Discussion

Robert White

It is a pleasure to be here and a great honour to comment on a paper prepared by such distinguished authors. This paper is well-written and ambitious, and it addresses a very important issue in an actively researched area. As we saw from the papers presented earlier today, the subject concerns an industry that is undergoing tremendous change. The authors have raised a lot of questions, and my comments are meant either to build on what they have presented or to serve as suggestions for future work.

The question addressed by the paper is whether market quality has been enhanced or degraded by making the consolidated order book more transparent. The authors used several measures of market quality: quoted spreads, effective spreads, and volatility. The quoted spread is the inside quote on the book, and the effective spread is the change in the traded price relative to an equilibrium price. The authors used a number of univariate and multivariate tests to answer this question. The event that they examine occurred on 12 April 1990, when the order book was opened up somewhat, not completely, but by four ticks. Therefore, instead of simply having the market quote being shown, it was four ticks out on either side. And at that time, the tick was one-eighth, which means 50 cents up and 50 cents down.

But why was it four ticks? I expect it was because market orders on the floor walked out the book four ticks.[1] That is, large market orders often consumed

---

1. Walking out the book four ticks refers to the quantity of limit orders consumed by a given market order, in price space. When buying or selling shares, a person can submit a limit order, i.e., a price-quantity combination, which is relatively passive, since it simply goes onto the order book. A market order is more aggressive, however; it transacts with existing limit orders at the best price(s). Thus, when a market order "walks out the book," it is of such a magnitude that it consumes existing limit orders and moves the market price out on one side of the market.

limit orders and moved the price by four ticks. It was simply a market-making mechanism. And I believe this is an important issue that was not raised.

The authors used a control group from the computer-aided trading system (CATS). However, market orders on CATS and market orders on the floor were treated differently. The former do not walk the book. They can only go right into the book. That can have quite a different impact on performance.

Instead of the authors' "control-group," I would use the term "quasi control group." For example, the stocks listed on the exchange in the CATS trading venue are quite different from those on the floor. They tend to be smaller, more illiquid, thinly traded firms. The more actively traded, larger firms are on the floor. Another difference is that the floor walks out four ticks and CATS walks out only one tick. Moreover, the Registered Trader, the specialist equivalent, differs on this venue (versus the floor venue).

The authors' main conclusions are that the costs have been augmented and that volatility has increased as a result of increasing the disclosure on the book by four ticks. The authors conclude that market quality is inversely related to the degree of transparency—the more opaque the market, the better. And I believe this to be a reasonable conclusion.

One should be cautious about interpreting these results, however. This was a single event, and other factors could have played a role. In other words, you have to control for everything but that single event, and that is what the authors attempt to do with the control sample. Consequently, the differential impacts are important.

The early 1990s were characterized by significant events such as the Gulf War. Canada was on the cusp of a major recession. These kinds of events would be impounded into the price and they would most likely have an impact on the spreads and on volatility. The question is, do they, in fact, have a differential impact across the two sets of securities—the authors' control group and their test group. The mix of the securities in the two groups and the qualities of those securities are very important.

While I haven't gone back and re-estimated all of their data, I have looked at data from the spring of 1990. Basically, the February-March period is the authors' pre-period, and May-June is their test period. The patterns of returns on the overall index are drastically different and that is why they want the control sample. But can one obtain this perfect control sample and make it work?

The sample can be split into one index for securities trading under two dollars and another index for trading over two dollars. (The small firms are typically under two dollars.) The price level and firm size are highly correlated. But the pattern of returns for the small firms with the low-priced securities is dramatically different in the test period than it is for the large firms with the high-priced securities. That would lead one to find the floor securities to be of a higher cost, which is what the authors found. It doesn't prove anything, but it does raise questions that the authors might want to explore.

The volume of trading also has some bearing. What is interesting is that the trading bottoms out in April, which is the time of their event. Now, there are two reasons why this might be significant. If there is a trend in the volume of trading, then the issue of the four-ticks market order might become important, and it could drive the volatility of the floor to a greater extent than the volatility of CATS because of the ability of the market orders to walk the book. So, if there is a trend, one might expect to find pressure in that direction. The other sort of thing that differs between CATS and the floor is that the floor has an open outcry. Thus, there is a possibility for imitation trading on the floor. Papers have been published indicating that imitation trading is real and, consequently, any trends would increase volatility.

I would like to make some suggestions regarding an additional test for the quoted spreads. The narrower the quoted spread, the more efficient the market is, theoretically, and the lower the potential transactions costs. Another way of testing for the quoted spreads is to examine the duration when the bid-ask spread is wide and when it is narrow. Industry calls this a time-weighted average, and it is used as a statistic to compare securities, to compare markets, and to look at the difference and means tests to discover which market is more efficient.

One other statistic that is widely used is price continuity. Trading is all about being able to make trades at the desired quantities without significantly moving the market. One measure that the industry uses is price continuity, which is the percentage of the trade that is within five cents of the previous trade. The higher that percentage, the better the market. The TSE, for example, runs about 97 per cent price continuity, which is quite efficient. Once again, a difference and means test could be set up to determine whether or not a change in price continuity is significant.

One could take a variant of the authors' model in order to use the spreads as the dependent variable and to use trade-by-trade data (because the authors want to test for volatility) to set up a GARCH[2] model with a variance

---

2. GARCH: generalized autoregressive conditional heteroscedasticity.

equation in it. One would put a shift variable in the price equation for the floor versus CATS, and then put a shift variable in the variance equation to indicate whether the variances have increased or decreased. It seems to me that this becomes a more powerful test, since one can now control explicitly for information effects, size of orders and firms, and price volatility on a trade-by-trade basis.

Another problem with the authors' results is that they are based on an average equation. The average price is typically correlated with the average firm size, but firm size is also highly correlated with average volume. One would therefore expect the average price and the average volume to be highly correlated. If they are, that indicates an econometric issue that could distort the results. A correlation matrix would indicate whether or not the use of the average price is valid.

These comments are simply meant to be suggestions on extensions. Overall, the paper is well-written and the conclusions are reasonable.