IT'S BAAACK! JAPAN'S SLUMP AND THE RETURN OF THE LIQUIDITY TRAP

In the early years of macroeconomics as a discipline, the liquidity trap - that awkward condition in which monetary policy loses its grip because the nominal interest rate is essentially zero, in which the quantity of money becomes irrelevant because money and bonds are essentially perfect substitutes - played a central role. Hicks (1937), in introducing both the IS-LM model and the liquidity trap, identified the assumption that monetary policy was ineffective, rather than the assumed downward inflexibility of prices, as the central difference between "Mr. Keynes and the classics". It has often been pointed out that the Alice-in-Wonderland character of early Keynesianism, with its paradoxes of thrift, widow's cruses, and so on, depended on the explicit or implicit assumption of an accommodative monetary policy; it has less often been pointed out that in the late 1930s and early 1940s it seemed quite natural to assume that money was irrelevant at the margin. After all, at the end of the 30s interest rates were hard up against the zero constraint: the average rate on Treasury bills during 1940 was 0.014 percent.

Since then, however, the liquidity trap has steadily receded both as a memory and as a subject of economic research. Partly this is because in the generally inflationary decades after World War II nominal interest rates stayed comfortably above zero, and central banks therefore no longer found themselves "pushing on a string". Also, the experience of the 30s itself was reinterpreted, most notably by Friedman and Schwartz (1963); emphasizing broad aggregates rather than interest rates or monetary base, they argued in effect that the Depression was caused by monetary contraction, that the Fed could have prevented it, and implicitly that even after the great slump a sufficiently aggressive monetary expansion could have reversed it. To the extent that modern

macroeconomists think about liquidity traps at all - EconLit lists only 21 papers with the phrase in title, subject, or abstract since 1975 - their view is basically that the liquidity trap can't happen, it didn't happen, and it won't happen again.

But it has, and to the world's second-largest economy. Over the past several years money-market rates in Japan have been consistently below 1 percent, and the Bank of Japan plausibly claims that it can do no more; yet the economy, which has been stagnant since 1991, is sliding deeper into recession.

Since Japan is such an important economy, and its slump threatens to shatter the already fragile prospects for economic recovery in the rest of Asia, understanding what is going wrong there has become quite urgent. And there is also a deeper reason for concern: if this can happen to Japan, perhaps it can happen elsewhere. In short, it is time to reexamine the theory of liquidity traps, which has turned out not to be irrelevant after all.

But don't we already understand liquidity traps well enough to formulate policy? Can't we just pull the old models out of the basement, dust them off, and put them to work? In effect that is what policymakers at the U.S. Treasury and elsewhere have done: drawing on the simple liquidity trap framework that used to appear in macroeconomics textbooks a generation or so ago, they have urged Japan to follow the classic recovery strategy of pump-priming fiscal expansion. (Since hardly anybody in the thoroughly urbanized societies of modern America and Japan has any idea what it means to prime a pump, I hereby suggest that we rename this the "jump-start" strategy). Macroeconomics has, however, moved on since those textbooks, in several ways that might require a rethinking of the issue.

We might, in particular, identify three strands of modern thought that are missing from the

classic IS-LM analysis. First is the intertemporal nature of decisions: we now understand, perhaps better than macroeconomists 50 years ago, that how one formulates expectations is a crucial matter in macroeconomic analysis, and that a good first-pass assumption is that these expectations are rational. Second is the openness of the economy. Although the Britain of Keynes and Hicks was actually a quite open economy - with a share of trade in GDP more than twice that of modern Japan - their analysis, and almost all subsequent analysis of the liquidity trap, ignored foreign trade and capital mobility. It was a justifiable strategic simplification; but since many of the disputes surrounding Japan's direction involve the future of the country's current account and exchange rate, we need to know what happens when this assumption is relaxed. Finally, traditional IS-LM analysis neglects the role of financial intermediaries. But how you interpret the experience of the 1930s hinges crucially on how broad a monetary aggregate you choose; and the same has turned out to be true in recent arguments over Japan. Furthermore, one school of thought about the Depression argues that a troubled banking system was the essence of the problem; a similar view has become near-orthodoxy about contemporary Japan. So we had better have at least a basic sense of how financial intermediation fits into the liquidity-trap picture.

This paper is in two major parts. The first, longer part is an extended generic discussion of the causes and consequences of liquidity traps; it uses a succession of small, highly stylized models to address both the traditional questions regarding liquidity traps and a number of novel issues. The central new conclusion of this analysis is that a liquidity trap fundamentally involves a credibility problem - but one that is the inverse of the usual one, in which central bankers have difficulty convincing private agents of their commitment to price stability. In a liquidity trap the problem is that the markets believe that the central bank *will* target price stability, given the chance - and

hence that any current monetary expansion is merely transitory. The traditional view that monetary policy is ineffective in a liquidity trap, and that fiscal expansion is the only way out, must therefore be qualified: monetary policy will be effective after all if the central bank can credibly promise to be irresponsible, to seek a higher future price level.

The theoretical analysis also appears to refute two widely held beliefs. First, international capital flows, which allow a country to export savings to the rest of the world, are not a sure-fire guarantee against a liquidity trap; because goods markets remain far from perfectly integrated, the required real interest rate in terms of domestic consumption can be negative even if capital is perfectly mobile and there are positive-return investments abroad. A corollary to this is that the extent to which a successful monetary expansion - in which the central bank succeeds in creating expectations of inflation - will be a beggar-thy-neighbor policy that expands demand at the rest of the world's expense will probably be less than widely imagined.

Second, putting financial intermediation into a liquidity-trap framework suggests, *pace* Friedman and Schwartz, that looking at monetary aggregates under these circumstances is quite misleading: in a liquidity trap the central bank may well find that it *cannot* increase broader monetary aggregates, that increments to the monetary base are simply added to reserves and currency holdings - and thus both that such aggregates are no longer a valid indicator of the stance of monetary policy, and that their failure to rise does not indicate that the essential problem lies in the banking sector.

The second, briefer part of the paper turns to some specific questions surrounding Japan; it essentially represents a survey of others' estimates rather than original empirical work. It considers four main issues. First is the size of Japan's output gap; I argue that it is probably

considerably larger than the standard estimates, and hence that the need for expansionary policy is even greater than commonly supposed. Second is the reason for the apparent large gap between saving and willing investment at full employment. Third is the relevance of Japan's banking woes to its macroeconomic malaise; although the conventional wisdom is that Japan's banks are at the center of the problem, I argue that banking problems have less of a causal role than is widely assumed. Finally, I make a stab at quantifying the size, duration, and side effects of the inflation that would be needed to lift Japan out of its trap.

REVISITING THE THEORY OF LIQUIDITY TRAPS

General considerations

It is useful, in considering the problem of Japan's liquidity trap, to begin at a high level of generality - to adopt what we might almost call a philosophical stance. The reason is that popular discussion of the current problem has a strong tendency to plunge quickly into the specifics, to cite one or another structural issue as *the* problem, missing the central point that whatever the specifics of its history Japan is now in a liquidity trap, and that the generic issues surrounding such traps apply whatever the details of the particular case.

A liquidity trap may be defined as a situation in which conventional monetary policies have become impotent, because nominal interest rates are at or near zero - so that injecting monetary base into the economy has no effect, because base and bonds are viewed by the private sector as perfect substitutes. By this definition, a liquidity trap could occur in a flexible-price, full-

employment economy; and although any reasonable model of the United States in the 1930s or of Japan in the 1990s must invoke some form of price stickiness, we can think of the unemployment and output slump that occurs under such circumstances as what happens when the economy is "trying" to have deflation - a deflationary tendency that monetary expansion is powerless to prevent.

This may seem a peculiar way of putting the issue, but it does highlight the central mystery of a liquidity trap - and the reason why "structural" explanations, in a fundamental sense, cannot by themselves resolve that mystery. For if there is one proposition everyone agrees with in macroeconomics it is that, aside from the possible role of price stickiness in causing monetary expansion to be reflected in output rather than prices, increases in the money supply raise the equilibrium price level. Indeed, the normal view is that money is roughly neutral - that an increase in the money supply produces a roughly equiproportional increase in the general price level.[1] Or to be more specific, an increase in *outside* money - the monetary base - must raise prices.

Putting it this way immediately reveals that many of the explanations one often hears for the ineffectuality of monetary policy in Japan are wrong, or at least inadequate. One often hears, for example, that the real problem is that Japan's banks are troubled, and hence that the Bank of Japan cannot increase monetary aggregates; but outside money is still supposed to raise prices, regardless of the details of the transmission mechanism. In addition to the problem of bad loans, one often hears that corporations have too much debt, that the service sector is overregulated and inefficient, and so on. All of this may be true, and may depress the economy for any *given* monetary base - but it does not explain why *increases* in the monetary base should fail to raise prices and/or output. One way to say this is to remember that the neutrality of money is not a

conditional proposition; money isn't neutral “if your banks are in good financial shape” or “if your service sector is competitive” or “if corporations haven't taken on too much debt”. Money (which is to say outside money) is supposed to be just plain neutral.[2]

So how is a liquidity trap possible? The answer lies in a little-noticed escape clause in the standard argument for monetary neutrality: an increase in the money supply in the current *and all future* periods will raise prices in the same proportion. There is no corresponding argument that says that a rise in the money supply *that is not expected to be sustained* will raise prices equiproportionally - or indeed at all.

In short, approaching the question from this high level of abstraction already suggests that a liquidity trap involves a kind of credibility problem. A monetary expansion that the market expected to be sustained (that is, matched by equiproportional expansions in all future periods) would always work, regardless of whatever structural problems the economy might have; if monetary expansion does not work, if there is a liquidity trap, it must be because the public does not expect it to be sustained.

To firm up this insight, of course, we need a specific model.

Money, interest, and prices: a minimalist model

Although the idea of a liquidity trap is normally bound up with the IS-LM model, there are several compelling reasons *not* to start with that model here. For one thing, many macroeconomists believe that IS-LM is too ad hoc to be worthy of serious consideration. Some of us do not share that view, and continue to regard Hicks's construction as a very useful heuristic

device. Still, it is important to stress that the possibility of a liquidity trap is not something that depends on the ad hockery of the IS-LM model, that it can occur in a model that dots its microeconomic "i"s and crosses its intertemporal "t"s. Also, a liquidity trap does, as we have already seen, fundamentally involve expectations and credibility; using models that explicitly recognize the intertemporal aspects of the problem helps clarify this point.

Let us therefore move immediately to an explicit intertemporal model that establishes relationships among output, money, prices, and interest rates; we will then be able to use this model as a base for a series of thought experiments and extensions.

We consider a one-good, representative agent economy (in which, however, agents must purchase their consumption from others); initially we suppose that the good is inelastically supplied, so that we can simply think of each agent as receiving a given endowment y_t in each period. For concreteness, the utility function is assumed to take the form

$$U = \frac{1}{1-\rho}\sum c_t^{1-\rho} D^t \qquad \textbf{(1)}$$

where c is consumption within a period, ρ is relative risk aversion, and D the discount factor.

The simplest way to introduce money into this model, one that has the added advantage of avoiding the suspicion that the conclusions are dictated by arbitrary assumptions about the way money enters utility, is to assume a cash-in-advance constraint. Specifically, within each period agents are assumed to go through a two-stage process. At the beginning of each period there is a capital market, in which individuals can trade cash for one-period bonds, with a nominal interest rate i_t. Their consumption during the period is then constrained by the cash with which they emerge from this trading: the nominal value of consumption, $P_t c_t$, cannot exceed money

holdings M_t. After the capital market is held, each individual purchases his desired consumption, while receiving cash from the sale of his own endowment.

Government policy can take two forms. First, the central bank is assumed able to engage in open-market operations during that beginning-of-period capital market, buying or selling bonds. Second, at the end of the period the government can collect or distribute lump-sum taxes and transfers. The government must obey its own intertemporal budget constraint, which takes into account the seignorage that results from money creation.

Analyzing this model in general requires a careful specification of the budget constraints of both individuals and the government, and of intertemporal choices. However, if we make some simplifying assumptions the model's implications can be derived with almost no algebra. Let us assume that from the second period onwards output (and therefore also consumption) will remain constant at a level y^*, and that the government will also hold the money supply constant at a level M^*. Then we can immediately guess at the solution from period 2 on: the price level will remain constant at $P^* = M^*/y^*$, and the interest rate will also be constant at a rate $i^* = (1-D)/D$. It is straightforward to confirm that this is indeed an equilibrium: one plus the real interest rate equals the ratio of marginal utility in any two successive periods; because the nominal interest rate is positive individuals have an incentive to acquire only as much cash as they need, so all money will indeed be spent on consumption.

All the action, then, goes into determining the price level and interest rate in the first period. Let us use un-subscripted letters to represent first-period output, consumption, interest rate, etc..

Our first relationship comes from the monetary side. Under normal circumstances - that is, when the nominal interest rate is positive - individuals will hold no more cash than they need to

make their consumption purchases. So the cash-in-advance constraint will be binding: $Pc = Py = M$, so that

$$P = M/y \qquad \textbf{(2)}$$

So under normal circumstances there is a simple proportional relationship between the money supply and the price level.

The second relationship comes from intertemporal choice. By holding one less yen in period 1, an individual gives up $1/P$ units of first-period consumption but allows himself to consume $(1+i)/P^*$ additional units in period 2. At an optimum this change must leave him indifferent. But the marginal utility of consumption in period 1, given the assumed utility function, is $c^{-\rho}$; the marginal utility in 2 is $D(c^*)^{-\rho}$. It follows that we must have

$$(c/c^*)^{-\rho} = DP(1+i)/P^* \qquad \textbf{(3)}$$

or, since consumption must equal output in each period,

$$1+i = \frac{P^*}{DP}(y^*/y)^{1/\rho} \qquad \textbf{(4)}$$

This says that the higher is the current price level, the lower the nominal interest rate. The easiest way to think about this is to say that there is an equilibrium *real* interest rate which the economy will deliver whatever the behavior of nominal prices. Meanwhile, since the future price level P^* is assumed held fixed, any rise in the current level creates expected deflation; hence higher P means lower i.

The two relationships are shown in Figure 1 as *MM* and *CC* respectively; as drawn, they intersect at point 1, simultaneously determining the interest rate and the price level. It is also immediately apparent that an increase in the first-period money supply will shift *MM* to the right, leading to a higher price level and a lower nominal (but not real) interest rate.

While this is surely the normal case, however, there is also another possibility, to which we now turn.

The liquidity trap in a flexible-price economy

Suppose that we start with an economy in the equilibrium described by point 1 in Figure 1, and then imagine an initial open-market operation that increases the first-period money supply. (Throughout we imagine that the money supply from period 2 onwards remains unchanged - or equivalently that the central bank will do whatever is necessary to keep the post-2 price level stable). Initially, as we have already seen, this operation will increase the price level and reduce the interest rate. And such a monetary expansion can clearly drive the economy down the *CC* curve as far as point 2 in Figure 1. But what happens if the money supply is increased still further - so that the intersection of *MM* and *CC* is at a point like 3, with a negative nominal interest rate?

The answer is clearly that the interest rate cannot go negative, because then money would dominate bonds as an asset. What must therefore happen is that any increase in the money supply beyond the level that would push the interest rate to zero is simply substituted for zero-interest bonds in individual portfolios (with the bonds being purchased by the central bank in its open-market operation!), with no further effect on either the price level or the interest rate. Because

spending is no longer constrained by money, the *MM* curve becomes irrelevant; the economy stays at point 2, no matter how large the money supply may be.

It is probably worth emphasizing here that the interest rate at point 2 is zero only on *one-period* bonds; it would not be zero on longer-term bonds, such as consols. This is important if one is trying to map the model onto the current situation in Japan, or for that matter the United States during the 1930s: long rates in Japan are still positive, but short-term rates are indeed very close to zero.

A good way to think about what happens when money becomes irrelevant here is to bear in mind that we are holding the long-run money supply fixed at M^*, and therefore also the long-run price level at P^*. So when the central bank increases the current money supply, it is lowering the expected rate of money growth M^*/M, and also - if it does succeed in raising the price level - the expected rate of inflation P^*/P. Now what we know is that in this full employment model the economy will have the same *real* interest rate whatever the central bank does. Since the *nominal* interest rate cannot become negative, however, the economy has a minimum rate of inflation or maximum rate of deflation.

Now suppose that the central bank in effect tries to impose a rate of deflation that exceeds this minimum - which it does by making the current money supply M large relative to the future supply M^*. What will happen is that the economy will simply cease to be cash-constrained, and any excess money will have no effect: the rate of deflation will be the maximum consistent with a zero nominal rate, and no more.

This may seem a silly thought experiment. Why would a central bank try to impose massive deflation? But the maximum rate of deflation need not be large, or even positive. Suppose that the

required real rate of interest is negative; then the economy "needs" inflation, and an attempt by the central bank to achieve price stability will lead to a zero nominal interest rate, and excess cash holdings.

The condition under which the required real interest rate is negative is straightforward in this simple endowment economy. Market-clearing will require a negative real interest rate if the marginal utility of consumption in period 2 is greater than that of consumption in period 1, which will be the case if the economy's future output is expected to be sufficiently less than its current output. Specifically, given the assumed utility function, the required real interest rate is negative if

$$(y/y^{*})^{1/\rho} < D \qquad \textbf{(5)}$$

This condition may seem peculiar. After all, we normally think of economies as growing rather than shrinking. One possible answer involves an equity premium; another involves demography; but let us reserve this issue for later discussion.

Of course, in a flexible-price economy even the necessity of a negative real interest rate does not cause unemployment. This conclusion may surprise economists who recall the tortured historical debate about the liquidity trap, much of which focussed on the question of whether wage and price flexibility were effective as a way of restoring full employment. In this model the problem does not arise - but the reason is a bit peculiar. What happens is that the economy deflates now in order to provide inflation later. That is, if the current money supply is so large compared with the future supply that the nominal rate is zero, but the real rate needs to be negative, P falls below P^*; the public then expects the price level to rise, and this provides the

necessary negative real interest rate. And to repeat, this fall in the price level occurs regardless of the current money supply, because any excess money will simply be hoarded without adding to spending.

At this point we have a version of the liquidity trap: money becomes irrelevant at the margin.[3] But aside from frustrating the central bank - which may have a thing about price stability, but finds itself presiding over inflation no matter what it does - this trap has no adverse real consequences. To turn this analysis into a real problem, in both senses, we need to introduce some kind of nominal rigidity.

The Hicksian liquidity trap

Suppose that the consumption good is produced rather than simply appearing, with a maximum productive capacity y^f in period 1. And suppose, also, that this productive capacity need not be fully employed. In particular, this paper will simply assume that the price level in period 1 is predetermined - so that the economy now acquires a Keynesian feel, and monetary policy can affect output. (In period 2 and subsequently output will still be assumed to take on the value y^*).

In this sticky-price world the level of period-1 consumption and output must still be equal, but now output adjusts to consumption rather than the other way around. Given the utility function, and the assumption that consumption will be y^* in period 2, we can immediately write an expression for current real consumption, which becomes the "IS curve" determining real output:

$$c = y = y^*(P^*/DP)^{1/\rho}(1+i)^{-1/\rho} \qquad \textbf{(6)}$$

Figure 2 illustrates the joint determination of the interest rate and output in this case. The curve IS, as just indicated, shows how output will be determined by consumption demand, which is decreasing in the interest rate. Meanwhile, as long as the nominal interest rate is positive, the cash-in-advance constraint will be binding, so we will have the *MM* curve

$$y = M/p \qquad \textbf{(7)}$$

Increasing the money supply can now increase output, up to a point - specifically, up to point 2. But what if productive capacity is at a point like 3? Then the same argument as in the previous section applies: since the nominal interest rate cannot go negative, any increase in money beyond the level that drives the rate to zero will simply be substituted for bonds, with no effect on spending. And therefore no open-market operation, no matter how large, can get the economy to full employment. In short, the economy is in a classic Hicksian liquidity trap.

Under what conditions will a liquidity trap occur? One possibility is that P is high compared with P^* - that people expect deflation, so that even a zero nominal rate is a high real rate. The other possibility, however, is that even if prices are expected to be stable, y^f is high compared with the future - or to put it differently, peoples' expected future real income is low compared with the amount of consumption needed to use today's capacity. In that case, to persuade people to spend enough now may require a negative real interest rate, and with downwardly inflexible prices that may not be possible.

Or to put it yet another way, one that is closer to the language of applied macroeconomics: if people have low expectations about their future incomes, then even with a zero interest rate they may want to save more than the economy can absorb. (In this case, of course, the economy

cannot absorb *any* savings - but we will come to that point below). And in that case, no matter what the central bank does with the current money supply, it cannot reflate the economy sufficiently to restore full employment.

So we have now seen that a fully specified model, which does not fudge either the role of money or the necessity of making intertemporal choices, can indeed generate a liquidity trap. The model does, however, clearly omit some important aspects of standard macro models; perhaps most notably, it has no investment; no foreign trade or capital mobility; and no financial intermediation, so that all money is outside. Can the same story still be told if these elements are introduced?

Investment, productive capital, and q

One way of stating the liquidity trap problem is to say that it occurs when the equilibrium real interest rate, the rate at which savings and investment would be equal at potential output, is negative. An immediate question is therefore how this can happen in an economy which is not the simple endowment economy described above, but one in which productive investment can take place - and in which the marginal product of capital, while it can be low, can hardly be negative.

An answer that may be extremely important in practice is the existence of an equity premium. If the equity premium is as high as the historic U.S. average, the economy could find itself in a liquidity trap even if the rate of return on physical capital is as high as 5 or 6 percent.

A further answer is that the rate of return on investment depends not only on the ratio of capital's marginal product to its price, but also on the expected rate of change of that price. An

economy in which Tobin's *q* is expected to decline could offer investors a negative real rate of return despite having a positive marginal product of capital.

This point is actually easiest to make if we consider an economy, not with capital, but with land (which can serve as a sort of metaphor for durable capital) - and also if we temporarily depart from the basic setup to consider an overlapping-generations setup, in which each generation works only in its first period of life but consumes only in its second. Let A be the stock of land, and L_t be the labor force in period t - that is, the number of individuals born in that period. Given the special assumption that the young do not consume during their working years, but use all their income to buy land from the old, we have a very simple determination of q_t, the price of land in terms of output: it must simply be true that

$$q_t A_t = w_t L_t \qquad \textbf{(8)}$$

where w_t is the marginal product of labor. So in this special setup q itself is not a forward-looking variable; it depends only on the size of the current labor force.

However, the expected rate of return on purchases of land *is* forward-looking. Let R_t be the marginal product of land, and r_t the rate of return for the current younger generation. Then we have that

$$1 + r_t = \frac{R_{t+1} + q_{t+1}}{q_t} \qquad \textbf{(9)}$$

Now suppose that demographers project that the next generation will be smaller than the current one, so that the labor force and hence (given elastic demand for labor) the real price of

land will decline. Then even though land has a positive marginal product, the expected return from investing in it can in principle be negative.

This is a highly stylized example, which begs many questions. However, it at least establishes the principle that a liquidity trap can occur despite the existence of productive investment projects.

International mobility of goods and capital

Many writers on Japan have assumed that one solution to the apparent excess of savings over investment, even at a zero interest rate, is simply for Japan to invest its excess savings abroad. The influential recent study by Smithers (1998) suggests that over the long term Japan should run capital account deficits (and hence current account surpluses) of no less than 10 percent of GDP. The general view in such discussions seems to be that an open economy can always extricate itself from a liquidity trap as long as there are profitable investment opportunities overseas; that the main problem is the political one of persuading the rest of the world to accept the corresponding trade surpluses.

Unfortunately, the economics of capital export are not as favorable as this analysis suggests. The reason is that the limited integration of markets for *goods and services* turns out to prevent capital flows from equalizing real interest rates *in terms of domestic consumption*, even when the mobility of capital itself is perfect. The fact is that in large economies like Japan (or the United States), the bulk of employment and value-added is in goods and services that remain nontradeable despite modern communication and transportation technology. And this large

nontradeable share may well mean that capital export, even at a zero interest rate, is not enough to spring a liquidity trap.

This argument can be made in the language of conventional open-economy IS-LM models. In such models it is usual to tie down the exchange rate by assuming that the market expects the real exchange rate to return to some normal value in the long run. The current real exchange rate is then determined off this long-run rate via the real interest differential between domestic and foreign bonds. So a monetary expansion that lowers nominal and hence real interest rates at home will produce a real depreciation, and this real depreciation will increase net exports at any given level of output. However, there is a limit to the size of the stimulus that this depreciation can generate: because the real exchange rate is expected to revert to its normal level, even a zero interest rate will produce only a finite real depreciation. If trade is a small share of GDP, and if price elasticities of imports and exports are also fairly small - both of which are true, if not in reality, at least in econometric models of large economies such as Japan and the United States - even near-perfect capital mobility may therefore provide only limited extra scope for monetary expansion.

But should we believe this story? While open-economy IS-LM may be a highly useful heuristic device for thinking about short- and medium-run macro issues, many economists doubt whether it is really trustworthy, especially about such fundamental questions as the scope for international capital flows. And in any case the thrust of this paper is to try to remove the stigma of ad hockery from the concept of the liquidity trap; so it may be helpful to supplement this conventional view with a restatement in terms of a variant of our basic intertemporal model.

Consider, then, a somewhat modified version of our basic model. We now suppose that the

economy produces and consumes two goods, a tradeable good *T* and a nontradeable *N*. Utility takes the form

$$U = \frac{1}{1-\rho}\sum_t D^t\left[c_{Tt}^{\tau}\ c_{Nt}^{1-\tau}\right]^{1-\rho} \qquad \textbf{(10)}$$

In general, we would want to give the economy a transformation curve between *N* and *T* at any point in time. For simplicity, I will assume that the transformation curve is right-angled - that is, that the economy receives exogenous endowments of the two good in each period. It can, however, borrow and lend on world markets at a given real interest rate r_T in terms of the tradeable good, so that consumption of that good need not be the same as production.

Does this assumed perfect capital mobility therefore imply that the domestic real interest rate must equal the world rate? Not if inflation is measured in terms either of the nontraded good or in terms of a consumption basket that includes both traded and nontraded goods. This is most easily seen if we consider the special case in which ρ=1; that is, in which (10) takes the special form

$$U = \sum_t D^t[\tau\ln(c_{Tt}) + (1-\tau)\ln(c_{Nt})] \qquad \textbf{(11)}$$

In this particular case, utility becomes separable between tradeables and nontradeables. For each good the relationship between consumption growth and the real interest rate must obey the rule $1+r_t = D^{-1}(c_{t+1} / c_t)$. However, whereas in the tradeable sector relative consumption is determined by the exogenous real interest rate, in the nontraded sector - assuming full employment - it will be the other way around: because consumption of nontradeables must equal production, the real interest rate in terms of nontraded goods will have to adjust to the path of production. And as a result it is entirely possible that the market-clearing real interest rate *in*

terms of nontraded goods will be negative, even with perfect capital mobility; and if the traded share in the consumption basket is small enough, the overall domestic real rate may be negative even if the world real rate is positive.

Now introduce the possibility of unemployment, by making the nominal price of nontradeables downwardly rigid, and consider the effects of a temporary monetary expansion (i.e., one that increases the money supply in the first period but does not change expectations about money supplies in later periods). Such a monetary expansion will lower the nominal interest rate, with different effects on the two sectors. In tradeables the real interest rate is tied down by world capital markets, so there must now be expected *deflation* in traded goods prices; but the future price is also tied down by the assumption that the monetary expansion is only temporary. So the current price of tradeables must rise, in order to allow for the subsequent fall. There must therefore be a nominal depreciation of the exchange rate.

In nontradeables the situation will be exactly as in the economy as a whole in our closed-economy model: the lower nominal rate will also be a lower real rate, and both consumption and production will increase.

The important point is that for both the exchange rate and for nontradeable production, the zero constraint on the nominal interest rate can be binding. That is, even at a zero interest rate the size of the output increase and of the nominal depreciation will be finite - and the economy may not be able to go all the way to full employment.

Incidentally, in this log-utility case monetary expansion has no effect on the current account. The reason is that the separability of the utility function means that consumers in effect have completely separate decisions to make on tradeable and nontradeable consumption over time; and

since the real interest rate on tradeables does not change, there is no reallocation between present and future consumption of those goods. This is obviously an artifact of the assumption that $\rho=1$; we will see the consequences of larger ρ shortly.

Financial intermediation and monetary aggregates

Attempts to make sense of the origins and persistence of the Great Depression hinge crucially on how one interprets the radical divergence between the growth of monetary base and that of broader aggregates during the period. Figure 3 shows the familiar picture, presented slightly differently (with monetary base and M2 both presented as indices, 1929=100): monetary base actually rose even during the early years of slump, and continued to rise steeply throughout the 1930s; yet M2 fell by more than a third, and did not surpass its 1929 level until 1939. These basic facts in effect underlie two influential views of the Depression. One, suggested by Friedman and Schwartz (1963), is that a broad aggregate like M2 is the proper measure of the money supply; that the Depression happened because the Fed allowed broad money to fall so much, and that recovery was so long delayed because the needed increase in broad money was equally long delayed. The other, associated with Bernanke (1994), Cooper and Corbae (1997), and others is that the dramatic decline in the money multiplier was the signature of a major episode of financial disintermediation; and that this disintermediation, which may be thought of more nearly as a supply-side than a demand-side phenomenon, was the cause of the sustained slump. On the other hand, monetary explanations of the Great Depression have been criticized, most notably by Temin (1976), who suggests that the decline in monetary aggregates was a result rather than a cause of

the slump, and perhaps one that the Fed could not have prevented.

Since the Depression is the main historical example of liquidity-trap economics, and since similar arguments are fairly often heard about contemporary Japan, it is therefore important to ask how financial intermediaries and monetary aggregates fit into the liquidity trap story.

Fortunately, it is quite easy to sketch out how this might be done, using a framework that might be described as "cash-in-advance meets Diamond-Dybvig". (A formal exposition of this framework is given in Appendix A). In the classic paper by Diamond and Dybvig (1982), a demand for liquidity is introduced by making individuals uncertain about their own consumption needs; they only discover whether they are type 1 consumers, who derive utility from consumption in period 1 but not 2, or type 2 consumers, who are the reverse, after they have made commitments to illiquid investments. This dilemma can be resolved by a class of financial intermediaries that allow individuals to withdraw funds on demand, but are able to make illiquid investments because the number of early withdrawals is predictable. Diamond and Dybvig were, of course, mainly concerned with showing how such a system could be vulnerable to self-fulfilling bank runs; but we can also use their approach as a device for putting intermediates and monetary aggregates into our basic model.

To do this, return again to a one-good endowment economy, but now suppose that at the beginning of each period there is actually a three-step process, as follows:

1. Individuals trade currency for bonds in a capital market; they are also able to make deposits at a class of banks.
2. Individuals discover whether they derive utility from consuming in the current period.

3. Those who do want to consume withdraw the necessary cash from their bank accounts.

The determination of the real interest rate is, of course, somewhat more complex here, because while the representative-agent assumption may hold ex ante, it does not hold ex post. However, given the equilibrium real rate, it is straightforward to see what must happen in the financial sector. As long as the nominal interest rate is positive, individuals will have no incentive to hold on to cash; instead, they will deposit enough in their bank accounts to cover their cash needs if they do turn out to be type-1 consumers. Banks in turn will have to hold enough of the deposits they receive in cash to cover such withdrawals; again, given a positive nominal interest rate, they will hold no more than the minimum required, putting the rest in bonds. So at the beginning of the period a monetary aggregate defined as currency plus deposits will actually consist of no currency, plus a volume of deposits that is a multiple of the base money held as reserves.And any increase in that base will, under conditions of full employment, lead to an equal proportional increase both in deposits and in the price level.

But what if the nominal interest rate is driven to zero? Then consumer and banks will both become indifferent between holding monetary base and holding bonds (and consumers will also be indifferent between both and bank deposits). Exactly what happens to an increase in the monetary base under these conditions is indeterminate: it could be absorbed by consumers, who might substitute cash either for bonds or for bank deposits in their portfolios; or the extra base could be absorbed by banks, who will simply hold excess reserves. Of these three possibilities, only one - consumers substitute cash for bonds, rather than deposits - will have any effect on a currency-plus-deposits measure of the money supply; either a substitution of cash for deposits or an

addition of base money to reserves will reduce bank credit but leave the monetary aggregate unchanged. And in any case there will be no effects on the price level, or on output if prices are sticky.

Applying what one of my colleagues calls the principle of "insignificant reason", we may surmise that an increase in monetary base will lead to substitution in all three directions. This means that under liquidity-trap conditions such a base expansion will (1) expand a broad aggregate slightly, but only because the public holds more currency (2) actually reduce deposits, because some of that currency substitutes for deposits (3) reduce bank credit even more, because banks will add to reserves.

The implications of this thought experiment should be obvious. If an economy is truly in a liqudity trap, then failure of broad monetary aggregates to expand is not a sign of inadequately expansionary monetary policy: the central bank may simply be unable to achieve such an expansion, because additional base is either added to bank reserves or held by the public in place of bank deposits. On the other hand, this inability to expand broad money does *not* mean that the essential problem lies in the banking system; it is what you would expect to see even if the banks were in perfectly fine shape.

This is sufficiently important to bear repeating: under liquidity trap conditions, the *normal expectation* is that an increase in high-powered money will have little effect on broad aggregates, and even lead to a decline in bank deposits and a larger decline in bank credit. This seemingly perverse result is part of the looking-glass logic of the situation, and need have nothing to do with the problems of the banks per se.

Fiscal policy

We are now prepared to consider possible policy responses when an economy is in a liquidity trap. The classic Keynesian answer is, of course, fiscal expansion - which clearly does work in an IS-LM framework. How does it look in a modernized version of liquidity-trap theory?

The framework we have been using is strongly biased against finding any useful role for fiscal policy, because the representative-agent, intertemporal optimization approach implies Ricardian equivalence. This bias does not represent an empirical judgement: it is an accidental byproduct of modeling decisions made for the sake of simplicity on other fronts. True, a number of commentators have suggested (mainly because of the apparent ineffectiveness of Japan's efforts at fiscal stimulus to date) that Japan may come closer to Ricardian equivalence than most countries, and it is interesting at least as an exercise to think through the implications of such equivalence.[4] But in reality fiscal policy would surely have some impact. There are then two questions about that impact: one qualitative, one quantitative.

The qualitative question is this: can a temporary fiscal stimulus have permanent effects? If current income has very strong impacts on spending - so much so that the marginal propensity to spend (consumption plus investment) is actually greater than one over some range - there can be multiple equilibria. A liquidity trap may therefore represent a low-level equilibrium, and a sufficiently large temporary fiscal expansion can jolt the economy out of that equilibrium into a region where conventional monetary policy works again.

Is this story plausible? It seems to be part of the folk wisdom in macroeconomics that this is in fact how the Great Depression came to an end: the massive one-time fiscal jolt from the war

pushed the economy into a more favorable equilibrium. However, Romer (1992) argues that most of the output gap created during 1929-33 had been eliminated before there was any significant fiscal stimulus.[5] The main explanation of that expansion, she argues, was a sharp decline in real interest rates, which she attributes to monetary policy - although most of the decline in her estimate of the real interest rate is actually due to changes in the inflation rate rather than in the nominal interest rate. Indeed, Romer estimates that for most of the recovery period ex ante real rates were sharply negative, ranging between -5 and -10 percent.[6]

The point here is that the end of the Depression - which is the usual, indeed perhaps the sole, motivating example for the view that a one-time fiscal stimulus can produce sustained recovery, does not actually appear to fit the story line too well; much though by no means all of the recovery from that particular liquidity trap seems to have depended on inflation expectations that made real interest rates substantially negative.

If temporary fiscal stimulus does *not* jolt the economy out of its doldrums on a sustained basis, however, then a recovery strategy based on fiscal expansion would have to continue the stimulus over an extended period of time. The question then becomes how much stimulus is needed, for how long - and whether the consequences of that stimulus for government debt are acceptable.

Credibility and monetary policy

It may seem strange even to have a subsection mentioning monetary policy, given that everything up to this point has stressed the ineffectuality of such policy in a liquidity trap. However, as we noted at the beginning, only *temporary* monetary expansions are ineffectual. If a

monetary expansion is perceived to be permanent, it will raise prices (in a full-employment model) or output (if current prices are predetermined). The mechanism may be seen immediately from equation (6): a rise in the expected future price level P^* will shift out the "IS curve" in the current period.

The ineffectuality of monetary policy in a liquidity trap, then, is really the result of a looking-glass version of the standard credibility problem: monetary policy does not work because the public expects that whatever it does now, given a chance the central bank will revert to type, and stabilize prices near their current level. If the central bank can "credibly promise to be irresponsible", that is, convince the market that it will in fact allow prices to rise sufficiently, it can bootstrap the economy out of the trap. (Again, although she does not put it this way, Romer's (1992) analysis of the U.S. recovery from 1933-1941 suggests that just such a bootstrap process was the main cause of the growth in output).

Since first being widely aired a few months ago, proposals for "managed inflation" have drawn a number of questions; we may as well go through the FAQs and their answers.

1. *Why inflation? Isn't an end to deflation good enough*? If you believe the analysis given here, price stability is actually not an option for a liquidity-trap economy. The economy "needs" inflation, because it needs a negative real interest rate; the deflationary pressures it is actually experiencing are the economy "trying" to generate that needed inflation by reducing current prices compared with the future price level. The only way to avoid lowering the current level is to raise that expected future level.

2. *Isn't inflation a bad thing*? Again, if you believe the analysis, a liquidity-trap economy is

"naturally" an economy with inflation; if prices were completely flexible it would get that inflation regardless of monetary policy, so a deliberately inflationary policy is remedying a distortion rather than creating one. We might also note that one might arrive at the recommendation of inflation from a quite different route: Friedman's (1969) famous theory of the optimum quantity of money. Although he actually says that the economy should deflate at the rate of time preference, the proper interpretation of the logic is that it should deflate at the market-clearing real rate of interest - which for a liquidity-trap economy, where that market-clearing rate is negative, means a negative rate of deflation, i.e., inflation.

3. *Won't expected inflation produce perverse incentives*? In terms of the models, at least, a fall in the real interest rate achieved via expected inflation is identical in its effects to one produced via a fall in nominal interest rates when that is possible. There is no reason in principle to expect the increase in spending generated by a commitment to inflation to be any different in character from that generated by a conventional monetary expansion in an economy that starts with positive nominal rates.

4. *Won't an inflationary policy lead to a plunge in the exchange rate and be in effect a beggar-thy-neighbor policy carried out at the rest of the world's expense*? Because expected inflation plays the same role in a liquidity-trap economy that interest-rate reductions do under more normal circumstances, inflating one's way out of a trap is no more (and no less) a beggar-thy-neighbor policy than any monetary expansion under flexible exchange rates. But what is the beggar-thy-neighbor aspect of monetary policy, anyway?

In the traditional open-economy IS-LM model developed by Mundell (1963) and Fleming (1962), and also in large-scale econometric models, the effect of monetary expansion is

unambiguously to lead to currency depreciation; but there are two offsetting effects on the current account balance. On one side the currency depreciation tends to increase net exports; on the other side the expansion of the domestic economy tends to increase imports. For what it is worth, policy experiments on such models seem on balance to suggest that these effects very nearly cancel each other out. Table 1 presents estimates from the comprehensive if somewhat elderly model comparison by Bryant et al (1988). It shows, for 11 models, the second-year effects on the exchange rate and the current account of a monetary expansion sufficient to raise real United States GDP by 1 percent. The exchange rate impacts were substantial, the current account impact negligible. To the extent that these estimates are correct, they suggest that the impact of expected inflation on a large economy with fairly small trade shares will be a significant currency depreciation with small impact on the current account.[7]

However, we have been trying to get beyond the IS-LM model; how does the result look in an intertemporal open-economy model?

Let us again assume the utility function given in equation (10), exogenous output of traded goods, and sticky prices/excess capacity in the nontraded sector. If the nominal interest rate is positive, ordinary monetary policy can raise output of the nontraded good; if the economy is in a liquidity trap, expectations of future monetary expansion can achieve the same result. What we now want is to know what impact this expansion has on the current account; in this framework, that amounts to asking what happens to consumption of the traded good.

We can take a shortcut here if we imagine that the expansion is “brief”, in the sense that we can ignore the effect of the current account on the future investment income of the country. As shown in Appendix B, removing this assumption would only reinforce the results. In the case of a

"brief" expansion, it is possible to calculate analytically a *beggar-thy-neighbor coefficient*, defined as the ratio of the increase in the expanding country's current account surplus, measured as a share of GDP, to the percentage increase in its GDP. Appendix B shows that

$$B = \frac{1-\rho}{1-\rho-\frac{1}{\tau}} \qquad \textbf{(12)}$$

where, again, ρ is relative risk aversion and τ the traded share of consumption (and hence value-added in the economy). We see immediately that in the special case $\rho=1$ a monetary expansion has no effect on the expanding country's current account - which is roughly what the econometric exercises in Table 1 indicated. If relative risk aversion is higher than 1, expansion via inflation will come to some extent via a widened current account surplus - but the extent will depend inversely on the economy's openness, as measured by τ. Table 2 shows beggar-thy-neighbor coefficients for a range of values of ρ and τ. If one believes the folk wisdom that relative risk aversion is something like 2, and judgementally assumes that the tradeable share in the Japanese economy is not much more than 0.2, the implication is that an inflationary policy that raised Japanese output by as much as, say, 5 percent relative to baseline would require an expansion of Japan's current account surplus by something like 1 percent of GDP - far short of the huge surpluses envisaged by Smithers (1998) and others.

We can also calculate the real depreciation - as measured by the change in a domestic price index with weights τ and $1-\tau$ relative to the price of traded goods - associated with a one percent increase in real GDP achieved via creation of inflationary expectations. As shown in the appendix, this depreciation is

$$1 + \frac{1-\tau}{\tau}\frac{1-\rho}{1-\rho-1/\tau} \qquad \textbf{(13)}$$

Again, a range of values is shown in Table 3. These numbers look generally small compared with the model simulations in Table 1; we will discuss the reasons why this might be the case in Part 2 of this paper.

Summary

This part of the paper has offered a quick tour of a rather extensive and unfamiliar territory, the land of the liquidity trap. Perhaps the most important lesson to be learned from this tour is the strangeness of that territory: once one accepts that an economy really is in a liquidity trap, much of the conventional wisdom of macroeconomics ceases to apply - indeed, applying basically conventional models to the liquidity trap universe implies some quite unconventional conclusions. Aside from the observation that international capital mobility makes less difference than most economists probably suppose - an observation that actually applies to open-economy macroeconomics in general - I would highlight two conclusions in particular.

The first is that one must be careful about making inferences from divergences between the growth of monetary base and broad monetary aggregates. The failure of aggregates to grow need not indicate dereliction in the usual sense on the part of the central bank; in a liquidity-trap economy the central bank *in principle* cannot move broad monetary aggregates. For the same reason, the observation that the central bank has slashed interest rates and pumped up monetary base, but that the broader money supply has not grown, does not necessarily imply that the fault

lies in the banking system; again, it is just what you would expect to see in a liquidity-trap economy.

Second, whatever the specifics of the situation, a liquidity trap is always a product of a credibility problem - the belief by the public that current monetary expansion will not be sustained. Structural factors can explain why an economy "needs" expected inflation; they can never imply that *credibly sustained* monetary expansion is ineffective.

JAPAN'S TRAP

The Japanese slump

Table 4 presents some standard summary statistics on Japanese economic performance since 1981. It shows the familiar point that after rapid growth to 1991, Japan has gone through an extended period of very slow growth. The breakpoint shown in the table, however, is actually 1992 rather than 1991. The reason is that the Japanese economy appeared to be overheated in 1991, so that part of the slowdown in growth as measured from that date can be viewed as simply a correction of an unsustainable boom. By 1992, however, inflationary pressures had clearly eased, so that the low growth rate thereafter is a better indicator of the economy's true shortfall. Of course, it is clear that Japan will have a significant decline in real GDP this year. And the unemployment rate has already risen above 4 percent.

There are two striking features of these familiar and dreary numbers. The first is the extent of the slowdown. In the period 1981-1992 Japan grew at an average rate of 3.7 percent, ending with

the same unemployment rate with which it started, and with a lower inflation rate than at the start; in short, potential as well as actual output would seem to have risen at about 3.7 percent annually over that period. If one had projected that growth rate forward, one would have over-predicted 1998 output by about 14 percent.

The second striking feature is the low interest rates of recent years; Japanese money-market rates have been below 1 percent since 1995. It is true that Japan has not pushed money market rates down to their absolute minimum - there were still, at the time of writing, 43 basis points to go - but the economy is clearly in a very good approximation to liquidity-trap conditions.

How important a role does this liquidity trap play in the growth slowdown and current slump? It might in principle be the case that the great bulk of the slowdown represents a reduction in the rate of potential output growth; in that case even a successful stimulative policy would have only a small payoff, so that freeing the economy from its liquidity trap is not a particularly urgent issue. So it is important to try to estimate the gap between actual and potential output.

In the United States, the output gap is usually estimated by combining an estimate of the natural rate of unemployment with an estimate of the Okun's Law coefficient between changes in unemployment and in real GDP. Although Japan's measured unemployment rate has traditionally moved much less than that of the United States, there is actually a surprisingly close Okun's Law relationship in the 1981-91 period, shown in Figure 4. The slope of the apparent relationship is about 3 times as steep as in the United States: it apparently took about 6 percentage points of excess growth to reduce the unemployment rate by one percentage point. If we were to take the average 2.5 percent unemployment rate in the pre-slump period as an estimate of the natural rate, the 3.4 percent unemployment rate in 1997 would therefore seem to imply an output gap of more

than 5 percent last year - and with potential output still presumably growing while output slumps, the gap by end-1998 could be as high as 10 percent.

Most published estimates of Japan's output gap are far smaller than this. Many of these estimates, notably those of the International Monetary Fund, are based on the Hodrick-Prescott filter, which minimizes a weighted sum of squared deviations of actual from potential output and squared changes in the growth rate of potential output. (See Giorno et. al. 1995 for a description). The main practical advantage of this method is that it can be used even when there are secular changes in both potential growth rates and the natural rate of unemployment. However, Hodrick-Prescott has severe disadvantages when applied to an economy that undergoes a sustained slump. For one thing, it imposes the assumption that average deviations from potential are zero over the whole period, so that when the economy slumps the filter automatically re-evaluates earlier periods as times of above-potential output, reducing the estimated shortfall. Also, any sustained drop in output gets built into the estimated potential growth rate. The result is to systematically understate the actual shortfall from potential. A stark if somewhat unfair way to make this point is to apply Hodrick-Prescott to the U.S. in the interwar period, as in Figure 5. This calculation was done with the smoothing parameter λ set at 25; but for a wide range of values of λ one comes up with the conclusion that output was in excess of potential by 1935.

The OECD, whose most recent estimates of Japan's output gap are shown in Figure 6, has adopted a more complex technique (described in Giorno et al 1995); nonetheless, the estimated output gap in 1997 is a remarkably small -1.2 percent. This result seems to be due to the fact that the OECD, while not engaging in simple Hodrick-Prescott filtering, does in fact update estimates of normal worker hours and worker productivity in such a way that possibly cyclical components

get reinterpreted as structural trends. A nice illustration of this process at work is to contrast the estimates of potential growth in the 1995 study by Giorno et al that introduced the OECD's current method, and the potential growth estimates that appear in the most recent *OECD Economic Outlook*, as shown in Figure 7. As recently as three years ago, the OECD estimated Japan's potential growth at 3 percent; now it has been marked down to 1.6. If the original 3 percent potential growth were applied to the period since 1994, the 1997 output shortfall rises to 4.6 percent; not too far short of the estimate suggested by the Japanese Okun's Law calculation.

If the 1997 output gap was 3-4 percent; if potential output growth in Japan is 2-3 percent; and if, as now seems certain, output falls throughout 1998, then the output gap at end year will be quite probably exceed 7 percent. Obviously there is no precision in this estimate; my personal guess is that in retrospect it will seem clear that Japan's 1998 output gap was 8 percent or more. But we can make a very strong case that it will exceed 5, making demand-side policies to close that gap of very real importance.

Savings and investment

A liquidity trap happens when desired savings exceed desired investment at full employment, even at a zero short-term interest rate. As argued in the first part of this paper, for some purposes it does not matter why this is the case, as long as it is. Still, the intepretation of Japan's problem, and to some extent the policy implications, do depend on how we view the apparent excess savings.

Table 5 shows consumption/GDP ratios for Japan and the United States since 1991. Two

familiar observations stand out. The first is that Japan's consumption ratio remains very low by comparison with the United States. A vast literature has attempted to explain this disparity; this paper is in no position to add to it. The other observation is that Japan's consumption ratio has not declined in the 1990s; on the contrary, if anything it has risen slightly. This at least suggests that the shift into liquidity-trap territory reflects declining investment demand rather than rising savings supply.

How significant is the difference between U.S. and Japanese consumption ratios? Consider the 1997 difference of approximately 7 percentage points of GDP. If US consumers were suddenly to start behaving like their Japanese counterparts, this would be the equivalent of a 7 percent of GDP negative fiscal impulse. Suppose that the Fed then tried to offset that contraction with looser monetary policy. Would it be able to do so, or would we find ourselves in liquidity-trap territory? One way to look at this is to notice that the emergence of U.S. budget deficits of approximately 3 percent of GDP in the 1980s was widely held to have raised real short-term interest rates by 3 or 4 percentage points; the difference between Japanese and U.S. consumption shares is more than twice as large a shock. It also is possible, once again using the comparison of standard econometric models in Bryant et al (1988), to ask what the impact of such a monetary-fiscal switch would be in variety of such models. Table 6 reports estimates of the impact of a monetary-fiscal switch equivalent to a *one* percent decline in the consumption ratio; both the mean and the median effect is a 3 percentage point decline in the short-term interest rate.

One can easily list reasons why such exercises might overstate the case; probably structural models tend to understate the spending impact of a sustained reduction in the interest rate. Nonetheless, even this crude comparison makes it substantially less surprising than one might

otherwise have supposed that Japan, with its low consumption, has indeed found itself in a liquidity trap.

In fact, this exercise suggests that the real puzzle is not why Japan is now in a liquidity trap, but why this trap did not materialize sooner. How was Japan able to invest so much, at relatively high real interest rates, before the 1990s? The most obvious answer is some version of the accelerator: investment demand was high because of Japan's sustained high growth rate, and therefore ultimately because of that high rate of potential output growth. In that case the slump in investment demand in the 1990s may be explained in part by a slowdown in the underlying sources of Japanese potential growth, and especially in *prospective* potential growth.

As noted above, there is considerable uncertainty about the actual rate of Japanese potential growth in the 1990s. Nonetheless, it is likely that there has been a slowdown in the rate of increase in total factor productivity, even cyclically adjusted. What is certain, however, is that Japan's long-run growth, even at full employment, must slow because of demographics. Through the 1980s Japanese employment expanded at x.x percent annually. However, the working-age population has now peaked: it will decline at x.x percent annually over the next xx years (OECD 1997), and - if demographers' projections about fertility are correct - at a remarkable x.x percent for the xx years thereafter. As suggested by the discussion of investment and *q* in the first half of this paper, such prospective demographic decline should, other things equal, depress expectations of future *q* and hence also depress current investment.

Of course, the looming shortage of working-age Japanese has been visible for a long time; indeed, the budgetary consequences of an aging population have been a preoccupation of the Ministry of Finance, and an important factor inhibiting expansionary fiscal policy. Why, then,

didn't this prospect start to affect long-term investment projects in the 1980s? One answer is that businesses may have believed that total factor productivity would grow rapidly enough to make up for a declining work force. However, the “bubble economy” of the late 1980s may also have masked the underlying decline in investment opportunities, and hence delayed the day of reckoning.

That bubble economy, of course, also left a legacy of large debts and troubled bank balance sheets, which are widely regarded as the main villains of Japan's current plight. Thus we turn next to the role of banking in Japan's malaise.

Banking problems

Japan clearly faces a huge problem of bad bank loans; the current conventional wisdom is a trillion dollars of bad loans. These bad loans are in part a legacy of the burst of the 80s asset bubble, reinforced by the consequences of the slow growth since. Clearly Japan will need to engage in a cleanup operation that will dwarf the U.S. thrift crisis, especially measured against Japan's smaller economy. Inevitably, also, the form and funding of that cleanup will be a central political preoccupation. But how central are the problems of banks to the country's macroeconomic difficulties?

This may seem an odd question to ask. Disruption of financial intermediation has clearly played a crucial role in many if not most historical financial crises, including the current crisis in emerging Asia. Why should Japan be different? Also, to many economists it seems *a priori* obvious that if conventional monetary policy has become ineffective, it must be because the

troubles of the banks have blocked the usual channels of central bank influence.

A casual look at the data seems to support the view that the problem with monetary policy lies in the banks. Table 7 shows developments in high-powered money, broad money, and bank credit since end-1994. Clearly a fairly rapid growth in monetary base has failed to produce an equivalent growth in broad monetary aggregates, and has actually been accompanied by stagnation in bank credit.

However, recall the discussion of financial intermediation under liquidity-trap conditions in the first part of this paper: given that an economy is in a liquidity trap, this sort of disconnect between monetary base, aggregates, and bank credit is what one would expect to see *even if the banks were financially healthy*. It provides no evidence that the problems of banks aggravated the problem.

It is important to realize that Japan has *not* (yet?) suffered from any widespread run by depositors: in this sense Japanese banks are like the U.S. thrifts, whose financial woes were widely recognized well before the cleanup began, but whose depositors remained calm because of the underlying government guarantee. As a result, Japanese banks have not been forced into the kinds of fire-sale liquidations of loans, abrupt removal of credit lines, etc. that produce a classic bank-centered financial crisis - the kind of crisis that has afflicted its emerging-economy neighbors.

In the absence of a bank run, however, how would we expect a bank of questionable solvency to behave? Would it restrict credit? The textbook answer (which now plays a major role in discussions of emerging Asia's troubles - see McKinnon and Pill (1997), Krugman (1998), Corsetti et al (1998)) is just the opposite: as long as an insolvent or near-insolvent bank is able to

hold on to deposits thanks to government guarantee, it has an incentive to overlend to risky projects, in effect playing “heads I win, tails the taxpayer loses”. Indeed, one could argue that Japan's financial institutions in the post-bubble years have actually been in the situation of U.S. thrifts *before* the crackdown, with the moral hazard of their position creating a bias toward too much rather than too little lending.

Nor is this merely abstract speculation. Japan has already gone through a miniature version of the system-wide bank cleanup it must now undertake, involving the *jusen* (non-bank subsidiaries of financial institutions specializing in housing loans). According to Cargill et al (1997), *jusen* lending actually grew rapidly in 1990-91, even as asset deflation was underway, “as a result of funds provided by agricultural cooperatives and their prefectural associations”. Because these agricultural cooperatives had strong political influence, they were able to take large risks while counting on an implicit government guarantee; and the result was behavior strongly reminiscent of the U.S. thrifts.

Cargill et al offer a striking example of moral-hazard-driven lending in the case of two credit cooperatives that failed in November 1994. The relevant authorities apparently knew that these cooperatives were insolvent more than a year before their actual closure; presumably the management knew considerably earlier. What actually happened in the two years before the institutions were closed, however, was a rapid *expansion* in both their deposits and their loans.

But how can the logic of excessive lending by banks be reconciled with tales of credit crunch? The immediate answer is that such tales are a very recent phenomenon. An informal search of news archives finds few allegations of credit rationing in Japan before the second half of 1997; even well into the fall of last year a number of observers questioned whether there was really any

credit crunch, or at least whether it was serious. Only by early 1998 did the credit squeeze become a widely accepted phenomenon.

The reasons for the emergence of credit constraints in late 1997 is also quite clear from a review of press reports. The immediate forcing event was the announcement, in tk 1997, of new capital adequacy standards, effective April 1998. To meet this standard, banks began cutting back on loans that would have required larger capital backing. In other words, the financial problems of the banks only became a drag on aggregate demand when the government began half-hearted efforts to come to grips with those problems.

More generally, we can argue that since late last year the prospect of eventual government seizures of some but not all banks has created a new incentive for banks near the edge to dress up their balance sheets, in order to make the cut. The payoff to banks successful in this endeavor is, loosely speaking, that they will then be able to live to make bad loans again; or to say it somewhat differently, they want to stay out of government hands at least for a while in order to capture the value of the put option implied by government deposit guarantees.

This should all sound familiar to U.S. economists. A mild form of the same ailment appeared in 1990-1992, when the size of the S&L bailout had become apparent and there was widespread discussion that commercial banks might be next. As in the Japanese case, the credit crunch appeared not during the years when banks were getting into financial trouble, but at the point when the government began to look likely to do something about it.

If the threat of bank closures or seizures is causing a credit crunch that has deepened Japan's slump, why engage in bank reform at all? The answer is that cleaning up bad banks is a *microeconomic* policy, undertaken to remove the distortion in the direction of investment that

results from moral hazard (and also to limit the eventual liability of the government, since - as both the S&L case and Japanese experience with credit cooperatives so graphically demonstrate - delay only multiplies the losses). If it reduces aggregate demand as a side consequence, so what? Under normal circumstances the macroeconomic effects of this (or any other) move toward microeconomic efficiency that happens to discourage spending can simply be offset with a looser monetary policy. Japan's problem, of course, is that because it is in a liquidity trap the normal disconnect between micro and macro policy no longer applies.

Policy options and their consequences

Given all that we have said, what are Japan's policy options, and how well would they work? Current discussion focusses on three basic alternatives (which are not, of course, mutually exclusive).

- *Fiscal expansion*: This is the classic remedy for a liquidity trap, and has been pursued by Japan in a sort of stop-go fashion for much of the post-92 period. At the time of writing the traditional emphasis on public works seems to have given way to a new emphasis on “permanent” tax cuts.

There are two major question about fiscal expansion as a remedy for Japan, one strictly economic, one political. The economic issue is whether an adequate expansion is possible without unacceptable impact on the government long-term fiscal position. Much discussion of fiscal stimulus in Japan seems to be predicated on some pump-priming (or, if my suggestion is adopted, jump-start) idea: a brief period of stimulus will jolt the economy back into a favorable equilibrium.

However, there is no good evidence for such a multiple-equilibrium view; indeed, Romer has argued that even the historical episode usually invoked in support of that view, the U.S. recovery from the Great Depression, has been misinterpreted. If one alternatively views Japan as having a long-term deficiency of demand due to low rates of time preference combined with negative-population-growth demographics - and also views the output gap as being on the order of 7 percent or more - the implied size and duration of the deficits implied would be very large. (A useful indication of the seriousness of the situation is that 10-year government bond rates in Japan are now less than 0.7 percent, suggesting that investors expect Japan to be in or near a liquidity trap for at least a decade). Of course, if one expects interest rates to stay near zero indefinitely, the level of government debt hardly matters. But if one expects that at a sufficiently distant date real rates will become strongly positive again, the eventual size of that debt becomes an important concern.

The political point is that Japan - like, we might note, the United States during the New Deal - appears to have great difficulty working up its political nerve for a fiscal package anywhere close to what would be required to close the output gap. Exactly why is an interesting question, beyond this paper's scope.

Does this mean that fiscal policy should be ignored as part of the policy mix? Surely not. On the general Brainard principle - when uncertain about the right model, throw a bit of everything at the problem - one would want to apply fiscal stimulus. (Even I wouldn't trust myself enough to go for a purely "Krugman" solution). However, it seems unlikely that a mainly fiscal solution will be enough.

- *Banking reform*: Japan clearly needs to clean up its financial system. Many commentators seem to believe that this urgent microeconomic step will also make a major contribution to solving the macroeconomic problem. However, as we have seen, if anything the financial problems of the banks have until recently biased them toward lending too much rather than too little.

The irony of the situation is that indications that the Japanese government is finally getting its nerve up to do something about the banks have probably been a significant factor in the economy's slide over the past year. From a macroeconomic as opposed to microeconomic view, a situation in which the government is expected to start seizing banks but has not yet done so is the worst of all possible worlds. Assuming that the government will eventually do a financial cleanup, the most important thing is to get on with the job and get it over with. If Japanese authorities behave true to form and carry out bank seizures and closures slowly, initially adopting excessively lenient criteria and only gradually tightening them, credit constraints could be a depressing factor on the economy for years to come.

And even a radical, forceful bank cleanup - one that basically settles the issue and leaves the remaining banks reasonably sure that they will not be taken over - would in principle still leave the banking system no more willing to lend, and in fact somewhat less so, than it was a year ago. The reason is that until the second half of 1997, at least some banks were driven by moral hazard to take excessive risks in their lending; once the system has been cleaned up, that extra boost to aggregate demand will be gone.

A financial cleanup, in short, is vital on micro grounds; and given that it must be done, on macroeconomic grounds 'twere best that it were done quickly. But it is unlikely to bootstrap Japan out of its liquidity trap.

- *Managed inflation*: Thanks to the Internet (Nouriel Roubini has become the Matt Drudge of the Asian crisis), proposals for Japan to adopt an inflation target as an answer to its liquidity trap have become the subject of widespread, if not always well-informed, discussion. The logic of such an approach was laid out in the first part of this paper. In Japan's case, there would be three main questions: implementation, the appropriate target, and the likely effects.

How can a country that is in a liquidity trap - that is, where increases in the money supply seem to have no effect - engineer inflation? As we have seen, the problem is essentially one of credibility. If the central bank can credibly commit itself to pursue inflation where possible, and ratify inflation when it comes, it should be able to increase inflationary expectations despite the absence of any direct traction on the economy via current monetary policy. Indeed, if one views monetary policy in terms of nominal interest rates, a credible commitment to inflation can seem to be a pure bootstrap policy: interest rates never actually need fall, all that is required is a promise not to raise them when the economy expands and prices begin to rise.

How to actually create these expectations is in a sense something outside the usual boundaries of economics. However, one obvious suggestion is that Japan deal with its inverted credibility problem by passing a law giving the Bank of Japan an inverted version of the price stability targets now in force in a number of countries: the BOJ would be enjoined to achieve an inflation rate of not less than “x” percent over “y” years. (What if this doesn’t work? Appendix C discusses several ways in which the necessary inflation expectations might nonetheless be generated).

Which brings us to the question of the appropriate inflation target. A key insight here is that the objective of the inflation target is not particularly exotic: it is simply to reduce the real interest rate sufficiently to bring the economy back to potential output. Although this real interest

reduction must be achieved via inflation because the nominal interest rate is up against the zero constraint, in other respects it should act just like a conventional monetary expansion. So we can estimate the size of the necessary inflation simply by asking how large a real interest rate reduction would normally be needed to eliminate an output gap of Japan's size.

We might also note that while the theoretical models of the first part of this paper were cast in terms of a one-period liquidity trap, we have no real idea how long a “period” is. However, Japan's liquidity trap looks like a fairly long-term problem; also, investment and exchange rates are generally believed to be driven by long-term interest rates. So Japan probably requires a sustained - at least a decade - period of inflation, sufficient to reduce the real long-term rate enough to close the output gap.

At this point matters become difficult. The size of Japan's output gap is, as we have seen, highly uncertain, although it is probably well over 5 percent. Worse yet, there is no consensus on the stimulative effect of a given interest rate reduction. As in earlier discussions, it may be useful to look not at the small number of estimates for Japan, but at the larger range of estimates for that other large, relatively closed advanced economy, the United States. Table 8 shows estimates of the reduction in long-term interest rates needed to expand real US GDP by 1 percent.

Given these uncertainties, any number is a matter of multiplicative guesswork. At this point I would suggest the following series of leaps of faith: although Japan's current output gap is probably well over 5 percent, the combination of fiscal stimulus and - if all goes well - a clarification of which banks will be taken over and which will not should reduce that gap by several percentage points. So managed inflation would need to close a remaining gap of, say, 4-5 points. Given the median estimate in Table 8, this would require an inflation target of 3-3.75

percent. So to give a bit of extra room (one can always raise nominal interest rates if the economy seems to be overheating - as long as the inflation target is met), how about 4 percent inflation for 15 years?

This target should not really be taken seriously. It should, instead, serve mainly to stimulate serious research; there is probably time for such research, since it will take some time before the idea of managed inflation overcomes the instinctive negative reactions of many policymakers.

However, once we have raised the possibility of such a solution to Japan's slump, what might we expect the side consequences to be? In particular, what would happen to Japan's current account and the value of the yen?

But we have already seen some numbers on that. Again, a policy of managed inflation is in principle simply a monetary expansion by other means. Typical estimates suggest that a monetary policy that expands output by 1 percent leads to a depreciation on the order of 5 percent. So the implied yen depreciation from such a policy would be on the order of 20-25 percent - a number that is actually probably less uncertain than the required inflation rate, although still more of a stimulus to debate than a serious estimate.

CONCLUDING REMARKS

Japan’s economic difficulties are widely viewed as essentially political: if only Japanese politicians would bite the bullet, they would get their country moving again. But in fact it has been far from clear what exactly Japan should be doing - which is to say that the problems are not so much political as conceptual.

In this paper I have argued that to understand Japan's problems we need to revive and modernize the theory of the liquidity trap - a concept that once played a major role in macroeconomics, but had virtually disappeared from economic discourse in the past 20 years. Taking liquidity traps seriously does not, it turns out, require a rethinking of the fundamentals of macroeconomics; liquidity traps can quite easily be generated in basically conventional models, models that meet the modern criteria of rational behavior and intertemporal consistency. It is even possible to have full-employment, flexible-price analyses of the liquidity trap. However, applying conventional modeling to liquidity-trap conditions produces unconventional conclusions and policy recommendations. My claim is that strange as they may seem, these conclusions are the best guide we have to dealing with Japan's malaise.

Nor is Japan, important as it is, the sole issue. Nobody thought that a liquidity trap could happen in Japan; now that it has, we should wonder whether it could happen elsewhere. Germany and France currently have short-term interest rates of only 3.5 percent, and Europe faces Japan-style demographics; could a liquidity trap happen to EMU? We now know that the liquidity trap is not a historical myth: it can and does really happen sometimes, and we had better try to understand it.

REFERENCES

Bernanke, B. (1994) "The macroeconomics of the Great Depressions: a comparative approach" *NBER* working paper # 4814.

Bryant, R., Henderson, D., Holtham, G., Hooper, P., and Symansky, S.,eds. (1988): *Empirical Macroeconomics for Interdependent Economies*, Washington: Brookings.

Cargill, T., Hutchinson, M., and Ito, T. (1997), *The Political Economy of Japanese Monetary Policy*, Cambridge: MIT Press.

Cooper, R., and Corbae, D. (1997), "Financial Fragility and the Great Depression", NBER Working Paper #6094

Corsetti, G., Pesenti, P. and Roubini, N. (1998), "Paper Tigers? A Preliminary Assessment of the Asian Crisis" (mimeo)

Diamond, D. and Dybvig, P. (1983), "Bank runs, deposit insurance, and liquidity", *JPE* 91, 401-19.

Fleming, M. (1962), "Domestic financial policies under fixed and under floating exchange rates", *IMF Staff Papers* 9, 369-79.

Friedman, M. (1969), *The Optimum Quantity of Money and other Essays*, Chicago: Aldine.

Friedman, M. and Schwartz, A. (1963), *A Monetary History of the United States*, Princeton: Princeton University Press.

Giorno, C., Richardson, P., Roseveare, D. and van den Noord, P. (1995), "Potential output, output gaps and structural budget balances", *OECD Economic Studies 24*, 167-208.

Hicks, J.R. (1937), "Mr. Keynes and the classics", *Econometrica*

Krugman, P. (1998), "What happened to Asia?", mimeo.

McKinnon, R. and Pill, H. (1997), "Credible economic liberalizations and overborrowing", *American Economic Review* 87, 189-93.

Metzler, L. (1951), "Wealth, saving, and the rate of interest", *Journal of Political Economy* 54, 93-116.

Mundell, R. (1963), "Capital mobility and stabilization policy under fixed and flexible exchange rates", *Canadian Journal of Economics and Political Science*, 29, 475-85.

Romer, C. (1992), "What ended the Great Depression?", *Journal of Economic History* 52, 757-84.

OECD (1997), *Economic Survey of Japan*, Paris: OECD

Smithers, A. (1998), *Japan's problems of debt and demography*, Smithers and Co.

Temin, P. (1976), *Did Monetary Forces Cause the Great Depression*? New York: W.W. Norton.

APPENDIX A: FINANCIAL INTERMEDIATION AND MONETARY AGGREGATES IN A LIQUIDITY TRAP

In the text I sketched out how one might think about the role of financial intermediaries and the behavior of monetary aggregates in a liquidity trap. This appendix describes the "cash-in-advance meets Diamond-Dybvig" approach suggested there more fully.

We consider for simplicity a full-employment endowment economy that lasts for only two periods, with each individual receiving an endowment y_1 in period 1, y_2 in period 2 . In the aggregate there is no uncertainty; however, each individual is uncertain ex ante about when he will want to consume. The assumed utility function takes the form

$$U = HU_1(c_1) + (1-H)U_2(c_2) \qquad \textbf{(14)}$$

where H takes on the value 1 with probability π, 0 with probability 1-π. So in the population there will turn out to be a fraction π of first-period consumers, 1-π of second-period consumers.

We want to make this a cash-in-advance economy. Finite horizons, however, pose problems for a fiat-money economy, while trying to have an infinite horizon would complicate the simple Diamond-Dybvig-type logic considerably. As a device for sidestepping these problems, we assume that each individual is issued with a quantity of money M^* at the beginning, which must be repaid at the end. The government may, however, inject additional money into the economy via open-market operations, as described in a moment.

Within each period, consumers must pay in cash before they receive income from selling their own endowment. As in the basic model in the text, they are able to trade cash for bonds at the

beginning of the first period (including bonds issued or purchased by the government in open-market operations). However, in order to motivate financial intermediaries, we assume that a consumer does not know his own type until *after* the capital market; so he can no longer simply acquire "just enough" cash for planned purchases within the period.

Here is where financial intermediaries come in. We assume that there exist banks which accept deposits during the initial capital market, then allow customers to withdraw their deposits if they turn out to have H=1. (Issues of bank runs are left on one side for this paper!) Deposits earn competitive interest if not withdrawn.

The sequence of events, then, looks like this:

1. Consumers come into existence, and receive the money supply M^*.
2. A capital market is held; consumers deposit money in banks, and open-market operations may increase or decrease the monetary base.
3. Consumers learn their type.
4. They withdraw their funds if necessary.
5. Consumers receive income from sale of their endowment, receive bonds and deposits, and pay/receive whatever tax/transfer is needed to
6. Consumers purchase second-period consumption.
7. They receive income from sales of endowment, and repay their money to the government.

In this setting, the real interest rate is determined independent of the money supply. Each individual gets to spend the present value of his endowment in the appropriate period. Thus a period-1 consumer will get to purchase $y_1 + y_2/(1+r)$ units of the good in period 1; but since a fraction π of consumers is type 1, we must have

$$\pi[y_1 + y_2/(1+r)] = y_1 \qquad \textbf{(15)}$$

implying that the real interest rate is

$$1 + r = \frac{\pi}{1-\pi}\frac{y_2}{y_1} \qquad \textbf{(16)}$$

Assume provisionally that the nominal interest rate is positive. Then the behavior of consumers and banks is straightforward. Consumers will borrow, establishing bank accounts equal to Pc_1, the amount they will spend if they are type-1; they will hold no cash. Banks, however, need hold only a fraction π of their deposits in reserves, and will hold no more than necessary; they lend the rest out (which is where consumers get the money for the deposits). So bank deposits will be a multiple $1/\pi$ of the monetary base; the velocity of base will be 1, that of deposits π. And from here on the model will work pretty much the same as the pure outside-money model in the text.

But what happens if the government increases M relative to M^* to such an extent that the nominal interest rate goes to zero - which can clearly happen here, just as in the simple endowment model with no uncertainty. First, consumers become indifferent between holding cash and holding deposits; second, they become indifferent between cash and bonds; finally, banks also become indifferent betwen cash and bonds. At this point any further open-market bond purchase by the government could be absorbed in three ways:

1. Consumers could create new bonds to sell to the government, and simply hold extra currency.
2. Banks could sell bonds to the government, and add the cash to their reserves.

3. Consumers could sell bonds to the government instead of borrowing from banks.

It is indeterminate which would happen, since none of these actions has any effect either on real variables or on the price level. Action 1 would lead to some increase in common definitions of the money supply; the others would not. Action 3 would lead to an actual decline in bank credit. So as stated in the text, it is actually normal for increases in the monetary base to have little effect on broader aggregates, and even to reduce bank credit, when the economy is in a liquidity trap.

APPENDIX B: CURRENT ACCOUNT AND REAL EXCHANGE RATE CONSEQUENCES OF MONETARY EXPANSION

In the text I introduced a simple traded-nontraded good model to discuss the possibility of a liquidity trap despite the possibility of capital movement. In that model a monetary expansion - current money in a positive-interest environment, expected future money in a liquidity trap - can raise output of the nontraded good. But what is the impact on the current account? In this model that reduces to the question of what happens to traded-good consumption.

We can simplify this issue by starting with an economy in which trade is balanced, and normalizing initial prices of both traded and nontraded goods to one. In that case, we have initially that

$$\frac{c_T}{c_N} = \frac{\tau}{1-\tau} \qquad \textbf{(17)}$$

We can further simplify the issue by supposing that the monetary expansion - which leads to an increase in the production and consumption of nontraded goods - is “brief”, in the sense that it does not have a significant effect on the country’s net investment income from abroad. In that case we know that the level of consumption of both traded and nontraded goods in later periods will be unchanged, and hence also that the *marginal utility* of each good in later periods will be unchanged. However, the real interest rate on traded goods is given by the world capital market. Hence the marginal utility of traded goods even in the current period will remain unchanged. This marginal utility may be written

$$\frac{\partial U}{\partial c_T} = \tau c_T^{\tau(1-\rho)-1} c_N^{(1-\tau)(1-\rho)} \quad \textbf{(18)}$$

Now suppose that there is a monetary expansion. This will lead to an increase in c_N, which is also the increase in GDP at initial prices; it may also lead either to a fall or a rise in c_T, which corresponds to a move toward current account surplus or deficit. The change in c_T associated with a small rise in GDP can be evaluated as follows. First, we note that

$$\frac{\partial^2 U}{\partial^2 c_T^2} = \tau(\tau(1-\rho)-1) c_T^{\tau(1-\rho)-2} c_N^{(1-\tau)(1-\rho)} \quad \textbf{(19)}$$

and that

$$\frac{\partial^2 U}{\partial c_T \partial c_N} = \tau(1-\tau)(1-\rho) c_T^{\tau(1-\rho)-1} c_N^{(1-\tau)(1-\rho)-1} \quad \textbf{(20)}$$

Finally, we have

$$\frac{\partial c_T}{\partial c_n} = -\frac{\frac{\partial^2 U}{\partial c_T \partial c_N}}{\frac{\partial^2 U}{\partial c_T^2}} = -\frac{(1-\tau)(1-\rho)}{\tau(1-\rho)-1}\frac{c_T}{c_N} = \frac{1-\rho}{1-\rho-\frac{1}{\tau}} \quad \textbf{(21)}$$

which is the beggar-thy-neighbor coefficient described in the text.

APPENDIX C: CREATING INFLATION EXPECTATIONS

Suppose that one believes that Japan needs a negative real interest rate on a sustained basis, but also believes that a pure bootstrapping policy - in which the announcement of an inflation target generates the expansion that eventually creates the inflation - is infeasible. Then what Japan needs to do is apply some temporary policy that moves the economy into a position where monetary policy does have traction, then use that traction to generate sustained inflation.

In this case the temporary fiscal jolt once again comes into its own. The strategy would work along the following lines: a large fiscal expansion would be applied, with interest rates kept at zero, and sustained even as the economy began to develop inflation. Ideally the fiscal stimulus would then be phased out gradually, just slowly enough for rising expectations of inflation to take up the slack. The important point would be that monetary policy would have to remain accommodating, not only up to the point of full employment, but as inflation rose to the necessary level.

What kind of fiscal policy would be appropriate? One answer might be an explicitly temporary investment tax credit, which would more or less encourage the same kind of spending that an immediate, successful creation of inflation expectations would achieve.

NOTES

1. Strictly speaking, in traditional models money is not quite neutral when the private sector holds nominal claims on outside agents, such as government debt, because changes in the price level then have wealth effects on these assets - a point emphasized by Metzler (1951). However - even

leaving aside empirical doubts about the importance of the Metzler effect, and theoretical questions about its relevance (with Ricardian equivalence the effect goes away) - this complication can at most dampen the effect of money on the price level, not eliminate it.

That said, many macroeconomists bristle at the mention of monetary neutrality. The reason for their disdain, however, is the widespread belief (which I share) that because prices are not perfectly flexible increases in the money supply often get reflected mainly in output rather than prices. However, this has nothing to do with the puzzle of a situation in which increases in outside money cannot raise either output or prices, and indeed seem powerless to prevent deflationary pressures.

2. This summary of the standard remarks about Japan does not contradict my earlier assertion that almost everyone believes that money is approximately neutral. The point is that to my knowledge nobody has made this connection - that is, nobody has noticed that to say that monetary expansion is ineffective at raising output is equivalent to saying that it is ineffective at fighting deflation, and that this conflicts with the almost universally held belief in the near-neutrality of money.

3. Some commentators on an earlier draft seemed to believe that this possibility of monetary irrelevance depends on the assumption that the central bank is expected to defend a future price *level* target as opposed to an *inflation rate* target - that money becomes irrelevant only because the central bank creates expectations of future deflation.But when the equilibrium real interest rate is negative the liquidity trap emerges even if all the central bank wants is to keep prices stable. And assuming that the central bank has an inflation target leads to even more paradoxical results than the price level target. Since the economy “needs” inflation, attempting to keep the rate of change of prices constant means that there is no equilibrium price level at all: prices simply fall without limit.

If one make the more realistic assumption that prices are downward sticky in the short to medium run, this paradox disappears; then a committment to price stability, measured either by a predetermined target level or by inflation from the current level, will still imply a liquidity trap when the full-employment real rate is negative.

4. Suppose that one really believed that Japan was Ricardian-equivalent, or nearly so.The first implication is the obvious one that changes in taxes and transfers should have no effect. In the practical discussion of Japanese policy there has been much concern over whether tax cuts should be temporary or permanent; if one really believes in Ricardian equivalence, this discussion is irrelevant unless one believes that a “permanent” tax cut will constrain future government purchases of goods and services.

A less obvious point is that under liquidity-trap conditions the multiplier on government expenditures, for example public works projects, should be exactly one: that is, such projects will generate exactly as much additional income as the government spends, no more, no less. This may be seen directly by noticing that in our basic model current consumption is tied down by the Euler condition; if current policy cannot either raise expected future consumption or change the real interest rate, it cannot change current consumption. Alternatively, we can note that the extra income generated by government spending will be matched by an exactly equal present discounted

value of future tax liabilities. Either way, the point is that government spending will not generate any second-round increase in private spending.

A third point that has not been appreciated in some recent discussion is that if temporary tax cuts will not raise consumption, any other policy that can be reinterpreted as a temporary tax cut or transfer will be equally ineffectual. For example, several foreign commentators have suggested that the Japanese government promote consumption by issuing vouchers that must be spent within some short period. But individuals could presumably use the vouchers for purchases they would otherwise have made with cash; and if they take the future tax liability implied by the vouchers into account, they will do so, with no increase in spending.

A surprising corollary is that what is normally regarded as the most extreme inflationary monetary policy possible - a "helicopter drop" of cash - is just as ineffective in a liquidity trap as an open-market operation. After all, in a liquidity trap money and bonds are perfect substitutes. So a helicopter drop of currency is no different from a lump-sum transfer of bonds to the public - which, by Ricardian equivalence, has no effect.

These extreme results are, of course, implications of the strong assumption of completely rational, forward-looking consumption behavior.

5.The question of whether the economy had largely recovered from the Great Depression prior to the onset of massive wartime spending partly depends on choice of denominators. Significant fiscal stimulus began in 1941 (before Pearl Harbor - a massive military buildup was already under way). By 1940 real GDP had risen 70 percent from its 1933 level, but was only 11 percent above its 1929 level, so that a significant output gap surely remained. The "half full or half empty" issue is apparent in the contrast between Romer's discussion and that of Gordon (1988). Gordon views the U.S. economy in 1939 as stuck; Romer emphasizes the more than 8 percent growth rates in 1939 and 1940.

6. Romer uses commercial paper rates, which did decline somewhat even in nominal terms, in this calculation. However, the commercial vs. T-bill spread presumably is to some extent endogenous. T-bill rates averaged 0.515 percent in 1933 - that is, roughly the same as Japanese rates today. While they did fall virtually to zero by the end of the decade, any fall in real rates using this measure of nominal interest would be almost entirely dominated by changes in inflation expectations.

Indeed, if one views Romer's evidence through the lens of this paper's analysis, they seem to suggest a somewhat different interpretation of events than the one she gives. One way to think of what she finds is that the real expansion of the economy - and the rise in prices associated with that expansion - was the result of a rise in inflation expectations, which reduced real interest rates when nominal rates were already at the floor. Without this expected inflation, the expansion of monetary base that she stresses would have been ineffectual.

7. Of course, many people want Japan actually to act as a "locomotive" - to run much smaller current account surpluses, thereby aiding the recovery of neighboring economies. Perhaps the important point to make here is that even a large recovery in Japanese output would have only a small locomotive effect unless accompanied by a substantial *strengthening* of the yen. Typical estimates of the short-run income elasticity of import demand are around 2; given Japan's import

share in GDP of approximately 0.1, this means that a 5-percentage-point recovery would, at an unchanged real exchange rate, reduce Japan's surplus by roughly 1 percent of GDP, some $35 billion. Since only a fraction of this swing would come vis-a-vis troubled emerging economies, we are actually talking about fairly small change here. The only way to get a much larger locomotive effect would be for Japan to have a Reagan-style expansion, in which the exchange rate appreciates substantially. However, given the great difficulty Japan is having achieving a recovery at all, advocating a currency *appreciation* seems rather strange.