pragma

NO. 5 | APRIL, 2013

The Limits of Price Prediction Algorithms

Agency trading algorithms are given an order and strive for best execution in the form of a good average price and high fill rate. Within the basic constraints of the order, they can make intelligent trading decisions based on historical and real-time market data as well as any number of external data sources. For example, they can decide whether to use limit or market orders, route to lit or dark trading venues, speed up or slow down the rate of trading, etc. Over the past few years, price prediction—often referred to as "HFT signals"—has become a staple of agency algorithm marketing pitches, presumably since high-frequency traders are widely reputed to have the best short-term price prediction in the marketplace. Almost universally these marketing pitches lack specifics, but the impression is of powerful price prediction over a period of minutes at least. When faced with questions, vendors often invoke the "secret sauce" defense—that providing specifics might compromise the intellectual property of the provider.

Price prediction is a broad term, and is central to many investment strategies. At the extreme, high-frequency traders use signals such as trade and quote arrivals and order book state in different exchanges to decide when and where to post or take liquidity over a timescale of seconds or milliseconds. However, execution algorithms fundamentally differ from their proprietary cousins in that they cannot choose what instrument to trade and on which side of the market they want to be.

In this note, we consider execution algorithms that use price prediction to speed up or slow down trading over time scales of a few minutes. We provide a framework to quantify the benefit that can be expected from such price prediction, and we find that some reasonable assumptions point to it being small. We also touch on another potential obstacle to properly evaluating the performance improvement provided by price prediction: the optionality introduced when completing an order is not mandatory.

Other kinds of signals—in particular very short-term signals that might help inform an algorithm where or how to route an order—are not considered here.

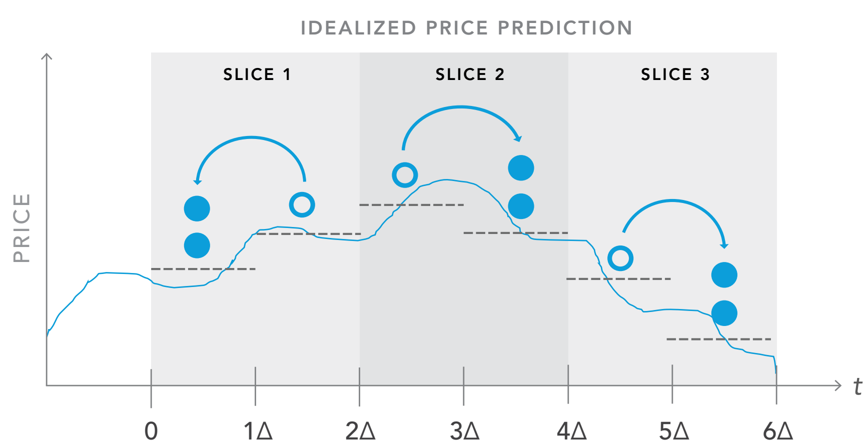

FIGURE 1
In our stylized trading scenario, a signal predicts prices over a horizon 2Δ. The execution algo accordingly trades (filled circles) or not (empty circles) on each half of the prediction horizon.

Bounding Signal Performance: Optimistic Case

Our approach is to obtain an upper bound on the performance improvement that can be achieved for an execution algorithm by predicting future price moves. For this purpose we devise an optimistic stylized trading scenario that captures the most important constraints faced by an execution algorithm. In our scenario, we have an algorithm that follows a baseline TWAP schedule with a target of Q shares every (for example) 30 minutes. (All the results are applicable to VWAP and other similar algorithms with very minimal changes.) We also have a signal that can predict returns 30 minutes into the future. In particular, it can predict whether the average price over the next 15 minutes will be lower or higher than the average price over the subsequent 15 minutes. The algorithm uses this information to decide whether it should trade Q shares over the first or the second half of the next 30-minute interval. This process is repeated every 30 minutes, as depicted in Figure 1 on the previous page.

We can motivate this scenario as follows. First, we assume most execution algorithms must stay reasonably close to a benchmark-oriented schedule such as VWAP or TWAP to limit execution risk. Although our exercise here is to presuppose that short-term price prediction is possible, we assume that long-term price prediction to the horizon of hours or a full day is very weak at best, and that departing too far from a benchmark-oriented trading schedule introduces unacceptable execution price risk for a large category of order flow. Thus, if a signal predicts a positive return, a buying algo might speed up a bit, but should not execute everything immediately—the risk of being wrong on the global timeframe is too high. Second, liquidity is finite. Even if our algorithm wanted to trade everything very quickly, it would incur significant price impact. Therefore we cap our peak trading rate at 2× the average without introducing any penalty for this speedup—unrealistically generous, but useful for the purpose of developing a bound on performance improvement.

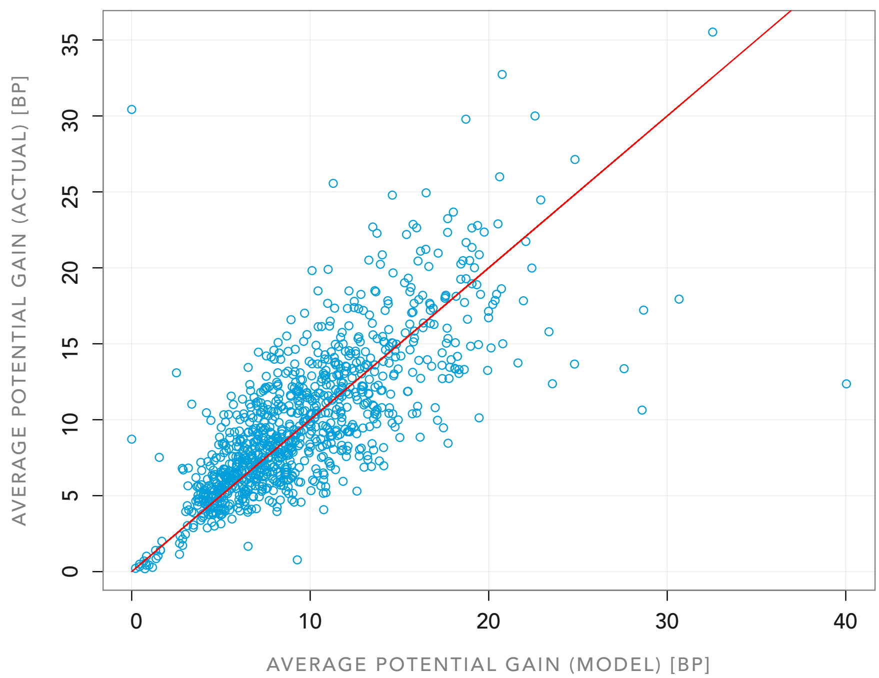

FIGURE 2
Gain predicted by the Brownian motion model vs. hypothetical gain computed from market data, using a perfect signal.

Now we can ask by how much the prediction-enhanced algorithm outperforms the TWAP benchmark. We model the price as a Brownian motion with volatility σ. Then the difference X between the average prices over the first and second halves has a Gaussian distribution with zero mean and a standard deviation $\sigma\sqrt{2T/3}$, where T = 15 minutes. Suppose for now that the prediction is always right. Then the average improvement over TWAP is the mean of $|X|/2$, which is $\sigma\sqrt{T/3\pi}$. For a daily volatility of 1.5% and our example of T = 15 minutes, this works out to a potential improvement of 9.6 basis points.

Figure 2 depicts the fit of actual potential improvements against this model. This data set includes the top 1000 US stocks by market cap in June-July 2012. We estimated their volatilities and used the model to compute the potential improvement. We also simulated our trading strategy using actual market data for the same data set, again assuming no market impact. The figure shows the model's estimated improvement and the empirical potential improvement on the X and Y axes respectively. There is good agreement—hardly a surprise, since we are in effect just testing the validity of the popular random walk model.

Introducing Prediction Errors: A More Plausible Bound

A prediction that is always right is, of course, an unachievable ideal. More realistically, a prediction will only be right p percent of the time. Our expected gain is as before, times a factor of (2p-1). Based on the prediction, we will always shift our trading as described above—since there is no penalty for doubling our rate of trading, as long as $p > .5$ using our signal will lead to an improved expected result. Figure 3, below, depicts the relationship between the accuracy of the prediction p and the performance improvement in basis points for three prediction horizons.

This illustrates graphically the fundamental insight of the model—for a given volatility, the two key parameters that govern the value of price prediction are the accuracy of the prediction and the time horizon over which it operates. What may be counterintuitive is that even given the optimistic assumptions of the model, for plausible ranges of values the benefits from price prediction are quite modest. Many experienced traders will find even a 55% accuracy across all stocks in all prediction periods an enviable record, yet with a 2 minute prediction horizon this corresponds to only ¼ basis point performance improvement. And this still relies on the unrealistic assumption that there is no price impact penalty for doubling the trading rate.

An independent check on the estimated value of price prediction comes from a simple no-arbitrage argument. If short-term price prediction were powerful enough to exceed the typical agency algorithm TWAP shortfall of 1-2bp, trade execution would actually be a profitable proposition and reliable source of alpha. The absence of such experience, or even credible specific claims in the marketplace for execution algorithms that reliably beat the benchmark are consistent with the range of performance improvements suggested by this model across intuitively reasonable values for p and prediction horizons.

This model is intended primarily to provide a quantitative framework to help buyside traders evaluate the credibility of the marketing claims of their agency algo providers. It's worth noting that this doesn't imply that proprietary trading strategies cannot make profitable use of short-term price prediction. In contrast with an execution algorithm that must complete every order regardless of the strength of the signal, proprietary-trading strategies can scan a large universe of securities and trade only those that have strong signals. In other words, they can focus on the tails of the distribution, whereas an agency algorithm must average out large benefits from strong signals with the much more common negligible benefits from weak signals.

FIGURE 3

Model-predicted gain for signals predicting returns over 2, 6 and 12 minutes for different levels of signal accuracy. We assume 1.5% daily volatility.

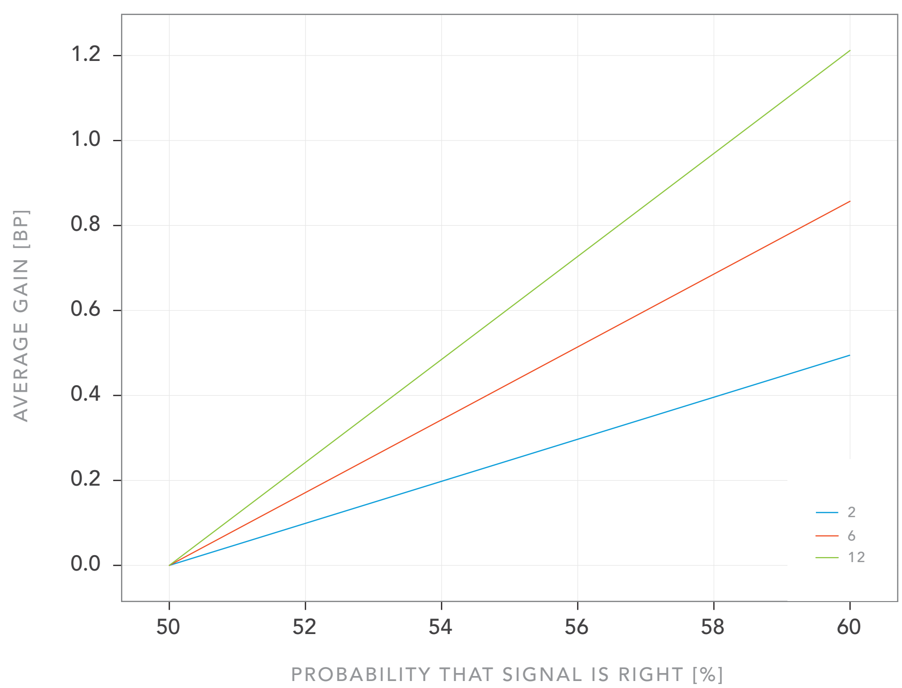

FIGURE 4

Hypothetical price trajectory of a buy order with no guaranteed execution. Whenever the price is worse than the interval VWAP benchmark, the algorithm does not trade. The algorithm beats the benchmark but the price runs away and the order does not complete.

The Illusion of Optionality

This last observation also highlights a pitfall of evaluating the performance of algorithms that claim to have predictive powers. There *is* a way an execution strategy can reliably beat a benchmark such as VWAP—by sacrificing mandatory completion of the order execution. A simple example of such a strategy is to start trading at the arrival price, but then to continue trading only if and when the market price is favorable relative to the benchmark. If instead the price runs away, the algorithm simply refuses to trade and may ultimately fail to complete the order unless the price comes back. Thus by construction the average price of the executed portion will always appear favorable relative to the benchmark. This is an extreme example but similar, more subtle effects can come from even perfectly well-intentioned price prediction algorithms which don't guarantee completion of the order. Such effects can create a powerful illusion of good performance because while completed trades look good against standard performance benchmarks, measuring the opportunity cost of letting the winners run away unexecuted is challenging. Though hard to measure, the negative effects on investment performance are very real.

This effect is similar to trading with limit prices. The existence of a limit price can effectively become an instruction to the broker to trade only when the execution looks favorable. This common practice in the institutional trading environment is a major obstacle to evaluating or comparing brokers, as shortfall gets hidden in the opportunity cost embedded between waves, which many TCA products don't properly track.

Conclusion

In this note we have provided a framework for quantifying the value that short-term price prediction can bring to agency-style execution algorithms, and shown that under reasonable assumptions of imperfect price prediction on the timescale of a few minutes, price prediction is unlikely to yield large expected performance improvement. We have also identified a dangerous illusion of performance improvement that may be introduced when execution is not mandatory.

This framework highlights critical questions that an algorithm consumer should require providers to answer when evaluating the credibility of their marketing claims—even if the details of the signal must remain secret, what is the time horizon of the prediction; what is the accuracy; and what is the average performance improvement across all orders when execution is guaranteed? The answers to these questions may be quite different than the impression conjured up by marketing terms like "HFT signals."

For questions or comments, please email Dr. Eran Fishler, Chief Operating Officer (technotes@pragmatrading.com).