An application of agent-based modeling to market structure policy: the case of the U.S. Tick Size Pilot Program and market maker profitability

Charles Collver[*]
U.S. Securities and Exchange Commission

Original version: July1, 2017
This version: December 26, 2017

Abstract

I demonstrate the feasibility of applying a heterogeneous agent-based model of a modern limit order market to inform U.S. equity market structure policy. The model specifically addresses matters related to the U.S. Tick Size Pilot Program as an example to demonstrate the utility of agent-based models in general. In this case, the model provides quantitative results on changes in market maker participation and profitability as well as quoted spreads when the minimum pricing increment is raised from one tick to five ticks. Perhaps more importantly, the very act of applying the agent-based model and analyzing the results generates additional insights into the nuances of the impact of tick size on small cap equity market quality. Agent-based modeling should be included in the regulatory toolbox because it can not only provide answers to specific policy-related questions but can also help policy-makers ask better questions before implementing the policy.

[*] collverc@sec.gov. This White Paper represents work in progress and is intended to elicit comments and further discussion. Thanks to Frank Hatheway, Laura Tuttle, my colleagues in the Office of Analytics and Research and an anonymous reviewer for helpful comments and thoughtful suggestions.

In a comment letter addressing the U.S. Securities and Exchange Commission's "Concept Release on Equity Market Structure," a group of scientists and business people recommended that the SEC develop and deploy an agent-based model test bed to investigate the policy initiatives set forth in the concept release.[2] The authors argued that the U.S. equity market could be especially susceptible to agent-based models because it is essentially a closed system and because real equity market behavior is driven primarily by algorithmic agents with humans acting as "meta-agents." They provide examples of several areas of potential interest including the "analysis of the impact of proposed regulations on market quality metrics" as well as the effects of proposed regulation on institutional transaction costs.

For a variety of reasons, the agent-based model test bed recommendation has gained little traction in the past seven years later. In 2010, there was some doubt that agent-based models could be reasonably executed in a regulatory setting. Agent-based modeling required a skill set that was not necessarily covered in standard economics and finance curricula. Modern agent-based models require a combination of advanced coding, software development and large scale computer technology skills along with a deep understanding of the actual workings of the markets to be modeled. Moreover, agent-based models are compute and storage intensive. Once the initial development stage is completed, they require a lot of computer firepower to run in production. In 2017, the question of whether agent-based models can be developed and deployed in a regulatory setting has largely been answered in the affirmative. Modern programming languages are more versatile and easier to use, software development is now taught in masters-level quantitative finance and economics courses, and, current cloud technologies have made dramatic improvements in compute power and storage capacity available at very low cost.

But should agent-based models be used in a regulatory setting? Currently accepted models and methodologies, such as theoretical equation-based and empirical statistics-based models of standard economics, are well understood while their shortcomings are often ignored or elided by practitioners and academics alike. Agent-based models are relatively new and less well understood. Their shortcomings are easy to spot while their potentially applicability remains to be seen. In Buchanan (2009), Chester Spatt, a former chief economist at the SEC, commented on the use of agent-based models in a regulatory setting: "It would be problematic for the rule-making process to use methods whose foundation or applicability were not established."

This paper provides a step in the direction of establishing a foundation for the use of agent-based models in a regulatory setting by applying an agent-based model to the U.S.

[2] See the Concept Release on Equity Market Structure and the comment letter of April 16, 2010 for more details.

Tick Size Pilot Program. The model provides answers to several questions posed by advocates for and opponents of the proposed pilot:

- Does increasing the minimum pricing increment (MPI) necessarily imply that spreads will widen, thereby making market makers more profitable?
- What impact do penny jumpers have on spreads, and by extension, market maker participation and profitability?
- Does increasing the MPI mitigate the impact of penny jumpers?

Section 106(b) of the JOBS Act required the Securities and Exchange Commission to conduct a study examining the effects of decimalization on initial public offerings and small and middle capitalization companies, specifically stating[3]

> "The study shall also examine the impact that this change has had on liquidity for small and middle capitalization company securities and whether there is sufficient economic incentive to support trading operations in these securities in penny increments."

The resulting study, Report to Congress on Decimalization, reviewed academic literature and suggested in one of their caveats to relying on academic literature that "… the effect of decimalization on capital formation has not been explored in the literature."[4] However, the study does provide a narrative describing the link between "sufficient economic incentive to support trading operations" and "the effect of decimalization on capital formation." In short, the lack of revenue from trading small and mid capitalization companies discourages intermediaries from underwriting and making markets for newly public companies. Will increased revenue from trading small and mid capitalization companies facilitate capital formation? Who knows? Time will tell.[5]

Whether a wider pricing increment will increase revenue for market making intermediaries depends on, in part, whether market makers earn the spread. If market makers were able to simultaneously buy and sell matched quantities, then they would always earn the spread and never carry inventory. But equal quantities of buy and sell orders rarely arrive at exactly the same time. These days, market makers act as dealers - buying when others want to sell, selling when others want to buy and carrying inventory between the mismatched arrival times of buyers and sellers. Consequently, market makers face the risk of unfavorable inventory revaluation when prices change. In the absence of

---

[3] JOBS Act § 106(b) (2012).
[4] See https://www.sec.gov/news/studies/2012/decimalization-072012.pdf
[5] Despite SEC Staff recommending against a tick size rulemaking (Report to Congress on Decimalization, page 22), the SEC approved a pilot "to study and assess the impact of increment conventions on the liquidity and trading of the common stocks of small capitalization companies." (SEC Order Approving a "Plan To Implement A Tick Size Pilot Program") There is no mention of capital formation. However, SEC Press Release 2014-126 states "The Commission plans to use the program to assess whether these changes would enhance market quality to the benefit of U.S. investors, issuers, and other market participants." Presumably market makers are considered other market participants.

information-based trading, these inventory revaluations would tend to cancel each other out on average and over long periods of time. But, in the presence of informed traders, market makers face the risk of building inventory precisely when revaluation is likely to be unfavorable. In modern anonymous markets, market makers cannot know for sure who is informed and who is not. To compensate for the additional risk of trading with better informed traders, market makers quote wider spreads than they would in the absence of informed counterparties.

But market makers do not necessarily capture their quoted spread because wider quoted spreads provide incentives for other sophisticated traders to employ a variety of quote matching strategies to step in front of the market maker's quotes.[6] Penny jumping is a form of quote matching where traders improve upon a standing displayed limit order by the permissible minimum pricing increment.[7] In the presence of penny jumpers, market makers provide a free option: if prices move favorably for the penny jumper, then they earn the (smaller) spread, if not they can minimize their losses to one price increment by leaning on the market maker's quote (i.e., reversing their position by selling into standing buy orders or buying against standing sell orders). To be successful penny jumpers must be faster than other liquidity traders, especially when prices are moving. Harris (2015) deems penny jumping a parasitic trading strategy because it harms market makers and other limit order traders. Penny jumping purportedly provides little benefit to liquidity takers while discouraging liquidity providers from quoting their best prices and/or depth, thereby harming the host exchange.

If parasitic penny jumping strategies are harming the liquidity of small and mid capitalization companies, then widening the MPI could be the chloroquine. A larger price increment potentially limits the number of inter-quote prices and/or limits the amount of time wider spreads are in force while simultaneously increasing the cost to penny jumpers of leaning on market maker quotes. Whether a larger MPI improves market maker profitability is an open question. Agent-based modeling provides an ideal way to test this question because it facilitates emphasizing particular aspects of the participants in a limit order market and the order-matching mechanism by which trades are facilitated. Like theoretical equation-based models and empirical statistics-based models, certain aspects of agent-based models of reality can be de-emphasized while particular areas of interest can be accentuated. Also like theoretical and empirical models, the choice of emphasis can impact any conclusions. Unlike theoretical and empirical models, agent-based models can

---

[6] See Trading and Electronic Markets: What Investment Professionals Need to Know (Harris, 2015) for a general definition of quote matching strategies and some examples.

[7] It is called penny jumping because the minimum price increment for displayed quotes in U.S. equity markets was one cent for most stocks in the 15 years prior to the Tick Size Pilot.

easily handle heterogeneity in types of agents and their rule sets, and, relatively simple rule sets can generate emergent behavior consistent with observed complex interactions.

In this study, I extend a zero intelligence agent-based model of a price-time priority limit order book proposed in Pries et al. (2006, 2007) and later extended by Bookstaber et al. (2015) by including two specialized liquidity providers: a penny jumper and a market maker. These agents have liquidity provision rules distinct from ordinary zero intelligence liquidity providers. The model also tests the impact of imposing a five tick MPI on stocks that typically trade at a spreads ranging from less than five ticks to greater than fifteen ticks. I make several contributions. First, when penny jumpers are successful, market maker participation and profitability decrease with the resultant decreasing quoted spreads. Second, imposition of a larger MPI reduces these symptoms of penny jumper success: market maker participation is partially restored and profitability improves along with widening spreads. This might be expected for stocks that typically trade at spreads less than the new wider MPI. However, these results also hold for stocks that trade at spreads much wider than five ticks in the one tick MPI regime. This is the third contribution. When inside prices are considered reservation prices, spreads widen after imposition of a larger MPI, even for stocks that previously exhibited quoted spreads larger than the new MPI and in regimes with no penny jumpers. A larger price increment improves market maker profitability in two ways. First, a larger MPI generates wider spreads, thereby improving profitability for all liquidity providers. And second, a larger MPI discourages penny jumpers, thereby improving participation and profitability for other liquidity providers, including market makers.

Empirical strategy

The strategy is to first calibrate the agent-based model simulations to a set of market structure metrics observed for the actual stocks in the U.S. Tick Size Pilot Program during the pre-pilot phase, then increase the MPI in a set of simulations and check for consistency in the same set of market structure metrics with the actual stocks in the U.S. Tick Size Pilot Program during the post-implementation phase and measure the impact of increasing the MPI on simulated penny jumper behavior and market maker participation and profitability. Exhibit A of the NMS Plan to Implement a Tick Size Pilot Program describes the construction of the test and control groups of U.S. stocks.[8] To be considered, NMS common stocks must have: a market capitalization of $3 billion or less on the last day of the measurement period; a closing price of at least $2.00 less on the last day of the measurement period and at least $1.50 on every day of the measurement period; a measurement period volume weighted average price of at least $2.00; and an average daily

[8] See Exhibit A of the Plan for details on the three test groups, their construction and the stratified random sampling procedure.

volume of 1 million shares or less for the measurement period. The measurement period – June, July and August, 2016 – also serves as the pre-pilot period for the calibrations. Stocks were phased in to the test regime during September and October, 2016. For consistency checking, the post-implementation period is December, 2016 – February, 2017.[9]

The model in a nutshell

The model is composed of one security, one limit order book and four types of traders: liquidity takers, liquidity providers, market makers and penny jumpers. The limit order book follows a standard price-time priority. Liquidity takers and liquidity providers arrive according to independent draws from agent-specific Poisson arrival time distributions (i.e., arrival times for liquidity takers are heterogeneous but drawn from the same distribution and similarly for liquidity providers). The single market maker arrives once every time step. Upon arrival, liquidity providers submit one limit order priced from an exponential distribution capped by (and excluding) the prevailing price on the other side of the market. This distribution can vary over time as a function of the information environment. Liquidity takers submit one market order. Market makers submit 12 orders from a choice of 60 prices on a discrete uniform interval determined by the prevailing best price on the same side of the market with the following rule: if there is only one share at the best price, shift the interval away by one price increment.[10] In other words, market makers never join a best price if there is only one share displayed at that price. Liquidity providers and market makers have an opportunity to randomly cancel orders at each time step. Penny jumpers have a probabilistic opportunity to step in front of posted prices after the arrival of each of the other agents during a time step. Penny jumpers formulate buy and sell orders such that they are the only liquidity provider at the inside price(s) upon arrival. If there is no opportunity to be alone at the inside, penny jumpers withdraw from the market.

There are no privately informed traders. However the information environment is modeled via the probability a liquidity taker submits a buy order. This probability follows a mean-reverting random walk over the 100,000 time steps in the simulation. At each time step the liquidity providers, market maker and a subset of the liquidity takers are randomized in a queue and selected in order. Each selected agent behaves according to their rules: liquidity providers and the market maker add orders according to their arrival times (i.e., if their arrival time interval matches the current time step) and potentially cancel a fraction of their outstanding orders regardless of their arrival time. Liquidity

---

[9] I exclude November, 2016, the month of the tumultuous U.S. elections.

[10] This strategy allows the penny jumper to be the price setter when the penny jumper is active. A slightly more complex decision rule could be to allow the market maker to join other liquidity providers but not the penny jumper. The results show that the market maker and the penny jumper (when present) participate in the vast majority of trades, suggesting that after a few time steps the standard liquidity provider is rarely the price setter.

takers submit market orders if they are selected (i.e., if their arrival time interval matches the time step). After any of these three agents arrives, the penny jumper has an opportunity, based on a parameterized probability function, to step in front of displayed quotes if possible. The penny jumper could participate zero, one or many times during a time step.

Eight distinct simulation groups vary along two dimensions: choice of MPI and a parameter (C_λ) to control the liquidity provider price distribution choice set. Four simulation groups have an MPI of one and four have an MPI of five. For each MPI choice, C_λ can be 1, 5, 10 or 50, with larger values of C_λ generating wider quoted spreads, on average. Each simulation group contains six sets of 100 simulations, with each set distinguished by the probability a penny jumper will participate (α_{PJ}), which ranges from 0% to 10%. Figure 1 provides a schematic of the simulation strategy and Table 1 summarizes the model parameters.

In the first four groups of simulations the MPI is one tick. The second four groups of simulations impose a five tick MPI. Liquidity providers choose from the same price sets as in the one-tick regime, but round away from the quote midpoint: ask prices are rounded up to the nearest five tick price and bid prices are rounded down. In this sense, the prices before rounding are reservation prices. The market maker does not observe the reservation price but chooses prices from a five-tick grid with probabilities based upon the observed posted prices.[11] To facilitate comparisons between simulation groups, I use the same set of 100 fixed random seeds for each of the 100 simulations in a simulation set. This means that for each choice of α_{PJ}, the differences between the eight simulation groups are determined solely by the choice of MPI and C_λ.

Model calibration

The agent-based model simulations are calibrated to quoted spreads and cancel-to-trade ratios observed in the Tick Size Pilot stocks during the pre-pilot period. Figure 2 shows average quoted spreads from the four simulation sets with the MPI set to one and α_{PJ} set to zero match the distribution of daily duration-weighted average quoted spread in cents for 2,397 symbols included in the Tick Size Pilot from June, 2016 to August, 2016.[12] The mean (median) average daily quoted spread for the Tick Size Pilot stocks is 15.7 (6.10) cents and the range is from 1.01 cents to $12.17. The mean (median) quoted spread for the simulation sets ranges from 4.54 (3.77) ticks with C_λ=1 to 15.2 (12.4) ticks with C_λ=50. Taken together, the quoted spread distributions from the simulations match the overall

---

[11] See the Methods section of the Appendix for more details.

[12] By June, 2016, the Limit Up – Limit Down regime was in place. The SIP data contains a message and field depicting the bid and ask quote limit bands. Quotes with the best bid or best ask (or both) outside of these bands are excluded from the average spread computations.

shape and range of the average quoted spreads observed in the Tick Size Pilot stocks prior to implementation of the pilot.

Using the same 100 random seeds ensures each simulation set has the same number of trades. The mean number of trades for the 100 simulations is 34,020. The daily average number of trades for the Tick Size Pilot stocks is 1,374. Based on these averages, a simulation with 100,000 steps represents 24.8 trading days, or, each simulation step represents 5.79 seconds of real trading time for Tick Size Pilot stocks, on average. The median number of trades in the simulation set is 25,773 and in the Tick Size Pilot stocks the median number of trades is 742. Based upon medians, each simulation step represents about 8 seconds of trading time.

Figure 3 shows the average cancel-to-trade ratio from a simulation set with the MPI set to one and α_{PJ} set to zero coincides with the distribution of average cancel-to-trade ratios for the 2,397 symbols included in the Tick Size Pilot from June, 2016 to August, 2016. The mean (median) average cancel-to-trade ratio for the 2,397 Tick Size Pilot stocks is 96.8 (30.3) and the range is from 0 to 198,990. For the 100 simulations, the mean (median) cancel-to-trade ratio is 69.1 (66.1) and the range is from 10.2 to 207. While the distributions differ somewhat, the cancel-to-trade ratios from the simulations cover the majority of the observed cancel-to-trade ratios in Tick Size Pilot stocks.[13]

For the simulation set, the ratio of trade volume to order volume is the reciprocal of the cancel-to-trade ratio because the order size is fixed at one unit. In real U.S. stock data, average trade size ranges between one and two round lots. Figure 4 shows the distribution of trade volume to order volume (in percent) for the simulation set and the Tick Size Pilot stocks. The mean (median) average percentage trade-order volume for the 2,397 Tick Size Pilot stocks is 2.55 (2.05) and the range is from 0 to 81.9. For the 100 simulations, the mean (median) percentage trade-order volume is 2.03 (1.49) and the range is from 0.478 to 8.87. While the simulated trade-order volume measures cover the majority of those observed in the Tick Size Pilot data, due to the lower average trade size (which is exactly 1) the simulated trade-order volumes are lower than the observed trade-order volumes.[14]

While the simulations are not calibrated to order book depth, a comparison of simulated and observed depths is informative. Table 2 provides some summary statistics for median order book depth observed at one minute intervals for the Tick Size Pilot stocks

[13] The x-axis in Figure 3 is truncated at 400, thereby excluding some symbols with large average cancel-to-trade ratios in the Tick Pilot stocks. From a practical calibration perspective, the goal is to have the simulations span a large portion of the real distribution while keeping the mean and median simulated values between the mean and median observed values.

[14] At the expense of some additional complexity, taker trade size can be sampled from a distribution. This would permit closer calibration for trade-order volume, but adds little to the proposed study of the impact of tick size on market-maker profits. As a practical matter, smaller trade sizes in the simulations should coincide with slightly smaller depths when compared with real data.

and for the simulations with MPI=1 and α_{PJ}=0 observed at the end of every time step. In general, the mean and median depths from the simulations are between the mean and median depths observed in the Tick Size Pilot stocks, but tend toward the lower end of that range. Applying a factor of 1.5 (approximating the average trade size in observed U.S. small cap equity data) would bring the simulated depth values into the middle of the observed range of depths.

Introducing a penny jumper

Spreads decline with the introduction of a penny jumper. Also, market maker participation, profitability and trade-order volume decline while the market maker cancel-to-trade ratio increases. All of these effects are more pronounced when the penny jumper is more successful.

Figure 5 shows the impact of increasing the probability the penny jumper will participate (α_{PJ}) on quoted spreads in the one tick MPI environment. As expected, spreads increase with increasing C_λ and decrease with increasing α_{PJ}. Quoted spreads decline dramatically when the penny jumper is more likely to succeed (i.e., when both C_λ and α_{PJ} are large).

Figures 6 through 9 show the impact of increasing α_{PJ} on the market maker trade participation rate and corresponding impacts on the cancel to trade ratio and profitability when the MPI is one tick. Figure 6 shows market maker participation increases with increasing C_λ in the absence of a penny jumper, and, market maker participation declines monotonically with increasing penny jumper success. Figure 7 shows the corresponding impact on the market maker cancel to trade ratio. As market maker participation declines with increasing α_{PJ}, market maker trades decline and the cancel to trade ratio increases. Figures 8 and 9 show the associated declines in market maker total cash flow and per share cash flow, respectively, when α_{PJ} increases. A comparison of Figures 6 and 9 shows the market maker tends to earn the spread when the MPI is one tick.

Figures 10 through 13 show the impact of increasing α_{PJ} on the penny jumper trade participation rate and corresponding impacts on the cancel to trade ratio and profitability when the MPI is one tick. Figure 10 shows the penny jumper participates in more trades when spreads are wider (C_λ is larger) and, of course, when α_{PJ} is larger. Figure 11 shows the penny jumper cancel to trade ratio first increases then decreases with increasing α_{PJ}, suggesting the cancel rate increases more than the trade rate for low α_{PJ} and less than the trade rate for high α_{PJ}. Figure 11 also shows the penny jumper cancel to trade ratio declines with increasing C_λ. The penny jumper cancels only if it is first alone at the inside and then joined by a liquidity provider before a liquidity taker arrives. This sequence of events is more likely when C_λ is lower.

Figure 12 shows penny jumper cash flow first increases then decreases with increasing α_{PJ}. From a total cash flow perspective, the optimal participation rate decreases with increasing C_λ, suggesting that for the penny jumper, there is a tradeoff involving the percentage of time spent at the inside and the resulting quoted spread. When the penny jumper participates too much, the spread narrows and the penny jumper earns less overall. Figure 13 shows the penny jumper earns less on a per share basis when α_{PJ} is larger. Taken together, the penny jumper cash flow results suggest that the optimal market outcome does not coincide with the optimal penny jumper outcome. In this model the penny jumper is infinitely fast when he is given a chance to act – he sees the state of the top of the book and is immediately given an opportunity to submit his orders. If the penny jumper agent is thought of as two or more competing penny jumpers then speed becomes paramount and a policy intended to "slow down" the market could be detrimental to market quality.

Increasing the minimum price increment to five ticks

Spreads increase with the imposition of a five tick MPI. Also, market maker participation, profitability and trade-order volume increase while the market maker cancel-to-trade ratio decreases.

Figure 14 depicts the change in quoted spreads upon increasing the MPI from one to five.[15] The increase in quoted spreads declines with increasing C_λ for each choice of α_{PJ}. Increasing the MPI exhibits its greatest impact on quoted spreads for stocks that would have been quoted at smaller spreads in the one MPI scenario. Figure 15 shows the monotonic decline in quoted spreads as α_{PJ} increases for each choice of C_λ.

Figure 16 shows the increase in market maker participation upon increasing the MPI to five. While market maker participation increases in each scenario, the change is most evident when α_{PJ} is large, with participation increasing by over 50 percentage points when α_{PJ} is 0.1 and C_λ is 1. Figure 17 shows market maker participation in the five tick MPI environment increases with increasing C_λ in the absence of a penny jumper, and, market maker participation declines monotonically with increasing penny jumper success, just as in the one MPI case.

Figure 18 shows the market maker cancel-to-trade ratio declines when the MPI is increased from one to five, and, consistent with the participation results, the change in the cancel-to-trade ratio is most dramatic for large α_{PJ}. Figure 19 shows the market maker cancel-to-trade ratio increases with increasing α_{PJ} and the effect of α_{PJ} is most pronounced when C_λ is larger.

Figure 20 shows the increase in net cash flow to the market maker when the MPI is increased from one to five. The market maker is more profitable in the five-tick environment for all choices of α_{PJ} and C_λ. Figure 21 shows the decline in market maker net cash flow as the penny jumper becomes more successful (α_{PJ} increases) is similar to that

---

[15] All of the charts depicting changes in a metric result from subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation. This strategy is facilitated by using the same 100 random seeds for each simulation set.

observed in the one-tick case. Figures 22 and 23 provide similar evidence on market maker profitability measured on a per-share basis. The market maker is more profitable in the five-tick regime, but is less profitable as the penny jumper becomes more successful.

Figures 24 shows that penny jumper participation rates decline dramatically upon increasing the MPI from one to five and the decline is most dramatic for $C_\lambda = 1$ (i.e., when spreads were lowest in the one tick regime), and, Figure 25 shows the penny jumper participation rate increases with increasing α_{PJ}, as expected.

Figure 26 shows the penny jumper cancel to trade ratio generally increases when the MPI is increased to five ticks, but, again the effect is most pronounced for $C_\lambda = 1$. When C_λ is 50, the penny jumper cancel to trade ratio decreases for small α_{PJ}. Figure 27 shows that the penny jumper cancel to trade ratio increases monotonically with increasing α_{PJ}, in contrast to the inverted U-shape observed in the one tick MPI scenario.

Figure 28 shows penny jumper net cash flow declines when the MPI is increased to five ticks, and, the decline is generally larger for increasing C_λ. The declines in net cash flow to the penny jumper are largest for α_{PJ} in the range of 0.01 to 0.05. Figure 29 shows net cash flow to the penny jumper is maximized when the penny jumper limits his participation, similar to the one-tick MPI case (Figure 12). However, Figure 29 also shows there is considerable variation in the net cash flow results obtained from the five-tick MPI scenario. Figure 30 shows the decline in penny jumper profitability on a per-share basis. The impact of increasing the MPI from one to five ticks is much less pronounced for higher values of α_{PJ}, because spreads increased more (Figure 14) while participation decreased more (Figure 24) and net cash flow decreased less (Figure 28). Figure 31 shows the penny jumper earns less on a per share basis when α_{PJ} is larger in the five-tick environment, which is similar to the results observed in the one-tick environment.

Comparing the model to the Tick Pilot Test Groups

Figure 32 shows average quoted spreads from the four simulation sets with the MPI set to five and α_{PJ} set to zero match the distribution of daily duration-weighted average quoted spread in cents for 1,183 symbols included in the Tick Size Pilot Test Groups from December, 2016 to February, 2017. The mean (median) average daily quoted spread for the Tick Size Pilot Test Group stocks is 19.2 (8.70) cents and the range is from 5.00 cents to $9.15. The mean (median) quoted spread for the simulation sets ranges from 6.69 (5.39) ticks with C_λ=1 to 16.4 (13.3) ticks with C_λ=50. Taken together, the quoted spread distributions from the simulations with the MPI set to five match the overall shape and range of the average quoted spreads observed in the Tick Size Pilot Test Group stocks following implementation of the pilot.

Table 3 provides some summary statistics for median order book depth observed at one minute intervals for the Tick Size Pilot Test Group stocks and for the simulations with MPI=5 and α_{PJ}=0 observed at the end of every time step. In general, the mean and median depths from the simulations are between the mean and median depths observed in the Tick

Size Pilot Test Group stocks for levels beyond the inside. Median inside depths from the simulations are a bit lower than depths observed in the Tick Size Pilot Test group stocks. Applying a factor of 1.5 (approximating the average trade size in observed U.S. small cap equity data) would bring the simulated inside depth values closer to the middle of the observed range of depths.

A caution, a limitation, and an extension

Equating particular agents with specific market participants is common in agent based models, but it doesn't necessarily have to be so. In many cases it is helpful to think of actual market participants as amalgams of simple rule-following agents. Actual market participants exhibit combinations of the caricatured behavior depicted in the zero intelligence agents. For example, registered market makers will occasionally employ a quote-matching strategy similar to the penny jumper or cross the spread and behave like a taker agent. In the current model, this real market maker would be a combination of the market maker agent, the penny jumper and the liquidity taker agent. Put another way, the market maker profits in the model are best interpreted as profits accruing to the market-making strategy. This does not alter the basic findings of the model, but merely imparts a cautionary note when applying the findings to real market participants.

In a similar vein, the taker agent is a stylized version of a noise (or uninformed) trader. In real markets, uninformed traders can choose to take the displayed price or post a limit order in an attempt to obtain a better price at a later time. Actual market makers, who compete with noise traders for precedence in the limit order book, optimally choose to join other liquidity providers or step in front of them. The zero-intelligence agent-based model does not capture the market maker's choice to quote match or the noise trader's decision to provide instead of take liquidity. Nor does the model capture how these choices change with different minimum price increments. This is a general limitation of all zero-intelligence models. By construction, zero-intelligence models remove strategic behavior from the agent's arsenal of decision rules.[16]

A zero-intelligence model could provide some insight into market maker participation and profitability in the presence of a noise trader with a simple decision rule to post or take liquidity. But it still would not evaluate the make versus take decisions of real market participants. An agent based model augmented with learning agents would endogenize the noise trader's make-take decision and the market maker's decision to quote match or cross the spread.[17] With such a model in hand, we could ask more insightful questions about the proposed Tick Pilot:

---

[16] See Ladley (2012) for an overview of zero-intelligence models.
[17] Farmer, Patelli and Zovko (2005) argue for a modeling strategy that starts with zero intelligence and is subsequently extended by adding moderately more intelligent agents.

- Will institutional traders increase their liquidity-providing behavior for five-tick stocks?
- Will other institutions optimally choose to provide liquidity acting as principal, thereby competing with designated market-makers in small cap stocks?
- If so, will this encourage more institutions to provide research services and otherwise "sponsor" small cap stocks as part of the public offering process? Or will they free-ride on the public offering services of other broker-dealers?

Recapitulation

I build an agent based model of a modern limit order market with four agents – liquidity provider, market maker, penny jumper and liquidity taker – and then calibrate to metrics observed in pre-Tick Size Pilot stocks. After calibrating on spreads and cancel to trade ratios and checking depth, I test the impact of first adding a penny jumper and then increasing the MPI from one tick to five ticks with four scenarios distinguished by the quoted spread (i.e., the choice of C_λ).

When the MPI is one, increasing C_λ generates larger quoted spreads, which is by design. Increasing the penny jumper's opportunity to participate (increasing α_{PJ}) results in smaller quoted spreads and decreased market maker participation in trades. As a result, the market maker cancel to trade ratio increases and profitability decreases. The penny jumper's participation in trades increases with both α_{PJ} and C_λ. As a result, the penny jumper cancel to trade ratio decreases and profitability increases with rising C_λ. Unlike the market maker, the penny jumper's cancel to trade ratio and total cash flow first increases, then decreases with increasing α_{PJ}. However, like the market maker, penny jumper cash flow per share declines with decreasing spreads as α_{PJ} increases. Overall, the penny jumper is more successful – and the market maker is less so – when the penny jumper has more opportunity to participate and/or when quoted spreads are wider.

Increasing the MPI to five generates larger quoted spreads and improves market maker participation in trades for all choices of α_{PJ} and C_λ. The market maker cancel to trade ratio decreases and profitability increases. Penny jumper participation decreases when the MPI is five. His cancel to trade ratio increases while his profitability decreases.

The agent based model provides answers to the policy questions posed in the opening section:

- Increasing the minimum pricing increment (MPI) necessarily implies that spreads will widen, thereby making market makers more profitable.
- Spreads decrease, and by extension, market maker participation and profitability decrease when penny jumpers are active.
- Increasing the MPI mitigates the impact of penny jumpers on market maker profitability.

This study demonstrates the feasibility of using agent-based models to inform market structure policy and provides evidence in support of adding agent-based modeling to the regulatory analytical toolbox. Extensions to this model include formulating a set of liquidity providers to match the price dynamics observed in real markets and adding learning agents to the model. Both of these extensions are areas of current research.

References

Bookstaber, Richard, Foley, Michael D., and Brian F. Tivnan, 2015, Market Liquidity and Heterogeneity in the Investor Decision Cycle, Office of Financial Research Working Paper 15-03.

Buchanan, Mark, 2009, Meltdown Modelling, *Nature*, Vol 460, 680 – 682.

Farmer, J. Doyne, Patelli, Paolo, and Illija I. Zovko, 2005, The Predictive Power of Zero Intelligence in Financial Markets, *PNAS*, 102 (6), 2254 – 2259.

Harris, Larry, 2015, Trading and Electronic Markets: What Investment Professionals Need to Know, *CFA Institute Research Foundation*.

Ladley, Dan, 2012, Zero Intelligence in Economics and Finance, *The Knowledge Engineering Review*, 27 (2), 273 – 286.

Preis, T., Golke, S., Paul, W., and J. J. Schneider, 2006, Multi-agent-based Order Book Model of Financial Markets, *Europhysics Letters*, 75 (3), 510 – 516.

Preis, Tobias, Golke, Sebastian, Paul, Wolfgang, and Johannes J. Schneider, 2007, Statistical Analysis of Financial Returns for a Multiagent Order Book Model of Asset Trading, *Physical Review E*, 76 (1), 1 – 13.

Methods

<u>Agents</u>

The zero intelligence (ZI) traders are either liquidity providers or liquidity takers. In the simplest Preis et al. (2006) model, liquidity takers trade with a probability μ per time step and buy with probability q_{take}. Both of these probabilities are held constant throughout the simulation. Liquidity providers send an add order with probability α per time step and submit a bid order with probability $q_{provide}$. Liquidity providers also cancel a fraction δ of their outstanding orders each time step. In this simplest of models, these three probabilities are held constant throughout the simulation as well. Orders are priced on a uniform grid such that an add order will never result in a trade. Limit buy orders are priced on the uniform interval $[p_a - 1 - p_{int}, p_a - 1]$ with equal probability and limit sell orders on $[p_b + 1, p_b + 1 + p_{int}]$, where p_a is the best ask price, p_b is the best bid price and p_{int} is the pricing interval. The order size is one for takers and providers.

In the first extension of the Preis et al. model, orders are priced on an exponential distribution of prices such that an add order will never result in a trade. Limit buy orders are priced on the interval $[p_a - 1 - \eta, p_a - 1]$ and limit sell orders on $[p_b + 1, p_b + 1 + \eta]$, where p_a is the best ask price, p_b is the best bid price and $\eta = \lfloor -\lambda_0 \ln(x) \rfloor$ with $0<x\leq 1$. Smaller values for λ_0 imply prices closer to the prevailing inside prices.

The second extension of the Preis et al. model includes a time-varying process for q_{take} and its impact on the exponential distribution of prices for liquidity providers. In previous models, q_{take} is held constant at 0.5. Preis et al. (2007) introduce a feedback random walk for q_{take} starting the simulation at its expected value of 0.5 and then "incremented and decremented by a value of Δs after (each step) … (where) the probability for returning to the average value of ½ is given by ½ + $|q_{taker}(t) - ½|$ and thus the probability for departing from the mean value is given by ½ - $|q_{taker}(t) - ½|$." Also, λ_0 is replaced with λ_t, a function of the deviation of q_{take} from its expected value:

$\lambda_t = \lambda_0 \left(1 + \frac{\left|q_{take,t} - \frac{1}{2}\right|}{\sqrt{\langle \left[q_{take,t} - \frac{1}{2}\right]^2 \rangle}} C_\lambda \right)$, where the average variation, $\langle \left[q_{take,t} - \frac{1}{2}\right]^2 \rangle$, is computed in a separate simulation.

In the third extension, due to Bookstaber et al. (2015), providers and takers are assigned random inter-arrival intervals drawn from an exponential distribution with means of α and μ, respectively. These intervals are held constant for each provider and taker throughout the simulation. Providers and takers also have heterogeneous order sizes taken from [1, 5, 10, 25, 50]. Once assigned, the inter-arrival intervals are adjusted to reflect the larger order sizes and these are held constant for each provider and taker throughout the simulation. The parameters α and μ are the means of the exponential distribution for arrival times.[18] In the current study, I limit order size to 1.

In the fourth extension, I add a new liquidity providing agent - called a market maker - with its own rule set. Market makers submit orders of size 1 and arrive at the market on each time step. Market makers submit 12 orders from a choice of 60 prices on a discrete uniform interval: [p_b - j - 59, p_b - j] for buys where $j = \begin{cases} 0, & best\ bid\ size > 1 \\ 1, & best\ bid\ size = 1 \end{cases}$ and [p_a + j, p_a + j + 59] for sells where $j = \begin{cases} 0, & best\ ask\ size > 1 \\ 1, & best\ ask\ size = 1 \end{cases}$. Like the basic liquidity providers, market makers cancel the same fraction δ of their outstanding orders each time step.

When the minimum price increment is set to five, the liquidity provider rounds away from the inside by rounding down to the nearest five-tick increment (nickel) for bids and up to the nearest nickel for asks. In this sense, the actual draws from the exponential distribution are reservation prices. The market maker does not observe the reservation price but must "infer" it from the posted price at the nearest nickel. To affect this in the model, the market maker chooses from a set of thirteen prices with an adjustment to the uniform probabilities (1/12) to account for the inferred reservation price. For example, a best bid price of 995 could result from an actual draw of 995, 996, 997, 998 or 999. The market maker infers the expected value of 997 and adjusts the uniform probabilities to [1/30, 1/12, 1/12, 1/12, 1/12, 1/12, 1/12, 1/12, 1/12, 1/12, 1/12, 1/12, 1/20] for prices in [935, 940, 945, 950, 955, 960, 965, 970, 975, 980, 985, 990, 995]. In other words, the market maker's (hypothetical) one tick price grid is shifted down by two ticks to reflect the expected upper limit in the liquidity provider's (hypothetical) one tick price grid.[19]

In the current model, I add another liquidity providing agent – called a penny jumper – with a simple rule set. If there is a permissible price inside the quoted spread and the penny jumper is not currently alone at the inside price(s), then the penny jumper adds a limit order for one share priced one minimum price increment better than the prevailing price(s) and cancels any orders no longer alone at the inside. In other words, when permitted by the minimum price increment and given the opportunity, the penny jumper will attempt to jump ahead of any other liquidity provider, except himself. The opportunity is modeled by α_{PJ}, a probability ranging from 0 to 0.1. After each arrival of any of the other agents, the penny jumper has a probability of α_{PJ} of receiving the opportunity to jump in front of posted quotes.

---

[18] The median inter-arrival time is less than the mean inter-arrival time for an exponential distribution. For example, if α is 0.1, then the mean inter-arrival time is $1/\alpha = 10$ and the median inter-arrival time is $\ln(2)/\alpha = 6.93$. More than half of the traders will have inter-arrival times less than the mean.

[19] A considerably more complex modeling solution would be to carry the actual one-tick prices in memory and use these prices to establish the uniform grid from which the market maker chooses and subsequently rounds. Since these are simulations, using the expected values in a probabilistic manner achieves the same effect.

The Matching Engine

Orders are matched by price then time. Incoming orders are processed in two steps. If the incoming order is a cancel order it is processed in the order book. If it is an add order, it is compared to the prevailing best bid and ask prices. If the incoming add order is not priced to trade (non-marketable), it is added to the order book. If it is marketable, the order is sent to the matching engine where one of two events can occur. If the entire order is priced such that the order size is less than the size available on the book at that price, then the order results in one or more trades, the order book is updated to reflect removal or modification of the resting orders that participated in the trade and confirmations are sent to the traders behind the resting orders. If the order is not priced to trade the total order size, then the remainder of the order is placed in the order book. In this study, takers always submit marketable orders.

Simulation

The simulation performs many functions. It creates the order book/matching engine (the exchange) and the traders, seeds the order book with one ask order and one bid order, primes the order book, runs the simulation loop and finally collects some output in permanent storage. To seed the order book, one ask order from the interval [1000001, 1000001 + p_{int}] and one bid order from the interval [999999 – p_{int}, 999999] with p_{int} = 2000 are placed on the book. Priming the order book is accomplished by running the simulation for 20 time steps with takers, penny jumpers and market makers excluded. For each time step, each liquidity provider is selected according to their arrival interval (determined by α) to submit one add order. The remaining simulation potentially includes market makers, takers and penny jumpers as well. For each time step, all of the providers and market makers as well as zero or more takers submit orders in a random sequence. Penny jumpers are (conceptually) inserted into this sequence according to a probability parameter. One time step works like this:

1. If it is the first time step, compute top of book prices and sizes.
2. Randomize all of the providers, market makers and any of the takers with an arrival interval consistent with the current time step (i.e., the current time step is evenly divisible by the arrival interval) in a queue.
3. For each agent in the queue:
    a. If the agent is a provider:
        i. If provider arrival interval is consistent with the current time step:
            1. Provider: Process top of book signal and submit order
            2. Exchange: Process submitted order
            3. Update top of book
        ii. Agent potentially cancels orders
        iii. If agent cancels orders:

            1. Exchange: Process submitted order(s) and send cancel confirmation
            2. Provider: Process confirmation, update agent orderbook
            3. Update top of book
    b. Else if agent is a market maker:
        i. If market maker arrival interval is consistent with the current time step:
            1. Market maker: process top of book signal and submit orders
            2. Exchange: process submitted orders
            3. Update top of book
        ii. Agent potentially cancels orders
        iii. If agent cancels orders:
            1. Exchange: Process submitted order(s) and send cancel confirmation
            2. Market maker: Process confirmation, update agent orderbook
            3. Update top of book
    c. Else (agent is a taker):
        i. Taker: process top of book signal and submit order
        ii. Exchange: process submitted order, send confirmation to liquidity provider or market maker
        iii. Liquidity Provider, Penny Jumper or Market Maker: updates own orderbook
        iv. Update top of book
    d. If penny jumper is active:
        i. Penny Jumper: process top of book signal and submit order(s)
        ii. Update top of book

<u>File Structure</u>

```
pyziabm/
    __init__.py
    docs/
    notebooks/
    pyziabm/
        __init__.py
        orderbook3.py
        runner2017mpi_r3.py
        runner2017mpi_r4.py
        runwrapper2017mpi_r3.py
        runwrapper2017mpi_r3x.py
        runwrapper2017mpi_r4py
```

trader2017_r3.py
trader2017_r3.py
tests/
__init__.py
testOrderbook3.py
testTrader2017_r3.py

See https://github.com/JackBenny39/pyziabm for the python code.

Data

Quoted Spread
Simulation: the average of the difference between the best (lowest) ask price and the best (highest) bid price observed at the end of each time step.
Tick Size Pilot Stocks: duration-weighted average difference between the best (lowest) ask price and the best (highest) bid price taken from the MIDAS SIP (top-of-book) feed. Quotes are included if they occur after the opening trade message and the first LULD band status message, the best ask and the best bid are inside or at the LULD bands, and the quoted spread is positive (locked and crossed markets are excluded).

Depth
Simulation: the mean cumulative depth, in shares, observed at the end of each time step.
Tick Size Pilot Stocks: the median cumulative depth, in round lots (100 shares), observed at the end of each minute. Depth is taken from the MIDAS Direct feeds for each exchange.

Cancel to Trade Ratio
Simulation: the ratio of cancel messages to trade messages.
Tick Size Pilot Stocks: the ratio of cancel messages to trade messages for trades involving displayed orders. Messages are counted from the MIDAS Direct feeds for each exchange except NYSE and NYSE-MKT.[20]

Trade – Order Volume
Simulation: the ratio of trade volume (in shares) to order volume (in shares), in percent.
Tick Size Pilot Stocks: the ratio of trade volume (in shares) for trades involving displayed orders to order volume (in shares), in percent. Volumes are counted from the MIDAS Direct feeds for each exchange, including NYSE and NYSE-MKT.

Participation
Simulation: percentage of trades involving a liquidity provider, market maker or penny jumper.

Total Net Cash Flow
Simulation: cash flow to the liquidity provider, market maker or penny jumper. This measure excludes the impact of ending inventory by matching buys and sells in sequence

[20] These two direct feeds are excluded because they used the level-book reporting method during the Tick Pilot sample date range used in this study. See Order Book Reporting Methods and Their Impact on Some Market Activity Measures for more details on the different types of MIDAS direct feeds.

and excluding any unmatched orders. In practice, this is achieved by capping the sample of orders by the minimum of the buy order count and sell order count.

Net Cash Flow Per Share

Simulation: Total net cash flow divided by the number of shares matched (excluding ending inventory).

Table 1
Model Parameters

Description	Parameter	Value
Simulations Per Trial Set		100
Number of liquidity providers	N_{LP}	38
Number of market makers	N_{MM}	1
Number of market maker quotes		12
Market maker quote range		60 ticks
Number of liquidity takers	N_{LT}	50
Maximum order quantity	Q_{max}	1
Number of time steps	T	100000
Liquidity provider arrival rate parameter	α	0.0375
Liquidity taker arrival rate parameter	μ	0.001
Liquidity provider and market maker cancel rate	δ	0.025
Liquidity provider and market maker bid probability	$q_{provide}$	0.5
Fixed scale parameter for the liquidity provider price distribution	λ_0	100
Adjustable scale parameter for the liquidity provider price distribution	C_λ	[1, 5, 10, 50]
Variance parameter for the q_{take} mean-reverting random walk	Δs	0.001
Minimum price increment	MPI	[1, 5]
Probability the penny jumper will have the opportunity to participate	α_{PJ}	[0, 0.001, 0.005, 0.01, 0.05, 0.1]

Table 2
Median Depth; MPI=1

			C_λ			
Level	Statistic	Tick Pilot Pre	1	5	10	50
Inside	Daily Mean	6.09	2.41	2.43	2.48	2.64
	Daily Median	3.05	2.5	2.5	2.5	2.5
	Daily Minimum	0.01	2	2	2	2
	Daily Maximum	2250	3	3	3	3
2	Daily Mean	12.3	5.57	5.50	5.61	5.97
	Daily Median	5.1	5.5	5.5	5.5	6
	Daily Minimum	0.01	4	4.5	4.5	5
	Daily Maximum	6181	7	6.5	6.5	6.5
3	Daily Mean	17.9	9.11	8.81	8.88	9.39
	Daily Median	6.76	9.25	9	9	9.5
	Daily Minimum	0.02	6	6.5	7	8
	Daily Maximum	7897	11	10	10	10.5
4	Daily Mean	23.6	12.9	12.3	12.3	12.9
	Daily Median	8.5	13	12.5	12.5	13
	Daily Minimum	0.02	8.5	9	9.5	11
	Daily Maximum	9300	15.5	14.5	14	14
5	Daily Mean	29.7	16.9	16.0	16.0	16.5
	Daily Median	10.3	17.5	16.5	16	16.5
	Daily Minimum	0.02	11	11.5	12	14
	Daily Maximum	10339	20	18.5	18	18
6	Daily Mean	36.5	21.1	19.8	19.7	20.2
	Daily Median	13	21.5	20.25	20	20.5
	Daily Minimum	0.02	14	14	15	17
	Daily Maximum	11284	25	22.5	22	22
11	Daily Mean	65.3	42.7	39.8	39.1	39.0
	Daily Median	24.4	43.5	40.5	40	39.5
	Daily Minimum	0.025	28.5	28	29	33.5
	Daily Maximum	13866	49	44	42.5	41.5

Tick Pilot Data: median depth in round lots observed at one minute intervals; 2,397 symbols from June, 2016 to August, 2016.
Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=0.

Table 3

Median Depth; MPI=5

			C_λ			
Level	Statistic	Tick Pilot Post	1	5	10	50
Inside	Daily Mean	54.2	9.3	8.25	8.02	7.98
	Daily Median	10.5	9.5	8.5	8	8
	Daily Minimum	0.02	5	5.5	6	7
	Daily Maximum	3286853	13	11.5	11	10
2	Daily Mean	54.5	9.3	8.25	8.02	7.98
	Daily Median	10.5	9.5	8.5	8	8
	Daily Minimum	0.02	5	5.5	6	7
	Daily Maximum	3286853	13	11.5	11	10
3	Daily Mean	54.5	9.3	8.25	8.02	7.98
	Daily Median	10.5	9.5	8.5	8	8
	Daily Minimum	0.02	5	5.5	6	7
	Daily Maximum	3286853	13	11.5	11	10
4	Daily Mean	54.5	9.3	8.25	8.02	7.98
	Daily Median	10.5	9.5	8.5	8	8
	Daily Minimum	0.02	5	5.5	6	7
	Daily Maximum	3286853	13	11.5	11	10
5	Daily Mean	54.5	9.3	8.25	8.02	7.98
	Daily Median	10.5	9.5	8.5	8	8
	Daily Minimum	0.02	5	5.5	6	7
	Daily Maximum	3286853	13	11.5	11	10
6	Daily Mean	114	30.2	27.5	26.8	26.7
	Daily Median	21.6	31.5	28.5	27.5	27
	Daily Minimum	0.06	18.3	18.5	19.5	22.5
	Daily Maximum	4139240	36.5	32.5	31	30
11	Daily Mean	165	51.9	47.8	46.6	46.0
	Daily Median	33.9	53.3	49.0	47.5	46.5
	Daily Minimum	0.06	33.5	33.5	34.5	39
	Daily Maximum	4356545	60.5	53.5	51.5	49.5

Tick Pilot Data: median depth in round lots observed at one minute intervals; 1,183 Test Group symbols from December, 2016 to February, 2017.

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=0.

Figure 1

Simulation Strategy

Eight groups vary by combinations of the MPI and C_λ, a parameter for the liquidity provider price distribution. Within each group, there are 100 simulation sets for each choice of α_{PJ}, the probability that the penny jumper will participate.

		C_λ			
		1	5	10	50
MPI	1	Group 1	Group 2	Group 3	Group 4
	5	Group 5	Group 6	Group 7	Group 8

Group X

α_{PJ}	# of Sims
0	100
0.001	100
0.005	100
0.01	100
0.05	100
0.1	100

Figure 2
Spread Calibration Results
Average quoted spread distributions from Groups 1 through 4 (C_λ = [1, 5, 10, 50] , MPI = 1) with no penny jumper (α_{PJ} = 0) match the overall shape and range of the average quoted spreads observed in the Tick Size Pilot stocks prior to implementation of the pilot.

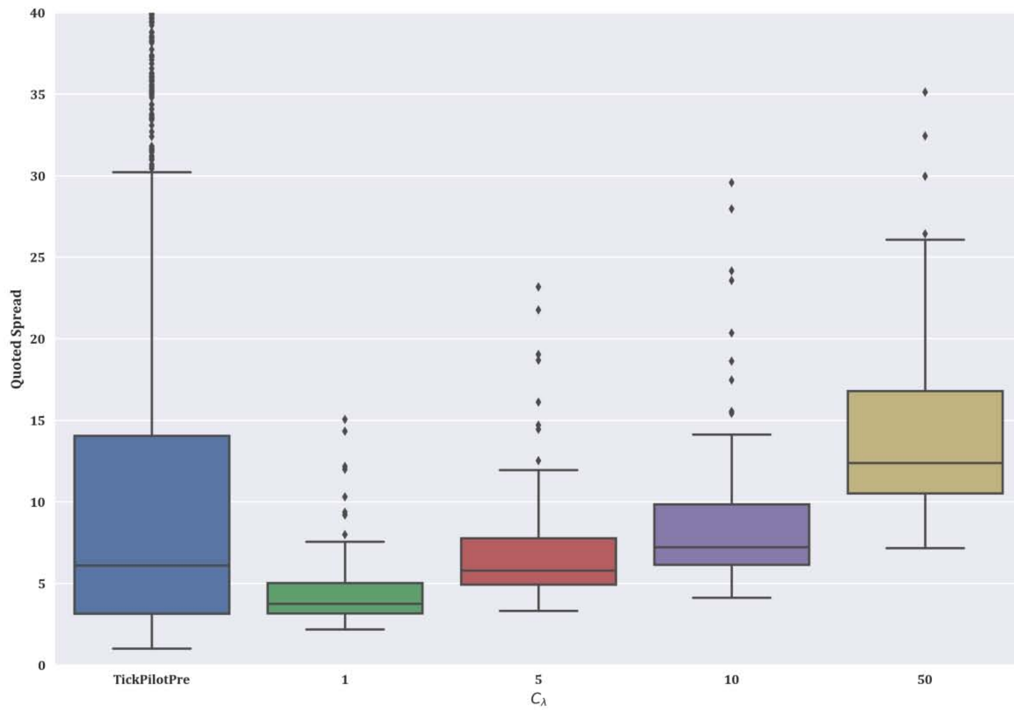

Tick Pilot Data: distribution of daily duration-weighted average quoted spread in cents; 2,397 symbols from June, 2016 to August, 2016.
Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=0.

Figure 3
Cancel-to-Trade Ratio Calibration Results
Average cancel-to-trade ratio distribution from Group 1 (C_λ = 1, MPI = 1) with no penny jumper (α_{PJ} = 0) coincide with the distribution of cancel-to-trade ratios observed in the Tick Size Pilot stocks prior to implementation of the pilot. Using the same 100 random seeds ensures each simulation set has the same number of trades.

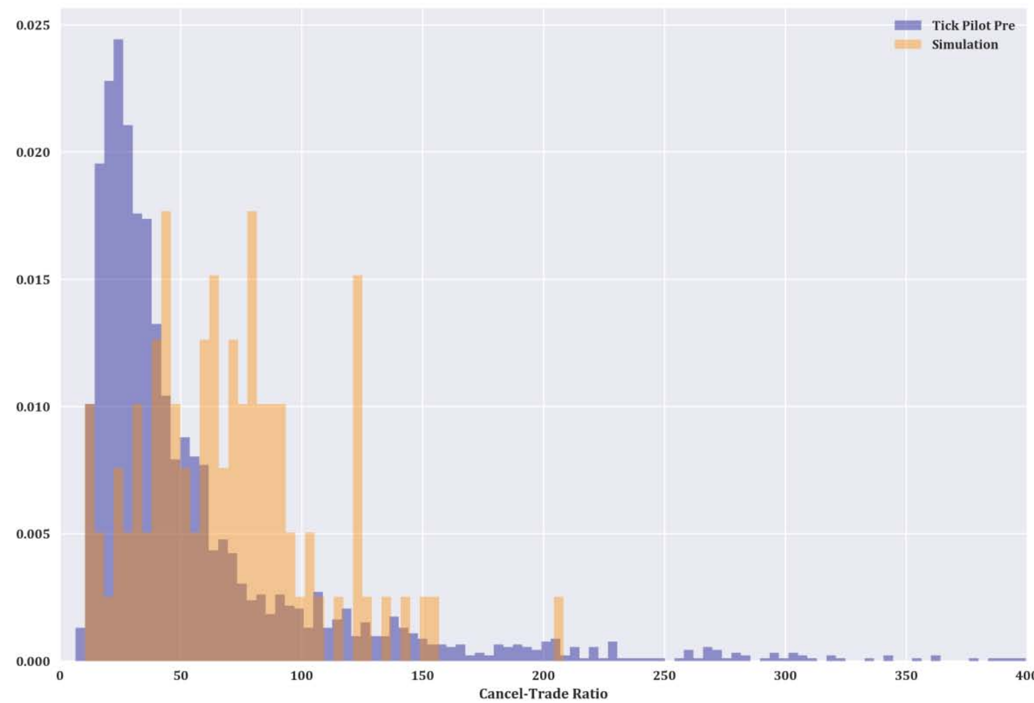

Tick Pilot Data: distribution of average daily cancel-to-trade ratios; 2,397 symbols from June, 2016 to August, 2016.
Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=1, Δs=0.001, MPI=1, α_{PJ}=0.

Figure 4

Trade-Order Volume (%) Calibration Results

Average trade-order volume (%) distribution from Group 1 (C_λ = 1, MPI = 1) with no penny jumper (α_{PJ} = 0) compared with the distribution of trade-order volume observed in the Tick Size Pilot stocks prior to implementation of the pilot. The simulations use a fixed order size of one. Tick Size Pilot stocks have an average trade size between one and two round lots. Using the same 100 random seeds ensures each simulation set has the same number of trades.

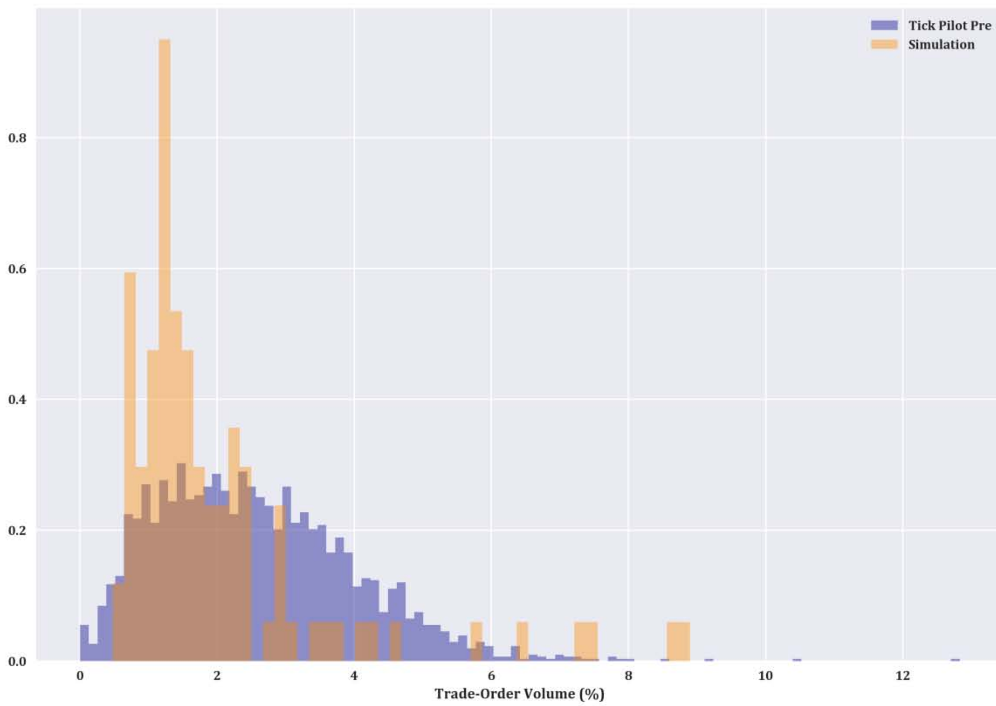

Tick Pilot Data: distribution of average daily trade-order volume (%); 2,397 symbols from June, 2016 to August, 2016.

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=1, Δs=0.001, MPI=1, α_{PJ}=0.

Figure 5
Impact of α_{PJ} on Quoted Spreads; MPI = 1
Quoted spreads increase as C_λ increases from 1 to 50 and decrease as α_{PJ} increases within each group. The impact of α_{PJ} on quoted spreads is most pronounced (in nominal ticks) when quoted spreads are large (Group 4, C_λ = 50).

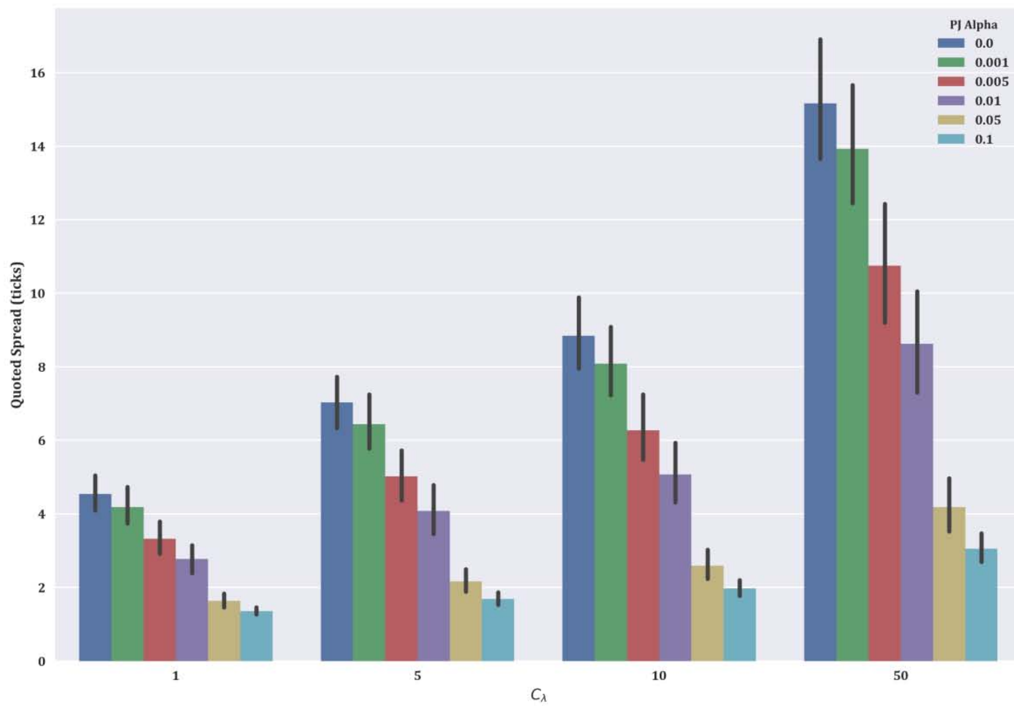

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 6

Impact of α_{PJ} on Market Maker Participation; MPI = 1

Market maker participation increases as C_λ increases from 1 to 50 for $\alpha_{PJ} \leq 0.01$ and decreases with increasing C_λ for $\alpha_{PJ} \geq 0.05$. Within each group, market maker participation declines monotonically with increasing α_{PJ}.

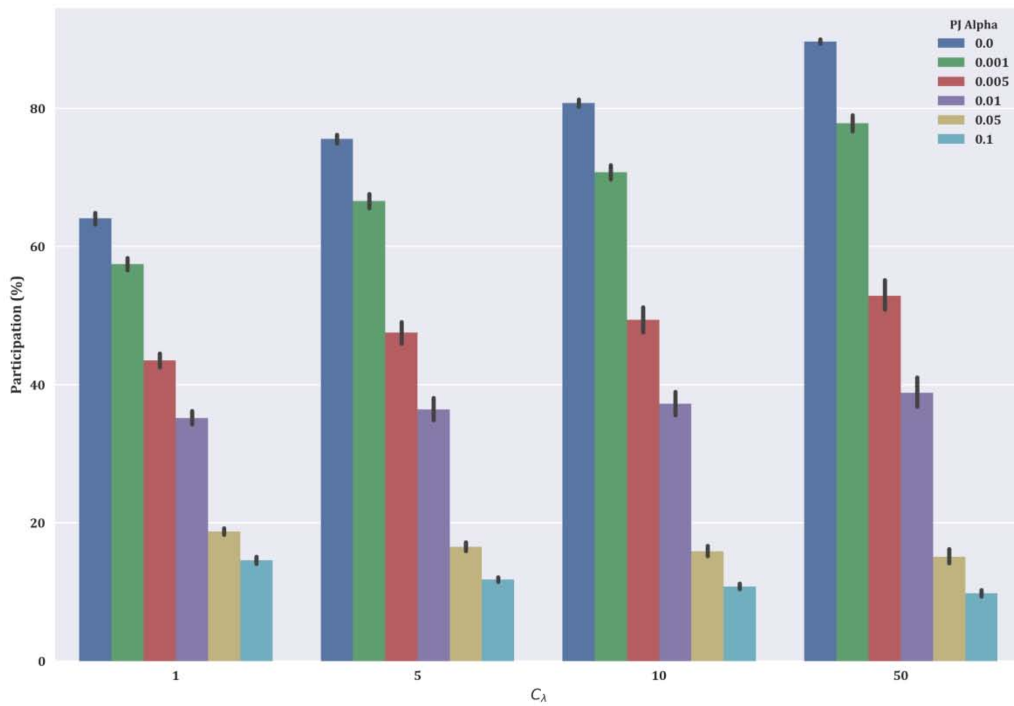

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 7

Impact of α_{PJ} on Market Maker Cancel-To-Trade Ratio; MPI = 1

The market maker cancel-to-trade ratio decreases as C_λ increases from 1 to 50 for $\alpha_{PJ} \leq 0.01$ and increases with increasing C_λ for $\alpha_{PJ} \geq 0.05$. Within each group, the market maker cancel-to-trade ratio increases with increasing α_{PJ}.

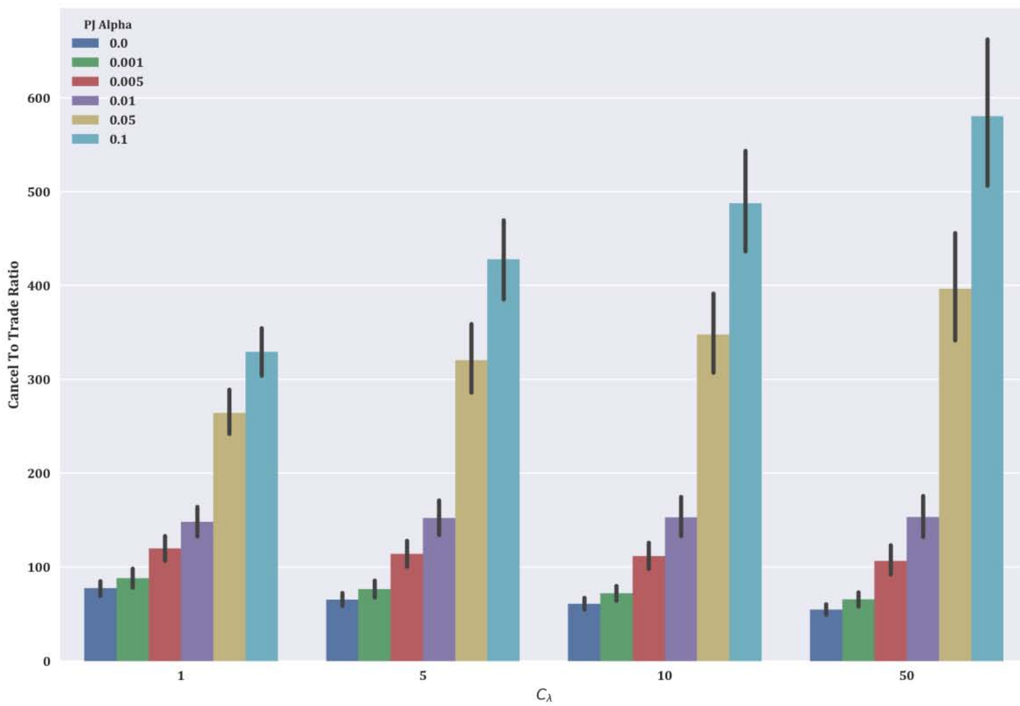

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 8
Impact of α_{PJ} on Market Maker Total Net Cash Flow; MPI = 1
Market maker total net cash flow increases as C_λ increases from 1 to 50. Within each group, market maker total net cash flow decreases with increasing α_{PJ}.

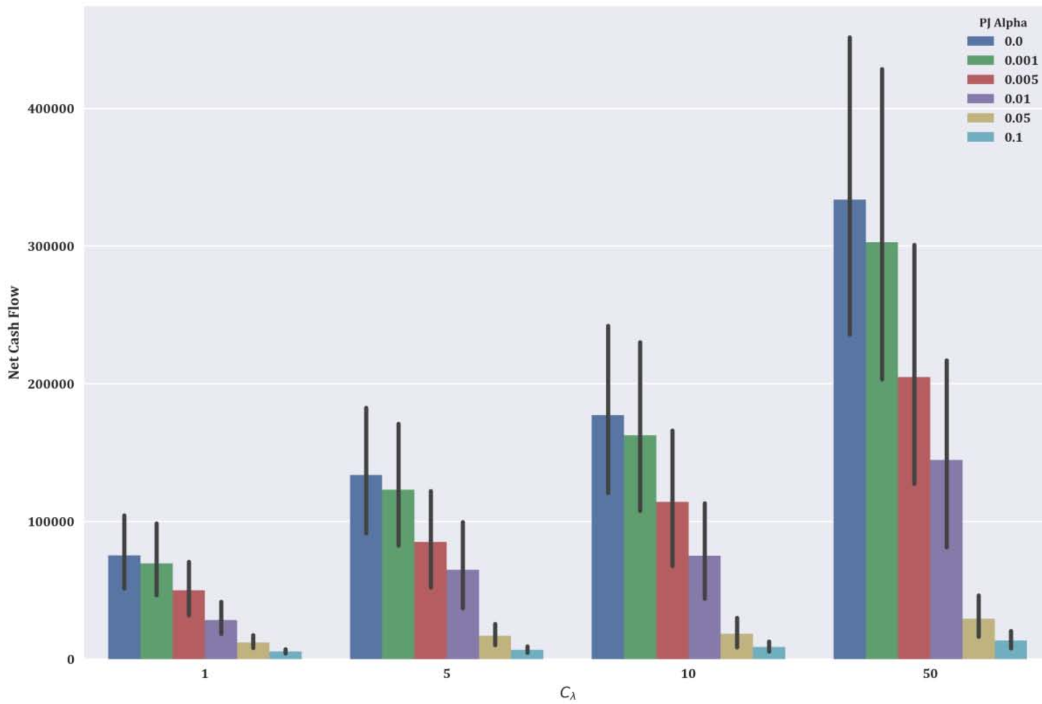

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 9
Impact of α_{PJ} on Market Maker Net Cash Flow Per Share; MPI = 1
Market maker net cash flow per share increases as C_λ increases from 1 to 50. Within each group, market maker net cash flow per share decreases with increasing α_{PJ}.

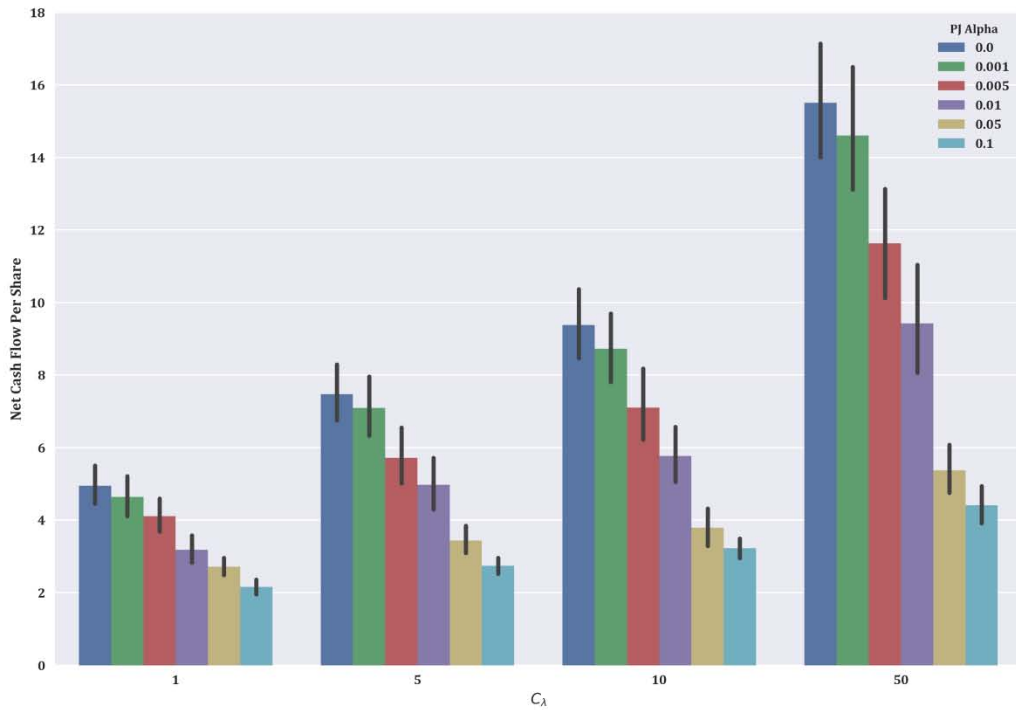

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 10
Impact of α_{PJ} on Penny Jumper Participation; MPI = 1
Penny jumper participation increases as C_λ increases from 1 to 50. Within each group, penny jumper participation increases monotonically with increasing α_{PJ}.

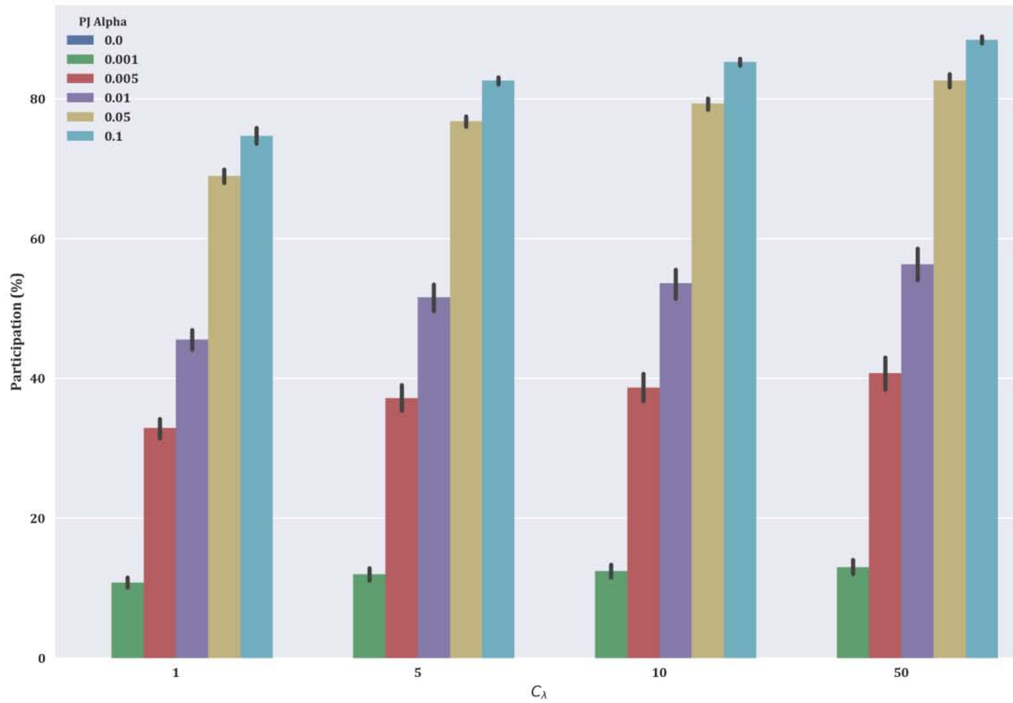

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 11
Impact of α_{PJ} on Penny Jumper Cancel-To-Trade Ratio; MPI = 1
The penny jumper cancel-to-trade ratio decreases as C_λ increases from 1 to 50. Within each group, the penny jumper cancel-to-trade ratio displays an inverted U-shape (i.e., first increases, then decreases) with increasing α_{PJ}.

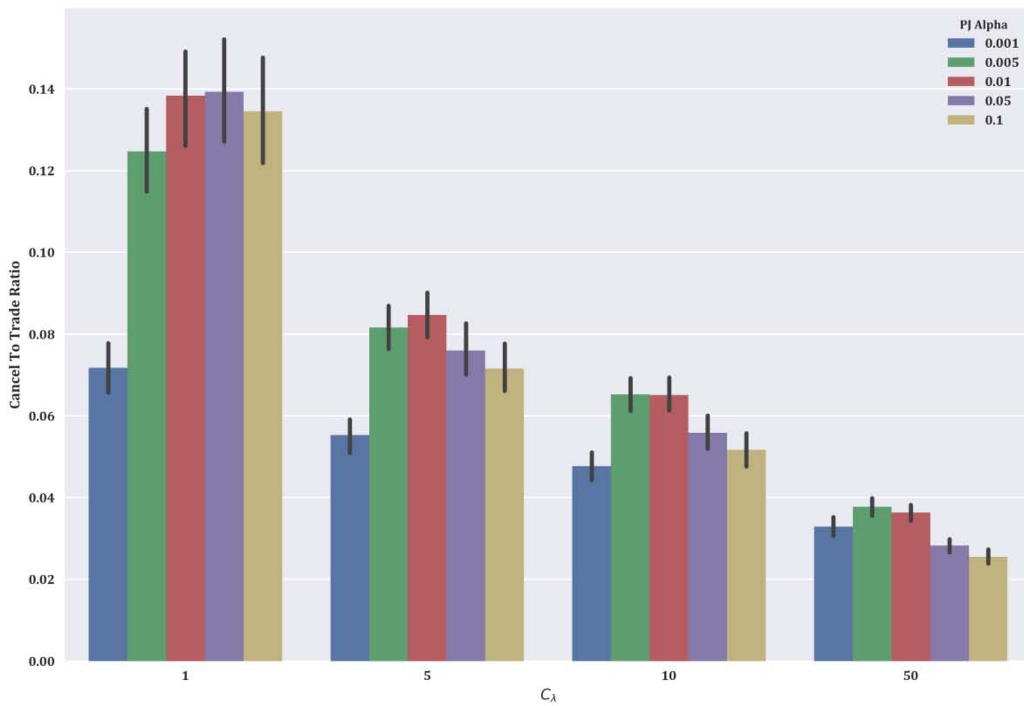

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 12

Impact of α_{PJ} on Penny Jumper Total Net Cash Flow; MPI = 1

Penny jumper total net cash flow increases as C_λ increases from 1 to 50. Within each group, penny jumper total net cash flow displays an inverted U-shape (i.e., first increases, then decreases) with increasing α_{PJ}.

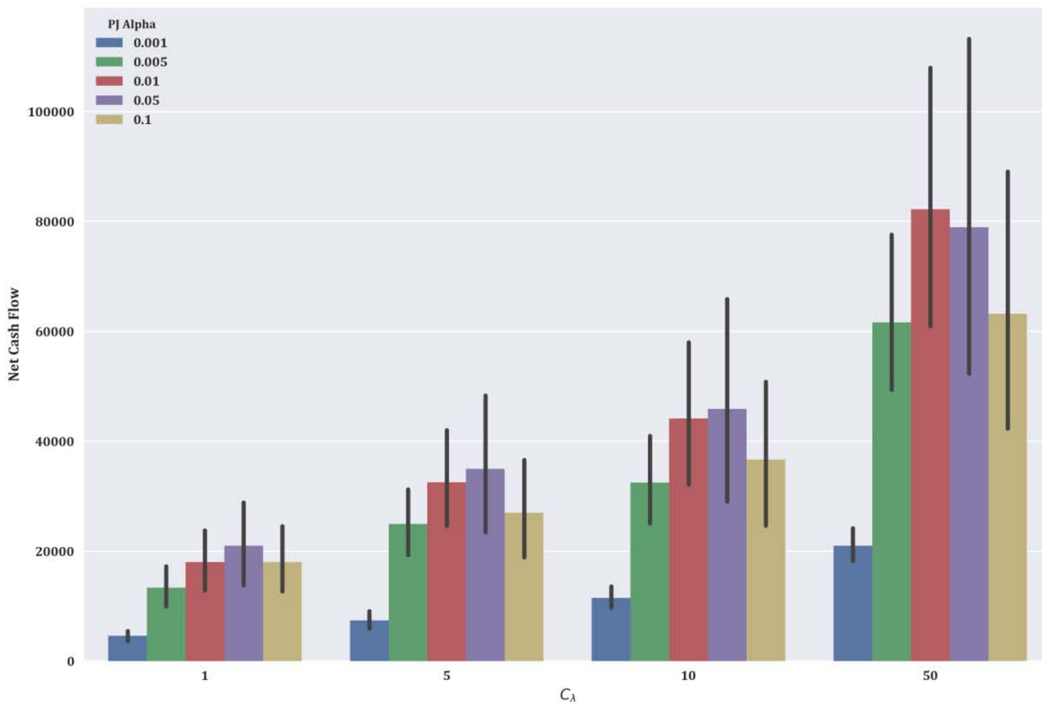

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 13
Impact of α_{PJ} on Penny Jumper Net Cash Flow Per Share; MPI = 1
Penny jumper net cash flow per share increases as C_λ increases from 1 to 50. Within each group, penny jumper net cash flow per share decreases with increasing α_{PJ}.

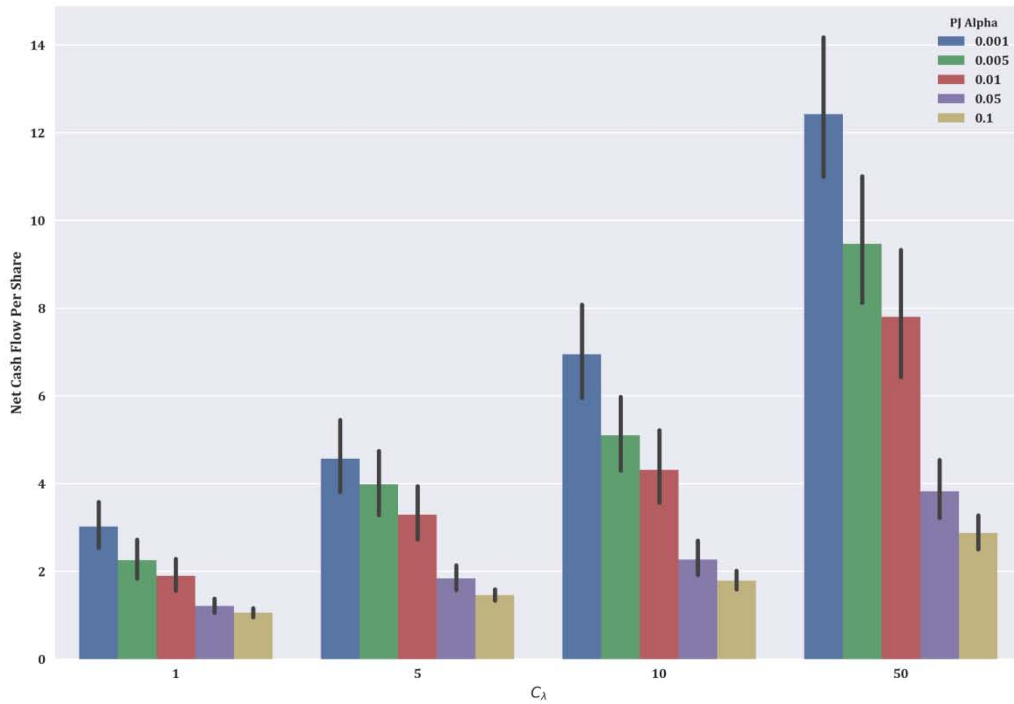

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=1, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 14

Change in Quoted Spreads Upon Increasing the MPI from 1 to 5

The change in quoted spread is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0 (the blue bar on the left) are determined by subtracting the average quoted spread for each of the 100 simulations in Group 1 from the average quoted spread for the paired simulation in Group 5. The difference between average quoted spreads observed for each group in the 5-tick regime and its paired group in the 1-tick regime is positive and decreases with increasing C_λ. The difference in average quoted spreads increases monotonically with increasing α_{PJ} when C_λ = 1, and exhibits a U-shape with increasing α_{PJ} when C_λ > 1.

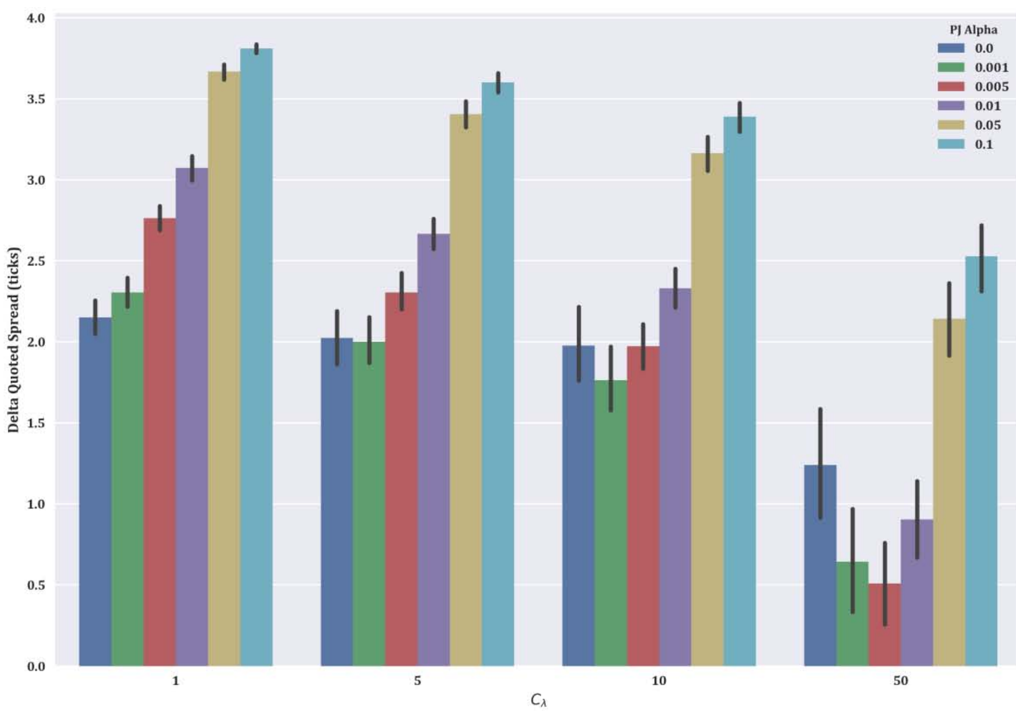

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 15

Impact of α_{PJ} on Quoted Spreads; MPI = 5

Quoted spreads increase as C_λ increases from 1 to 50 and decrease as α_{PJ} increases within each group. The impact of α_{PJ} on quoted spreads is most pronounced (in nominal ticks) when quoted spreads are large (Group 8, C_λ = 50).

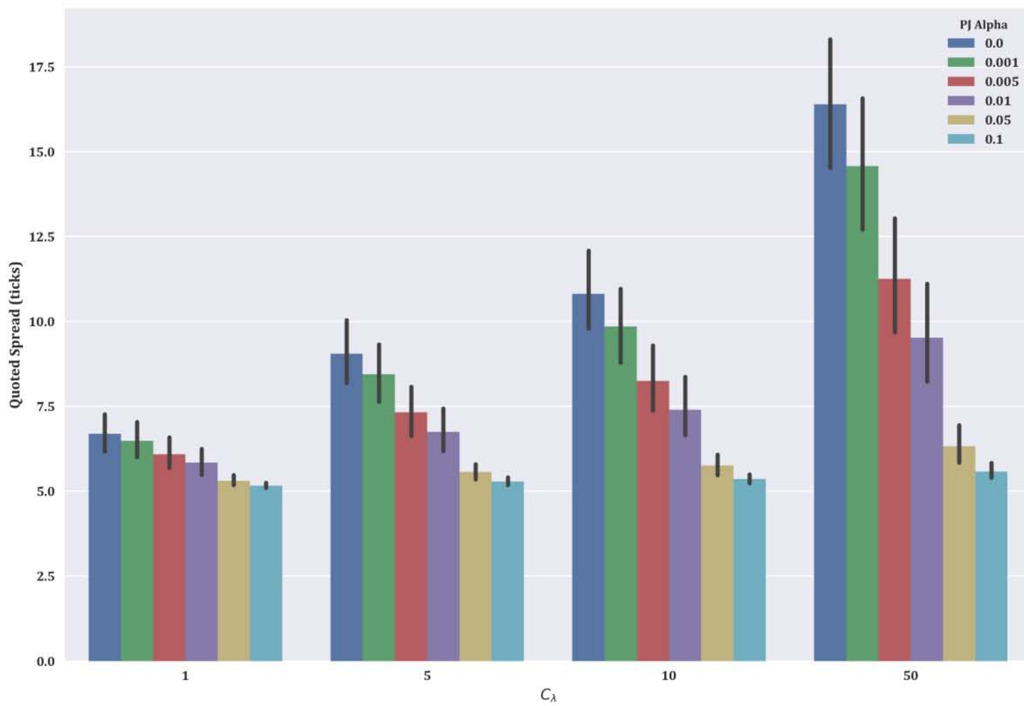

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 16
Change in Market Maker Participation Upon Increasing the MPI from 1 to 5
The change in market maker participation is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0 (the blue bar on the left) are determined by subtracting the market maker participation for each of the 100 simulations in Group 1 from the market maker participation for the paired simulation in Group 5. The difference between market maker participation observed for each group in the 5-tick regime and its paired group in the 1-tick regime is positive and decreases with increasing C_λ. Within each paired group (holding C_λ constant), the difference in market maker participation increases monotonically with increasing α_{PJ}.

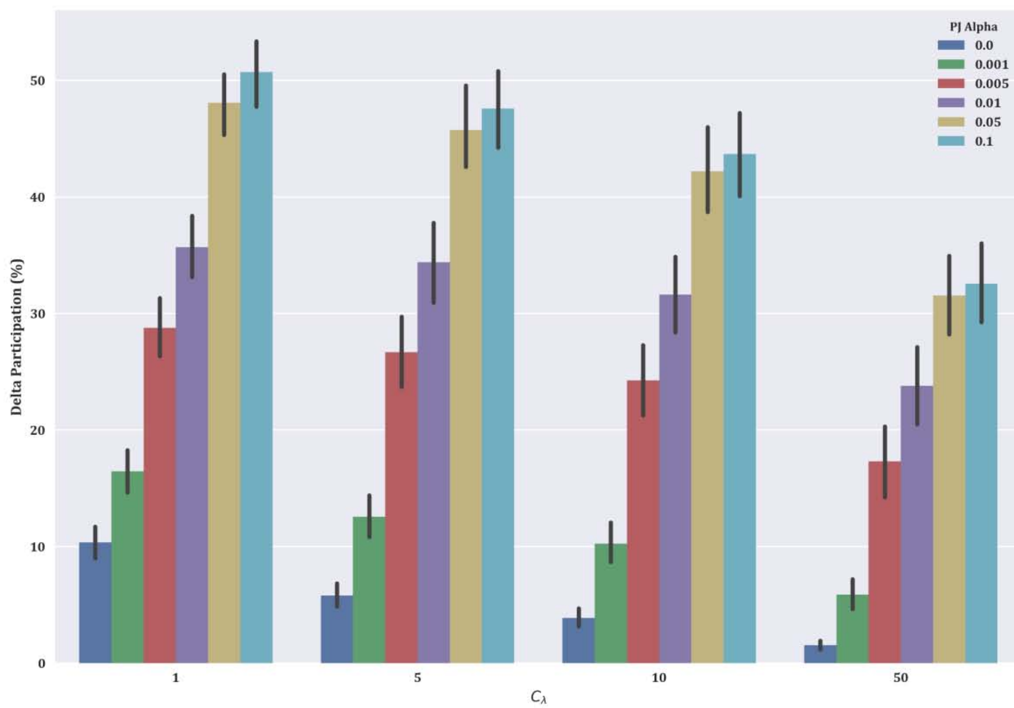

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 17
Impact of α_{PJ} on Market Maker Participation; MPI = 5
Market maker participation increases as C_λ increases from 1 to 50 for $\alpha_{PJ} \leq 0.001$, exhibits an inverted U-shape for $\alpha_{PJ} = 0.005$ and decreases with increasing C_λ for $\alpha_{PJ} \geq 0.01$. Within each group, market maker participation declines monotonically with increasing α_{PJ}.

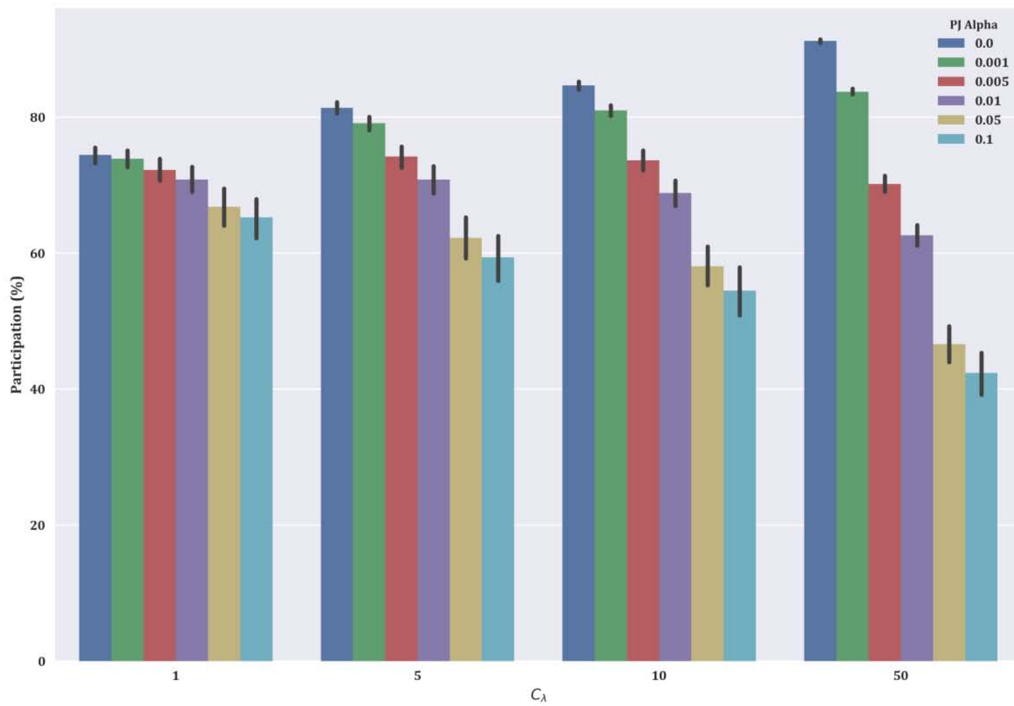

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 18

Change in Market Maker Cancel-To-Trade Ratio Upon Increasing the MPI from 1 to 5

The change in the market maker cancel-to-trade ratio is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0 (the blue bar on the left) are determined by subtracting the market maker cancel-to-trade ratio for each of the 100 simulations in Group 1 from the market maker cancel-to-trade ratio for the paired simulation in Group 5. The difference between the market maker cancel-to-trade ratios observed for each group in the 5-tick regime and its paired group in the 1-tick regime is negative. The difference increases (becomes less negative) with increasing C_λ for $\alpha_{PJ} \leq 0.005$ and decreases with increasing C_λ for $\alpha_{PJ} \geq 0.05$. Within each paired group (holding C_λ constant), the difference in market maker cancel-to-trade ratios decreases (becomes more negative) monotonically with increasing α_{PJ}.

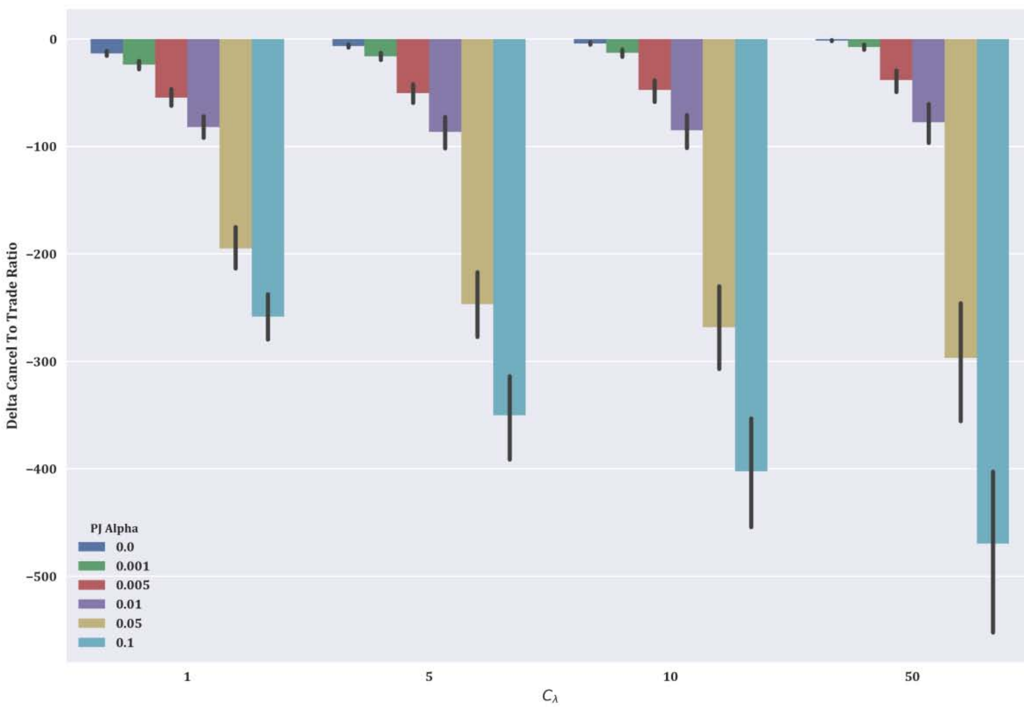

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 19

Impact of α_{PJ} on Market Maker Cancel-To-Trade Ratio; MPI = 5

The market maker cancel-to-trade ratio decreases as C_λ increases from 1 to 50 for $\alpha_{PJ} \leq 0.001$ and increases with increasing C_λ for $\alpha_{PJ} \geq 0.01$. Within each group, the market maker cancel-to-trade ratio increases with increasing α_{PJ}.

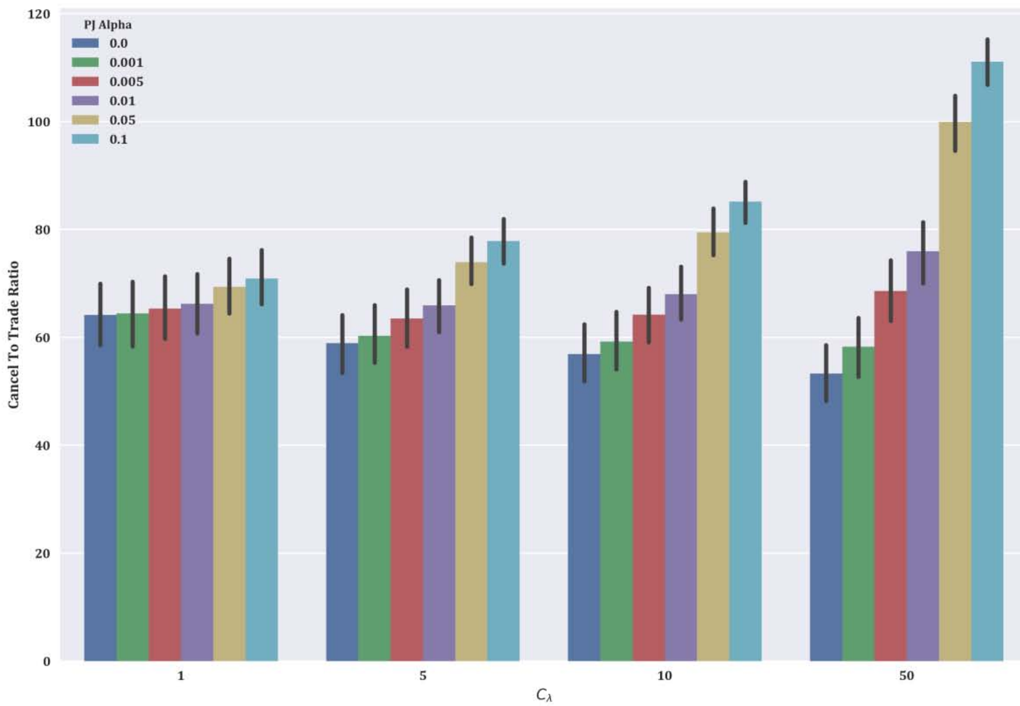

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 20

Change in Market Maker Total Net Cash Flow Upon Increasing the MPI from 1 to 5

The change in the market maker total net cash flow is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0 (the blue bar on the left) are determined by subtracting the market maker total net cash flow for each of the 100 simulations in Group 1 from the market maker total net cash flow for the paired simulation in Group 5. The difference between the market maker total net cash flow observed for each group in the 5-tick regime and its paired group in the 1-tick regime is positive.

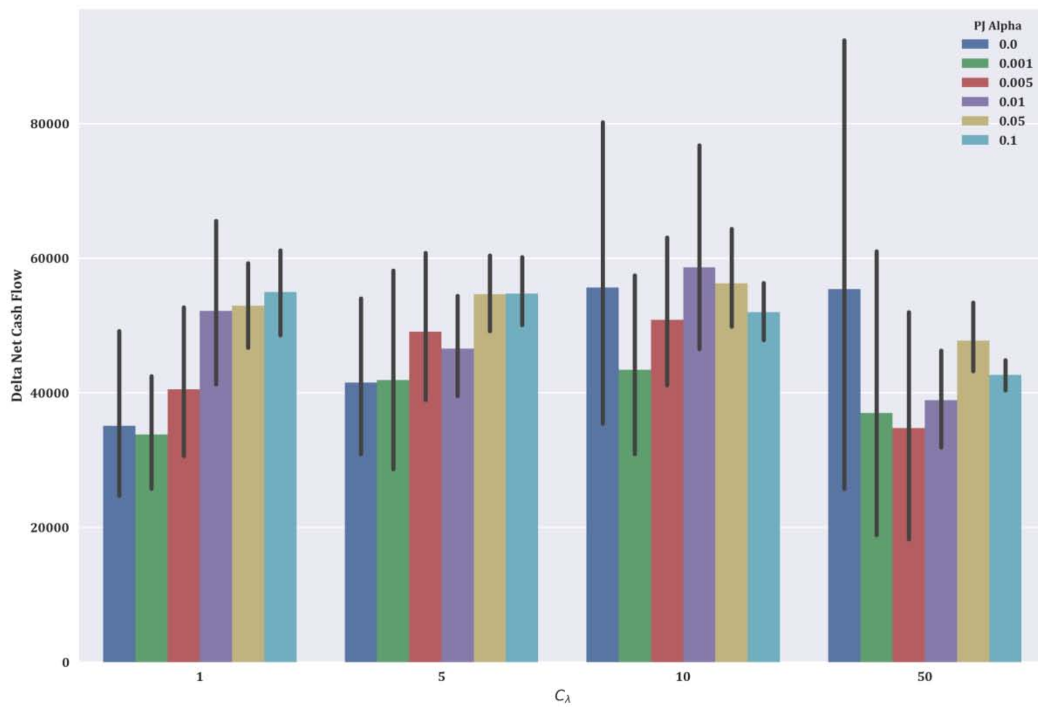

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 21
Impact of α_{PJ} on Market Maker Total Net Cash Flow; MPI = 5
Market maker total net cash flow increases as C_λ increases from 1 to 50 for $\alpha_{PJ} \leq 0.05$, and does not vary much for $\alpha_{PJ} = 0.1$. Within each group, market maker total net cash flow decreases with increasing α_{PJ}.

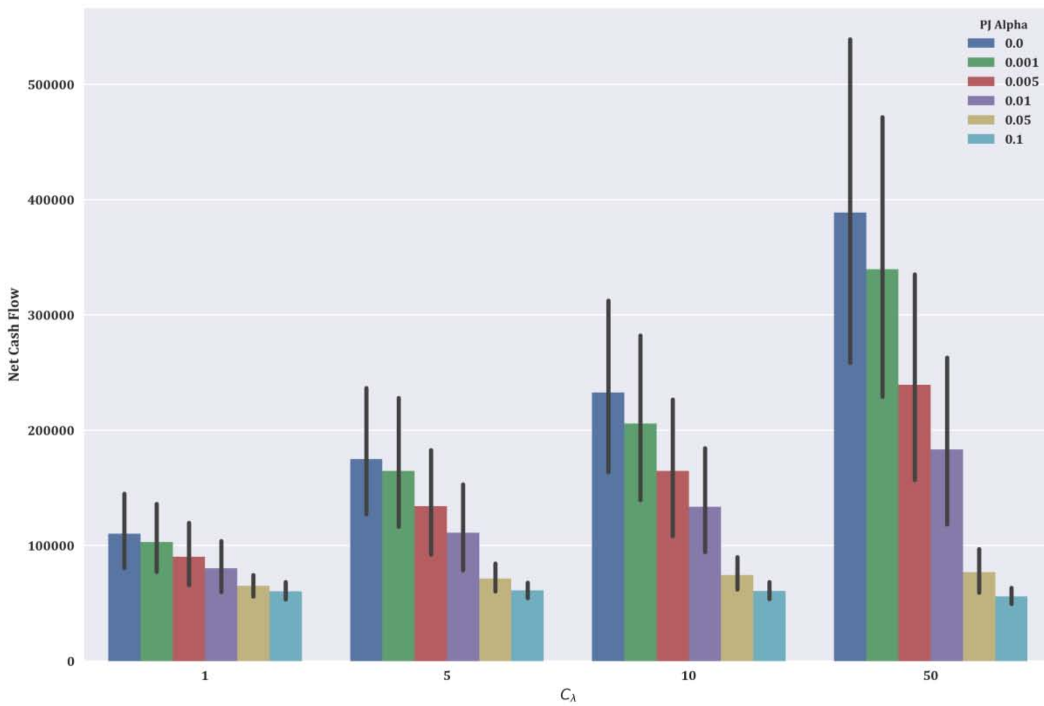

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 22

Change in Market Maker Net Cash Flow Per Share Upon Increasing the MPI from 1 to 5

The change in the market maker net cash flow per share is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0 (the blue bar on the left) are determined by subtracting the market maker net cash flow per share for each of the 100 simulations in Group 1 from the market maker net cash flow per share for the paired simulation in Group 5. The difference between the market maker net cash flow per share observed for each group in the 5-tick regime and its paired group in the 1-tick regime is positive.

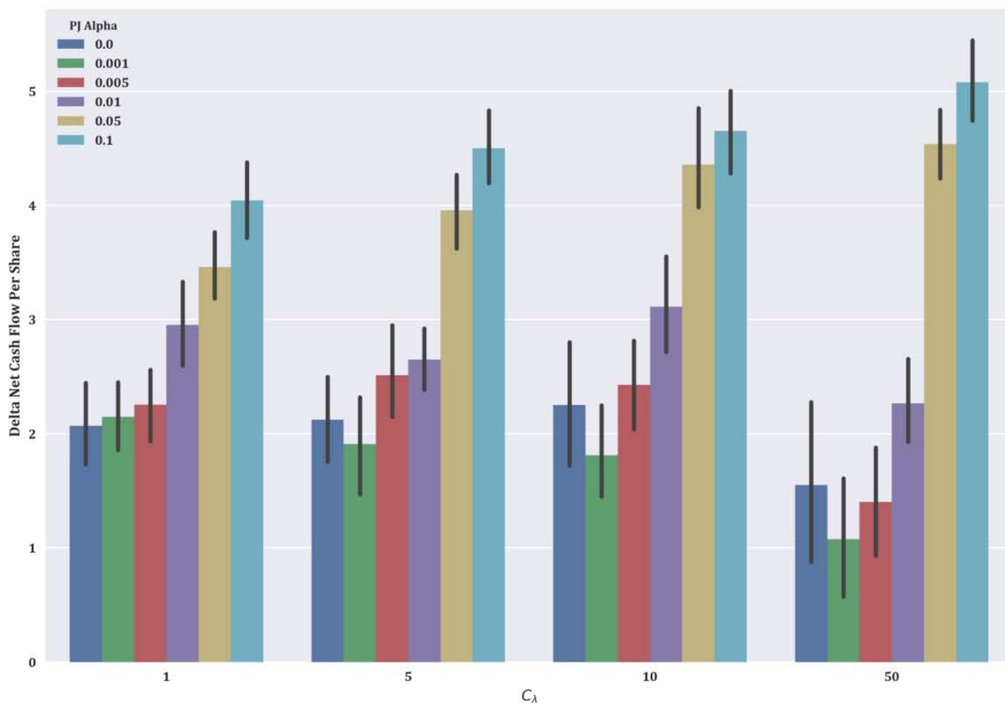

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 23
Impact of α_{PJ} on Market Maker Net Cash Flow Per Share; MPI = 5
Market maker net cash flow per share increases as C_λ increases from 1 to 50. Within each group, market maker net cash flow per share decreases with increasing α_{PJ}.

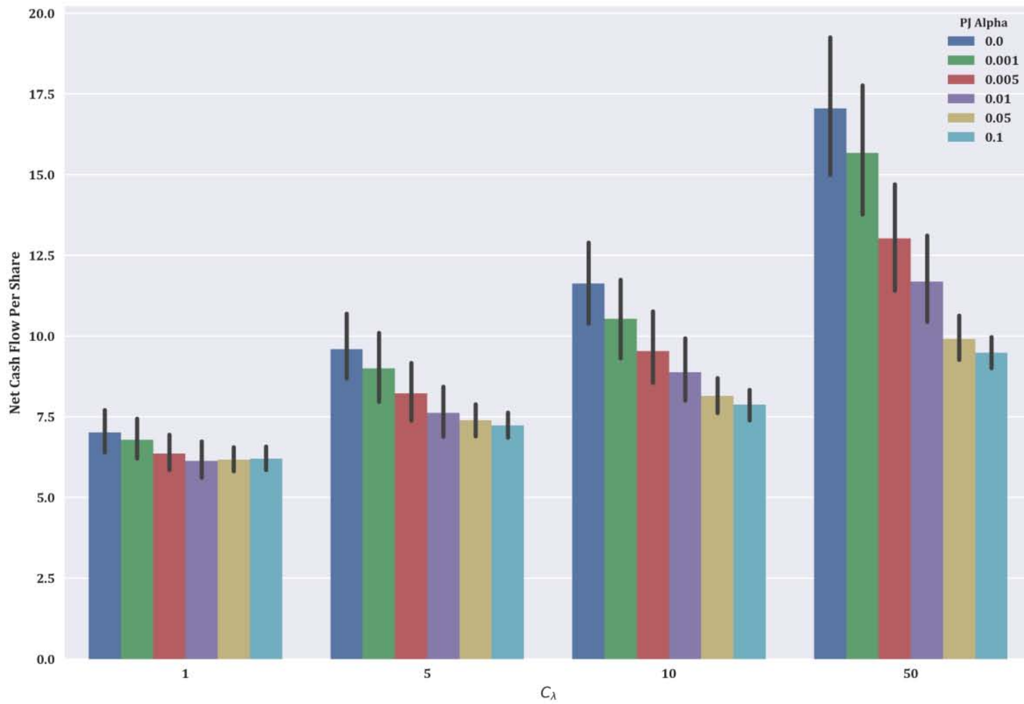

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 24

Change in Penny Jumper Participation Upon Increasing the MPI from 1 to 5

The change in penny jumper participation is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0.001 (the green bar on the left) are determined by subtracting the penny jumper participation for each of the 100 simulations in Group 1 from the penny jumper participation for the paired simulation in Group 5. The difference between the penny jumper participation observed for each group in the 5-tick regime and its paired group in the 1-tick regime is negative and increasing (less negative) for increasing C_λ. Within each paired group (holding C_λ constant), the difference in penny jumper participation decreases (becomes more negative) monotonically with increasing α_{PJ}.

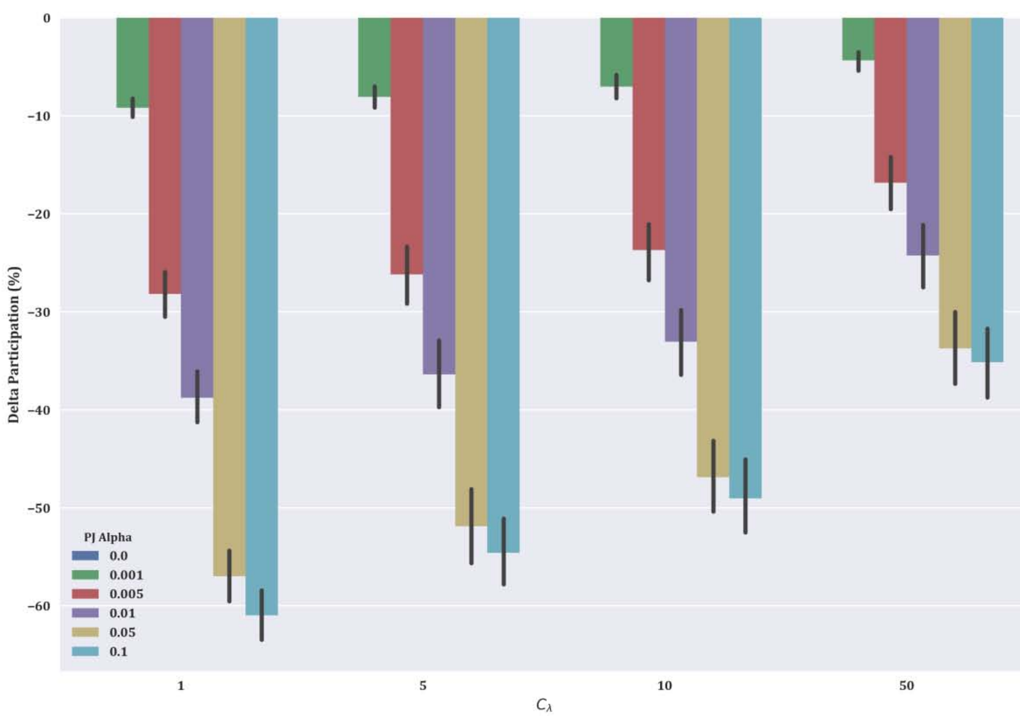

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 25

Impact of α_{PJ} on Penny Jumper Participation; MPI = 5

Penny jumper participation increases as C_λ increases from 1 to 50. Within each group, penny jumper participation increases monotonically with increasing α_{PJ}.

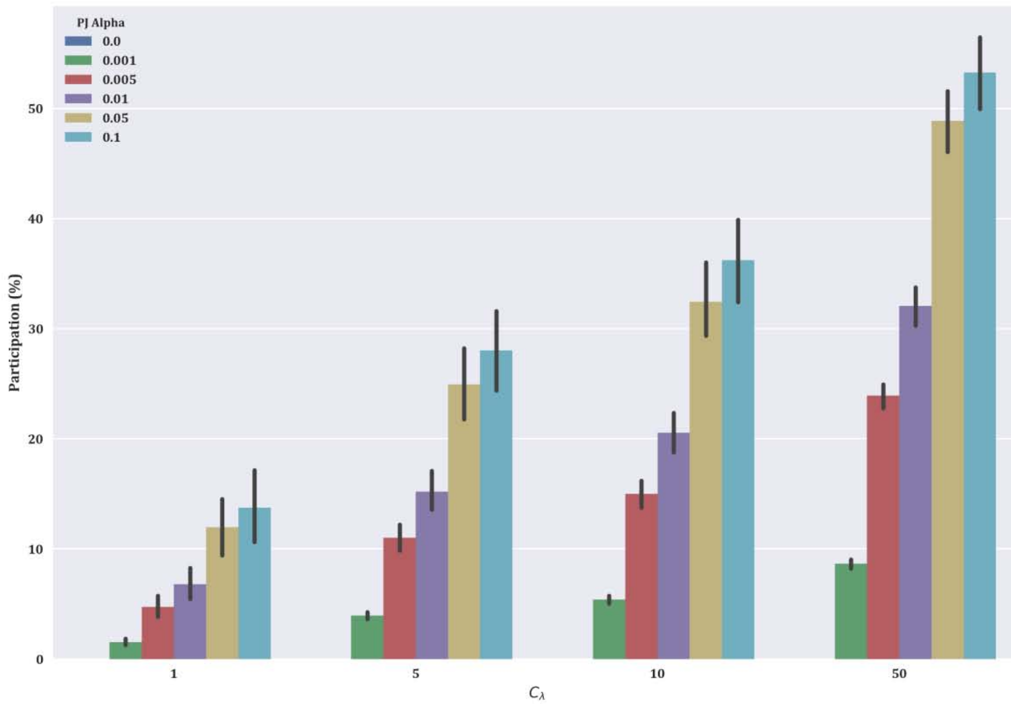

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 26

Change in Penny Jumper Cancel-To-Trade Ratio Upon Increasing the MPI from 1 to 5

The change in the penny jumper cancel-to-trade ratio is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0.001 (the blue bar on the left) are determined by subtracting the penny jumper cancel-to-trade ratio for each of the 100 simulations in Group 1 from the penny jumper cancel-to-trade ratio for the paired simulation in Group 5. The difference between the penny jumper cancel-to-trade ratios observed for each group in the 5-tick regime and its paired group in the 1-tick regime is positive except for α_{PJ} = 0.001 and C_λ > 1. The difference decreases with increasing C_λ. Within each paired group (holding C_λ constant), the difference in penny jumper cancel-to-trade ratios increases monotonically with increasing α_{PJ}.

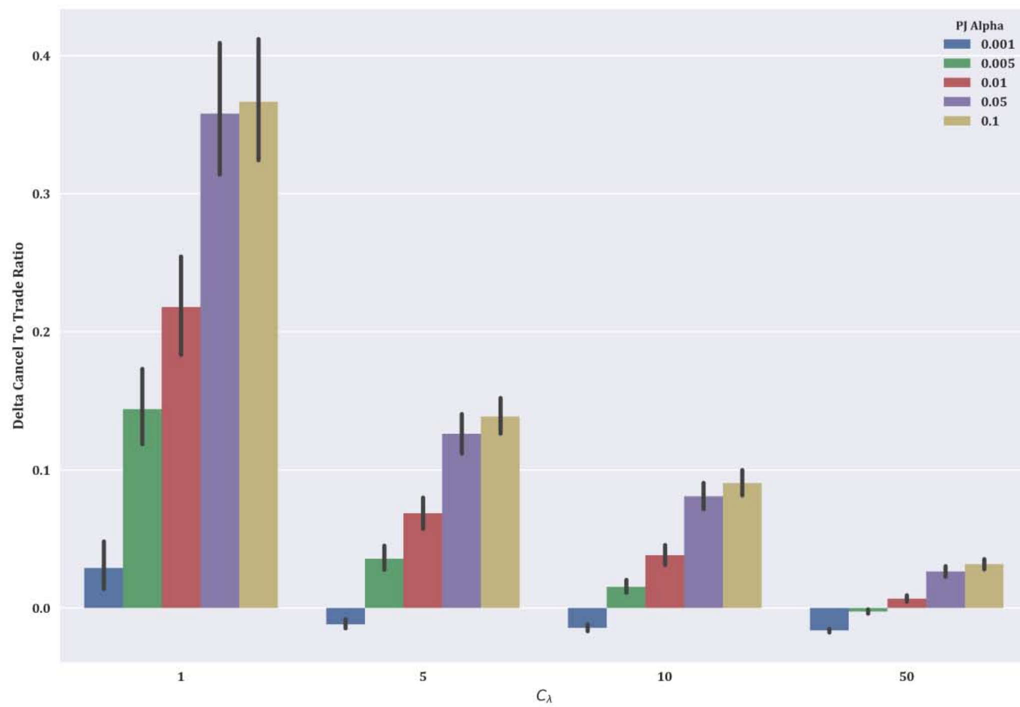

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 27
Impact of α_{PJ} on Penny Jumper Cancel-To-Trade Ratio; MPI = 5
The penny jumper cancel-to-trade ratio decreases as C_λ increases from 1 to 50. Within each group, the penny jumper cancel-to-trade ratio inceases with increasing α_{PJ}.

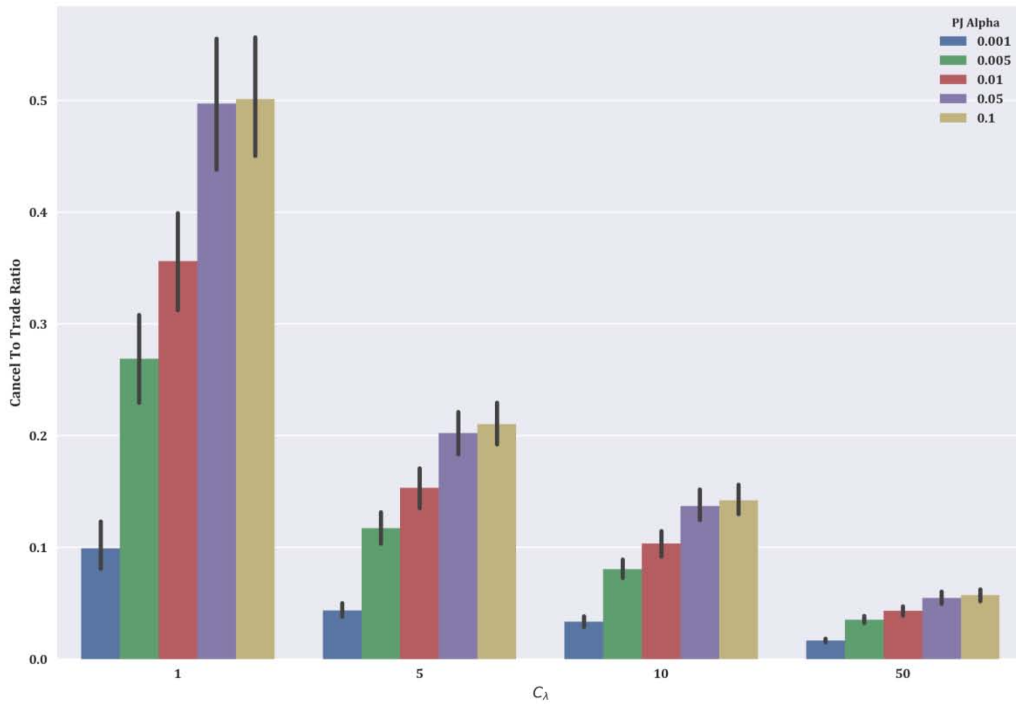

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 28
Change in Penny Jumper Total Net Cash Flow Upon Increasing the MPI from 1 to 5
The change in penny jumper total net cash flow is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0.001 (the green bar on the left) are determined by subtracting the penny jumper total net cash flow for each of the 100 simulations in Group 1 from the penny jumper total net cash flow for the paired simulation in Group 5. The difference between the penny jumper total net cash flow observed for each group in the 5-tick regime and its paired group in the 1-tick regime is negative and generally decreasing (more negative) for increasing C_λ. Within each paired group (holding C_λ constant), the difference in penny jumper total net cash flow exhibits a U-shape with increasing α_{PJ}.

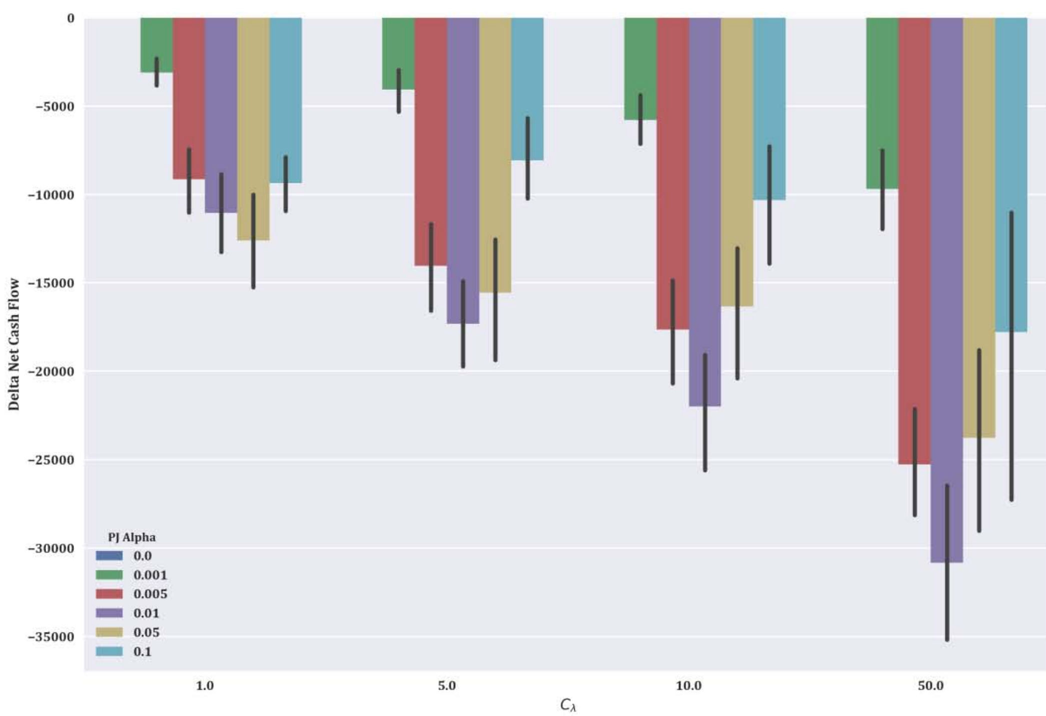

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 29

Impact of α_{PJ} on Penny Jumper Total Net Cash Flow; MPI = 5

Penny jumper total net cash flow increases as C_λ increases from 1 to 50. For $C_\lambda > 1$, penny jumper total net cash flow displays an inverted U-shape (i.e., first increases, then decreases) with increasing α_{PJ}. For $C_\lambda = 1$, penny jumper total net cash flow increases with increasing α_{PJ}.

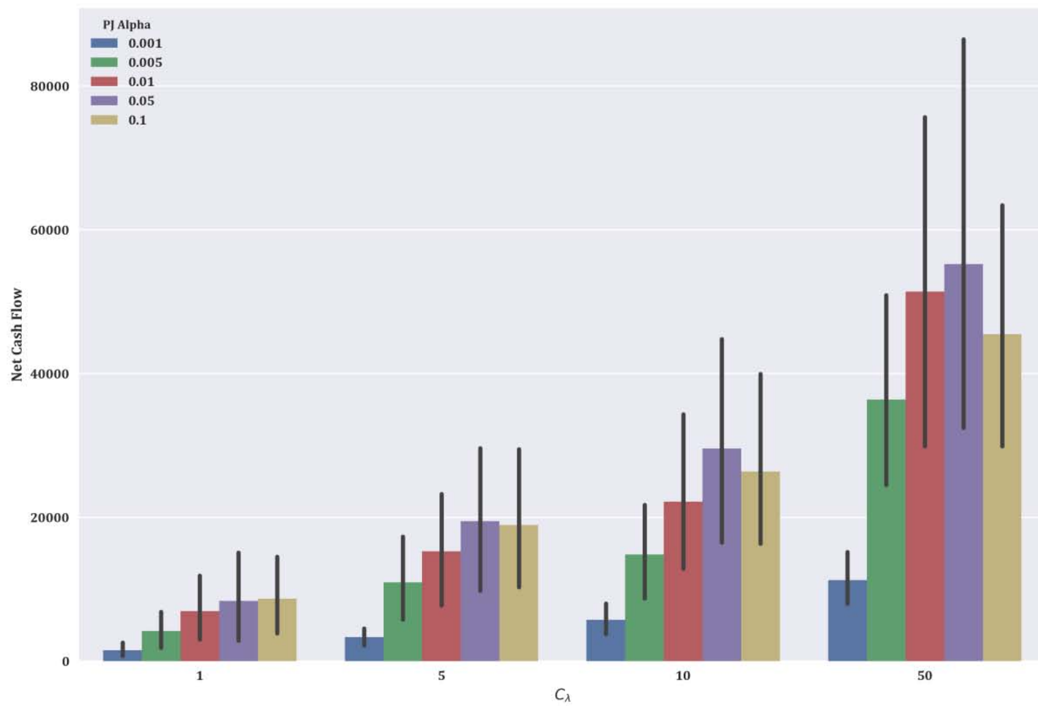

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 30

Change in on Penny Jumper Net Cash Flow Per Share Upon Increasing the MPI from 1 to 5

The change in penny jumper net cash flow per share is measured by subtracting the value obtained from the 1-tick MPI from the value obtained from the 5-tick MPI for each paired simulation (i.e., holding α_{PJ} constant). For example, the results for C_λ = 1and α_{PJ} = 0.001 (the green bar on the left) are determined by subtracting the penny jumper net cash flow per share for each of the 100 simulations in Group 1 from the penny jumper net cash flow per share for the paired simulation in Group 5. The difference between the penny jumper net cash flow per share observed for each group in the 5-tick regime and its paired group in the 1-tick regime is negative.

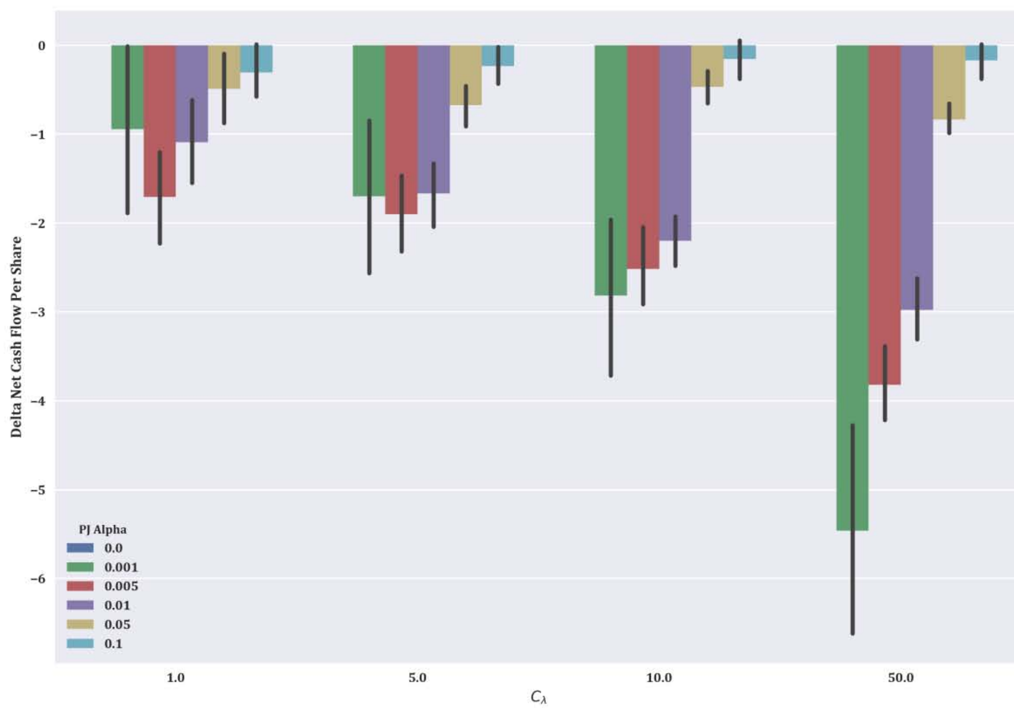

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=[1,5], α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 31

Impact of α_{PJ} on Penny Jumper Net Cash Flow Per Share; MPI = 5

Penny jumper net cash flow per share increases as C_λ increases from 1 to 50. For $C_\lambda > 1$, penny jumper net cash flow per share generally decreases with increasing α_{PJ}.

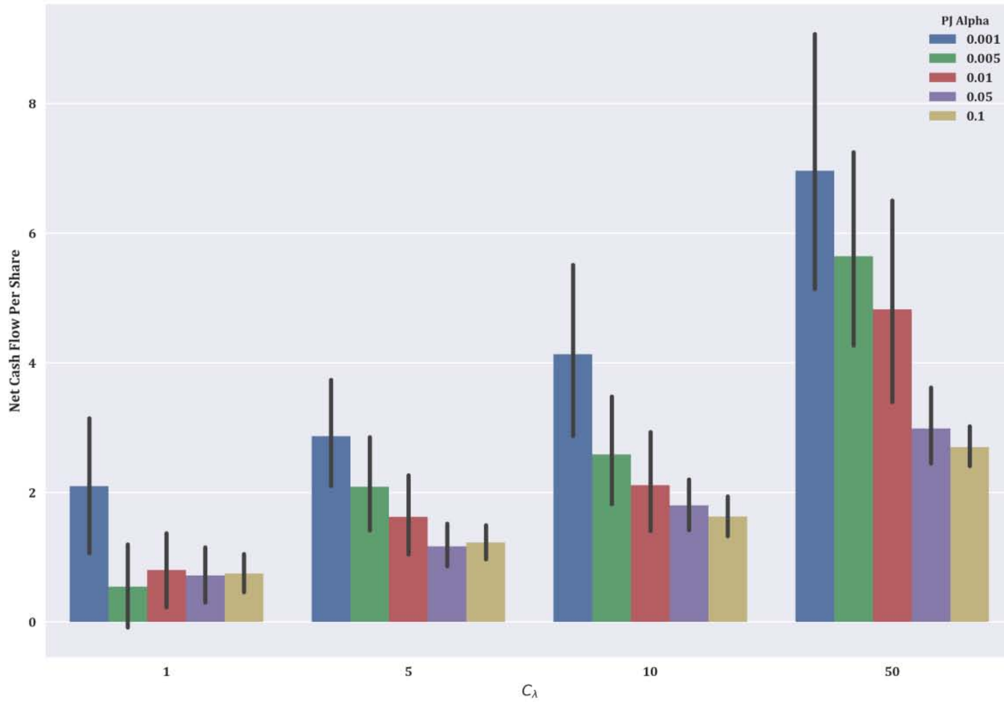

Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=[0, 0.001, 0.005, 0.01, 0.05, 0.1].

Figure 32
Spread Comparison Results; MPI=5
Average quoted spread distributions from Groups 5 through 8 (C_λ = [1, 5, 10, 50] , MPI = 5) with no penny jumper (α_{PJ} =0) match the overall shape and range of the average quoted spreads observed in the Tick Size Pilot Test Group stocks after the implementation of the pilot (December, 2016 to February, 2017).

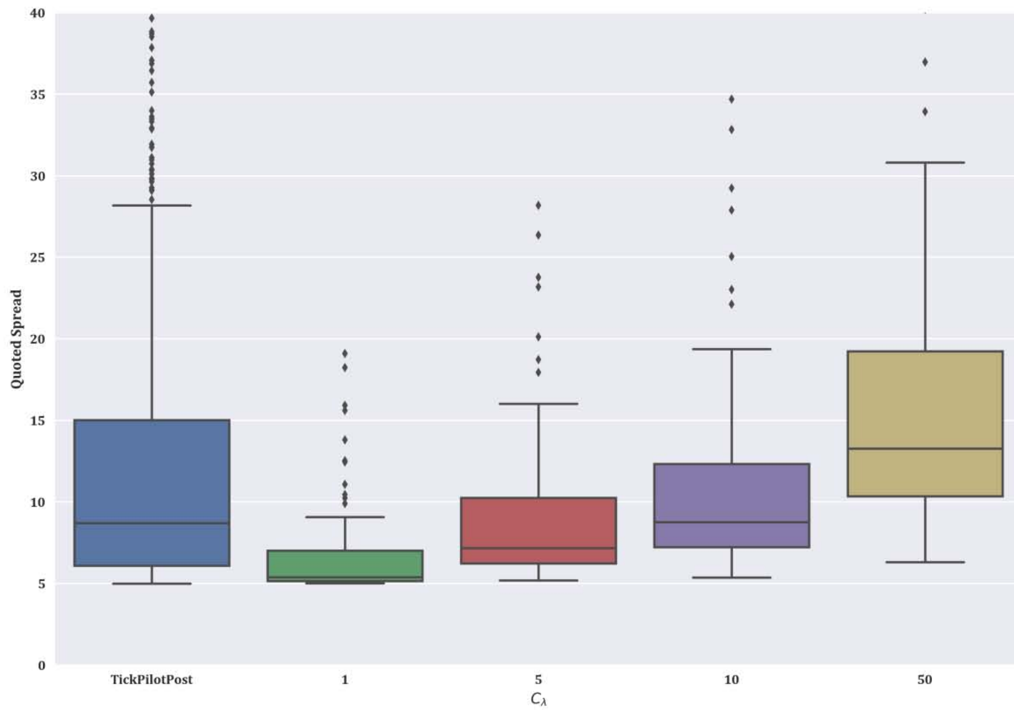

Tick Pilot Data: distribution of daily duration-weighted average quoted spread in cents; 1,183 symbols from December, 2016 to February, 2017.
Simulations: 100 simulations, N_{LP}=38, N_{MM}=1, N_{LT}=50, Q_{max}=1, T=100,000, α=0.0375, μ=0.001, δ=0.025, $q_{provide}$=0.5, λ_0=100, C_λ=[1, 5, 10, 50], Δs=0.001, MPI=5, α_{PJ}=0.