CITADEL | Securities

OCTOBER 2021

Market Lens

Why Restricting Cancel Rates Can Increase Bid-Ask Spreads

Last year, Citadel Securities published "Why Do Electronic Traders Cancel Orders? What ever-increasing speeds for issuing and cancelling orders tell us about today's market structure." This paper is a follow up to that article with new analysis and data.

While increasing electronification of markets has brought many benefits to market participants, some have expressed concern about the speed and automation of electronic trading, including rapidly changing price quotations and order cancellations. Upon deeper examination, however, these features have become an integral part of our market structure that reduce volatility and tighten bid-ask spreads.

The decrease in bid-ask spreads has resulted in a boon for investors in the form of lower trading costs. The lower bid-ask spreads also result in lower issuing costs for issuers, which benefits the broader economy in addition to the direct impact on savers.

This paper examines how order cancellations fit into modern markets and how various factors have contributed to the rise of – and benefits from – this activity. Readers familiar with last year's article and the market structure discussed may wish to skip the first three sections of this paper and proceed directly to new data presented in the fourth section of this paper.

THE ROLE OF AUTOMATED TRADERS

The basic automated trading model has been consistent throughout the evolution of modern markets. Automated traders, including market makers, set the prices at which they are willing to buy and sell any given security based on available pricing information and predictive analysis. The more confidence that automated traders have in the accuracy of this information and analysis, the tighter the bid-ask spread at which they can quote prices in competition with others. To ensure that securities trade at fair and competitive prices, automated traders continually update their prices in response to market movements and changes in information.

While this process has remained the same, information now changes faster than ever before and competition among automated traders to offer better pricing has become more vigorous. In the face of these developments, order cancellations allow automated traders to dynamically adjust their prices to reflect rapid changes in supply and demand, which results in tighter bid-ask spreads and better execution for all market participants.

THE ROLE OF TECHNOLOGY

Several features of modern trading and markets cause high levels of order cancellations as a normal and beneficial course of business. A common characteristic across today's markets is the use of computers to calculate desired prices — as well as to route, execute and communicate the status of orders — with far greater speed, scale, transparency, and efficiency than was possible in manual markets.

Investors and traders now employ sophisticated quantitative tools that allow them to consider a variety of inputs simultaneously and in real time when determining the price at which to buy or sell a security. Once they determine the right prices, they can route orders using computer algorithms and fast communications technology to exchange matching engines that receive, fill, and confirm execution of orders in fractions of a second.

This evolution of the market has resulted in far more accurate price discovery, lower bid-ask spreads, and lower transaction costs, which have greatly benefited all market participants. Investors are therefore able to retain more of the profits from their investments.

In 2010, for example, Japan Exchange Group (JPX) introduced a new trade matching platform, called Arrowhead, which dramatically reduced messaging and execution latency. According to a 2014 JPX working paper[1] on the introduction of Arrowhead, "the launch of Arrowhead boosted liquidity provision in volatile stocks, contributing to reduced transaction costs."

THE ROLE OF RISK MANAGEMENT

To operate successfully in a modern system that delivers these substantial benefits, market participants have to fundamentally alter their approach to risk management. To manage risk in real-time and by extension keep their displayed quotes as tight and competitive as possible, market participants must frequently cancel quotes and update their prices. In other words, a faster system requires greater cancellations.

Consider a market maker who decides to submit a displayed limit order to an exchange. These orders commit the user to buy or sell a set amount of a particular stock at a specified price. As such, limit orders and quotes are like automatically executable options posted for all other market participants to exercise through the exchange's order book. The "premium" a market maker receives for providing this option to the market primarily comes from the bid-ask spread. This "compensation" is required since market makers who place limit orders — the foundation of public price discovery — are exposed to the risk that their quotations will be executed at an inopportune time, leading to potential losses. The greater the risk of an inopportune execution, the more compensation is required, which leads to wider bid-ask spreads. Conversely, anything the market maker can do to lower the risk of an inopportune execution will lower the compensation required, which leads to narrower bid-ask spreads. That is true not only for official market makers, but also for proprietary traders providing two-sided liquidity without an official mandate from an exchange and agency brokers representing customer interest.

Consequently, a wide array of market participants seek to lower their risk of inopportune executions by constantly updating their orders to reflect changing market conditions. In automated markets, this means frequently cancelling and replacing firm orders resting on electronic order books. A firm posting two-sided liquidity in an individual security, for instance, might need to raise both its bid and offer after an outstanding sell order is executed in order to adjust for supply and demand changes in that security. In such a case, it would have two firm orders on the book – one of which it would execute and another that it would cancel, resulting in a 50% cancellation rate.[2] With modern-day exchange order books able to process messages and execute transactions in thousandths of a second, and execution priority for resting orders determined according to price-time priority,[3] quote updates need to keep pace, leading to high quote cancellation rates.

As an example, consider 2 simplified, hypothetical markets: a market where order cancels are restricted and a market that does not restrict healthy levels of order cancellations. A market maker in the first market who places a tight bid-ask spread but cannot cancel orders will be continuously selling as markets rise and continuously buying as they fall. This market maker would face significant losses as it will frequently find itself on the wrong side of a moving market. This is shown in *Panel 1* of the below exhibit.[4] In this simplified example, the market maker has a 4 penny wide spread. As soon as one side of its quotes is executed, it reprices its quotes so that the midpoint in its quote is the executed price.[5] In this brief example, the market marker buys 10 times and sells 11 times. While the market maker was nearly flat at the end of the period, its buys were consistently in declining markets and sells in rising ones because it could not cancel orders and reprice. The result is that the market maker would have lost substantial money by buying as prices drop and selling as they increase. Clearly this is an unsustainable strategy.

1. Hosaka, Go: "Analysis of High-Frequency Trading at Tokyo Stock Exchange." JPX Working Paper, March, 2014
2. Levine, Matt: "Why Do High-Frequency Traders Cancel So Many Orders?" Bloomberg View, October 8, 2015
3. Price-time priority means that limit orders at the best prices — the highest bids and the lowest offers — are filled first; In the event of a tie on price, order-entry time determines priority — the first order to arrive gets executed first, and so on for later-arriving orders.
4. Market maker actions are entirely illustrative and simplified
5. Assuming the market maker is obligated to provide quotes and does not withdraw from the market which would decrease liquidity

Given the above fact pattern, and in the absence of the ability to cancel orders, the market maker will widen its bid-ask spreads so that it typically remains on the right side of moving markets. This is shown in *Panel 2* of the below exhibit. In this illustrative example, the market maker from *Panel 1* changes its spread from approximately 4 pennies to 10. While it will still sell in a rising market and buy in a declining one, with a sufficiently large spread, these variances are minimized and the market maker is compensated for the risk by the larger spread. In this example, the market maker ends the period with 2 buys and 2 sells and a quote consistent with where it started the day. However, the impact on the market is quite clear: the spread that the market maker requires is much larger, transferring profits from investors to the market maker to make up for its increased risk because it cannot cancel orders. At the same time, the number of trades completed declines substantially (from 21 to 4 in this simplified example). Thus, we witness an overall impairment of both price discovery and liquidity.

Now consider a market which does not restrict healthy levels of order cancellations. A market maker in this market will be able to offer tight bid-ask spreads, but it will cancel quotes as prices change to avoid the fate of the market maker in *Panel 1*. The market maker in this market will offer the same bid-ask spread as in the first panel but will not face the same adverse selection of constantly buying in declining markets and selling in rising markets because it will cancel orders as required in light of changing market conditions. This market maker is shown in *Panel 3* of the below exhibit. As can be seen, this hypothetical market maker increases the total number of quotes made relative to the first example and significantly increases trading relative to the second example (8 trades relative to 4).

As is shown in these examples, a market that artificially limits cancellations also artificially widens tick sizes in the process. These wider tick sizes increase transaction costs for investors and impair liquidity.

CASE EXAMPLES: THE TICK SIZE PILOT AND THE SUB-PENNY RULE

As discussed above, a higher cancellation rate is part and parcel of increased market efficiency. All else equal, when bid-ask spreads widen or narrow, cancellation and message rates move in the opposite direction. This can be seen time and again when there is an exogenous change in a security which causes spreads to either narrow or widen. Two examples are seen in

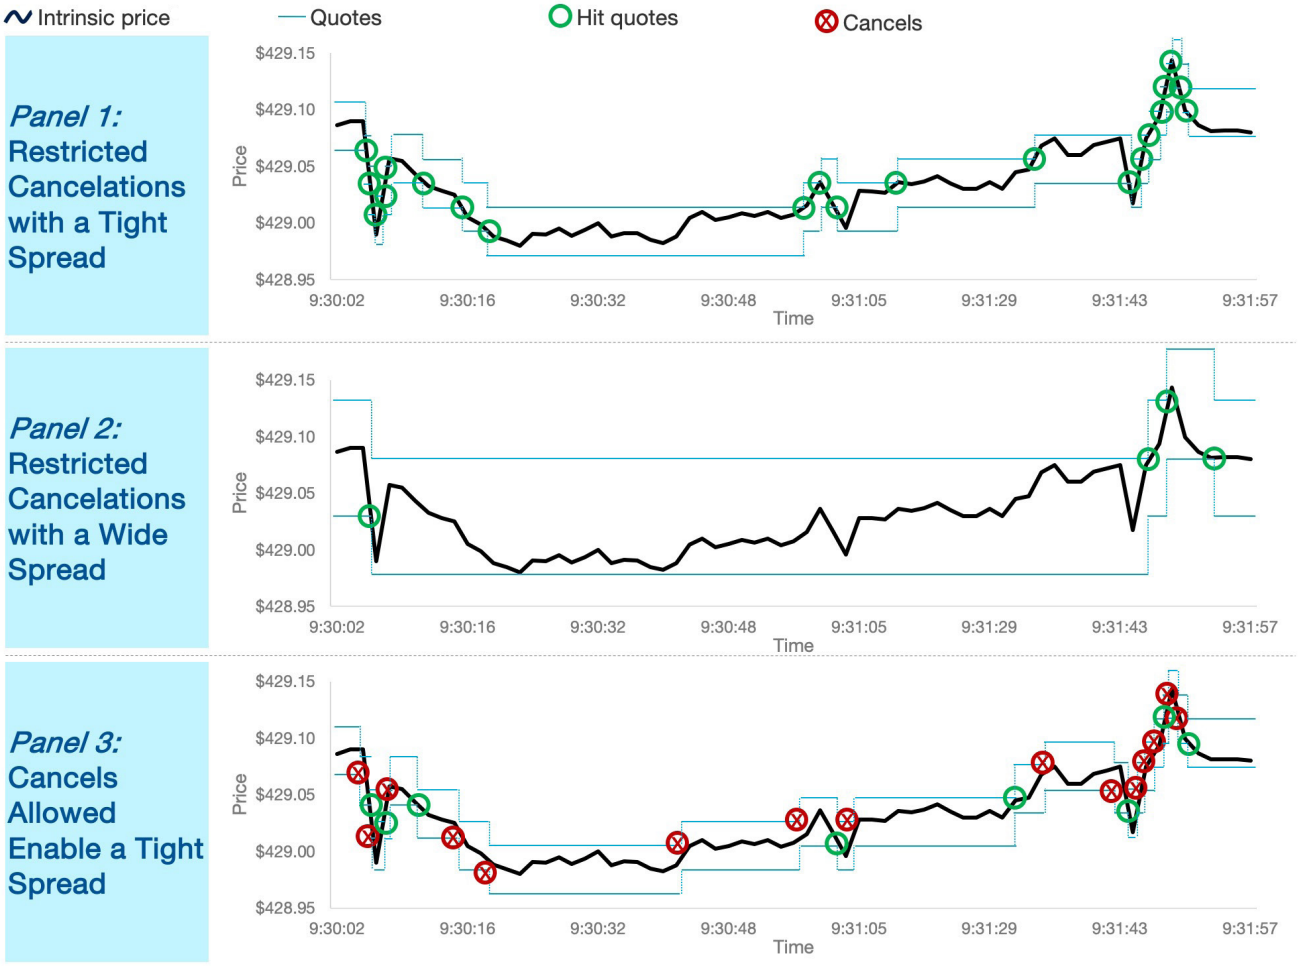

the United States' "Tick Size Pilot" and stocks that change between below and above $1. These examples are particularly illustrative because, in the first, the pilot provided an ideal test and control setup to see the effects of a change in spreads where the comparison between the test and control groups should eliminate other potential factors. In the second, the "Sub-Penny Rule" in the United States constrains any stocks priced above $1 to a minimum spread of $0.01 while those priced below $1 can have a spread as small as one one-hundredth of that size. This means that the spreads for any stocks that are tick constrained when barely above $1 and which move below $1 can collapse only due to that price change.

The U.S. Securities and Exchange Commission (SEC), created the Tick Size Pilot (TSP) "to allow the Commission, market participants, and the public to study and assess the impact of wider minimum quoting and trading increments – or tick sizes – on the liquidity and trading of the common stocks of certain small-capitalization companies."[6] The SEC created pilot groups of stocks where the minimum tick size was moved from 1 cent to 5 cents (along with certain other specifics which are beyond the scope of this paper) as well as a control group which was left unchanged. While the pilot itself was not successful in creating more liquidity and trading in pilot groups,[7] it did create an environment with pilot and control groups to show the effects of wider spreads on messaging and cancellation rates.

As expected, when the TSP went into effect, the bid-ask spreads of the pilot groups relative to the control group increased substantially (by definition).[8] At the same time, cancel messages decreased substantially for the pilot group relative to the control (by ~20%). This is shown in the below exhibit:

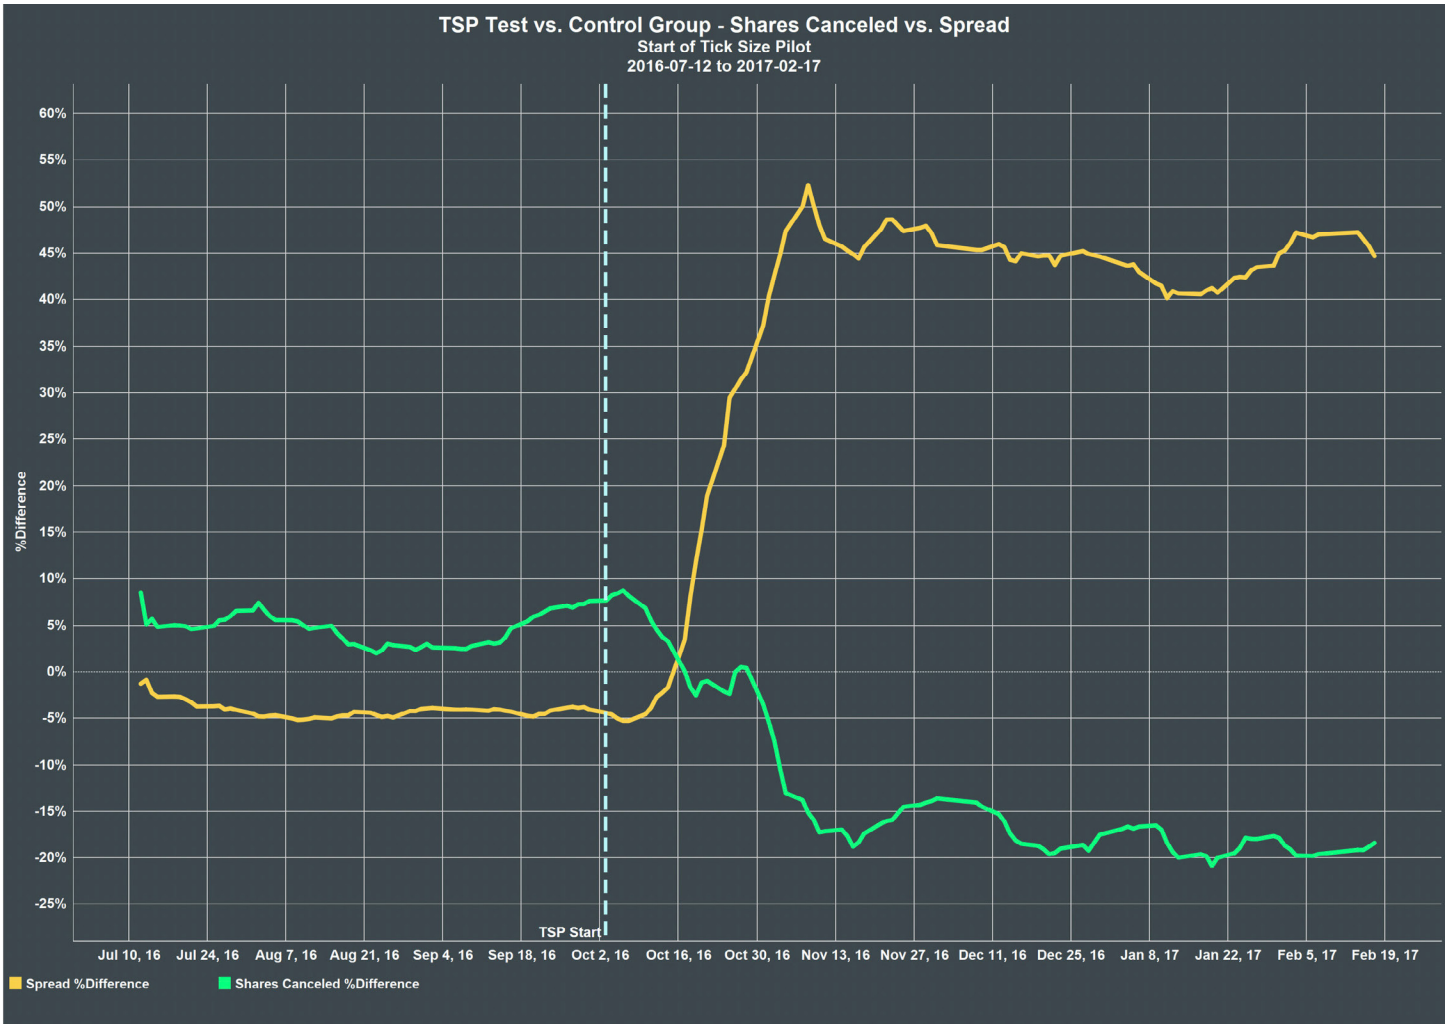

6. https://www.sec.gov/ticksizepilot

7. https://www.sec.gov/files/TICK%20PILOT%20ASSESSMENT%20FINAL%20Aug%202.pdf

8. Note: all spreads shown in this paper are relative spreads (i.e., spreads as a percent of stock price)

Because the pilot was implemented for different stocks over the course of several weeks, there is a slight lag between the start of the pilot and when spreads increased and cancels decreased. Nonetheless, the trend is clear: the increasing spreads resulted in fewer cancels as market makers no longer had to adjust their prices as frequently around a changing intrinsic price.

When the TSP was ended 2 years later, and the spreads of the pilot group returned to a minimum of 1 penny, the opposite effect occurred: the spread difference between the pilot and control groups collapsed back towards 0 and cancels of the pilot group increased to around parity with the control group:

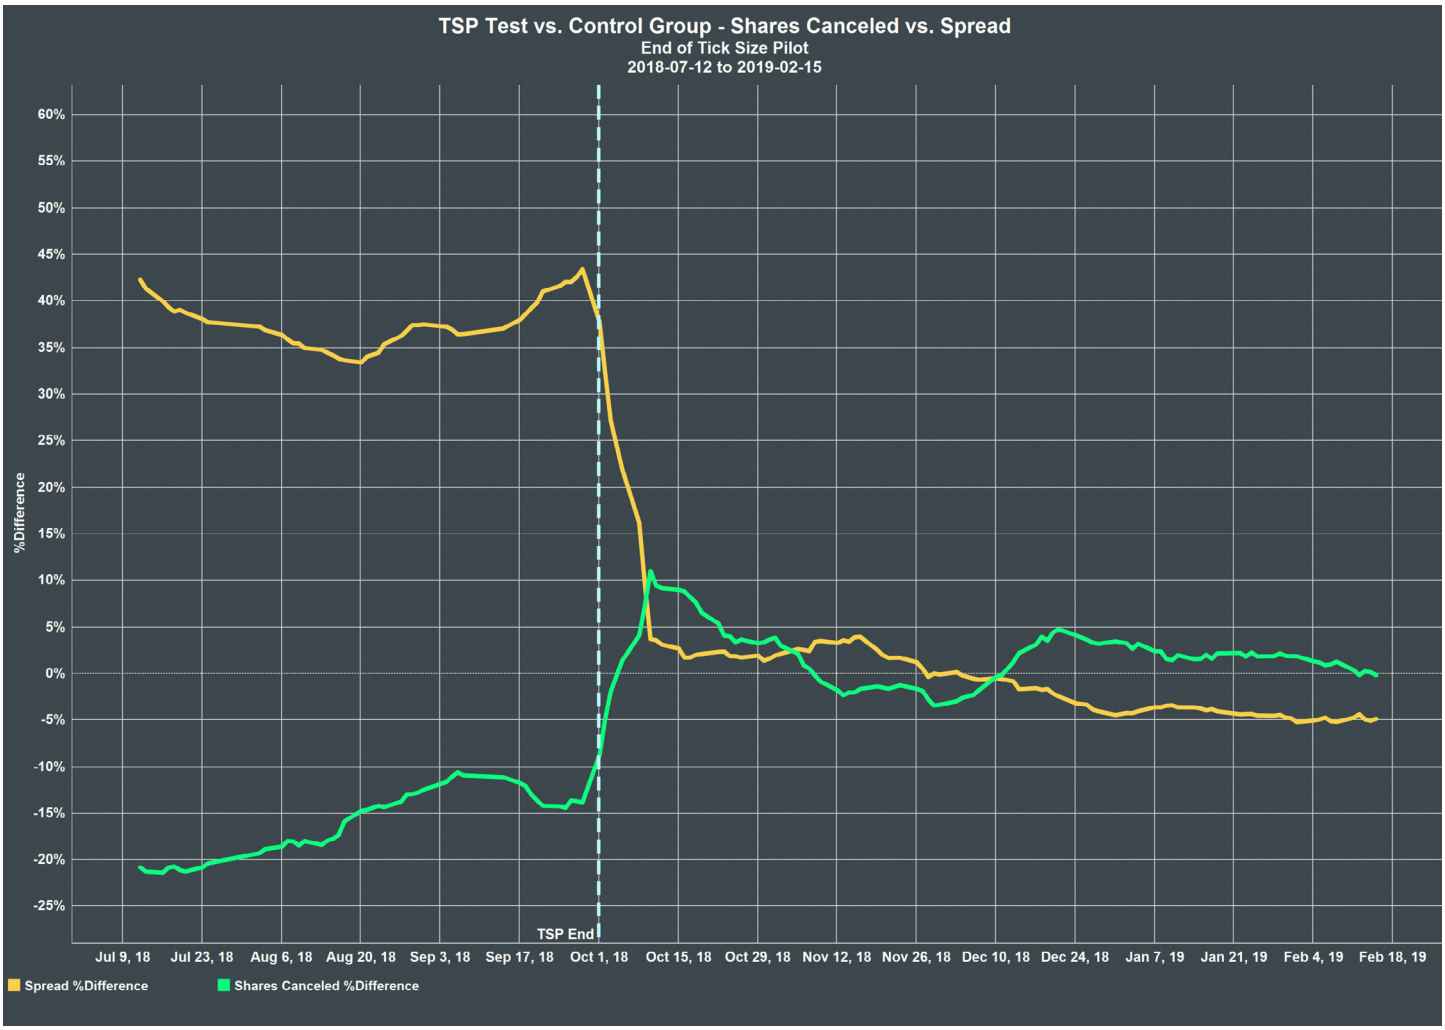

Because the pilot ended on a set date for all stocks (unlike the experience in the launch of the pilot), the spreads declined rapidly and the cancel messages increased back to the level of the control group.

A second example is driven by US equity market microstructure. When the SEC implemented Regulation NMS in the United States in 2005, it included a rule that the minimum price increment for stocks trading above $1 on an exchange was a single penny and that stocks trading below $1 could trade at one one-hundredth of a penny (Rule 612 or the “Sub-Penny Rule”). For stocks trading at exactly $1, this would mean their spread is at least 1% (100bps), which is significant. As a result of this rule, the spread for liquid stocks trading just above $1 which move below that threshold can decrease substantially (to as low as $0.0001 from $0.01) only due to the price movement and vice versa.

A good illustration of this phenomenon occurs in Evofem Biosciences (EVFM) in mid-2021. In April and early May 2021, EVFM was trading above $1 with a spread of ~80bps in early May. On May 18, EVFM lost approximately 1/3 of its value and moved from a closing price of $1.28 on May 17 to $0.8426 on May 18. Because the stock moved below $1, the spread could decline and did indeed collapse (to 5-10bps from 80bps). As expected, the cancellation rates increased dramatically: from ~30 million a day to 40-60 million, an increase of 33-100%. Then, less

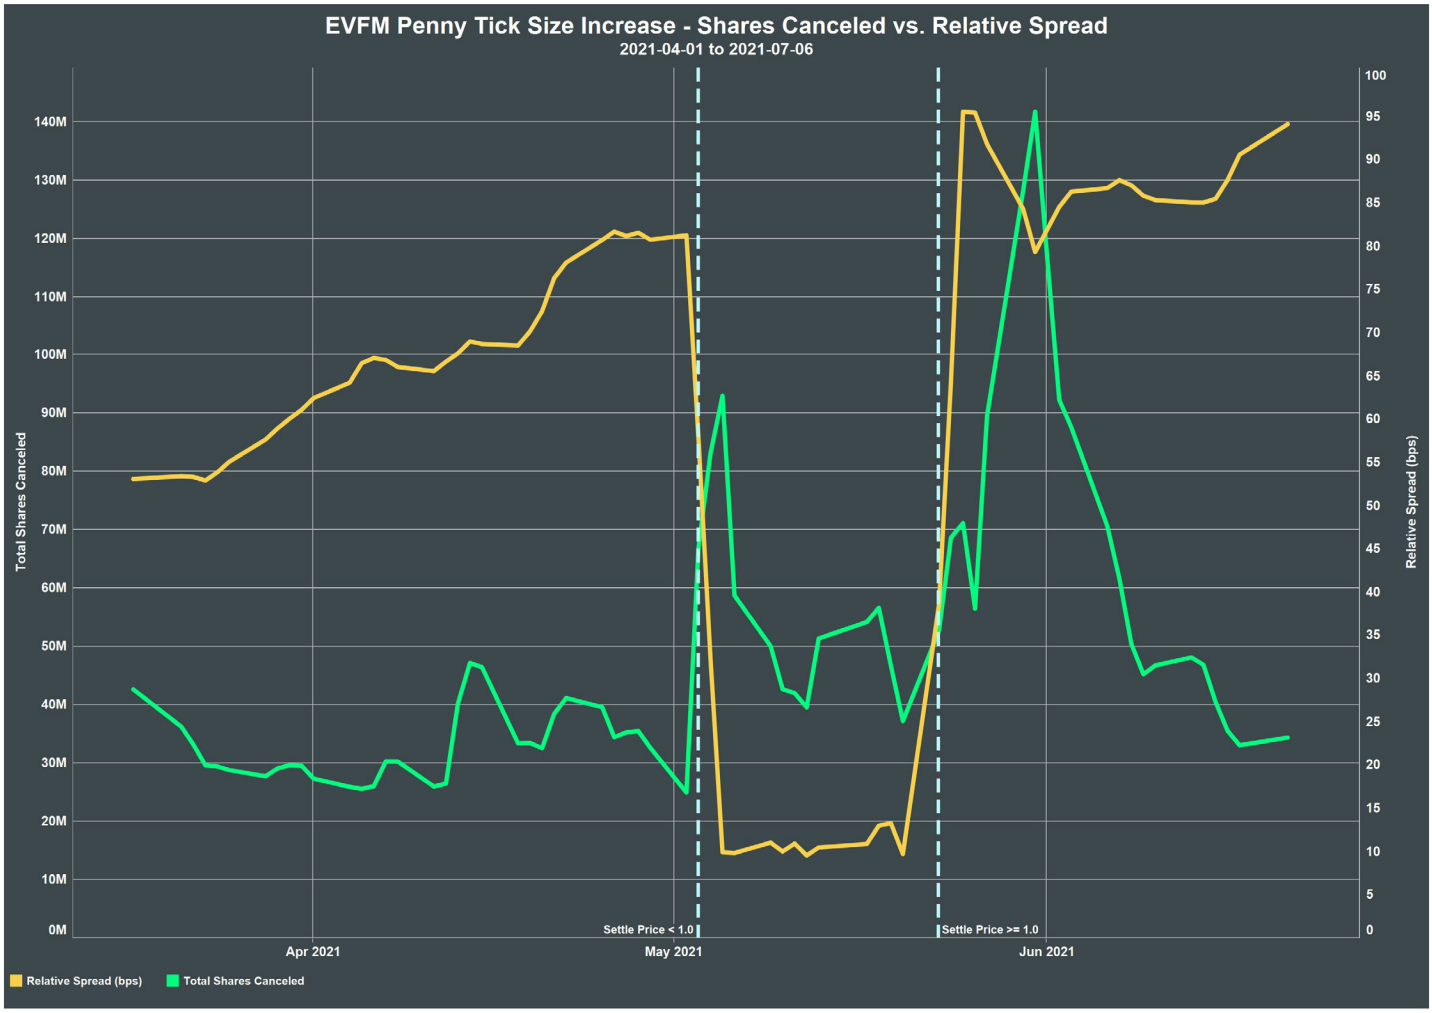

than one month later (on June 7), EVFM moved back above $1 and the spread moved to nearly 100bps. After trading volumes subsided a few days later (trading volume was ~10x of average volume after it moved above $1), the cancellation rates collapsed back down towards its starting point.

These three examples illustrate the point that cancel rates and spreads are two sides to the same coin. When market makers have the opportunity to compete more aggressively they do so. Tighter spreads inherently drive higher cancellation rates, but the value created in the form of lower bid-ask spreads, better price discovery, and ultimately a lower cost of capital for issuers are a huge benefit to both investors and the broader economy.

CONCLUSION

As we outline in this paper, rather than being indicative of any problems, healthy quote cancellation rates have become not only normal, but also integral to the proper functioning of modern markets, resulting in greater efficiency, narrower bid-ask spreads, and more robust price discovery. The narrower spreads benefit investors and the market as a whole in the form of better pricing and therefore more profitable investing. These benefits result in lower capital costs for issuers, creating value for the whole economy.

This paper provides additional evidence on the inextricable relationship between bid-ask spreads and cancel rates, why high levels of order cancellations fit into a healthy market structure, and how this ultimately benefits investors.