

PAUL R. KRUGMAN

Massachusetts Institute of Technology

It's Baaack: Japan's Slump and the Return of the Liquidity Trap

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 in the early years of macroeconomics as a discipline. John Hicks, in introducing both the IS-LM model and the liquidity trap, identified the assumption that monetary policy is ineffective, rather than the assumed downward inflexibility of prices, as the central difference between Mr. Keynes and the classics.¹ It has often been pointed out that the Alice in Wonderland character of early Keynesianism—with its paradoxes of thrift, widows' cruises, 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 1930s interest rates were hard up against the zero constraint; the average rate on U.S. 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. In part, this is because in the generally inflationary decades after World War II nominal interest rates have stayed comfortably above zero, and therefore central banks have no longer found themselves “pushing on a string.” Also, the experience of the 1930s itself has been reinterpreted, most notably by

1. Hicks (1937).

Milton Friedman and Anna Schwartz.² Emphasizing broad aggregates rather than interest rates or the monetary base, Friedman and Schwartz argue, in effect, that the Depression was caused by monetary contraction; that the Federal Reserve could have prevented it; and implicitly, that even the great slump could have been reversed by sufficiently aggressive monetary expansion. To the extent that modern macroeconomists think about liquidity traps at all (the on-line database EconLit lists only twenty-one papers with that phrase in title, subject, or abstract since 1975), their view is basically that a liquidity trap cannot happen, did not happen, and will not happen again.

But it has happened, and to the world's second-largest economy. Over the past several years, Japanese money market rates have been consistently below 1 percent, and the Bank of Japan plausibly claims that it can do no more; yet the Japanese 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 surely economists already understand liquidity traps well enough to formulate policy. Can we not 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 appeared 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 in several ways that require a rethinking of the issue.

In particular, one might identify three strands of modern thought that are missing from the classic IS-LM analysis. First is the intertemporal

2. Friedman and Schwartz (1963).

nature of decisions. Economists now understand, perhaps better than fifty 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, ignores foreign trade and capital mobility. It is 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, one needs to know what happens when this assumption is relaxed. Finally, traditional IS-LM analysis neglects the role of financial intermediaries. But how one interprets the experience of the 1930s hinges crucially on how broad a monetary aggregate one chooses; 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 lay at the heart of the problem; a similar view has become near orthodoxy about contemporary Japan. So one needs at least a basic sense of how financial intermediation fits into the picture of the liquidity trap.

There are two major parts to this paper. The first is an extended generic discussion of the causes and consequences of liquidity traps. I use 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 it 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 in fact be effective if the central bank can credibly promise to be irresponsible, to seek a higher future price level.

My 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 surefire 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 is that a successful monetary expansion, in which the central bank does create expectations of inflation, will probably be less of a beggar thy neighbor policy, expanding demand at the rest of the world's expense, than is widely imagined.

Second, putting financial intermediation into a liquidity trap framework suggests, *pace* Friedman and Schwartz, that it is quite misleading to look at monetary aggregates under these circumstances: 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 valid indicators of the stance of monetary policy and that their failure to rise does not indicate that the essential problem lies in the banking sector.

In the second part of the paper, I turn to some specific questions surrounding Japan. I survey other analysts' estimates to consider four main issues. First is the size of Japan's output gap. I argue that this is probably considerably larger than the standard estimates, and hence that the need for expansionary policy is even greater than is 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 Japanese banks are at the center of the problem, I argue that they have played less of a causal role than is widely assumed. Finally, I make a first attempt at quantifying the size, duration, and side effects of the inflation that would be needed to lift Japan out of its trap.

The Theory of Liquidity Traps

It is useful, in considering Japan's liquidity trap, to begin at a high level of generality, to adopt what one might almost call a philosophical stance. Popular discussion of the current situation has a strong tendency to plunge too quickly into the specifics, to cite one or another structural issue as *the* problem, missing the central point that whatever the details

of its history, Japan is now in a liquidity trap, so that the generic issues surrounding such traps apply.

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: 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 Japan in the 1990s must invoke some form of price stickiness, one can think of the unemployment and output slump that occurs under such circumstances as what happens when an 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 with which everyone in macroeconomics agrees it is that, aside from the possibility that price stickiness will cause 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.³ Or to be more specific, an increase in outside money—the monetary base—must raise prices.

Putting the issue this way immediately reveals that many of the common explanations of why Japanese monetary policy is ineffectual are wrong, or at least inadequate. One often hears, for example, that

3. 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). 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, but cannot eliminate it.

That said, many macroeconomists bristle at the mention of monetary neutrality. The reason for their disdain 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 in prices. However, this has nothing to do with the puzzle of a situation in which increases in outside money can raise neither output nor prices, and indeed seem powerless to prevent deflationary pressures.

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 supposed to raise prices regardless of the details of the transmission mechanism. Aside from the bad loans, one also 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, or output, or both. Recall that the neutrality of money is not a conditional proposition; it does not depend on banks being in good financial shape, or the service sector being competitive, or corporations not taking on too much debt. Money (which is to say, outside money) is supposed to be just plain neutral.⁴

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 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 suggests that a liquidity trap involves a kind of credibility problem. A monetary expansion that the market expects to be sustained (that is, matched by equiproportional expansions in all future periods) will always work, 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, one needs 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. Many macroeconomists believe that IS-LM is too ad

4. This summary of the standard remarks about Japan does not contradict my earlier assertion that almost everyone believes that money is approximately neutral. My point here 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.

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 does not depend on the ad hoc nature 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, as shown above, a liquidity trap fundamentally involves expectations and credibility; using models that explicitly recognize the intertemporal aspects of the problem helps to clarify this point. Let me therefore move immediately to an explicit intertemporal model that establishes relationships among output, money, prices, and interest rates. I then use this model as a base for a series of thought experiments and extensions.

Consider a one-good, representative agent economy (in which, however, agents must purchase their consumption from others). Suppose, initially, that the good is inelastically supplied, so that one 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

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

where c is consumption within a period, ρ is relative risk aversion, and D is 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 nominal interest rate i_t . Their consumption during the period is 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, it is assumed that the central bank is able to engage in open market operations during the beginning of period capital market, by 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 careful specification of the budget constraints of both individuals and the government, and of intertemporal choices. However, if one makes some simplifying assumptions, the model's implications can be derived with almost no algebra. Assume that from the second period onward, 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 one can immediately guess at the solution from period two 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 of the action, then, goes into determining the price level and interest rate in the first period (I use letters without subscripts to represent first period output, consumption, interest rate, and so forth). The 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. Thus the cash in advance constraint will be binding: $Pc = Py = M$, so that

$$(2) \quad P = M/y.$$

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

Figure 1. Relationships between Prices and the Interest Rate



Source: Author's model as described in text.

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

or, since consumption must equal output in each period,

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

This says that the higher is the current price level, the lower is 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.

The Liquidity Trap in a Flexible Price Economy

Suppose that one starts with an economy in the equilibrium described by point 1 in figure 1 and imagines an initial open market operation that increases the first period money supply. (Throughout, one imagines that the money supply remains unchanged from period two onward—or equivalently, that the central bank will do whatever is necessary to keep the price level stable from period two onward.) Initially, as I have shown above, 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. But what happens if the money supply is increased still further, so that the intersection of *MM* and *CC* is at point 3, with a negative nominal interest rate?

The answer clearly is that the interest rate cannot go negative, because money would then dominate bonds as an asset. Therefore it must be 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 (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.

Note 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 in 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 under such circumstances is to bear in mind that one is 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 lowers 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 . One knows 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 a maximum rate of deflation.

Now suppose that the central bank in effect tries to impose a rate of deflation that exceeds this maximum, by making the current money supply, M , large relative to the future supply, M^* . In this case 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 two is greater than that in period one, 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

$$(5) \quad (y/y^*)^p < D.$$

This condition might seem peculiar. After all, one normally thinks of economies as growing rather than shrinking. One possible answer involves an equity premium, another involves demography; but I reserve this issue for discussion below.

In a flexible price economy, 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 focused on whether wage and price flexibility were effective means of restoring full employment. In this model the problem does not arise, but for a reason that is a bit unusual: 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, which 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, rather than added to spending.

At this point one has a version of the liquidity trap: money becomes irrelevant at the margin.⁵ But aside from frustrating the central bank—which finds itself presiding over inflation no matter what it does—this trap has no adverse real consequences. To turn the analysis into a real problem, in both senses, one must 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' in period one. And suppose, also, that this productive capacity need not be fully employed. In particular, this paper assumes simply that the price level in period one is predetermined, so that the economy now acquires a Keynesian feel, and monetary policy can affect output. In period two and subsequently, output will still be assumed to take on the value y^* .

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

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

5. Some commentators on an earlier draft of this paper 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 the assumption that the central bank has an inflation target leads to even more paradoxical results. Since the economy needs inflation, attempting to keep the rate of change of prices constant means that there is no equilibrium price level: prices simply fall without limit.

If one makes the more realistic assumption that prices are downward sticky in the short to medium run, this paradox disappears. In this case, a commitment 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.

Figure 2. Relationships between Output and the Interest Rate

Source: Author's model as described in text.

Figure 2 illustrates the joint determination of the interest rate and output in this case. The IS curve, 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, giving the *MM* curve

$$(7) \quad y = M/p.$$

Increasing the money supply can now increase output, up to a point; specifically, up to point 2. But what if productive capacity is at point 3? The same argument as in the previous section applies: since the nominal interest rate cannot become 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, however large, can get the economy to full employment. In short, the economy is in a classic Hicksian liquidity trap.

Under what conditions will such a liquidity trap occur? One possi-

bility is that P is high compared with P^* ; that is, 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 equivalently, people's expected future real income is low compared with the amount of consumption needed to use today's capacity. In that case, it may take a negative real interest rate to persuade people to spend enough now, and with downwardly inflexible prices that may not be possible.

To put it yet another way, closer to the language of applied macroeconomics, if people have low expectations about their future incomes, even with a zero interest rate they may want to save more than the economy can absorb. (In this case, the economy cannot absorb *any* savings; I address that point below.) And therefore, no matter what the central bank does with the current money supply, it cannot reflate the economy sufficiently to restore full employment.

So I have now shown that a fully specified model, fudging neither the role of money nor the necessity of making intertemporal choices, can indeed generate a liquidity trap. The model does, however, omit some important aspects of standard macroeconomic models. Perhaps most notable, it has no investment, no foreign trade or capital mobility, and no financial intermediation, so that all money is outside. Can the same story be told if these elements are introduced?

Investment, Productive Capital, and Tobin's 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 saving and investment would be equal at potential output—is negative. An immediate question is how this can happen in an economy in which, in contrast with the simple endowment economy described above, productive investment can take place and the marginal product of capital, while it can be low, can hardly be negative.

One 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 easiest to make if one considers an economy with not capital but land (which can serve as a sort of metaphor for durable capital); and also if one temporarily departs from the basic framework to consider an overlapping generations model, 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 , defined as 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, one has a very simple determination of q_t , the price of land in terms of output: it must be true that

$$(8) \quad q_t A_t = w_t L_t,$$

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

The expected rate of return on purchases of land, however, 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

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

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 land can, in principle, be negative.

This is a highly stylized example, which begs many questions. Nevertheless, it at least establishes 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 saving over investment, even at a zero interest rate, is simply

to invest the excess savings abroad. In a recent influential study, Andrew Smithers 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.⁶ The general view seems to be that an open economy can always extricate itself from a liquidity trap as long as there are profitable investment opportunities overseas. 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 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 nontradable despite modern communications and transportation technology. And this large nontradable share may well mean that capital export, even at a zero interest rate, is not enough to escape 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 the price elasticities of imports and exports are also fairly small—both of which conditions are true in econometric models of large economies, if not in reality—even near perfect capital mobility may provide only limited extra scope for monetary expansion.

But should one believe this story? While the open economy IS-LM model may be a highly useful heuristic device for thinking about short-

6. Smithers (1998).

and medium-run macroeconomic issues, many economists doubt that it is really trustworthy, especially in considering such fundamental questions as the scope for international capital flows. And in any case, the thrust of this paper is to remove the stigma of the ad hoc nature of the liquidity trap concept. It may therefore be helpful to supplement this conventional view with a restatement in terms of a variant of my basic intertemporal model.

Consider a somewhat modified version of the basic model, in which the economy produces and consumes two goods, one tradable (T) and the other nontradable (N). Utility takes the form

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

In general, one would want to give the economy a transformation curve between N and T at any point in time. For simplicity, I assume that the transformation curve is right-angled; that is, the economy receives exogenous endowments of the two goods in each period. It can, however, borrow and lend on world markets at a given real interest rate r_T in terms of the tradable good, so 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 of either the nontraded good or a consumption basket that includes both traded and nontraded goods. This is most easily seen by considering the special case in which ρ is equal to one; that is, in which equation 10 takes the special form

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

In equation 11, utility becomes separable between tradables and nontradables. For each good, the relationship between consumption growth and the real interest rate must obey the rule $1 + r = D^{-1}(c_{t+1}/c_t)$. Whereas in the tradable sector relative consumption is determined by the exogenous real interest rate, however, in the nontraded sector (assuming full employment) it will be the other way around: because consumption of nontradables must equal production, the real interest rate in terms of nontraded goods will have to adjust to the path of production. 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 nontradables downwardly rigid, and consider the effects of a temporary monetary expansion—that is, 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 tradables, 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 tradables must rise, in order to allow for the subsequent fall. There must therefore be a nominal depreciation of the exchange rate.

The situation in nontradables will be exactly as in the economy as a whole in the 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 both for the exchange rate and for nontradable production, the zero constraint on the nominal interest rate can be binding. That is, even at a zero interest rate, the output increase and the nominal depreciation will have finite magnitudes—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. This is so because the separability of the utility function means that consumers in effect must make completely separate decisions on tradable and nontradable consumption over time; and since the real interest rate on tradables does not change, there is no reallocation between present and future consumption of those goods. This is obviously an artifact of the assumption that ρ is equal to one; I discuss the consequences of larger ρ below.

Financial Intermediation and Monetary Aggregates

Attempts to make sense of the origins and persistence of the Great Depression in the United States hinge crucially on how one interprets the radical divergence between the growth of monetary base and that

Figure 3. U.S. Monetary Trends, 1929–39

Index, 1929 = 100



Source: Temin (1976, p. 5)

of broader aggregates during the period. Figure 3 shows the familiar picture presented slightly differently from the standard representation (with both monetary base and M2 presented as indexes constructed so that 1929 equals 100). It shows that monetary base actually rose during the early years of slump and continued to rise steeply throughout the 1930s. By contrast, M2 fell by more than a third and did not surpass its 1929 level until 1939. These basic facts underlie two influential views of the Depression. One, suggested by Friedman and Schwartz, is that a broad aggregate like M2 is the proper measure of the money supply, that the Depression occurred 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.⁷ The other view, associated with Ben Bernanke and Russell Cooper and Dean Corbae, among others, is that the dramatic decline in the money mul-

7. See Friedman and Schwartz (1963).

tiplier 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.⁸ However, monetary explanations of the Great Depression have been criticized, most notably by Peter Temin, who suggests that the decline in monetary aggregates was a result rather than a cause of the slump, and perhaps could not have been prevented by the Federal Reserve.⁹

Since the Depression is the main historical example for liquidity trap economics, and since one quite often hears similar arguments made about contemporary Japan, it is 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 could 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 their classic paper, Douglas Diamond and Philip Dybvig introduce a demand for liquidity by making individuals uncertain about their own consumption needs; only after they have made commitments to illiquid investments do they discover whether they are “type one” consumers, who derive utility from consumption in period one but not period two, or “type two” consumers, who do the reverse.¹⁰ 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. Although Diamond and Dybvig are mainly concerned with showing how such a system could be vulnerable to self-fulfilling bank runs, one can also use their approach as a device for putting intermediaries and monetary aggregates into the basic model of this paper.

To do this, return to a one-good endowment economy, but now suppose that at the beginning of each period a three-step process takes place, as follows: (1) individuals trade currency for bonds in a capital market and 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.

8. See Bernanke (1994); Cooper and Corbae (1997).

9. See Temin (1976).

10. Diamond and Dybvig (1983).

The determination of the real interest rate is somewhat more complex in this setup, 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 money in their bank accounts to cover their cash needs if they do turn out to be type one 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, but 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 in both deposits and the price level.

But if the nominal interest rate is driven to zero, consumers and banks will become indifferent between holding monetary base and bonds—and consumers will also be indifferent between both of these 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 for either bonds or bank deposits in their portfolios; or the extra base could be absorbed by banks, which will simply hold excess reserves. Of these three possibilities, only the one in which 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, nor on output if prices are sticky.

Applying what one of my colleagues calls the principle of insignificant reason, one 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; and (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 liquidity trap, failure of broad monetary aggregates to expand is not a sign of insufficiently 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. However, this inability to expand broad money does not mean that the essential problem lies in the banking system; it is to be expected even if the banks are in perfectly fine shape.

The point is important and bears repeating: under liquidity trap conditions, the normal expectation is that an increase in high-powered money will have little effect on broad aggregates, and may 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, irrespective of the problems of the banks, *per se*.

Fiscal Policy

One can now consider possible policy responses to an economy in a liquidity trap. The classic Keynesian answer is 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 developed above 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 judgment: it is an accidental by-product 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 Japanese 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.¹¹ But in reality, fiscal

11. Suppose that one really believed that Japan was Ricardian equivalent, or nearly so. The first and most obvious implication is 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 second, less obvious, point is that under liquidity trap conditions the multiplier on

policy would surely have some impact. Two questions about that impact follow, one qualitative and one quantitative.

The qualitative question is whether a temporary fiscal stimulus can have permanent effects. If current income has very strong impacts on spending, 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 could jolt the economy out of that equilibrium into a region where conventional monetary policy worked again.

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, Christina Romer contends that most of the output gap created during 1929–33 had been eliminated before there was any significant fiscal stimulus.¹² She argues that the main explanation

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. This may be seen directly, by the fact that in the 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, 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, government spending will not generate any second-round increase in private spending.

A third point, which 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 that 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. Since in a liquidity trap money and bonds are perfect substitutes, it 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.

12. Romer (1992). Significant fiscal stimulus began in 1941 (before Pearl Harbor—a massive military buildup was already under way). One's assessment of whether the economy had largely recovered from the Great Depression prior to the onset of massive wartime spending partly depends on the choice of denominators. By 1940, real GDP

of that expansion 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 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.¹³

My point 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 the doldrums, however, a recovery strategy based on fiscal expansion would have to continue the stimulus over an extended period. Which raises the quantitative question of how much stimulus is needed, for how long—and whether the consequences in terms of government debt are acceptable.

Credibility and Monetary Policy

It may seem strange to have a subsection mentioning monetary policy, given that up to this point the paper has stressed the ineffectuality of such policy in a liquidity trap. However, as I noted at the beginning,

had risen 70 percent from its 1933 level, but it 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). While Gordon views the U.S. economy in 1939 as stuck, Romer emphasizes growth rates of more than 8 percent in 1939 and 1940.

13. For this calculation, Romer uses commercial paper rates, which did decline somewhat even in nominal terms. However, the spread between commercial paper and Treasury bills is presumably to some extent endogenous. T-bill rates averaged 0.515 percent in 1933—roughly the same as Japanese rates today. While they did fall to virtually 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, seen through the lens of the analysis in the present paper, Romer's evidence seems to suggest a somewhat different interpretation of events. One might think of her findings as showing that the real expansion of the economy—and the associated rise in prices—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 Romer emphasizes would have been ineffectual.

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 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 the central bank may do now, given the chance, it 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 analysis of the U.S. recovery over 1933–41 suggests that just such a bootstrap process was the main cause of the growth in output.

Proposals for “managed inflation,” first widely aired a few months ago, have since drawn a number of questions.¹⁴ One may as well go through those most frequently asked, and their answers.

Why inflation— isn't an end to deflation good enough? In terms of the analysis given above, price stability is not an option for an economy in a liquidity trap. The economy needs inflation, because it needs a negative real interest rate; the deflationary pressures actually being manifested represent 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 the expected future level.

Isn't inflation a bad thing? Again, in terms of my 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. One might also arrive at the recommendation of inflation by a quite different route: Friedman's famous theory of the optimum quantity of money.¹⁵ Although he says that the economy should deflate at the rate of time preference, the proper interpretation

14. Managed inflation gained widespread attention following a posting on the author's worldwide web site, “Japan's Trap,” in May 1998.

15. Friedman (1969).

of this logic is that the economy should deflate at the market-clearing real rate of interest. For a liquidity trap economy, where that market-clearing rate is negative, this means a negative rate of deflation—that is, inflation.

Won't expected inflation produce perverse incentives? In terms of the models, at least, a fall in the real interest rate achieved through expected inflation is identical in its effects to one produced through 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.

Won't an inflationary policy lead to a plunge in the exchange rate and become a beggar thy neighbor policy at the rest of the world's expense? Because expected inflation plays the same role in a liquidity trap economy as do interest rate reductions 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 Robert Mundell and Marcus Fleming, and also in large-scale econometric models, monetary expansion unambiguously leads to currency depreciation.¹⁶ 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 to suggest that these effects very nearly cancel each other out. Table 1 presents estimates from the comprehensive, if somewhat elderly, comparison of eleven models by Jeffrey Frankel.¹⁷ For each model, it shows the second year effects on the exchange rate and the current account of a monetary expansion sufficient to raise real United States GNP by 1 percent. The exchange rate impacts are substantial, the current account impacts negligible. To the extent that these estimates are correct, they suggest that in a large economy with fairly small

16. Mundell (1963); Fleming (1962).

17. Frankel (1988).

Table 1. Second Year Effects on the U.S. Exchange Rate and Current Account after a Monetary Expansion to Raise Real GNP by 1 Percent

Percent		
<i>Model^a</i>	<i>Exchange rate</i>	<i>Current account^b</i>
DRI	-8.1	-0.02
EEC	-4.0	-0.07
EPA	-5.3	-0.03
LINK	-2.3	-0.01
LIVERPOOL	-39.0	-3.1
MCM	-4.0	-0.05
MINIMOD	-5.7	-0.07
MSG	-6.7	-0.21
OECD	-1.6	-0.13
VAR	-7.6	-0.04
WHARTON	-1.4	-0.17
<i>Summary statistic</i>		
Median	-5.3	-0.03

Source: Frankel (1988).

a. Models are fully identified by Frankel.

b. As a percentage of GNP.

trade shares, expected inflation will produce a significant currency depreciation but have small impact on the current account.¹⁸

I have been trying to get beyond the IS-LM model, however. How does the result look in an intertemporal open economy model? Assume the utility function given in equation 10, exogenous output of traded goods, and sticky prices or 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.

18. In the case of Japan, many people want the economy 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). And only a fraction of this swing would come vis-à-vis the troubled emerging economies of Asia. The only way to get a much larger locomotive effect would be for Japan to pursue a Reagan-style expansion, in which the exchange rate appreciates substantially. However, given that Japan is having great difficulty achieving any type of recovery, advocating a currency appreciation seems rather strange.

Table 2. Beggar Thy Neighbor Coefficients

ρ	τ		
	0.2	0.3	0.4
2	0.167	0.231	0.286
3	0.286	0.375	0.444
4	0.375	0.474	0.545

Source: Author's calculations as described in text.

The question is what impact this expansion has on the current account, which in this framework amounts to asking what happens to consumption of the traded good.

One can take a shortcut here, if one imagines that the expansion is "brief," in the sense that one can ignore the effect of the current account on the future investment income of the country. 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

$$(12) \quad B = - \frac{1 - \rho}{1 - \rho - (1/\tau)},$$

where ρ is relative risk aversion and τ is the traded share of consumption (and hence value added in the economy). One sees immediately that in the special case of ρ equal to 1, a monetary expansion has no effect on the expanding country's current account, which is roughly what the econometric exercises in table 1 indicate. If relative risk aversion is higher than 1, there will to some extent be expansion by means of inflation, as a result of a widened current account surplus, but the extent will depend inversely on how open the economy is, 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 judgmentally assumes that the tradable 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 5 percent relative to baseline would require an expansion of the current

Table 3. Depreciation Coefficients

ρ	τ		
	0.2	0.3	0.4
2	1.67	1.54	1.43
3	2.14	1.88	1.67
4	2.5	2.11	1.82

Source. Author's calculations as described in text.

account surplus of something like 1 percent of GDP. This is far short of the huge surpluses envisaged by Smithers and others.¹⁹

One 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 1 percent increase in real GDP achieved through inflationary expectations. It can be shown that this depreciation is

$$(13) \quad 1 + \frac{1 - \tau}{\tau} \cdot \frac{1 - \rho}{1 - \rho - \frac{1}{\tau}}$$

Table 3 shows depreciation coefficients for a range of values of ρ and τ . These numbers look generally small, compared with the model simulations in table 1. I discuss why this might be so in the second part of the paper.

Summary

I have 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 the territory: once an economy really is in a liquidity trap, much of the conventional wisdom of macroeconomics ceases to apply—indeed, applying 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.

19. Smithers (1998).

First, one must be careful about making inferences from divergences between the growth of monetary base and of broad monetary aggregates. The failure of aggregates to grow need not indicate dereliction on the part of the central bank; in a liquidity trap economy the central bank in principle cannot move broad monetary aggregates. Likewise, the observation that although the central bank has slashed interest rates and pumped up monetary base, the broader money supply has not grown, does not necessarily imply that the fault lies in the banking system; it is just what one would expect in a liquidity trap economy.

Second, whatever the specifics of the situation, a liquidity trap is always the product of a credibility problem: the public believes 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

Table 4 presents some standard summary statistics on Japan's economic performance since 1981. It makes the familiar point that following rapid growth to 1991, the economy 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 simply viewed as the correction of an unsustainable boom. Inflationary pressures had clearly eased by 1992, though, and the low growth rate thereafter is a better indicator of the economy's true shortfall. It is clear that Japan will have a significant decline in real GDP for 1998; and the unemployment rate has already risen above 4 percent.

There are two striking features of these dreary numbers. The first is the extent of the slowdown. In the period 1981–92 Japan grew at an average rate of 3.7 percent. It ended the period with the same unemployment rate as it had started with, and with a lower inflation rate; in short, potential as well as actual output would seem to have risen at about 3.7 percent annually over the period. If one had projected that growth rate forward, one would have overpredicted 1998 output by about 14 percent.

Table 4. Economic Performance in Japan, 1981–97

Percent

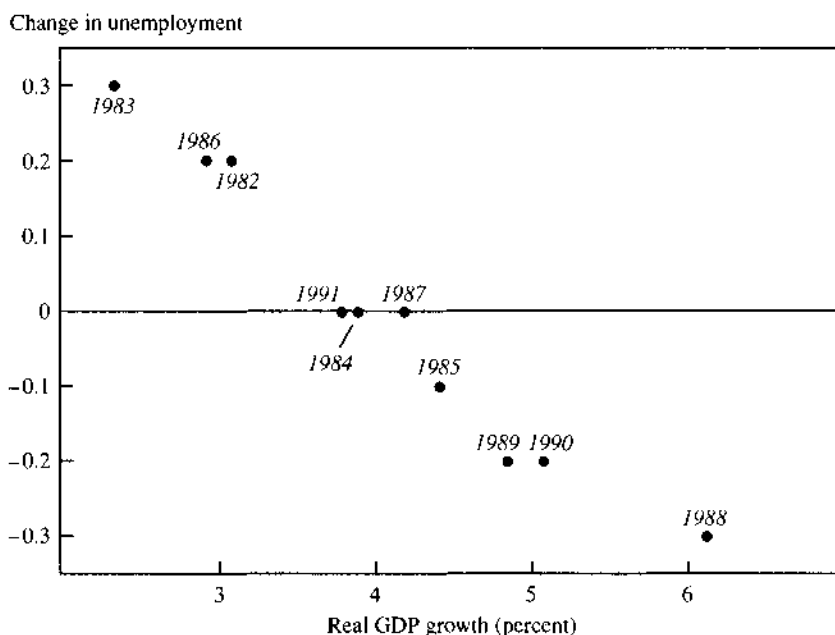
<i>Year</i>	<i>Real GDP growth</i>	<i>Inflation</i>	<i>Unemployment rate</i>	<i>Money market interest rate</i>
1981	3.2	4.1	2.2	...
1982	3.1	1.8	2.4	...
1983	2.3	1.8	2.7	...
1984	3.9	2.6	2.7	6.5
1985	4.4	2.1	2.6	6.6
1986	2.9	1.8	2.8	5.1
1987	4.2	0.1	2.8	4.2
1988	6.2	0.7	2.5	4.5
1989	4.8	2.0	2.3	5.4
1990	5.1	2.3	2.1	7.7
1991	3.8	2.7	2.1	7.2
1992	1.0	1.7	2.2	4.3
1981–92 average	3.7	2.0	2.5	5.7
1993	0.3	0.6	2.5	2.9
1994	0.6	0.2	2.9	2.3
1995	1.5	-0.6	3.1	1.2
1996	3.9	-0.5	3.4	0.6
1997	0.9	0.6	3.4	0.6
1993–97 average	1.4	0.1	2.4	1.5

Source: *International Financial Statistics*, 1998.

The second 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—at the time of writing, there are still 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? In principle, the great bulk of the slowdown might represent 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. It is therefore important to estimate the gap between actual and potential output.

For 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

Figure 4. Okun's Law for Japan, 1982–91


Source: *International Financial Statistics*, 1998.

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 1982–91 period, as shown in figure 4. The slope of the apparent relationship is about three times as steep as for the United States: it apparently took about 6 percentage points of excess growth to reduce the unemployment rate by 1 percentage point. If one were to take the average 2.5 percent unemployment rate in the period before the slump 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 in that year—and with potential output presumably still growing while output slumps, by the end of 1998 the gap could be as great as 10 percent.

Most published estimates of Japan's output gap are far smaller. 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

Figure 5. U.S. Gross Domestic Product, 1919–39

ln GDP




Source: Author's calculations based on data from U.S. Bureau of Economic Analysis (1973).

changes in the growth rate of potential output.²⁰ 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. First, it imposes the assumption that average deviations from potential are zero over the whole period, so that when the economy slumps, the filter automatically reevaluates earlier periods as times of above-potential output, reducing the estimated shortfall. Second, any sustained drop in output gets built into the estimated potential growth rate. As a result, it systematically understates the actual shortfall from potential. A stark, if somewhat unfair, way to make this point is to apply the Hodrick-Prescott filter to the United States in the interwar period, as shown in figure 5. For this figure, the smoothing parameter λ is set at 25; but a

20. See Giorno and others (1995) for a description.

Figure 6. Output Gap for Japan, 1982–97

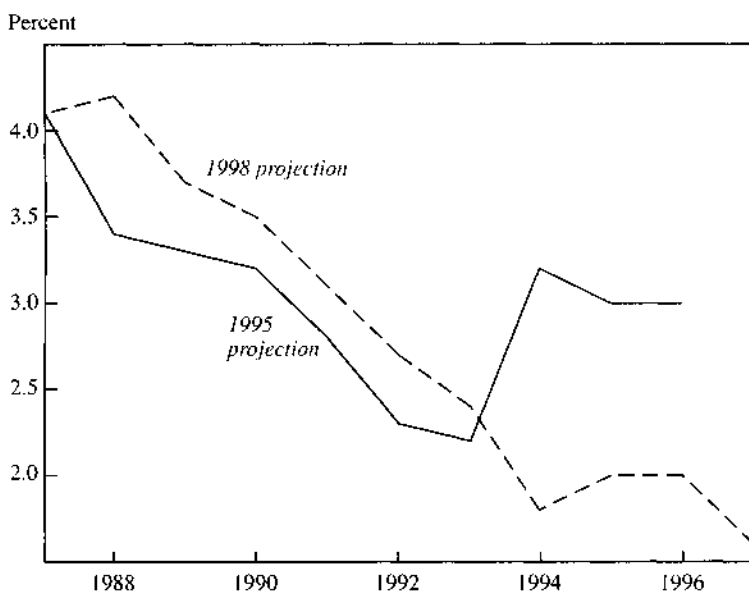


Source: *OECD Economic Outlook*, June 1998

wide range of values of λ yields the conclusion that output was in excess of potential by 1935.

The Organisation for Economic Co-operation and Development (OECD) has adopted a more complex technique to assess Japan's output gap; its most recent estimates are shown in figure 6.²¹ Nonetheless, the estimated output gap in 1997 is remarkably small: -1.2 percent. This seems to be due to the fact that although the OECD does not engage in simple Hodrick-Prescott filtering, it updates estimates of normal worker hours and worker productivity in such a way that possibly cyclical components get reinterpreted as structural trends. Figure 7 provides a nice illustration of this process at work, by contrasting the estimates of potential growth in the 1995 study by Giorno and others that introduced the OECD's current method with the potential growth estimates that appear in the most recent *OECD Economic Outlook*. As recently as

21. The OECD methodology is described in Giorno and others (1995).

Figure 7. OECD Estimates of Japanese Potential Growth, 1987-97

Source: Giorno and others (1995), *OECD Economic Outlook*, June 1998.

three years ago, the OECD estimated Japan's potential growth at 3 percent; now it has marked it down to 1.6 percent. Applying the earlier 3 percent potential growth to the period since 1994, the 1997 output shortfall rises to 4.6 percent; not too far short of the estimate suggested by the Okun's Law calculation for Japan.

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

Saving and Investment

A liquidity trap occurs when desired saving exceeds desired investment at full employment, even at a zero short-term interest rate. As

Table 5. Private Consumption as a Share of GDP, Japan and the United States, 1991-97

Percent

<i>Year</i>	<i>Japan</i>	<i>United States</i>
1991	57.1	67.1
1992	57.8	67.6
1993	58.6	68.0
1994	59.7	67.8
1995	60.1	68.2
1996	59.9	68.2
1997	60.6	67.9

Source: *International Financial Statistics*, 1998.

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 interpretation of Japan's problem, and to some extent the policy implications, do depend on how one views the apparent excess saving.

Table 5 shows ratios of consumption to GDP for Japan and the United States since 1991. Two familiar observations stand out. One 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 has nothing to add to it. The other observation is that Japan's consumption ratio has not declined in the 1990s; if anything, it has risen slightly. This suggests that the shift into liquidity trap territory might reflect declining investment demand rather than rising saving 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 U.S. 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 the U.S. find itself in liquidity trap territory? Recall that the emergence of 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. Another approach is to ask what the impact of such a monetary-fiscal switch would be in a variety of standard econometric models, such as those considered in table 1. Table 6 reports estimates of the impact of

Table 6. Effect on U.S. Short-Term Interest Rates of a Fiscal-Monetary Switch Equivalent to a 1 Percent Decline in the Consumption Ratio

Percent	
<i>Model^a</i>	
DRI	-4.3
EEC	-4.4
EPA	-5.3
LINK	-1.9
LIVERPOOL	-2.2
MCM	-4.3
MINIMOD	-2.9
MSG	-3.3
OECD	-2.3
TAYLOR	-0.7
VAR	-0.4
WHARTON	-5.3
<i>Summary statistic</i>	
Median	-3.1

Source: Frankel (1988, pp. 21, 23)

a. Models are those used in table 1

a monetary-fiscal switch equivalent to a 1 percent decline in the consumption ratio; both the mean and the median effects are a 3 percentage point decline in the short-term interest rate.

It is easy to find reasons why such exercises might overstate the case; structural models probably tend to understate the spending impact of a sustained reduction in the interest rate. But even this crude comparison makes it substantially less surprising than one might 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 the 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 the 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. It is certain, however, that Japan's long-run growth must slow, even at full employment, because of demographics. Through the 1980s, Japanese employment expanded at 1.2 percent annually.²² However, the working-age population has now peaked: it will decline at 0.7 percent annually over the next thirty years, and—if demographers' projections about fertility are correct—at a remarkable 1.0 percent for the twenty-five years thereafter.²³ As suggested by the discussion of investment and Tobin's q in the first part of this paper, such prospective demographic decline should, other things equal, depress expectations of future q , and hence also depress current investment.

The looming shortage of working-age Japanese people 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. One reason why this prospect did not start to affect long-term investment projects earlier is the "bubble economy" of the late 1980s. Businesses may have believed that total factor productivity would grow rapidly enough to make up for a declining work force. However, the bubble economy may also have masked the underlying decline in investment opportunities, and hence delayed the day of reckoning. Moreover, that bubble economy left a legacy of large debts and troubled bank balance sheets, which are widely regarded as the main culprits of Japan's current plight.

Banking Problems

Japan clearly faces a huge problem of bad bank loans; the current conventional wisdom places their value at a trillion dollars. These bad loans are in part a legacy of the burst of the asset bubble of the 1980s, reinforced by the consequences of the subsequent slow growth. Clearly Japan will need to engage in a cleanup operation dwarfing that of the U.S. thrift crisis, especially as measured against Japan's smaller economy. Inevitably, also, the form and funding of that cleanup will be central political preoccupations. 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 inter-

22. *OECD Economic Outlook*, December 1997, p. A23.

23. Organisation for Economic Co-operation and Development (1997, p. 113).

Table 7. Japanese Financial Data, 1994–97

Index, 1994 = 100

<i>Year</i>	<i>Monetary base</i>	<i>M2 plus certificates of deposit</i>	<i>Bank credit</i>
1994	100.0	100.0	100.0
1995	107.8	103.3	100.8
1996	117.0	106.5	100.6
1997	125.6	110.6	100.9

Source: International Financial Statistics, 1998

mediation has clearly played a crucial role in many if not most historical financial crises, including the current crisis in the emerging economies of Asia. Also, to many economists it seems a priori obvious that if conventional monetary policy has become ineffective, the reason must be that the troubles of the banks have blocked the usual channels of central bank influence.

A casual look at the data does seem to support the view that the Japanese problem with monetary policy lies in the banks. Table 7 shows developments in high-powered money, broad money, and bank credit since the end of 1994. It is evident that 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 an economy in a liquidity trap, this sort of disconnect between monetary base, aggregates, and bank credit is to be expected even if the banks are financially healthy. It is not evidence that the banks' troubles 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 an 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, and so forth, 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 one expect a bank

of questionable solvency to behave? Would it restrict credit? The textbook answer 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.²⁴ In effect, the game is "heads I win, tails the taxpayer loses." Indeed, one could argue that since the bubble burst, Japan's financial institutions 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.

This is not merely abstract speculation. Japan has already, in regard to the *jūsen* (nonbank subsidiaries of financial institutions, specializing in housing loans), gone through a miniature version of the systemwide bank cleanup that it must now undertake. According to Thomas Cargill, Michael Hutchinson, and Takatoshi Ito, *jūsen* 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."²⁵ Because these agricultural cooperatives had strong political influence, they were able to take large risks while counting on implicit government guarantees. The result was behavior strongly reminiscent of that of the U.S. thrifts. Indeed, Cargill, Hutchinson, and Ito offer a striking example of lending driven by moral hazard 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. Nonetheless, in the two years before the institutions were closed, both their deposits and their loans expanded rapidly.

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 that 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 widely accepted.

24. This line of argument now plays a major role in discussions of the troubles in the emerging economies of Asia; see McKinnon and Pill (1997); Krugman (1998); Corsetti, Pesenti, and Roubini (1998).

25. Cargill, Hutchinson, and Ito (1997, p. 121).

A review of press reports also makes the reasons for the emergence of credit constraints in late 1997 quite clear. The immediate forcing event was the announcement, in October 1997, of new capital adequacy standards, to be effective from 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, one can argue that since late 1997, the prospect that the government would eventually seize 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 those successful in this endeavor is, loosely speaking, that they will 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 economists in the United States. A mild form of the same ailment appeared in 1990–92, when the size of the savings and loans bailout had become apparent, and it was widely said 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 when it began to look likely that the government would do something about the situation.

But 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 U.S. savings and loans case and Japanese experience with credit cooperatives so graphically demonstrate) delay only multiplies the losses. That it might reduce aggregate demand as a side effect is of little relevance. 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 is that because it is in a liquidity trap, the normal disconnect between microeconomic and macroeconomic policy no longer applies.

Policy Options and Their Consequences

Given all that I have said, it is useful to review Japan's policy options and ask how well they would work. Current discussion focuses on three basic alternatives, which are not 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 period since 1992. 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 questions about fiscal expansion as a remedy for Japan, one strictly economic, one political. The economic issue is whether an adequate expansion is possible without creating an unacceptable impact on the government's long-term fiscal position. Much discussion of fiscal stimulus in Japan seems to be predicated on some form of pump-priming (or jump-starting): the idea that a brief period of stimulus will jolt the economy back into a favorable equilibrium. However, there is no good evidence for such a multiple-equilibria 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.²⁶ An alternative view is that Japan has a long-term deficiency of demand due to low rates of time preference combined with negative population growth—and also an output gap of 7 percent or more—so that the size and duration of the deficits implied would be very large.²⁷ 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 the United States during the New Deal—appears to have great difficulty in working up political nerve for a fiscal package anywhere close to that required to close the output gap. Exactly why this is so is an interesting question, but beyond this paper's scope.

26. Romer (1992).

27. A useful indication of the seriousness of the situation in Japan is that ten-year government bond rates are now less than 0.7 percent, suggesting that investors expect the country to be in or near a liquidity trap for at least a decade.

This surely does not, however, mean that fiscal policy should be ignored as part of the policy mix. 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. (Not even I would trust myself enough to go for a purely Krugman solution.) But 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 shown above, the financial problems of the banks have until recently biased them toward lending too much rather than too little.

Ironically, 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. 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.

Moreover, a radical, forceful bank cleanup—which basically settles the issue and leaves the remaining banks reasonably sure that they will not be taken over—would in principle 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.

In short, a financial cleanup is vital on microeconomic grounds; and given that it must be done, on macroeconomic grounds it is best 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 that Japan 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 is laid out in the first part of this paper. In Japan's case, there would be three main issues: implementation, the appropriate target, and the likely effects.

How can a country in a liquidity trap—that is, where increases in the money supply seem to have no effect—engineer inflation? As I have shown, the problem is essentially one of credibility. If the central bank can credibly commit 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 by means of 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 need never fall; all that is required is a promise not to raise them when the economy expands and prices begin to rise.

How in fact to create these expectations is, in a sense, outside the usual boundaries of economics. However, one obvious suggestion is that Japan deal with its inverted credibility problem through legislation giving the Bank of Japan an inverted version of the price stability targets now in force in a number of countries: it would be enjoined to achieve an inflation rate of not less than x percent over y years. (If this does not work, appendix C discusses several ways in which the necessary inflation expectations might nonetheless be generated.)

And this raises the question of the appropriate inflation target. A key insight 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 one 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 as large as Japan's.

One 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, economists have no real idea of 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. Therefore Japan probably requires a sustained period—

Table 8. Reduction in U.S. Long-Term Interest Rates That Would Expand Real GNP by 1 Percent

Percent	
<i>Model^a</i>	
DRI	-0.82
EEC	-1.80
LINK	-0.64
MCM	-0.68
MINIMOD	-0.30
OECD	-0.94
TAYLOR	-0.30
WHARTON	-2.70
<i>Summary statistic</i>	
Median	-0.75

Source: Bryant and others (1988).
 a. Models are those used in table 1

at least a decade—of inflation, to reduce the real long-term rate sufficiently to close the output gap.

At this point, matters become difficult. The size of Japan's output gap is 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 U.S. GDP by 1 percent, using the various standard econometric models introduced in table 1.

Given the uncertainties, any number is a matter of multiplicative guesswork. 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—clarification of which banks will be taken over and which will not, should reduce that gap by several percentage points. Therefore managed inflation would need to close a remaining gap of, say, 4 to 5 percentage points. Looking at the median estimate in table 8, this would require an inflation target of 3 to 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 fifteen years?

This target should not really be taken seriously. Rather, it should serve mainly to stimulate serious research. And 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.

What side consequences might one expect from such a solution to Japan's slump? In particular, what would happen to the current account and the value of the yen? Recall that 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 to 25 percent—a number that is probably less uncertain than the required inflation rate discussed above, 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 the 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 one needs to revive and modernize the theory of the liquidity trap, a concept that once played a major role in macroeconomics, but has virtually disappeared from economic discourse in the past twenty 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 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 available for 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, one should wonder whether it could happen elsewhere. Germany and France currently have short-term interest rates of only 3.5 percent, and Europe faces Japanese-style demographics; could a liquidity trap happen to the European Monetary Union? Economists 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.

APPENDIX A

Financial Intermediation and Monetary Aggregates in a Liquidity Trap

IN THE TEXT, I sketch out how one might think about the role of financial intermediaries and the behavior of monetary aggregates in a liquidity trap. This appendix describes that “cash in advance meets Diamond-Dybvig” approach more fully.

Consider, for simplicity, a full-employment endowment economy that lasts for only two periods, with each individual receiving an endowment y_1 in period one, y_2 in period two. 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

$$(14) \quad U = HU_1(c_1) + (1 - H)U_2(c_2),$$

where H takes on the value 1 with probability π , value 0 with probability $1 - \pi$. So in the population there will be fractions π of first period consumers, $1 - \pi$ of second period consumers.

One wants to make this a cash in advance economy. But finite horizons 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, 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 below.

Within each period, consumers must pay in cash before they receive income from selling their own endowments. 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, 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.

This is where financial intermediaries come in. 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 equal to 1. (Bank runs are left on one side for this paper!) Deposits earn competitive interest if not withdrawn.

Thus the sequence of events looks like this:

- Consumers come into existence and receive the money supply M^* .
- A capital market is held; consumers deposit money in banks, and open market operations may increase or decrease the monetary base.
- Consumers learn their type.
- They withdraw their funds if necessary.
- Consumers receive income from sale of their endowment, receive bonds and deposits, and pay or receive whatever tax or transfer is needed.
- Consumers purchase second period consumption.
- 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 one consumer will get to purchase units of the good in period one; but since a fraction π of consumers is type one,

$$(15) \quad \pi[y_1 + y_2/(1 + r)] = y_1,$$

implying that the real interest rate is

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

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 one; 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 how 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 in pretty much the same way 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 between cash and bonds. At this point any further open market bond purchase by the government could be absorbed in any combination of 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 on either real variables or 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 introduce 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 or 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 question reduces to the question of what happens to traded-good consumption.

One 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, one initially has that

$$(17) \quad \frac{c_T}{c_N} = \frac{\tau}{1 - \tau}.$$

One 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, one knows that the levels 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 even in the current period the marginal utility of traded goods will remain unchanged, as

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

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 to either 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, note that

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

and that

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

Finally,

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

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 that a pure bootstrapping policy—in which the announcement of an inflation target generates the expansion that eventually creates the inflation—is infeasible. Then Japan should apply some temporary policy that moves the economy to a position where monetary policy does have traction and use that traction to generate sustained inflation.

In this case, the temporary fiscal jolt 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 is 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 encourage more or less the same kind of spending as the immediate creation of inflation expectations.

Comments and Discussion

Kathryn M. Dominguez: This paper analyzes the efficacy of various macroeconomic stabilization policies for a recessionary economy with a nominal short-term interest rate close to zero. This economic situation was first discussed by Keynes, and subsequently dubbed the liquidity trap, in the context of the Great Depression. Paul Krugman now re-examines the same issues in a new context, that of present day Japan. What makes the paper novel and provocative is that Krugman's policy prescription differs substantially from the Keynesian one—and, for that matter, from any of those usually given by well-trained economists. He recommends that policymakers facing a liquidity trap engage in sustained monetary expansion, that is, “credibly promise to be irresponsible, to seek a higher future price level.” My comments fall into two categories. First, I assess Krugman's (purposely) provocative policy advice; and second, I examine some issues arising from the models that he presents.

A liquidity trap is defined as a situation in which, because nominal interest rates are at or near zero, investors are indifferent between holding bonds and money, and as a consequence, monetary policy is ineffective at boosting demand. Under normal circumstances, monetary policy will be neutral, or ineffective, in the long run. But if one allows for some price stickiness, in the short run a monetary expansion may increase output and will, in any case, lead to an equiproportionate increase in prices. However, when interest rates are already “too low,” so that the economy is in a liquidity trap, monetary policy is powerless to influence output or prices (in the IS-LM paradigm, the LM curve is

horizontal). It is for this reason that the traditional way out of a liquidity trap has been to rely on fiscal policy (an IS shift).

Krugman largely dismisses fiscal policy as a solution to the liquidity trap on the basis of a Ricardian equivalence argument. If the public understands that current government expenditures or tax cuts will eventually need to be reversed, it will undo the government actions—save rather than spend—in anticipation of the future fiscal policy reversal. So, if the public is to some degree Ricardian, the argument goes, fiscal policy will do little to lift an economy out of a liquidity trap. (I have more to say on this below.) Krugman then returns to monetary policy, which Keynes and Hicks suggest will be ineffective, and suggests that if central bankers can credibly commit to being inflationary, it may work after all.

A key insight highlighted in this paper is the role of expectations in a liquidity trap world. If one starts with the assumption that the public believes that the central bank is committed to price stability, and if as a consequence of expansionary monetary policy prices rise above their current level, the central bank will be expected to contract in order to stabilize prices. Fundamentally it is this belief, that central banks are credible in the usual sense of the word (that is, they are committed to price stability), that renders monetary policy ineffective. A monetary expansion in this context will always be considered temporary. And it is for this reason that Krugman suggests that the central bank commit to letting prices rise in the future: if the public believes that the central bank will not reverse an expansion, the economy *can* escape from the trap by means of expansionary monetary policy.

Before discussing some of the subtleties of Krugman's arguments, it is worth asking whether it is possible or wise for central banks to commit to being irresponsible. There exists a vast literature examining the reverse proposition of whether a central bank can commit to being responsible. And my sense of this literature is that in theory the answer is yes, and in practice the answer is often no. Game theory provides all kinds of "commitment mechanisms" for central bankers; yet one rarely sees such mechanisms put in place, and even more rarely enforced. If central banks have had a hard time convincing the public that they will be responsible and maintain price stability, will it be more or less difficult to convince us that they will steadfastly pursue an inflationary policy? I would venture that such a promise would be difficult to sell,

just as any promise to hold to one policy is unconvincing, because of time inconsistency.

For the sake of argument, assume that a central bank in a liquidity trap goes ahead with Krugman's advice and announces a sustained monetary expansion. And assume that the public is convinced that the central bank will not reverse its program. What is the likely outcome? In Krugman's scenario, the expansion will steadily raise prices, lower real interest rates, and depreciate the exchange rate in an orderly fashion. These changes will, in turn, provide the public with incentives to consume and invest more now, and thereby jump start the economy.

But there remain some pitfalls. If the public is convinced that the central bank will not contract, what is to stop people from expecting accelerating inflation? What if, instead of moderate inflation, expectations drive prices too high? Can a central bank commit to being moderately irresponsible? And even if all goes well with moderate inflation, will consumers necessarily spend more? If real interest rates turn negative, consumers will have an incentive to invest abroad. This, in turn, would rob domestic banks of deposits and potentially limit bank lending. Higher inflation will also lead to higher long-term interest rates for corporate borrowers, again dampening the demand stimulus. Finally, there are potential repercussions from a currency depreciation. In the current postcrisis environment in Asia, it is difficult to predict the impact of a further depreciation of the yen. A fall in the value of the yen might well set off further declines in the currencies of other Asian countries, increasing the burden of their foreign debts and leading to even greater economic turmoil in the region.

Although attention is likely to focus on Krugman's unorthodox policy prescription, the bulk of the paper is not about policy, but rather, a reexamination of the causes and consequences of liquidity traps. In a series of simple models, Krugman considers the roles of investment, capital mobility, and financial intermediation in this context. The purpose of these models is to show that liquidity traps can be created without resorting to the ad hoc features of Hicks's original IS-LM analysis, and that certain features of modern economies—features that Keynes and Hicks did not consider—do not much change the original analysis. I expect that Krugman's main conclusion from this exercise, that liquidity traps can exist in modern economies, is correct. I do,

however, take issue with some of the simplifying assumptions that he makes to illustrate his point.

The basic model presents an endowment economy with cash in advance constraints. Consumers can hold money or one-period bonds and the nominal value of consumption in each period cannot exceed money holdings. When interest rates are positive, the cash in advance constraint will bind, and prices will be proportional to the money supply. Starting with this simple model, the first interesting result is that even in a full-employment flexible price economy, a version of the liquidity trap will arise when nominal interest rates fall to zero. In this case, people become indifferent between holding money and bonds, the cash in advance constraint is no longer binding, and prices and the interest rate no longer reflect current changes in the supply of money. Money becomes irrelevant. This is a rather strange version of the trap, however, because it is, by construction, a trap with no negative consequences (in terms of output). The consequences of a trap get a bit more interesting when there exists some nominal rigidity—in this case, a predetermined price level in the first period. But in either case, one issue that becomes apparent with these models is that it is hard to think seriously about stabilization policy in the context of an endowment economy. Further, if the models are to apply to Japan, one must surely wonder how the enormous stock of past savings might change the model dynamics. In Japan's case, it is hard to think of consumption in each period being constrained by current money holdings or income levels.

The next issue that Krugman addresses is the role of productive investment. Are the negative real interest rates that must exist when an economy is in a liquidity trap still possible when productive investment is introduced? The insight provided by the model is that even with positive marginal product of capital, if Tobin's q is expected to fall, the rate of return on investment can be negative. In a simple example using land and a declining population, Krugman shows that the expected return on land might be negative. But in this setup Tobin's q is essentially exogenously determined; in a less stripped-down model it may not be so easy to produce negative returns.

Virtually free capital mobility is one of the features that distinguishes modern economies from the world that Keynes and Hicks knew. Can negative real interest rates coexist with capital mobility? Or in other

words, if capital is free to move across national boundaries, under what circumstances (if any) can the domestic real interest rate differ from the world rate? In order to analyze this question, Krugman modifies his model so that the economy produces two goods, one tradable and the other nontradable. Utility is assumed to be separable between tradable and nontradable goods. This assumption also allows consumption decisions and real interest rates in the two sectors to be determined independently. It is then possible to concoct a scenario—when production of the nontradable good falls exogenously over time—in which the real interest rate as measured by nontradable goods is negative even as the real interest rate as measured by tradable goods is positive (and equal to the world interest rate). Noting that the share of traded goods in consumption is relatively modest, Krugman argues that the weighted real interest rate might therefore be less than zero.

It is worth reflecting on the relevance of an interest rate constructed in such a way, and to do so it is useful to return to the issue of international capital mobility. International capital mobility has two implications: first, that consumption is constrained by the present value of income rather than by national output; and second, that investment funds are allocated internationally so as to equate real marginal products of capital at the world interest rate. Krugman's disarmingly simple model considers only the first of these, the implications of capital mobility for consumers' Euler equations, and disregards the productive aspects of capital mobility (since the model is essentially an endowment economy). Were productive capital movements permitted, the domestic nontradable sector would be a prime candidate for investment, because the prices of its output are expected to rise. Indeed, if it were possible merely to store nontradable goods, it would make sense to do so in this model, given anticipations of higher prices in the future. One cannot store haircuts; but as I look around my block in Ann Arbor, I see that one can certainly store houses, and housing, with one-third of the entire capital stock, represents a substantial part of the nontradable sector. More generally, the scenario in which real interest rates are negative requires that prices are expected to rise faster than nominal interest rates, that the good cannot be stored, and that it is not possible to invest in this sector so that competitive forces can drive down future prices. I do not know to how much of the economy such a scenario applies; it may not be very much.

If the model's analysis of interest rates that differ from world rates is not meant to apply to investment, it must therefore apply to consumption. The difficulty with this application is that consumers face the world interest rate (r) in selecting consumption paths of tradable goods. Even if the nontradable sector is large, there is always the option of consuming one fewer tradable good today in return for $(1 + r)$ tradable goods next year. So consumers do, in fact, face the world interest rate in making saving decisions. Indeed, in Krugman's model, saving and dissaving occur through exchanges of tradable goods; the nontradable sector is entirely irrelevant to intertemporal decisionmaking. Hence an interest rate that is constructed by weighting interest rates on tradable and nontradable goods by their consumption shares is not relevant to consumption and saving; the relevant rate is the world interest rate. The significance of the nontradable sector in this model is rather that, since tradable goods represent a modest fraction of total consumption, a given change in interest rates may have a small effect on total consumption. But this is very different from saying that the relevant real interest rate is negative.

The last issue that Krugman addresses is the role of financial intermediation. He makes the case that when nominal interest rates are close to zero, so that both consumers and banks are indifferent between holding money and bonds, the influence of an expansion of the monetary base on broader monetary aggregates will depend on both consumer and bank behavior. Banks therefore should not be blamed when central banks are unable to influence broad money. In a liquidity trap, neither banks (regardless of their condition) nor consumers have the proper incentives. This point is important, and it undercuts many recent analyses of the causes of Japan's current economic situation.

Having made this point, and before presenting his novel policy prescription, Krugman considers the traditional cure for a liquidity trap: a fiscal stimulus. He is not optimistic that fiscal policy can successfully stimulate the Japanese economy, due to the mysterious properties of Ricardian equivalence. The view that economies will completely offset any fiscal actions of the government is not traditionally associated with the Massachusetts Institute of Technology (MIT)—nor, for that matter, associated with Japan, where the traditional budgetary austerity is commonly thought to leave ample scope for stimulatory fiscal policy if desired. Certainly that is the position of the United States in recent

years, as advanced in part by Lawrence Summers, a onetime MIT professor. The theoretical limitations of models of Ricardian equivalence are well-enough known that it is hardly worth cataloguing them here. Do they apply to modern Japan? No one knows for sure. But it stands to reason that in an economy in which consumers are reluctant to spend and the balance sheets of banks and industrial conglomerates make it difficult for them to invest in positive net present value projects, a fiscal stimulus provided by the government might have more effect than under other circumstances.

Turning to the current situation in Japan, Krugman makes a convincing case that the economy is in a liquidity trap, with a widening output gap, near zero nominal interest rates, deflation, a low ratio of consumption to GDP, and high growth of outside money without corresponding increases in broader money aggregates. And he presents some intriguing back of the envelope calculations as to just how much inflation the Bank of Japan would need to create in order to “untrap” the economy. Krugman’s rather brave bottom-line policy pronouncement is that Japan ought inflate at about 4 percent for fifteen years. On the one hand, I admire his willingness to stick out his neck and actually suggest a specific policy action. On the other hand, I hope that the Bank of Japan does not decide to follow his prescription. Perhaps more important, I hope that the market has confidence that the Bank of Japan will not do so.

Whatever one’s opinion of the provocative policy advice, this paper makes an important contribution. Krugman has single-handedly revitalized the liquidity trap and provided the economics profession with many interesting and urgent research questions. My expectation is that his bold prescription will provoke economists and policymakers to think creatively about alternative solutions to liquidity traps in general, and Japan’s trap in particular.

Kenneth Rogoff: This is a truly inspired paper on Japan’s ongoing “Great Recession,” although I have to keep pinching myself to ask if its main thesis can really be true. Is the equilibrium (full-employment) medium-term real interest rate for Japan actually negative, so that unless the Bank of Japan (BOJ) resigns itself to sustained inflation, the zero bound on nominal interest rates will present serious problems? Has the BOJ so thoroughly convinced the public of its anti-inflation credi-

bility that it has lost the power to rekindle inflation now that Japan needs it?

The idea that the non-negativity constraint on nominal interest rates may pose problems in a world of low inflation has been receiving a growing amount of attention. Lawrence Summers has warned that there may be times when optimal stabilization policy calls for temporarily inducing negative nominal interest rates, but that this may be impossible for a central bank that has successfully drained all inflationary expectations out of the economy.¹ Recent papers that explore this issue in more detail (without necessarily calling it the liquidity trap) include those by Jeffrey Fuhrer and Brian Madigan, Alexander Wolman, and Athanios Orphanides and Volker Wieland.² All of these authors, like Krugman, use well-specified maximizing models to understand the importance of the zero bound. What distinguishes the present paper from the others (aside from its open economy perspective and the extraordinary clarity of its prose) is Krugman's contention that in Japan, negative real interest rates are not merely a useful weapon in the arsenal of monetary policy but an absolute necessity. Even if Japan were not in a recession, he argues, generational imbalances would still result in a negative real interest rate, at least in the short to medium run. If this is true, any full-employment equilibrium must have expected inflation (at least over the horizon that the equilibrium real rate is negative), and monetary policy is powerless to stop it. Thus the BOJ's efforts to maintain price stability are not merely neutral; they are actually contractionary in an economy badly in need of stimulation.

Few academic economists would disagree with Krugman's general conclusion that after seven years of deep recession, the time has come for the BOJ to stop trying to stabilize prices and to allow at least a bit of inflation. This part of the story is conventional wisdom, right or wrong. But Krugman's specific recommendation is far more unorthodox: he would have the Bank of Japan try to bring inflation (and inflationary expectations) up to 4 percent and keep it there for fifteen years. In a world where central banks are still congratulating themselves for conquering the inflation of the 1970s and 1980s, this is a truly radical suggestion. But if the full-employment equilibrium (medium-term) real

1. Summers (1991).

2. Fuhrer and Madigan (1997); Wolman (forthcoming); Orphanides and Wieland (1998).

interest rate is indeed negative, then, as Krugman elegantly demonstrates, inflation is eventually going to express itself in some form, no matter what the BOJ does. For example, Krugman's first model illustrates that an attempt to target next period's price level will tend to drive down the current price level (so the economy can have the inflation it needs to achieve a negative real interest rate).

This is an interesting and remarkable insight, but the prescription for long-term doses of inflation is predicated on the assumption that the full-employment real interest rate should be negative. Is this plausible in a country that is still investing well over 20 percent of GDP? What about the fact that Japanese savers can lend their surplus savings to the rest of the world, rather than accept negative real rates at home? I admit that Krugman makes a forceful case that "crazy" just might be "right." His casual argument is that the aging Japanese population, desperate to provide for its own retirement, is saving so much that it would take a negative rate of return to clear the market. He goes on to offer a simple overlapping generations model in which land yields a negative real return, even though its marginal product is positive. I should note that while this result turns on the empirically plausible assumption (at least for Japan) that future working generations will be smaller than the current one, labor-augmenting technological progress could substantially mitigate or even eliminate this problem.

As for why Japan does not simply lend its surplus to the rest of the world, where equilibrium real rates are presumably still positive, an obvious answer is that international capital markets are far from fully integrated. Moreover, Krugman notes that even if capital markets were fully integrated, imperfect integration of goods markets can still lead to real, consumption-based, interest differentials. Admittedly his model, in which the relative size of traded and nontraded goods production is exogenous, exaggerates the prospects for negative real rates. If nontraded goods are really going to be so scarce in the future, there should be a strong incentive to shift investment in that direction. This would raise future nontraded output and therefore raise the consumption-based real interest rate. Still, all in all, Krugman builds an interesting case that equilibrium real interest rates may be much lower in Japan than is suggested by historical norms.

Whether or not real interest rates need to be negative for fifteen years or just for two or three years, it is hard to argue with the view that the

time has come for Japan to risk some inflation. No one should seriously believe that the BOJ would face any significant technical problems in inflating if it puts it mind to the matter, liquidity trap or no. For example, one can feel quite confident that if the BOJ were to issue a 25 percent increase in the current supply and use it to buy back 4 percent of government nominal debt, inflationary expectations would rise. The real obstacle is that the BOJ does not want to blemish its record of price stability. As Krugman's formal analysis shows, in fact, if the BOJ does not realize that it needs to let go of its long-term price level anchor, it might as well forget about even short-term stabilization policy—but that would seem a very second order issue in the midst of a record recession. The real problem is that the BOJ does not have the big picture right. It does not realize that a good conservative central bank should be willing to let the price level rise on a rainy day—and Japan is experiencing a typhoon.

Toward the end of the paper, Krugman intimates that the new European Central Bank, with its mandate to keep inflation low, may soon face similar problems, since European demographics are similar to those of Japan. This is an interesting observation, although the European Central Bank has a sufficiently flexible mandate that it could easily target an inflation rate of 1 or 2 percent for an indefinite period—if it were to perceive that such a policy was necessary.

I have glibly asserted that the BOJ can always inflate if it wants to, simply by increasing the rate of base money growth. Compared with the "normal" situation of positive interest rates, however, an inflation-happy BOJ would be flying partly blind. That is, with the short-term nominal interest rate temporarily stuck near zero, the BOJ would have to try to engineer its monetary expansion without the benefit of a very crucial feedback variable. This increases the risk that in trying to engineer a 4 percent inflation, the BOJ might find prices going up by 20 percent. Given the dire straits that Japan currently finds itself in, however, this small risk seems worth taking, for all the reasons that Krugman argues.

It is interesting to contrast Krugman's prescription for Japan with the conventional wisdom. The conventional wisdom is that, in addition to cleaning up its banks, what Japan needs most is real fiscal stimulus (as opposed to phony accounting). Krugman rightly notes that modern models of how government spending affects output do not necessarily

yield a significant multiplier effect, even in the presence of Keynesian price rigidities—a result which does not in fact depend on whether one believes that the country is in a liquidity trap.³ But he neglects to point out that even if fiscal stimulus does not have a multiplier effect on output, it could still serve to raise the real interest rate, thereby greatly simplifying the task of the monetary authorities. And while tax cuts do not provide any stimulus if Ricardian equivalence holds—though in the model Krugman uses to demonstrate why real interest rates might be negative, Ricardian equivalence does not hold—fiscal stimulus can also be applied by increasing government spending on, say, infrastructure. Certainly, if having more government infrastructure investment means that big construction firms compete to bribe politicians and then build yet another bridge with a \$50 toll, it does not sound appealing. But considering that 30 percent of the houses in the greater Tokyo area do not have access to sewage, Narita airport is inadequate, the hospitals are awful, the university system weak, it should be possible to come up with something.

So a combination of temporary government spending and increases in money supply would solve the liquidity trap problem—if there is one. Moreover, using fiscal policy in conjunction with monetary policy might help to temper any depreciation of the yen that a monetary expansion would cause.

Krugman correctly argues that expansionary monetary policy in Japan would most likely benefit the country's neighbors and trading partners, even if it does lead to a significant depreciation of the yen exchange rate. This perspective is quite consistent with recent theoretical research on "new open economy macroeconomics."⁴ It is also quite consistent with the interpretation of the Great Depression proposed by Barry Eichengreen and Jeffrey Sachs.⁵ They argue that those countries that abandoned gold early and inflated did themselves a lot of good at relatively little cost to the rest of the world.

The lessons of recent research can be carried one step further. Uni-

3. Admittedly, in modern sticky price intertemporal models, the impact of government spending on the real interest rate may be quite different than in flexible price models. If a temporary increase in government spending leads to a concomitant increase in output, there is no tilt in the output available to consumers and no change in the equilibrium real rate. See Obstfeld and Rogoff (1995).

4. See Obstfeld and Rogoff (1996).

5. Eichengreen and Sachs (1985).

lateral Japanese monetary expansion would almost certainly be a good thing. But on top of a big monetary stimulus from Japan, it would be helpful to have a moderate level of stimulus from the United States and Europe, both to mitigate the depreciation of the yen and to enhance the global effects of the expansion. While I agree completely with Krugman that the BOJ should inflate, however, I find the prescription of 4 percent inflation for fifteen years too exotic. A shorter, sharper boost would seem to make more sense—say, three years of inflation cumulating to 20 percent. But then, I do not quite buy the view that short- and medium-term full-employment real interest rates for Japan are negative. And even if they are negative, the right policy is probably to raise the real interest rate through expansionary fiscal policy, which would then free monetary policy from its supposed liquidity trap.

Before closing, I should mention a couple of points about the modeling, which is certainly masterful. First, the theoretical results on the costs and existence of liquidity traps can be quite sensitive to the way in which money is introduced; shopping time models and money-in-the-utility function models can have a Pigou-type effect and yield different results. Second, it should be noted that if one subscribes to the Leeper-Sims-Woodford fiscal theory of the price level—which I do not, but it is darn clever—then there are reasons other than a liquidity trap why the central bank might lose short-term control over the price level.⁶

Although I have taken issue with some of the more unorthodox prescriptions in this paper, let me conclude by reiterating that it is a stunning piece of work. And it is going to make a lot of economists think harder about a problem that we should have been thinking hard about already.

General discussion: Benjamin Friedman applauded the paper, hoping that it would push central bankers and the central banking literature away from the past decade's mania that inflation is the only problem a central bank should worry about. He paraphrased James Tobin's words about the U.S. economy in an earlier period: "there are things worse than inflation, and Japan has them." However, he suggested that Krugman overstated the difference between his notion of the liquidity trap and the original Keynesian liquidity trap; Friedman himself saw the

6. See Woodford (1998).

original liquidity trap as hinging just as much on intertemporal issues as Krugman's modern version. The Keynesian trap reflected individuals' preference for holding money rather than long-term bonds, because of expected future capital losses on the bonds. Policymakers could not push bond prices any higher, nor the interest rate any lower, because expected future bond prices do not move one to one with current prices; having been pushed up, they are expected to fall, thus resulting in capital losses. The focus on long-term bonds made Keynes's original discussion implicitly intertemporal, while Krugman's paper is explicitly intertemporal with one-period bonds. Friedman also took issue with Krugman's uncritical acceptance of monetary neutrality. If monetary neutrality is to apply to something other than changes in the unit of account, it depends on a number of well-known assumptions, each of which would need to be verified and some of which are likely not to be true in the context of application to a real world economy such as Japan's.

Martin Baily noted that a risk-free interest rate near zero does not mean that the borrowing rate for investment in private capital is low. Just as in the Great Depression, there is now in Japan a substantial risk premium on private borrowing, and this may have increased as the environment has worsened. Consequently, he noted it is wrong to infer from the risk-free rate that there is no constraint on borrowing or that the rate of return on capital investment is negative. Edmund Phelps expressed doubt that the liquidity trap was central to Japan's problems. He thought that the excess supply of nontradable capital goods and the associated fall in their price below the reproduction cost was the key problem.

While accepting the view that Japan is in a trap, William Dickens was skeptical about the paper's policy recommendations. He believed that in Japan today, people save and hold on to real balances because of uncertainty about the future, fear of job loss, and the instability of the economic, and even the political, systems. Given these fears, he was not convinced that inflation would induce people to reduce real balances or spend more, rather than to increase saving for the uncertain future. He also wondered how many years of 4 percent inflation it would take for the Bank of Japan to convince people that it would not eventually deflate, given its strong perceived commitment to price stability. In sum, he found it questionable whether either consumption or invest-

ment spending would rise if Japan followed the paper's prescription. He feared that the only mechanism that might work with Krugman's policy would be the further depreciation of the yen, which could have significant negative effects on the East Asian region as a whole, particularly if it led China to devalue. He concluded that greater fiscal stimulus would be a preferable course of action, leading to a domestic demand-driven recovery that would not have the effect of increasing Japanese trade surpluses at the expense of other East Asian countries.

Robert Gordon also criticized the assumption in Krugman's models that the monetary authorities can easily change inflationary expectations for the future—that the announcement of a policy will change expectations despite present slack in the economy. He believed that agents' expectations depend largely on actual experience, and that they will experience increased inflation only when there is pressure in the markets for goods, services, and labor. Alan Blinder agreed. He thought that Krugman's inflationary policy would work if it could be implemented; but that would require the Bank of Japan to create expected inflation, which, in turn, would require persuading people that the future was going to be fundamentally different from the past. Japan had zero inflation in the past six years, and the average in the previous decade was 1.8 percent per year. Thus to create expected inflation of 4 percent, with actual inflation lagging behind, would be difficult. Baily concurred, observing that it would be easy for Russia to be credible in announcing inflationary policy but hard for Japan.

Baily and Gordon also agreed with Dickens that stimulating the Japanese economy through increases in the trade surplus looked a terrible idea in the current East Asian situation, just as it was in the 1930s. Krugman responded that a yen depreciation would probably lead only to a modest increase in the Japanese current account surplus and would not have a major impact on the rest of Asia. The reason people fear a yen depreciation is that they think this might set off a speculative panic, which would cause massive depreciations of other Asian currencies. Krugman suggested that this kind of reasoning held hostage the macroeconomic policy of the second largest world economy to what one thinks speculators might believe, even though they should not.

Gordon did not see why Krugman dismissed fiscal policy on the basis of Ricardian equivalence. He thought that the proper policy was a fiscal expansion financed by money creation. Since the government would

not be issuing any debt, Ricardian equivalence would not apply. Gordon believed that the U.S. evidence clearly showed that fiscal expansion resulting from preparation for war was critical to ending the Depression. Joseph Stiglitz thought the paper's discussion of Ricardian equivalence was useful in view of the recent disagreement between the United States and Japan, when the former advocated a permanent tax cut while the latter argued that a cut should only be temporary. If the government has a strong dislike for fiscal deficits, it is difficult to sound credible in announcing a permanent tax cut that will lead to a continuing budget deficit and more credible to seek temporary cuts. Stiglitz also thought that the intertemporal substitution effects of those policies may be as important as their income effects. A temporary consumption tax cut or a subsidy could provide strong incentives to increase current expenditures. Phelps commented that fiscal policy, too, could affect intertemporal prices, agreeing with Kenneth Rogoff that one could use it to increase the marginal efficiency of capital. Phelps thought this could be achieved by subsidies to either investment or employment.

Baily suggested that it was useful to think about the policy problem in Japan as how to shift out an IS curve, whether by government expenditures or by policies designed to increase private demand. But he wondered about the implications of the IS curve being so far to the left in the first place. He suggested that poor investment policies led to very high investment in some areas of the economy, drove up the capital-to-output ratio, and drove down the returns to capital. This has left overinvestment and excess capacity in industrial parts of the economy. At the same time, there has been substantial underinvestment in other areas, such as retailing, which offer Japanese investors high returns. He also saw a potential to expand residential housing despite demographic trends, since most existing housing units are so small, and reasoned that increasing the ability of consumers to borrow and use credit cards would reduce the saving rate.

Stiglitz discussed some of the political economy aspects of increasing government spending in Japan. Most observers think that increases would take the form of spending on construction, possibly further depressing real estate prices. In addition, since the construction industry is widely perceived as having too much influence under the current regime, giving it more support would be unpopular. He also noted that the Japanese Ministry of Finance believes strongly that it is irrespon-

sible to have large budget deficits, and noted that this view was reinforced by some interpretations of the recent U.S. experience, according to which a reduced federal deficit has led to economic expansion. In Stiglitz's own view, the deficit reduction helped the U.S. economy mainly by reinvigorating the banking system. In the early 1990s banks held a large amount of long-term government debt, and when interest rates came down, the banks effectively got a big injection of capital that increased their net worth.

Stiglitz believed that the credit crunch was more important than indicated by the paper. Indeed, he thought that fear of worsening credit constraints was one reason why the government had not moved more quickly to restructure the banking system. However, he was bewildered by the belief of some in Japan that a yen depreciation would reduce the ability of banks to make loans. He argued that weakening the yen would increase the net worth of Japanese banks that owned a lot of assets abroad, and would thus alleviate the problem. Blinder agreed with Stiglitz that the paper was too quick to dismiss the presence of a credit crunch. He observed that weak economies weaken banks, while weak banks weaken economies, so that a crunch was not surprising. But he acknowledged that the existence of a crunch did not alter the main argument of the paper. Friedman agreed with Krugman that insolvent banks lend too much, not too little, in the absence of regulation or supervision. He noted that there are important constraints, however, and that the weight Japanese authorities placed on sticking to the Basle capital requirements might explain why one observed signs of a credit crunch.

References

- Bernanke, Ben S. 1994. "The Macroeconomics of the Great Depression: A Comparative Approach." Working Paper 4814. Cambridge, Mass.: National Bureau of Economic Research (August).
- Bryant, Ralph C., and others, eds. 1988. *Empirical Macroeconomics for Interdependent Economies*, supplemental volume. Brookings.
- Cargill, Thomas, Michael M. Hutchinson, and Takatoshi Ito. 1997. *The Political Economy of Japanese Monetary Policy*. MIT Press.
- Cooper, Russell, and Dean Corbae. 1997. "Financial Fragility and the Great Depression." Working Paper 6094. Cambridge, Mass.: National Bureau of Economic Research (July).
- Corsetti, G., Paolo Pesenti, and Nouriel Roubini. 1998. "Paper Tigers? A Preliminary Assessment of the Asian Crisis." Unpublished paper. New York University.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91(3): 401-19.
- Eichengreen, Barry, and Jeffrey D. Sachs. 1985. "Exchange Rates and Economic Recovery in the 1930s." *Journal of Economic History* 45(4): 925-46.
- Fleming, J. Marcus. 1962. "Domestic Financial Policies Under Fixed and Under Floating Exchange Rates." *IMF Staff Papers* 9: 369-79.
- Frankel, Jeffrey A. 1988. "Ambiguous Policy Multipliers in Theory and in Empirical Models." In *Empirical Macroeconomics for Interdependent Economies*, edited by Ralph C. Bryant and others. Brookings.
- Friedman, Milton. 1969. *The Optimum Quantity of Money and Other Essays*. Chicago: Aldine.
- Friedman, Milton, and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States*. Princeton University Press.
- Fuhrer, Jeffrey C., and Brian F. Madigan. 1997. "Monetary Policy when Interest Rates Are Bounded at Zero." *Review of Economics and Statistics* 79(4): 573-85.
- Giorno, Claude, and others. 1995. "Potential Output, Output Gaps and Structural Budget Balances." *OECD Economic Studies* 24: 167-209.
- Gordon, Robert J. 1988. "Back to the Future: European Unemployment Today Viewed from America in 1939." *BPEA*, 1:1988, 271-304.
- Hicks, John Richard. 1937. "Mr. Keynes and the 'Classics.'" *Econometrica* 5(2): 147-59.
- Krugman, Paul R. 1998. "What Happened to Asia?" Unpublished paper. Massachusetts Institute of Technology.
- McKinnon, Ronald I., and Huw Pill. 1997. "Credible Economic Liberaliza-

- tions and Overborrowing." *American Economic Review, Papers and Proceedings* 87(2): 189–93.
- Metzler, Lloyd A. 1951. "Wealth, Saving, and the Rate of Interest." *Journal of Political Economy* 59(2): 93–116.
- Mundell, Robert A. 1963. "Capital Mobility and Stabilization Policy under Fixed and Flexible Exchange Rates." *Canadian Journal of Economics and Political Science* 29(4): 475–85.
- Obstfeld, Maurice, and Kenneth Rogoff. 1995. "Exchange Rate Dynamics Redux." *Journal of Political Economy* 103(3): 624–60.
- . 1996. *Foundations of International Macroeconomics*. MIT Press.
- Orphanides, Athanios, and Volker Wieland. 1998. "Monetary Policy Effectiveness when Nominal Interest Rates Are Bounded at Zero." Unpublished paper. Board of Governors of the Federal Reserve System (June).
- Organisation for Economic Co-operation and Development. 1997. *Economic Surveys: Japan*. Paris.
- Romer, Christina D. 1992. "What Ended the Great Depression?" *Journal of Economic History* 52(4): 757–84.
- Smithers, Andrew. 1998. *Japan's Problems of Debt and Demography*. London: Smithers and Co.
- Summers, Lawrence. 1991. "How Should Long-Term Monetary Policy Be Determined?" *Journal of Money, Credit, and Banking* 23(3, pt. 2): 625–31.
- Temin, Peter. 1976. *Did Monetary Forces Cause the Great Depression?* W.W. Norton.
- U.S. Bureau of Economic Analysis. 1973. *Long-Term Economic Growth, 1860–1970*. Department of Commerce.
- Wolman, Alexander. Forthcoming. "Staggered Price Setting and the Zero Bound on Nominal Interest Rates." Federal Reserve Bank of Richmond *Economic Quarterly*.
- Woodford, Michael. 1998. "Public Debt and the Price Level." Unpublished paper. Princeton University (May).