

THINKING ABOUT THE LIQUIDITY TRAP

Paul Krugman

December 1999

We live in the Age of the Central Banker - an era in which Greenspan, Duisenberg, and Hayami are household words, in which monetary policy is generally believed to be so effective that it cannot safely be left in the hands of politicians who might use it to their advantage. Through much of the world, quasi-independent central banks are now entrusted with the job of steering economies between the rocks of inflation and the whirlpool of deflation. Their judgement is often questioned, but their power is not.

It is therefore ironic as well as unnerving that precisely at this moment, when we have all become sort-of monetarists, the long-scorned Keynesian challenge to monetary policy - the claim that it is ineffective at recession-fighting, because you can't push on a string - has reemerged as a real issue. So far only Japan has actually found itself in liquidity-trap conditions, but if it has happened once it can happen again, and if it can happen here it presumably can happen elsewhere. So even if Japan does eventually emerge from its slump, the question of how it became trapped and what to do about it remains a pressing one.

In the spring of 1998 I made an effort to apply some modern, intertemporal macroeconomic thinking to the issue of the liquidity trap. The papers I have written since have been controversial, to say the least; and while they have helped stir debate within and outside Japan, have not at time of writing shifted actual policy. Moreover, too much of that debate has been confused, both about what the real issues are and about what I personally have been saying.

The purpose of this paper is twofold. First, it is a restatement of what I believe to be the essential logic of liquidity-trap economics, with an emphasis in particular on how the "modern" macro I initially used to approach the problem links up with more traditional (and still very useful) IS-LM-type thinking. Second, it attempts to examine in a more or less coherent way the various alternative policies that either are in place or have been proposed to deal with Japan's liquidity trap, ranging from fiscal stimulus to unconventional open-market operations (and it tries in particular to make clear the difference between the latter and the expectations-focused inflation targeting I have

proposed).

1. *The liquidity trap: an IS-LM view*

Consider the sort of economy introduced a few chapters into most undergraduate macroeconomics books: an economy in which prices are for the moment assumed fixed, meaning both that there can be unemployment because of inadequate nominal demand, and that we need not make a distinction between the nominal and real interest rates. Since the classic 1937 paper by Hicks, it has been usual to summarize short-run equilibrium in such an economy by looking at two curves: a downward sloping IS that shows how lower interest rates increase the demand for goods and hence real output y ; and an upward-sloping LM curve that shows how increased output, by increasing the demand for money (whatever exactly that means in the modern world), drives up the interest rate. Monetary policy shifts LM, fiscal policy shifts IS.

Literally from the beginning of IS-LM analysis, however, Hicks realized that monetary policy might in principle be ineffective under "depression" conditions. The reason is that the nominal interest rate cannot be negative - otherwise, cash would dominate bonds as an asset. So at an interest rate near zero the demand for money must become more or less infinitely elastic, implying that the leftmost parts of the LM curve must actually be flat. And suppose that the IS curve happens to intersect LM in that flat region, as it does in [Figure 1](#) . Then changes in the money supply, which move LM back and forth, will have no effect on interest rates or output; monetary policy will be ineffective.

An alternative way to state this possibility is to say that if the interest rate is zero, bonds and money become in effect equivalent assets; so conventional monetary policy, in which money is swapped for bonds via an open-market operation, changes nothing.

I think that it is fair to say that for around two generations - from the point at which it became clear that the 1930s were not about to reemerge, to the belated realization circa 1997 that Japan really was back in a 30s-type monetary environment - nobody thought much about the deeper logic of the liquidity trap. But once it became clear that the Bank of Japan really did consider itself unable to increase demand in an economy that badly needed it, it also became clear (to me at least) that the theory of the liquidity trap needed a fresh, hard look.

I started with a preconception: that the idea of the liquidity trap was basically a red herring, that surely a determined central bank could always reflate the economy. Partly this preconception represented wishful thinking: having engaged in sometimes bitter arguments with "vulgar Keynesians" (e.g. the journalist William Greider (1997)) who believed that spending is always good and saving always bad, I was reluctant to concede that there might be circumstances under which they were right. But it also reflected my intuition - which turned out to be wrong - that the apparent possibility of such a trap in the IS-LM model was an artifact of that model's intellectual corner-cutting.

The IS-LM framework is, of course, an ad hoc approach that is strategically careless about a number of issues, from price determination, to the consequences of capital accumulation, to the determinants of consumer behavior. Most of the violence in the macro wars of the last generation has focused on aggregate supply; but since one must assume some kind of nominal price rigidity even to get into the discussion of Japan's demand-side problems, that was not the issue here. Rather, the apparent weakness of IS-LM was in its modeling of aggregate demand.

Here's how my initial argument - not that different from the debates between Keynes and Pigou - went. In the IS-LM model both the money supply and the price level enter in only one place: on the left-hand side of the money demand equation, which defines a demand for real balances M/P . Monetary policy and changes in the price level therefore affect aggregate demand through the same channel. And to say that increases in M were ineffective beyond some point was therefore equivalent to saying that reductions in P were ineffective in raising demand - that the aggregate demand curve looked something like AD in [Figure 2](#), downward-sloping over some range but vertical thereafter. And in that case even full price flexibility might not be enough to restore full employment.

But as Pigou pointed out, that simply cannot be right. If nothing else, a fall in the overall price level increases the real value of the public's holdings of money, and this wealth effect will increase consumption. If the IS-LM model seems to suggest that no full employment equilibrium exists, it is only because that model does not really get the budget constraints right. And it seemed to me that what went for P must go for M ; just as a sufficiently large fall in P would always expand the economy to full employment, so must a sufficiently large rise in M . It seemed to me to be a truism that increases in M always raise the equilibrium price level, and hence given a downwardly inflexible price level will always increase output.

To demonstrate the truth of that supposed truism, all that was needed was to write down a model that got the budget constraints right, that did not fudge the individual's decision problem. So I set out to write down the simplest such model I could. And it

ended up saying something quite different.

2. *Manna and money*

Instead of the rather complicated world of the IS-LM model, imagine a pure exchange economy. There is a single consumption good, which drops as manna from heaven, so that consumption in each period is a given; the representative individual sets out to maximize a utility function of the form

$$U = (1/\rho) \sum_t D^t c_t^\rho \quad (1)$$

In order to introduce money in a minimalist way, suppose that in order to purchase goods consumers must have cash in hand. Thus there is a cash-in-advance constraint each period of the form

$$P_t c_t \leq M_t \quad (2)$$

However, we simplify matters by assuming that additional cash may be acquired, or excess cash disposed of, in a money-for-bonds market that takes place at the beginning of each period. As long as there is no uncertainty, this implies that under normal circumstances the cash-in-advance constraint (2) will be binding. It also implies an Euler condition on consumption, the nominal interest rate, and prices:

$$(1 + i_t)(P_t / P_{t+1}) = (1/D)(c_{t+1}/c_t)^{-\rho} \quad (3)$$

It's not much of a model; but it does have endogenous interest rates and prices in it, has rational consumers, and fully respects budget constraints. If my intuition that the liquidity trap is an artifact of the IS-LM model's incompleteness were right, we ought to find that monetary policy always raises the price level in this model. Does it?

Suppose that we consider a change in the money supply only for the current period, via an open-market operation during that beginning-of-period asset market. We assume that the money supply in subsequent periods is restored to its original level, and that any implications of the operation for the government's budget constraint are taken care of via lump-sum taxes and transfers. Then we can take the evolution of the price level and the interest rate in subsequent periods as given, and focus only on equilibrium in the current period given P_{t+1} .

Equilibrium can then be illustrated, Patinkin-style, via a sort of IS-LM diagram in i, P space ([Figure 3](#)). The "IS" curve is defined by the Euler condition (3); the easiest way to think about it is to say that the real interest rate is given, so any rise in i must be offset by a fall in P relative to its future value, generating the expected inflation needed to keep the real rate constant. The "LM" curve is under normal circumstances defined by setting the cash-in-advance constraint (2) to equality, so that $Pc = M$. The intersection of the two curves then simultaneously determines the interest rate and the price level.

Unless, that is, the implied nominal interest rate is negative, as illustrated in Figure 3. What must happen in that case is that the cash-in-advance constraint ceases to be binding; in effect, some money is now held merely as a store of value, indistinguishable from bonds. And in that case any further increase in the money supply will have no effect; the economy will, in fact, be in the liquidity trap.

How should one think about this case? Perhaps the first thing to say is that the proposition of monetary neutrality, as usually stated, is not quite right. We normally say that if you double the money supply you double the equilibrium price level. The correct statement is that if you double the current *and all expected future* money supplies, you double the current price level; a monetary expansion perceived as temporary may have no effect at all.

The reason is that the economy has a maximum rate of deflation, equal to the "natural" real rate of interest. A temporary monetary expansion that tries to raise the current price level so high that it would have to be followed by deflation more rapid than this maximum rate will end up spilling over into excess, unused liquidity instead.

Now you may ask, why would a central bank try to impose deflation? But there is nothing in the logic of this exercise that says that the maximum rate of deflation must be positive. If, for whatever reason, the natural real rate of interest is negative, then the economy "wants" inflation. You may well ask why and how it should happen that the natural real rate is negative; but just for the moment suppose that it is. Then the economy will find itself in a liquidity trap as long as the private sector does not expect sufficiently rapid inflation.

It is important to notice that this does not mean that the old Keynesian idea that no full-employment equilibrium exists is validated. With fully flexible prices the economy will still manage to achieve full employment - but the mechanism is a bit unusual. Namely, the economy will get this inflationary expectations it needs via deflation - that is, by reducing the current price level compared with its expected future.

But in real life this won't be easy. If we suppose that there is some downward inflexibility of prices, the economy's need for inflation will manifest itself as a real slump, one that persists even if the nominal rate is reduced to zero.

Now in standard macroeconomics it is possible to compensate for downward price inflexibility by increasing the money supply instead: the economy doesn't have enough M/P , so if P won't go down just raise M instead. In the liquidity-trap case illustrated by our little model, however, raising current M is ineffective, essentially because it's a different ratio - P^e/P , where P^e is the expected future price level - that is out of line. But all is not lost for monetary policy. A credible commitment to expand not only the current but also *future* money supplies, which therefore raises expected future prices - or, equivalently, a credible commitment to future inflation - will still succeed in raising the equilibrium current price level and hence current output.

Let me say this perhaps more forcefully than I have in the past. Inflation targeting is not just a clever idea - a particular proposal that might work in fighting a liquidity trap. *It is the theoretically "correct" response* - that is, inflation targeting is the way to achieve in a sticky-price world the same result that would obtain if prices were perfectly flexible. Of course in policy the perfect is the enemy of the good, and I would not oppose trying a variety of tactics to fight Japan's stagnation. But it is inflation targeting that most nearly approaches the usual goal of modern stabilization policy, which is to provide adequate demand in a clean, unobtrusive way that does not distort the allocation of resources.

So the intertemporal approach led me to a different destination than I expected. I thought it would show that the liquidity trap was not a real issue, that without the inconsistencies of the IS-LM model it would become clear that it could not really happen. Instead it turns out that a liquidity trap can indeed happen; but that it is in a fundamental sense an expectational issue. Monetary expansion is irrelevant because the private sector does not expect it to be sustained, because they believe that given a chance the central bank will revert to type and stabilize prices. And in order to make monetary policy effective, at least in a simple model, the central bank must overcome a credibility problem that is the inverse of our usual one. In a liquidity trap monetary policy does not work because the markets expect the bank to revert as soon as possible to the normal practice of stabilizing prices; to make it effective, the central

bank must credibly promise to be irresponsible, to maintain its expansion after the recession is past.

3. Investment at home and abroad

For me, at least, the pure-exchange, manna and money model of the previous section was important as a mind-opener, as a way of laying bare the fundamental issues. But IS-LM remains in use for a reason: it is a much more convenient way to do back-of-the-envelope analysis of policy questions. Now that we know that the IS-LM version of the liquidity trap is, properly interpreted, more or less right we can return to address some further complications - in particular, two non-monetary challenges to the whole idea of a liquidity trap.

The first challenge is embodied in the old rhetorical question, why not fill in the Gulf of Mexico? (Or, in this case, the Sea of Japan – e.g., Kansai Int'l Airport). The point is that it is almost impossible to think of an economy in which there are literally no investment projects with a positive real rate of return; so how is it possible for the natural rate of interest to be negative?

The second challenge is a more modern, subtle one: why not invest abroad? If we think of a liquidity trap as a problem of excess desired saving (see below), why can't a country invest that excess saving in other countries with positive rates of return? Of course the whole world could find itself in a liquidity trap (assuming some answer is found to the first challenge), but that is manifestly not the case at present. So doesn't the possibility of capital export make a liquidity trap impossible?

It turns out that for expositional purposes it is easier to address these questions in reverse order: first supposing that the business about filling in the Gulf of Mexico is for some reason not a killing objection, and analyzing the open-economy problem, then using the results of that analysis to help make sense of the domestic investment puzzle.

So let us suppose that we have an economy with a negative "natural" real rate of interest. [Figure 4](#) shows the savings and investment schedules at the full employment level of output. If this were a closed economy this would mean that at full employment those schedules would cross below zero. Thus at a zero real interest rate there would be an incipient excess supply of savings, which would then translate via a multiplier process into a depressed real economy.

But the basic accounting identity of international economics is that $S-I = NX$ (net exports). So one might think that the economy can invest the excess savings abroad; the counterpart of that overseas investment would be a current account surplus, which would provide the additional demand the economy needs. Indeed, analysts such as Smithers (1998) have argued that the essence of the Japanese dilemma is not the liquidity trap per se but the political problems raised by the implied trade surplus.

However, it is not quite that simple. While economists sometimes fall into the trap of supposing that savings-investment gaps are automatically translated into trade surpluses without any need for micro-level incentives to sell or buy more - a view John Williamson once referred to as the "doctrine of immaculate transfer" - in reality *something* has to effect the changes in exports and imports. If we restrict ourselves to considering incipient gaps at full employment, the adjusting must be done via relative prices, which for simplicity we can summarize by the real exchange rate. Thus, ignoring some basically unimportant complications, we can write net exports as a function of the real exchange rate:

But what determines the real exchange rate? We could do this carefully and correctly, but the basic insight comes through via an ad hoc approach. Investors will be aware that a country cannot run current account surpluses forever; in many empirical models it is assumed that they have in mind some long-run equilibrium real exchange rate, perhaps the rate at which net exports would be zero, and expect the actual rate to regress toward that rate over time:

$$de^e / dt = \lambda(e_I - e) \quad (4)$$

The crucial point, then, is that given this expected change in the real exchange rate even risk-neutral investors will not equate real interest rates across countries; rather, arbitrage will set the expected depreciation of a currency equal to the real interest differential vis-as-vis the rest of the world:

$$r - r^* = de^e / dt \quad (5)$$

And this will therefore make both the real exchange rate and the level of net exports functions of the real interest rate:

$$e = e_I - (1/\lambda)(r - r^*) \quad (6)$$

In other words, the level of net foreign investment is, like the level of domestic investment, a function of the interest rate. If an economy is very open on the goods side, small changes in the interest rate and hence in the real exchange rate will have

large effects on the current account; but if it is relatively closed (like Japan), this effect will be much smaller.

The implication can be seen in [Figure 5](#), which takes the saving-investment balance from Figure 4 and adds the net exports implied by (6). We are assuming that at a zero real interest rate there will be an excess supply of domestic saving; if the long-run expected exchange rate corresponds to a zero current balance there will be positive net exports at a zero r ; but they may not be large enough, and a liquidity trap can still exist despite investment opportunities abroad.

Here is the intuition: suppose that in order to reach full employment Japan would, as Smithers might argue, need to run a current account surplus of nearly 10 percent of GDP. To run such a surplus would require a very weak yen - say 200-250 to the dollar. Yet investors might well surmise that the yen would not stay that weak indefinitely: eventually both a recovery in Japan and the investment income from massive cumulative current account surpluses would imply a strong, not weak yen. But given that prospect, yen assets would be an attractive buy even with a zero interest rate - in other words, the yen cannot be gotten sufficiently weak to effect the needed current account surplus.

In short, we are once again faced with an expectational aspect to the liquidity trap; in this case, the expectation of future appreciation prevents the necessary export of capital.

Now that we have seen the limits to investment abroad, we can ask why a zero interest rate cannot generate unlimited investment at home. And one can immediately see a possible parallel. Suppose that we have a "Tobin's q " theory of investment. If a country is currently in a period of high saving, to invest that saving would require a high value of Tobin's q today, with an eventual return to something like normal levels. In that case it would be possible to have a positive marginal product of capital, yet a negative expected real rate of return, just as on overseas assets.

As a practical matter this is probably not the whole story. After all, even in non-liquidity-trap economies the rate of return on bonds is normally well below any reasonable estimate of the rate of return on investment; in the U.S., for example, it is a good bet that real investment has a rate of return in double digits, while the return on index bonds is less than 4 percent. Exactly why this differential exists is an interesting and possibly complex question, but it is not too hard to imagine that Japan might have a real rate of return on capital that is as high as 5 percent while still having a negative natural rate of interest.

But why is the rate of return so low? You could argue - and I have - that this need not matter. Inflation targeting is the appropriate policy response regardless of exactly why the economy seems to need a negative real interest rate. But other policies - fiscal expansion, unconventional open-market operations, etc. - are either in use or being proposed as the next step, and the prospects for those alternatives do (as we will see) depend on the causes of the trap. So let us turn next to three stories about the reasons for a negative natural interest rate.

4. Three causes of a depressed economy

In popular accounts of Japan's problems one often hears a litany of supposed causes. Some argue that the problem is structural, rooted in both demography (ageing and a declining working-age population) and in waning technological vigor. Others suggest that specific events - in particular, the severity of the collapse of the bubble economy - have jolted Japan into a self-reinforcing spiral of pessimism. Finally, one sometimes hears that the bubble left problems of a more tangible nature, namely large debts that burden enterprises and leave them unable to take advantage even of promising investment opportunities.

The purpose of this paper is to discuss the general problem of the liquidity trap, rather than get too much into Japanese specifics, so I can remain agnostic about these differing claims (although as a practical matter I would argue that the cases for both self-fulfilling pessimism and balance-sheet constraints do not hold up very well under critical scrutiny). The point I want to make here is that these are three distinct stories, with different implications for what sorts of policy might work.

The "structural" story is simplest: for whatever reason this economy currently has a high propensity to save, offers limited investment opportunities, and therefore looks like Figure 4. Structural is not a synonym for "immutable", so policy actions could conceivably shift these curves in a favorable direction. However, one would not expect aggregate demand policy to change the fact that the natural rate of interest is negative; all it could do would be to provide a way for the economy to cope with that reality better, that is, without unemployed resources.

The story that attributes a liquidity trap to self-fulfilling pessimism is very different. It is in a fundamental sense a multiple equilibrium story, with the liquidity trap corresponding to the low-level equilibrium. It is easiest to think about this story in terms of a version of the Keynesian cross - a much-maligned device that becomes

very useful when the interest rate is fixed because it is hard up against the zero constraint. [Figure 6](#) illustrates a simple multiple-equilibrium story: over some range spending rises more than one-for-one with income. (Why should the relationship flatten out at high and low levels? At high levels resource constraints begin to bind; at low levels the obvious point is that gross investment hits its own zero constraint. There is a largely forgotten literature on this sort of issue, including Hicks (194?), Goodwin (194?), and Tobin (1947)).

The important point about multiple equilibria is that they allow for permanent (or anyway long-lived) effects from temporary policies. There may be excess desired savings even at a zero real interest rate given the pessimism that now prevails in the economy, and that is sustained by the continuing stagnation; but if some policy could push the economy to a high level of output for long enough to change those expectations, the policy would not have to be maintained indefinitely. As we will see, this enlarges the range of policies that could "solve" the problem.

Finally, balance-sheet problems are somewhat different yet again. They may involve an element of self-fulfilling slump: a firm that looks insolvent with an output gap of 10 percent might be reasonably healthy at full employment. But aside from this, balance-sheet problems may be self-correcting given time. If the economy can be put on life support through some kind of temporary policy, this will give firms the chance to pay down their debts, and possibly therefore to regain the ability to invest without support at a later date.

As we will see next, the prospects for many policy options (but not for inflation targeting) depend on which of these stories is most nearly true.

5. Fiscal policy

"Pump-priming" fiscal policy is the conventional answer to a liquidity trap. The classic case is, of course, the way that World War II apparently bootstrapped the United States out of the Great Depression. And in either the IS-LM model or a more sophisticated intertemporal model fiscal expansion will indeed offer short-run relief from a liquidity trap. So why not consider the problem solved? The answer hinges on the government's own budget constraint.

You might suspect that we are about to talk about Ricardian equivalence here. But that is not the crucial issue. True, if consumers have long time horizons, access to

capital markets, and rational expectations tax cuts will not stimulate spending. However, real purchases of goods and services will still create employment, albeit perhaps with a low multiplier. (In a fully Ricardian setup the multiplier on government consumption will be exactly 1: the income generated by the purchases will not lead to higher consumption, because it will be matched by the present value of future tax liabilities). The problem instead is that deficit spending does lead to a large government debt, which will if large enough start to raise questions about solvency.

One might ask why government debt matters if the interest rate is zero in any case. But the liquidity trap, at least in the version I take seriously, is not a permanent state of affairs. Eventually the natural rate of interest will turn positive, and at that point the inherited debt will indeed be a problem.

So is fiscal policy a temporary expedient that cannot serve as a solution to a liquidity trap? Not necessarily: there are two circumstances in which it can work.

First, if the liquidity trap is short-lived in any case, fiscal policy can serve as a bridge. That is, if there are good reasons to believe that after a few years of large deficits monetary policy will again be able to shoulder the load, fiscal stimulus can do its job without posing problems for solvency. This might be the case if there were clear-cut external factors that one could expect to improve - say if the domestic economy was currently depressed because of a severe but probably short-lived financial crisis in trading partners. Or - a possibility argued by some defenders of the current Japanese problem - temporary fiscal support might provide the breathing space during which firms get their balance sheets in order.

If you listen to the rhetoric of fiscal policy, however - all the talk about pump-priming, jump-starting, etc. - it becomes clear that many people implicitly believe that only a temporary fiscal stimulus is necessary because it will jolt the economy into a higher equilibrium. Thus in Figure 5 a policy that shifts the spending curve up sufficiently will eliminate the low-level equilibrium; if the policy is sustained long enough, when it is removed the economy will settle into the high-level equilibrium instead.

If this is the underlying model of how fiscal policy is supposed to succeed, however, one must realize that the criterion for success is quite strong. It is not enough for fiscal expansion to produce growth - that will happen even if the liquidity trap is deeply structural in nature. Rather, it must lead to large increases in private demand, so large that the economy begins a self-sustaining process of recovery that can continue without further stimulus.

It is in this light that one should read economic reports about Japan today, and perhaps about other troubled economies in the future. For what it is worth, at the time of writing there is nothing in the data that would suggest that anything like the supposed shift to a higher equilibrium is in progress. Indeed, private demand is actually falling, with more than all the growth coming from government demand.

None of this should be read as a reason to abandon fiscal stimulus - in fact, one shudders to think what would happen if Japan were not to provide further packages as the current one expires. But fiscal stimulus is a solution, rather than a way of buying time, only under some particular assumptions that are at the very least rather speculative.

6. *Varieties of monetary policy*

If fiscal policy is not a definitive answer, we turn to monetary policy. As I have tried to argue, the most basic models of a liquidity trap already imply that a credible commitment to future monetary expansion is the "correct" answer to a liquidity trap, in the sense that – like monetary expansion in the face of a conventional recession – it is a way of replicating the results the economy would achieve if it had perfectly flexible prices. But this notion of monetary policy has become confused with two other monetary proposals, "quantitative easing" and unconventional open-market operations; it is important to be aware that these are not the same thing, and rest on different assumptions about what is needed.

Quantitative easing: There has been extensive discussion of "quantitative easing" , which usually means urging the central bank simply to impose high rates of increase in the monetary base. Some variants argue that the central bank should also set targets for broader aggregates such as M2. The Bank of Japan has repeatedly argued against such easing, arguing that it will be ineffective – that the excess liquidity will simply be held by banks or possibly individuals, with no effect on spending – and has often seemed to convey the impression that this is an argument against any kind of monetary solution.

It is, or should be, immediately obvious from our analysis that in a direct sense the BOJ argument is quite correct. No matter how much the monetary base increases, as long as expectations are not affected it will simply be a swap of one zero-interest asset for another, with no real effects. A side implication of this analysis (see Krugman 1998) is that the central bank may literally be unable to affect broader monetary

aggregates: since the volume of credit is a real variable, and like everything else will be unaffected by a swap that does not change expectations, aggregates that consist mainly of inside money that is the counterpart of credit may be as immune to monetary expansion as everything else.

But this argument against the effectiveness of quantitative easing is simply irrelevant to arguments that focus on the expectational effects of monetary policy. And quantitative easing could play an important role in changing expectations; a central bank that tries to promise future inflation will be more credible if it puts its (freshly printed) money where its mouth is.

Unconventional open-market operations: A second argument on monetary policy is that while conventional open-market operations are ineffective, the central bank can still gain traction by engaging in unconventional operations – with the most obvious ones being either currency-market interventions or purchases of longer-term securities. The argument of proponents of such moves, for example Alan Meltzer, is that in reality foreign bonds and long-term domestic bonds are not perfect substitutes for short-term assets, and hence open-market operations in these assets can expand the economy by driving the currency and the long-term interest rate down.

Clearly there is something to this argument: perfect substitutability is a helpful modeling simplification, but the real world is more complicated. And in the absence of perfect substitution, these interventions will clearly have some effect. The question is how much effect – or, to put it a bit differently, how large would the BOJ's purchases of dollars and/or JGBs have to be to make an important contribution to economic recovery. (You might say that it doesn't matter – the BOJ can print as many yen as it likes. And perhaps that is the right thing to say in principle. But if supporting the economy requires that the BOJ acquire, say, 100 trillion yen in assets over the next four years – and if it is likely to lose money on those assets – the policy is going to be difficult to pursue).

A rigorous model of monetary policy in the face of imperfect substitution is difficult to construct (if only because one must derive that imperfection somehow). But a shortcut may be useful. Consider, then, the case of foreign exchange intervention – purchasing foreign bonds in an effort to bid down the currency. And let us look back at Figure 5, which illustrates how a liquidity trap can occur even in an open economy, because the desired capital export even at a zero interest rate will be less than the excess of domestic savings over investment.

What would the central bank be doing if it engages in exchange-market intervention in such a situation? The answer is that in effect it would be trying to do through its own operations the capital export that the private sector is unwilling to do. So a minimum estimate of the size of intervention needed per year is "enough to close the gap" – that is, the central bank would have to buy enough foreign exchange, i.e. export enough capital, to close the *ex ante* gap between S-I and NX at a zero interest rate. In practice, the intervention would have to be substantially larger than this, probably several times as large, because the intervention would induce private flows in the opposite direction. (An intervention that weakens the yen reinforces the incentive for private investors to bet on its future appreciation).

Here is some sample arithmetic: suppose that you believe that Japan currently has an output gap of 10 percent, which might be the result of an *ex ante* savings surplus of 4 or 5 percent of GDP. Then intervention in the foreign exchange market sufficient to close that gap would have to be several times as large as the savings surplus – i.e., it could involve the Bank of Japan acquiring foreign assets at the rate of 10, 15 or more percent of GDP, over an extended period. (Incidentally, does it matter if the interventions are unsterilized? Well, an unsterilized intervention is a sterilized intervention plus quantitative easing; the latter part makes no difference unless it affects expectations).

Purchases of long-term bonds would work similarly. In this case the central bank would in effect be competing with private investors as a source of investment finance (this would be true even if the intervention itself were in government bonds). Again, there would be an offset – with lower yields, private investors would divert some of their savings from bonds into short term assets or, what is equivalent under liquidity trap conditions, cash. So again the central bank would have to sustain purchases at a rate several times the *ex ante* savings-investment gap; in this case the BOJ might find itself purchasing long-term bonds at a rate of 10-15 or more percent of GDP.

There are obvious political economy problems with such actions. The prospect of having a quasi-governmental institution owning a trillion dollars of overseas assets, or most of the Japanese government's debt, is a bit daunting. Of course this would not happen if a relatively brief period of unconventional monetary policy led to a self-sustaining recovery. But to believe in this prospect you must, as in the case of fiscal policy, believe that the economy is currently in a low-level equilibrium and can be jolted back to prosperity with temporary actions – a fairly exotic, though not necessarily wrong, view on which to base policy.

The same remarks applied to fiscal policy also apply here: while unconventional open-market operations are less certain a cure than their proponents seem to think, they could help, and might well be part of a realistic strategy.

Expectations: Finally, we return to the issue of inflation targeting. The basic point, once again, is that a credible commitment to expand the future money supply, perhaps via an inflation target, will be expansionary even in a liquidity trap. There are two problems, however, with this view. One is that it is not enough to get central bankers to change their spots; one must also convince the market that the spots have changed, that is, actually change expectations. The truth is that economic theory does not offer a clear answer to how to make this happen. One might well argue, however, that one way to help make a commitment to do something unusual credible is to do a lot of other unusual things, demonstrating unambiguously that the central bank does understand that it is living in a different world. Market participants are pretty much unanimous in their belief that unsterilized intervention would have a much bigger effect than sterilized, essentially because it would convey news about future BOJ policy; the same could be said of other actions, including quantitative easing. My personal view is that a country deep in a liquidity trap should try everything, even if careful analysis says that some of the actions should not matter; if, in the precise if annoying phrase I used in my first paper on the liquidity trap, a central bank must "credibly promise to be irresponsible", it should waste no opportunity to demonstrate its new spirit.

The other problem is that the policy shift must not only be credible but sufficiently large. A too-modest inflation target will turn into a self-defeating prophecy. Suppose that the central bank successfully convinces everyone that there will henceforth be 1 percent inflation – but that a real interest rate of minus 1 percent is not low enough to restore full employment. Then despite the expectational change, the economy will remain subject to *deflationary* pressure, and the policy will fail. Half a loaf, in other words, can be worse than none.

7. Concluding remarks

The whole subject of the liquidity trap has a sort of Alice-through-the-looking-glass quality. Virtues like saving, or a central bank known to be strongly committed to price stability, become vices; to get out of the trap a country must loosen its belt, persuade its citizens to forget about the future, and convince the private sector that the government and central bank aren't as serious and austere as they seem.

The strangeness of the situation extends to policy discussion. Because the usual rules do not apply, conventional rules of thumb about policy become hard to justify. We usually imagine that policy is more or less based on conventional models – in particular, that normally policy will be based on the simple, rather dull models in the textbooks rather than exotic stories that might be true but probably aren't. In the case of the liquidity trap, however, conventional textbook models imply unconventional policy conclusions – for inflation targeting is not an exotic idea but the natural implication of both IS-LM and modern intertemporal models applied to this unusual situation. To defend the conventional policy wisdom one must therefore appeal to various unorthodox models – supply curves that slope down, demand curves that slope up, multiple equilibria, etc.. So unworldly economists become defenders of analytical orthodoxy, while the dignified men in suits become devotees of exotic theories.

What I hope that I have done in this paper is to make clear how conventional the logic behind seemingly radical proposals like inflation targeting really is, and conversely how hard it is to rationalize what still passes for sensible policies among many officials. Let's see if it works this time around.




Figure 1




Figure 2

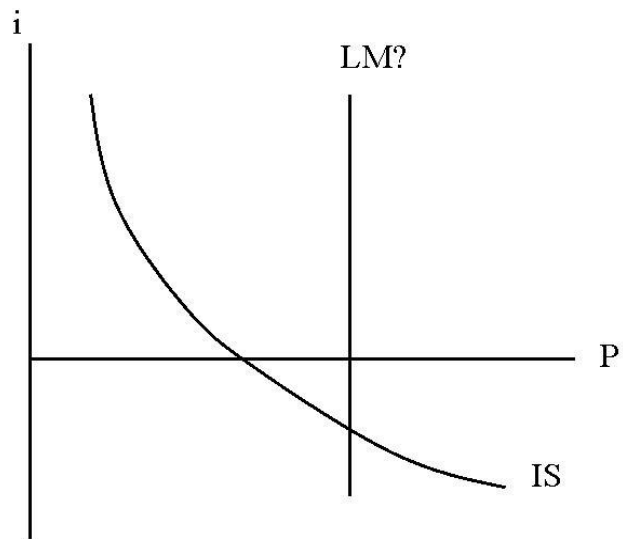


Figure 3




Figure 4

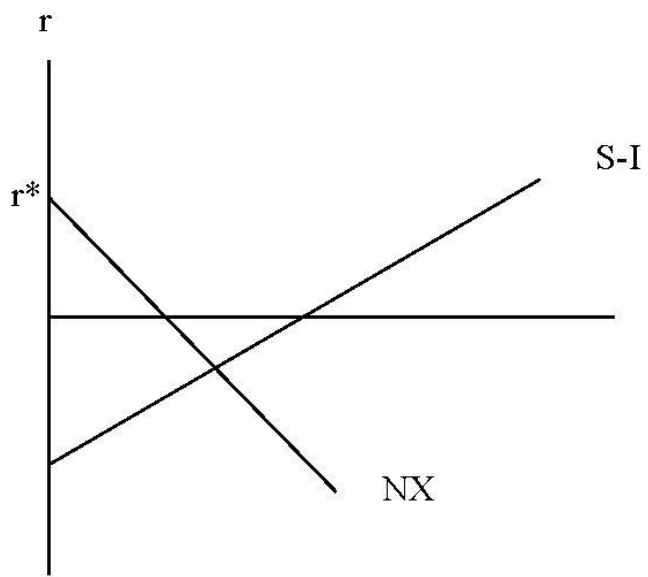


Figure 5




Figure 6