1 Policy Issues - Microfounding the objectives of macroeconomic stabilisation policy

1.1 The general problem

After the great recession of the 1930s, a key objective of macroeconomic policy was to avoid widespread involuntary unemployment. As inflationary tendencies grew in the 1960s and 1970s, a central aim of macroeconomic policy was to control inflation. We can formalise these ideas by writing down an *objective function* for policy makers. The policy problem can be posed in the following way: move the *instruments* of policy to achieve the best value of the objective function.[1] This process is termed *optimal control*, and it essentially involves a maximisation problem for policy makers, where the objective function is being maximised, and the constraint is the behaviour of the economy.

There are two problems with this approach which have motivated very recent research, by Woodford in particular (see Woodford (2003) Interest and Prices: Foundations of a Theory of Monetary Policy, Princeton: Princeton University Press.). First, if we have more than one objective, how do we trade-off these objectives against each other? How important is inflation relative to output stability (and hence unemployment) , for example? And if our goal is to stabilise output, what level should it be stabilised around: some trend, or something more sophisticated? Second, as macroeconomics became microfounded, models typically used representative agents with explicit utility functions. How did the objectives of policy makers relate to agent's utility? Surely there should be a connection. This is a specific way of asking a more general question: what guidance can macroeconomists give policy makers about what their objectives should be? How do objective functions relate to *social welfare*?

The lack of any connection between policy objectives in macroeconomic models and agent's utility also made it difficult to relate stabilisation policy to other economic objectives. Suppose, for example, that income or sales taxes were used to respond to macroeconomic shocks. While this might reduce inflation and stabilise output, these taxes are also distortionary. How do we compare the distortionary effects of taxation with any stabilisation benefits?

So why not simply have macro policy aiming to maximise utility? I think there were two important reasons why this was not done in the past. The first, which remains unresolved, is that standard macroeconomic models do not capture the disutility of involuntary unemployment. The second is that it appeared, in most macroeconomic models, that agent's utility was independent of inflation. This was the issue that was addressed by Woodford.

1.2 Deriving social welfare from agents' utility

Suppose the i'th individual agent maximises

[1] Alternatively, we can talk about minimising a loss function, but the process is formally identical.

$$\max_{\{C_{i,s},n_{i,s}\}_{s=t}^{\infty}} U_{i,t} = E_t \sum_{s=t}^{\infty} \beta^{s-t}[u(C_{i,s},\xi_{i,s}) - v(n_{i,s},\xi_{i,s})] \quad (1)$$

Here C is consumption of a basket of goods, where agents are free to choose between goods, n is the supply of labour, and ξ is a preference shock (so an increase in ξ_s.implies a desire to consume more goods *and leisure* in period s). We will only need to make use of one first order condition from this maximisation problem, which is

$$\frac{v_n(n,\xi)}{u_C(C,\xi)} = \frac{W}{P_C} \quad (2)$$

where P_C is the consumer price index, and W the wage (which we assume is identical across goods). This is the standard work/leisure trade off. Here $u_C > 0, v_n > 0, u_{CC} < 0, v_{nn} > 0$. (Note that with log utility for consumption, we have $u_C = \frac{\xi}{C\xi} = \frac{1}{C}, u_{C\xi} = 0, u_{CC} = -\frac{1}{C^2} = -\frac{u_C}{C}$. As always, assuming log utility simplifies matters, but we do not want to restrict ourselves to a particular utility function here, but instead derive some general propositions about policy.) Here and below we drop the time script if all variables are coincident.

If agents are uniformly distributed over the unit interval, then the social welfare function is

$$\Gamma_t = \int_0^1 U_{i,t} di$$

The first question we need to address is how to agents differ. In his baseline exposition, Woodford makes the realistic assumption that each agent supplies a unique type of labour, and there is a one to one mapping between labour types and types of good. As a result, the equation above involves integrating over goods. However, the key points can be made in the simpler, but less realistic case, that each agents supplies some labour to every good produced. In that case, all agents are completely identical, and so social welfare is indentical to individual welfare. (We note the implications of the more realistic alternative examined by Woodford below.)

Lets assume that the aggregate consumption basket is given by

$$C = \left[\int_0^1 c(z)^{\frac{\epsilon-1}{\epsilon}} dz\right]^{\frac{\epsilon}{\epsilon-1}} \qquad \epsilon > 1 \quad (3)$$

where z is an index of goods, uniformly distributed along the unit interval. This is a CES aggregator. It gives rise to a demand curve

$$c(z) = \left(\frac{p(z)}{P}\right)^{-\epsilon} C \quad (4)$$

where P is a measure of competitors prices defined as (see appendix)

$$P = [\int [p(z)^{1-\epsilon}]dz]^{1/1-\epsilon}$$

We want to define a measure of aggregate output. In a closed economy without government, Y=C, so a natural way to define aggegate output (see appendix) is

$$Y = [\int_0^1 y(z)^{\frac{\epsilon-1}{\epsilon}} dz]^{\frac{\epsilon}{\epsilon-1}}$$

This equation actually contains an important message. Suppose we have a fixed labour force - how should this be distributed across goods to maximise Y (and C)? The answer is to produce equal amounts of each good. You might think that this will be what happens in equilibrium anyway, if we assume all firms are alike. However, this ignores inflation. Lets assume that firms are monopolistic competitors, who set prices according to Calvo contracts. That is, each period there is a fixed probability that each firm will change its price, and each firm faces a demand curve for its product. Suppose we have a steady state with a positive inflation rate. In that steady state, some firms will change their price each period, but others will not, because of Calvo contracts. As a result, relative prices will be changing. These changes in relative prices have nothing to do with changes in the supply or demand for individual goods, but instead reflect a combination of inflation and Calvo contracts. Changes in relative prices imply that the outputs of individual goods will not be equal. This will in turn reduce aggregate output below what it might otherwise have been.

Thus Calvo contracts mean that inflation generates variability in output *across goods,* which means that labour is inefficiently used. We now have a direct link between inflation and utility. If inflation was zero, this inefficiency would disappear, and individuals would have to work less hard to produce goods to consume. This occurs, even though all workers work in all firms. As I noted above, Woodford actually assumed that workers were tied to the production of a specific good. This produces a second reason why variability in output (and therefore inflation) is costly. Not only is output inefficiently low, but some workers will be working too much, and others two little, and this will have utility costs.

This idea, that one important cost of inflation is the relative price variability that it induces, is far from new. Woodford's achievement was to show how it could be captured directly in models that involved Calvo contracts. What Woodford did was to take a second order Taylor expansion of equation (1) around a zero inflation steady state. First, lets focus on the 'per period' part of equation (1)[2]:

$$\Theta_s = u(C_s, \xi_s) - v(n_s, \xi_s) \tag{5}$$

A second order Taylor expansion of the first term (and droping the common

[2] If each worker produced just one good, this would become $\Theta_s = u(C_s, \xi_s) - \int_0^1 [\tilde{v}(n_s(z), \tilde{\zeta}_s)]dz$

time subscript) is given by (see appendices)

$$\begin{aligned} u(C,\xi) &= u(\bar{C},1) + \hat{C}\bar{C}u_C(\bar{C},1) + \xi u_\xi(\bar{C},1) + \hat{C}\xi\bar{C}u_{\xi C}(\bar{C},1) + \\ &\frac{1}{2}\hat{C}^2\{u_C(\bar{C},1)\bar{C} + u_{CC}(\bar{C},1)\bar{C}^2\} + \frac{1}{2}\hat{\xi}^2\{u_\xi(\bar{C},1) + \frac{1}{2}u_{\xi\xi}(\bar{C},1)\} + O[(\hat{C},\hat{\xi})^3] \end{aligned}$$

where we use the following notation: $\bar{X}$ for the steady state value of X (the steady state value of the preference shock is unity), $\hat{X}$ is the log deviation from steady state, and $O[(\hat{C},\hat{\xi})^3]$ (or simply $O[^3]$) to denote terms that are third order (in $\hat{C}$ and $\hat{\xi}$) or above. As we are looking at small deviations from a steady state, we assume that third order or above terms are sufficiently small that we can ignore them.

We can simplify this expression in two ways. First, all terms in utility and its derivatives are evaluated at steady state levels, so we can use $u(.)$ rather than $u(\bar{C},1)$. Second, stabilisation policy has no influence over steady state levels or shocks, so we have no interest in these terms from the point of view of policy action. We therefore collect such terms, plus terms of order greater than two, under the heading *tip* . This expression can therefore be rewritten as

$$\begin{aligned} u(C,\xi) &= \hat{C}\bar{C}u_C(.) + \hat{C}\hat{\xi}\bar{C}u_{\xi C}(.) + \frac{1}{2}\hat{C}^2\{u_C(.)\bar{C} + u_{CC}(.)\bar{C}^2\} + tip \quad (6) \\ &= \bar{C}u_C(.)\{\hat{C}(1 + \frac{\hat{\xi}u_{\xi C}(.)}{u_C(.)}) + \frac{1}{2}\hat{C}^2(1 + \frac{u_{CC}(.)\bar{C}}{u_C(.)})\} + tip \end{aligned}$$

Thus our general functional form for the utility of consumption can be represented by terms in the level of consumption and a quadratic terms in consumption, provided we are correct in being able to ignore all terms collected under *tip*.

We can do the same for the disutility of labour in equation (5):

$$v(n,\xi) = \bar{n}v_n(.)\{\hat{n}(1 + \frac{\hat{\xi}v_{\xi n}(.)}{v_n(.)}) + \frac{1}{2}\hat{n}^2(1 + \frac{v_{nn}(.)\bar{n}}{v_n(.)})\} + tip$$

We can relate labour supply to output in the following way

$$\hat{n} = \hat{Y} + \frac{\epsilon}{2}Var_i\{p(i)\} + O[^3]$$

where $Var_i\{p(i)\}$ is the variance of prices across goods[3]. This captures the point made earlier, that variability in prices has an efficiency cost, in that it 'uses up' labour at the cost of less output. In substituting this into the previous

[3]The proof is quite complex. However, variances and integrals are simply connected through the definitions of expectations i.e.

$$E_z\{x\} = \int xdz, Var_z\{x\} = \int (x - E_z\{c\{x\})^2 = \int x^2dz - (E_z\{x\})^2 \quad (7)$$

expression, note that $Var_i\{p(i)\}$ is itself second order, so $\hat{n}^2$ will only involve $\hat{Y}^2$ to second order. We then get

$$v(n,\xi) = \bar{n}v_n(.)\{\hat{Y}(1+\frac{\hat{\xi}v_{\xi n}(.)}{v_n(.)}) + \frac{1}{2}\hat{Y}^2(1+\frac{v_{nn}(.)\bar{n}}{v_n(.)}) + \frac{\epsilon}{2}Var_i\{p(i)\}\} + tip \quad (8)$$

The analysis so far has introduced two distortions: Calvo contracts, and monopolistic competition. The latter on its own will imply that output is inefficiently low, even if inflation is zero (see earlier handout on imperfect competition). It is convenient to abstract from this monopolistic distortion (for reasons that will become clear in the handout on time inconsistency), which we can do by assuming a production subsidy which is set at a level that induces firms to produce at the efficient, perfectly competitive level. This allows a further simplification. In a closed economy, $C = Y$. In addition, under perfect competition the real wage in steady state will be unity, so

$$u_c(.) = v_n(.)$$

This, coupled with $n = y$, means that the first term in equation (6) cancels with the first term in (8).

We now define *natural* levels of variables. This is the level the variable would be at if prices were completely flexible. With a production subsidy that eliminates the monoploy distortion, the natural level is also the efficient level. Natural levels are *tip*, and because they are efficient they reflect a natural target for stabilisation policy. We can show that the following is true:

$$\hat{Y}^n\bar{Y}(v_{nn}(.) - u_{CC}(.)) + v_{n\zeta}(.)\hat{\xi} - u_{C\varepsilon}(.)\hat{\xi} + O[^2] = 0 \quad (9)$$

where $\hat{Y}^n$ is the natural level of output (see appendix). This allows us to eliminate the preference shock from our second order expansions. We are left with terms in $\hat{Y}\hat{Y}^n$, but noting that $(\hat{Y}-\hat{Y}^n)^2 = \hat{Y}^2 - 2\hat{Y}\hat{Y}^n + tip$ we can simplify to obtain the following expression

$$\Theta = -\frac{\bar{Y}^2}{2}\{(\hat{Y}-\hat{Y}^n)^2(v_{nn}(.) - u_{CC}(.)) + Var\{p(i)\}\frac{u_C(.)}{\epsilon\bar{Y}}\} + tip$$

Thus we have reduced per-period social welfare to a quadratic term in $\hat{Y}-\hat{Y}^n$, and a term in the variability in prices over goods. Finally, Woodford shows that when you put this back into the measure for discounted social welfare, this variance term can be replaced by a quadratic term in inflation. (We noted above why it was inflation that generated relative price variability.) Thus social welfare to second order is a quadratic function of the difference between output and its natural level, and a quadratic term in inflation.

1.3 Implications for policy

We therefore get a remarkable result: the kind of objective functions that economists typically assumed on an ad hoc basis can actually be derived as social welfare functions derived from the utility of representative agents. This is clearly reassuring, although we should note that Calvo contracts play a crucial role in this, so the derivation is model specific in this sense. However, there are further implications.

In the introduction I noted a number of questions associated with ad hoc objective functions. We can now address these directly. First, what should output be stabilised around? The answer supplied above is the natural level of output. Recall that this is the level of output that would occur under flexible prices, which is also the efficient level of output with a subsidy in place. What does the natural level of output depend on? It is clearly not a simple trend. First, it will be influenced by technology: a technology shock will raise $\hat{Y}^n$. So it would be quite wrong to try and suppress an increase in output if it was the result of a positive technology shock. (We typically call this a supply side shock. Note that a technology shock will still mean that $\hat{Y}$ may not equal $\hat{Y}^n$, so policy will still have something to do.) More interestingly, $\hat{Y}^n$ will also depend on preference shocks. Thus, if people want to produce less because they want to enjoy more leisure, it would again not be sensible for stabilisation policy to fight this.

All this has a very natural interpretation. Stabilisation policy should attempt to reproduce the flex price equilibrium. This follows directly from the fact that the only distortion in the model is price stickiness. Note, however, that this assumes the existence of a subsidy to knock out the monopoly distortion. In more realistic settings we will have a *distorted steady state*, and so the efficient level of output will not be the flex price level.

A second question was what the trade-off between output stabilisation and control of inflation should be. Again, this analysis provides an answer - it depends on the parameters in the model (not just preference parameters, but also demand curve parameters). In most applications analysed to date, using standard calibrations, it appears that the inflation term is dominant. However, this could easily change as models incorporate additional elements. We should note that this analysis, while successfully incorporating inflation into social welfare, has not addressed the issue of the hetrogeneity of unemployment.

The analysis also has a very clear implication about what the optimum rate of inflation should be - zero. Only when prices are zero do we avoid the sticky price distortion. While this may seem an obvious conclusion, it can be criticised in both directions. Friedman, in a celebrated analysis, suggested that the optimum rate of inflation should be negative. This was because he argued the optimum nominal rate of interest should be zero, which with positive real interest rates implied negative inflation. The reason for wanting zero nominal interest rates was that this eliminated the cost of holding money. On the other hand, many argue that the optimal inflation rate should be small and positive, because this avoids liquidity traps following shocks, and because it may avoid other rigidities associated with falling prices.

1.4 Extensions and Applications

A key point about the analysis above is that it is model specific. If we use a different model, the social welfare function may change (or may not). As a result, it has now become standard to augment any new model with a derivation of social welfare implied by that model. For example, would the social welfare function derived for an open economy be similar to that for a closed economy? The answer appears to depend on the particular open economy model used. However, one result does appear reasonably general, and that is that policy should attempt to minimise a measure of inflation based on output prices, rather than consumer prices.

A consumer price index contains prices of imported goods as well as home produced goods, whereas an output price index only contains the prices of goods produced domestically, including those sold for export. In the analysis above, inflation mattered because it lead to variability in relative prices, which in turn lead to variability in output, which meant that labour was used inefficiently. It follows directly that the relevant measure of inflation from this welfare point of view is output price inflation, not consumer price inflation. However, when governments have an explicit inflation target, this is always based on a consumer price index. There appears to be an interesting contrast between theory and practice here.

Woodford's analysis, and subsequent extentions, are interesting in part because they relate utility maximisation to more traditional objective functions, but also because it helps us understand what factors influence social welfare. However, it remains equivalent to simply maximising the utility of the representative agent. Many recent studies of macroeconomic policy simply maximise utility directly, without going through the transformations outlined above. This is often discribed as solving the *Ramsey problem*, and links directly to similar techniques used in the public finance literature on optimal taxation. In contast, Woodford's technique is often described as the *linear quadratic* approach. An advantage of the linear quadratic approach is that it suggests some general implications for what policy makers should do, that do not require us to choose some particular form for agents utility. However, some of the assumptions required to obtain these results are quite strong (in particular, the assumption of a non-distorted steady state), and we do not know yet how robust they are to relaxing these assumptions.

1.5 Mathematical appendices

1.5.1 The price index implied by a CES aggregator

$$PC = \int p(z)c(z)dz$$

Using the demand curve this becomes

$$\begin{aligned} PC &= \int [p(z)\left(\frac{p(z)}{P}\right)^{-\epsilon} C]dz \\ &= CP^{\epsilon}\int [p(z)^{1-\epsilon}]dz \\ P^{1-\epsilon} &= \int [p(z)^{1-\epsilon}]dz \\ P &= [\int [p(z)^{1-\epsilon}]dz]^{1/1-\epsilon} \end{aligned}$$

1.5.2 The aggregate output index

Y is an index of output. We should be able to write the demand curve as

$$y(z) = (\frac{p(z)}{p})^{-\epsilon} Y$$

which implies

$$p(z) = P(\frac{y(z)}{Y})^{-1/\epsilon}$$

For consistency

$$PY = \int_0^1 p(z)y(z)dz$$

Substituting the demand curve into the above gives

$$\begin{aligned} PY &= \int_0^1 y(z)P(\frac{y(z)}{Y})^{-1/\epsilon}dz \\ &= PY^{1/\epsilon}\int_0^1 y(z)^{1-1/\epsilon}dz \\ Y^{1-1/\epsilon} &= \int_0^1 y(z)^{\frac{\epsilon-1}{\epsilon}}dz \\ Y &= [\int_0^1 y(z)^{\frac{\epsilon-1}{\epsilon}}dz]^{\frac{\epsilon}{\epsilon-1}} \end{aligned}$$

1.5.3 Second order expansions

A second order Taylor expansion of a function of a single variable x around the point x_0 is given by

$$f(x) = f(x_0) + f_x(x_0)(x - x_0) + \frac{1}{2}f_{xx}(x_0)(x - x_0)^2 + O[(x)^3]$$

where f_x is the first derivative, f_{xx} is the second derivative, $f(x_0)$ is the value of the function at $x = x_0$, and $O[(x)^3]$ stands for terms of order x^3 or higher. As a simple example, take $f(x) = x^2$, evaluated at $x = 2$. The expansion gives

$x^2 = 4 + 4(x-2) + (x-2)^2$. Evaluate this at $x = 3$. We know the answer must be 9, and it is.

A second order Taylor expansion of a function of two variables $f(x,y)$ around the point (x_0, y_0) is given by

$$\begin{aligned} f(x,y) = &f(x_0,y_0) + f_x(x_0,y_0)(x-x_0) + f_y(x_0,y_0)(y-y_0) \\ &+ \frac{1}{2}f_{xx}(x_0,y_0)(x-x_0)^2 + \frac{1}{2}f_{yy}(x_0,y_0)(y-y_0)^2 \\ &+ f_{xy}(x_0,y_0)(x-x_0)(y-y_0) + O[(x,y)^3] \end{aligned} \tag{10}$$

1.5.4 Log linear expansions

We are interested in log linear deviations from steady state, so we use the following second order approximation:

$$\begin{aligned} Let \quad y &= e^x, \ln y = x \\ e^x &= e^{x_0}[1 + (x-x_0) + \frac{1}{2}(x-x_0)^2 + O[(x-x_0)^3] \\ so \quad y &= y_0[1 + \ln(\frac{y}{y_0}) + \frac{1}{2}\ln(\frac{y}{y_0})^2 + O[\ln(\frac{y}{y_0})^3]] \\ y - y_0 &= y_0[\ln(\frac{y}{y_0}) + \frac{1}{2}\ln(\frac{y}{y_0})^2 + O[\ln(\frac{y}{y_0})^3]] \end{aligned}$$

Let us denote, for any y, $\hat{y} = \ln(y/y_0)$.Thus

$$y = y_0 + y_0\{\hat{y} + \frac{1}{2}\hat{y}^2 + O[\hat{y}^3]\} \tag{11}$$

Appling this to 10 gives

$$f(x,y) = f(x_0,y_0) + f_x(x_0,y_0)x_0\{\hat{x} + \frac{1}{2}\hat{x}^2 + O[\hat{x}^3]\} + f_y(x_0,y_0)y_0\{\hat{y} + \frac{1}{2}\hat{y}^2 + O[\hat{y}^3]\} \tag{12}$$

$$\begin{aligned} &+ \frac{1}{2}f_{xx}(x_0,y_0)x_0^2\{\hat{x} + \frac{1}{2}\hat{x}^2 + O[\hat{x}^3]\}^2 + \frac{1}{2}f_{yy}(x_0,y_0)y_0^2\{\hat{y} + \frac{1}{2}\hat{y}^2 + O[\hat{y}^3]\}^2 \\ &+ f_{xy}(x_0,y_0)y_0x_0\{\hat{x} + \frac{1}{2}\hat{x}^2 + O[\hat{x}^3]\}\{\hat{y} + \frac{1}{2}\hat{y}^2 + O[\hat{y}^3]\} + O[(x,y)^3] \\ &= f(x_0,y_0) + \hat{x}x_0f_x(x_0,y_0) + \hat{y}y_0f_y(x_0,y_0) + \hat{x}\hat{y}y_0x_0f_{xy}(x_0,y_0) \\ &+ \frac{1}{2}\hat{x}^2\{f_x(x_0,y_0)x_0 + f_{xx}(x_0,y_0)x_0^2\} + \frac{1}{2}\hat{y}^2\{f_y(x_0,y_0)y_0 + \frac{1}{2}f_{yy}(x_0,y_0)y_0^2\} + O[(x,y)^3] \end{aligned} \tag{13}$$

1.5.5 Preference shocks and natural output

A first order condition from the agent's maximisation of utility is

$$\frac{v_n(n),\xi)}{u_C(C,\xi)} = W \tag{14}$$

where W is the real wage. The log linear expansion of $f(x,y) = x/y$ is given by

$$\begin{aligned} x/y &= x_0/y_0 + \hat{x}x_0/y_0 - \hat{y}y_0x_0/y_0^2 + \frac{1}{2}\hat{x}^2x_0/y_0 - \frac{1}{2}\hat{y}^2x_0y_0/y_0^2 \\ &= x_0/y_0[1 + \hat{x} - \hat{y} + \frac{1}{2}\hat{x}^2 - \frac{1}{2}\hat{y}^2] \end{aligned} \tag{15}$$

Letting $y = n$, an applying this to (14) up to first order only, implies

$$\ln(W) = \frac{v_n(.)}{u_C(.)}(1 + v_{nn}(.)\hat{y}\bar{y} + v_{n\zeta}(.)\hat{\xi} - u_{CC}(.)\hat{C}\bar{C} - u_{C\varepsilon}(.)\hat{\xi}) + O[^2] \tag{16}$$

Finally, if prices were flexible $W = 1$,so $\ln(W) = 0$.