

Incomplete:
Conclusions and References missing

CREDIT CHAINS

by

Nobuhiro Kiyotaki[†] and John Moore^{††}

January 1997

[†]University of Minnesota and Federal Reserve Bank of Minneapolis

^{††}London School of Economics and Heriot-Watt University

* This is a major revision of our July 1995 mimeo, "Credit Chains". We have benefitted from feedback from a number of colleagues and seminar audiences. We would particularly like to thank Martin Hellwig and Bengt Holmstrom for their help, and we are grateful to the University of Basel for hospitality. Financial assistance is acknowledged from the US National Science Foundation, the UK Economic and Social Research Council, and the Financial Markets Group at the London School of Economics.

1. Introduction

This paper is a theoretical study into how shocks propagate through a network of firms who borrow from, and lend to, each other. In particular, we investigate how a small, temporary shock to the liquidity of some firms may cause a chain reaction in which other firms get into financial difficulties, thus generating a large, persistent fall in aggregate activity. Also, we use our model to look at the aggregate effects of creditors' postponing the debts of delinquent debtors, rather than liquidating their assets. We show that, although it may be bilaterally efficient for a creditor and debtor in one link of a credit chain to reschedule debt, postponement can have serious social consequences because no new liquidity is injected into the system. Finally, we demonstrate that firms may privately decide not to insure against the risk of default by their debtors, even though such insurance may be available.

The economy we consider has many small firms owned and run by entrepreneurs who are unable to raise finance from outside investors, because they cannot precommit not to default, and because they do not have any collateral to offer potential investors by way of security. However, an entrepreneur can borrow from his suppliers, for the reason that they can withhold their supply if he defaults. That is, suppliers have additional leverage over him because the specific input that they supply to him is necessary for his production.

In our model we will demonstrate that supply contracts have to involve an element of lending, otherwise a supplier will be unable to compete against other suppliers. Since an entrepreneur sells to other entrepreneurs, he is thus forced to lend, even though he is short of funds. In other words, he simultaneously lends to his customers (other entrepreneurs) and borrows from his suppliers (who might be other entrepreneurs or unconstrained firms). Because his balance sheet has financial assets (accounts receivable from his customers), as well as liabilities (accounts payable to his suppliers), he is exposed to the risk of default by his debtors/customers.

Crucially, the entrepreneur cannot net out these gross positions in order to shed the default risk. In particular, he cannot raise money against

the accounts receivable, because only he is in a position to supply his customers with their specific input. If unconstrained third parties -- deep pockets -- were to buy the debt owed to him by his customers they wouldn't have the leverage necessary to get the customers to pay up. As a result, the entrepreneur cannot securitize the debt owed to him. If his customers experience a negative liquidity shock and default, he himself may run into financial difficulties, and he may have to default against his suppliers, thus causing further difficulties along the credit chain. This can be thought of as systemic risk.

To fix ideas, consider the following simple example. Suppose entrepreneur A has ordered 100 units of specific input from entrepreneur B at \$1 a unit. That is, A owes \$100 to B at the date the input is due for delivery, by which time she expects to have \$100 in cash with which to pay him. At the same time, B has ordered 100 units of specific input from entrepreneur C at \$1 a unit: B owes \$100 to C on delivery. B has no cash of his own, but expects to use the \$100 that A owes him to pay his debt to C. These may be two links in a longer credit/supply chain, which in turn might be part of a larger network.

Now suppose that A finds her cash holding is only \$60, rather than the \$100 she expected. She is unable to meet her obligations to B. If, following A's default, B continues to charge \$1 for each unit he delivers, then A can only afford to take delivery of 60 units. B can liquidate the remaining 40 undelivered units, but, since they are specific to A, they fetch less than \$1 a unit, say just \$0.5. Thus B gets \$80 in total: \$60 in cash from A plus \$20 in liquidation receipts.

This is less than the \$100 B had expected to receive from A, and so he is unable to meet his obligations to C. If C charges \$1 for each unit delivered, then B can only afford to take delivery of 80 units. C can liquidate the remaining 20 undelivered units, again say for \$0.5 each. Thus C gets \$90 in total: \$80 in cash from B plus \$10 in liquidation receipts. And if C is a debtor, there may be further default along the credit chain.

Notice that the secondary defaults by B, C, ... would be avoided if A's debt contract were with a deep pocket, who is not credit-constrained, rather

than with B. That is, the wrong people are being exposed to A's risk. The example suggests that the economy with credit chains reacts to shocks more than an economy where gross financial positions can be netted out (i.e. where the market for loans is anonymous on the supply side).

This is confirmed in Section 2 of the paper. Our formal model is of a network of credit rather than of a line of credit. Each person is owed money by several people; and each person owes money to several people, not all of whom are credit constrained. We exhibit an aggregate multiplier effect: everyone could in principle be made better off if people forgave each others' debts. We show that the greater are people's gross positions (the longer the credit chains), the stronger is the multiplier.

In practice, liquidation appears to be the exception not the rule, even though creditors usually have the right to liquidate the assets of a delinquent debtor. Typically, unpaid debts are postponed until such time as the debtor has enough cash to pay. How would the rescheduling of debt affect our analysis? This is the subject of Section 3 of the paper.

The reason why unpaid debts are postponed is that it may be in the joint interests of a creditor and debtor not to liquidate. That is, it may be privately efficient for them to reschedule the debt. Socially, however, postponement may be more damaging than liquidation if the two parties are a link in a credit chain, because postponement injects less liquidity into the system than does liquidation.

To see this, consider again the A, B, C, ... chain. When B liquidates 40 units of A's product, he injects \$20 cash into the system. (He defaults less than A: he pays \$80 of his debt, whereas A pays only \$60.) And when C liquidates 20 units of B's product, a further \$10 cash is injected into the system. As we move along the chain, the initial \$40 shock (to A's cash holding) converges to zero.

In fact, however, B may choose to postpone the unpaid portion of A's debt. Rather than liquidate the undelivered stock of 40 units today, B may be better off storing it until tomorrow, by which time A will have more cash (viz., the cash return from the 60 units of input that she can afford to buy

today). Liquidation generates \$20 today, which B can use to buy 20 more units of his own input from C. But if tomorrow's return on B's investment is less than \$2 per unit of input, this strategy yields him less than the \$40 he gets from simply postponing delivery of the 40 units until tomorrow. Moreover, A is also better off if B decides to postpone the debt, because she enjoys the gains from trade on those 40 units tomorrow, which she doesn't enjoy if B liquidates. Hence, provided the return on B's investment is less than \$2 per unit of input, postponement is bilaterally efficient.¹

Postponement imposes a cost on C, however. Now C is in the same position as B: namely, of being owed \$100 by a debtor who only has \$60 cash. By the same argument, then, provided the return on C's investment is less than \$2 per unit of input, C will choose to postpone \$40 of B's debt rather than liquidate 40 units. And so on. Notice that there is no diminution of the shock along the credit chain if debts are postponed, because there is no liquidation to inject new cash into the system. In fact the social loss can be arbitrarily large, given a long enough chain! This extremely stark conclusion is tempered if, as in our formal model, each entrepreneur has many suppliers/creditors, not all of whom are credit constrained. Nevertheless, it is clear that postponement can be socially worse than liquidation, even though it is privately optimal. The root cause of the inefficiency is that a decision by a creditor to postpone debt imposes a negative externality on the creditor's creditors, and in turn on their creditors.

Thus far in the paper, the analysis assumes that the liquidity shocks are unexpected. Section 4 places the model of Section 2 in a fully-specified stochastic framework. Given rational expectations, shocks are anticipated, which has implications for the form of contract. In particular, we ask the question: If default risk has aggregate consequences, and so is publicly observable, why don't entrepreneurs insure against it? That is, why don't entrepreneurs insure themselves against fluctuations in their accounts receivable?

¹More generally, if the delivery price is p per unit, if liquidation yields l per unit, and if the return on investment is α per unit, where $l < p < \alpha$, then B will postpone A's debt if $l(\alpha/p) < p$.

An entrepreneur faces two kinds of uncertainty in his income: uncertainty in the returns from his own investment; and uncertainty in his accounts receivable arising from fluctuations in his customers' incomes. If the entrepreneur has many customers, fluctuations in his accounts receivable will be highly correlated with the aggregate state of the economy, and so can be insured. (We assume that, because individual returns are unverifiable, the entrepreneur cannot insure them.) The benefit of insurance is that the entrepreneur would obtain better terms from his own suppliers if he could demonstrate that he is less likely to default. Moreover, if all entrepreneurs were to take out such insurance, then the aggregate multiplier we have identified would go away: default risk would be absorbed by deep pockets via insurance companies. We find, however, that the private costs of insurance may outweigh the benefit.

There are two reasons why insurance is unattractive. First, given the limited enforcement of contracts, the entrepreneur cannot commit to pay the insurance premium out of the returns from his investment. That is, his "ideal" policy -- in which he pays out if and only if the economy booms (in return for being subsidized in a recession) -- is infeasible.² Instead, he has to pay the premium upfront, which eats into the funds he has available for undertaking his own investment. And, given that he is credit constrained, his internal rate of return will dominate that of the insurance company, so that investing in an insurance policy yields a poor return.³ Second, the fluctuations in accounts receivable may be relatively small and not strongly correlated with the fluctuations in individual returns. So in a recession, when more of his customers default, the entrepreneur's own return may well be high enough to more than compensate for the shortfall in accounts received. In which case, he doesn't need any insurance in order to meet his own debt obligations. In other words, the payout from an insurance policy conditioned on the aggregate state may not be sufficiently correlated with

²We assume that there are no enforcement problems relating to the deep pockets: they can commit to pay out.

³A similar argument shows that self-insurance, in the form of investing enough in liquid assets to avoid default, also has too high an opportunity cost.

his own needs for cash to make the policy worth buying.

The first conclusion of Section 4, then, is that the aggregate multiplier may not be insured away in a full rational expectations model.

Another conclusion is that positive shocks cause an economy to respond less, in percentage terms, than negative shocks. The reason is that default or postponement is essentially asymmetric: a shortage of cash leads to the multiplier we have identified, whereas there is no counterpart if entrepreneurs have surplus cash. An implication of this asymmetry is that increased variance across entrepreneurs -- some defaulting, others with surplus cash -- leads to lower levels of activity in aggregate.

In the A-B-C chain, B is in effect playing the role of a financial intermediary. One is tempted to say that B stands for "bank". However we are far from modelling banks. As will become clear in the next section, ours is a primitive model of small entrepreneurial firms. We have found it easier to model the kinds of basic decision that small firms face -- such as how much to buy, produce, and sell -- rather than to model complex financial institutions like banks. In fact, though, most small firms do simultaneously hold gross credit and debt positions. And, arguably, systemic risk is just as important for inter-firm lending as it is for inter-bank lending. Trade credit is an important part of commercial life, and constitutes a significant fraction of total borrowing and lending. However, this paper should not be seen as an attempt to model trade credit as such, but rather as a framework for thinking about the aggregate behavior of an economy where agents simultaneously hold gross financial assets and liabilities.

In Section 5, we relate our paper to the literature, and make some final remarks.

2. The basic model: systemic risk

The economy has three dates, 0, 1 and 2; and there are two populations, entrepreneurs and deep pockets, which are both large. All agents are risk neutral and only consume a general commodity at date 2. We refer to this general commodity simply as "goods", and we take it as the numeraire. Goods can be stored without depreciation. Initially, the entrepreneurs and deep pockets are respectively endowed with M and \bar{M} goods in aggregate, where \bar{M} is large. Between dates 0 and 1, the entrepreneurs and the deep pockets respectively have an aggregate endowment of N and \bar{N} labor, where \bar{N} is large.

Deep pockets have access to a constant-returns, one-period technology for making goods at date 1 directly from their own labor: one unit of labor makes one unit of goods.

A typical entrepreneur, E , has access to two technologies for producing goods, one short-term and the other long-term. Both exhibit constant returns, and are specific to him. The short-term technology is effectively a superior storage technology: E can use one unit of goods at either date $t = 0$ or 1 to make $\sigma > 1$ units of goods at date $t+1$.

E 's long-term technology has two stages, each stage taking one period to complete. In the first stage, between dates 0 and 1, labor is used to make an intermediate product, which is specific to E and has to be custom-made. We assume that E cannot use his own labor for this purpose. Instead, at date 0 he places orders with suppliers -- other entrepreneurs and deep pockets -- for intermediate product to be delivered at date 1. A supplier works between dates 0 and 1 with a blueprint that E gives her at the time of placing the order: one unit of her labor makes one unit of intermediate product. In the second stage, between dates 1 and 2, E uses his specific technology to make goods from the intermediate product: one unit of intermediate product at date 1 yields $\alpha > 1$ units of goods at date 2.

At date 1, since the intermediate product is specific, it is of less value to a supplier than it is to E . We assume that the only use for the supplier is to liquidate it. Liquidation is instantaneous and exhibits constant returns: each unit of intermediate product can be liquidated for

$l < 1$ goods at date 1.

The model is symmetric: E uses his labor to make intermediate products for other entrepreneurs. At date 0 there is a competitive market: any entrepreneur is free to place an order with any other entrepreneur or any deep pocket to supply intermediate product at date 1.

To avoid counter-trade, however, we also assume that if E makes intermediate product for some entrepreneur E^* , then this precludes E^* producing for E.

Restrictions on contracting

In Section 3 we will appeal to some auxiliary assumptions about an entrepreneur E's technology to justify a number of restrictions on contracting.⁴ The first restriction prevents E from obtaining outside funds:

Restriction 1: E cannot commit to hand over any of the return from either his short-term or his long-term technology.

That is, since E has no collateral to offer, a deep pocket would be unwilling to put up funds to finance E's investment: E has to use his own resources.

The second restriction concerns the relationship between E and his suppliers:

Restriction 2: E and his suppliers cannot contract over the terms of delivery of the intermediate product.

⁴Note that these additional assumptions do not affect the analysis in any way other than to restrict contracting.

This means that at date 1 the terms of delivery will be a matter of negotiation. We suppose the outcome of the bargain is that at date 1 E pays a fixed amount, p goods, for each unit of intermediate product delivered. The parameter p is taken to satisfy $\ell < p \leq 1$; i.e., there is an interior division of surplus.⁵ In particular, then, the negotiated delivery price p is assumed to be strictly greater than the suppliers' valuation ℓ .

It is important to notice that not all the units of intermediate product ordered at date 0 will be delivered at date 1 if E does not have sufficient funds to pay p for all of them. In this event, E "defaults", and the suppliers liquidate the undelivered units, losing $p - \ell$ on each. Although the liquidated intermediate product isn't being put to best use, the outcome of the bargain is nevertheless constrained efficient, taking into account the fact that E has limited funds with which to pay his suppliers.

The third restriction is that suppliers have equal bargaining power in the event of default:

Restriction 3: If E defaults against his suppliers, then he does so on a pro-rata basis.

In particular, the bargaining outcome does not depend on the identities of E's suppliers: entrepreneurial suppliers and deep pocket suppliers are treated alike.

The final restriction deals with E as a supplier:

⁵Strictly speaking, since $\alpha > 1$, an interior division of surplus doesn't require that $p \leq 1$; but this tighter upper bound on p helps with the interpretation of the model.

Generally, p could depend on the funds held by E and his suppliers, and could vary from supplier to supplier. Here we treat p as a fixed parameter in order to simplify the analysis.

Restriction 4: E cannot raise funds at date 0 against the future receipts from any intermediate product that he makes for other entrepreneurs.

In other words, E's earnings as a supplier cannot be securitized.

Implicit debt/supply contracts

In the competitive market at date 0, an entrepreneur agrees an implicit contract with his suppliers. Inter alia, the implicit contract deals with what fraction, λ , of the intermediate product that the entrepreneur orders at date 0 is expected to be liquidated at date 1. (In principle, each entrepreneur could have a different λ ; although, as we shall see, in the absence of any shocks, all λ 's will equal zero.)

The market clears by means of transfers of goods at date 0. Specifically, an entrepreneur makes a downpayment of q goods for each unit that he orders. (q is personal to the entrepreneur, insofar as λ is too.)


Thus the implicit contract between the entrepreneur and one of his suppliers takes the following form. At date 0, the entrepreneur gives the supplier his blueprint and orders a quantity x of intermediate product. At the same time, the entrepreneur makes a downpayment of qx goods. Then at date 1, the entrepreneur pays $p(1 - \lambda)x$ goods for the delivery of $(1 - \lambda)x$ units of intermediate product. The supplier liquidates λx units. In total, the supplier gets a return of qx goods at date 0 and $px - (p - \ell)\lambda x$ goods at date 1.

The implicit contract can be thought of as a debt contract and a supply contract bundled together. The deferment of part of the payment until date 1 can be construed as the entrepreneur borrowing from the supplier. (But this can just as well be turned round the other way: because the supplier is paid something in advance by the entrepreneur, she borrows from him. It is moot as to who is the borrower and who is the lender in a context where the product takes time to build.)

To recap: the reason why the deal between the entrepreneur and his supplier takes the above form is that formal contracts are difficult to enforce. By staggering the payments -- a downpayment of q goods per unit at the time of ordering, and a final payment of p goods per unit delivered -- the parties protect their respective interests. If the balance between q and p were shifted towards q (more downpayment), then the entrepreneur would be at risk from the supplier demanding more on delivery than was implicitly agreed. And if the balance were shifted towards p (more payment on delivery), then the supplier would be at risk from the entrepreneur refusing to pay. Unlike in a conventional spot transaction, the payments do not all coincide with delivery. And unlike in an Arrow-Debreu forward contract, the payments do not all occur at the time of ordering. Crucially, the fact that the implicit contract involves some form of debt contract is inescapable. Entrepreneurs who are credit-constrained are forced to offer credit to their customers, otherwise they will be unable to compete with other suppliers.

Credit chains


This opens the door to a model of credit/supply chains. Consider an entrepreneur E who supplies entrepreneur E^* . E purchases intermediate product both from another entrepreneur E^{**} ($\neq E^*$) and from a deep pocket D . The chain of supply of intermediate product is thus



a supply chain: E^{**} and D supply E

E supplies E^*

and the corresponding credit chain is



a credit chain: E* is in debt to E

E is in debt to E** and D

Here, the goods that E* owes to E are part of E's accounts receivable; and the goods that E owes to E** and D are part of E's accounts payable.

It is important to bear in mind that these links are part of an intricate network of such credit/supply relationships: E has many customers like E*, and he has many suppliers like E** and D. For future reference, it helps to draw the "canonical" credit network, viz. the triangle given in Figure 1. (Note that the figure may be misleading insofar as the actual credit/supply network comprises a jumble of bilateral trades.)




Figure 1: The canonical credit network

E is in debt to suppliers E** and D

E* is in debt to suppliers E and D*

E** is in debt to suppliers E* and D**

Equilibrium

We now find the equilibrium value of q , for a given λ . Since the deep pockets' endowment of goods, \bar{M} , is large, in equilibrium some will be stored until date 2. That is, the deep pockets value goods equally at all three dates: their rate of interest is zero. Also, since the deep pockets' endowment of labor, \bar{N} , is large, in equilibrium some will be used to make goods using their own technology. For a deep pocket D to be indifferent between using her own technology and making intermediate product for some entrepreneur E, the q for that entrepreneur must satisfy

$$(2.1) \quad q + p - (p - \ell)\lambda = 1,$$

where λ is the fraction of E's intermediate product that is expected to be liquidated at date 1. The left hand side of (2.1) is the expected present value of the return to D of producing one unit of intermediate product for E,

and the right hand side is D's opportunity cost of withholding one unit of labor from her own technology. Our earlier assumption that $p \leq 1$ ensures that $q \geq 0$.

Note that in equilibrium the aggregate demand for intermediate product by the entrepreneurs will exceed their aggregate supply, and the deep pockets supply the difference. This is why another entrepreneur E^{**} who supplies E with intermediate product charges E the same downpayment q that D charges -- namely, the q given by (2.1).⁶

It is clear from (2.1) that if E doesn't anticipate any shocks to his flow-of-funds at date 1, then he has an incentive to avoid default, since q rises with λ . The point is that liquidation is wasteful: the deep pocket D is more productively employed on her own technology (yielding one unit of goods per unit of labor) than on making intermediate product which is liquidated (effectively yielding only ℓ goods per unit of labor).

Since E's λ equals zero, he has to pay a downpayment of $q = 1 - p$ goods for each unit of intermediate product that he orders at date 0. And by symmetry, each of E's purchasers, E^* , will pay E a downpayment price of $1 - p$.

To avoid defaulting at date 1, E must arrange to have at least as many funds flowing in as he plans to spend paying for the delivery of his own intermediate product. His planned expenditure is pX , where X , say, is the quantity of intermediate product that he orders at date 0. (pX corresponds to his accounts payable.) His planned income comes from two sources. First, he expects a return σ on the short-term investment Y , say, that he undertakes at date 0.⁷ Second, he expects to be paid by his customers (other entrepreneurs) for the delivery of intermediate product that they order from him at date 0. As he is endowed with N units of labor between dates 0 and 1, he can make N units of intermediate product. Moreover, none of his customers

⁶Because E defaults equally against all his suppliers, E^{**} cannot offer a supply contract different from D's.

⁷Since the short-term technology strictly dominates storage, E does not store any goods at date 0.

is planning to default, and so E expects to be paid pN goods at date 1. (pN corresponds to his accounts receivable.) Hence, in order to avoid defaulting at date 1, E must choose X and Y at date 0 so that

$$(2.2) \quad pX \leq \sigma Y + pN. \quad (\text{date 1 flow-of-funds})$$

\uparrow
 accounts payable

\uparrow
 revenue from short-term investment

\uparrow
 accounts receivable

If E chooses X and Y so that this inequality is strict, then he will have surplus goods at date 1 to reinvest in his short-term technology, the return from which will augment the date 2 return αX from his long-term technology. On the other hand, if E chooses X and Y so that (2.2) is an equality, E has just enough funds at date 1 to meet his debt obligations. This latter case, which we call a balanced investment strategy, yields E a date 2 consumption of αX .

It is straightforward to show that the balanced investment strategy is optimal if the long-term return α is sufficiently large relative to the short-term return σ . A sufficient condition is that

$$(A.1) \quad \alpha > \sigma^2.$$

Assumption (A.1) says that using the long-term technology yields more than using the short-term technology twice.

We can solve for E's balanced investments X and Y from (2.2) as an equality, together with his budget constraint at date 0:

$$(2.3) \quad (1 - p)X + Y = M + (1 - p)N \quad (\text{date 0 flow-of-funds})$$

Here, the left-hand side comprises E's downpayment $(1 - p)X$ goods on the X units of intermediate product that he orders from his suppliers for delivery at date 1, plus the Y goods he invests in his short-term technology. The right-hand side comprises his initial goods endowment M , plus the $(1 - p)N$ goods that he receives as downpayment for the delivery of N units of intermediate product to his customers at date 1.

Solving for X and Y , we obtain

$$(2.4a) \quad X = N + \sigma M / [p + \sigma(1-p)]$$

$$(2.4b) \quad Y = pM / [p + \sigma(1-p)].$$

This leads to E having a date 2 consumption C equal to αX ; i.e.,

$$(2.5) \quad C = \alpha N + \alpha \sigma M / [p + \sigma(1-p)].$$

In sum, the equilibrium is characterised by all entrepreneurs choosing X and Y given by (2.4), and consuming C given by (2.5). The excess demand for intermediate product, $X - N$ per entrepreneur, is supplied by the deep pockets.

The effects of an unanticipated shock at date 1

Now suppose that, at date 1, σ unexpectedly turns out to be smaller for all entrepreneurs, say equal to $\hat{\sigma} < \sigma$.⁸

⁸In Section 4 we look at the model where the shock is rationally expected at date 0.

Now E no longer has enough funds at date 1 to be able to take delivery of all X units of intermediate product: he must default. But then so too must all entrepreneurs, since they also experience the shock. Let \hat{X} and \hat{X}^* denote the respective quantities of intermediate product that E and a typical other entrepreneur E^* (who buys from E) can afford to pay for at date 1. (By symmetry, $\hat{X}^* = \hat{X}$. But it clarifies the discussion if we use an asterisk to denote other entrepreneurs' activities.) Both \hat{X} and \hat{X}^* are less than the amounts that E and E^* ordered at date 0, X and X^* respectively.

Recall that E^* defaults against his suppliers (including E) on a pro-rata basis. That is, E^* defaults on a fraction $\lambda^* = (X^* - \hat{X}^*)/X^*$ of the order he placed with E at date 0. Since this is true of all the entrepreneurs that E supplies, E's sales revenues are reduced by $(p - \ell)\lambda^*N$ at date 1. This indirect shock is over and above the direct productivity shock $(\sigma - \hat{\sigma})Y$ to E's flow-of-funds:

$$(2.6) \quad p\hat{X} = \hat{\sigma}Y + pN - (p - \ell) \left(\frac{X^* - \hat{X}^*}{X^*} \right) N.$$

Equation (2.6) lies at the heart of the paper. A multiplier process is at work: E defaults on his suppliers, some of whom are entrepreneurs. As a result, these entrepreneurs default more -- that is, even more than the direct productivity shock dictates -- on their suppliers. And so on, until we find that the entrepreneurs that E is supplying default more on him; etc. The credit chains amplify the shock.

This can be seen most easily in terms of Figure 1. E defaults on E^{**} and D on a pro-rata basis. This means that E^{**} defaults more on E^* and D^{**} than he otherwise would. E^* thus defaults more on E and D^* , which worsens E's flow-of-funds, etc.

There is an easy way for the economy to effect a Pareto improvement. If each entrepreneur were to forgive some of the "debt" of the entrepreneurs he supplies -- that is, if he were to charge them less than p per unit of intermediate product that he delivers -- and if his suppliers who are

entrepreneurs were to forgive some of his debt, then all the entrepreneurs would be strictly better off (and none of the deep pockets would be affected). In terms of Figure 1, E can charge E* less than p per unit, given that E** is charging E less than p per unit (and E* is charging E** less than p per unit). Another way of saying this is that instead of E, E* and E** liquidating intermediate product that they have respectively custom-made for E*, E** and E, the three entrepreneurs could pass it around the triangle for free, and they would all be better off.

The root cause of the inefficiency is the lack of liquidity in the system, and the difficulty of coordination. Coordination is difficult because each entrepreneur can gain by continuing to insist that he be paid p goods per unit of intermediate product that he supplies, even though other entrepreneurs are forgiving his debt to them.

The longer the credit chains, the greater the multiplier. Think of a parallel economy in which entrepreneurs are able to make their intermediate product from their own labor. In such an economy, there would be no chains. In terms of Figure 1, the only links would be between E and D, E* and D*, and E** and D** (the deep pockets satisfy the three entrepreneurs' residual demand). Hence there would be no multiplier. Following the shock, the entrepreneurs would still default, but the deep pockets would absorb the loss without passing it on to others, and without having to cut back on any profitable investment of their own.

Appealing to symmetry (dropping the asterisks), we can solve (2.6) for \hat{X} in terms of its proportional deviation from X:

$$(2.7) \quad \frac{X - \hat{X}}{X} = \frac{\sigma - \hat{\sigma}}{\sigma} \left(\frac{\sigma p M}{\sigma p M + \ell N [p + \sigma(1-p)]} \right).$$

This is also the proportional fall in the entrepreneurs' consumption at date 2.

A numerical example

Suppose the returns on the short-term and long-term technologies, σ and α , are 1.2 and 1.8 respectively; and the liquidation return ℓ is 0.5. (Note that assumption (A.1) is satisfied.) Let E's initial endowments of goods and labor, M and N, be 5 and 24. And let $p = 1$.

A 5% fall in the short-term productivity at date 1 -- which corresponds to a 1% fall in the entrepreneurs' net worth -- leads to a $\frac{5}{3}\%$ fall in the entrepreneurs' consumption at date 2. Had the entrepreneurs been able to make their own intermediate product -- so that there were no credit chains -- then the fall in their consumption would have been 1%; i.e. the same as the drop in their net worth. The extra $\frac{2}{3}\%$ fall is due to the multiplier.

3. A three period model: postponement

In practice, liquidation is the exception not the rule. Typically, unpaid debts are rescheduled. In this section, we extend our basic model by a period, so as to investigate the aggregate implications of postponement.

There are now four dates, 0, 1, 2, and 3. Consumption only occurs at date 3. The entrepreneurs' aggregate endowments are unchanged: they initially have M goods, and are endowed with N units of labor between dates 0 and 1. An entrepreneur can use his short-term technology at each of dates 0, 1 and 2. And he has the opportunity to use his long-term technology twice: between dates 0 and 2, and between dates 1 and 3.

The deep pockets have an aggregate endowment of \bar{M} goods initially, and \bar{N}_0 units of labor between dates 0 and 1. In addition, they have \bar{N}_1 units of labor between dates 1 and 2. \bar{M} , \bar{N}_0 and \bar{N}_1 are all large. A deep pocket can use her own technology for producing goods directly from her own labor. Or she can use her labor to make intermediate product for entrepreneurs. Since only the deep pockets are endowed with labor between dates 1 and 2, only they are in a position to supply intermediate product to entrepreneurs at date 2.

We continue to assume that implicit contracts for the supply of intermediate product take the form given in Section 2: entrepreneurs pay q goods for each unit of intermediate product that they order, and pay p goods per unit delivered. A supplier can liquidate any undelivered intermediate product for ℓ goods per unit.

We assume that intermediate product not delivered at date 1 need not be liquidated, however. Instead, a supplier can store it without depreciation until date 2, when it can be delivered to the entrepreneur for whom it was made. We also assume that the entrepreneur will pay a price of p . That is, if delivery is postponed, the terms of delivery are negotiated afresh at date 2, and p is the outcome of the bargain between the entrepreneur and his supplier.

Using arguments similar to those in Section 2, one can show that the entrepreneurs will choose balanced investment plans. Specifically, an

entrepreneur E draws up a plan at date 0 under which he invests Y in his short-term technology at date 0, orders X_1 intermediate product to be delivered at date 1, and then at date 1 orders X_2 intermediate product to be delivered at date 2 -- where Y, X_1 and X_2 satisfy

$$(3.1) \quad (1 - p)X_1 + Y = M + (1 - p)N \quad (\text{date } 0)$$

$$(3.2) \quad (1 - p)X_2 + pX_1 = \sigma Y + pN \quad (\text{date } 1)$$

$$(3.3) \quad pX_2 = \alpha X_1 \quad (\text{date } 2)$$

The first expenditure terms on the left-hand sides of (3.1) and (3.2) reflect the fact that, since E plans not to default at either dates 1 or 2, he pays price $q = 1 - p$ for each of the units of intermediate product that he orders at dates 0 and 1. The final income terms on the right-hand sides (3.1) and (3.2) reflect the fact that the other entrepreneurs also do not plan to default at date 1: E receives price $1 - p$ for each of the N units of intermediate product that is ordered from him at date 0, and he expects to deliver all N units for price p at date 1.

The idea behind (3.1)-(3.3) is that E invests as much as he can long-term -- using the returns αX_1 at date 2 from the first round of long-term investment to pay for the delivery of intermediate product X_2 in the second round. However, at date 1 he has to pay for the delivery of X_1 and the downpayment for X_2 partly out of the return from short-term investment Y at date 0.

E's planned consumption at date 2 equals

$$(3.4) \quad C = \alpha X_2.$$

The effects of an unanticipated shock at date 1

Now suppose that, for all the entrepreneurs, the return from the short-term investment unexpectedly turns out to be $\hat{\sigma} < \sigma$ at date 1.⁹ As a result, entrepreneurs are unable to execute the plans they drew up at date 0, because they are short of funds.

On the income side, an entrepreneur E finds that his customers (other entrepreneurs) take delivery of only $(\hat{X}_1^*/X_1^*)N$ units of the intermediate product he made for them -- where \hat{X}_1^* is the total amount that they take delivery of, out of the X_1^* units they ordered at date 0. E has a choice. Either he can liquidate the $N - (\hat{X}_1^*/X_1^*)N$ undelivered units for l goods each. Or he can postpone delivery of these units until date 2, by which time his customers will have more goods with which to pay him.

E's choice boils down to this. Is it more profitable to liquidate and invest in his own long-term technology -- i.e. to use the liquidation receipts l (per unit) to take delivery of an additional l/p units of his own intermediate product, yielding $\alpha l/p$ goods at date 2? Or is it more profitable to "invest in his customers" -- i.e. to wait until they can pay p at date 2? If

$$(A.2) \quad l(\alpha/p) < p$$

then the latter strategy will be preferred. Note that E's customers are only too happy to go along with this, given that, if E postpones, they get intermediate product delivered at date 2 without their having to make any more downpayment to a new supplier at date 1. (They get nothing if E liquidates.) Hence, under assumption (A.2), postponement is bilaterally efficient between E and each of his customers.

⁹This is only a temporary shock. If an entrepreneur invests in his short-term technology at date $t = 1$ or 2 then he earns σ at date $t+1$.

Notice that E's own suppliers face the same choice as he does: whether to liquidate the intermediate product they have made for him, or whether to postpone delivery until date 2. By a symmetric argument, the entrepreneurs who are supplying E will choose to postpone, given assumption (A.2). And, a fortiori, the deep pocket suppliers will postpone too, since postponement yields a price p at date 2, whereas liquidation yields only ℓ at date 1 (and they have a zero rate of interest).

On the expenditure side, at date 1 E has to decide how to allocate his limited resources between taking delivery of \hat{X}_1 intermediate product and placing orders for delivery of $\hat{X}_2 - (X_1 - \hat{X}_1)$ additional intermediate product at date 2. (If \hat{X}_2 is E's total input of intermediate product at date 2, and his date 1 suppliers are to postpone delivery of $X_1 - \hat{X}_1$ units until that date, then at date 1 he only needs to make the downpayment $1 - p$ on the difference, $\hat{X}_2 - (X_1 - \hat{X}_1)$.) It can be shown that E will adopt a balanced investment plan:

$$(3.5) \quad (1 - p) \left(\hat{X}_2 - (X_1 - \hat{X}_1) \right) + p \hat{X}_1 = \hat{\sigma} Y + p \left(\hat{X}_1^* / X_1^* \right) N \quad (\text{date 1})$$

$$(3.6) \quad p \hat{X}_2 = \alpha \hat{X}_1 + p \left(N - (\hat{X}_1^* / X_1^*) N \right). \quad (\text{date 2})$$

This new plan entails ordering just enough additional intermediate product $\hat{X}_2 - (X_1 - \hat{X}_1)$ for delivery at date 2 that he can afford to take delivery of \hat{X}_2 at price p using the returns $\alpha \hat{X}_1$ plus the receipts $p \left(N - (\hat{X}_1^* / X_1^*) N \right)$ from the postponed delivery to his date 1 customers.

E's new consumption at date 3 equals

$$(3.7) \quad \hat{C} = \alpha \hat{X}_2.$$

\hat{C} can be found by appealing to symmetry (dropping the asterisks), and solving (3.1), (3.5), (3.6) and (3.7).

Comparison with liquidation

The point of this analysis is to compare \hat{C} with what an entrepreneur would consume if everyone liquidated rather than postponed. (This is what would occur if (A.2) did not hold.) Let a tilde denote this possibility. (3.5)-(3.7) would read:

$$(3.8) \quad (1 - p)\tilde{X}_2 + p\tilde{X}_1 = \hat{\sigma}Y + pN - (p - \ell) \left(\frac{X^* - \tilde{X}^*}{X^*} \right) N \quad (\text{date 1})$$

$$(3.9) \quad p\tilde{X}_2 = \alpha\tilde{X}_1 \quad (\text{date 2})$$

$$(3.10) \quad \tilde{C} = \alpha\tilde{X}_2.$$

The final income term in (3.8) reflects the lost revenue $p - \ell$ on each unit that E liquidates at date 1.

The formulae for \hat{C} and \tilde{C} are messy, so it helps to consider the simple case $p = 1$, and to look at the proportional deviation of consumption from planned consumption, given postponement and liquidation:

$$(3.11) \quad \frac{C - \hat{C}}{C} = \frac{\sigma - \hat{\sigma}}{\sigma} \frac{\alpha\sigma M + (\alpha - 1)N}{\alpha\sigma M + \alpha N} \quad (\text{postponement})$$

$$(3.12) \quad \frac{C - \tilde{C}}{C} = \frac{\sigma - \hat{\sigma}}{\sigma} \frac{\sigma M}{\sigma M + \ell N} \quad (\text{liquidation})$$

An inspection of (3.11) and (3.12) reveals that the proportional drop in \hat{C} may be more than the drop in \tilde{C} .

Consider again the numerical example from the end of Section 2: $\sigma = 1.2$, $\alpha = 1.8$, $\ell = 0.5$, $M = 5$, $N = 24$ and $\hat{\sigma} = 0.95\sigma$. In this example, assumption (A.2) is satisfied, and so debts are postponed. However, \hat{C} is $(25/9)\%$ lower than C , whereas \tilde{C} is only $(5/3)\%$ lower. In other words, postponement may be socially inferior to liquidation, even though it is bilaterally efficient.

The reason is that each bilateral credit/supply relationship is part of a credit chain. When an entrepreneur E privately decides to postpone the debt (i.e. chooses not to liquidate the intermediate product) of one of his customers, he does not take into account the fact that it imposes a negative externality on his creditors (suppliers). From E 's perspective, more return at date 2 (p per unit postponed) is better than less at date 1 (ℓ per unit liquidated) -- given assumption (A.2). But E 's suppliers are better off if E has more goods today with which to pay them.

In short, there is too little liquidity in the system at date 1, and postponement does nothing to relieve the situation, whereas liquidation helps. Were the agents able to participate in a grand Coasian bargain, then the inefficiency could be avoided, but we are considering an economy where only bilateral negotiations are feasible.

4. A stochastic model: insurance

Thus far in the paper, the date 1 shocks have been unanticipated. In this section, we place the two-period model of Section 2 in a fully-specified stochastic framework, where agents have rational expectations.

There are two aggregate states at date 1. In the boom, which occurs with probability $1 - \pi$, all the entrepreneurs receive a deterministic return $\sigma > 1$ from their short-term technologies. In the recession, which occurs with probability π , a fraction θ of the entrepreneurs receive a return of only $\sigma^- < \sigma$; and a fraction $1 - \theta$ receive $\sigma^+ > \sigma$. We take $\sigma^+ - \sigma$ and $\sigma - \sigma^-$ to be small. It is not known at date 0 who will be the (relatively) unproductive entrepreneurs. There are no further shocks; in particular, the return on all the entrepreneurs' short-term technologies between dates 1 and 2 equals σ .

Notice that we allow for the case where $\theta\sigma^- + (1 - \theta)\sigma^+ = \sigma$. That is, if there is a mean-preserving spread of the entrepreneurs' short-term returns, then default by a fraction of the entrepreneurs is enough to bring about a recession. Moreover, as we shall see, the size of the downturn is exacerbated by a chain reaction of default.

We assume that at date 1 it cannot be publicly verified if an entrepreneur E has productivity σ^+ or σ^- . For this reason, E cannot insure against fluctuations in his individual return in a recession. However, E can insure against the aggregate state, because this is publicly verifiable at date 1.

E's ideal insurance policy would be of the form: in return for agreeing to pay out at date 1 if one particular aggregate states occurs, he receives payment in the other state. However, the restrictions on contracting discussed in Section 2 (in particular Restriction 1) mean that E cannot commit to hand anything over at date 1. Hence the only feasible insurance policy is one in which he pays a premium in advance, at date 0, and he is paid according to the aggregate state at date 1. This is less than ideal for E because he spends valuable funds on the premium at date 0 which could otherwise be used for investment.

Deep pockets can set aside funds at date 0 as security. Thus in a competitive equilibrium, they can provide insurance to the entrepreneurs. The rate of return on insurance equals zero, the deep pockets' rate of interest.

If an entrepreneur E pays a premium of $Z > 0$ goods at date 0, he can buy a policy that pays out Z/π goods in the recession. Equally, he can buy a different policy that pays out $Z/(1 - \pi)$ goods in the boom.

At date 0, E has to decide what amount, X , of intermediate product to order at date 0; what amount, Y , to invest in his short-term technology; and what amount, Z , to invest in insurance. Let (X^*, Y^*, Z^*) denote the other entrepreneurs' choices. In a symmetric equilibrium, $(X^*, Y^*, Z^*) = (X, Y, Z)$.

We are interested in an equilibrium in which no entrepreneur buys insurance ($Z = Z^* = 0$), and each arranges to have just enough funds in the boom not to default:

$$(4.1) \quad pX = \sigma Y + pN.$$

(This corresponds to the anticipated date 1 flow-of-funds constraint (2.2) from section 2, as an equality.) In the boom, then, there is no fall in accounts received: the final term in (4.1) is pN , the revenue E gets from delivering all his output of intermediate product to his customers at price p .

In the recession, however, the unproductive entrepreneurs will default, which means that everyone suffers a loss in accounts received. First, take the case where E himself is unproductive. Let \hat{X} denote the amount of intermediate product that he can afford to buy in this case; and let \hat{X}^* denote the equivalent for other unproductive entrepreneurs. (\hat{X}^* equals \hat{X} in a symmetric equilibrium.) Then, from E's flow-of-funds at date 1,

$$(4.2) \quad p\hat{X} = \sigma^-Y + pN - (p - \ell)\theta\left(\frac{X^* - \hat{X}^*}{X^*}\right)N.$$

(This corresponds to the realised date 1 flow-of-funds constraint (2.6) from Section 2.) The final term in (4.2) is the loss in accounts received arising from liquidation: the unproductive fraction, θ , of E's customers default on a fraction $(X^* - \hat{X}^*)/X^*$ of their order, incurring E a loss of $(p - \ell)$ per unit.

Next, take the case where E is productive in the recession. In this case, we want the equilibrium to be such that E has spare funds available at date 1, in that the revenue σ^+Y from his short-term investment more than makes up for the loss in accounts received, thereby enabling him to pay for all X units of intermediate product that he ordered at date 0 from his suppliers:

$$(4.3) \quad \sigma^+Y + pN - (p - \ell)\theta\left(\frac{X^* - \hat{X}^*}{X^*}\right)N > pX.$$

E reinvests these surplus funds into his short-term technology between dates 1 and 2.

Turn now to date 0. There is a probability $\pi\theta$ that E will default at date 1 and only be able to take delivery of a fraction (\hat{X}/X) of the X units of intermediate product he orders. Hence E has to make a downpayment of q goods on each unit, where q solves

$$(4.4) \quad q + p - (p - \ell)\pi\theta\left(\frac{X - \hat{X}}{X}\right) = 1.$$

(This corresponds to the equilibrium condition (2.1) from section 2.) By symmetry, all other entrepreneurs have to make the same downpayment q on the intermediate product that they order.

Given that he chooses $Z = 0$, E's date 0 flow-of-funds constraint

reduces to

$$(4.5) \quad qX + Y = M + qN.$$

(This corresponds to (2.3) from Section 2.)

E's expected consumption at date 2 equals

$$(4.6) \quad C = (1 - \pi)\alpha X + \pi\theta\hat{X} + \pi(1 - \theta) \left[\alpha X + \sigma \left(\sigma^+ Y + pN - (p - \ell)\theta \left(\frac{X^* - \hat{X}^*}{X^*} \right) N - pX \right) \right],$$

where the first term corresponds to E's consumption following the boom, which happens with probability $1 - \pi$; the second term corresponds to E's consumption following the recession if he is unproductive, which happens with probability $\pi\theta$; and the third term corresponds to E's consumption following the recession if he is productive, which happens with probability $\pi(1 - \theta)$. Notice that the third term includes the return from investing surplus funds in his short-term technology between dates 1 and 2.

To sum up: we are interested in a symmetric equilibrium (dropping the asterisks) where q , X , Y , \hat{X} and C solve (4.1)-(4.6). One can show that this is the unique equilibrium if the following condition holds:

$$(A.3) \quad \alpha\sigma > \theta\alpha\sigma(1 - \ell) + \theta\alpha + (1 - \theta)\sigma. \quad 10$$

¹⁰ An additional condition on σ^+ is needed, to ensure that (4.3) holds. Details are available from the authors.

Note that for $\theta = \frac{1}{2}$, (A.3) is satisfied by the numerical example at the end of Section 2.

The left-hand side of (A.3) reflects the marginal opportunity cost of insurance: the fact that an entrepreneur E has to pay the premium out of his date 0 investment funds.

The right-hand side of (A.3) gives the marginal benefits to E of taking out insurance against the recession. The first two terms reflect the fact that, with insurance, he will not default as much: the first term reflects the concomitant drop in his downpayment q at date 0, and the second term reflects his date 2 return from being able to pay for more intermediate product to be delivered. The final term on the right-hand side reflects the fact that if he is productive in the recession then he will spend the insurance payout on additional short-term investment between dates 1 and 2.

It is useful to look at the comparative statics of (A.3). The condition is more likely to hold as α , σ or l rise, or as θ falls. The dominant effect of a rise in either the long-term return α or the short-term return σ is to push up the opportunity cost of paying the insurance premium at date 0. If the liquidation value l is high (near to p), then insurance doesn't cause the downpayment q to drop by much. Finally, if the probability θ of being unproductive in the recession (and hence of defaulting) is low, then insurance doesn't bring much benefit either.

To sum up this section: there are circumstances where agents will choose not to insure against fluctuations in their accounts receivable, even when such insurance is available. As a result, there can be a chain reaction of default in a recession, which serves to exaggerate the extent to which the economy responds to negative shocks.

Appendix: Justifying the restrictions on contracting

In this Appendix, we make some additional assumptions about an entrepreneur E's technologies to justify the four restrictions on contracting made in Section 2.

First, we assume that the final output of both E's short-term and his long-term technologies accrues directly to him. Moreover, E can divert, or steal, this output: there is no sanction against theft. In consequence, he is unable to mortgage any of his investment returns. (Restriction 1.)

We make a number of assumptions about E's long-term technology. These are perhaps easiest to understand if we start with the second stage, between dates 1 and 2, where the technology is Leontief: if, at date 1, E inputs an equal amount, X/k say, of k intermediate products then at date 2 he makes αX goods, where $\alpha > 1$.

The k intermediate products are made at the first stage, between dates 0 and 1. Each requires the labor of someone other than E to make. To this end, at date 0, E provides k other people -- his suppliers -- each with a key, which is necessary to make intermediate product. Aside from the key itself, production exhibits constant returns: one unit of labor produces one unit of intermediate input at date 1. E is only endowed with k keys. No-one can work with more than one key, and no key can be shared. In effect, E's k intermediate products can only be supplied by k different people, and none by E himself. We assume that k is large enough for divisibility not to be an issue, but small compared with the sizes of the two populations.

We make one further assumption. Instead of delivering the intermediate product at date 1, a supplier S can costlessly make and deliver a fake product, which is of no value to E, but which outsiders (such as the courts) cannot distinguish from the real thing. As a result, there is no point in E signing a supply contract with S at date 0 specifying an amount that he must pay S on delivery at date 1 -- because S could always deliver the fake product and appeal to the courts to collect the contractually-agreed payment from E. Thus explicit contracts are useless. (Restriction 2.)

This does not mean that E and S cannot do business, however. At date 1, S and the other $k-1$ suppliers to E each has some amount of intermediate product, X/k , which, collectively, is worth αX to E, but only ℓX to them. We assume that the parties bargain over the quantity and terms of trade. Specifically, they negotiate a price p per unit, and E takes delivery of as much intermediate product as he can afford at that price (but of course no more than the supply X that has been made between dates 0 and 1). If E cannot afford to buy all of X at this price, then the sellers liquidate the remainder. We are assuming equal treatment across the sellers: given E's Leontief technology, all k the sellers are in same bargaining position. (Restriction 3.) We make no attempt to specify the details of the bargaining other than to suppose that p is a fixed parameter p satisfying $\ell < p \leq 1$; that is, we assume an interior bargaining solution. At date 0, anticipating the bargain that will be struck at date 1, E and S agree on some additional amount, qX/k goods, that E must pay S, where the price q reflects market conditions.

E and S have an implicit contract. At date 0, E gives S a key and orders a quantity X/k of one of his k intermediate products. At the same time, E pays S an amount qX/k goods. At date 1, E buys up to X/k units of the intermediate product at a price p . Any remaining intermediate product can be liquidated by S for ℓ per unit.

Finally, consider the position of E as a supplier of intermediate product (to some other entrepreneur). Any attempt by E to raise funds at date 0 from a third party (deep pocket) in return for handing over the intermediate product at date 1 is thwarted by the fact that E can deliver a fake product. (Restriction 4.)