Benchmarking passive limit orders

By Dr Darren Toulson, Head of Research, LiquidMetrix

A key ingredient in most modern electronic trading algorithms is the effective use of passive lit limit orders. The aim is to capture a portion of the bid offer spread. Doing this consistently at the fill level will help Broker algorithms match or over-perform certain TCA benchmarks, particularly VWAP related ones, where crossing the spread leads to almost inevitable underperformance.

The increased use of passive lit strategies by Broker algorithms has blurred the distinction between traditional market makers and traders. The old academic models of simple utilitarian market makers providing two way liquidity to a mixture of informed and uninformed traders as a 'service' for which they are paid the bid-offer spread has been replaced by a new reality with many different types of participants competing to receive the benefit of being liquidity providers.

Given this new complexity, it's important that Broker algorithms maximise the potential benefits of using limit orders (spread capture) whilst minimising

Dr Darren Toulson

the risks (opportunity costs, increased price risk, adverse selection). To do this, the more sophisticated Broker algorithms will use a mixture of short term alpha/volatility models and BBO chasing/crossing strategies to optimise the net contributions of passive lit orders to performance.

However, as with any complex trading algorithm, post trade measurement of the performance contribution of passive limit order strategies is vital. For the Sell Side this post trade analysis provides a feedback loop that can be used to further optimise limit order strategies. For Buy Sides, having visibility of the 'real' contribution of passive limit orders to overall algorithm performance will help assure them that their Brokers aren't over-using passive limit orders in algorithms for fee reductions/rebates.

Measuring limit order performance

The goal of using passive limit orders is to capture spread. However, if we simply measure the spread captured at the time of execution, we get a misleading view of the contribution of limit orders to overall order performance. Consider a practical example:

Imagine running a VWAP algorithm that needs to buy 500 shares in some time segment. Assume current EBBO prices are bid 99/offer 100 and consider four possible trading outcomes:

1. Immediately trade 500 shares at the aggressive lit market price of **100**.
2. Post 500 shares passively at the best bid price of 99 and get filled almost immediately – achieving an improved price of **99**.
3. Post a bid at 99 but the bid subsequently rises. We 'chase' (repost) up to a

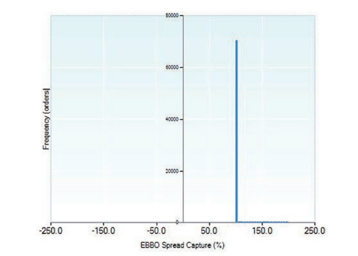

(A) Spread capture: 100%

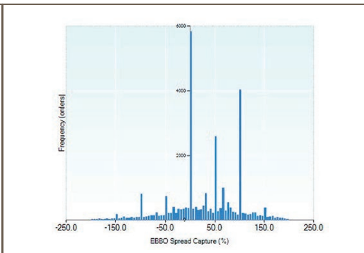

(B) Spread Capture: 29%

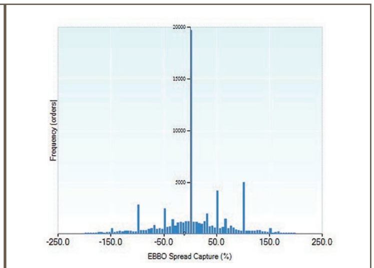

(C) Spread Capture: 6.9%

price of 101, before finally getting filled at **101**.

4. Post at 99, the market bid rises to 102, but after 10 minutes we're still unexecuted. To maintain our VWAP schedule, we're forced to cancel our passive order and aggress the market, crossing the spread and executing at **103**.

The key to properly assessing the performance of a limit order posting strategies is:

- The final execution price of any limit order should be **compared against (1. above)**, i.e. the price we could have obtained by aggressing the market at the point of order submission; NOT at the time of execution (which always looks good!).
- If any of our limit orders fail to fully fill, we must account for opportunity cost. The best way to do this is to assume that remaining unexecuted shares must be purchased at the current aggressive EBBO price.

By taking these two factors into account, we can measure the true benefits of using passive limit orders.

The effect of being realistic is shown in Diagrams (A) to (C) taken from real trading data.

(A) shows EBBO spread capture measured at the time of execution for a passive only execution strategy, which results in **100%** spread capture.

(B) shows the spread capture if we use the market BBO at the time of order submission rather than execution, but still ignoring unexecuted limit orders. Already the average spread capture is lower (**29%**).

(C) shows the spread capture histogram when also including opportunity costs from unexecuted limit orders. The result is now very different: just **6.9%** spread capture. However, this improvement is **real** – any positive spread capture using this methodology means that use of limit orders was a net positive in terms of overall trading performance.

Conclusion

Clever use of limit orders to capture bid-offer spread is a vital part of modern electronic trading algorithms. However, there are pitfalls to using limit orders, in particular opportunity costs when limit orders fail to execute. Properly measuring the performance of limit orders, post trade, allows Sell Sides to optimise their trading algorithms and Buy Sides to be assured that use of limit orders is contributing to order performance and not simply reducing fees. ■

LiquidMetrix
Best Execution, Quantified