Beyond Broker Scorecards: What's Next for Algo Wheels?

"Everyone gets the experience, some get the lesson"

-T.S. Elliot

Trading involves a bewildering number of individual choices. Each order could be traded in a huge number of different ways and every market participant gets to witness the outcomes. Everyone gets the experience, but do they also get the lesson? As an industry, do we structure our decision making so that we can best evaluate which strategies are working?

As you might suspect our answer to this question would be a firm no. Empirical methods for continuously improving processes and decision making have not been nearly as widely adopted by the trading community as they have in industries like healthcare, advertising, and communications.

To date, algo wheels have been the most visible evidence that the philosophy of continuous process improvement has been making inroads into the trading community. By submitting broker selection to randomized controlled trials (RCTs), buy-side firms can generate unbiased broker performance statistics. This leads both to improved trading performance in the short term, and in the long term, stronger incentives for brokers to invest in their execution capabilities. This article surveys the current algo wheel landscape; looks at what the trading community can learn from applications of RCTs in other industries; and anticipates future developments in this space.

Surveying the Algo Wheel Landscape

At FlexTrade we've been implementing custom and productized wheels for the better part of a decade, and our FlexAlgoWheel product has been designed based on our extensive experience across a highly diverse range of clients and workflows. While there are several different kinds of vendor wheels available, (and we've been asked to improve on a wide assortment of in-house solutions), they mostly aim at the same results: unbiased assessment of brokers' implementation of narrowly defined algo strategies. While the details differ significantly across vendors and trading desks, all wheel implementations broadly follow the same structure (see Figure 1):

1. **Standardization:** The algo offerings of different brokers are standardized and classified into comparable groups (algo strategies), such as IS, Dark, Liquidity Seeking, etc. In some products this standardization is performed by the wheel vendor, in others by the trading desk themselves (there are advantages and disadvantages to both approaches).

2. **Sampling:** Orders are then routed to an algo strategy and the broker selected by a random process.

Figure 1:
Standard Algo Wheel Implementation

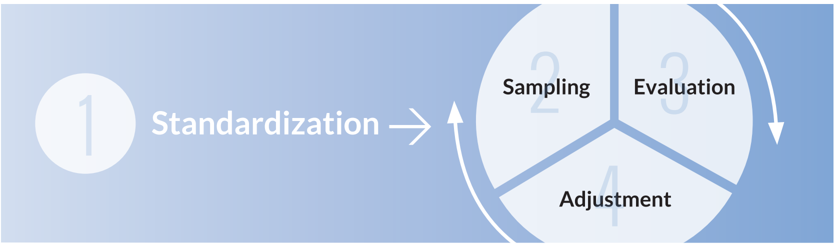

FLEXTRADE
Trade your best.

CONTINUED ON NEXT PAGE

Beyond Broker Scorecards:
What's Next for Algo Wheels?

CONTINUED

3. **Evaluation:** The performance of the orders is evaluated using a TCA platform and a picture obtained of the relative performance of each broker.

4. **Adjustment:** The analysis is then translated into a re-weighting of the wheel (i.e., the fraction of order flow that each broker receives is adjusted up or down depending on how they compare to their peers).

Steps 2-4 are continuously repeated, with a clearer picture of relative broker performance being built up over time as a larger sample of data is collected. Periodically, step 1 is revisited. New brokers are "added to the wheel" and the classification of broker algos into strategies is refined.

Common Pitfalls

There are many different approaches to each of these stages. Even restricting the use of this technology to the (relatively simple) problem of broker evaluation, there are a few common pitfalls:

False Positives: There is often a temptation to make significant adjustments to brokers' allocations based on an insufficient sample of data. Without careful evaluation it is easy to mistake noise (a broker getting lucky in the orders they received, or the opposite) for signal (systematic differences between brokers).

Re-Biasing the Data: Various difficulty adjustments and cost models can be profitably applied during the evaluation stage to reduce the noise in the performance measurements. It is crucial that these models be constructed so that they do not adjust for factors caused by the selection of the broker. For example, a model that adjusts for security specific returns will bias the data in favor of brokers who cause more market impact (the choice of the broker affects the returns on the security).

Overly-Complex Performance Metrics: Algo wheels create powerful incentives for brokers, but this can be a double-edged sword. Broker evaluation can include multiple criteria (e.g., cost versus arrival, completion rate, etc.), however it's important that these be combined in predictable ways to avoid creating perverse incentives. Equally important is that brokers have a relatively clear idea of what they can do to improve their ranking.

Elephant in the Room?

Before going further I'd like to address the elephant in the room: automation. Algo wheels seem to automate what was previously considered one of the traders' core responsibilities: deciding which broker to route an order to. Taking a step back, we see that this task hasn't really been "automated away". True, the trader no longer needs to make this decision at the level of individual orders, but instead their time is spent deciding how to re-weight the algo wheel.

Thousands of small decisions, each taking maybe a second or two, have been exchanged for one decision that could take hours or days (depending on how involved the broker evaluation process is). This pattern is familiar from similar changes in other industries and will likely be repeated as an empirical approach is adopted for more and more trading processes. Small, reflexive tasks will be traded for larger, reflective ones. Unless these reflective tasks are consolidated there will likely be few overall labor-saving benefits from this technology.

Instead the adoption of algo wheels will lead to increased specialization on the trading desk. With low-touch orders handled through a dedicated algo wheel process, traders' attention can be diverted to focus on more difficult orders. With many funds looking to differentiate themselves by increasing their investment in less liquid assets, the importance of reducing transaction costs for these assets can only be expected to increase. Meanwhile the specialization of the low-touch trading function can be expected to increase engagement from the buy side in the creation of customized algos and routing logic.

Learning from Other Industries

The real benefits from this kind of automation will be found in reduced trading costs – ultimately resulting in better fund performance. High-profile websites typically have thousands of experiments running at once, with each user being shown a slightly different design and placement of ads. Data on the user's behavior is then captured, analyzed and used to optimize future decision making. This intensity of experimentation might seem excessive, but it's become standard practice because it contributes significantly to attracting and retaining customers. Given the prevalence of standardization and controlled experimentation, it is reasonable to ask what trading operations could learn from other industries.

One of the most obvious practices that trading operations could adopt from the tech industry is that of using reinforcement learning when running trading experiments, rather than the more traditional A/B test. An A/B test for a website works roughly as follows (see Figure 2, below):

1. Create two versions of the site (A and B).
2. Divide users evenly between the versions, some get shown A and others B.
3. After a set period of time (e.g., three months) evaluate the performance (e.g., in converting visits to sales) of the two versions.
4. Switch to using the version with the best performance, displaying only A or B to users.

The issue that has been found with this approach is that it leaves a lot of money on the table. Information about the relative performance of A and B is being gathered constantly over the three-month period in the example above, but this is only being exploited after the experimentation period is over. Reinforcement learning is designed to solve this problem. There is an extensive research literature in this area but here's a quick summary of one reinforcement learning algorithm called epsilon-greedy (see Figure 3, below):

1. Create two versions of the site (A and B).
2. With some probability epsilon (between 0% and 100%), direct users to the version with the best performance. The rest of the time pick a version at random (with equal probability).
3. Start with epsilon equal to 0%, as the total number of visitors to the site increases, and then increase epsilon.
4. When epsilon reaches 100% the experiment is over and users will only be shown one version of the site.

As you can see this is a little more complex than a simple A/B test. The advantage is that it allows the operator of the website to capitalize on the information they receive as soon as they get it, rather than having to wait until the end of the experiment. Care must be taken translating this approach to the broker selection problem but doing so successfully can significantly boost an algo wheel's contribution to overall trading performance.

Figure 2:
A/B Testing Workflow

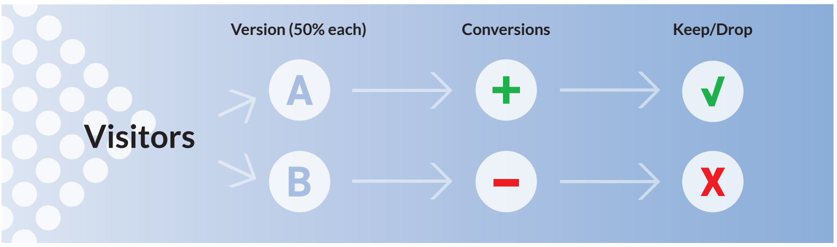

Figure 3:
Reinforcement Learning Workflow

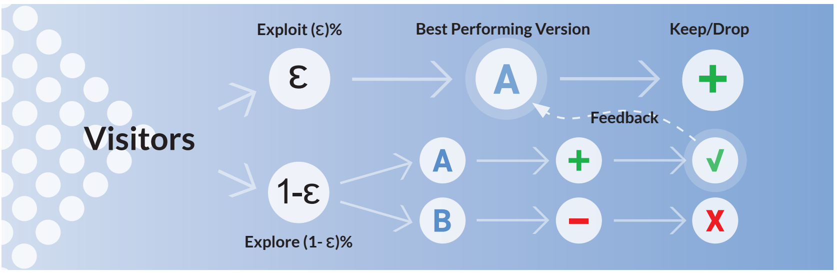

FLEXTRADE
Trade your best.

CONTINUED ON NEXT PAGE

Beyond Broker Scorecards:
What's Next for Algo Wheels?

CONTINUED

Stratified Sampling

Another best practice that should be adopted is that of using stratified sampling. We've already discussed how using insufficiently large samples can lead to false conclusions when analyzing algo wheel data. Post-trade cost models are one way of reducing the amount of data needed, but they come with their own pitfalls. Another, complementary, approach is to use stratified sampling within the algo wheel routing itself. For example, imagine that there are only two kinds of orders: "easy" and "hard". In a simple random sample one broker may receive more "hard" orders than another, leading to a false conclusion that they have worse performance.

However, if the algo wheel ensures that all brokers receive an equal proportion of "easy" and "hard" orders then this is no longer a possibility. The noise in the data due to order difficulty has been eliminated. Of course, implementing these techniques in practice is more complicated but the same logic applies. Stratifying order routing by factors such as the estimated pre-trade cost, volatility, and order size reduces the risk of false positives in the data. It is unsurprising that this has become standard practice among FlexAlgoWheel clients.

Looking Ahead: What Next for Algo Wheels?

Much of the low-hanging fruit for further enhancing performance lies in areas immediately adjacent to the current use of algo wheels. For instance, similar workflows can be used for high touch and program trading. RFQ and principal bid trading can also, with appropriate changes, benefit from an experimental approach. By rotating which counterparties are included in the RFQ or bidding process, unbiased measurements of execution quality and information leakage can be obtained. Again, the dual benefits of improved trading performance and stronger incentives apply. The application of algo wheels to RFQ and other bidding processes opens the door to applying these processes more extensively to FX, Fixed Income and Options trading.

Looking further afield, it's possible to see how an experimental approach can be applied to decisions upstream of broker selection in the trading process. While many brokers claim to be able to guide client's choice of algo strategy to improve performance, the truth is that without using well controlled experiments this advice is likely to be hopelessly inaccurate and may be actively detrimental to performance. Rather than relying on observations of past performance, setting up workflows that randomly allocate flow to different strategies can generate unbiased datasets, which are then used to accurately determine criteria for discriminating between strategies. The same approach can also be applied to algo parameters, such as minimum fill sizes and urgency. Clearly care needs to be taken around segmenting trading flow so that only appropriate strategies are included in the trial, but for most firms the gains to performance would likely be greater than those realized from improved broker selection.

Two of the most common questions that get asked of trading analysts are "Can you demonstrate/debunk the value of block trading?" and "Can you show how PM limits harm/improve performance?" Block liquidity can be a much more efficient way to execute but typically commands premium commission rates. Similarly, disciplined use of limits can be highly advantageous, while poor use can be extremely damaging. In both cases the answers to these questions are: a) likely to be highly contingent; and b) best answered using controlled trials. Again, much of the infrastructure used in algo wheels can be repurposed to address these very different questions. Not only would this result in the settling of some long-running disagreements (much to the relief of those trading analysts), but it would probably also result in some significant cost savings.

Summing Up

The rise of algo wheels has so far followed the familiar pattern of most new technologies. At first, they existed primarily as systems built in-house by a few innovative trading desks. However, in the last two years the first vendor-provided solutions have emerged with wheels now well on their way to becoming a standard tool of most trading desks and included in most RFPs for new Execution Management Systems. We can expect the revenues derived from mass market adoption to fuel a second wave of innovation in algo wheels, resulting in expansion to some of the use cases predicted here.

At FlexTrade our mission is to deliver superior trading technology to our clients, providing them with a significant competitive edge. Getting the lesson rather than just the experience is a significant advantage, and one that we intend to deliver to our clients.

Americas
+516.627.8993
sales@flextrade.com

EMEA
+44 203 757 9320
sales_eu@flextrade.com

Asia
+65 6829 2569
sales_asia@flextrade.com