Transaction Cost Analysis for Futures

By Greg Wood

A futures broker sits down with his pension fund client for a quarterly review of their recent trading activity. The conversation starts off in the usual way with a brief discussion of market conditions, but soon turns to topics such as "shortfall," "negative trend," "trading alpha" and "reversion." A trader writes a complicated formula on the back of a napkin, confers with his boss and passes it across the table. "This is how you calculate the impact of our order on the market, isn't it?"

This isn't a conversation about trading equities, where such terms are used frequently. This is a conversation about trading futures—indices, bonds, commodities, the full spectrum of products available electronically—and the subject is being discussed more often as the buy-side evaluates more efficient execution.

There are two main reasons for this. First, clients rely less on brokers to manage the execution of their orders. As futures markets have become electronic, customers can now use algorithms to do what brokers have traditionally done for them.

Second, hedge funds, commodity trading advisors and real money managers are consolidating their trading desks, so it is now common for individual traders to execute across multiple asset classes. When this happens, they expect to use the same toolset regardless of asset class. Tools include execution management systems that provide routing to multiple brokers as well as the execution algorithms that the brokers provide.

Execution algorithms are designed to help preserve trading alpha—broadly defined as a return in excess of the market during the execution timeframe. Execution algorithms are not designed to generate trading decisions—the "what," "why" and "when" —but rather the "how" so as to minimize execution risk that could negate any trading alpha identified with the trade idea.

Measuring Market Impact

Every trade has an impact on the market. Regardless of the size of the trade, it is going to incur a cost simply by its presence in the marketplace and how it accesses liquidity. It is also possible to signal intentions and pay more for a trade than intended. This is especially true in the futures markets where there is no concept of trading in the dark. For example, an order to buy 1,000 contracts of Comex gold futures, shown in the central limit book on Globex, will influence the market. The offer will rise as market participants become aware of an order that is larger than normal. This increase in the offer represents a cost to the trade. Liquid markets can absorb the impact of trades without the cost registering until the trade size becomes noticeably large relative to other orders in the market. Less liquid markets will not absorb the impact so easily, which can leave a lasting effect by moving the market from its prior level. To minimize this impact a trader needs to reduce any signalling of intent. They can work the order manually in smaller clips or use an execution algorithm.

The trader can chose to schedule the execution of the 1,000 contracts over a period of time using VWAP or TWAP, or use a mechanical algorithm such as an iceberg. Algorithms that target VWAP (volume-weighted average price) will feed more orders to the market during high volume periods such as the open and close and work less at quieter periods such as lunch time. An algorithm targeting TWAP (time-weighted average price) will feed orders to the market evenly over a designated period of time.

Another approach is to use an intelligent stealth algorithm that minimizes signalling risk by never posting to the central limit book. Instead it takes prices that are within range and works hard to blend in with other participants' activity in the market.

Another type of algorithm is designed to reduce implementation shortfall (also known as "slippage"). Implementation shortfall is a measure of performance versus the "arrival price" or the price at which the order entered the market. For futures, this is typically defined as the far side of the market because spreads are often just one tick wide. A trader who buys at the offer "pays" the spread (the difference between the bid and the offer) to trade; his shortfall is zero. If they buy at a tick above the arrival price, then the shortfall is negative since the market has moved away from the arrival price. If they bought at the bid, they actually "captured" the spread and incurred positive shortfall. An example of an algorithm that seeks to reduce implementation shortfall is one that has the expectation that prices are mean-reverting, or will typically fluctuate within a small range. This will become more aggressive when the market is favorable compared to the arrival price and

Alpha Profile

A sample trade profile for a buy order demonstrating the concepts of implementation shortfall and reversion in graphical form.

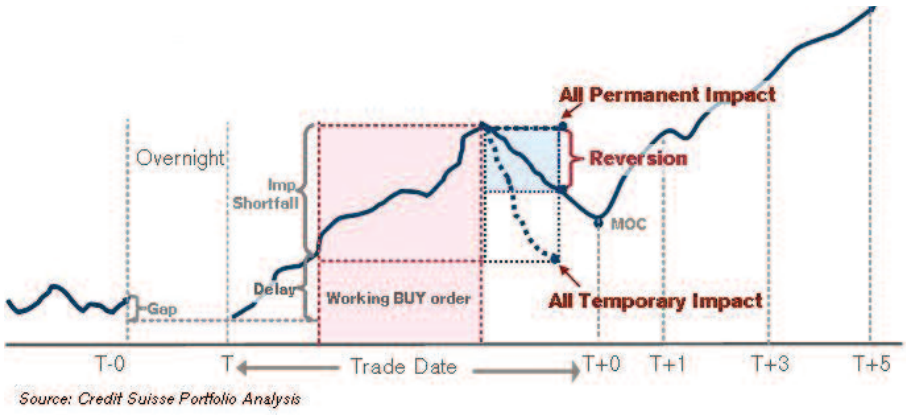

Source: Credit Suisse Portfolio Analysis

more passive when the market is moving away. The danger here is that the algorithm may not be able to complete the order if the market continues trending away from the arrival price.

The concepts of shortfall and reversion apply to any type of financial market. Because these concepts can be quantified, algorithms can be used to attempt to minimize any of these factors, and their performance measured.

Implementation Shortfall

Implementation shortfall is a well established industry standard approach to measuring transaction costs. It includes anything that has an effect on a financial instrument's price—notably trend and impact, amongst other factors.

If the underlying trend during your trade is negative (i.e. a rising market when buying, a falling market when selling) then part of shortfall will be a cost attributable to trend. Similarly, if the trend is positive (i.e. a falling market when buying, a rising market when selling) then this part of shortfall will be a saving.

Impact is always a cost created by a trade's presence in the market. An example of how impact can be defined would be the difference between the trade's arrival price and the sum of all other trades on the market —a relatively straightforward formula that could be written on a napkin.

Different types of algorithms approach trend and impact in different ways. For example, an algorithm can be instructed to avoid crossing the spread or to be a very small part of the overall volume traded on the market. This stretches out the duration of the order, minimizes impact and increases trading alpha—but there is a risk that the order may not be completed in time. If it is imperative to complete the order then you may need to risk higher impact by using a more aggressive algorithm or instead select a VWAP or TWAP algorithm that will trade when the market is unfavourable just to remain on

Test-Driving an Algo

Transaction Cost Analysis, commonly known as TCA, is spreading from equities across all asset classes. The idea of being able to quantify execution and measure how it performed against the market is not new, but the tools to provide such quantification are now becoming common in all asset classes. Here is a hypothetical example to show how a customer might apply TCA to futures trading.

You call your broker and ask about the algorithms it provides. They sound enticing, but scary. How do I know which to use? Wasn't there bad press last year about algorithms during the Flash Crash? How do I know that my trade is safe and won't cause some major impact on the market?

Your broker runs through the algos available, what they are supposed to do, and in which markets they work best. Then comes the dealmaker—you can review the performance on a TCA report. Start off with some smaller orders, look at the reports, learn what parameters work best and build confidence.

Back at the trading desk, you place some small orders through the system and discuss the results with the broker. Then it's time to really give it a trial. You need to sell a few thousand long gilt futures on Liffe from the open. It's a largish trade, but not too large, but just to be safe you spread it over a few hours and ensure that it is never more than 20% of the market. Let's see how this works. While the order is in motion you're wondering whether it is doing what you want it to. Is it selling the highs and avoiding those temporary dips in the market? Is it participating when there's more volume? Is it sitting back when there's less volume? Is it signalling what you're trying to do? What will it do if the market trends away?

The following day you get an email from your broker with the TCA report. The report is normalized so as to allow ready comparison across different markets. Values are quoted in basis points, monetary amounts in U.S. dollars. The average bid/ask spread on the long gilt futures contract is shown as 1.09 basis points of the notional. Makes sense—you do the math and it comes out as 1 tick, or £10. You look at the other values. Implementation shortfall is shown as -5.1 bps—that's the measure of execution versus the arrival price (the price at the start of the trade which in this case was the opening print of the day). Your shortfall is made up of trend, impact and trading alpha. OK the market trended down immediately after the open, so that explains the large negative value. "I was expecting that," you think.

Here's the clincher though. How did you compare to the volume-weighted average price over the duration of the trade? It shows that you sold at a price 0.5 basis points lower than the VWAP benchmark.

Hmm. Hadn't expected that. Let's take a closer look. You chose a VWAP algorithm for your trade because you wanted to stretch out the order, guarantee completion but never be more than 20% of the market volume. And because you knew that the market was likely to trend down, you told it to be more aggressive than usual. An aggressive VWAP algorithm is typically more likely to pay the spread when trading rather than trying to capture spread.

Your broker walks you through what the market did during your trade. You can see the overall trend down across the duration of your order, with a little bit of choppiness here and there. The volume profile on the day of your order was comparable with other days, so the VWAP algorithm was not disadvantaged by any unusual volume on the day. It participated fairly evenly across the entire duration according to the expected volume profile, but was never more than 2% of the traded volume because of the lengthy duration of the order.

schedule. Such an algorithm, however, will not look to minimize implementation shortfall (see "Test-Driving an Algo" below).

Reversion

An interesting measure of the impact of a trade on the market is how the market reverts afterwards. Reversion can be measured over different periods of time to identify if the impact of a trade was temporary or permanent.

If the price reverts all the way back to the price at the start of the trade, then impact was temporary and likely to be influenced by this trade. In this case reversion represents a cost since it had too much impact, either because it was a large order that could not be easily absorbed by market supply and demand, or because it was trading too aggressively by crossing the spread and pushing the price away.

If the price after a finite period remains the same as the price at completion of the trade, then impact can be considered permanent. This is likely to happen when trading occurs based on genuine news or other events that revalue the market, rather than a short-term demand for liquidity.

Most reversion falls somewhere in-between. If a trader can minimize reversion, along with implementation shortfall, then they can maximize the alpha preserved by the trade. The optimum time to wait to measure reversion after the trade has completed depends on the size of the trade. In practice, the larger the trade the more impact it has on the market, so for larger trades the timeframe for measuring reversion should be lengthened in order to find where the equilibrium occurs between temporary and permanent impact.

Putting It into Practice

The concept of measuring execution is not new and traders have been applying these principles for many years. The ability to quantify a trade compared with the market before, during and after its time in the market is an extraordinarily powerful tool. It can provide insight for the buy-side on how they are trading and whether they are using the right approach for their objectives, i.e., whether they are using the best type of algorithm for a particular trade.

By evaluating results over time, clients can refine their usage and gain knowledge of which tools to use. An aggressive stealth algorithm may be good for getting a reasonably sized Treasury futures trade done efficiently without showing intent, but a passive implementation shortfall algorithm is better suited to a choppy commodity market where the trader wants to participate more when the market is favourable and less when not. VWAP may be passé in cash equities, but it is quite serviceable for the well-defined and consistent volume profile on the E-mini S&P 500 futures—especially when hedging a portfolio of stocks.

As the buy-side looks to quantify its futures execution, it means that the sell-side has to understand these concepts as well. As satisfying as it is to chat with customers about golf, it's even more satisfying to tell a client that they saved 0.57 basis points in shortfall over a month. If that sounds geeky then think about this … that number comes from a real-world example of a money manager that used algorithmic tools to improve its executions in fixed income futures. And the 0.57 basis points saved equated to a difference of $1 million in the value of its portfolio. That's almost as good as an eagle on the golf course. ■

Greg Wood is vice president and business development manager for futures within the Advanced Execution Services business of Credit Suisse.

TCA

A sample Transaction Cost Analysis report demonstrating the various execution measurements broken down by side, order duration and sector.

Type	Contracts	Tickets	Percent Complete [%]	% ADV	Participation Rate [%]	Imp Shtfall [bps]	Trend Cost [bps]	Impact Cost [bps]	Trading Alpha [bps]	vs. Interval VWAP [bps]	Bid-Ask Spread [bps]
Business Breakdown											
Total	**121,666**	**261**	**98.95**	**1.93**	**28.06**	**0.57**	**0.56**	**-0.11**	**0.11**	**0.21**	**0.96**
Futures	121,666	261	98.95	1.93	28.06	0.57	0.56	-0.11	0.11	0.21	0.96
Side Breakdown											
Total	**121,666**	**261**	**98.95**	**1.93**	**28.06**	**0.57**	**0.56**	**-0.11**	**0.11**	**0.21**	**0.96**
SELL	60,918	133	99.43	0.11	27.60	0.75	0.71	-0.12	0.16	0.35	0.95
BUY	60,748	128	98.47	3.73	28.52	0.39	0.42	-0.09	0.07	0.08	0.96
Duration Breakdown											
Total	**121,666**	**261**	**98.95**	**1.93**	**28.06**	**0.57**	**0.56**	**-0.11**	**0.11**	**0.21**	**0.96**
0 - 5 mins	107,118	236	99.41	2.20	29.90	0.49	0.56	-0.11	0.04	0.02	0.94
5 - 10 mins	3,712	8	98.29	0.09	12.95	0.95	0.37	-0.06	0.63	0.68	1.15
10 - 30 mins	9,086	15	94.15	0.13	18.43	1.29	0.75	-0.06	0.60	1.69	1.01
0.5 - 1 hrs	1,750	2	100.00	0.14	4.53	0.57	-0.24	-0.06	0.87	1.98	0.95
Sector Breakdown											
Total	**121,666**	**261**	**98.95**	**1.93**	**28.06**	**0.57**	**0.56**	**-0.11**	**0.11**	**0.21**	**0.96**
BOND	118,336	243	98.89	2.03	26.84	0.59	0.58	-0.10	0.12	0.22	0.98
INTEREST RATE	3,330	18	100.00	0.15	50.02	0.16	0.27	-0.13	0.02	0.05	0.58

Source: Credit Suisse Portfolio Analysis

That 0.5 bps cost compared to the actual VWAP could largely be attributed to crossing the spread a little more than usual. A less aggressive order would have sat back more and tried to capture spread rather than paying spread. Still, had you paid the spread on every one of those several thousand contracts in your order then you would have paid 1.09 bps on average. On a trade of £1 billion notional that's a cost of £109,000 to get your trade done. You actually paid only 0.5 bps. You effectively saved half a tick across your trade. Isn't that what you try to do with a block trade? Of course, a block trade gives you the certainty of a single execution price, but it comes at a premium in commission. Here is a quantified alternative for executing a large trade.

Lessons learned from the experience? You didn't need to be as aggressive as you thought. The VWAP kept up with market trend over the duration and you were never more than 2% of the volume. Had it been allowed to sit back more you could have reduced cost by a few fractions of a basis point—and every one of those equates to precious dollars saved in execution costs.