirable to treat a rise in debt resulting from

an increase in investment differently from one caused by an increase in government consumption; but in practice governments have considerable flexibility to reclassify spending from one to the other.

The problem of designing rules based entirely around the imperative of eliminating deficit bias is that such rules may lead to severely sub-optimal outcomes. The most obvious example is a balanced budget rule. Even if monetary policy was perfectly able to stabilise demand, so that switching off the automatic stabilisers was not a concern, any remaining shocks that impacted on the government's finances would have immediate consequences for either spending and taxation, leading to variability in one or both that would damage social welfare. Balanced budgets are the antithesis of tax smoothing. For the same reason, attempting to balance the cyclically adjusted budget deficit each year would be seriously sub-optimal.

One response is to say that the benefits of eliminating deficit bias always outweigh the costs of sub-optimal policy. Even if this is true, which we very much doubt (see, for example, Kirsanova et al (2007)), it might be short-sighted from a political point of view. Fiscal rules that are severely sub-optimal may not be durable. Another, more cynical, view might be that all governments are always bad, and that this rather than optimality should be our focus. However the empirical evidence also suggests that this is incorrect. While on average government debt in the OECD area almost doubled in the 30 years up until 2007, this trend was not uniform across countries or across time. For example the UK debt to GDP ratio fell until the mid-1970s, and showed no clear trend thereafter until 2008. A more general, but perhaps more contentious, example might be the drive for austerity since 2010, which many would argue illustrated that governments could be too anxious to cut deficits even when it was not optimal to do so. (For further discussion, see section 5).

Equally, it is probably a mistake to ignore the issue of deficit bias and focus exclusively on optimality. An optimal rule that allowed governments to let debt increase following negative shocks could easily be exploited, because the impact that particular shocks might have on the public finances is difficult to judge. Much more complex rules that attempted to build in conditionality also increase the scope for politicians to game the system. (For example, rules that aim to balance the budget over the course of an economic cycle can be manipulated by selectively choosing the start and end dates of the cycle.)

To illustrate these trade-offs, suppose governments can be of two types. They can be 'good' (benevolent) governments, which means that they want to implement the policy that maximises social welfare. These governments are not subject to deficit bias. However there can also be 'bad' (non-benevolent) governments that are subject to deficit bias. In the previous section we just considered policies designed for good governments, but we also want to allow for bad governments, because we have good evidence from trends in debt over the last 30 odd years that these governments exist. It seems sensible to try and consider how rules that work for good governments might perform if governments were bad.

This dilemma can be summed up in the table below.

Rule/Policy maker	Benevolent	Prone to Deficit Bias
Optimal rule	Policy is optimal, but discretionary action might achieve the same outcome	Leaves scope for deficit bias
Rule designed to prevent deficit bias	Sub-optimal policy, but less deficit bias. May lack political durability.	

There are two immediate implications. Rules that are effective at reducing deficit bias, yet do not stray too far from optimality, are clearly desirable if they can be found. Second, the choice of fiscal rule may legitimately vary between countries (or perhaps over time), if there are known differences in the type of government likely to be in power, and how much other institutional checks on deficit bias exist. For example, if the fiscal forecasts on which budget decisions were made was contracted out to an independent body, we might be tempted to put greater weight on the optimality of rules.

Section 4. Alternative types of fiscal rule

The previous sections looked at two rather different approaches to formulating fiscal rules. In section 2 we looked at what might influence the choice of an optimal fiscal rule that a benevolent government might adopt. In contrast section 3 considered the problem of deficit bias, which motivated the usefulness of a fiscal rule as a disciplining device for a non-benevolent government. In this section we see how this tension between design principles might play out in terms of specific forms of fiscal rule.

Simple debt feedback rules

Academic work on fiscal rules have typically adopted rules of this type

$$g_t - g^* = -a(d_{t-1} - d^*)$$

or

$$\tau_t - \tau^* = c(d_{t-1} - d^*)$$

where d and d^* are the actual and long run target for the debt to GDP ratio, g^* and τ^* are the spending and tax ratios associated with that optimal level, and 'a' and 'c' are positive parameters. We can describe rules of this form as 'simple debt feedback rules'. If debt is above its long run target, government spending falls below its long run level, and/or taxes rise above this level. The further debt is away from the target, the larger the adjustment in terms of taxes and spending should be. This form of rule follows naturally from the discussion in Section 2.

Simple debt feedback rules of these types have been used to proxy optimal government behaviour in many academic papers (e.g. Schmitt-Grohe and Uribe (2007)). As Kirsanova and Wren-Lewis (2012) clearly show, if the adjustment parameters are small, such rules can come close to mimicking optimal fiscal adjustment. This reflects the key result discussed in section 2 that adjustment towards the target level of debt should be slow.

Most academic studies do not find a large difference in the optimal speed of adjustment for spending and taxes (i.e. in the parameters 'a' and 'c'). In fact most assume only one instrument is used for debt feedback. We could therefore combine the two rules in the following way;

$$\tau_t - g_t = rd_{t-1} + f(d_{t-1} - d^*)$$

where, as in section 2, τ and g are the ratios of tax and spending to GDP, and r is the *growth corrected* real interest rate. Here 'f' is our feedback parameter. Thus the primary surplus as a ratio to GDP responds to deviation of the debt to GDP ratio from its long run desired level, but if debt is at this level the primary surplus needs to cover growth corrected interest payments on debt.

There are two major problems with rules of this type as a means of disciplining governments, or providing an incentive for a government to behave in a benevolent manner. The first is that the rule does not provide any incentive for the government to fulfil its plans. In any year the government may promise to follow the rule, but in practice spend more, or tax less, than the rule implies. As a result, debt may get no nearer its target, or could even move further away from it,

but this would have little consequence for future fiscal actions. The following year, the government would again promise to follow the rule. We can call this the absence of an 'implementation incentive'.

The second problem is that the debt target d^* is unlikely to be achieved within the lifetime of a government (Dolphin, 2013). This is a consequence of the optimum adjustment parameters being small. Governments naturally do not like targets that they cannot come close to achieving. We can call this the 'realisable target' problem. It is convenient to consider each problem in reverse order.

The realisable target problem

As we discussed in Section 2, it is quite possible that the optimal long run level of debt could be a considerable distance away from current debt levels. It might even be optimal for debt to be negative. Equally it is almost certainly going to be the case that optimal adjustment to this target should be slow, involving decades rather than years. As a result, the long run optimal level of debt is unlikely to be achieved within the lifetime of a government.

This will cause obvious political problems. If the optimal level of debt is 30% of GDP, say, and current levels of debt are 80% of GDP, it may be optimal to only reduce debt to 75% of GDP within five years. Yet to the public, and of course political opponents, it may appear as if the government has failed to come close to 'its target'.

One way of overcoming this problem is to calculate the levels of debt, and associated deficits, that would be consistent with the optimal path of debt reduction *in the absence of any unforeseen shocks*. We could use the debt feedback equation involving the primary deficit to make this calculation. The debt or deficit level at the end of the government's term could then become what we might term the government's 'operational target', which it would have a clear incentive to achieve. The problem with this strategy is that unforeseen shocks will inevitably occur. As Section 2 made clear, it is sensible to allow debt (and the deficit) to be a shock absorber. The simple debt feedback rules outlined above do that, because d^* is a long run target, and the parameters of the rule imply weak feedback. An intermediate operational target for debt would not. We can see this most clearly if we imagine an adverse shock that occurred just before the end of the government's term. Attempting to achieve the level of debt specified by the optimal pre-shock path in the following year could involve sharp cuts in spending or increases in taxes, which is directly contrary to the smoothing principle outlined in section 2.

The implementation incentive

If the fiscal operations of governments were completely transparent, the implementation incentive would not be necessary. Each year the government would outline plans that were consistent with a rule, and the public could then check that its actions were consistent with that plan. In reality this is not the case.

Of course the public does have access to data on tax receipts and spending levels, although this data is often not as transparent as it could be. In addition we know what the government is

borrowing. However realisations for these variables could be the result of unforeseen shocks (or equivalently, forecast errors) as well as deliberate actions by government. So a government could say that its failure to deliver on plans was due to 'bad luck'. This will often be the case. As a result, a non-benevolent government that was subject to deficit bias could spend more and tax less than it had promised to do, and pretend that this was instead due to adverse unforeseen shocks. There is nothing in the form of the simple debt feedback rules to 'punish' such behaviour. So, to return to our numerical example, a government that came into office with a debt level of 80% of GDP, and had a long run target of 30%, could easily pass off a debt level of 80% of GDP at the end of its period of office as the result of bad luck, even if no adverse shocks had in reality occurred.

The implementation incentive would be much greater if the government was *required* to hit targets for debt or deficits in particular years i.e. if it had a fixed operational target. If the policy rule had a fixed target for government debt that had to be achieved by the end of the government's term in office, any overspending in the early years through failures to implement policy would be 'punished' by either failing to hit the target, or having to under spend in later years. But as we have already noted, if the economy had actually been hit by adverse shocks, fixed operational targets would be severely sub-optimal.

The problem is well illustrated by the experience of the current UK government, which had a target that debt should be falling in 2015-16. This target is now not forecast to be met, primarily because economic circumstances have been much less favourable than originally forecast. It would clearly have been sub-optimal for the government to have adjusted fiscal policy so as to be sure of hitting the target. However, simply ignoring the rule (given that, while the economy has been weak, there has been no renewed financial panic or other obviously extraordinary development since 2010) clearly undermines the credibility of the government's commitment to any future rule.

Debt or deficit targets

It is probably more common for rules to be specified in terms of deficits rather than debt. Of course it is just a matter of arithmetic to move from one to another. We should note that there is no 'natural' target for the budget deficit: the appropriate deficit target will depend on both the target level of debt and the extent of any initial excess debt. Once we consider both the need to adjust gradually towards any debt target and the sub-optimality of operational targets following shocks to the government's finances, then the distinction between rules based on debt and deficits becomes more interesting.

A deficit target has two advantages over a debt target. First, it makes it easier to achieve the target over the lifetime of the government, so it helps tackle the realisable target problem. Second, and more importantly, it is more robust if adverse shocks to the economy occur. Returning to our example where the initial level of debt was 80% of GDP, suppose the deficits required to achieve a debt target of 75% after 5 years were around 3% of GDP. (This would be consistent with nominal GDP growing by around 5% each year.) As a result, the debt target of 75% could be replaced with a deficit target of 3%. Now imagine an adverse shock that lasted just one year, which occurred in the penultimate year. The shock raised debt to 77% of GDP, rather than the expected 76%. To move from 77% to 75% in a single year would require quite rapid fiscal contraction. However if the target

had been specified in terms of the 3% deficit, any departure from the optimal policy would be much smaller – indeed, because the shock was just temporary, achieving a 3% deficit would still be very close to the policy implied by the simple debt feedback rule.

Outcomes are less benign if the adverse shock is persistent. Suppose, for example, that the shock was expected to persist into the final year, so that if no action was taken the deficit would be 4% of GDP in the final year. Policy changes would be required to reduce the deficit by 1% of GDP to hit the 3% deficit target, which is almost certainly a more rapid correction than would be implied by a simple debt feedback rule with optimally small parameters. However it is still the case that a rule based on deficits would be less suboptimal than one based on debt: reducing debt from 77% of GDP to 75% of GDP would require a much sharper fiscal contraction.

This example also shows why any operational deficit target has to be dependent on the extent to which current debt levels deviate from the long run goal for government debt. Although a 3% deficit target might imply the right degree of debt adjustment when debt was 80% of GDP, if it was retained indefinitely we would never achieve the ultimate goal of 30%. With 5% nominal GDP growth, a 3% deficit would converge to a debt to GDP ratio of 60%: we eventually need a deficit of 1.5% to move to and retain a debt target of 30%. So any deficit target would need to be periodically (every five or ten years) revised as debt falls. We could of course start off with a deficit target of 1.5%, but this would produce a much more rapid debt reduction in the early years. (Debt would fall from 80% to 70% rather than 75% in the first five years.)

One possibility that has been suggested is that any deviation from the deficit target should be credited in a control account, and that this control account has to be returned to zero within five or ten years if it is in deficit. This is the Swiss debt brake (Geier, 2011), which has been subsequently adopted in Germany (Wyplosz, 2012). It is straightforward to see that this also effectively controls the level of debt, but much more rapidly than would be implied by an optimal policy. If shocks raised the level of debt and deficits above target paths (thereby putting the control account into the red), deficits *below* target would be required to bring debt back to target and the control account to balance within 5 or 10 years. So in practice a debt brake is likely to lead to a large departure from optimal policy, which seems undesirable unless deficit bias in the recent past has been persistent and substantial.

We could generalise this discussion in the following way. The closer rules are to optimal policy, specifying rules in terms of debt rather than deficits would seem more appropriate and straightforward, as the simple debt feedback rules illustrate. The closer rules are designed to prevent non-benevolent behaviour, and therefore involve fixed operational targets, specifying targets in terms of deficits rather than debt is more robust to unexpected shocks. Reducing the deficit below target to correct past deficit overruns (the 'debt brake') implies very rapid debt stabilisation, and this large deviation from optimal policy should only be considered if governments are highly non-benevolent.

If deficits rather than debt are the basis of operational targets in fiscal rules, two further issues arise. The first is whether the target should be the actual or primary deficit. The second is what date these targets should apply to.

The primary deficit is the actual deficit excluding interest payments on government debt. If the ultimate goal of policy is to achieve some target for government debt, or more weakly sustainability (a stable level) in government debt, then targets for the primary deficit will not be robust to long term shifts in real interest rates. However, as targets for actual deficits are also not robust to past persistent errors for reasons already discussed, this is not a new problem. An advantage that operational targets for primary deficits have is that they remove a potential interaction between monetary and fiscal policy. However, this interaction could be helpful. In a boom, higher interest rates would tend to worsen the actual budget deficit if a significant amount of debt was short term, and if spending was cut as a result, this would help dampen the boom, acting as an automatic stabiliser. Perhaps a more important potential advantage of focusing on the primary deficit is that it reduces the extent to which market panics can become self-fulfilling. If interest rates on government debt rise because of fears of default, and the government therefore has to cut back further on spending to meet a deficit target, this may increase the incentive for the government to default.

Fixed or rolling deficit targets

Suppose that because the government wants a fiscal target it can achieve within its lifetime, but it wants a target that is also to some extent robust to shocks, it specifies a target in the form of the actual or primary deficit. The question then arises as to the date that this target should apply to. Having a fixed target applied to each year does not make sense, because it removes the shock absorber property that was stressed in Section 2. There seem to be two main possibilities. Either the government aims to achieve the deficit target by at least a set date in the future (e.g. once it takes office, in 5 years time at most), or it always aims to achieve the target over a 5 year time horizon at most. An example of the latter, rolling target, is the current primary UK fiscal mandate which aims at achieving structural current balance (i.e. excluding investment, but including interest payments on the existing debt) within a five year time horizon.

The advantage of a rolling target, compared to a fixed operational target, comes once we allow shocks to occur. If the economy is hit by adverse shocks, but the target remains unchanged, the government has five years to smooth any adjustment. A rolling target therefore allows debt to be used as a shock absorber to some degree, while a fixed operational target does not *once we approach the target date*. This, as noted above, is precisely what has happened in the case of the UK government's secondary (debt) target, which was why it was abandoned.

In some sense a rolling deficit target could be described as an intermediate step on the way to the optimal debt target described at the beginning of this section. The latter is based on the principle that, at any point, the government should target a level of debt, and set the deficit consistent with achieving that level slowly, over time, while accommodating shocks. The former is then based on moving to the latter, but again slowly, over time, while accommodating shocks.

A rolling target clearly does not solve the implementation problem. A government can continue to fail to implement plans and pass this off as bad luck, just as it can with a simple feedback rule. The rolling target goes some way, however, to solving the realisable target problem. If there are no adverse shocks, and the government acts benevolently, it can take credit for achieving its rolling

target after five years, even if it formally need not do so. On the other hand if it is actually hit with adverse shocks, the rolling target gives it more flexibility than a fixed target would.

Thus a rolling target does share two important characteristics of optimal policy, which is that fiscal adjustment should be both slow and state contingent. If the economy is hit by a bad economic shock that raises deficits over a temporary period (say a temporary downturn in the housing market that reduced tax revenues), then it would be inappropriate to try and correct for those immediately, and a rolling target allows the government to do this.

Cyclical adjustment and other forms of conditionality

One way of tackling the implementation problem, and making fixed operational targets more robust, is to incorporate conditionality into the rule. Conditionality helps increase implementation incentives, because it formalises elements of 'bad luck' in ways that can be verified. For example, if a rule says that a recession should mean that the debt target to be achieved at a certain date should be revised up from 75% to 77%, the government cannot then claim that the recession excuses debt increasing to 80%. Conditionality can also mitigate the problem that fixed operational targets depart from the principle that debt and deficits should be shock absorbers. If the target is adjusted for some specific shock, it can help absorb that shock.

Cyclical adjustment is just one form of conditionality. It is widely adopted as part of the fiscal rule architecture in many countries. It is relevant to a deficit rather than a debt based rule. There are two ways of allowing for the economic cycle. One is to calculate cyclically adjusted deficits. This is a difficult calculation which could potentially be manipulated if done by the government. The alternative is to apply rules 'over the course of a cycle', as the Labour government's rules created in 1998 did. Although this approach appears less complex, it can prove equally controversial if subsequent data appears to alter cycle dates. In addition, there is the danger that large shocks that occur early in the cycle could lead to inappropriate policy in the latter part of the cycle. As Wren-Lewis (2013) discusses, both problems occurred over the ten years that Labour's fiscal rules were applied.

A fiscal rule based on achieving very short term targets for the deficit is clearly better if it focuses on the cyclically adjusted deficit, because this lets the automatic stabilisers work. Now in theory an omniscient monetary authority could always stabilise the cycle such that automatic stabilisers were unnecessary, but in practice monetary policy is not that good. However it is less clear why rules that focus on deficits in five or more year's time, like the UK's fiscal mandate, should bother with cyclical adjustment. This is because, liquidity traps apart (see section 5), any forecast of the economy in five years time is almost bound to assume that output will be close to the natural rate, so the projected actual and cyclically adjusted deficits will be very similar. (The current cyclically adjusted deficit will be useful in assessing whether that target is likely to be achieved, of course. The argument above is about the form of rolling targets: calculating the current cyclically adjusted deficit is still useful.)

There is a clear cost to conditional rules, which is the added complexity that they introduce. It is generally the case that the more complicated the rule, the easier it is for a bad government to manipulate it. In a downturn governments may be tempted to overestimate the impact of the cycle

on taxes and spending, so as to produce a more favourable cyclically adjusted deficit. When the Labour government's fiscal rules involved achieving targets 'over the course of the cycle', there was considerable controversy when the dating of the cycle was changed (Wren-Lewis, 2013). One way to reduce this possibility is to engage an independent fiscal council in operationalising the conditionality (e.g. calculating any cyclical adjustment), which we consider in section 6.

It is also important to recognise that cyclical shocks are just one of many forms of shock that can impact on the government's finances. It is arguably of particular importance because of the need to maintain automatic stabilisers (see section 2), but the principle that debt should be a shock absorber applies to all shocks. For example, the experience of the Labour government from 1997 to 2007 was that tax receipts were unexpectedly high for the first half of the period, but then subsequently unexpectedly low (Wren-Lewis, 2013). It is therefore not practical to allow for every type of shock using conditionality. Conditionality is not as practical solution to the implementation problem. A more general solution to this problem is forecast transparency, which we consider in the section on fiscal councils below.

Government investment

Should a fiscal rule make special provision for government investment? For example, the 'golden rule' brought in by the Labour government in 1997 involved balancing the current balance over the economic cycle, and therefore excluded government investment spending. (Their second fiscal rule, which involved an upper limit on government debt, was therefore called by many the sustainable investment rule, because only this prevented the government going on a borrowing spree to fund investment, although as Kell (2001) noted, the implied level of sustainable investment implied by the two rules combined was implausibly high.)

The motivation for excluding investment spending from a rule that requires a balanced budget over some period is straightforward enough, coming from the observation that firms often use borrowing of some form to finance investment. However public investment, unlike private sector investment, will not normally generate a direct return. If the investment increases economic growth, and therefore future taxes, then there may be an indirect return, but not all worthwhile investment will do that. So taxation will have to rise at some point to pay for the investment. As a result, we cannot exclude public investment entirely from any fiscal rule.

There are nevertheless clear differences between government consumption and investment in terms of intergenerational equity. If we do not trust the private sector to appropriately adjust their bequests to compensate for future tax liabilities, or if the tax system discourages them from doing so, then it makes sense for the government to relate the timing of taxes to the timing of benefits received. As many forms of public investment are likely to benefit future generations, then deferring taxes by borrowing makes sense on equity grounds.

This argument is reinforced if we follow the idea that rules are designed to prevent particular forms of non-benevolent behaviour by governments. At the time that the Labour government introduced its fiscal rules there was a strong belief that previous governments trying to bring down public borrowing had cut public investment much more than public consumption (as indeed has happened since 2010). It is easy to see how this might happen: future generations who

would be harmed as a result cannot vote. So treating public investment differently from public consumption may be desirable to protect future generations.

Arguably a better way in dealing with the problems discussed above is to supplement the main fiscal rule (in terms of debt or deficits), with a subsidiary rule specifically involving public investment. For example, governments could aim to achieve overall levels of public investment of at least $x\%$ of GDP. This would be a much more transparent way of having the discussion about what appropriate levels of investment were required for intergenerational equity, and how any targets related to the rates of return on these investments.

Summary

In Section 3, we argued that there was a tension between designing rules for optimal behaviour, and designing rules that were effective at preventing non-benevolent government actions (leading to deficit bias). In this section we suggested that if governments were benevolent (so that there was no need for implementation incentives), and were happy to not achieve targets during their lifetime (there was no realisable target problem), then optimal rules would take the form of a simple and very gradual feedback on the deviation of debt from its long run target.

In a more realistic setting where implementation incentives were required, we argued that operational targets in the form of deficits were more robust to shocks. Deficit targets might limit the extent to which deficits could act as shock absorbers, but they to an extent still allowed debt to play this role. Whether deficit targets were fixed or rolling would depend on the extent of non-benevolent behaviour by the government. Targets for the deficit that had to be achieved by a fixed date could (if achieved) lead to severely sub-optimal policy if shocks occurred just before that date. Rolling targets avoided that problem, but increased the likelihood that a non-benevolent government would continually put off achieving the target.

A similar trade-off existed with conditionality. Conditional targets could more closely approximate optimal policy, but could also allow a non-benevolent government to game the system. For this reason the desirability of conditionality would depend on the dating of targets. For example targets to be achieved by a fixed date would need to be cyclically adjusted to avoid turning off the automatic stabilisers, but if rolling targets were for a date that exceeded the length of a normal economic cycle, cyclical adjustment would not be necessary. Finally we argued that public investment should not be excluded from deficit targets, but instead separate targets for the ratio of public investment to GDP could ensure that any fiscal retrenchment did not lean too heavily on public investment projects.

Section 5. Fiscal policy when and monetary policy is constrained

Kirsanova et al (2009) argue that, prior to the Great Recession, there was a widespread consensus as to the respective roles of monetary and fiscal policy. Monetary policy was assigned to the control of inflation and the stabilisation of demand, while the objective of fiscal policy was reasonably stable tax rates, consistent with the control of government debt, broadly in line with the conclusions of the section 2. In practice this meant handing over monetary policy to independent central banks, who were given inflation targets and expected to follow some version of a Taylor rule, implying that the central bank would largely offset any impact of fiscal policy changes on demand; this in turn would ensure that those responsible for tax policy would focus on medium-term stability, since they would have little power and hence incentive to use fiscal policy for demand management.

Most policy assignments of this type involve some deviation from fully optimal policy, where both instruments would be used to achieve both goals. However that paper discusses at length why this might not be the case for the goal of demand stabilisation, because here there are reasons for believing that monetary policy may dominate fiscal policy. (The paper suggests the same is not true for the goal of debt stabilisation: if debt is high monetary policy may be quite efficient at reducing debt without sacrificing too much in terms of demand stabilisation. As a result, this aspect of the consensus assignment probably requires a different justification, such as the desire to avoid fiscal dominance based on political economy concerns.)

The emergence of this consensus assignment undoubtedly owes a great deal to the end of the Bretton Woods era of fixed exchange rates, which gave countries greater scope to use monetary policy for demand stabilisation. It also reflected increasing scepticism about the impact of fiscal policy on demand, following the "rational expectations" revolution in macroeconomics. Ricardian Equivalence implies that cuts in taxes would have no income effect on consumption spending if consumers were free to borrow and lend at the same rate as the government. (There would still be important incentive effects, unless taxes were 'lump sum'.) However temporary changes in government spending would still have an important impact on aggregate demand, however they were financed. So Ricardian Equivalence and rational expectations certainly does not imply that all fiscal policy changes would have no impact on short run demand.

Reliance on monetary policy for demand stabilisation clearly does not apply to countries within a fixed exchange rate regime, or who are members of a monetary union. Here monetary policy is not available to deal with asymmetric shocks across the union, or asymmetric economic structures giving different responses to common shocks. As fiscal policy is able to respond to these shocks, then any optimal use of fiscal policy is inevitably going to involve fiscal policy playing a discretionary countercyclical role, unless nominal prices and wages are completely flexible. There is strong empirical evidence that this is not the case, not least the behaviour of nominal and real exchange rates themselves.

It might nevertheless be the case that asymmetric shocks are sufficiently small or uncommon, or prices are sufficiently flexible, that this countercyclical role for fiscal policy is unnecessary. This was perhaps the reason why even the limited role for countercyclical policy envisaged in the original Maastricht Treaty was circumscribed by the Stability and Growth Pact

agreed on the formation of the Euro area. Subsequent events clearly show that this view is incorrect. Easier monetary conditions in periphery countries, along with other country specific factors, led to a substantial and sustained period of excess demand and relatively high inflation in these economies, such that over a period of seven years their real exchange rates became seriously overvalued within the Eurozone (Wren-Lewis (2013b)).

This strongly suggests that any fiscal rule for an economy within a monetary union or fixed exchange rate regime should contain an important countercyclical element. However, the same logic also suggests that the focus should be on the cyclical condition of the economy relative to the other union members, rather than some absolute measure of cyclicity. If, for example, there was excess demand in all union members, then the consensus assignment suggests this is best handled by monetary policy at the union level, and not by fiscal policy operating within all the union members.

There are a number of ways relative divergence could be captured within a fiscal rule. Relative inflation rates are easy to measure. Relative output gaps are less so, but could also be used. The third possibility is to look at the relative price level, rather than its rate of change i.e. the real exchange rate. All three possibilities were examined in Kirsanova et al (2007), with no clear 'winner' in terms of welfare gains, but this is an important area where further work is required. This work could also look at whether current account balances could play a similar role.

The second straightforward caveat to the consensus assignment, largely ignored prior to the Great Recession but highly relevant at present, is when interest rates hit their lower bound. (We will refer to this as the Zero Lower Bound or ZLB problem, but the term liquidity trap is also widely used and any distinction between the two is not important for our purposes.) In this situation monetary policy is not able to reduce real interest rates to the level required to eliminate the output gap and stabilise inflation, so any macroeconomic or practical advantages that monetary policy may have in this regard disappear. Fiscal policy can therefore play an important stabilisation role.

There might be two objections to this conclusion, besides the point which we have already dealt with about fiscal policy effectiveness. The first is that alternative monetary policy instruments could substitute for the loss of the short term interest rate tool. The instrument that has been widely used in the Great Recession has been Quantitative Easing. The general consensus among central banks is that this works by reducing longer term interest rates via portfolio effects. While there is some evidence that these measures have significant effects, they are also subject to an equivalent ZLB. More importantly, whereas it is reasonable to assume that monetary authorities can have a fairly precise control over short term interest rates, the impact of QE on longer term rates and hence on aggregate demand is much more uncertain. As a result, any superiority that conventional monetary policy has over fiscal policy as a stabilisation tool is unlikely to apply to QE.

The second objection is that the ZLB constraint can be partially overcome by the monetary authority making commitments to raise output and inflation above desired levels in the future, which if credible would lead forward looking agents to raise spending and inflation expectations today. (This mechanism is discussed in detail in Eggertsson and Woodford, 2003.) This promise is time inconsistent, but this problem could be reduced through mechanisms such as price level or nominal GDP targets (see Woodford, 2012). However, such measures are a clear second best to

conventional monetary policy, in that they involve costs in terms of above target inflation in the future; and again, their impact on aggregate demand is subject to considerable uncertainty. Once again, any superiority that conventional monetary policy has over fiscal policy as a stabilisation instrument is likely to disappear, and it may well be the case that fiscal policy is now the superior stabilisation instrument (see Eggertsson and Woodford, 2004). There is also the prospect implicitly raised by Larry Summers; that if the desired level of real interest rates (that is, consistent with a small output gap/the natural rate of unemployment) is below zero, the authorities will be tempted to promote financial “bubbles”, so as to raise private borrowing levels and hence demand; it seems clear that expansionary fiscal policy would be a more efficient and less uncertain approach.

Consider, for example, a monetary policy regime involving targets for the level of nominal GDP. It remains quite possible that a sufficiently large negative demand shock might mean that monetary policy is not able to achieve this target in the short term. As a result, achieving these targets subsequently will involve inflation going above desired levels. If fiscal policy is able to help achieve the nominal GDP target in the short run, it will also avoid this cost in terms of excess future inflation.⁴

The importance of using fiscal policy to help stimulate demand at the ZLB poses an obvious problem for a fiscal policy rule. While the routine need for countercyclical policy in a monetary union can easily be incorporated into a fiscal rule by, for example, adding terms in relative inflation or the real exchange rate, the ZLB problem makes fiscal policy actions conditional on the state of monetary policy. Occasions in which the ZLB applies are obviously asymmetric (there is no mirror case where fiscal policy needs to be tighter because interest rates reach an upper limit), but they are (hopefully) infrequent. It would as a result be inappropriate to simply add a stabilisation term into a fiscal rule, so that fiscal rules took on the character of Taylor rules for monetary policy.

One possibility is to treat the ZLB problem as one of those unusual events that imply a departure from any normal rule, and to treat such departures on an ad hoc basis. The signal that the normal fiscal rule should be temporarily suspended is clear enough, because short term interest rates can be directly observed. Unfortunately in practice, as Wren-Lewis (2010) notes, uncertainty coupled with policy lags mean that it may be prudent to implement expansionary fiscal policy before the ZLB limit is reached, but instead when there is a significant possibility that it might be reached.

There is a danger in an ad hoc approach that the importance of setting aside the normal fiscal rule when the ZLB is likely to be hit becomes forgotten over time, and when it does occur there is resistance to abandoning the rule. After all, even though the Great Depression is widely recognised as both a period in which the ZLB constrained monetary policy, and where fiscal policy was instrumental in ending the Depression, it is still the case that there has been widespread opposition to using fiscal policy in similar circumstances during the Great Recession (Wren-Lewis, 2011). As a result, it may be important to ‘hard wire’ this role for fiscal policy by incorporating it in some way into a fiscal policy rule.

⁴ This point, and the associated point that austerity in a ZLB recession and NGDP targets will raise future inflation, is discussed further here: <http://mainlymacro.blogspot.co.uk/2012/05/austerity-nominal-gdp-targets-and-zero.html>

The possibility of a monetary union hitting a ZLB might also call into question the conclusion above that any routine countercyclical terms in a fiscal policy rule for a union member should be relative to other union members. If these terms were absolute rather than relative, then they could also capture the need for fiscal policy to be expansionary across the union as a whole if the ZLB had been hit. However this is bound to be a second best approach: for example it would not capture the asymmetric nature of the ZLB problem.

How far should fiscal policy be used to try and stabilise the economy when the ZLB constraint to monetary policy applies? A simple answer is straightforward: fiscal expansion should be large enough such that monetary policy is no longer constrained. For example, the monetary authority could be asked at what point fiscal expansion would be sufficient that their forecast suggested there was a more than 50% chance that they would raise interest rates above zero. This would imply a target, rather than a limit, for the fiscal deficit for the short term (assuming that monetary policy is aimed at targeting inflation or NGDP).

It would seem appropriate to involve the monetary authority (i.e. central bank) in this forecasting capacity for a number of reasons. First, they have the forecasting expertise. Although the fiscal council and government might be involved in fiscal forecasting, their focus will probably be on the longer term. Second, it ensures cooperation between monetary and fiscal policy. Third, if the central bank provides some detail about what impact it thinks each element of any fiscal package will have, it provides an incentive for the government to undertake policy changes that have a large impact on demand, rather than to focus on changes that might be politically popular.

This ZLB ‘knockout’⁵ for the fiscal policy rule could well involve a substantial increase in government debt. As the ZLB is almost certainly likely to have been hit as a result of a large increase in private sector saving, this is not a problem in the short term. Interest rates on government debt should not rise as long as the fiscal authority has the appropriate monetary policy backing. (Of course the whole point of the policy is to allow the monetary authority to raise short term interest rates above zero, so there should be some increase in long term interest rates.) This increase in debt will almost certainly mean that previous fiscal targets will become outdated, and so it makes sense for the government to say at the same time how they think the fiscal rule will change once the ZLB constraint no longer operates. Indeed it would be positively desirable for it to do so. Raising the level of debt to help counteract a recession must imply that taxes will be higher and/or government spending will be lower once the recession is over. It is important to calculate what those tax and spending levels are, if only to demonstrate that they are politically feasible.

This section therefore suggests two clear conclusions about the extent to which fiscal rules should respond to aggregate demand. First, within a monetary union (or any fixed exchange rate regime), fiscal rules need to be augmented to include terms that capture the relative cyclical position of individual union members. Second, for any economy fiscal rules should be suspended whenever the central bank believes that there is more than a 50% chance that nominal interest rates will hit the zero lower bound. When that happens, the fiscal authority should cooperate with the central bank in devising a fiscal stimulus package that is expected to allow interest rates to rise above this

⁵ The term comes from the forward guidance introduced by the Bank of England in 2013 for monetary policy.

lower bound. This will imply a significant increase in debt, and the fiscal authority will at the same time need to demonstrate how its fiscal rule will change once interest rates are expected to rise again.

Section 6. Fiscal Councils

Over the last ten years or so, there has been a substantial growth in what are sometimes called Fiscal Councils, or Independent Fiscal Institutions. These are generally advisory bodies, set up and funded by governments, who are mandated to provide information and often advice on aggregate fiscal policy. Calmfors and Wren-Lewis (2011) surveyed eleven such councils, but since then councils have also been created in Australia, Portugal, Ireland, Slovakia, Slovenia and South Korea.

Calmfors and Wren-Lewis note that most fiscal councils work with governments that also have fiscal rules. They are not alternatives to fiscal rules, but complements to them. Calmfors and Wren-Lewis suggest two main ways this complementarity might work. First, councils could help monitor the implementation of rules. If rules involve forecasts, both economic and fiscal, then councils can evaluate whether those forecasts are realistic, and might even prepare the forecasts themselves (as is the case with the OBR for example). If rules involve cyclical adjustment, then a council could check that the adjustment was reasonable. In terms of our earlier discussion, they could provide a strong implementation incentive. If governments broke promises on spending, but attempted to explain these in terms of bad luck (adverse shocks), a fiscal council would have the resources to establish what the government was actually doing.

Delays in information will mean that fiscal councils will not be able to eliminate the implementation problem completely. When data for government borrowing is published, it may not be clear what the reasons are for any unexpected developments. It may take time to establish why changes are occurring, and crucially whether such developments are likely to be persistent or not. This may reduce the ability of the fiscal council to signal that a government should modify its rule because it has suffered bad luck. On the other hand, delays in information will only delay the time at which a fiscal council accuses a government of deficit bias, so the political costs resulting from deficit bias may still arise.

Second, a fiscal council could validate departures from rules, where the rules were too simple to allow an optimal response to changing conditions. A clear example was advice from the Swedish fiscal council after the 2008/9 recession that departures from the government's fiscal rules were justified.

The existence of fiscal councils can therefore have two implications for the design of fiscal rules. First, rules that involve a high degree of complexity or forward looking behaviour might have been criticised because they allowed too much scope for governments prone to deficit bias to engage in manipulation. The presence of a fiscal council can reduce those risks, to the extent that the council has sufficient political capital that any rebuke would involve the government in political costs. Second, there might be a concern that some rules might not cover particular eventualities (such as hitting the zero lower bound), yet the government might be forced to rigidly adhere to these rules because not to do so would involve political costs. In this case a fiscal council can provide the government with political cover to depart from rules when it is appropriate to do so, to the extent that the council is seen as independent.

So how might a fiscal council change the matrix of pay-offs shown in section 3? If the government was benevolent, but was forced to follow a sub-optimal rule, then the fiscal council could help sanction departures from that rule when following the rule became severely sub-optimal. If the government was not benevolent, it could reduce the extent of non-benevolent behaviour whatever the form of the rule, by increasing implementation incentives. A fiscal council does not avoid the need to compare the gains and losses which each type of rule generates, but it does influence the size of these gains and losses.

More specifically, we can consider the debate over whether to have operational targets that are for a fixed point in time, or are for a rolling period (i.e. always 5 years hence). As noted above, the latter has a number of desirable properties, but may not prevent deficit bias. With a rolling end-date and a bad government, having a fiscal council could mean that deficit bias is no longer inevitable. The fiscal council can alert the public to the fact that, in its judgement, the government is deliberately not implementing its plans, or just putting off adjustment into subsequent years. The government therefore would pay a political cost in pursuing deficit bias.⁶ In the case of a fixed end-date policy, the fiscal council can signal when it thinks the fiscal situation has changed because of 'bad luck', and therefore when the government would be right to modify its rule.

There is one final way in which a fiscal council could have an important role to play, which is in adjusting the fiscal rule itself. For example, we have noted above that a fiscal rule that involved an operational target for budget deficits (whether fixed in time or rolling), will almost certainly have to be revised if shocks to public sector finances mean that deficit targets have been systematically missed (either deficits have been persistently too high or too low) in the past. Debt is now above target, and so the operational target for the deficit needs to be reduced. The fiscal council can play a very useful role in helping to calculate what that adjustment should be. For a benevolent government, this would help avoid accusations that any adjustment to rules was being governed by political advantage. For a non-benevolent government, it could avoid precisely that outcome.

This all suggests that the design of fiscal rules should be considered jointly with the design of fiscal councils. For example if fiscal rules rely heavily on forecasts of the public accounts, and those forecasts are not produced by independent public bodies, then it makes sense to have the fiscal council involved in producing or monitoring these forecasts. An example would be the UK, where the fiscal mandate involves a rolling five year target, but the OBR judges whether that target will be hit. Equally if cyclical adjustment is an important part of a fiscal rule, then a fiscal council should either undertake the cyclical adjustment (as is the case with the OBR), or have the means to ensure that cyclical adjustment calculations are not abused. A further implication is that, going back to the tradeoffs discussed before, the more flexibility and conditionality attached to a rule, and the more that rules aim at optimality rather than averting deficit bias, then the more important it is that the fiscal council has the ability and responsibility to pass judgement on whether the government is abusing either the letter or spirit of the rule, and producing analysis explaining its views.

⁶ Note, however, that in the case of the UK the very constrained mandate of the Office of Budget Responsibility means it has not expressed a view on this.

The UK may illustrate a case where this integration is absent. Currently the remit of the OBR is solely to assess whether the government is meeting its rolling five year target, which as set out above generally does not impose a binding constraint in the short term. However, the OBR's remit explicitly prevents it from commenting on the desirability of alternative fiscal policies or rules, in the short or long term, or indeed from even computing scenarios involving different government policies. It is also severely constrained in practice from commenting on the extent to which the government is observing the spirit (as opposed to the letter) of the rule, for example by artificially switching spending or revenue between years.

By contrast, the fiscal rules embodied in the original SGP can be seen as focusing on the prevention of deficit bias, with few concessions to optimality, and hence a relatively limited role for judgement in their application. For this reason, the designers of the SGP did not see any role for national fiscal councils or indeed any genuinely independent monitoring authority. In retrospect, however, the SGP rules proved ineffective, being both too rigid and too lax. Instead, monitoring the SGP was top down, undertaken by the Commission. This proved largely ineffective when the rules were actually breached, in part because governments could invoke the national interest when justifying breaking the SGP rules against Commission advice. Moreover, and in retrospect even more seriously, the view that fiscal policy was not a useful tool of demand management in a monetary union meant that countries not in deficit were unconstrained by the SGP's focus on deficits: so there were no independent public institutions warning that fiscal policy was not tight enough when some peripheral countries suffered a large monetary stimulus after the Euro was created, leading to significant departures from optimal policy.

There is little evidence that these deficiencies have been addressed by the new governance arrangements within the Eurozone. Although the latest version of the Eurozone's fiscal compact does acknowledge a role for national fiscal councils, their ability to identify non-compliance with any of the many rules of the compact remains minor at best; nor, perhaps more importantly, will they have any significant influence in cases where the rules are leading to sub-optimal policy.

Section 7 Conclusions

One of our central themes is the potential conflict in designing rules between the need to mimic optimal policy, where debt is a shock absorber and is adjusted only very slowly, and the need to prevent deficit bias by non-benevolent governments. It may therefore make sense to make different recommendations depending on what past evidence implies about the nature of governments.

If governments appear not to be subject to deficit bias, perhaps because of the existence of a powerful and well established fiscal council, then academic studies have clearly shown that simple debt feedback rules come close to reproducing the optimal fiscal policy. A key qualification is that the exchange rate is floating and there is little risk of hitting the ZLB, points we address below. A fiscal council, with a wide monitoring mandate, can provide a backstop against the possibility that a government may start trying to manipulate policy for political ends .

If past history suggests that there is a *substantial* risk that the government will behave in a non-benevolent manner, then a fiscal rule needs to have operational targets that are fixed in time. Year by year targets are too inflexible, but targets fixed for the end of the natural term of the government should provide strong incentives. These targets should be for cyclically adjusted primary deficits, rather than for debt, for reasons outlined above. An intermediate case, where there is a significant but not substantial risk of deficit bias, is discussed below in the case of the UK.

These targets should be set in cooperation with a fiscal council, which would monitor progress towards the target. The fiscal council could also, in extreme circumstances, be able to suggest that the target be revised before the term of the government is over. Persistent over or under shooting of targets because of past shocks can be dealt with as part of this process of setting targets – it makes little sense to incorporate short periods of debt correction as in a ‘debt brake’.

Within a fixed exchange rate regime or monetary union, the rule should also include some form of relative stabilisation term, such as the level of inflation compared to the union average. So if inflation was above the union average, the target for the deficit should be undershot, for example. There should be a ‘knockout’ if there is a chance of hitting the Zero Lower Bound (ZLB), the details of which are discussed below. This type of regime would appear appropriate for a number of Eurozone governments, and could help balance the need for strong, externally monitorable constraints on fiscal policy (given past history and the current political dynamics within the Eurozone between "surplus" and "deficit" countries) and the need for fiscal policy to accommodate (at least) rather than exacerbate already severe contractions in some countries.

For reasons outlined in section 5, these rules should not apply if interest rates have hit the zero lower bound. In that case fiscal policy should focus on demand stabilisation whatever the exchange rate regime. Exactly how this might happen is discussed above, but they represent the suspension of the rule while the ZLB constraint applies, and not its modification. This ‘knock out’ should be an explicit part of the rule. Problems of inadequate public investment are best dealt with by specifying separate targets for that investment, rather than adapting the fiscal rule or exempting investment spending from it.

We can use the UK as an example of how these rules might be applied to a particular country with its own monetary policy. Experience over the last few decades suggests that UK governments have not been subject to sustained deficit bias: the increase in debt that has occurred since 2008 was largely the result of an extraordinary shock and quite legitimate attempts to mitigate its impact. This suggests an appropriate framework for the UK, once interest rates rise above the zero lower bound, would be a rolling medium-term target, for the actual rather than cyclically adjusted primary deficit. This rolling target would be more robust to shocks than a target that specified a fixed date, and it also avoids the need for cyclical adjustment. However historical experience also suggests periods in which UK governments have tried to circumvent their own rules (either the actual rules or the intention of those rules), so the UK's fiscal council, the OBR, should be given a much wider remit to certify that the government is not manipulating the rolling nature of the target. Experience also suggests a tendency to put too much weight on cuts in public investment during fiscal retrenchment, so there should be a specific target for the ratio of public investment to GDP. The OBR should also be involved in comparing the implications of alternative deficit targets for the long run level of government debt.

References

- Barro, R. J. (1979): "On the Determination of Public Debt," *Journal of Political Economy*, 87, 940—971.
- Benigno, P., and M. Woodford (2004): "Optimal Monetary and Fiscal Policy: A Linear-Quadratic Approach," in *NBER Macroeconomics Annual 2003*, pp. 271—333.
- Ricardo J. Caballero and Emmanuel Farhi (2013), *A Model of the Safe Asset Mechanism (SAM): Safety Traps and Economic Policy*, [mimeo](#), Harvard University.
- Calmfors, L. , and Wren-Lewis, S (2011), *What Should Fiscal Councils Do?*, *Economic Policy*, Vol 26, pp 649-695
- Dolphin, T (2013) [How to conduct fiscal policy in the next parliament](#), IPPR, March.
- Eggertsson, G.B and Woodford, M (2003) "[Optimal Monetary Policy in a Liquidity Trap](#)," [NBER Working Papers](#) 9968, National Bureau of Economic Research, Inc.
- Eggertsson, G.B. and Woodford, M. (2004), 'Optimal monetary and fiscal policy in a liquidity trap', *NBER International Seminar on Macroeconomics*.
- Friedman, M (1993) The 'plucking model' of business fluctuations revisited, *Economic Inquiry* 31, 2
- Geier, A (2011), The Debt brake – the Swiss fiscal rule at the federal level, *Federal Finance Administration Working Paper* No.15
- Kirsanova, T, Satchi, M, Vines, D and Wren-Lewis, S (2007), Optimal Fiscal Policy Rules in a Monetary Union, *Journal of Money, Credit and Banking*, vol 39 pp 1759-1784
- Kirsanova, T, Leith, C and Wren-Lewis, S (2007), Optimal Debt Policy, and an Institutional Proposal to help in its Implementation, *European Economy Economic Papers No 275, April*
- Kirsanova, T and S. Wren-Lewis (2012) 'Optimal Fiscal Feedback on Debt in an Economy with Nominal Rigidities', *Economic Journal*, Vol. 122, pp 238-264
- Kirsanova, T., Leith, C. and Wren-Lewis, S. (2009), 'Monetary and fiscal policy interaction: the current consensus assignment in the light of recent developments', *Economic Journal*, 119, pp. 482–96
- Kopits, G and Symansky, S.A. (1998) [Fiscal Policy Rules](#), [International Monetary Fund](#)
- Leith, C, Moldovan, I and Wren-Lewis, S (2011) *Debt Stabilisation in a Non-Ricardian Economy*, *Oxford Discussion Paper* No. 542
- Marcet, A. and Scott, A. (2008). Debt, deficits and the structure of bond markets, *Journal of Economic Theory*, vol. 144(2), pp. 473–501.
- Mash, Richard (2010) "Tax smoothing and asymmetric shocks" *mimeo*, New College, Oxford.

Schmitt-Grohe, S., and M. Uribe (2004): “Optimal fiscal and monetary policy under sticky prices,” *Journal of Economic Theory*, 114, 198—230.

Schmitt-Grohe, S., and M. Uribe (2007): “Optimal Simple and Implementable Monetary and Fiscal Rules,” *Journal of Monetary Economics*, 54(6), 1702—1725.

Stern, N. (2006), *The Economics of Climate Change*, Cambridge University Press.

Woodford, M (2012) [“Methods of Policy Accommodation at the Interest-Rate Lower Bound.”](#) presented at the Jackson Hole symposium, August 2012.

Wren-Lewis, S (2010), Macroeconomic Policy in light of the credit crunch: the return of counter-cyclical fiscal policy? *Oxford Review of Economic Policy*, vol 26 pp 71-86

Wren-Lewis, S (2013) Aggregate fiscal policy under the Labour government, 1997–2010, *Oxford Review of Economic Policy* 29 (1): 25-46

Wren-Lewis, S (2013b) Macroeconomic Stabilisation in the Eurozone: Lessons from Failure, *Global Policy*, volume 4, [Issue Supplement s1](#), pages 66–73.

Wren-Lewis, S (2013c) Comparing the Delegation of Monetary and Fiscal Policy in [Restoring Public Debt Sustainability: The Role of Independent Fiscal Institutions](#) ed George Kopits, Oxford University Press.

Wyplosz, C (2012) Fiscal Rules: Theoretical Issues and Historical Experiences, NBER Working Papers 17884.