OPTIMAL FISCAL POLICY IN A LIQUIDITY TRAP

Paul Krugman Dec. 29, 2008

One thing that's been conspicuously missing from the back-and-forth among economists about fiscal stimulus is anything resembling a fully specified model. To some extent that's OK – you can't really demand that every aspect of policy, especially in response to a crisis, be based on careful analysis of first-order conditions. But there are dangers in relying entirely on implicit theorizing; you can find yourself saying things you think must be true, but that turn out not to be true even in a simple model with maximizing agents.

Case in point: I went into the liquidity trap business believing that the concept of such a trap was nonsense – that you can always drive up prices by printing money. It wasn't until I wrote down a very simple maximizing model that I realized that this was wrong, that a monetary expansion perceived as temporary may be entirely ineffective (see http://web.mit.edu/krugman/www/japtrap.html).

In my home field of international economics, the so-called New Open Economy Macroeconomics, which builds on New Keynesian economics, has had a big impact, mainly because it gives you a systematic way to talk about the welfare effects of monetary policy. The models aren't realistic: they basically assume full intertemporal maximizing behavior by everyone, except for some short-run price stickiness that gives rise to short-run real effects of monetary shocks. But even unrealistic models can do a lot to clarify your thoughts.

What I want to do in this note is sketch out a New Keynesian-type analysis of the role of fiscal policy in a liquidity trap. The bottom line is quite striking: aside from some qualifications I'll discuss at the end, when the economy is in a liquidity trap *government spending should expand up to the point at which full employment is restored*. That's not a guess or a statement of personal preferences, it's a result.

The basic intuition behind this result is that when the economy is in a liquidity trap, the social marginal cost of government spending is low, because there isn't enough private demand to fully employ the economy's resources. This means that we would normally expect more government spending to raise welfare, right up to the point that full employment (a concept that needs a bit of explanation here) is restored. At that point the marginal cost of government spending jumps up, because it's diverting resources from private spending.

OK, let me sketch out the model. I'll try a more careful and better written version later, but right now I just want to get this out there.

Start with a typical New Keynesian model of the type used in, say, Obstfeld and Rogoff's 1996 paper "Exchange rate dynamics redux", (which is in turn heavily based on Blanchard and Kiyotaki's 1987 "Monopolistic competition and the effects of aggregate demand".) In this kind of model consumers maximize an intertemporal utility function, like this one (specific functional forms aren't that important):

$$(1) \quad U = \sum_{t=0}^{\infty} \beta^t \left[\mathrm{s_t} \ln(C_t) + \chi \ln\left(\frac{M_t}{P_t}\right) + v(G_t) - \kappa {\ell_t}^2 \right]$$

Here C is private consumption, M the quantity of money held, G real government spending on public goods, and the last term is labor effort. I've introduced a shift factor s on consumer demand, as a crude way to allow for the kind of negative demand shock we're experiencing right now. G is paid for with lump-sum taxes; as I'll explain later, lifting this assumption offers one way to soften the results.

We think of each worker as being the sole producer of a differentiated product – I think of them as "yeoman entrepreneurs" – who are monopolistically competitive. Consumers and the government purchase an aggregate of these differentiated products:

$$(2)\ C_t + G_t = \left[\int_0^{\infty} {\ell(z)_t}^{\rho} \, dz \right]^{1/\rho}$$

What the monopolistic competition thing does is make price stickiness comprehensible: the gains to an individual of correction small deviations of his price from the optimum are only second-order in the deviation, which means that small costs of changing prices can make prices sticky. But the aggregate welfare implications of price stickiness are first-order: reduce the money supply, and if prices don't fall there are first-order welfare losses.

Private behavior in this model can be summarized by three first order conditions:

$$(3) \quad \frac{M_t}{P_t} = \chi C_t \frac{i_t}{1 + i_t}$$

(Money demand depends on consumption and the nominal interest rate)

$$(4)\ C_t = \left(\frac{1}{\beta}\right) \frac{s_t}{s_{t+1}} C_{t+1} \frac{P_{t+1}}{P_t} \frac{1}{1 + i_t}$$

(Euler condition relating consumption today to future consumption and real interest rates)

$$(5)\ \frac{s_t}{C_t} = \rho \kappa \ell_t$$

(Labor supply condition setting marginal disutility of labor equal to $\rho<1$ times marginal utility of consumption -- wedge is because of monopolistic competition)

The "Keynesian" thing about the model is the assumption that (3) holds only in expectation – that producers must set their prices before they know the money supply and possibly other things, such as the consumption shift term. This temporary price stickiness gives monetary policy real effects.

Now, the way we typically do policy experiments in this model is to assume that from period 2 onwards everything is in a steady state – constant M, s constant, G constant. This ties down P and C in period 2 and later; this in turn makes (4) a kind of IS curve, with real consumer demand depending on the real interest rate. And since future P is tied down while current P is predetermined, it's also an IS curve in terms of the nominal interest rate.

Meanwhile, (3) is a sort of LM curve; with P predetermined, the monetary authority is able to set the interest rate by changing the money supply.

Under normal circumstances, the effectiveness of monetary policy means that the central bank can always ensure "full employment" – which I take to mean that condition (5) holds, that the level of employment is the same as it would be if price-setters had known in advance what the state of the economy would be. Actually, this kind of model implies that the central bank can raise welfare by creating "over-full" employment; that's because of the monopolistic competition wedge. But this in turn raises issues of time consistency, a la Barro-Gordon: if the central bank is known to seek over-full employment, this will cause inflationary expectations, etc.. So let me just assume that what the central bank does is to adjust the money supply to achieve full or normal employment, and no more.

Can the central bank do this? Take the level of government purchases G as given; from (2) and (5) this will tell you the level of C needed to achieve full employment; (4) will tell you the real interest rate needed to get that level of C; and since we already have both the current and future price levels tied down, this implies a necessary level of the nominal interest rate. So all the central bank has to do is increase the money supply until the rate is at the desired level.

But what if the required nominal rate is negative? In that case monetary policy can't get you there: once the interest rate hits zero, people will just hoard any additional cash – we're in the liquidity trap. The only way to make monetary policy effective once you're in such a trap, at least in this framework, is to credibly commit to raising future as well as current money supplies. That's the point I discovered back in my Japan's trap paper.

Which finally brings me to fiscal policy.

Under normal conditions, the central bank ensures that the economy is at full employment, so that increased government purchases come at the expense of reduced private consumption. Optimal fiscal policy sets the marginal utility of G equal to the marginal utility of C.

But suppose that the economy is operating below full employment, and is in a liquidity trap. C is tied down by the Euler condition (4): an increase in G does *not* crowd out private spending (there's also no multiplier effect, by the way – using this kind of perfect markets, immortal consumer model loads the dice against fiscal policy.) Increasing G does require more work effort, with the resulting disutility. But even at full employment, in these models, the disutility from an additional hour of work is less than the utility added by the extra consumption this work makes possible, and if the economy is depressed the gap is larger.

What this says is that if the economy is in a liquidity trap, the social marginal cost of G is substantially lower than it is under normal conditions. And that has a clear policy implication, illustrated here:

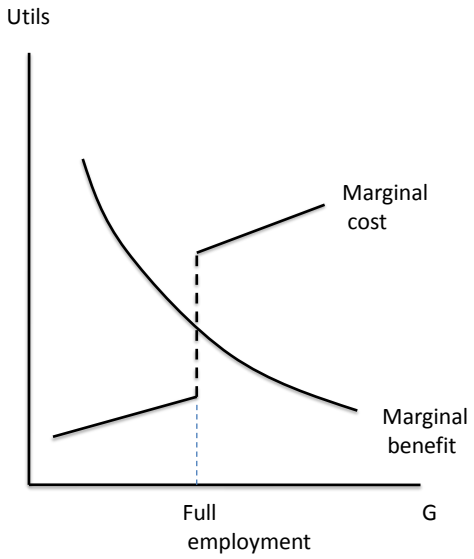

What I've illustrated here is the marginal cost and benefit of government purchases of public goods in and near a liquidity trap. The marginal benefit is presumably a downward-sloping curve. If G is low, so that monetary policy cannot achieve full employment, the marginal cost of an additional unit of G is low, because the additional government purchases don't crowd out private spending. Once G is high enough to bring full employment, however, any further rise in government purchases will be offset by a rise in the interest rate, so that extra G does come at the expense of C, implying a jump in the marginal cost.

As the figure suggests, there should be a fairly wide range of situations in which the optimal level of G is precisely the level at which the marginal cost jumps – that is, the optimal fiscal expansion is one that brings the economy right to full employment.

What could undo this result? I mentioned above that the assumption of lump-sum taxes may be important. What this assumption means is that we don't have to worry about government debt – the government can always raise taxes enough to service whatever debt is had, without our having to concern ourselves with any deadweight losses caused by future taxes. Obviously that's wrong, and debt concerns give a reason to limit fiscal expansion. On the other hand, that logic applies in good times as well as bad – i.e., debt concerns raise the marginal cost of G, but don't change the fact that there's a jump in the marginal cost at full employment. So there's still a presumption that G should rise when you're facing a liquidity trap, and probably rise enough to restore full employment.

The bottom line here is that while we usually think of Keynesianism as the preserve of ad hoc models, in this case doing it "right" – using a macromodel with maximizing agents and a proper concern for intertemporal constraints – actually suggests a very strong case for big government spending in the face of a liquidity trap. When the economy is depressed and monetary policy can't set it right, the true opportunity cost of government spending is low. So let's get those projects going.