

An Algorithmic Approach

Banc of America Securities offers optimal access to dark liquidity

By Michael Kearns

IN THE PAST FEW YEARS, the development of so-called ‘dark pools’ – electronic marketplaces that allow traders to anonymously cross orders without displaying their intentions to other market participants – has had a significant impact on U.S. equity markets. There now exist a large number of such venues, including broker-dealer internal crossing systems and consortium- and exchange-owned dark pools, as well as independent destinations. The volumes traded in them are significant, and new venues are appearing regularly. How many will survive is unclear, but what is certain is that they are here to stay, and it is increasingly important for both buy- and sell-side traders to tap into this type of flow.

The initial aim of dark pools was to facilitate trading of large blocks of stock that have become difficult to trade on traditional, fully-displayed exchanges (light) without incurring significant market impact. However, it is now commonplace for any algorithmically traded order, large or small, to interact with these venues, in addition to traditional exchanges. The sheer number of dark pools – each with its own rules, liquidity characteristics and order flow composition – introduces confusion into traders’ decision-making process.

Which venues should be used? How should an order be split across competing dark pools? How should traders react dynamically to changing liquidity conditions across destinations? Is a given venue better suited for trading some names but not others? To navigate the modern equity marketplace in an intelligent fashion, one must know the answer to all of these questions.

At Banc of America Securities, we employ an integrated solution to this particular problem, which forms the basis of our dark pool trading algorithm. We first estimate the most likely distribution of available volume in each dark pool based on historical and real-time proprietary trading data. We then use such estimates to opti-

mally split an order across multiple dark pools in a way that maximizes the expected number of shares to be executed (or a variety of other user-selected optimization criteria). As time passes, we receive new information from our fills and lack of fills, and this process is repeated continuously throughout the life of an order. Our trading algorithm is a clear leap beyond ‘liquidity aggregators,’ as we are not simply offering connectivity to as many exchanges as possible but providing a complete front-to-back logic of optimal interaction with dark liquidity.

The Problem with Popular Methods

As dark pools have been around for several years, a number of ad-hoc methods of accessing dark liquidity have been developed. First, without having any insight into availability of liquidity in individual dark pools, it is not unreasonable to split an order evenly across venues.

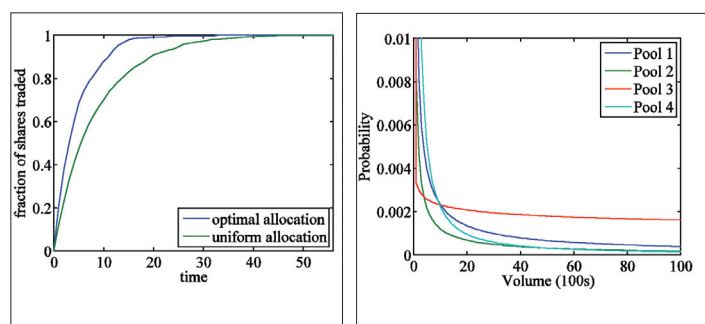



Figure 1 Splitting an order uniformly will result in slower execution (left panel) because volumes are distributed differently in different dark pools (right panel).

In Figure 1, the left panel shows how a uniform split can underperform dramatically compared to an allocation that uses estimates of the full volume distribution: having a reliable prediction of volume in each venue allows us to execute the entire order much faster. The reason for this is that different dark pools are not at all alike (Figure 1, right panel). Therefore, somebody who can estimate the full shape of the volume distribution will have a clear edge over somebody who assumes homogeneity across all dark pools.

Another popular allocation scheme is to send more shares to a venue with recent fills. This is an altogether more sensible approach, as our research verifies that liquidity is indeed persistent. However, it is not sufficient to just forecast ‘some volume’ or ‘no volume’ availability - optimal trading requires knowledge of the probability of executing at all possible volumes.

Figure 2 underscores this point. Our research and historical data show that the distribution of volume in dark pools is heavy-tailed: often, there is no volume available, but sometimes very large volume is present (10,000+ shares) – a feature not seen on light exchanges. More traditional distributional estimates cannot capture this reality, therefore they result in poor fits to the data, sub-optimal allocation of shares across venues and a drag on execution speed.

Figure 2 A heavy-tailed distribution is the best model for representing dark pool liquidity. For a particular stock and dark pool, we show the probability of at least that volume being available. A (parametric) heavy-tailed model clearly yields a far superior fit to the observed data than more traditional models (such as an exponentially decaying distribution).



Incorporating all the realities above into our algorithm is what we believe sets us apart from currently available solutions.

The BofA Way

Our approach to algorithmic interaction with dark pools consists of two main iterative processes: estimation and allocation. We have a robust way of predicting a full distribution of possible volumes available for trading in each venue. We treat each possible stock, venue and trade side as unique – our research shows that this is important (see Figure 1 again). Therefore, we maintain more than a dozen unique volume distributions for each stock – a buy and a sell distribution for each dark pool we access. When there is not enough data to make a reliable prediction, we assemble sensible ‘default distributions’ from similar stocks, which are specialized and refined over time. We continuously update our distributional estimates: both trades and lack of trade activity give us valuable information and make our predictions more accurate.

Given our best estimates of volume in each dark pool (most probable given all the data we have), we spread shares across venues according to their marginal liquidity. This entails computing the estimated probability that each successive share of an order will be filled in each venue and allocating to the venue that maximizes execution probability. The resulting total numbers of shares allocated to each venue are then submitted, and the resulting (possibly partial) executions in turn refine our distributional estimates. With minimal assumptions, our method provably maximizes the expected number of shares executed each time we make an allocation decision. Just like the estimation, this allocation process is dynamic; as we receive fills and our estimates change, we rebalance child orders across the appropriate venues.

At the top level, this iterative predict-allocate-predict process effectively explores all dark pools, searching for liquidity that best ‘fits’ a given order. This highly adaptive system is made possible by combining our optimized real-time data and trading infrastructure with an abundance of our own internal and proprietary historical data, which we have collected through trading and experience data.

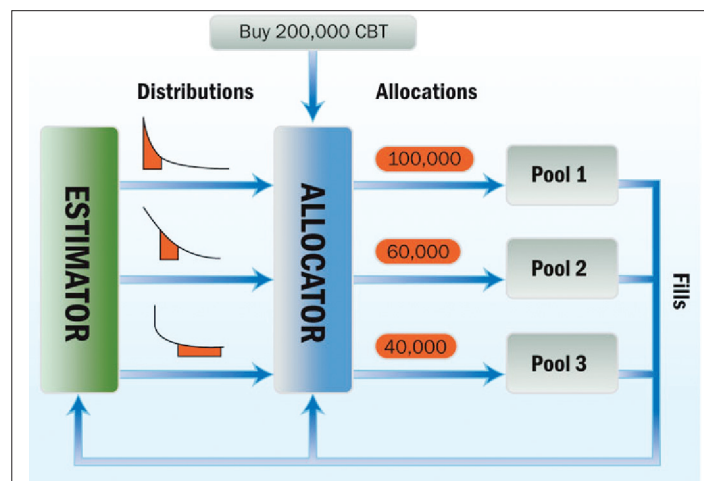


Figure 3 Dark Pool Trading System Workflow

Our framework is designed to handle interactions with all types of dark liquidity: trading in dark pools with continuous and periodic crossing; taking and posting liquidity in a pool; taking advantage of hidden liquidity on light exchanges; and both internal and external crossing. We believe that by reacting dynamically to every liquidity event and by performing sophisticated real-time forecasts of available volume, we give our clients a unique cutting-edge trading tool that cannot be replicated through human intervention. We also offer our clients economies of scale, as serving as an aggregator of flow for so many buy- and sell-side customers gives us the most accurate real-time estimate of liquidity across exchanges.

This algorithm will be first released on a stand-alone basis, but it will soon become a gateway to dark pools for all other algorithmic strategies. The entire suite of Banc of America electronic trading products has been created around the same principles: sound fundamental theory, rigorous backtesting on historical data, implementation in a reliable framework and continuous improvement based on experience.

About the Company:

Banc of America Securities' innovative Electronic Trading Services (ETS) platform provides institutional investors with automated access to diverse, leading-edge liquidity systems and venues. Whatever your trading objective, our electronic trading products can meet virtually any need, including single stock, block options, direct market access and program trading. ETS turns smart trading ideas into simple trading solutions that are easy to access and use.