Discussion

Douglas Gale

In this elegant and well-crafted paper, Gobert, González, Lai, and Poitevin address an important subject, the allocation of liquidity and its impact on the stability of the economy. In an Arrow-Debreu world with complete markets, every commodity is perfectly liquid, and liquidity can be taken for granted. By contrast, in a world with incomplete markets, the allocation of liquidity may be far from optimal. The authors highlight a particular source of market failure, arguing that when markets are incomplete, the market value of a firm does not reflect the value of future liquidity services it can provide to the market. As a result, the decision to continue or terminate a firm may be inefficient. Furthermore, when markets are incomplete, inefficient bankruptcy decisions have multiplier effects that can be interpreted as a form of financial fragility.

One of the paper's features that I particularly liked is that it provides a genuine general-equilibrium analysis of liquidity provision. This is essential, because the termination decisions of individual firms help to determine and are, in turn, determined by the aggregate supply of liquidity. Another attractive feature of the paper is the central role played by the valuation of the firm. It is obvious that a firm's market value is crucial in determining whether it will fail or continue, but the future solvency of the firm is also an important determinant of its market value. Moreover, the firm's continuing solvency depends on the future supply of liquidity, which, in turn, depends on expectations about the solvency of the firm in the still more distant future. So, to determine what happens in the present, we have to consider possibilities of bankruptcy, valuation, and liquidity supply in an infinite regress. It is only in a general-equilibrium setting that one can properly study the interaction of liquidity and asset pricing.

A third attractive feature of the paper is its emphasis on the normative as well as the positive aspects of policy. Welfare economics is the microeconomic foundation of good policy, and here we are provided with a thorough analysis of the welfare economics of liquidity provision.

General-equilibrium analysis can become very intractable, and quickly, so in the interests of tractability, the authors strip their model down to the essentials, eliminating institutional details as far as possible. This seems to be the right strategy for a first cut at the problem: keep it simple and add complications one at a time. Their model captures some of the essential features of liquidity provision and the exposition is clear. At the same time, one may wonder whether essential parts of the story are missing. In particular, the authors eliminate all the usual functions of “financial institutions” and represent them by revenue streams, what I call “firms,” rather like Lucas’s “fruit trees.” Whether they have gone too far in simplifying the story is a point to which I shall return.

The paper is organized around the analysis of equilibrium valuation and bankruptcy in four institutional settings. The first is autarky, in which firms have no access to external finance. The second corresponds to perfect capital markets, which provide unlimited liquidity subject to a present-value budget constraint. The third is identified with centralized decision making and the fourth with decentralized decision making in incomplete markets. The focus of the paper is on the last case, which represents the closest approximation to actual economies, in which firms have limited access to external funds. In this case, Gobert et al. can show that bankruptcy decisions are inefficient and argue that they represent a form of market failure or financial fragility.

To clarify what I see as the essential features of the model and the argument of the paper, it will be helpful to introduce a simple matchbox-sized example in section 1 and use it in section 2 to illustrate each of the four regimes mentioned above. In section 3, I change the example to illustrate the possibility of “financial fragility.”

In section 4, I return to questions raised by the model and the analysis. In particular, I will suggest that there may be simple decentralized solutions to the market failures discussed in the paper, even if one is restricted to the use of spot markets.

1 The Model

The authors assume that there is a finite number of firms $i = 1, \ldots, n$ operating at an infinite sequence of dates $t = 1, 2, \ldots$. The revenue of firm i at date t is denoted by $y_i(s_t)$, where s_t is a random variable representing exogenous uncertainty at that date. The “state of nature” in this economy is a realization of the sequence of identically, independently distributed (i.i.d.) random variables $\{s_t\}_{t=1}^{\infty}$ and determines the firms' revenues at each date.

Revenues can be positive or negative. If $y_i(s_t) < 0$, then firm i cannot continue unless it can obtain external finance to make up the deficit. It is important to note that a negative value of $y_i(s_t)$ does not represent a debt that must be paid. It is, rather, an investment that must be made in order for the firm to continue in operation. The investment can be avoided if the firm is closed down.

There is neither accumulation nor storage and there are no assets other than the firms, so the only source of external finance for a firm in deficit is the revenues of firms that are not in deficit. The aggregate liquidity of the economy is sufficient to keep all firms going if (and only if)

$$\sum_{i=1}^{n} y_i(s_t) \geq 0 .$$

If this inequality is violated, some firms will have to be closed down. Once a firm is closed, it is gone forever and its future revenue will be lost. The important question is whether the right firms are closed down.

To fix ideas, consider the following matchbox-sized example. At each date, Nature tosses a fair coin to determine the state. The state takes two values, Heads (*H*) or Tails (*T*), with equal probabilities, so

$$s_t = \begin{cases} H & \text{w. pr. } 0.5 \\ T & \text{w. pr. } 0.5 \end{cases}$$

at each date *t*. There are two firms, labelled $i = H, T$. The revenues of firm $i = H, T$ at date *t* are given by

$$y_i(s_t) = \begin{cases} 2 & if \quad s_t = i \\ -1 & if \quad s_t \neq i. \end{cases}$$

In other words, if the coin comes up Heads, firm *H* earns \$2 and firm *T* loses \$1. If the coin comes up Tails, the situation is reversed. The states (and hence the revenues) are i.i.d. over time. Note that the revenues of the two firms are *perfectly negatively correlated*—when *H* makes a loss, *T* makes a profit, and vice versa—so there is never an aggregate shortage of liquidity.

2 Liquidity Provision

Now let us consider our example in a series of institutional settings.

2.1 Autarky

The first institutional setting considered by Gobert et al. is autarky, in which firms have no access to external finance and must be self-sufficient at each date. Since the revenues are i.i.d., the law of large numbers ensures that, with probability one, the firm eventually makes a loss. By definition, there is no source of external finance in autarky, so a firm with negative revenue must go bankrupt. Thus, each firm goes bankrupt in finite time.

2.2 Unlimited liquidity

Now suppose that there is a perfectly elastic supply of capital, meaning that a firm can obtain external finance as long as it can repay it at the fair interest rate $R = (1-\beta)/\beta$. The expected revenue in each period is

$$E[y_i(s_t)] = (0.5)\times 2 + (0.5)\times(-1) = 0.5,$$

so the present value of expected future earnings is

$$\frac{0.5}{R} = \frac{\beta}{1-\beta}(0.5).$$

The firm's expected net present value (NPV), calculated at the interest rate *R*, is given by

$$NPV = y_i(s_t) + \frac{\beta}{1-\beta}(0.5).$$

For β close to 1 (*R* close to 0), the *NPV* must be positive in any state: even if the firm loses $1 this period, it can borrow the money and repay the loan with interest. It is therefore optimal and feasible to continue. In this setting, because the individual firms are always solvent (in a present-value sense), they continue operating forever.

Note that this argument, based on a calculation of *NPV*, does not specify the form of the financial contract that will be used. It cannot be a standard debt contract, however. Debt accumulates every time the firm makes a loss, and since there is a positive probability of a very long sequence of losses, we cannot rule out the possibility that the accumulated debt and interest will exceed the firm's *NPV* in finite time, in which case the firm is technically bankrupt. So the assumption of an unlimited supply of liquidity implies the use of complex state-contingent financial contracts.

2.3 A centralized solution to the allocation problem

The third institutional setting considered by the authors is a centralized allocation. Suppose that the two firms form a coalition. The combined revenues of the coalition are

$$\sum_{i = H, L} y_i(s_t) = 2 - 1 = 1$$

in each state, at every date, so it is certainly feasible for the two firms to continue forever. Furthermore, it is optimal for the firms to continue indefinitely, because the *NPV* of each firm is positive (assuming, as usual, that R is not too large). Once again, therefore, it is feasible and optimal to finance both firms forever.

The internal structure of the coalition is not examined; once again we do not know what the financial contract is, but we can be sure that it requires complex, state-contingent transfers between the two "firms."

The coalition can be given different institutional interpretations. The authors' preferred interpretation of this setting appears to be a "central planning" model, but in this example, it could equally well be the result of a merger or an acquisition. There are also decentralized interpretations (see below).

2.4 A decentralized solution to the allocation problem

The last setting corresponds to an imperfect capital market. Suppose that firms can borrow and lend at a rate $R = \dfrac{\beta}{1-\beta}$ at each date.

At this rate, the successful firm should be willing to lend to the unsuccessful firm, and the unsuccessful firm can repay the loan in terms of present expected value.

It is implied that both firms survive forever and the decentralized allocation is the first best. Note, however, that simple debt contracts cannot be used, for the reasons pointed out above.

3 Financial Fragility

In the example we have been discussing, both firms have positive *NPV* at every date and in every state, so there is no ambiguity about whether the firms should continue at each date. Regardless of the institutional setting, both firms should and will continue. If we change the numbers slightly, however, we can see differences among the outcomes in different settings and the possibility of inefficient bankruptcies or terminations.

Suppose that firm *H* has the same revenue function as before, but that firm *T* has the new revenue function defined by

$$y_T(s_t) = \begin{cases} 1 & \text{if } s_t = T \\ -1 & \text{if } s_t = H. \end{cases}$$

Consider the revenues of the "merged" firm in the centralized solution:

$$y_H(s_t) + y_T(s_t) = \begin{cases} 1 & \text{if } s_t = H \\ 0 & \text{if } s_t = T. \end{cases}$$

Clearly, it is still feasible and optimal to maintain both firms in the central planning solution. In a decentralized economy, firm *T* must fail, however. To see this, suppose to the contrary that the firm survives in both states. Then the expected revenue of firm *T* is 0, so now firm *T*'s *NPV* in state $s_t = H$ is negative:

$$NPV = -1 + \frac{\beta}{1-\beta} \times 0 < 0 .$$

Thus, it cannot be possible for firm *T* to survive indefinitely.

Things are not good for firm *H*, either. Once firm *T* has failed, firm *H* is left in autarky and, as we have seen, neither firm can survive there. Consequently, both firms go bankrupt with probability 1 in finite time.

Note that the analysis of the *NPV* of the firm assumes that the bankruptcy decision is stationary (time-invariant). If firm *T* survives in state *H* at one date, it must survive in that state at every date. If we conclude that the firm always fails in state *H*, then the value of the firm is *increased*, because we can eliminate the losses in state *H*. Recall that the deficit in state *H* is not a debt that must be repaid but rather an investment that is required for the firm to continue. So, conditional on failure in state *H*, firm *T*'s revenues are given by

$$\hat{y}_T(s_t) = \begin{cases} 1 & \text{if } s_t = T \\ 0 & \text{if } s_t = H, \end{cases}$$

and the expected present value of future revenue is

$$\beta\frac{1}{2} + \beta^2\left(\frac{1}{2}\right)^2 + \ldots = \frac{\beta/2}{1-\beta/2} = \frac{\beta}{2-\beta}$$

and the *NPV* in state *H* is

$$NPV = -1 + \frac{\beta}{2-\beta} < 0,$$

as long as $\beta < 1$. Thus, although the option value of keeping the firm going is positive, it is not enough to repay the investment of $1 in the event of a loss at date *t*. This is because the valuation of firm *T* pays attention only to the present value of future revenues (using the risk-neutral discount factor $\beta/(1-\beta)$), and does not take into account the value of providing future liquidity to firm *H*.

The market "undervalues" the weaker firm *T*, because it does not take account of the value of its liquidity provision to firm *H*. This is what the authors characterize as *financial fragility.* Because the market does not value the contribution of firm *T* to the liquidity of the economy, it does not provide finance to firm *T* when it is insolvent, according to the negative *NPV* calculation. But without firm *T*, firm *H* cannot survive, and in the long run, there is a loss of value that is greater than what is required to keep firm *T* going.

4 Questions

4.1 In what sense is this a model of financial fragility?

There is a resource constraint that requires some firms to close in some states of nature. This would be true under any institutional framework as long as the aggregate liquidity constraint is binding. The problem identified by the authors is that, under some circumstances, the bankruptcy decision may be inefficient. The fact that this is a general-equilibrium phenomenon and that the failure or anticipated failure of a firm has implications for the solvency of other firms gives it the flavour of a contagion model. If firm *A* fails, there may not be enough liquidity to keep firm *B* going, but if firm *B* fails, there may not be enough liquidity to keep firm *C* going, and so on. But this phenomenon is characteristic of any other general-equilibrium system with incomplete markets, where one firm's failure has "multiplier" effects on other firms.

More importantly, although the inefficiency of bankruptcy has "financial" roots, in the sense that it arises from the incompleteness of markets available for risk sharing and intertemporal smoothing, it is not clear in what sense we should see the firms themselves as representing a financial system. What is it about the firms in this model that identifies them as "financial institutions" rather than, say, industrial firms? If there is something that makes the financial system more fragile than, for example, the logging industry or the car manufacturing industry, it is not modelled here. Stripping a model down to the essentials can be a fruitful strategy, but in doing so, there is a risk of eliminating essential features.

4.2 What defines the set of admissible financial contracts?

The argument of the paper runs in terms of present values without much thought to the form of the financial contracts that are being used to finance the firms' deficits. This may seem innocuous in view of the risk-neutrality assumption, but it needs to be handled with care. For example, suppose that one tries to use simple debt to finance a deficit. If firm T has a loss in some period, it will have to borrow \$1 and will owe $\$(1+R)$ the next period. If it experiences another loss, it will have to repay $\$(1+R)+\$(1+R)^2$ the next, and so on. With positive probability, the amount the firm owes can be made as large as desired, certainly greater than the expected present value of the future revenues. What is going on here? When we use expected present-value calculations to decide whether a firm can afford to borrow and finance a current deficit, we are implicitly assuming that the firm can write a complex state-contingent contract in which the repayments depend on the firm's ability to repay. Simple debt contracts will not do this: either the lender will not receive enough, or the borrower will be forced into bankruptcy in some states. Why should this matter? Well, if firms are able to write complex, state-contingent financial contracts, it is not clear why the markets in the decentralized solution should be incomplete in the first place. A more formal statement of the financial contracting problem would help to make clear what contracts are allowed, what contracts are implicitly ruled out, and what the justification is for drawing the boundary where it lies. This might be a useful extension to the analysis in the paper.

A related question is why a firm's value is *assumed* to be the present value of the revenue stream. A present-value calculation is, after all, just a special case of the valuation that occurs in an Arrow-Debreu economy. If the theory assumes that state-contingent contracts are being used to provide external finance to a deficit firm, why aren't state-contingent prices used to value the firm's future revenue stream, as well?

4.3 What happens if accumulation is allowed?

One of the special features of the model is that accumulation of assets is not allowed. This is a simplifying assumption, obviously, but one that is somewhat restrictive. Firms do retain earnings and hold reserves precisely to avoid liquidity problems. One interpretation of the model might be as a theory of corporate finance. If there is an external finance constraint, one would expect firms to hold reserves as part of their corporate financial strategy. This will not solve the problem of incomplete markets entirely (though it is well known that, as the discount rate R approaches zero, self-insurance is equivalent to complete markets), but it may be an important element of the story.

4.4 What happens if equity can be traded?

Gobert et al. characterize the decentralized allocation as the allocation that can be achieved using spot markets only, whereas the centralized allocation corresponds to what can be achieved with complete forward markets. As I have indicated, the capital market that exists is not exactly a spot market: it contains elements of a forward market or, at least, of a long-term financial contract. I want to suggest, however, that introducing a spot market for equity might be sufficient in the present context to improve the allocation of liquidity substantially and perhaps achieve the first best.

Consider a firm that is insolvent and is about to be closed because it cannot borrow enough to survive. If the bankruptcy is inefficient, that is, if the remaining firms can potentially be made better off by financing the deficit and keeping the firm going as a source of liquidity, why don't they buy the firm's equity, make the necessary investment, and share the future revenues? It is not obvious that this will capture all the gains from liquidity, because the future demand for liquidity is state-contingent, and it is not clear that the firms that buy the equity today will match their own demand perfectly without further trade. However, there do seem to be gains from trade that are not exploited here, and it is not clear that a world in which all gains from trading equity in spot markets were exploited would not approximate the first best.

4.5 How would such equity be priced?

In any competitive model, the market price of an asset is determined by the agents who value it most, and the marginal value of an asset to an individual is calculated using the individual's personal state-contingent prices (marginal rates of substitution). To revert to our earlier example, if firm *H* were thinking of buying the equity of firm *T*, it is not clear why the price at which it would be willing to buy the equity would be given by the present-value rule. Even if we assume that firm *H* is maximizing the present value of its revenues from all sources, by investing in firm *T* today, it is changing its future revenues in two ways. First, there is the revenue received directly from firm *T*. Second, there is the revenue that comes from extending the life of firm *H* beyond the point where the bankruptcy constraint would otherwise bind. Any sensible buyer would add the shadow value of relaxing this constraint to the market value of firm *T*'s equity.

This example may be too simple: there are only two (types of) firms, so the buyer can internalize the liquidity value of the equity it is buying. The argument is not quite so simple when there are more firms. With three firms, for example, the firm with the cash to buy the equity of the distressed firm today might not have a need for liquidity in the future and so will not value the equity of the firm at more than its present value. However, if there is a third firm that is willing to buy the equity for more than the present value at some time in the future, the first firm should discount this "liquidity premium" in the price that it is willing to pay today. So, properly functioning spot markets for equity should be able to price all of the liquidity services provided by the distressed firm, without resorting to forward markets.