Our clients would never know the difference anyway

Peter A. Forsyth*

May 7, 2018

If you have used an on-line broker, you will have observed the following situation. Suppose you put in an order to sell 1000 shares of company X at 10:00 am. A few minutes after you put in the order, you click on the *order status* button. You see the following

10:01:20 200 shares X sold at $20.02

10:01:50 100 shares X sold at $20.00

10:03:04 400 shares X sold at $20.05

10:04:20 300 shares X sold at $19.99

10:04:20 Order completed

So, what's going on here? You are observing the result of using an *optimal trade execution* algorithm. If you dump 1000 shares on the market all at once, you will depress the price (temporarily) and the average price you get for your 1000 shares will be less than the pre-trade price (the arrival price). This is due to the price impact caused by rapid trading. This is also known as *walking down the order book*. So, the idea is to break up your order into smaller lots, and spread out the sells. This way, you hope to get a higher average price for your shares, by minimizing the price impact.

However, if you trade slower, there is always the risk that the stock you are selling will move randomly lower. So, a trade execution algorithm attempts to balance the reward (reduced price impact from slow trading) and risk (random price movements).

I got interested in this topic a few years ago. It seems clear that the stock seller wants to (i) maximize the total amount obtained from selling her stocks and (ii) minimize the effect of random movements in the stock price, relative to the arrival price. The obvious way to measure (ii) is by the variance of the total amount of cash received from the sale of the stocks. At the end of the day, that's all the seller really cares about: maximize the total cash received, and minimize the uncertainty in the total cash.

However, when I looked at the actual algorithms used in practice, I noticed something odd. Instead of looking at the variance of the cash received at the end of trading, the risk measure used was based on the average volatility of the stock holding during the trading period. Now, obviously, these two concepts are related, but they are not quite the same thing.[1]

I developed an algorithm based on measuring risk in terms of variance of the total cash received. In many cases, the results were fairly close to the industry standard approach. However, in some

*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, paforsyt@uwaterloo.ca, +1 519 888 4567 ext. 34415.

[1]Technically, this is difference between variance of the total cash received at the end of trading, and the quadratic variation of the trading portfolio.

cases, if you all you really cared about was the total cash received, the industry standard method was significantly sub-optimal.

I was very excited about this, and then talked to one of the quants[2] involved in trade execution at a major bank. He agreed with the basic mathematics: the average volatility was in fact different from the variance of the total cash. However, this guy made an interesting point. He liked the average volatility measure, because he had to manage a large book of trade orders. It was easier to keep track of what was going on during the day if he kept an eye on the average volatility of all his unfilled orders. This was a convenient measure of risk from his perspective.

I managed to get on the technical program of a *practitioner* conference, i.e. mostly bankers, not academics. [3] I gave a talk pointing out the fact that the standard algorithm was sub-optimal from the client point of view. During the question period after the talk, one of the quants from a well-known investment bank reiterated the point I had heard before, i.e. the average volatility was useful from the bank's perspective. Then, in a room full of about 200 quants, he summarized his comments:

> *"Sure, I agree that in general this is not the optimal strategy for the client. But most of our clients are not sophisticated enough to notice the difference."*

Note to self: don't assume that investment banks are acting in your best interests.

[2]Short for Quantitative Analyst. Usually requires an MMath or PhD, and the ability to act knowledgeable even if you don't have a clue as to what is going on.

[3]Here is the citation for the final paper: S.T. Tse, P.A. Forsyth, J.S. Kennedy, H. Windcliff, "Comparison between the mean variance optimal and mean quadratic variation optimal trading strategies,"Applied Mathematical Finance 20 (2013) 415-449.