Anti-Gaming in the OnePipe Optimal Liquidity Network

Pragma Financial Systems

I. Introduction

With careful estimates[1] suggesting that trading in so-called dark pools now constitutes more than 7% of US equity volume, tools to use these new liquidity sources safely and effectively are more important than ever. Of particular concern is the "gaming" of crossing orders — aggressive trading in the open market in order to affect the midpoint quote at which a cross occurs. Figure 1 illustrates the main conclusions of this whitepaper, that gaming does indeed occur frequently, that it can be detected algorithmically and that the Lifeguard™countermeasures employed by OnePipe™are successful in thwarting gamers while preserving high crossing rates from optimal allocation across more than 30 hidden liquidity pools.

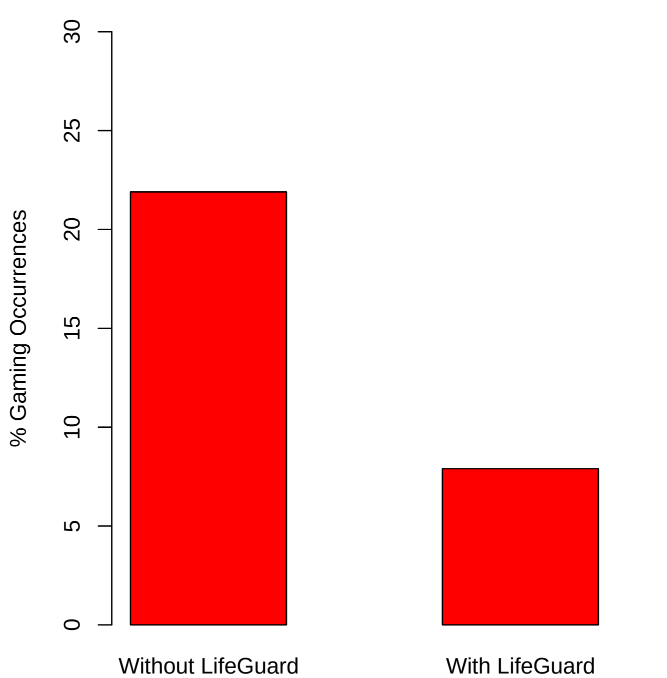

Fig. 1. Gaming rate (fraction of orders experiencing gaming) in OnePipe, with and without Lifeguard. In both cases, crossing rate is approximately 85%.

[1]See Rosenblatt Securities, "Trader Talk", June 2008

II. Classic Gaming – Market Impact Arbitrage

In Edwin Lefèvre's *Reminiscences of a Stock Operator*, the narrator (a thinly disguised Jesse Livermore) describes the operation of "bucket shops" at the beginning of the 20th century. In these betting parlors, clients could place buy or sell orders that settled at whatever price came over the tape from the New York Stock Exchange without any actual shares changing hands. In general, the house profited handsomely by charging a large spread, but sometimes even more money could be extracted by, for example, misreporting the quotes from New York. Livermore describes a time that he got even with a dishonest bucket shop by placing a large bet just after sending an opposite order to the Exchange through a legitimate broker: "Well, you can imagine what happened when the selling order got to the floor of the Exchange; a dull inactive stock that a commission house with out-of-town connections wanted to sell in a hurry. Somebody got cheap stock. But the transaction as it would be printed on the tape was the price that I would pay on my five buying orders."

Modern crossing networks are highly regulated; the trades on them result in a legitimate exchange of assets; and they serve an invaluable role facilitating large block liquidity. In order to provide this liquidity without advertising demand, however, the networks must rely on settlement prices snapped from the open market midprice, and they thus share one critical vulnerability with the old bucket shops.

Gamers generally begin by sending to the crossing network a small "ping" order, which, should it be filled, will reveal the presence of our resting order of the opposite side. While the gamer does not know the exact size of our order, he can make a reasonable inference and begin accumulating a countervailing position in the open market. In doing so, he takes no pains to avoid market impact. To the contrary, if it's a buy, he wants to push the price up as much as possible and then cross the shares back to us, pocketing the difference between his average accumulation price and the peak at which he sells. Figure 2 depicts a real example of a sell order being gamed. The gamer pings an illiquid stock that trades infrequently and detects a resting sell order. He then begins selling the stock rapidly beginning at around 12:34 to push the price down, then buys the stock back from us at 12:36. The gamer then repeats the process two times. Notice that for the third attempt there are no longer enough shares in the crossing network to replace the shares the gamer sold in the market, so he must buy them back in the open market after the cross.

There are other sorts of manipulation possible, for example by placing small quotes inside the spread to briefly change the

Fig. 2. Trading Example: Example of gaming of a crossing order. In the upper pane, the grey band represents the open market NBBO; the grey dots are open market prints; the blue dots are cross fills of our order; and pink dots are other cross prints (reported on the tape as FINRA trades). In the lower pane, the grey line shows open market volume in shares per second, while the blue line shows our crossing rate.

national best bid/offer. This technique is outside the scope of the present paper.

It is important to keep in mind that most of the crossing network trades with which one might be disappointed have nothing whatsoever to do with gaming. Figure 3 shows a trade that could be considered a bad execution but was almost certainly not gamed. In this particular case, the entire market began rallying at around the time the order was placed. In other cases, part of an order may be worked in open markets while the remainder rests in a crossing network. Depending on the liquidity profile of the stock and the trading urgency, this may be a very legitimate strategy, but it will result in statistically high execution shortfalls for the crosses that do occur. When developing a mechanical means of identifying crossing trades that have been gamed, it is valuable to be able to distinguish between trades that have actually been gamed and trades that are merely unfortunate.

III. Quantitative Gaming Metric

The first step in creating anti-gaming defenses is to develop a quantitative measure of whether a completed crossing order has been gamed.

In essence, the metric answers the question, how likely is it that the average trade price we got could have occurred by chance. Imagine that we randomly perturb the observed history of fills, making some of them occur earlier, others later and then recalculating the trade cost by snapping the mid-prices at the new, perturbed fill times. Ideally, we'll find that our actual cost falls somewhere in the middle of a distribution of these imaginary outcomes. If, on the other hand, we're in a significantly adverse percentile, there's reason to be suspicious. This Monte-Carlo experiment can be written as an integral

$$M = \sum_{i=1}^{n} \int_{-\infty}^{\infty} \phi(t_i + \tau, \Phi, \Phi_i) \lambda(p_{t_i+\tau}, p_{t_i}) d\tau, \qquad (1)$$

where the sum is over fills at time t_i, p_t is the mid-price at time t, and ϕ is a weighting function that can depend on an arbitrary set of information Φ associated with the sequence of fills as well as information related only to a specific fill via Φ_i, and λ is a measure of price distance. The value of our metric thus lies in the choice of a weighting function ϕ and distance measure λ. The choices are not trivial, as equation 1 encompasses everything from an infinitesimal time jitter to an arbitrary shuffling of all fill times and sizes.

To develop confidence in our gaming metric, we compare the results to assessments by traders who, like Potter Stewart, know gaming when they see it. Figure 4 shows a histogram of gaming metric values for 326 OnePipe orders; the orders in red were identified by traders as highly likely to be gaming, and as we see they cluster at high values of the gaming metric. It is evident that our mechanical assessment closely matches the traders' intuition.

IV. Gaming Protection

Having developed a means of identifying orders that have been gamed, we now prepare candidate sets of trading rules, which will dictate when an order is permitted to cross. While the gaming metric utilized all data available before, during and after the life of the order, the actual anti-gaming rules applied at any moment will only have access to information available at that moment. Obviously we cannot have a rule that says we refuse to trade if the price will be going our way afterwards; we can only look at the past. The rule can be formalized as a running "fair price", beyond which we will curtail trading,

Fig. 3. Trading Example: Disappointing, but not gaming. See figure 2 for graph key.

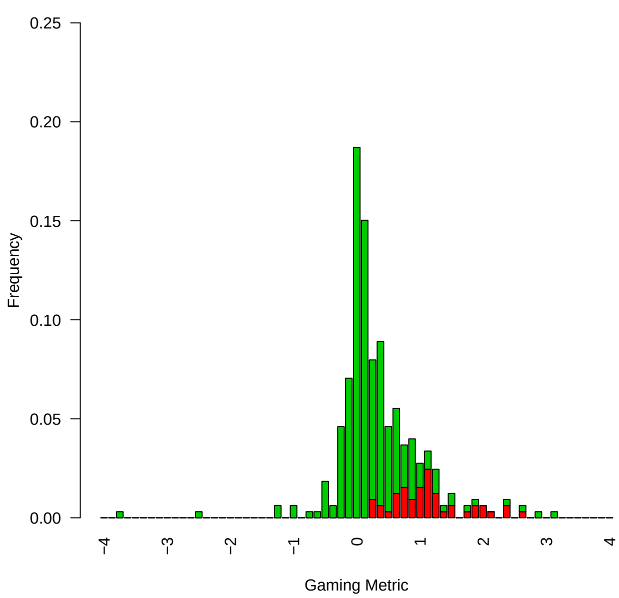

Fig. 4. Histogram of gaming metric values for orders executed without Lifeguard anti-gaming protection. Orders shown in red were identified by traders as highly likely to be gaming.

either by withdrawing from a pool or by setting a limit price. The fair price is a filter

$$p_f(t) = \mathcal{F}(\{p_i, q_i : t_i < t\}, \{\Psi(t') : t' < t\}) \qquad (2)$$

that can depend on all prior fills of quantity q_i and price t_i, and all prior market events $\Psi(t')$. In a sense, the point of equation 2 is to dictate behavior *ex ante* that will result in favorable *ex post* evaluations from equation 1.

We can now assess a candidate trading rule by simulating its application to past crossing orders, admitting only those fills that occurred when the market was better (lower for a buy, higher for a sell) than the fair price $p_f(t)$. For each order, we then compute the gaming metric and the adjusted fill rate. Since this simulation only allows us to accept or suppress a trade, the fill rate can only be reduced; it should in fact be an underestimate of what we would see in live trading, since the suppressed cross could well have occurred later at a more favorable price.

Through experiments like these, we arrived at the proprietary choice of ψ that constitutes Lifeguard. Figure 5 shows a distribution of metric values for actual orders that were protected by Lifeguard. Comparing this histogram to figure 4, the protective effect is clear; with Lifeguard, the prevalence of orders beyond a gaming metric of 1.0 is negligible, and the number of orders judged qualitatively to be gamed is also sharply reduced.

It is important to note that Lifeguard's gaming protection is dynamic, inhibiting trading only when it is important to do so, and that it does not rely on strong assumptions about the way gamers infer the existence of crossing positions. Some competing anti-gaming measures reduce to a policy of setting minimum fill sizes in order to discourage (but not eliminate) pinging. OnePipe does take advantage of minimum fill settings at those destinations that offer them, but it confers active protection regardless of the availability and efficacy of site-specific measures.

Figure 6 is an interesting example of Lifeguard in action. The red bands show periods where the price is above the rolling fair limit price from equation 2; since we do not trade here, there are no blue dots in this region. The pink dots show the reported crossing fills of another buyer, who apparently did not have Lifeguard. As in figure 2, we see small pings, followed by open market accumulation, concluded with a large cross at a local extremum in price. The OnePipe order is

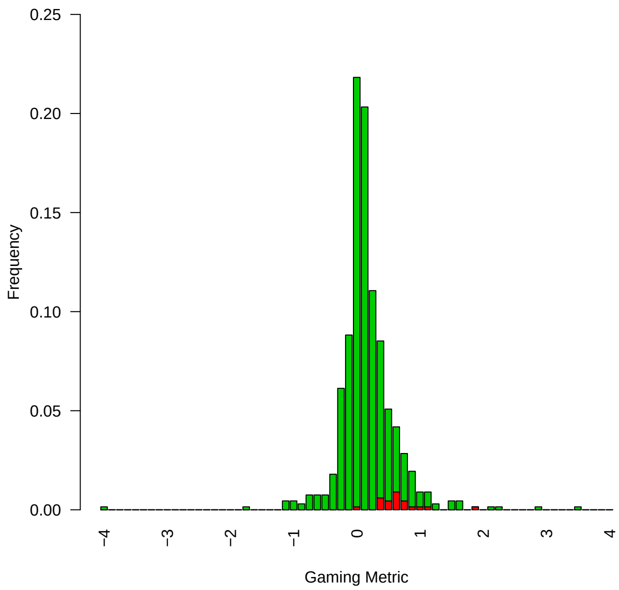

Fig. 5. Histogram of gaming metric values for orders executed with Lifeguard anti-gaming protection.

protected from these adverse fills by Lifeguard, but the non-OnePipe order is gamed.

V. Discussion

Lifeguard dramatically reduces the incidence of gaming while only slightly changing fill rates. Moreover, because it is applied at the aggregation level, it protects all orders in OnePipe, irrespective of the protection or lack thereof at any particular destination. With OnePipe's optimal allocation logic and Lifeguard's anti-gaming protection, traders can access an unprecedented amount of liquidity, and they can do so safely.

Fig. 6. Illustration of Lifeguard at work. The red bands show periods where trading was suppressed; the pink dots show the large crossing fills of a different participant, who got gamed.