



Liquidity, Business Cycles, and Monetary Policy

Nobuhiro Kiyotaki (Princeton University)
and John Moore (London School of Economics)

Paper presented at the Financial Cycles, Liquidity, and Securitization Conference
Hosted by the International Monetary Fund
Washington, DC—April 18, 2008

The views expressed in this paper are those of the author(s) only, and the presence of them, or of links to them, on the IMF website does not imply that the IMF, its Executive Board, or its management endorses or shares the views expressed in the paper.

Clarendon Lectures

Lecture 2

LIQUIDITY, BUSINESS CYCLES, AND MONETARY POLICY

by

Nobuhiro Kiyotaki
London School of Economics

and

John Moore
Edinburgh University and London School of Economics

27 November 2001

As I said yesterday, my lectures are based on joint research with Nobu Kiyotaki of the L.S.E..

In case some of you couldn't be here yesterday, today's lecture will be self-contained. But occasionally I'll need to recap.

Economists' views on money

Money. Economists' attitudes towards money vary a great deal. As a rough classification, there are three groups. The first group might be described as "nonmonetarists". A nonmonetarist is someone who thinks that money doesn't matter.

Nobu spent last year at M.I.T. He got into a discussion about money and the payments system. One of his colleagues said, "Oh, money, the payments system -- it's all just plumbing." Thus speaks a nonmonetarist.

Actually, the plumbing analogy is revealing. In a well-functioning plumbing system, the flow is all in one direction. The same could be said of much of modern macroeconomics. Nobu's M.I.T. colleague is a signed-up member of the S.E.D. -- the Society for Economic Dichotomists. S.E.D. members work out quantities first, and then, if they feel in the mood, back out asset prices. There's a one-way flow from quantities to asset prices.

Of course if the plumbing system fails -- if there is a blockage -- the system becomes unpleasantly two-way. When it comes to plumbing, feedback is not good news.

When it comes to the macroeconomy, however, we contend that there are rich two-way interactions between quantities and asset prices. We believe that these interactions are of first-order importance. It's inadequate to think of money in terms of plumbing. A better analogy is the one I gave yesterday: the flow of money and private securities through the economy is like the flow of blood. And prices are like the nervous system. Just as there is a complex interaction between the body, the nervous system, and the

flow-of-blood, so there is a complex interaction between quantities, asset prices, and the flow-of-funds.

Our model is of an economy in which money is essential to the allocation of resources. Let me define such an economy as a "monetary economy". There will be no nominal rigidities, and cash will not be imposed on the economy as a necessity.

I want to show you that, in the context of such a monetary economy, a number of famous puzzles can be better understood. Among the anomalies I have in mind are: the excess volatility of asset prices; the equity premium puzzle and its flip-side, the low risk-free rate puzzle; the anomalous savings behaviour of certain households, and their low rates of participation in asset markets. I want to persuade the nonmonetarists among you -- perhaps you should be called "realists" -- that these apparent anomalies of the "real economy" are in fact normal features of a monetary economy. It is precisely because there is an essential role for money that these so-called puzzles arise.

The second group might be described as "pragmatists". A pragmatist is someone who wants to get on with the job of analysing and advising on monetary policy, monetary union, and macroeconomic management generally. He or she needs a model of money to use. The leading off-the-shelf models these days seem to be cash-in-advance and dynamic sticky price models.

There are well-known concerns about those models. Money can be seen more as grit-in-the-system than a lubricant in the models, so they aren't models of a monetary economy as I have defined it. The peculiar role of money is imposed rather than explained, so the models do not satisfy the Wallace Dictum. In his dictum, Neil Wallace exhorts us not to make money a primitive in our theories. Equally, he would argue that a firm should not be a primitive in industrial organization theory, and that bonds and equity should not be primitives in finance theory.

The Wallace Dictum doesn't cut much ice with the pragmatists. After all, they would argue, industrial economics and finance theory have been remarkably successful in taking firms, bonds and equity as building blocks --

without opening up the contractual foundations. So why not assume cash in advance to get on with our macroeconomics? It's fair to say that monetary policy analysis would be in a bad shape were it not for the cash-in-advance short cut.

Nevertheless, we want to know about the effectiveness of monetary policy in a context where money is essential rather than grit in the system, and where there are no nominal rigidities. The medium run, perhaps. The model this evening will show that open market operations are indeed effective, but, interestingly, in a way that depends on the full time path of policy.

More generally, we want to have a broader understanding of liquidity. Keynes, Tobin, and even Friedman, weren't focussed on the narrow money/bonds tradeoff; they were concerned with policy across the entire spectrum of assets: money, bonds, equity, physical capital, and human capital -- each differing in its degree of liquidity. Cash-in-advance or dynamic sticky-price models are not well suited to answering larger questions to do with liquidity. By the end of my talk, I hope I will have convinced the pragmatists among you we have made some progress on this front.

The third group might be described as "fundamentalists". A fundamentalist is someone who cares deeply about what money is and how it should be modelled. A fundamentalist builds pukka models that satisfy the Wallace Dictum.

In recent years, the model on which the fundamentalists have lavished most attention is based on a random matching framework. A matching model captures the ancient idea that money lubricates trade in the absence of formal markets. Without money, opportunities for bilateral trade would be rare, given that a coincidence of wants between two people is unlikely when there are many types of good.

The matching models are without doubt ingenious and beautiful. But it's quite hard to integrate them with the rest of macroeconomic theory -- not least because they jettison the basic tool of our trade, competitive markets. The jury is out on what they will eventually deliver. But I am reminded of a commercial from the early days of Scottish television. The

commercial was for a strong beer, known as "ninety shilling" in Scotland. The woman at the bar sips her glass of ninety shilling, winces, and says: "Oh it's too strong for me. But I like the men who drink it." I guess that's how I feel about the random matching model.

Recap on lecture 1

Let me briefly recap on yesterday's lecture. Nobu and I see the lack of coincidence of wants as an essential part of any theory of money. But not necessarily over types of good. Rather, the emphasis should be on the lack of coincidence of wants over dated goods. For example, suppose you and I meet today. What day is it? Tuesday. I may want goods from you today to invest in a project that yields output in two days' time, on Thursday. You have goods today to give me, but unfortunately you want goods back tomorrow, Wednesday. Thus we have a lack of coincidence of wants in dated goods: I want to borrow long-term; you want to save short-term.

With this switch of emphasis, from the type dimension to the time dimension, comes a change in modelling strategy. We no longer need to assume that people have difficulty meeting each other, as in a random matching model. Without such trading frictions, we can breathe the pure oxygen of perfectly competitive markets. In fact, you'll see that in this evening's model there is only one departure from the standard dynamic general equilibrium framework.

Instead of assuming that people have difficulty meeting each other, we assume that they have difficulty trusting each other. There is limited commitment. If you don't fully trust me to pay you back on Thursday, then I am constrained in how much I can borrow from you today. And tomorrow, you may be constrained if you try to sell my IOU to a third party, possibly because the third party may trust me even less than you do. Both kinds of constraint, my borrowing constraint today and your resale constraint tomorrow, come under the general heading of "liquidity constraints", and stem from a lack of trust. We think that the lack of trust is the right starting point for a theory of money.

You will see that these two kinds of liquidity constraint are at the heart of the model. Not only do entrepreneurs face constraints when trying to raise funds, to sell paper; but also, crucially, the initial creditors, the people who buy the entrepreneurs' paper, face constraints when passing it on to new creditors. That is, not only am I constrained borrowing from you today, Tuesday, but also you are constrained reselling my paper tomorrow, Wednesday. It's your "Wednesday constraint" that is unconventional, and adds the twist to the model.

The model I presented yesterday was deterministic, both in aggregate and at the individual level. Also, I focussed on inside money -- the circulation of private debt. Only at the end of yesterday's lecture did I touch on the fact that outside money (non-interest-bearing fiat money) might circulate alongside inside money -- provided the liquidity shortage is deep enough. For most of the lecture, there was no fiat money.

The advantage of such an approach is that it teaches us that money and liquidity may, at root, have nothing to do with uncertainty or government. Of course, the disadvantage of yesterday's model is that it is a hopeless vehicle for thinking about government policy in a business cycle setting. That is the purpose of this evening's lecture: to model fiat money explicitly, in a stochastic environment.

The model

The model is an infinite-horizon, discrete-time economy. At each date t , in aggregate there are Y_t goods produced from a capital stock K_t . Goods are perishable. Capital is durable.

In addition, there is a stock of money, M . Money is intrinsically useless. Later I will be introducing a government, which adjusts the money supply, so M will have a subscript t . Indeed, at that point, you could reinterpret M_t as government bonds. But for now, think of M as the stock of seashells.

There is a continuum of agents, with measure 1. Each has a standard

expected discounted logarithmic utility over consumption of goods:

$$E_t \sum_{s=0}^{\infty} \beta^s \log c_{t+s}.$$

β is the discount factor. Whenever I use a Greek letter it refers to an exogenous parameter lying strictly between 0 and 1.

All agents use their capital to produce goods. If an agent starts date t with k_t capital, by the end of the date he will have produced $r_t k_t$ goods:

$$\begin{array}{ccc}
 k_t \text{ capital} & \xrightarrow{\hspace{2cm}} & \left\{ \begin{array}{l} r_t k_t \text{ goods} \\ \lambda k_t \text{ capital} \end{array} \right. \\
 \text{start of} & & \text{end of} \\
 \text{date } t & & \text{date } t
 \end{array}$$

λ is the depreciation factor. Notice that depreciation happens during the period, i.e. during production, not between periods.

Individually, production is constant returns: the productivity r_t is parametric to each agent. But in aggregate there are decreasing returns:

$$r_t = a_t K_t^{\alpha-1}$$

which is decreasing in the aggregate capital stock K_t . Aggregate output is of course increasing in K_t :

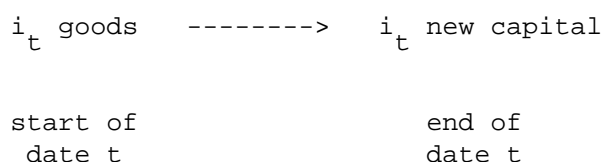
$$Y_t = r_t K_t = a_t K_t^{\alpha}.$$

One interpretation to have in mind is that there is a missing factor of

production, such as labour. The underlying technology has constant returns to capital and labour. The expression for r_t here is a reduced form, taking into account the aggregate labour supply. Our written paper models workers explicitly, but in this lecture let's keep them in the background.

The technology parameter a_t follows a stationary Markov process in the neighbourhood of some constant level a .

So all the agents produce goods from capital. But in addition, some of the agents produce capital from goods. Specifically, at each date t , a fraction π of the agents have what we call an "investment opportunity": i_t goods invested at the start of the period make i_t units of new capital by the end of the period:



Notice that the technology has constant returns -- in fact it is 1 for 1. Also, notice that new capital cannot be used for the production of goods until the next period.

An agent learns whether or not he has an investment opportunity at the start of the day, before trading. The point to stress here is that the chance to invest comes and goes. Investment opportunities are i.i.d. across people and through time. The problem facing the economy is to funnel resources quickly enough from the hands of those agents who don't have an investment opportunity into the hands of those who do -- that is, to get goods from the savers to the investors. Of course, to implement this in a decentralised environment, investors must have something to offer savers in return -- and that will prove to be the nub of the problem.

It simplifies the dynamic analysis later on to make the mild assumption that the fraction of investors, π , is greater than the depreciation rate, $1-\lambda$, which in turn is greater than the discount rate, $1-\beta$:

$$\pi > 1-\lambda > 1-\beta.$$

Capital is specific to the agent who produced it. But he can mortgage future returns by issuing paper. Normalise one unit of paper issued at date t so that it is a promise to deliver r_{t+1} goods at date $t+1$, λr_{t+2} goods at date $t+2$, $\lambda^2 r_{t+3}$ goods at date $t+3$, on so on. In other words, the profile of returns matches the return on capital. The returns depreciate by λ each period. And, viewed from the date of issue, they are stochastic. One can think of paper as an equity share.

At each date t , there are competitive markets. Let q_t be the price of a unit of paper, in terms of goods. And let p_t be the price of money, in terms of goods. Beware that this is upside down: usually p_t is the price of goods in terms of money. But we don't want to prejudge whether or not money will have value. Indeed, for a range of parameter values, money will not have any value. So it's sensible to make goods the numeraire.

I want to rule out insurance. That is, an agent cannot insure against having an investment opportunity. Since all agents are essentially the same, what I am really ruling out is some kind of mutual insurance scheme. A variety of assumptions could be made to justify this. For example, it may be impossible to verify whether an agent has an investment opportunity. Or it may take too long to verify -- by the time verification is completed, the opportunity will have gone. With asymmetric information, self-reporting schemes would have to be part of an incentive-compatible long-term multilateral contract: agents would have to have an incentive to tell the truth. Recent research suggests that truth-telling may be hard to achieve when agents have private information not only about their investment opportunities but also about their asset holdings.

Anyway, we believe that, in broad terms, our results would still hold even if partial insurance were feasible. But for now I want simply to rule out all insurance.

Now to the two central assumptions. First, an investing agent can mortgage at most a fraction θ_1 of (the future returns from) his new capital production.

an investing agent can mortgage at most a fraction θ_1 of (the future returns from) his new capital

As a result, investment may not be entirely self-financing. An investing agent may face a borrowing constraint. A variety of moral hazard assumptions could be appended to justify θ_1 . For example, if an agent commits too great a fraction of his future output he will default. (As we have defined it, paper is default-free.) Note that we must also assume some degree of anonymity, to rule out the possibility that social sanctions can be used to deter default. We don't want to get into supergame equilibria where agents can be excluded from the market. Anyway, without further ado, I make the crude assumption that θ_1 is the most an agent can credibly pledge of the output from new capital at the time of the investment.

The second central assumption is just as crude, but is non-standard. I want to assume that at each date, an agent can sell at most a fraction θ_2 of his paper holdings.

at each date, an agent can resell at most a fraction θ_2 of his paper holdings
--

The point is that if an agent turns out to have an investment opportunity at some date, then, before the investment opportunity disappears, he can exchange only a fraction θ_2 of his paper holdings for goods to be used as input. This does not mean that he is lumbered with holding the residual fraction, $1 - \theta_2$, for ever. He can sell a further fraction θ_2 of that

residual at the next date. In other words, he could eventually sell off his entire paper holding, but it would take time, because he would have to run it down geometrically, at the rate θ_2 . Think of this as peeling an onion slowly, layer by layer.

θ_2 measures the liquidity of paper, and is to be distinguished from the liquidity of money (whose θ_2 equals 1).

One natural justification for θ_2 is that a potential buyer of paper needs to verify that the paper is secured against a bona fide investment project. He needs to inspect the project's assets. But this takes time. By the time the buyer has finished inspecting, it may be too late for the seller of the paper to take advantage of his investment opportunity. In this race between verifying the existing assets and investing in new assets, θ_2 is the probability that the verification finishes first.

A better model would have the sale price of paper be a function of how fast it is sold -- on the grounds that anything can be sold quickly, as long as the price is low enough. Fascinating though that is, I want to stick to the crude assumption that agents face a resaleability constraint that precludes them from divesting more than a fraction θ_2 of their paper holdings per period. At the end of the lecture I will review the assumption. But for now let's see where it leads.

Both constraints, the borrowing constraint θ_1 and the resaleability constraint θ_2 , come under the heading of "liquidity constraints". They are the twin pillars of the model. Were θ_1 equal to 1, new investment would be self-financing, and the liquidity of agents' portfolios would be immaterial. And were θ_2 equal to 1, there would be no difference in liquidity between money and paper, and the purpose of our analysis would be lost.

Recall from yesterday's lecture the mnemonic: The subscript 1 on θ_1 denotes a constraint on the initial sale of paper by an investing agent to a saver. And the subscript 2 on θ_2 denotes a constraint on the resale by this saver to another saver at a later date.

In terms of the Tuesday/Wednesday/Thursday example I gave earlier, θ_1

corresponds to my borrowing constraint on Tuesday. And θ_2 corresponds to your resaleability constraint on Wednesday.

In a world where θ_1 and θ_2 are both strictly less than 1, an agent has three kinds of asset in his portfolio: money, paper and unmortgaged capital. We don't really need or want to have a model with three assets: two would be enough to get us going. Moreover, the three-asset model would be extremely hard to analyse because aggregation would be impossible by hand. We don't want to have to keep track of the distribution of asset holdings -- remember that although the agents are intrinsically identical, they have individual histories of investment opportunities.

With this all in mind, it helps enormously to make the following simplifying assumption: at every date, an agent can mortgage up to a fraction θ_1 of his unmortgaged capital stock. In other words, the onion analogy applies to the mortgaging of capital as well as to the sale of paper. Also, let us assume that θ_1 and θ_2 equal some common value, θ . The upshot is that now paper and unmortgaged capital are perfect substitutes as means of saving. They yield a common return stream, declining by a factor λ . And they have the same degree of liquidity: a fraction θ can be sold for goods in each period.

Thanks to this simplifying assumption, an agent in effect holds only two assets: a liquid asset, money; and an illiquid asset, paper plus unmortgaged capital. Paper and unmortgaged capital might better be described as semi-liquid, but let me use the adjective illiquid, in contrast to perfectly liquid money. At the start of date t , let m_t denote the money an agent holds, and let n_t denote the quantity of paper plus unmortgaged capital that he holds.

The simplification also enables us to collapse the borrowing constraint θ_1 and the resaleability constraint θ_2 into a single liquidity constraint (*):

$$\begin{array}{rcl}
 n_{t+1} & \geq & (1 - \theta)(i_t + \lambda n_t) \quad (*) \\
 \vdots & & \vdots \\
 \vdots & & \vdots \\
 \text{paper holding} & & \text{new capital} \quad \text{paper holding} \\
 \text{plus unmortgaged} & & \text{production} \quad \text{plus unmortgaged} \\
 \text{capital stock} & & \text{during } t \quad \text{capital stock} \\
 \text{at start of } t+1 & & \text{(if any)} \quad \text{at end of } t
 \end{array}$$

The paper plus unmortgaged capital that an agent holds at the start of period t depreciates to λn_t by the end of the period, but may have been augmented by new capital production i_t if the agent was lucky enough to have an investment opportunity. The borrowing constraint says that only a fraction θ of i_t can be sold, and the resaleability constraint says that only a fraction θ of λn_t can be sold. All in all, the agent must hold at least $(1 - \theta)(i_t + \lambda n_t)$ of paper plus unmortgaged capital at the start of period $t+1$.

It is cumbersome to keep saying "paper plus unmortgaged capital" every time, so let me simply say "paper" as a shorthand for the sum of the two.

So that is the set-up of the model. Let's turn to some preliminary results.

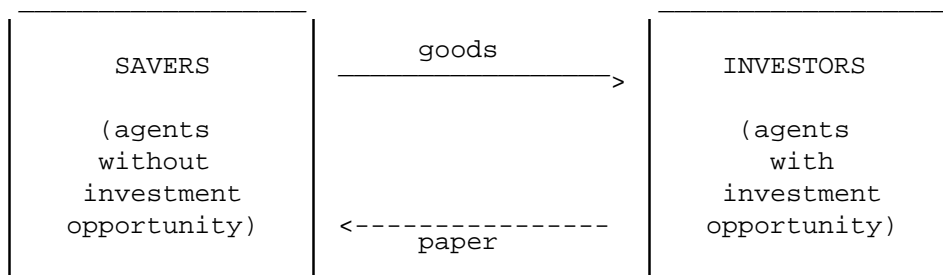
Preliminary results

First, if θ is large enough, the single liquidity constraint (*) does not bind in the neighbourhood of steady state, and the economy runs at first-best. Specifically, if θ is above some critical level θ^* , which is strictly less than 1, then at each date t the price of paper, q_t , equals the production cost of capital, 1. That is, Tobin's q equals unity. And the rate of return on paper -- i.e. tomorrow's return r_{t+1} plus depreciated value λq_{t+1} divided by today's price q_t -- equals the subjective rate of return:

$$\frac{r_{t+1} + \lambda q_{t+1}}{q_t} = \frac{1}{\beta}$$

(This is for $a_t \equiv a$.) Since q_t and q_{t+1} equal 1, this pins down the value of $r_{t+1} = (1 - \beta\lambda)/\beta$, and we can invert the aggregate production function to find the first-best level of the aggregate capital stock, K^* .

There is no role for money here: p_t equals zero. The paper market is sufficiently liquid that enough resources -- goods -- can be transferred from the savers to the investors:



The fact that the paper line is pecked rather than continuous connotes the idea that the flow of paper is subject to the liquidity constraint (*). Suppose θ lies strictly between θ^* and 1. Then although paper is not perfectly liquid, investing agents are indifferent between investing and not, and in equilibrium they are inside their liquidity constraint. Enough resources are being funnelled from savers to investors to maintain the first-best level of aggregate capital stock K^* . The exact value of π , the fraction of agents who have an investment opportunity is immaterial. Since the production of capital has constant returns, it does not matter exactly how many people do it. Also insurance would be irrelevant, since people are indifferent about whether or not they have an investment opportunity: the rates of return on saving and investment are the same.

However, π cannot be too small. The smaller the fraction of investors, the narrower the funnel, and the greater the strain on the flow of paper. To put this another way, as π falls, the critical degree of liquidity needed to sustain first-best, θ^* , rises.

We are interested in the regime where the first-best cannot be sustained. For θ below θ^* , the strain on the paper market is too great. The

supply of paper that investing agents are able to sell is too small. The price of paper, q_t , is thereby raised -- Tobin's q is above unity. Each investing agent is liquidity constrained: he has access to a constant-returns technology for the production of new capital whose unit cost, 1, is strictly less than the market price of the return flow, q_t . Unfortunately he can only mortgage a fraction $\theta_1 = \theta$ of that flow, and the value of q_t times θ is strictly less than the cost 1. Hence his scale of operation is determined by his flow of funds. To maximize production, he sells as many of his assets as he can, subject to the binding liquidity constraint (*).

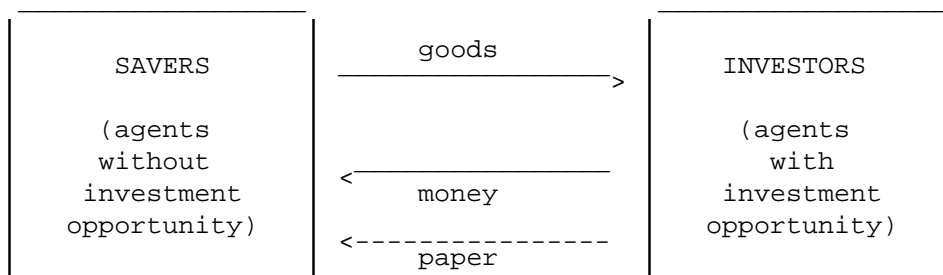
If θ is sufficiently far below θ^* , there is a role for money. Specifically, we learn from Proposition 1 that if θ is strictly below some smaller critical value $\hat{\theta}$, then there is a monetary equilibrium with p_t greater than zero.

Proposition 1: For θ below some critical value $\hat{\theta} < \theta^*$,

$$p_t > 0 \quad \text{and} \quad q_t > 1$$

in the neighbourhood of steady state.

At its simplest, money is providing an additional lubricant for the flow of goods between savers and investors.



To put this more sharply, savers are holding money in their portfolios because they realise that when an investment opportunity comes their way in

the future they will be glad of the extra liquidity that money provides relative to paper.

There is the obvious caveat that even when θ is less than $\hat{\theta}$ there is always a nonmonetary equilibrium, given that no-one wants to hold money if no-one else wants to. The nonmonetary equilibrium is not of interest to us. What matters for our purposes is that there exists a monetary equilibrium if and only if θ is strictly less than $\hat{\theta}$. This meets the definition of a "monetary economy" that I gave at the start. Money is indeed essential to the smooth allocation of resources: it is not grit in the system, and it has not been imposed on the economy.

Characterisation of a monetary equilibrium

From now on, I want to focus on a monetary equilibrium, so let's assume θ is strictly below $\hat{\theta}$. To characterise equilibrium, we need to find out the agents' optimal behaviour: consumption, investment and saving. This is not entirely straightforward, because the illiquidity of paper introduces some delicacies into the analysis. The details are spelt out in our written paper. Here let me simply report the answers, and try to convince you that they are correct.

Consider a typical agent at the start of period t who happens to hold m_t money and n_t paper. First, suppose he has an investment opportunity. Put a superscript i on his consumption c_t^i , to indicate that he is an investor. Thanks to the logarithmic utility, his optimal consumption is a fixed fraction $1-\beta$ of his net worth:

$$c_t^i = (1-\beta)(r_t n_t + \lambda n_t + p_t m_t).$$

net worth

His net worth is made up of the return on his paper, plus the capital worth of his paper (after depreciation), plus the worth of his money. The first and last terms are straightforward. But the middle term poses the problem:

how should the agent price an asset for which the resaleability constraint is binding (remember this is an agent with an investment opportunity)? The answer is to value all of his paper/unmortgaged capital -- even the inframarginal units that he can resell -- using the replacement cost, 1, i.e. the cost of producing an additional unit. He doesn't use the market price q_t . The point is that when an investing agent hits the wall of the resaleability constraint, the only relevant margin is whether to invest more or consume (he doesn't want to keep any money). That's why λn_t has the coefficient 1, not q_t , in the expression for c_t^i .

The rest of the agent's available liquid funds are ploughed into investment:

$$i_t = \frac{r_t n_t + q_t \theta \lambda n_t + p_t m_t - c_t^i}{(1 - q_t \theta)} = \frac{\text{funds at hand}}{\text{downpayment}}.$$

He exchanges all his money for goods. And he exchanges the maximum fraction θ of his paper for goods. Then, having consumed c_t^i , he makes levered investment. The gross cost of each unit of new capital is 1, but he can mortgage a fraction $\theta_1 = \theta$ at price q_t , so the required downpayment per unit of investment is $1 - q_t \theta$.

Notice that the expression for i_t is sensitive to the prices p_t and q_t , especially q_t . The higher these prices are, the more funds the agent has at hand with which to invest, and the smaller the downpayment required for each unit of investment. Here, then, is the feedback from asset prices to quantities that I spoke about at the start, which is absent in many standard models such as the real business cycle model.

Next, suppose the agent does not have an investment opportunity, indicated by a superscript n . Now his optimal consumption is

$$c_t^n = (1-\beta)(r_t n_t + q_t \lambda n_t + p_t m_t).$$

net worth

For a saver the resaleability constraint is not binding -- in fact, savers buy paper, not sell. So it is appropriate for him to value his paper holdings at the market price q_t -- not at the unit cost of new capital, 1, as an investor does.

It remains to find the saver's optimal portfolio of money and paper. The first-order condition is given by

$$\pi E_t \left(\frac{p_{t+1}/p_t - (r_{t+1} + \lambda)/q_t}{c_{t+1}^i} \right) = (1-\pi) E_t \left(\frac{(r_{t+1} + \lambda q_{t+1})/q_t - p_{t+1}/p_t}{c_{t+1}^n} \right).$$

In many ways, this equation (together with the earlier investment equation) is the key to understanding the model. In choosing between money and paper at date t , a saver has to balance rate of return against liquidity.

With probability π , he will have an investment opportunity tomorrow, at date $t+1$. In which case the difference in the expected rates of return between money and paper will be the numerator on the LHS. Money always has a rate of return p_{t+1}/p_t . Given that he is an investor at date $t+1$, paper has a rate of return for him of $(r_{t+1} + \lambda)/q_t$ -- because after receiving r_{t+1} he values each unit of depreciated paper at its replacement cost 1, not at its market price q_{t+1} .

However, with probability $1-\pi$ he will not have an investment opportunity at date $t+1$. He will be a saver again. In which case the rate of return on paper for him is $(r_{t+1} + \lambda q_{t+1})/q_t$ -- because he uses the market price q_{t+1} to value it at date $t+1$. That is, the difference in expected rates of return between paper and money is given by the numerator on the RHS.

The expectations operator E_t is with respect to the underlying

uncertainty at date t over r_{t+1} (and hence also over the prices p_{t+1} , q_{t+1}).

An optimal portfolio balances these differentials, weighted by the probabilities and the marginal utilities -- which given logarithmic utilities equal the reciprocals of the relevant consumption levels.

Notice that both sides of this portfolio equation can be positive only if q_{t+1} is strictly greater than 1.

The great merit of the expressions for consumption, investment and savings portfolio is that they are linear in an agent's start-of-period money and paper holdings, m_t and n_t . Hence we can aggregate without needing to look at the evolution of the distribution across individuals.

Recall that, although we have been describing n_t as an agent's paper holding, it in fact denotes an agent's holding of paper plus unmortgaged capital. Hence the sum of everyone's n_t equals the aggregate capital stock K_t . Also the sum of everyone's money holdings m_t equals the aggregate money stock M , which is fixed. Therefore in aggregate the only endogenous state variable is K_t . We can write down the equilibrium as a recursive system.

The state variable is the technology parameter a_t , which follows a stationary Markov process, and K_t . Aggregate investment I_t , asset prices p_t and q_t , and tomorrow's capital stock K_{t+1} , are four unknowns that solve the four equations here. I will not go into details, but you should recognize the middle two equations. First, goods market clearing: output equals investment plus consumption:

$$r_t K_t = I_t + (1-\beta) \left(r_t K_t + [\pi + (1-\pi)q_t] \lambda K_t + p_t M \right).$$

Second, the equation for investment, which is the aggregate version of our earlier expression:

$$I_t = \frac{\pi \left(\beta r_t K_t + \theta q_t \lambda K_t + \beta p_t^M - (1-\beta) \lambda K_t \right)}{(1 - q_t \theta)}$$

Third, the portfolio equation, which is exactly as before except that aggregates are in the denominators:

$$\begin{aligned} \pi E_t \left(\frac{p_{t+1}/p_t - (r_{t+1} + \lambda)/q_t}{C_{t+1}^i} \right) \\ = (1-\pi) E_t \left(\frac{(r_{t+1} + \lambda q_{t+1})/q_t - p_{t+1}/p_t}{C_{t+1}^n} \right). \end{aligned}$$

$$\text{where } C_{t+1}^i = (r_{t+1} + \lambda)(\theta I_t + [1-\pi+\pi\theta]\lambda K_t) + p_{t+1}^M$$

$$\text{and } C_{t+1}^n = (r_{t+1} + q_{t+1}\lambda)(\theta I_t + [1-\pi+\pi\theta]\lambda K_t) + p_{t+1}^M$$

Fourth, an expression for tomorrow's capital stock in terms of today's depreciated stock plus investment:

$$K_{t+1} = \lambda K_t + I_t.$$

This four-equation, four-unknown recursive system is both backward and forward looking. It is backward looking insofar as consumption and investment are functions of the inherited capital stock. It is forward looking insofar as the asset prices p_t and q_t are jump variables, and we restrict attention to non-explosive paths.

Most important, there is the feedback from asset prices to quantities through the investment equation. The system is decentralised in an essential way: this is not the solution to an obvious planning problem.

Properties of a monetary equilibrium

From Proposition 2 we learn that there is a unique steady state.

Proposition 2: $a_t \equiv a$ implies a unique steady state I, p, q, K .

Let us consider a neighbourhood of that steady state. Proposition 3(i) concerns rates of return dominance.

Proposition 3:

$$\begin{array}{ccc}
 \text{(i)} & E_t \frac{p_{t+1}}{p_t} < & E_t \frac{r_{t+1} + \lambda q_{t+1}}{q_t} < & \frac{1}{\beta} \\
 & \vdots & \vdots & \\
 & \text{expected} & \text{expected} & \text{subjective} \\
 & \text{rate of return} & \text{rate of return} & \text{rate of return} \\
 & \text{on money} & \text{on paper} & \\
 & \vdots & & \\
 & \text{liquidity} & & \\
 & \text{premium} & &
 \end{array}$$

(ii) liquidity premium = nominal interest rate $\cong \pi(q_t - 1)$.

The left-hand inequality tells us that money is dominated by paper. The difference in the two rates of return is a liquidity premium.

The simple explanation for the liquidity premium is that savers have to be compensated for holding paper because it is less liquid than money. Notice that agents only hold money because they anticipate that they will face borrowing constraints in the future. To put this in broad terms: borrowing constraints are an integral part of a monetary economy.

A few moments' thought shows that the liquidity premium is same as the nominal interest rate on paper. Proposition 3(ii) tells us that this is roughly π times the gap between q_t and 1. This can be sizeable.

We would argue that this indirectly sheds light on the equity premium puzzle and the low risk-free rate puzzle: both puzzles can be seen as natural features of a liquidity-constrained economy -- provided one is willing to accept that at least certain kinds of equity are not as liquid as government bonds.

The right-hand inequality in Proposition 3(i) says that although the rate of return on paper is higher than money, it is still lower than the subjective rate $1/\beta$ -- because of the high price of paper. This implies that an agent's consumption and net worth shrink during an episode of saving, and only expand again after he gets an investment opportunity. There is a myriad of such individual histories, and we don't see all this fine grain in the aggregate picture. (Note that if the economy were running at first-best -- i.e. if $\theta \geq \theta^*$ -- then there would be no differences in individual histories. These differences only arise because of the gap in the rates of return between saving and investing.)

Earlier I alluded to the possibility of including workers in the model. Suppose workers, unlike our entrepreneurial agents, do not have investment opportunities: they simply supply labour to the entrepreneurs for the production of goods from capital. The right-hand inequality says that a worker will not save. He will just consume his wage. This may help explain low rates of participation in asset markets. It is not that some people face barriers to trading assets. Rather, the returns are too low.

Proposition 4 tells us that the economy is too small relative to first-best.

Proposition 4: The aggregate stock of capital K_t is strictly less than K^* .

This is not too surprising, given the difficulty the economy has in

funnelling resources from savers to investors. One expects there to be too much consumption by savers and too little investment. Remarkably, it can often be quite difficult to get underinvestment in models of this kind. However, underinvestment is a feature of the present model.

Dynamics

Let's turn now to dynamics. Rather than analyse the discrete-time model directly, it is easier to look at a continuous time approximation, found by taking the length of period to zero. These charts are drawn for the limit economy. I should say that they represent qualitative solutions calculated by hand. By Nobu's hand, to be precise. As he says, you never know how much faith to put in a computer simulation. And a computer is not good at finding general qualitative answers.

Start with productivity shocks. Recall our assumption that a_t follows a stationary Markov process. Suppose in fact that it follows a 2-state Markov process -- high and low productivity -- where transitions from one state to the other are infrequent. See Figures 1(a) and 1(b).

Because the return on capital increases with productivity, a jump in productivity causes the price of paper to jump up. So too does the liquidity premium -- recall Proposition 3(ii). Anticipating a tight liquidity constraint in the future, entrepreneurs without an investment opportunity want to hold more liquid assets, which leads to a jump in real balances: the price of money jumps up. The jumps in the prices of paper and money, q_t and p_t , raises the investing agents' available funds, and raises leverage, so investment jumps up too.

With greater investment, capital stock starts accumulating. Aggregate output, which rose instantaneously with the jump in productivity, continues to rise with capital accumulation. The return on capital falls with the higher capital stock, and the price of paper falls back towards normal levels, as does the liquidity premium. The value of money continues to increase as the economy expands.

At some point in the future, productivity jumps back down, and these processes reverse: the price of paper jumps down, the liquidity constraint loosens. Real balances and investment jump down too, and the stock of capital starts to fall.

Overall, we conclude that if productivity shocks are driving the fluctuations, then the price of paper, the liquidity premium, and real balances are all procyclical, moving together with output. Investment is procyclical and quite volatile, because it is affected by the net worth of the investing entrepreneurs and the required downpayment, which means that the movements in the prices of paper and money combine to magnify the fluctuation. Consumption is also procyclical, given that the consumption of workers is equal to their wage income and the consumption of the entrepreneurs is proportional to their net worth.

Government

It is time to introduce government into the model. We make no attempt to explain government behaviour. Our goal is simply to explore the effects of an exogenous government policy.

Unlike the private entrepreneurs, the government is unable to produce goods from capital, capital from goods. However, it has sole access to a costless money-printing technology. Let M_t denote the stock of money outstanding at the start of date t .

The government can buy paper but, like everyone else, cannot resell more than a fraction θ at each date. Let N_t^g denote the government's holding of paper at the start of date t .

Finally, the government has an expenditure of G_t goods at date t . One might think of this as transfers to workers. $G_t < 0$ corresponds to lump-sum taxation (of workers).

The government's budget constraint is

$$G_t + q_t(N_{t+1}^g - \lambda N_t) = r_t N_t^g + p_t(M_{t+1} - M_t).$$

That is, government expenditure plus paper purchases must equal the return on paper plus seignorage revenues. Since the government is large, changes in its paper or money holdings will affect prices p_t and q_t .

With an active government, the steady state is indexed by the rate of growth of the money stock, μ :

$$M_{t+1}/M_t \equiv \mu.$$

In steady state, real money balances, $p_t M_t$, are held constant by means of a fall in the price p_t :

$$p_{t+1}/p_t \equiv 1/\mu.$$

μ may be greater or less than 1. $\mu > 1$ corresponds to inflation (remember p_t is the price of money in terms of goods, not vice versa). Productivity a_t , government paper holdings N_t^g , and government expenditure G_t , are all constant in steady state.

In our written paper we compare steady states, and the long run effects of government policy. There is not time to report our findings here, but I should remark that the "Friedman Rule" -- deflating at the rate $\mu = \beta$ -- achieves first-best, provided of course that it can be adequately financed through lump-sum taxation on workers.

Here, let us concentrate on shorter run dynamics. First, all proportional "helicopter drops" of money -- anticipated or not; today or in the future -- are neutral: they simply lead to inflation. By the same token, paying nominal interest on money doesn't affect anything except the future

prices of money.

That said, we are not primarily concerned with changing the money supply by helicopter drops or by paying nominal interest on money. Our focus is on the effects of open market operations.

A simple way to investigate open market operations is to suppose that the government's holding of paper, N_t^g , follows an exogenous 2-state Markov process. For the moment, set government expenditure G_t at a constant level. Then, for the government to meet its budget constraint, it has to adjust the money supply M_t .

In the continuous time approximation, M_t jumps when N_t^g changes. Between times, M_t adjusts continuously. See Figures 2(a) and 2(b).

Consider an upward jump in N_t^g . That is, there is a policy shock: the government purchases paper, paid for by printing money. Looking ahead, this paper will bring in a future stream of additional revenue, which the government will use to retire money. The price of money will therefore rise over time -- equivalent to paying real interest on money.

Hence, at the time of the shock, anticipating the higher future return, entrepreneurs demand higher real balances. (The direction of jump in the price of money is ambiguous, because the demand for real balances may or may not increase as much as the money supply.)

With larger real balances, the liquidity constraint is looser: the liquidity premium and the price of capital jump down, and investment jumps up.

After the policy shock, capital stock starts accumulating, and output rises. Real balances and the price of money also rise. The price of paper falls, because the return on capital falls with the higher capital stock -- and, by Proposition 3(ii), the liquidity premium also falls.

The expansion continues until the next policy shock, when the government reduces its paper holding.

Overall, when the government uses the return stream from its paper purchase to retire money, open market operations lead to persistent expansion in investment and output. The liquidity premium (the nominal interest rate), and the price of paper, are countercyclical, whereas real balances are procyclical.

A simple way to understand these expansionary effects is that the government is acting as a banker to the entrepreneurial sector. It is transforming a partially liquid stream of revenue on paper into a fully liquid stream of interest on money. Being more liquid, the latter income stream is a more effective instrument for funnelling resources from savers to investors.

Interestingly, there is a closely-related policy experiment that we might have considered that gives quite different answers.

Start with the same open market operation: the government purchases paper using money. Now suppose the government were expected to use the revenue stream from its paper purchase to make transfers to the workers. Then a partially liquid stream would be transformed into a nontradeable stream -- workers cannot borrow against their future income. The group of entrepreneurs would be deprived of an income stream which, although only partially liquid, would otherwise help to lubricate their resource allocation. As a result, at the time of the open market operation, the liquidity premium would jump up and investment would drop. The policy would be contractionary!

In other words, we find that the effect of an open market operation depends heavily on what the government does next: how it spends the additional stream of revenue from its paper purchase. This perspective is reminiscent of Lloyd Metzler's work in the early 1950's.

It may help to think of the initial open market operation as being akin to the government simply expropriating paper from the entrepreneurial sector as a whole. After all, the initial injection of money (used to pay for the paper) is neutral. Whether expropriation by the government is expansionary

or contractionary depends on what the government does with the additional revenue stream.

Of course, what we would like to do is to look at a world with productivity shocks and active government policy -- i.e. where the government pursues a monetary policy rule that reacts to the state of the economy. Agents have rational expectations and know the government's policy rule. A number of classic questions could be then answered. For example, if the objective were to stabilise some weighted combination of output and inflation, what kind of monetary rule would be needed? And what would be the implied interest rate policy?

Our model is well suited to answer such questions, but unfortunately it is hard to analyse active policy by hand. We have recently started work on a calibrated version of the model.

Assessment

This is a good point to step back and assess the model.

Everything hinges on the liquidity constraints, so let's start with the two θ 's.

θ_1 relates to the borrowing constraint. This is central. As I have said, if there were no borrowing constraint, investment would be self-financing and the liquidity of agents' asset portfolios wouldn't matter. θ_1 is by now a standard kind of assumption in the literature on credit constraints in macroeconomics, and needs no defence.

The only really new, and unconventional, component in the model is θ_2 -- the fraction of an agent's paper holding that he can sell per period. θ_2 captures something that people think is an important measure of the liquidity of an asset: the speed with which it can be sold.

Against the θ_2 assumption is the fact that it is too reduced form. Although we think the underlying idea makes sense -- that it may be difficult

to resell private claims -- θ_2 is nothing more than a peculiar transaction cost: zero for the first fraction θ_2 sold, and infinite thereafter. This is manifestly not deep theory. It is simply a device to differentiate the liquidity of paper from the liquidity of money.

Our next task is to endogenize θ_2 in an interesting way. We hope to be able to make rich predictions based on cross-sectional variations in θ_2 -- across firms, industries and countries. Equally, we hope to be able to exploit the fact that θ_2 may be cyclical. We believe that a model of θ_2 based on adverse selection in the secondary market may allow us to explain the so-called "flight to quality" that occurs during financial crises.

So much for the research questions we wish to pursue in future. For now, the question is: Does a model with an exogenous θ_2 deliver interesting predictions or useful insights?

Let's review the predictions. When θ_1 and θ_2 are less than some critical $\hat{\theta}$, it turns out that money plays an essential role in allocating resources. The model tells us something about what to expect in such a "monetary economy". The return on money is very low, and is dominated by the return on paper. The gap -- the nominal interest rate -- can be sizeable. Despite this, entrepreneurs choose to hold some money in their savings portfolios, because they anticipate facing liquidity constraints when an investment opportunity arrives later on. That is, liquidity constraints are an integral part of a monetary economy.

On the other hand, workers -- who don't have investment opportunities, and so don't anticipate facing liquidity constraints -- won't choose to hold money, or even paper for that matter, because the return on both is too low. (This is provided the shocks to the system are not too large or frequent.)

In contrast to a standard real business cycle model, the model has a feedback from asset prices to quantities: the prices of money and paper both affect the entrepreneurs' flow of funds, which in turn affects their investment. Aggregate investment and output are too low: the economy fails to transfer enough resources from savers to investors because of the liquidity constraints.

By including liquidity constraints, we have taken a step on from the general equilibrium asset pricing model. One can show that asset prices are volatile, and fluctuate with the tightness of those liquidity constraints.

We think all these features are normal to a monetary economy.

The model also tells us something about dynamics and policy. If the government purchases paper in an open market operation and then uses the stream of income to retire money -- pays a dividend on money -- then the economy expands, even in the long run. In effect, the government is acting as a banker, converting an illiquid stream of income on paper into a liquid stream of income on money. By contrast, if the same initial open market operation were followed by the government using the stream of income on paper to pay for additional expenditure, the effect would be the opposite: the economy would shrink.

There is a nagging worry that, although qualitatively these predictions look reasonable, the effects may not be quantitatively significant, despite the feedback from asset prices to quantities. In practice, an open market operation constitutes a tiny change in the composition of asset holding in the economy, so it is difficult to see why this change should have significant effects. The answer may lie in a more layered model of banking, where the government supplies extremely liquid assets for banks to use, who in turn supply somewhat less liquid assets for use by the rest of the economy. We conjecture that the effects of government policy may be amplified in such a multi-layered model.

Another source of amplification would be to have chains of credit, where default or delay at one point in the chain causes damage further along. I will talk about this in tomorrow evening's lecture. Notice that in the present model, there is no default or delay in meeting payment obligations.

Concluding remarks

I started my lecture this evening with a discussion of the different

ways economists think about money. Let me end by asking: How does our paper fit in?

It has been said that there are two ways of getting fiat money into a model. One is to endow money with a special function -- for example, cash in advance. The other is to starve the agents of alternative means of saving. This happens in the original Bewley and Townsend models, and in most overlapping generations and matching models.

Models of Money

Special Role for money	Starvation of Alternatives to money
cash-in-advance	Bewley, Townsend overlapping generations matching

Implicit in this two-way taxonomy is the idea that both ways flawed. The crime on the left is in shutting down a market for the direct trading of certain pairs of commodities -- e.g. goods against bonds. The crime on the right is in shutting down a market for direct trading between certain agents. Arguably, this second crime is the lesser of the two, because one can justify why a certain pair of agents may not be able to trade by assuming that they are separated in time or space. E.g., in an overlapping generations model one cannot trade with the unborn; in a matching model one cannot trade with someone outside one's own match.

How guilty of these crimes are we? I think we are innocent of the first crime. Remember that money only has value in our model if θ_1 and θ_2 are below some critical value $\hat{\theta}$. So money is not a logical necessity. We are not imposing a special role for it. Indeed, we can say something about

why and when money might eventually stop being used. Ours is a model of liquidity in advance, not cash in advance.

What of the second crime? Unlike in the early Bewley, Townsend, overlapping generations and matching models, in our model agents do have an alternative to money as a means of saving: there is private paper. They are not starved. Admittedly, we have restricted the liquidity of this paper, but then that was central to our purpose. Our goal was to take a context where different assets have different degrees of liquidity, to examine the behaviour of liquidity premia, to understand the interactions between asset prices and aggregate activity, and to examine policy in dynamic context.

We believe that one of the strengths of our model is that it is in many respects Walrasian. There are markets between all pairs of commodities and all pairs of agents. This is what brings our model close to the real business cycle model.

A criticism of the model as presented this evening is that there is money but no government bonds. In fact, though, the model hangs together just fine if money is reinterpreted as government bonds. Nothing substantive changes. But such a reinterpretation does presuppose that government bonds are as liquid as money. This is an old question: where do government bonds lie in the liquidity spectrum?

Finally, let me mention a line of enquiry related to the one I have been discussing this evening. In my slides, I assumed that the technology for producing output exhibits decreasing returns in aggregate. I waved my hands a little about the possibility of some missing factor of production, such as labour. In the written paper you have, we are explicit about workers.

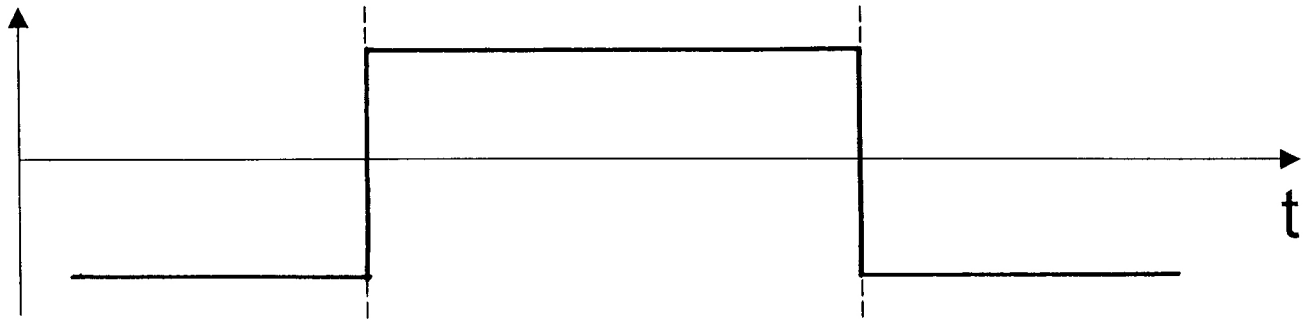
An interesting alternative is to model the missing factor of production as a second capital good, with its own degree of liquidity. Suppose the second capital good is something tangible, like land, or the assets of a well-established old-economy firm. Arguably, the θ for such assets may be closer to 1. In which case, claims on the income stream that the second capital good generates -- equity, or bonds issued by a land bank -- may be

used as money. Non interest-bearing fiat money would be driven out, and the crucial liquidity margin would then be between the less liquid, low- θ , capital good, and the more liquid, high θ , capital good. This two-capital model is the subject of a companion paper. In it we discuss how the government might manage liquidity more generally, other than at the narrow money/bonds margin.

My expectation is that over the next few years theories in which real assets serve as money, and assets are distinguished by their degree of liquidity, will assume a greater importance than theories of fiat money -- not least because cash may start to disappear. As I suggested yesterday, Monetary Economics may be displaced by Liquidity Economics -- which is what I guess Keynes and Tobin would want.


Figure 1(a)

a_t aggregate productivity (exogenous)

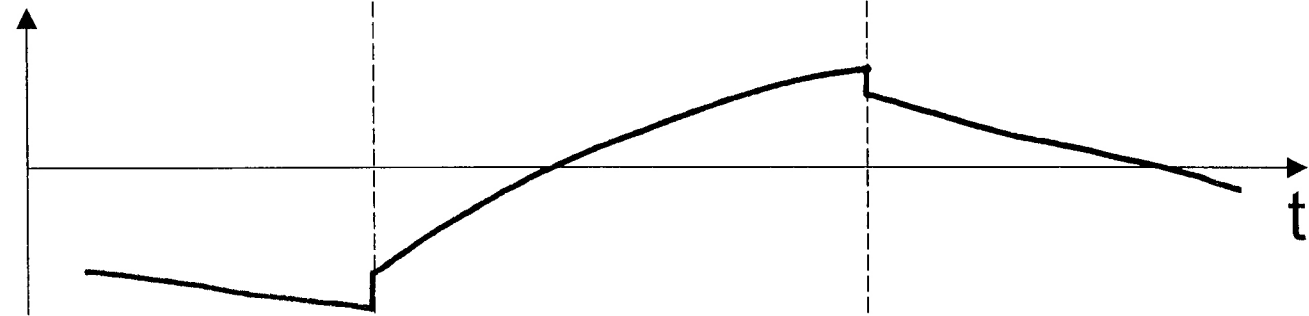


Q_t price of capital

$\propto \pi(q_t - 1) \cong$ liquidity premium (nominal interest rate)



$p_t M$ real balances



I_t investment

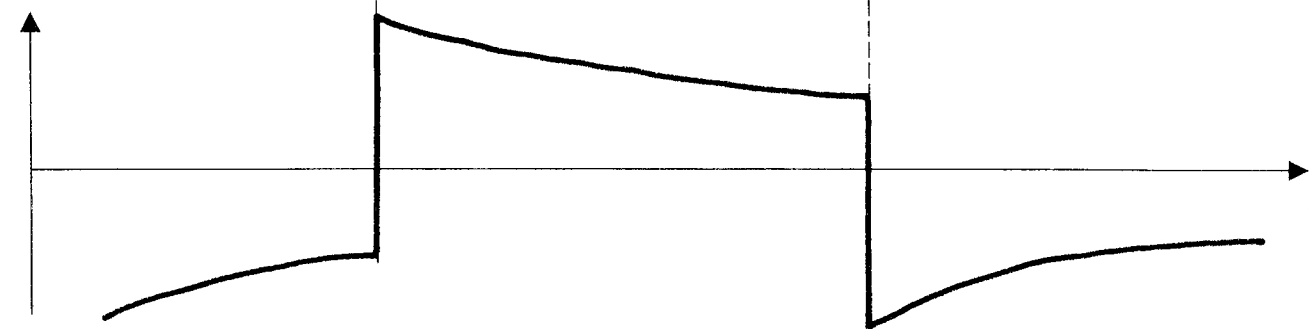




Figure 1(b)

K_t capital stock



Y_t output



C_t aggregate consumption


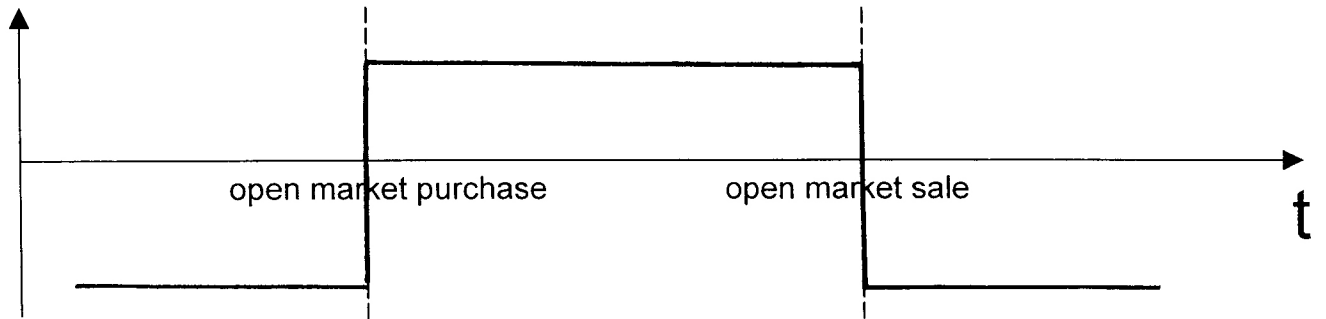



Figure 2(a)

N_t^g government paper holding (exogenous)



$\ln M_t$ money supply



$\ln p_t$ price of money



$\ln(p_t M_t)$ real money balances




Figure 2(b)

Q_t price of paper

$\propto \pi(q_t - 1) \cong$ liquidity premium (nominal interest rate)

