

Clarendon Lectures

Lecture 1

EVIL IS THE ROOT OF ALL MONEY

by

Nobuhiro Kiyotaki  
London School of Economics

and

John Moore  
Edinburgh University and London School of Economics

26 November 2001

It is an honour to have been invited to give the Clarendon lectures. Thank you. My lectures are based on joint research with Nobu Kiyotaki of the L.S.E. I was originally scheduled to perform here a year ago, but we were still in the midst of our research then. You know the old saying: research is like sausages. The finished article can be delicious. But people don't want to see what goes into the making of it. Had I told you about our research last November, it would have been just offal. So I asked Andrew Schuller of O.U.P. to postpone the lectures, and he kindly agreed. The extra time has allowed us to develop the research, especially for tomorrow's lecture. The sausages are now seasoned, and ready for the table.

The overall title for the lecture series is "Money and Liquidity". But let me say straight away that, of these two, we think liquidity is the key. Money -- that is, non-interest-bearing fiat money -- is merely the consequence of a liquidity shortage. As I shall explain later, money is not a logical necessity. Indeed, money may eventually disappear. It may be driven out by ultra-liquid, privately-issued securities that earn interest. In our view, Monetary Economics should be displaced by Liquidity Economics.

That said, it is useful to start with money. From the title of this evening's lecture, "Evil is the Root of all Money", some of you may have come expecting me to talk about morality as well as money. Well, you won't be disappointed. I will. The ratio of morality to money will be low, but the title is apt. It expresses what we think should be at the heart of a theory of money.

I should put my cards on the table. I am a microeconomist. So please forgive my temerity in coming before you to lecture about these macroeconomic topics: money and liquidity. Tomorrow evening's lecture will touch on government, and monetary policy. Wednesday's lecture concerns systemic risk, and the role a government might play in supplying liquidity to avert a financial crisis.

My only defence for venturing into these areas is that my collaborator, Nobu Kiyotaki, is a distinguished macroeconomist. It is a great privilege to work with Nobu. It is very good fun too. Research with him is an example of those curious technologies for which input and consumption are one and the

same thing. Together, he and I are engaged in a pincer movement: he brings his money-macro expertise, and I bring my experience of working with Oliver Hart on matters to do with power and control in financial contracting.

Money. Money is strange stuff. Take these Scottish pound notes. They're useless: they have no intrinsic value. So why should anyone be willing to hold them? That is the classic question economists ask about money. The answer seems obvious: people find it difficult to barter. I don't offer my dentist an economics lesson in exchange for fixing my teeth. It's hard for people to find a "double coincidence of wants". Instead, they use money to buy goods. The dentist accepts money, not because she wants money as such, but because she anticipates that she can use it later to buy what she does want. Money is the medium of exchange.

Notice that for this argument to hold together, there has to be set of mutually-sustaining beliefs, stretching off to infinity. I was willing to hold money yesterday because I believed the dentist would accept it today. She is willing to hold money today because she believes someone else will accept it tomorrow. And so on. If there were a known end-point to history, the entire structure of beliefs would collapse back from the end.

Nevertheless, with infinity on our side, we have arrived at the classic answer to our classic question: people are willing to hold money because it helps them do business with each other. It is the oil that lubricates the economic machine.

Unfortunately, this classic answer to our question about money is completely at odds with the classic answer we give to nearly all other questions in economics! One of the most useful tools of our trade is the notion of perfect competition. In a perfectly competitive framework, there are no frictions to impede trade, so we don't need money as a lubricant. The story goes like this: the dentist, the economist, and everyone else, get together in a marketplace, and deals are conducted through an auctioneer. In the pristine world of perfect competition, it doesn't matter whether there is a lack of coincidence of wants between any two people, because people don't trade in pairs. Rather, everyone trades with the auctioneer, who ensures that supply matches demand. In such a world, money isn't needed, because the

economic machine runs without frictions.

To put this in historical context, the core paradigm in economics is general equilibrium theory -- a theory that can be traced back to Adam Smith, and which has been refined by generations of economists, reaching its fullest expression in the work of Kenneth Arrow and Gerard Debreu, in the 1950's. A great deal of modern economics rests on the Arrow-Debreu foundations. Their framework is justifiably regarded as one of the highest achievements in the science. Yet, remarkably, there is no role for money in the Arrow-Debreu theory of perfect competition.

"So what?", you may ask. "Whoever thought that markets were perfectly competitive in the first place? In the real world, aren't there lots of frictions that impede trade?"

Well, some of the best brains in the profession have succeeded in building physical trading frictions into their models of the market. Their theories are ingenious and beautiful. But, regrettably, models of trading frictions usually require a lot of special tricks. With notable exceptions, the models are too rarefied to contribute to mainstream debate. Moreover, to us, it's not clear that physical trading frictions are really essential to monetary theory.

For John Maynard Keynes, the role of money was as central to economic theory as it was to economic policy. Money was the branch of the subject where people held views with religious fervour. In fact, money and religion have much in common. They both concern beliefs about eternity. The British put their faith in an infinite sequence: this pound note is a promise to pay the bearer on demand another pound note. Americans are more religious: on this dollar bill it says "In God We Trust". In case God defaults, it is countersigned by Larry Summers.

Just as religion has sparked some of the worst conflicts in history, so the subject of money has led to some of the fiercest battles in economics. John Hicks wrote in the 1930s that "... it is with peculiar diffidence and even apprehension that one ventures to open one's mouth on the subject of money." The battles continued through to the disputes between the Keynesians

and the Monetarists in the 1970's and 80's.

Things have gone suspiciously quiet now. Monetary theory has gone into the wings. In the drama of modern macroeconomics, money plays only a bit-part. For most of the actors, monetary theory is merely a side-show compared to the high drama of value theory.

Nobu and I think this is wrong. The flow of money and private securities through the economy is analogous to the flow of blood. In the body of the economy, prices are the nervous system, signalling the needs of different parts of the body. Money is the blood that dispatches resources in response to those signals. No doctor would be content with a model of the body that ignored the flow of blood. Value theory and monetary theory need to be integrated.

We decided early on that the best way to think about money is not to think about money. There is no point in assuming what you are trying to explain. Better to build a model in which something naturally emerges that plays the role of money. We decided to focus on the circulation of private IOUs.

Let me explain. A pound note -- cash -- is only one particular kind of money. Cash is known as "outside money", because it is issued by the government, and the government is outside the private economic system. But there are many other forms of money, that come from inside the system. Let's go back to the dentist. Suppose I pay her by debit card, instead of cash. To keep things simple, let's say that she and I happen to hold accounts at the same bank. I find it clearest to suppose that she and I both have nothing in our accounts at the start of today. When I pay her by debit card, funds are automatically transferred from my account to hers. That is, by the end of today, after she's fixed my teeth, I owe the bank because I'm now overdrawn, and the bank owes her because her account has a positive balance. In essence, what's happened is that I have given the bank one of my IOUs, and the bank has given the dentist one of its IOUs.

Now: Why didn't I just give the dentist one of my IOUs directly? Why do we need the bank as a go-between? The answer may simply be that the

dentist doesn't trust me to repay my debt. But there is a more subtle answer. My dentist may trust me to repay her -- not least because if I don't she can always threaten to do something nasty to my teeth next time they need her attention. But in the meantime she can't use my IOU for her own purchases, perhaps because no-one else trusts me, apart from her (and maybe the bank). Of course, if she were generally trustworthy, it wouldn't matter. She could simply hold on to my IOU, and at any time issue her own IOUs to make her purchases. But suppose no-one else trusts her either. Then she can't issue her own IOUs. Nor can she endorse my IOU (like adding a signature to a bill of exchange). Because if neither of us is trusted by other people, even the combination of her signature and mine won't be enough to get my IOU to circulate.

In these circumstances, the only way she can make purchases before the repayment date on my IOU is if she's paid with the bank's IOU (and, in return, the bank holds my IOU). She can use the bank's IOU to make purchases at any time, because, we suppose, everyone trusts the bank. In short, she gets more benefit from being paid with the bank's IOU than from being paid with my IOU.

This brings out a central idea. The bank's IOU is used by me and the dentist to lubricate our transaction. Why? Because the bank's IOU can freely circulate around the economy. Like blood, it is liquid. In fact, it is functionally equivalent to cash. But, unlike cash, it doesn't come from outside the private system, it comes from inside. For this reason, bank debt is called "inside money". Quantitatively, inside money dwarfs outside money, by a ratio of around 30:1 in Britain today, depending on how you measure it. That is, circulating private debt is extremely important, much more important than cash.

The reason why the dentist and I have to use the bank's IOU is that, unlike the bank's, my IOUs are not liquid. They are illiquid. They are sticky. They couldn't pass from my dentist on to anyone else. My IOUs are definitely not inside money. They are just plain IOUs.

Inside money is not a modern invention. Nobu and I have been inspired by reading about the medieval economy. In essence, the issues back then were

the same as they are today. And it's easier to see the wood, because there were fewer trees. The economic historian Raymond de Roover's famous study of medieval Bruges detailed the importance of banking. In those days, there were no debit cards. Instead, the dentist and I would walk together to the banker's premises, and both witness as he wrote in his ledger, deducting from my account, and adding to the dentist's.

When I borrow from you, I give you an IOU. To use the jargon, I "issue my paper" to you. Some people's paper circulates and others' doesn't. My paper doesn't circulate. Bank paper does. Typically bank paper doesn't literally circulate in the form of paper notes, but in the form of ledger entries or, nowadays, as entries in a computer database. But there have been historical episodes when private bank paper has circulated as money -- most notably during the free banking era in Scotland in the eighteenth and early nineteenth centuries. And, as you can see from these pound notes, certain Scottish banks still issue notes today.

Before I present a formal model of these matters, let me start with a little three-date example. [SLIDE 1.] There are three days to focus on: today, Monday; tomorrow, Tuesday; and the day after, Wednesday.

Let's think about someone hypothetical called Ian. This morning, Ian woke up with a great idea for an investment project. Draw the project like an aeroplane taking off. Ian's project is long-term. It's not going to come to fruition tomorrow, but in two days time, on Wednesday. The aeroplane lands. Two days may not sound very long-term to you. If you don't like days, then think of years or decades. Assuming Ian's project is a good one, it should land on Wednesday with more funds on board than when it takes off today.

Unfortunately, Ian doesn't have the funds needed to get the full project off the ground. So who will lend to him? Perhaps Jim. Jim does have funds today -- one of his projects has just been completed. But there is a difficulty. Jim is only willing to lend short-term -- in fact, overnight -- because he has an idea for another project that he wants to start tomorrow. And so he doesn't want to tie up his funds any longer than one night, tonight.

There is thus a lack of coincidence of wants between Ian and Jim, just as there was between me and my dentist. But whereas between me and my dentist the lack of coincidence of wants was over the types of good being offered, here it is over time -- the times of giving and receiving. Today, Ian wants to borrow long-term, but Jim only wants to lend short-term.

What we are doing in this example is to recast the classic idea of lack of coincidence of wants from the type dimension to the time dimension.

If the world comprised just Ian and Jim, they would be in a pickle. But there is a third actor in the drama, Kevin. He has no funds today. But he does have a project that he started yesterday. And this project will pay out funds tomorrow, Tuesday. Kevin wants to save those funds until Wednesday, when his next project starts up.

They form a curious trio. Ian needs funds today, but won't get any until Wednesday. Jim has funds today, but needs them back tomorrow. Kevin will have funds tomorrow, and won't need them until Wednesday. No pair of them has a coincidence of wants. But collectively they could do business. The efficient allocation would be for the person whose project is finishing on any given day to hand over funds to the person whose project is starting. In an Arrow-Debreu world, this would be the outcome: today, the three parties would each contract, through an auctioneer, to implement the efficient allocation. In fact, the Arrow-Debreu market need open only once, today, Monday.

If you're not keen on the idea of a centralized auction, then think instead of a decentralized marketplace today where people write bilateral contracts. Kevin contracts to deliver to Jim tomorrow. In exchange, Jim pays funds to Kevin today. Kevin uses these funds to pay for a contract from Ian promising to deliver on Wednesday. Both these deals are bilateral, and agreed on today. There is no need for more deals to be struck tomorrow or on Wednesday. The marketplace doesn't need to reopen. In particular, there is no need for Ian's paper to change hands: Ian borrows from Kevin today, and he pays Kevin back on Wednesday.

However, there may be a problem. Who is going to enforce these contracts? Although today Ian may promise to hand over funds on Wednesday, can he be trusted to do so? Is his promise credible? And can Kevin be trusted to deliver tomorrow? They certainly look an untrustworthy bunch.

Notice that the question of trust arises so starkly in this example because we have switched from the type dimension to the time dimension. The crucial twist that time introduces is that if I borrow from you today and then, later on, I fail to repay you, at that point you can't undo the initial loan. Time is irreversible. In the 1920's the Cambridge mathematical astronomer Arthur Eddington coined the phrase "time's arrow". Well, given that time's arrow can't fly backwards, in economic relationships the question of trust is critical.

Surprisingly, in Arrow-Debreu the time dimension is treated on a par with the type dimension. Trust is ignored. Implicitly, it is assumed either that all economic agents are entirely trustworthy, or that the auctioneer can wield a stick that is so big no-one dare renege on a promise.

Nobu and I think that factoring in a lack of trust -- placing a limitation on the degree of commitment -- is of primary importance. In particular, we think that it is the right starting point for a theory of money. Hence the title of this evening's lecture: "Evil is the Root of all Money".

Evil is a strong word. You may find the moral category too severe for something as mild as breaking a promise. In which case, you may want to change the title to "Distrust is the Root of all Money". But that wouldn't have quite the same ring.

To get back to Ian, Jim, and Kevin. If they can't trust each other, and there are no mechanisms available to enforce promises, then they are doomed to autarky. That is, each will have to do his own thing, as best he can. Ian will have to scale down his project if he can't borrow today. Jim will stuff his funds under his proverbial mattress overnight tonight. And Kevin will stuff his funds under his mattress tomorrow night. All this is highly inefficient. Funds shouldn't be stuffed under mattresses, they should

be put to good use.

Is there any remedy to this sorry state of affairs? Well, even though they don't trust each other, there may be a mechanism to enforce promises.

Remember that Ian is starting a new investment project today. We might reasonably suppose that his project has assets. But assets can be mortgaged! Although Ian personally can't be trusted, he can offer the assets as collateral. So he may, after all, be able to commit to pay out at least part of his Wednesday return, by issuing paper secured against the assets. The important distinction here is between human and physical capital. Ian's human capital is inalienable, but his physical capital can be seized in the event of default. This makes his pledge credible. The project is in part self-financing. Ian can borrow today to finance a bigger project.

What about Kevin and Jim? Can they also borrow today? Presumably not. For Kevin, it may be too late to borrow against tomorrow's return. Yesterday he may have borrowed as much as could, to invest in a bigger project, and he has no spare collateral today. As for Jim, his next project doesn't start until tomorrow, and people may be unable to "pre-mortgage" their future projects.

Let's agree that only Ian can issue paper today, pledging part of his Wednesday return. Think of it as long-term paper, given that it doesn't mature until two days time.

The question is: How might Ian's paper help the three of them to do business with each other? On the face of it, it doesn't help. Jim doesn't want to hold the paper, because he wants to be repaid tomorrow, and can't wait until Wednesday. And Kevin can't buy the paper today, because he hasn't got any funds readily available and he can't borrow. That is, even though Ian can issue paper, there remains the problem that no pair of them has a coincidence of wants in dated goods.

But there is solution. Suppose Ian's paper circulates -- from Ian, to Jim, to Kevin, and finally back to Ian. Then they can do business with each other, as follows. Today, Ian borrows from Jim. That is, Ian sells his

paper to Jim -- even though Jim doesn't really want to hold an IOU promising to repay on Wednesday. Never mind. Because when tomorrow comes, Jim can sell the paper on to Kevin. Kevin is happy to buy this "second-hand" paper because he wants to save over tomorrow night. On Wednesday, Kevin takes the paper to Ian, who redeems the IOU. It is important to see that, although Ian borrows from Jim, Ian repays Kevin.

Taken together, the two short valley-shaped lines represent the journey taken by Ian's paper. They're drawn in red, to remind us of the metaphor of the circulation of blood.

Ian's paper provides the economy with liquidity, the means of short-term saving: Jim uses it to save overnight tonight, and Kevin uses it to save overnight tomorrow night. The paper is functionally equivalent to money. Jim buys it today not because of its maturity value but because of its exchange value. Notice how we have arrived at a theory of money through the back door. We haven't explained money by assuming money. Instead, we have shown how this little economy can work well if private debt circulates, serving as inside money.

Another way to think about this is to use the language of commitment. In general terms, what we have shown is that if two people want to transact who don't trust each other -- who are unable to commit -- then they can make use of a third party's ability to commit. The third party's commitment power acts as a lubricant to the transaction. In our example, Ian is the "third party" to Kevin and Jim. Today, Kevin can't commit to deliver to Jim tomorrow. No matter. Because, thanks to Ian's commitment power, Kevin doesn't need to commit. When tomorrow arrives, Kevin simply engages in a spot transaction.

Roughly speaking, we might say that Ian is acting as a banker to Kevin and Jim. Deep down, of course, Ian is no different from the other two. The only reason why he can act as a banker is that today he has spare collateral.

I want to draw out two central ideas from this example. First, since Ian supplies the economy with money, the tighter is Ian's borrowing constraint, the more likely there is to be a shortage of liquidity. See the

paradox: Although it is Ian who is doing the borrowing, nevertheless he is the one who is supplying the economy with money. That's the peculiar nature of inside money -- someone's debt circulates.

Second, for Ian's paper to act as money, resale must be possible. The paper must be negotiable. It is not enough that Ian can work within his borrowing constraint to sell his paper today, to Jim. It must also be possible for Jim to resell the paper tomorrow, to Kevin. Unlike in Arrow-Debreu, the market must open twice: for the initial sale today, and for the resale tomorrow.

Let me put this second idea another way. For Ian's paper to circulate as inside money, he must be able to make not merely a bilateral commitment to the initial purchaser (Jim), but a multilateral commitment to any subsequent bearer of the paper (Kevin).

The distinction between bilateral and multilateral commitment is so important to us that I'd like to step away from the example for a few minutes.

It seems clear to us that multilateral commitment is a lot more demanding than bilateral commitment. And this has implications for which kinds of paper can circulate, and which can't.

Start with private debt. I can borrow from Nobu, given our close working relationship, but he would have difficulty passing my debt on to a third party. Earlier I gave the example of my dentist, who would be unable to use my IOUs for her own purchases.

More generally, when a supplier extends trade credit to her customers, she has special leverage over them because they will need to buy more from her in the future. In practice, suppliers have difficulty offloading this kind of debt: it's not easy selling trade credit to a third party at a fair price.

When a bank extends overdrafts to its customers, it too has special leverage over them because they need to keep the bank sweet for the future.

Banks have difficulty reselling this kind of debt. And the problem is exacerbated by the problem of asymmetric information. The bank has a good idea which of its customers are safe, and which dubious. But this knowledge is private. In the secondary market, it's in the bank's interest to resell off what it knows are the bad loans. Potential buyers are aware of this, and may be wary of buying -- to the point that the market can collapse.

In our written paper, we tell a different kind of story to rationalise why paper may not resaleable. Our story is one of moral hazard. When an entrepreneur issues paper secured against a project, he gives the initial creditor some control over access to the project. Now suppose the creditor were planning to resell the paper before the project matures. Then she and the entrepreneur could collude to strip the project, leaving just a shell that delivers no output, but which cannot be distinguished from an intact project by outsiders. Aware of this possibility, no-one will be willing to buy the paper in the secondary market.

Whatever the story -- be it special leverage, adverse selection, or moral hazard -- the conclusion is broadly the same. Namely: it may not be possible for an initial creditor to resell paper. Or perhaps the paper may take a long time to resell. And even if there is an active secondary market, the price may not reflect the true value of the paper to the initial creditor.

In this evening's lecture, I want to consider just the two extremes. Either paper can be resold -- it's liquid, and can circulate. Or it can't be resold -- it's illiquid, and can't circulate. In tomorrow evening's lecture, we will look at intermediate cases, where paper is partially resaleable.

Incidentally, in our example, if Ian's paper couldn't be resold, then there would be no inside money and trade would completely break down. The three of them would be back to autarky. In general, the outcome need not be that bleak. But to achieve efficiency, it's clear that the circulation of inside money is crucial.

Before I turn to the full model, let me stress again the two ideas that came out of this example. First, any constraint on Ian's ability to borrow

today matters, because it is he who supplies the economy with liquidity. Second, any constraint on Jim's ability to resell tomorrow matters, because unless Ian's paper circulates it cannot act as money.

These two potential constraints need to be thought about separately. The first constraint, a borrowing constraint, has received attention in the macroeconomics literature. Any number of moral hazard stories can be invoked to rationalise why people face borrowing constraints. The second constraint, a resaleability constraint, has received much less attention in the formal literature, but we think is just as important.

In the full model, it is interesting to see how these two constraints feed into each another. On the one hand, remember that in this three-date example if there were no borrowing constraints, resaleability wouldn't matter -- just as in Arrow-Debreu, where, because people never renege on their promises, paper doesn't need to circulate. More generally, we'll see that even though there may be less than full bilateral commitment, if there is enough, then multilateral commitment isn't needed; the economy works well without inside money.

On the other hand, we'll also see that if paper does circulate as money because there are no resaleability constraints, then the economy can work well even though people may not be able to borrow very much. A little multilateral commitment goes a long way.

We've spent a long time on the three-date example. But fortunately we are 90% of the way to the full model. Here it is. [SLIDE 2.]

Spot the differences between this figure and the last. First, there is an infinite time horizon, because I want to consider an economy in steady state. There is one homogeneous good at each date. The good can be stored, i.e. stuffed under a mattress. I've made every Sunday a day of rest. What you see in front of you, then, is a typical six-day working week.

Now there are many Ians, many Jims, and many Kevins. In fact, there is a continuum of each, with total measure 3. What a terrifying thought. The economy is competitive: there are no trading frictions.

Actually, the names here are arbitrary, because the entire population is homogeneous. Everyone can choose when to start their production cycle. We focus on the symmetric equilibrium where start-times are staggered evenly. Each investment project takes three days, from start to finish inclusive -- i.e. two nights. No-one can operate overlapping projects.

There is no uncertainty, either in the aggregate or at the individual level. We hope that one of the strengths of the model is that it can be used to discuss money and liquidity in a deterministic setting. We have come to think that money and liquidity may not have anything inherently to do with uncertainty. Rather, they are to do with the twin constraints -- borrowing and resaleability.

However, there is a serious downside to applying Occam's razor to get rid of uncertainty: one cannot sensibly talk about business cycles. In case you feel put off by the stylised nature of today's model, let me reassure that in tomorrow's lecture I will present a fully stochastic model which can be used to think about liquidity and monetary policy in the presence of shocks.

How should we model the crucial borrowing and resaleability constraints?

Start with the borrowing constraint. Remember the idea is that someone -- let's call him an entrepreneur -- borrows to finance a new project by issuing paper secured against the project's assets. The assets serve as collateral. That is, if the entrepreneur defaults ex post, then his creditor can seize the assets and liquidate them. But without the entrepreneur's specific human capital, the return from liquidated assets will be lower -- let's say only a fraction  $\theta_1$  of what it would have been had the assets not been seized. In effect, the creditor has an outside option worth  $\theta_1$  times the "inside return".

Now suppose the entrepreneur can always push the creditor's payoff down to this outside option -- no matter what formal contract has been written. Then, ex ante, the entrepreneur can't credibly promise to repay more than a

fraction  $\theta_1$  of the project return. [SLIDE 3.]

$\theta_1$  = fraction of a project's return that can be mortgaged,  
by selling paper at the time of investment

We assume that people cannot raise funds at any other time, and cannot mortgage future projects.

Because production takes two nights, the paper will mature two days after it has been issued. This choice of time horizon is deliberate. It is the shortest horizon we can get away with in order to model liquidity. If production took only one night, paper would mature the day after it is issued, and the question of resaleability wouldn't arise. With production taking two nights, we can ask: Can paper be resold when it is "middle-aged" -- i.e. on the middle day? We want to consider both possibilities. Let  $\theta_2$  be an index of resaleability.

$\theta_2 = \begin{cases} 1 & \text{if paper can be resold the day after investment} \\ 0 & \text{if paper cannot be resold} \end{cases}$

If  $\theta_2$  equals 1, paper can be resold. If  $\theta_2$  equals 0, paper can't be resold.

So:  $\theta_1$  corresponds to the borrowing constraint; and  $\theta_2$  corresponds to the resaleability constraint. They are the heart of the model. By the way, the mnemonic here is that the subscript 1 on  $\theta_1$  denotes the initial sale of paper, and the subscript 2 on  $\theta_2$  denotes the resale, a day later.

The rest of the model is quite standard. [SLIDE 4.] We assume that everyone consumes every day, has a logarithmic utility function, and a common

discount factor  $\beta$ :

$$\sum_{s=0}^{\infty} \beta^s \log c_{t+s} \quad \text{where } 0 < \beta < 1.$$

Also, the cost of a project is strictly convex in the output  $y$ . Per capita:

$C(y)$ , the cost of producing  $y$

$$\equiv y^\lambda \quad \text{where } \lambda > 1.$$

We like to think of this as an economy populated by lots of fast-moving "ants" -- ants called Ian, Jim and Kevin, I guess. So we think of the discount factor  $\beta$  as being close to 1. Also, we think of the technology as having close to constant returns, so that  $\lambda$  is close to 1 too. In what follows, when I say "approximately", I refer to the case where  $\beta$  and  $\lambda$  are both close to 1.

To provide a benchmark, let's look at a first-best allocation in steady state. [SLIDE 5.] Productive efficiency requires that the marginal cost of investment equals the discounted marginal return. Since production takes two nights, this means that, per capita, the efficient level of output,  $y^*$ , satisfies:

$$C'(y^*) = \beta^2.$$

Also in the first-best, consumption is perfectly smoothed. That is, everyone has the same per capita consumption,  $c^*$  say, irrespective of what point they are in their individual production cycles.  $c^*$  is given by

$$c^* = [y^* - C(y^*)]/3.$$

As I promised, we want to use this model to see how the borrowing constraint and the resaleability constraint feed into each other. To do so, let's start with the case where paper can be resold, i.e. where  $\theta_2$  equals 1. [SLIDE 6.] Just as in our earlier three-date example, resaleable paper provides the economy with liquidity. The paper acts as inside money.

Take one of the Ians, on a Monday. He borrows long-term -- by selling new paper to some Jim in the competitive market. The next day, Tuesday, Jim resells Ian's paper, which is now second-hand, to some Kevin, in the same paper market. Notice that because there is no uncertainty, new and second-hand paper are perfect substitutes as means of short-term saving. The next day, Wednesday, Kevin takes the paper back to Ian, who redeems it. Each day, the three of them rotate roles. As you can see, the whole diagram is filled with red lines. Think of this as a "red economy", because the paper is like blood: it is red and circulates. It is inside money.

Thanks to money, the first-best can be attained relatively easily. [SLIDE 7.]

#### Proposition 1 (red economy)

If paper is resaleable ( $\theta_2 = 1$ ), then the first-best is attained iff  $\theta_1$  is greater than approximately 1/3.

The reason why paper is resaleable -- why  $\theta_2$  equals 1 -- is that agents are able to make multilateral commitments to repay any bearer of their paper. Proposition 1 confirms something I said earlier: A little multilateral commitment --  $\theta_1$  as low as 1/3 -- goes a long way. It's enough to attain first-best.

However, if there is no multilateral commitment, if paper can't be resold, then matters are very different. We saw before that the three-date example collapsed to autarky if Jim couldn't resell Ian's paper to Kevin.

Nobu and I thought that the same would be true in this stationary model. After all, there appears to be a complete lack of coincidence of wants in dated goods: borrowers want to borrow long-term, over two nights; whilst savers want to save short-term, over only one night. But without resaleability -- without money -- all deals have to be bilateral. This is why we thought that, in the apparent absence of coincidences of wants, there could be no gains from trade.

We were wrong! When  $\theta_2 = 0$ , the infinite-horizon economy doesn't collapse to autarky. Instead, something much more interesting happens. [SLIDE 8.] Even though paper is now illiquid, the economy finds a way of creating coincidences of wants. What happens is this. Consider one of the Jims. On a Monday he lends to some Ian. That is, Ian issues an IOU, which Jim has to hold through to Wednesday because it is illiquid and can't be resold on Tuesday. On Wednesday, when the debt is due, Ian repays Jim. Jim promptly lends to some Kevin, who is starting a project. That is, Kevin issues an IOU, which Jim holds through to Friday. The two long valley-shaped lines represent Ian and Kevin's paper. They're drawn in blue, to denote illiquid paper. On Friday, when Kevin's debt is due, he repays Jim, who can now invest in his own project! Jim's project completes on the following Monday. In effect, Jim's Monday/Wednesday/Friday/Monday budgets are linked -- by holding blue paper (twice) and investing (on Fridays).

It is important to realise that Jim doesn't miss his Tuesday investment opportunity -- because his Saturday/Tuesday/Thursday/Saturday budgets are linked too. However, there is no link in the budgets of consecutive days. Think of two parallel turnpikes, with no cross roads to join them.

We could go on and fill up the diagram with blue lines like these. As they say in cookery programmes, here's one I prepared earlier. [SLIDE 9.]

In this economy, paper is illiquid and so cannot circulate, but the economy is making the very best of a bad job. Let's call this a "blue economy", to distinguish it from the earlier "red economy" where paper was liquid and could circulate. The blue economy has no inside money to act as a lubricant. And yet, ingeniously, the economy manages to create coincidences

of wants, where savers as well as borrowers are willing to use the illiquid paper. I must stress that the ingenuity here is not ours. It is the economy that is ingenious. The economy succeeds in finding gains from trade. I think this is a great example of Adam Smith's invisible hand at work.

But there is a problem. Just compare the blue economy with the red economy. You can tell at a glance that the blue economy uses much more paper than the red. If you think about it, the demand for paper in the blue economy is four times as great as it is in the red economy.

This puts a greater strain on the paper market. For the economy to run efficiently, there has to be a greater supply of paper. Agents must be able to issue more paper when it is their turn to invest. That is,  $\theta_1$  must be higher. We have the following proposition. [SLIDE 10.]

Proposition 2 (blue economy)

If paper cannot be resold ( $\theta_2 = 0$ ), then the first-best is attained iff  $\theta_1$  is greater than approximately  $2/3$ .

Given that the demand for paper is four times greater in the blue economy than in the red, you may be puzzled as to why the critical value of  $\theta_1$  only doubles. The answer is that, ceteris paribus, doubling  $\theta_1$  simultaneously doubles the supply of paper and halves the demand.

It's interesting to note that the critical threshold for  $\theta_1$  in Proposition 2 is strictly less than 1. It is approximately  $2/3$ . This confirms something else I said earlier: Even though there may be less than full bilateral commitment, if there is enough, then multilateral commitment isn't needed; the economy works well without inside money.

Taken together, Propositions 1 and 2 tell us that when  $\theta_1$  lies between  $1/3$  and  $2/3$ , the red economy attains the first-best, but the blue economy doesn't.

What is the general message to take away from this? Simply that if paper is less liquid, it is less convenient to use, and the economy needs more of it. This increased demand puts greater strain on the paper market. There may be a shortage of paper. I'd like to adopt the phrase "liquidity shortage". For  $\theta_1$  lying between  $1/3$  and  $2/3$ , the blue economy has a liquidity shortage, but the red economy doesn't.

There is not time to give details, but my third Proposition lists some of the symptoms of a liquidity shortage in the blue economy. [SLIDE 11.]

Proposition 3 (blue economy)

Assume  $\theta_2 = 0$  (paper cannot be resold)  
 and  $\theta_1 < 2/3$  ( $\Rightarrow$  liquidity shortage)

Then symptoms of a liquidity shortage include:

- price of two-period paper  $> \beta^2$   
 (implied one-period rate of return on paper  $< 1/\beta$ )
- borrowing constraints bind on day of investment
- consumption is jagged  
 (lowest on day of investment; highest the day before)
- investment and output are lower than in first-best

Moreover, the lower is  $\theta_1$ , the worse are these symptoms

For low enough  $\theta_1$ , there can be inefficient storage, even though

rate of return on storage	< . : : : .	implied one-period rate of return on paper
	liquidity premium	

The shortage of paper causes the price to be high -- which means that the implied rate of return on paper is low. Agents starting new projects would like to borrow more at these low rates, but they face binding borrowing constraints. The other agents are discouraged from saving. There is thus a log-jam in the paper market. Too few resources are transferred from savers to borrowers. As a result, investment is low; and output is low. Roughly speaking, the economy runs "too slowly".

These symptoms are worse in an economy with a lower  $\theta_1$ .

[SLIDE 12.] For low enough  $\theta_1$ , savers are being offered such a low return on their illiquid paper holdings that, at the margin, they resort to storage, i.e. they stuff their mattresses. Storage offers an even lower rate of return, but at least has the virtue of being short-term and hence liquid. Think of the gap between the return on storage and the implied one-period return on paper as a liquidity premium.

When  $\theta_1$  lies between  $1/3$  and  $2/3$ , all these symptoms are experienced by the blue economy, but not the red. In other words, we can blame all these bad things solely on the fact that paper cannot circulate in the blue economy -- that there is no inside money.

Your reaction to all this might be to ask: "Why doesn't the blue economy somehow create inside money?"

Well, one way to create inside money is to add some wealthy agents to the model. Let's add some Scottish lairds, who each own a castle. A laird doesn't necessarily produce anything. But as long as his castle is publicly visible, he will be in a position to make multilateral commitment. He can issue paper secured against his castle. Because of the multilateral commitment, his paper can circulate as inside money -- it is red.

[SLIDE 13.] Here is an economy where blue paper and a small amount of red paper coexist. Consider one of the Jims, on Monday the day he completes a project. A fraction  $1 - \theta_1$  of the project's output has not been mortgaged. He has a choice about what he does with these unmortgaged funds. Either he

can follow what one might call a "fast strategy": buy red paper so as to fund investment on the next day, Tuesday. Or he can follow a "slow strategy": buy blue paper twice in succession, so as to fund investment four days later, on Friday. When blue and red paper coexist, Jim will mix between buying red and blue paper on Monday.

In order for him to be indifferent between the fast and slow strategies, blue paper has to offer him a premium over red paper. That is, the implied interest rate on illiquid blue paper has to be greater than the interest rate on liquid red paper. The interest rate differential, the liquidity premium, is the compensation Jim demands for the inconvenience of holding illiquid paper.

Of course the trick would be to make a profit by buying blue paper and selling red! Imagine setting up your stall at Carfax, offering to lend to private people whose IOUs are illiquid. Since their paper is illiquid, blue, they have to pay you a relatively high rate of interest. Meantime, you raise funds by taking in deposits. Because you are sitting at your stall all day and every day, you become quite a public figure, and your IOUs, the paper you hand to your depositors, is liquid; it's red. So you can get away with offering your depositors a lower rate of interest than you receive from your debtors. You are making a profit, merely by sitting there! Do you know what you are? You are a bank. You effectively transform blue paper into red.

There is a more direct way you might transform blue paper into red: you could simply certify it. That is, you could add your signature underneath the signature of the issuer -- as happens on a bill of exchange. Once certified, the paper can circulate, because your signature -- your guarantee -- is commonly recognized.

As a banker, you don't necessarily produce anything. In this respect, you are on a par with the Scottish lairds. There is a difference, though. Whereas the lairds issue red paper secured against their castles, you issue red paper secured in part against other people's blue paper. Unlike the lairds, you have to work a little to earn your profit -- sitting out there in the cold, rather than ensconced in a castle.

We can incorporate banks into the model. Here is a (very!) rudimentary theory of banking. Dispense with lairds and castles, or people sitting out at Carfax. Instead, let's go back to some Ian, starting a project on Monday. In the blue economy, Ian mortgages a fraction  $\theta_1$  of his Wednesday output by issuing illiquid paper, paper that cannot be resold on Tuesday. In the blue economy, Ian can only make a bilateral commitment -- to the initial purchaser of the paper, Jim. The story is that on Monday night, Ian and Jim could collude to asset-strip the project. And, knowing this, no Kevin is willing to buy the paper on Tuesday.

But now suppose Ian can ring-fence his project in a way that limits the potential for asset-stripping. By erecting a fence, Ian can turn at least part of his paper from blue to red. He is his own bank.

One can think of this fence either literally or metaphorically. Literally, the cost of fencing is strictly convex in the height of the fence. The higher it is, the stouter have to be the posts to hold it up. And the higher the fence, the greater the fraction of Ian's paper that is red. In general, Ian will therefore choose an interior value of  $\theta_2$ , the fraction of liquid paper that he issues.

That, in the briefest possible terms, is our theory of banking. A theory of fencing. I'll be returning to fencing again in my final lecture.

The interesting question arises: does a private banking system create enough liquidity to maximize welfare? We find that, whenever storage is used, the answer is no: left to itself, the banking system is too small. Given that people store, there is too little fencing: at the margin, the direct cost to society of erecting taller fences would be outweighed by the indirect benefits of extra liquidity -- in social terms, banks transform too little blue paper into red.

This is provocative. But I must add a caveat. In arriving at our conclusion, we have used a crude "Modified Golden Rule" welfare criterion. We haven't considered transition dynamics, or distributional issues, which ought to be part of a fully-fledged welfare analysis. So our conclusion is only tentative.

Nevertheless, it suggests a role for government, or a central bank, to support the banking system. I'll return to this theme in tomorrow evening's lecture, when I introduce government money and bonds into a stochastic model.

Banks are not the only way of turning blue paper into red. In May 1970 there was a banking strike in Ireland, which lasted over six months. It was feared that the economy might collapse -- because, without banks, there would be too little liquidity. In the event, even though over 80% of the money supply was frozen, the economy hardly blipped! Do you know why? The pub landlords took over, and started circulating IOUs through their bars. This is a nice illustration of the resilience of an economy, to create alternative supplies of liquidity. Guinness is not only good for you, it's good for the economy. A paper on this episode (and two other banking strikes in Ireland) was published in the Manchester School (March 1978) by Antoin Murphy. As Professor Murphy pointed out: "one does not after all serve drink to someone for years without discovering something of his liquid resources".

We can see the same creativity in medieval times. From the work of the economic historian M.M.Postan, we learn that in the Middle Ages the legal system threw up obstacles to the transfer of debt. To use our language, the legal system artificially prevented blue paper from being red. And what happened? People devised new forms of contract, like the bill of exchange, to wriggle round the law. Notice again, just as in the Irish banking strike, the economic drive to create liquidity is hard to stop. The invisible hand at work once again.

In reality, there is every shade of colour between blue and red. The approach I will adopt in tomorrow evening's lecture is to measure liquidity in terms of the speed with which assets can be resold at a fair price. Measured this way, private debt is the bluest of blue paper, because typically it can never be resold. Trade credit and bank loans are also fairly blue -- in the sense that suppliers and banks can only resell this kind of paper at less than the true value.

A stake in a small firm might be resaleable; but to find a buyer would take time. So this kind of paper is only moderately liquid: think of it as

purple. Corporate bonds and equity are examples of fairly red paper: they can typically be resold quickly, and so are quite liquid.

With the latest innovations, paper that was blue is becoming redder. Mortgages used to be blue, but are now red, thanks to the creation of mortgage-backed securities. I may think that I make mortgage payments to the Royal Bank of Scotland. But in fact my mortgage has almost certainly been bundled with lots of other people's, sliced up in clever ways, and then resold. It has probably been resold many times since. In other words, my mortgage has become like inside money. And thanks to new electronic trading mechanisms, stocks and shares can be sold more quickly than before: they are now redder.

In this evening's model,  $\theta_2$  was the counterpart to the colour of paper. The higher the value of  $\theta_2$ , the redder the paper. I am suggesting that recent developments in financial markets have pushed up  $\theta_2$ 's, and are continuing to do so.

There are reasons to suppose that  $\theta_1$ 's have gone up too. It is argued that loans can now be more accurately targeted at certain groups of borrowers. And, perhaps more importantly, borrowers have more to lose from cheating. Credit scoring has become the norm. If I want to borrow today, potential lenders don't consider so much what I want to use the funds for. Nor do they even consider what collateral I have to offer. Mostly, they look at my credit history, and other aspects of my past life. The reason is that much more information can be stored about my past, which can be used to assess the likelihood of my defaulting. Also, crucially, this information is much more widely available. And that puts a greater premium on my maintaining my financial reputation. It's not that I am inherently any more trustworthy, but rather that the scope for my getting away with things is less. In terms of the model, this means that my  $\theta_1$  is higher. And so the supply of liquidity is higher too -- which, as we have seen, improves the state of the economy.

Let me end by coming back to where I started: these pound notes. Cash. The reddest paper of all. It earns no interest. (With inflation, it actually earns a negative return, but let's ignore that.) In the year 2001, people

hold cash. So, we may conclude that, in the current spectrum of interest rates, zero must be the rate of interest on the reddest paper. But there is no reason why this should always be true. Cash is not a logical necessity.

To see why not, consider this evening's model. Cash -- i.e. non-interest-bearing fiat money -- can circulate alongside inside money. But only if the interest rate on red paper drops to zero, which happens if  $\theta_1$  is low enough. In other words, fiat money can only augment the aggregate stock of liquidity if there is demand for it. Fiat money will play a role in tomorrow evening's model.

In the future, as the  $\theta_1$ 's and  $\theta_2$ 's rise further, so the spectrum of interest rates will shift up. Eventually, it may be the case that no-one will be content to earn zero interest on the reddest of papers. Fiat money may disappear, crowded out by ultra-liquid private securities that earn interest. Before long, I may be using a Merrill Lynch card to pay my dentist in MicroSoft shares.

The thesis that fiat money may disappear is controversial. There are several interesting arguments against it. First, we may be focussing here too much on the supply of liquidity. As the pace of the modern world quickens, people need to respond more quickly to opportunities. To put it grandly, as the time interval in people's lives shrinks, the problem of finding coincidences of wants in dated goods becomes more severe. Our demand for liquidity may be rising in line with the supply.


Next, our discussion presupposes that (fiat) money and other assets are substitute means of saving. It can be argued that in fact money is complementary. After all, assets such as bonds are promises to pay in money. This may be the point to bring back the idea that money lubricates trade in the absence of markets. We may need to model trading frictions after all.

Finally, cash will always be useful to people who want to conceal their nefarious activities, like drug dealers, because cash leaves no electronic trail. If, in due course, crime turns out to be the only reason why people hold money, then evil will still be the root of all money, but for different reasons than the ones I have outlined this evening.

MON


TUES

WED



Slide 1

--- red



Slide 2

$\theta_1 =$  fraction of a project's return that can be mortgaged,  
by selling paper at the time of investment

$\theta_2 = \begin{cases} 1 & \text{if paper can be resold the day after investment} \\ 0 & \text{if paper cannot be resold} \end{cases}$

Slide 3

PREFERENCES:

$$\sum_{s=0}^{\infty} \beta^s \log c_{t+s} \quad \text{where } \beta < 1$$

TECHNOLOGY:

$C(y)$ , the cost of producing  $y$

$$\equiv y^\lambda \quad \text{where } \lambda > 1$$


Slide 4

## FIRST-BEST:

productive efficiency:  $C'(y^*) = \beta^2$

smoothed consumption:  $c^* = \frac{1}{3} [y^* - C(y^*)]$

Slide 5




Slide 6 (red economy)


PROPOSITION 1 (red economy)

If paper is resaleable ( $\theta_2 = 1$ ), then first-best is attained  
iff  $\theta_1$  is greater than approximately  $1/3$ .

Slide 7



Slide 8 (blue economy)



Slide 9 (blue economy)

PROPOSITION 2 (blue economy)

If paper cannot be resold ( $\theta_2 = 0$ ), then first-best is attained  
iff  $\theta_1$  is greater than approximately  $2/3$ .

Slide 10

PROPOSITION 3 (blue economy)

Assume  $\theta_2 = 0$  (paper cannot be resold)

and  $\theta_1 < 2/3$  ( $\Rightarrow$  liquidity shortage)

Then symptoms of a liquidity shortage include:

- price of two-period paper  $> \beta^2$   
(implied one-period rate of return on paper  $< 1/\beta$ )
- borrowing constraints bind
- consumption is jagged  
(lowest on day of investment; highest the day before)
- investment and output are lower than in first-best

Moreover, the lower is  $\theta_1$ , the worse are these symptoms

cont...

Slide 11


PROPOSITION 3 (continued)

For low enough  $\theta_1$ , there can be inefficient storage,  
even though

$$\begin{array}{ccc} \text{rate of return} & < & \text{implied one-period} \\ \text{on storage} & & \text{rate of return} \\ & \uparrow & \text{on paper} \end{array}$$

liquidity premium

Slide 12



Slide 13

--- blue  
 ..... red