Discussion

Blake LeBaron

In securities markets, we observe a wide range of trading institutions. Have these markets evolved to handle specific types of buyers and sellers, or does the observed variety in markets reflect other institutional constraints or accidents of history? To address this question, we need to know something about the optimality of market institutions themselves, and who prefers particular market characteristics. It would be appealing if one optimal market structure existed for most trading situations, but as this paper reminds us, this is often not the case. Furthermore, the paper describes the conditions under which different market institutions should appear.

Trading situations currently exist between two extremes. One case involves highly transparent markets with generally open order books. Orders are fully revealed, and traders operate with full information on market location and depth. Many implications of increasing electronic trading appear to be moving towards this type of system. At the other extreme are dealer markets where trades are handled primarily through intermediaries. Here, information about order flow tends to remain hidden from the market. Traders can observe only the current set of spreads from the set of broker dealers. They do not receive information on the other orders in the system.

It may initially appear that more information, and therefore the order-book market, would be preferred, but as the authors demonstrate, this is not always the case. They model these two contrasting markets using a limit-order double-auction framework and a broker-dealer setup. Both markets are placed in a similar multi-round trading environment where agents receive an individual inventory shock that they bring to the market. Comparisons can then be made between the two types of market.

In many other studies, expediency and analytic tractability have forced the authors to use a small number of agent trading strategies. In this case, there might be one agent strategy and one dealer. The Audet, Gravelle, and Yang paper breaks from tradition by using a computational agent-based approach. These types of models analyze the overall economic dynamics from the bottom up, starting with strategies for the individuals and then analyzing the macroeconomic outcomes through computer simulations of the strategies dropped into the market institutions. The authors use a nonlinear functional form—an artificial neural network—to describe the agents' strategies. This flexible form should be able to capture many possible strategies that the players might try to implement.

What form these strategies should take, however, is an important and difficult question. The nonlinear functions have many parameters to be determined. The authors do this by using a very impressive algorithm for finding a Nash equilibrium in the nonlinear trading strategies. Technologically, this algorithm is an interesting piece of work. Once the strategies are determined, the agents can be sent into the various markets, where prices and performance are determined.

The markets are compared by changing the number of outside traders entering the market. For each number of traders, the authors ask which market would be preferred in terms of agent welfare. As the number of traders increases, the limit-order-book market is generally preferred, and as the number of traders falls, dealer markets are preferred.

The break-even point, at which traders are indifferent, is recorded and is used as an indicator for how likely it is that a dealer market would occur. If this break-even number is low, then dealer structures are probably unlikely, and if it is high, then it is much more probable that the dealer structure will occur.

Much of the paper examines the impact on this break-even number for differing parameter values. The authors find that markets with large correlated order flows tend to use a dealer-based system. Reducing the number of dealers will tend to move markets in the direction of order-book markets. Finally, reducing the risk aversion of dealers will increase the likelihood of a dealer-based market. Certainly, the first of these conclusions seems to connect with some real-world behaviour. One example might be the upstairs block trading markets on the New York Stock Exchange. The other results seem to make sense as well.

I think this framework is extremely useful, and I have some ideas for extensions. It would be helpful to know more about how nonlinear the strategies are. In the end, do traders generally stay with fairly simple linear

structures for their strategies? The experiments are run with fixed numbers of dealers. Clearly, the dealer numbers matter. Would it be possible to model potential entry and exit of dealers, making the dealer numbers endogenous? These simulations handle only liquidity shocks. Market institutions must also be good at handling general information shocks when the value of the asset has changed. This is a thorny issue, but it will eventually need to be tackled.

The performance of these markets is evaluated in terms of utility. Obviously, this is not a commonly used measure of market quality in the real world. However, these simulated markets generate values on spreads and depth, which can then be compared with real-world estimates. It would be interesting to see how closely the utility measures line up with other measures of market quality.

Finally, in terms of agent-based modelling strategies, it is important to note that the Nash equilibrium framework used here differs from many agent-based approaches. In many agent simulations, the market dynamics are analyzed out of equilibrium. It is often the actual convergence or lack of convergence that is the key issue under study. This could be done in these markets as well. However, in games where agents participate frequently, there is a case to be made for initially examining only Nash equilibria.

Market microstructure may be one of the critical areas for the use of agent-based technologies in finance. Breaking out of traditional moulds of homogeneous (or limited heterogeneous) agent worlds into a rich set of behaviours that better mimics the real world is crucial to obtaining a good picture of the properties of these market institutions. The computer also allows exploration of market conditions and settings that are not limited by analytic tractability, and are therefore closer to reality. The policy questions on market structure are pressing and important. One hopes that studies such as this will begin to give policy-makers new tools for designing markets that are both efficient and robust to a changing world.