Required fields are shown with yellow backgrounds and asterisks.

SECURITIES AND EXCHANGE COMMISSION
WASHINGTON, D.C. 20549
Form 19b-4

File No. * SR 2022 - * 079
Amendment No. (req. for Amendments *)

Filing by The Nasdaq Stock Market LLC

Pursuant to Rule 19b-4 under the Securities Exchange Act of 1934

Initial *	Amendment *	Withdrawal	Section 19(b)(2) *	Section 19(b)(3)(A) *	Section 19(b)(3)(B) *
✓	☐	☐	✓	☐	☐

Pilot	Extension of Time Period for Commission Action *	Date Expires *
☐	☐	

Rule

☐ 19b-4(f)(1)	☐ 19b-4(f)(4)
☐ 19b-4(f)(2)	☐ 19b-4(f)(5)
☐ 19b-4(f)(3)	☐ 19b-4(f)(6)

Notice of proposed change pursuant to the Payment, Clearing, and Settlement Act of 2010

Section 806(e)(1) *	Section 806(e)(2) *
☐	☐

Security-Based Swap Submission pursuant to the Securities Exchange Act of 1934
Section 3C(b)(2) *
☐

Exhibit 2 Sent As Paper Document	Exhibit 3 Sent As Paper Document
☐	☐

Description

Provide a brief description of the action (limit 250 characters, required when Initial is checked *).

A proposal to amend Rule 4702 of the Exchanges Rulebook

Contact Information

Provide the name, telephone number, and e-mail address of the person on the staff of the self-regulatory organization prepared to respond to questions and comments on the action.

First Name * Brett
Last Name * Kitt
Title * AVP Principal Associate General Counsel
E-mail * brett.kitt@nasdaq.com
Telephone * (240) 459-2221
Fax

Signature

Pursuant to the requirements of the Securities Exchange of 1934, The Nasdaq Stock Market LLC has duty caused this filing to be signed on its behalf by the undersigned thereunto duty authorized.

Date 12/21/2022

By John Zecca (Name *)

(Title *) EVP and Chief Legal Officer

NOTE: Clicking the signature block at right will initiate digitally signing the form. A digital signature is as legally binding as a physical signature, and once signed, this form cannot be changed.

Date: 2022.12.21 15:02:05 -05'00'

Required fields are shown with yellow backgrounds and astericks.

SECURITIES AND EXCHANGE COMMISSION
WASHINGTON, D.C. 20549

For complete Form 19b-4 instructions please refer to the EFFS website.

Form 19b-4 Information * Add Remove View SR-NASDAQ-2022-079 19b-4.doc	The self-regulatory organization must provide all required information, presented in a clear and comprehensible manner, to enable the public to provide meaningful comment on the proposal and for the Commission to determine whether the proposal is consistent with the Act and applicable rules and regulations under the Act.
Exhibit 1 - Notice of Proposed Rule Change * Add Remove View SR-NASDAQ-2022-079 Exhibit 1.doc	The Notice section of this Form 19b-4 must comply with the guidelines for publication in the Federal Register as well as any requirements for electronic filing as published by the Commission (if applicable). The Office of the Federal Register (OFR) offers guidance on Federal Register publication requirements in the Federal Register Document Drafting Handbook, October 1998 Revision. For example, all references to the federal securities laws must include the corresponding cite to the United States Code in a footnote. All references to SEC rules must include the corresponding cite to the Code of Federal Regulations in a footnote. All references to Securities Exchange Act Releases must include the release number, release date, Federal Register cite, Federal Register date, and corresponding file number (e.g., SR-[SRO]-xx-xx). A material failure to comply with these guidelines will result in the proposed rule change being deemed not properly filed. See also Rule 0-3 under the Act (17 CFR 240.0-3)
Exhibit 1A - Notice of Proposed Rule Change, Security-Based Swap Submission, or Advanced Notice by Clearing Agencies * Add Remove View	The Notice section of this Form 19b-4 must comply with the guidelines for publication in the Federal Register as well as any requirements for electronic filing as published by the Commission (if applicable). The Office of the Federal Register (OFR) offers guidance on Federal Register publication requirements in the Federal Register Document Drafting Handbook, October 1998 Revision. For example, all references to the federal securities laws must include the corresponding cite to the United States Code in a footnote. All references to SEC rules must include the corresponding cite to the Code of Federal Regulations in a footnote. All references to Securities Exchange Act Releases must include the release number, release date, Federal Register cite, Federal Register date, and corresponding file number (e.g., SR-[SRO]-xx-xx). A material failure to comply with these guidelines will result in the proposed rule change being deemed not properly filed. See also Rule 0-3 under the Act (17 CFR 240.0-3)
Exhibit 2- Notices, Written Comments, Transcripts, Other Communications Add Remove View	Copies of notices, written comments, transcripts, other communications. If such documents cannot be filed electronically in accordance with Instruction F, they shall be filed in accordance with Instruction G. ☐ Exhibit Sent As Paper Document
Exhibit 3 - Form, Report, or Questionnaire Add Remove View SR-NASDAQ-2022-079 Exhibit 3.pdf	Copies of any form, report, or questionnaire that the self-regulatory organization proposes to use to help implement or operate the proposed rule change, or that is referred to by the proposed rule change. ☐ Exhibit Sent As Paper Document
Exhibit 4 - Marked Copies Add Remove View	The full text shall be marked, in any convenient manner, to indicate additions to and deletions from the immediately preceding filing. The purpose of Exhibit 4 is to permit the staff to identify immediately the changes made from the text of the rule with which it has been working.
Exhibit 5 - Proposed Rule Text Add Remove View SR-NASDAQ-2022-079 Exhibit 5.doc	The self-regulatory organization may choose to attach as Exhibit 5 proposed changes to rule text in place of providing it in Item I and which may otherwise be more easily readable if provided separately from Form 19b-4. Exhibit 5 shall be considered part of the proposed rule change
Partial Amendment Add Remove View	If the self-regulatory organization is amending only part of the text of a lengthy proposed rule change, it may, with the Commission's permission, file only those portions of the text of the proposed rule change in which changes are being made if the filing (i.e. partial amendment) is clearly understandable on its face. Such partial amendment shall be clearly identified and marked to show deletions and additions.

1. Text of the Proposed Rule Change

(a) The Nasdaq Stock Market LLC ("Nasdaq" or "Exchange"), pursuant to Section 19(b)(1) of the Securities Exchange Act of 1934 ("Act")[1] and Rule 19b-4 thereunder,[2] is filing with the Securities and Exchange Commission ("SEC" or "Commission") a proposal to amend Rules 4702(b)(14) and (b)(15) of the Exchange's Rulebook to replace the static holding period requirements for Midpoint Extended Life Orders and Midpoint Extended Life Orders Plus Continuous Book with dynamic holding periods.

A notice of the proposed rule change for publication in the Federal Register is attached as Exhibit 1.

The text of the proposed rule change is attached as Exhibit 5.

(b) Not applicable.

(c) Not applicable.

2. Procedures of the Self-Regulatory Organization

The proposed rule change was approved by the Board of Directors of the Exchange (the "Board") on October 26, 2022. No other action is necessary for the filing of the rule change.

Questions and comments on the proposed rule change may be directed to:

Brett Kitt
AVP and Principal Associate General Counsel
Nasdaq, Inc.
(301) 978-8132

---

1 15 U.S.C. 78s(b)(1).

2 17 CFR 240.19b-4.

3. Self-Regulatory Organization's Statement of the Purpose of, and Statutory Basis for, the Proposed Rule Change

a. Purpose

The Exchange proposes to amend Rules 4702(b)(14) and (15) of the Exchange's Rulebook to replace the static 10 millisecond holding period requirements for its Midpoint Extended Life Order ("M-ELO") and Midpoint Extended Life Order Plus Continuous Book ("M-ELO+CB") Order Types with dynamic holding periods ("Dynamic M-ELO and M-ELO+CB" or collectively, "Dynamic M-ELO").

Background

In 2018, the Exchange introduced the M-ELO, which is a Non-Displayed Order priced at the Midpoint between the National Best Bid and Offer ("NBBO") and which is eligible for execution only against other eligible M-ELOs and only after a minimum of one-half second passes from the time that the System accepts the order (the "Holding Period").[3] In 2019, the Exchange introduced the M-ELO+CB, which closely resembles the M-ELO, except that a M-ELO+CB may execute at the midpoint of the NBBO, not only against other eligible M-ELOs (and M-ELO+CBs), but also against Non-Displayed Orders with Midpoint Pegging and Midpoint Peg Post-Only Orders ("Midpoint Orders") that rest on the Continuous Book for at least one-half second and have Trade Now enabled.[4]

---

3 See Securities Exchange Act Release No. 34-82825 (March 7, 2018), 83 FR 10937 (March 13, 2018) (SR-NASDAQ-2017-074) ("M-ELO Approval Order").

4 See Securities Exchange Act Release No. 34-86938 (September 11, 2019), 84 FR 48978 (September 17, 2019) (SR-NASDAQ-2019-048) ("M-ELO+CB Approval Order").

When the Exchange designed M-ELO, it originally set the length of the Holding Period at one-half second because it determined that this time period would be sufficient to ensure that likeminded investors would interact only with each other, and with minimal market impacts. The Exchange believed that the longer length of the M-ELO Holding Period and its simplicity in design would provide greater protection for participants than they could achieve through competing delay mechanisms.

In 2020, however, the Exchange shortened the length of the Holding Period to 10 milliseconds.[5] The Exchange did so after studying two years of actual use and performance of M-ELOs, as well as customer feedback. That is, the Exchange came to understand that, while users of M-ELO and M-ELO+CB are less concerned with achieving rapid executions of their Orders than are other participants, they are not indifferent about the length of time in which their M-ELOs and M-ELO+CBs must wait before they are eligible for execution. Indeed, participants informed the Exchange that in certain circumstances, such as when they sought to trade symbols that on average had a lower time-to-execution than a half-second, they were reticent to enter M-ELOs or M-ELO+CBs. They indicated that the associated Holding Periods for these Order Types were longer than necessary to achieve the desired protections and that, during the residual portion of the Holding Periods, they risked losing out on favorable execution opportunities that would otherwise be available to them had they placed a non-MELO order.

5 See Securities Exchange Act Release No. 34-88743 (April 24, 2020), 85 FR 24068 (April 30, 2020) (SR-NASDAQ-2020-011) ("M-ELO Timer Approval Order").

Based upon this feedback, the Exchange studied the potential effects of reducing the length of the Holding Periods for both M-ELOs and M-ELO+CBs (as well as for Midpoint Orders that would execute against M-ELO+CBs). Ultimately, the Exchange determined that it could reduce the Holding Periods to 10 milliseconds without compromising the protective power that M-ELO and M-ELO+CB are intended to provide to participants and investors.[6] Thus, the Exchange determined that shortening the Holding Periods to 10 milliseconds for M-ELOs and M-ELO+CBs would increase the efficacy of the mechanism while not undermining the power of those Order Types to fulfill their underlying purpose of minimizing market impacts. At the same time, the Exchange determined that a reduction in the Holding Periods to 10 milliseconds would dramatically add to the circumstances in which M-ELOs and M-ELO+CBs would be useful to participants. In its M-ELO Timer Approval Order, the Commission agreed with the Exchange:

> The Commission notes that, with the proposed ten-millisecond Holding Period and Resting Period, M-ELOs and M-ELO+CBs would continue to be optional

6 The Exchange examined each of its historical M-ELO executions to determine at what Midpoints of the NBBO the M-ELOs would have executed if their Holding Periods had been shorter than one-half second (500 milliseconds). After examining the historical effects of shorter Holding Periods of between 10 milliseconds and 400 milliseconds, the Exchange determined that a reduction of the M-ELO Holding Period to as short as 10 milliseconds would have caused an average impact on markouts of only 0.10 basis points (across all symbols). In other words, compared to the execution price of an average M-ELO with a one-half second Holding Period, the Exchange found that a M-ELO with a 10 millisecond Holding Period would have had an average post-execution impact that was only a tenth of a basis point per share – a difference in protective effect that is immaterial. See Nasdaq, "The Midpoint Extended Life Order (M-ELO); M-ELO Holding Period," available at https://www.nasdaq.com/articles/the-midpoint-extended-life-order-m-elo%3A-m-elo-holding-period-2020-02-13 (analyzing effects of shortened Holding Periods on M-ELO performance).

order types that are available to investors with longer investment time horizons, including institutional investors. The Commission also believes that the proposal could make M-ELOs and M-ELO+CBs more attractive for securities that on average have a time-to-execution of less than one-half second and, for investors who currently do not use M-ELOs and M-ELO+CBs for these securities, provide optional order types that could enhance their ability to participate effectively on the Exchange. The Commission notes that, if market participants determine that the proposal would make M-ELOs and M-ELO+CBs less attractive for their particular investment objectives, such market participants may elect to reduce or eliminate their use of these optional order types. Moreover, as noted above, the Exchange will continue to conduct real-time surveillance to monitor the use of M-ELOs and M-ELO+CBs to ensure that such usage remains appropriately tied to the intent of the order types. If, as a result of such surveillance, the Exchange determines that the shortened Holding Period does not serve its intended purpose or adversely impacts market quality, the Exchange would seek to make further recalibrations.[7]

For similar reasons and with even better potential results for participants, the Exchange now proposes to further refine the length of the Holding Periods for M-ELOs and M-ELO+CBs, this time through the application of innovative and patent pending machine learning technology.

Dynamic M-ELO

---

7 M-ELO Timer Approval Order, supra, at 85 FR 24069.

After receiving feedback from participants that even 10 millisecond Holding Periods for M-ELO and M-ELO+CB may, at times, exceed what is necessary to accomplish the underlying intent of these Order Types, the Exchange began to experiment with making further refinements to the duration of the Holding Periods. Ultimately, the Exchange concluded that shorter Holding Periods could achieve the same, if not better results for participants in terms of mark-outs, but not in all circumstances. That is, where prices of the underlying securities are stable, and not subject to imminent unfavorable changes, M-ELOs and M-ELO+CBs face lower risks of confronting spread-crossing orders, such that shorter Holding Periods could suffice to protect M-ELOs and M-ELO+CB from such orders. In periods of heightened price volatility, however, M-ELOs and M-ELO+CBs also face heightened risks, such that longer Holding Periods would continue to be beneficial in protecting M-ELOs and M-ELO+CBs from such risks. Thus, the Exchange determined that another across-the-board reduction of the static 10 millisecond Holding Periods would be sub-optimal because it could impact the performance of the M-ELO and M-ELO+CB Order Types during periods of heightened volatility.

In light of these observations, the Exchange tasked its artificial intelligence and machine learning laboratory (the "AI Core Development Group") to explore whether it could employ these innovative technologies to optimize the length of M-ELO and M-ELO+CB Holding Periods during various states of price volatility, and then to vary the lengths of the Holding Periods dynamically during the lifecycles of M-ELOs and M-ELO+CBs, with the objectives of improving the performance of these Order Types while also further reducing opportunity costs.

As the Exchange explains in greater depth in the attached white paper,[8] the AI Core Development Group proceeded to develop an artificial intelligence-based timer control system that will achieve these objectives.[9] The AI Core Development Group did so by using reinforcement learning techniques – machine learning paradigms which develop optimal solutions to problems over time by taking actions to solve them, generating feedback on the results of such actions, applying that feedback to direct and improve the next round of solutions, and then repeating the feedback loop until the paradigm achieves optimized solutions.

In this instance, the AI Core Development Group applied reinforcement learning techniques to a simulation of the M-ELO Book that it constructed using a representative data set from the first quarter of 2022 (the "Training Period"). The Training Period data consisted of 380 out of the 6,257 symbols on the M-ELO Book (accounting for approximately 67 percent of M-ELO volume). The symbols chosen reflect both actively-traded and thinly-traded securities, and both low-priced and high-priced securities.

---

8 See Diana Kafkes et al., "Applying Artificial Intelligence & Reinforcement Learning Methods Towards Improving Execution Outcomes," SSRN, October 19, 2022, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4243985 (attached hereto as Exhibit 3) (the "White Paper").

9 Although the AI Core Development Group acknowledges that an optimal Holding Period would update with every incoming order, it determined that training a reinforcement learning model on every order would be too difficult to program and too difficult to implement given the nanosecond latency requirements of the Exchange. The Group then investigated more feasible update cadences and determined the point at which optimal outcomes were best balanced with the level of programming and implementation difficulty to be between 15 and 30 second updates. Ultimately, the Group chose a 30 second update cadence to give the model the greatest opportunity to learn between potential actions.

The AI Core Development Group then developed a machine learning model with more than 140 features[10] and applied it to the Training Period data. The Group programmed the model to value the achievement of higher fill rates or lower mark-outs than that which occurred in a historical simulation of M-ELOs and M-ELO+CBs involving the Training Period data.[11] The Group then programmed the model to seek to achieve its goals by taking one of five possible actions with respect to the duration of the Holding Periods at 30 second intervals[12] for each symbol during each trading day of the Training Period. That is, at each 30 second internal, the model evaluated market conditions for each symbol over the prior 30 second period and either kept the Holding Periods the same, increased/decreased them by 0.25 milliseconds, or increased/decreased them by 0.50 milliseconds.[13] After each decision-making round, the model utilized the results to inform its actions at the next 30 second increment.

In making its decisions, the model considered 142 categories of data points. A confluence of data points that correlated with an increase in volatility tended to cause the model to increase the durations of Holding Periods, including increases in the standard deviation of NBBO prices, the number of unique participants placing sell orders on M-

---

10 See White Paper, supra, at 31, for a description of these features.

11 As the White Paper explains, the Group developed a model to simulate activity on the Exchange involving M-ELOs and M-ELO+CBs during the Training Period. See White Paper, supra, at 10.

12 See id.

13 The AI Core Development Group experimented with a range of permissible Holding Period durations. Ultimately, it concluded that it could produce better outcomes for M-ELO and M-ELO+CB participants than the existing approach using Holding Periods as low as 0.25 milliseconds and as high as 2.5 milliseconds, under normal market conditions.

ELO and M-ELO+CB, and the volume-weighted average of the NBBO spread. Conversely, a confluence of data points that correlated with greater price stability tended to cause the model to decrease the durations of Holding periods, such as an increase in the median and max number of shares per trade and the number of resting bids left in the M-ELO and M-ELO+CB Book.

The AI Core Development Team produced variations of its model that prioritized achievement of the lowest mark-outs, the highest fill rates, and a blend of these two objectives. [14] Through a process of learning and experimentation, the AI Core Development Group settled on a Dynamic M-ELO model that achieved substantial performance improvements for users of M-ELO and M-ELO+CB – both in terms of markouts and fill rates – as compared to the static 10 millisecond Holding Periods. As the White Paper explains in greater detail, Dynamic M-ELO yielded an average combined volume-weighted improvement of 31.7 percent, including a 20.3 percent increase in fill rates and a 11.4 percent reduction in mark-outs. [15] The White Paper provides a more fulsome explanation of these improvements.[16]

---

14 The AI Core Development Group also applied to the model a paradigm called "retraining" to combat the degradation of model performance that can otherwise occur as the reference data it uses for initial comparison becomes stale. Finally, the AI Core Development group added a stability protection mechanism to the model to provide maximum production to participants in the event that the model observes extraordinary levels of instability in the National Best Bid and Offer during the prior three seconds as compared to reference data. When the model detects such instability, it is programmed to increase the length of the Holding Period to 12 milliseconds for a period of 750 milliseconds.

15 See White Paper, supra, at 22.

16 See id.

Based upon these exciting results, the Exchange now proposes to amend Rule 4702(b)(14) and (15) to replace the static 10 millisecond timers applicable to M-ELO and M-ELO+CB with Dynamic M-ELO Holding Periods. Using the Exchange's proprietary and patent pending technology, the Dynamic M-ELO system will evaluate and, as it deems necessary, adjust the length of the Holding Periods for each symbol comprising M-ELOs and M-ELO+CBs (and Midpoint Orders on the Continuous Book that opt to interact with M-ELO+CBs after resting on the Book) every 30 seconds throughout the Market Hours (each such 30 second interval, a "Change Event"). In so doing, Dynamic M-ELO will help participants to achieve a more optimized blend of the underlying purposes of the M-ELO and M-ELO+CB Order Types: protection against adverse selection (low mark-outs) without sacrificing opportunities to achieve high-quality executions (high fill rates).

A proposed M-ELO or M-ELO+CB with a Dynamic Holding Period will operate as follows. At the outset of Market Hours (approximately 9:30:00 AM), the Exchange will impose initial Holding Periods of 1.25 milliseconds for M-ELOs and M-ELO+CBs in all symbols. Thereafter, Holding Periods for a given symbol will become eligible to change dynamically from the initial duration beginning at 9:30:30AM and then at 30 second intervals thereafter during Market Hours. The Exchange will then apply to the M-ELO or M-ELO+CB Order a Holding Period that is of the duration that prevailed at the time of entry. For example, if participant A enters a M-ELO for symbol XYZ at 9:30:25 AM, then Holding Period for that M-ELO will be 1.25 milliseconds. If at 9:30:30:00 AM, the System decides to lower the duration of the Holding Period by 0.50 milliseconds, and then participant B enters a M-ELO for symbol XYZ at 9:30:45 AM, then the System will assign a 0.75 millisecond Holding Period to participant B's M-ELO. To be clear, the System will determine Dynamic M-ELO Holding Periods independently for M-ELOs and M-ELO+CBs in each symbol.

During normal market conditions, the range of potential Holding Period durations for M-ELOs and M-ELO+CBs will be between 0.25 – 2.50 milliseconds, with the Holding Period duration being eligible to change by increments of either 0.25 or 0.50 milliseconds at each Change Event. Thus, if the Holding Period for a M-ELO in symbol XYZ is set at 0.75 milliseconds at 2:22:11 PM, and at 2:22:41 PM, the System determines to increase the duration of the Holding Period, it may do so only by 0.25 or 0.50 milliseconds during that event.

When a Change Event occurs, and the System determines to adjust the duration of a Holding Period for a symbol, that adjustment will apply, not only to all M-ELOs and M-ELO+CBs for that symbol entered within the 30 second period after the Change Event occurs, but also to M-ELOs and M-ELO+CBs entered prior to the Change Event with unexpired Holding Periods (with applicability retroactive to the time of Order acceptance). Thus, if a participant enters a M-ELO in symbol XYZ at 1:14:299 PM, and the prevailing Holding Period applicable to that M-ELO is 2 milliseconds, and at 1:14:30 PM, the System modifies the Holding Period to be 1.5 milliseconds, then the M-ELO will become eligible to execute at 1:14:3005 P.M. This is the case because the M-ELO will have already expended 1 millisecond of its Holding Period as of the time of the Change Event; thereafter, the M-ELO will need to rest only another 0.5 milliseconds to become eligible to execute under the new 1.5 millisecond Holding Period (as measured from 1:14:299 PM). This last feature ensures that the M-ELO Book maintains time priority among M-ELOs and M-ELO+CBs in a dynamic environment. That is, it ensures that no M-ELO or M-ELO+CB with an unexpired Holding Period at the time of a Change Event will end up becoming eligible to execute later than a M-ELO entered after the Change Event which has a shorter Holding Period applicable to it.

If at any time, the System detects extraordinary instability in a symbol, then the System will activate a "stability protection mechanism" to provide an extra layer of protection to M-ELO and M-ELO users from the heightened risks of adverse selection that exists during such periods of instability.[17] The stability protection mechanism will override the prevailing Holding Periods for M-ELOs and M-ELO+CBs in a symbol experiencing extraordinary instability and immediately increase the duration of those Holding Periods to 12 milliseconds for a period of 750 milliseconds. The System may activate the stability protection mechanism even between Change Events. The System will evaluate, at each NBBO update, whether market conditions remain extraordinarily unstable and, if so, it will restart the 750 millisecond Stability Protected Period and maintain the 12 millisecond Holding Period until conditions stabilize. Once the System

---

[17] For purposes of this Rule, the System determines that "extraordinary instability" for a symbol exists through observations it makes following every change in the NBBO for that symbol that occurs during the trading day. When the NBBO changes, the System looks back at the prior three seconds of trading and measures the difference between the highest and the lowest NBBO midpoint values that occurred during that period, and then it compares that measurement to a threshold value for the symbol. The System concludes that extraordinary instability exists for a symbol if the measurement exceeds the threshold value.

The threshold value for a symbol, in turn, is the difference between the highest and the lowest NBBO midpoint values for the symbol that, if applied to its trading activity during the prior trading day, would have caused the System to deem trading in the symbol to be extraordinarily unstable for as close to one percent of that day as possible.

determines that market conditions have stabilized (i.e., all measurements for the symbol are at or below the threshold value throughout the duration of the prevailing Stability Protected Period), the System will revert the duration of the Holding Periods to that which prevailed as of the Change Event that occurred immediately prior to the activation of the stability protection mechanism or, if the stability protection mechanism was active when a Change Event occurred, to the duration selected at the immediately preceding Change Event. The System will then proceed to reevaluate the duration of the Holding Periods as per the regular schedule of Change Events.

The following is an illustration of the operation of the stability protection mechanism. At 11:10:04 AM, the prevailing Holding Period for M-ELOs in symbol XYZ is 1.5 milliseconds. At the same time, the NBBO for symbol XYZ updates. The System looks back at the prior three seconds of trading in symbol XYZ and finds that during that period, the highest observed NBBO midpoint was $10.05, and the lowest was $10.00, such that the difference between these two values is a range of $0.05. The System then looks back at trading behavior for symbol XYZ during the immediately preceding trading day. In doing so, the System calculates the value of the threshold that would have caused the symbol to be deemed extraordinarily unstable for one percent of the trading day; the System determines that this threshold value is a range of $0.03. The System then compares the $0.03 threshold to its measurement of the prior three seconds of NBBO changes ($0.05), and concludes that over these past three seconds, the symbol is extraordinarily unstable. Accordingly, the System activates the stability protection mechanism and the Holding Period for M-ELOs in symbol XYZ immediately increases to 12 milliseconds for a period of 750 milliseconds. However, 5 milliseconds after the

Stability Protection Period commences, the NBBO updates again, thus prompting the System to repeat its assessment of the stability of the symbol in light of the update. This reassessment reveals that the symbol remains unstable, such that a new Stability Protection Period of 750 milliseconds begins at that time (overriding the pre-existing Period). Over the course of this new Stability Protection Period, the NBBO shifts two more times, but each of the ensuing reassessments indicate that the NBBO ranges for the symbol have fallen below the $0.03 threshold. The Stability Protection Period elapses 750 milliseconds after it began with the symbol remaining stable. Thus, the Holding Period reverts to 1.5 milliseconds.

If the Exchange halts trading in a symbol, then upon resumption of trading, any new M-ELO or M-ELO+CB in that symbol and any pending M-ELO or M-ELO+CB in that symbol with an unexpired Holding Period will be subject to a new 12 milliseconds Holding Period (running from the time when trading resumes) until the next scheduled Change Event, at which point the System may determine to adjust that Holding Period to a duration within the range applicable under normal market conditions.[18] If, however, the System determines that extraordinary instability in the symbol exists, it may instead determine to activate the stability protection mechanism and maintain the duration of the Holding Period at 12 milliseconds for another 750 milliseconds. This design will help to

18 Prior to commencement of a new 12 millisecond Holding Period for a new or pending M-ELO or M-ELO+CB following a Halt, the System will first determine whether the M-ELO or M-ELO+CB is or remains eligible for execution. That is, the Holding Period will commence only if, upon commencement of trading following the Halt, the midpoint price for the Order is within the limit set by the participant. If not, the System will hold the Order until the midpoint falls within the limit set by the participant, at which time the 12 millisecond Holding Period will commence.

ensure that M-ELOs and M-ELO+CBs receive added protection coming out of halt conditions.[19]

The Exchange notes that same dynamic process described above will also apply to and govern the time periods during which Midpoint Orders on the Continuous Book must rest before they will become eligible to interact with M-ELO+CBs (provided that participants have opted for their Midpoint Orders to interact with M-ELO+CBs). Thus, the same Holding Period duration that the System sets for a M-ELO+CB in a symbol during Regular Market Hours will also be the length of time that a Midpoint Order must rest on the Continuous Book must rest before it may interact with a M-ELO+CB.

Apart from these impacts of Dynamic Holding Periods, M-ELOs and M-ELO+CBs will continue to behave as they do now in all respects, and as set forth in Rules 4702(b)(14) and (15).

It is important to note that within the parameters discussed herein and in the White Paper, the Exchange will continue to re-train Dynamic M-ELO and M-ELO+CB regularly so that the model will continue to learn from and act upon the basis of new data, and further improve its performance over time. However, the Exchange will not modify the underlying structure of Dynamic M-ELO and M-ELO+CB without first obtaining the Commission's approval to do so, including modifications to the conditions under which the model will adjust the duration of Holding Periods, the frequency with which the

---

19 Also as a safeguard, the System will apply a default Holding Period of 12 milliseconds to a M-ELO or M-ELO+CB if ever it fails to receive a signal during a Change Event as to whether the System should adjust or maintain the duration of the prevailing Holding Period. The System will continue to apply the default 12 millisecond Holding Period until the next Change Event where the signal is restored and the System is able to act dynamically again.

model my adjust the Holding Periods, and the range of Holding Period durations available to M-ELOs and M-ELO+CBs.[20]

Implementation

The Exchange intends to make the proposed change effective for M-ELOs and M-ELO+CBs in the Second or Third Quarter of 2023, but that time frame is subject to change. The Exchange will publish a Trader Alert in advance of making the proposed change effective.

b. Statutory Basis

The Exchange believes that its proposal is consistent with Section 6(b) of the Act,[21] in general, and furthers the objectives of Section 6(b)(5) of the Act,[22] in particular, in that it is designed to promote just and equitable principles of trade, to remove impediments to and perfect the mechanism of a free and open market and a national market system, and, in general to protect investors and the public interest, by allowing for more widespread use of M-ELOs and M-ELO+CBs.

When the Commission approved the M-ELO and the M-ELO+CB, it determined that these Order Types are consistent with the Act because they "could create additional and more efficient trading opportunities on the Exchange for investors with longer

---

[20] In addition to the proposed changes described above, the Exchange proposes to delete an extraneous reference in Rule 4702(b)(15) to M-ELO+CB being eligible to execute against a Midpoint Order on the Continuous Book if the Continuous Book order has the "Midpoint" Trade Now Attribute enabled. In a prior filing, the Exchange folded the concept of "Midpoint Trade Now" into the general "Trade Now" Attribute. See Securities Exchange Act Release No. 34-92180 (June 15, 2021), 86 FR 33420 (June 24, 2021)(SR-NASDAQ-2021-044).

[21] 15 U.S.C. 78f(b).

[22] 15 U.S.C. 78f(b)(5).

investment time horizons, including institutional investors, and could provide these investors with an ability to limit the information leakage and the market impact that could result from their orders."[23] Nothing about the Exchange's proposal should cause the Commission to revisit or rethink this determination. Indeed, the proposal will not alter the fundamental design of these Order Types, the manner in which they operate, or their effects.

Even with Dynamic M-ELO Holding Periods, M-ELOs and M-ELO+CBs will continue to provide their users with protection against information leakage and adverse selection – and they will do so at levels which are substantially undiminished from that which they provide now.[24]

At the same time, however, the proposal will benefit market participants and investors by reducing the opportunity costs of utilizing M-ELOs and M-ELO+CBs. The proposal, in other words, will re-calibrate the lengths of the Holding Periods so that M-ELOs and M-ELO+CBs will operate in the "Goldilocks" zone – their Holding Periods will not be so short as to render them unable to provide meaningful protections against information leakage and adverse selection, but the Holding Periods also will not be too long so as to cause participants and investors to miss out on favorable execution opportunities. Nasdaq believes the proposal will render M-ELOs and M-ELO+CBs more useful and attractive to market participants and investors, and this increased utility and attractiveness, in turn, will spur an increase in M-ELO and M-ELO+CB use cases on the

---

[23] M-ELO Approval Order, supra 83 FR at 10938–39; M-ELO+CB Approval Order, supra, 84 FR at 48980.

[24] See note 6, supra.

Exchange, both from new and existing users of M-ELOs and M-ELO+CBs. Ultimately, the proposal should enhance market quality by increasing opportunities for midpoint executions on the Exchange.

The Exchange notes that use of Dynamic M-ELOs and M-ELO+CBs remains voluntary for all market participants. Accordingly, if any market participant feels that the dynamic Holding Periods are still too long or too short or because competing venues offer more attractive delay mechanisms, then the participants are free to pursue other trading strategies or utilize other trading venues. They need not utilize Dynamic M-ELOs or M-ELO+CBs.

Finally, the Exchange notes that it will continue to conduct real-time surveillance to monitor the use of M-ELOs and M-ELO+CBs to ensure that such usage remains appropriately tied to the intent of the Order Types. If, as a result of such surveillance, the Exchange determines that the Dynamic M-ELO Holding Periods do not serve their intended purposes, or adversely impact market quality, then the Exchange will seek to make further re-calibrations.

4. Self-Regulatory Organization's Statement on Burden on Competition

The Exchange does not believe that the proposed rule change will impose any burden on competition not necessary or appropriate in furtherance of the purposes of the Act. To the contrary, the Exchange believes that this proposal will promote the competitiveness of the Exchange by rendering its M-ELO and M-ELO+CB Order Types more attractive to participants.

The Exchange adopted the M-ELO and M-ELO+CB as pro-competitive measures intended to increase participation on the Exchange by allowing certain market participants that may currently be underserved on regulated exchanges to compete based on elements other than speed. The proposed change continues to achieve this purpose. With Dynamic M-ELO Holding Periods, both M-ELOs and M-ELO+CBs will afford their users with a level of protection from information leakage and adverse selection that is better from what is achievable at present.[25] At the same time, the Dynamic Holding Periods will increase opportunities to interact with other like-minded investors with longer time horizons while also lowering the opportunity costs for participants that utilize M-ELOs and M-ELO+CBs, particularly for securities that trade within the "Goldilocks" zone. In sum, the proposed changes will not burden competition, but instead may promote competition for liquidity in M-ELOs and M-ELO+CBs by broadening the circumstances in which market participants may find such Orders to be useful. With the proposed changes, market participants will be more likely to determine that the benefits of entering M-ELOs and M-ELO+CBs outweigh the risks of doing so.

The proposed change will not place a burden on competition among market venues, as any market may adopt an order type that operates similarly to a M-ELO or a M-ELO+CB with Dynamic M-ELO Holding Periods.

5. Self-Regulatory Organization's Statement on Comments on the Proposed Rule Change Received from Members, Participants, or Others

No written comments were either solicited or received.

6. Extension of Time Period for Commission Action

The Exchange does not consent to an extension of the time period for Commission action.

---

25 See White Paper, supra.

7. Basis for Summary Effectiveness Pursuant to Section 19(b)(3) or for Accelerated Effectiveness Pursuant to Section 19(b)(2)

Not applicable.

8. Proposed Rule Change Based on Rules of Another Self-Regulatory Organization or of the Commission

Not applicable.

9. Security-Based Swap Submissions Filed Pursuant to Section 3C of the Act

Not applicable.

10. Advance Notices Filed Pursuant to Section 806(e) of the Payment, Clearing and Settlement Supervision Act

Not applicable.

11. Exhibits

1. Notice of Proposed Rule Change for publication in the Federal Register.

3. "Applying Artificial Intelligence & Reinforcement Learning Methods Towards Improving Execution Outcomes," SSRN, October 19, 2022.

5. Text of the proposed rule change.

EXHIBIT 1

SECURITIES AND EXCHANGE COMMISSION
(Release No. ; File No. SR-NASDAQ-2022-079)

December__, 2022

Self-Regulatory Organizations; The Nasdaq Stock Market LLC; Notice of Filing of Proposed Rule Change to Amend Rules 4702(b)(14) and (b)(15) of the Exchange's Rulebook

Pursuant to Section 19(b)(1) of the Securities Exchange Act of 1934 ("Act")[1], and Rule 19b-4 thereunder,[2] notice is hereby given that on December 21, 2022, The Nasdaq Stock Market LLC ("Nasdaq" or "Exchange") filed with the Securities and Exchange Commission ("SEC" or "Commission") the proposed rule change as described in Items I, II, and III, below, which Items have been prepared by the Exchange. The Commission is publishing this notice to solicit comments on the proposed rule change from interested persons.

I. Self-Regulatory Organization's Statement of the Terms of Substance of the Proposed Rule Change

The Exchange proposes to amend Rules 4702(b)(14) and (b)(15) of the Exchange's Rulebook to replace the static holding period requirements for Midpoint Extended Life Orders and Midpoint Extended Life Orders Plus Continuous Book with dynamic holding periods.

---

1 15 U.S.C. 78s(b)(1).

2 17 CFR 240.19b-4.

The text of the proposed rule change is available on the Exchange's Website at https://listingcenter.nasdaq.com/rulebook/nasdaq/rules, at the principal office of the Exchange, and at the Commission's Public Reference Room.

II. Self-Regulatory Organization's Statement of the Purpose of, and Statutory Basis for, the Proposed Rule Change

In its filing with the Commission, the Exchange included statements concerning the purpose of and basis for the proposed rule change and discussed any comments it received on the proposed rule change. The text of these statements may be examined at the places specified in Item IV below. The Exchange has prepared summaries, set forth in sections A, B, and C below, of the most significant aspects of such statements.

A. Self-Regulatory Organization's Statement of the Purpose of, and Statutory Basis for, the Proposed Rule Change

1. Purpose

The Exchange proposes to amend Rules 4702(b)(14) and (15) of the Exchange's Rulebook to replace the static 10 millisecond holding period requirements for its Midpoint Extended Life Order ("M-ELO") and Midpoint Extended Life Order Plus Continuous Book ("M-ELO+CB") Order Types with dynamic holding periods ("Dynamic M-ELO and M-ELO+CB" or collectively, "Dynamic M-ELO").

Background

In 2018, the Exchange introduced the M-ELO, which is a Non-Displayed Order priced at the Midpoint between the National Best Bid and Offer ("NBBO") and which is eligible for execution only against other eligible M-ELOs and only after a minimum of one-half second passes from the time that the System accepts the order (the "Holding

Period").[3] In 2019, the Exchange introduced the M-ELO+CB, which closely resembles the M-ELO, except that a M-ELO+CB may execute at the midpoint of the NBBO, not only against other eligible M-ELOs (and M-ELO+CBs), but also against Non-Displayed Orders with Midpoint Pegging and Midpoint Peg Post-Only Orders ("Midpoint Orders") that rest on the Continuous Book for at least one-half second and have Trade Now enabled.[4]

When the Exchange designed M-ELO, it originally set the length of the Holding Period at one-half second because it determined that this time period would be sufficient to ensure that likeminded investors would interact only with each other, and with minimal market impacts. The Exchange believed that the longer length of the M-ELO Holding Period and its simplicity in design would provide greater protection for participants than they could achieve through competing delay mechanisms.

In 2020, however, the Exchange shortened the length of the Holding Period to 10 milliseconds.[5] The Exchange did so after studying two years of actual use and performance of M-ELOs, as well as customer feedback. That is, the Exchange came to understand that, while users of M-ELO and M-ELO+CB are less concerned with achieving rapid executions of their Orders than are other participants, they are not

---

3 See Securities Exchange Act Release No. 34-82825 (March 7, 2018), 83 FR 10937 (March 13, 2018) (SR-NASDAQ-2017-074) ("M-ELO Approval Order").

4 See Securities Exchange Act Release No. 34-86938 (September 11, 2019), 84 FR 48978 (September 17, 2019) (SR-NASDAQ-2019-048) ("M-ELO+CB Approval Order").

5 See Securities Exchange Act Release No. 34-88743 (April 24, 2020), 85 FR 24068 (April 30, 2020) (SR-NASDAQ-2020-011) ("M-ELO Timer Approval Order").

indifferent about the length of time in which their M-ELOs and M-ELO+CBs must wait before they are eligible for execution. Indeed, participants informed the Exchange that in certain circumstances, such as when they sought to trade symbols that on average had a lower time-to-execution than a half-second, they were reticent to enter M-ELOs or M-ELO+CBs. They indicated that the associated Holding Periods for these Order Types were longer than necessary to achieve the desired protections and that, during the residual portion of the Holding Periods, they risked losing out on favorable execution opportunities that would otherwise be available to them had they placed a non-MELO order.

Based upon this feedback, the Exchange studied the potential effects of reducing the length of the Holding Periods for both M-ELOs and M-ELO+CBs (as well as for Midpoint Orders that would execute against M-ELO+CBs). Ultimately, the Exchange determined that it could reduce the Holding Periods to 10 milliseconds without compromising the protective power that M-ELO and M-ELO+CB are intended to provide to participants and investors.[6] Thus, the Exchange determined that shortening the

6 The Exchange examined each of its historical M-ELO executions to determine at what Midpoints of the NBBO the M-ELOs would have executed if their Holding Periods had been shorter than one-half second (500 milliseconds). After examining the historical effects of shorter Holding Periods of between 10 milliseconds and 400 milliseconds, the Exchange determined that a reduction of the M-ELO Holding Period to as short as 10 milliseconds would have caused an average impact on markouts of only 0.10 basis points (across all symbols). In other words, compared to the execution price of an average M-ELO with a one-half second Holding Period, the Exchange found that a M-ELO with a 10 millisecond Holding Period would have had an average post-execution impact that was only a tenth of a basis point per share – a difference in protective effect that is immaterial. See Nasdaq, "The Midpoint Extended Life Order (M-ELO); M-ELO Holding Period," available at https://www.nasdaq.com/articles/the-

Holding Periods to 10 milliseconds for M-ELOs and M-ELO+CBs would increase the efficacy of the mechanism while not undermining the power of those Order Types to fulfill their underlying purpose of minimizing market impacts. At the same time, the Exchange determined that a reduction in the Holding Periods to 10 milliseconds would dramatically add to the circumstances in which M-ELOs and M-ELO+CBs would be useful to participants. In its M-ELO Timer Approval Order, the Commission agreed with the Exchange:

> The Commission notes that, with the proposed ten-millisecond Holding Period and Resting Period, M-ELOs and M-ELO+CBs would continue to be optional order types that are available to investors with longer investment time horizons, including institutional investors. The Commission also believes that the proposal could make M-ELOs and M-ELO+CBs more attractive for securities that on average have a time-to-execution of less than one-half second and, for investors who currently do not use M-ELOs and M-ELO+CBs for these securities, provide optional order types that could enhance their ability to participate effectively on the Exchange. The Commission notes that, if market participants determine that the proposal would make M-ELOs and M-ELO+CBs less attractive for their particular investment objectives, such market participants may elect to reduce or eliminate their use of these optional order types. Moreover, as noted above, the Exchange will continue to conduct real-time surveillance to monitor the use of M-ELOs and M-ELO+CBs to ensure that such usage remains appropriately tied to

---

midpoint-extended-life-order-m-elo%3A-m-elo-holding-period-2020-02-13 (analyzing effects of shortened Holding Periods on M-ELO performance).

the intent of the order types. If, as a result of such surveillance, the Exchange determines that the shortened Holding Period does not serve its intended purpose or adversely impacts market quality, the Exchange would seek to make further recalibrations.[7]

For similar reasons and with even better potential results for participants, the Exchange now proposes to further refine the length of the Holding Periods for M-ELOs and M-ELO+CBs, this time through the application of innovative and patent pending machine learning technology.

Dynamic M-ELO

After receiving feedback from participants that even 10 millisecond Holding Periods for M-ELO and M-ELO+CB may, at times, exceed what is necessary to accomplish the underlying intent of these Order Types, the Exchange began to experiment with making further refinements to the duration of the Holding Periods. Ultimately, the Exchange concluded that shorter Holding Periods could achieve the same, if not better results for participants in terms of mark-outs, but not in all circumstances. That is, where prices of the underlying securities are stable, and not subject to imminent unfavorable changes, M-ELOs and M-ELO+CBs face lower risks of confronting spread-crossing orders, such that shorter Holding Periods could suffice to protect M-ELOs and M-ELO+CB from such orders. In periods of heightened price volatility, however, M-ELOs and M-ELO+CBs also face heightened risks, such that longer Holding Periods would continue to be beneficial in protecting M-ELOs and M-ELO+CBs from such risks.

[7] M-ELO Timer Approval Order, supra, at 85 FR 24069.

Thus, the Exchange determined that another across-the-board reduction of the static 10 millisecond Holding Periods would be sub-optimal because it could impact the performance of the M-ELO and M-ELO+CB Order Types during periods of heightened volatility.

In light of these observations, the Exchange tasked its artificial intelligence and machine learning laboratory (the "AI Core Development Group") to explore whether it could employ these innovative technologies to optimize the length of M-ELO and M-ELO+CB Holding Periods during various states of price volatility, and then to vary the lengths of the Holding Periods dynamically during the lifecycles of M-ELOs and M-ELO+CBs, with the objectives of improving the performance of these Order Types while also further reducing opportunity costs.

As the Exchange explains in greater depth in the attached white paper, [8] the AI Core Development Group proceeded to develop an artificial intelligence-based timer control system that will achieve these objectives. [9] The AI Core Development Group did so by using reinforcement learning techniques – machine learning paradigms which

---

[8] See Diana Kafkes et al., "Applying Artificial Intelligence & Reinforcement Learning Methods Towards Improving Execution Outcomes," SSRN, October 19, 2022, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4243985 (attached hereto as Exhibit 3) (the "White Paper").

[9] Although the AI Core Development Group acknowledges that an optimal Holding Period would update with every incoming order, it determined that training a reinforcement learning model on every order would be too difficult to program and too difficult to implement given the nanosecond latency requirements of the Exchange. The Group then investigated more feasible update cadences and determined the point at which optimal outcomes were best balanced with the level of programming and implementation difficulty to be between 15 and 30 second updates. Ultimately, the Group chose a 30 second update cadence to give the model the greatest opportunity to learn between potential actions.

develop optimal solutions to problems over time by taking actions to solve them, generating feedback on the results of such actions, applying that feedback to direct and improve the next round of solutions, and then repeating the feedback loop until the paradigm achieves optimized solutions.

In this instance, the AI Core Development Group applied reinforcement learning techniques to a simulation of the M-ELO Book that it constructed using a representative data set from the first quarter of 2022 (the "Training Period"). The Training Period data consisted of 380 out of the 6,257 symbols on the M-ELO Book (accounting for approximately 67 percent of M-ELO volume). The symbols chosen reflect both actively-traded and thinly-traded securities, and both low-priced and high-priced securities.

The AI Core Development Group then developed a machine learning model with more than 140 features[10] and applied it to the Training Period data. The Group programmed the model to value the achievement of higher fill rates or lower mark-outs than that which occurred in a historical simulation of M-ELOs and M-ELO+CBs involving the Training Period data. [11] The Group then programmed the model to seek to achieve its goals by taking one of five possible actions with respect to the duration of the Holding Periods at 30 second intervals[12] for each symbol during each trading day of the Training Period. That is, at each 30 second internal, the model evaluated market conditions for each symbol over the prior 30 second period and either kept the Holding

---

10 See White Paper, supra, at 31, for a description of these features.

11 As the White Paper explains, the Group developed a model to simulate activity on the Exchange involving M-ELOs and M-ELO+CBs during the Training Period. See White Paper, supra, at 10.

12 See id.

Periods the same, increased/decreased them by 0.25 milliseconds, or increased/decreased them by 0.50 milliseconds.[13] After each decision-making round, the model utilized the results to inform its actions at the next 30 second increment.

In making its decisions, the model considered 142 categories of data points. A confluence of data points that correlated with an increase in volatility tended to cause the model to increase the durations of Holding Periods, including increases in the standard deviation of NBBO prices, the number of unique participants placing sell orders on M-ELO and M-ELO+CB, and the volume-weighted average of the NBBO spread. Conversely, a confluence of data points that correlated with greater price stability tended to cause the model to decrease the durations of Holding periods, such as an increase in the median and max number of shares per trade and the number of resting bids left in the M-ELO and M-ELO+CB Book.

The AI Core Development Team produced variations of its model that prioritized achievement of the lowest mark-outs, the highest fill rates, and a blend of these two objectives.[14] Through a process of learning and experimentation, the AI Core

---

13 The AI Core Development Group experimented with a range of permissible Holding Period durations. Ultimately, it concluded that it could produce better outcomes for M-ELO and M-ELO+CB participants than the existing approach using Holding Periods as low as 0.25 milliseconds and as high as 2.5 milliseconds, under normal market conditions.

14 The AI Core Development Group also applied to the model a paradigm called "retraining" to combat the degradation of model performance that can otherwise occur as the reference data it uses for initial comparison becomes stale. Finally, the AI Core Development group added a stability protection mechanism to the model to provide maximum production to participants in the event that the model observes extraordinary levels of instability in the National Best Bid and Offer during the prior three seconds as compared to reference data. When the model detects such instability, it is programmed to increase the length of the Holding Period to 12 milliseconds for a period of 750 milliseconds.

Development Group settled on a Dynamic M-ELO model that achieved substantial performance improvements for users of M-ELO and M-ELO+CB – both in terms of markouts and fill rates – as compared to the static 10 millisecond Holding Periods. As the White Paper explains in greater detail, Dynamic M-ELO yielded an average combined volume-weighted improvement of 31.7 percent, including a 20.3 percent increase in fill rates and a 11.4 percent reduction in mark-outs. [15] The White Paper provides a more fulsome explanation of these improvements.[16]

Based upon these exciting results, the Exchange now proposes to amend Rule 4702(b)(14) and (15) to replace the static 10 millisecond timers applicable to M-ELO and M-ELO+CB with Dynamic M-ELO Holding Periods. Using the Exchange's proprietary and patent pending technology, the Dynamic M-ELO system will evaluate and, as it deems necessary, adjust the length of the Holding Periods for each symbol comprising M-ELOs and M-ELO+CBs (and Midpoint Orders on the Continuous Book that opt to interact with M-ELO+CBs after resting on the Book) every 30 seconds throughout the Market Hours (each such 30 second interval, a "Change Event"). In so doing, Dynamic M-ELO will help participants to achieve a more optimized blend of the underlying purposes of the M-ELO and M-ELO+CB Order Types: protection against adverse selection (low mark-outs) without sacrificing opportunities to achieve high-quality executions (high fill rates).

A proposed M-ELO or M-ELO+CB with a Dynamic Holding Period will operate as follows. At the outset of Market Hours (approximately 9:30:00 AM), the Exchange

---

15 See White Paper, supra, at 22.

16 See id.

will impose initial Holding Periods of 1.25 milliseconds for M-ELOs and M-ELO+CBs in all symbols. Thereafter, Holding Periods for a given symbol will become eligible to change dynamically from the initial duration beginning at 9:30:30AM and then at 30 second intervals thereafter during Market Hours. The Exchange will then apply to the M-ELO or M-ELO+CB Order a Holding Period that is of the duration that prevailed at the time of entry. For example, if participant A enters a M-ELO for symbol XYZ at 9:30:25 AM, then Holding Period for that M-ELO will be 1.25 milliseconds. If at 9:30:30:00 AM, the System decides to lower the duration of the Holding Period by 0.50 milliseconds, and then participant B enters a M-ELO for symbol XYZ at 9:30:45 AM, then the System will assign a 0.75 millisecond Holding Period to participant B's M-ELO. To be clear, the System will determine Dynamic M-ELO Holding Periods independently for M-ELOs and M-ELO+CBs in each symbol.

During normal market conditions, the range of potential Holding Period durations for M-ELOs and M-ELO+CBs will be between 0.25 – 2.50 milliseconds, with the Holding Period duration being eligible to change by increments of either 0.25 or 0.50 milliseconds at each Change Event. Thus, if the Holding Period for a M-ELO in symbol XYZ is set at 0.75 milliseconds at 2:22:11 PM, and at 2:22:41 PM, the System determines to increase the duration of the Holding Period, it may do so only by 0.25 or 0.50 milliseconds during that event.

When a Change Event occurs, and the System determines to adjust the duration of a Holding Period for a symbol, that adjustment will apply, not only to all M-ELOs and M-ELO+CBs for that symbol entered within the 30 second period after the Change Event occurs, but also to M-ELOs and M-ELO+CBs entered prior to the Change Event with unexpired Holding Periods (with applicability retroactive to the time of Order acceptance). Thus, if a participant enters a M-ELO in symbol XYZ at 1:14:299 PM, and the prevailing Holding Period applicable to that M-ELO is 2 milliseconds, and at 1:14:30 PM, the System modifies the Holding Period to be 1.5 milliseconds, then the M-ELO will become eligible to execute at 1:14:3005 P.M. This is the case because the M-ELO will have already expended 1 millisecond of its Holding Period as of the time of the Change Event; thereafter, the M-ELO will need to rest only another 0.5 milliseconds to become eligible to execute under the new 1.5 millisecond Holding Period (as measured from 1:14:299 PM). This last feature ensures that the M-ELO Book maintains time priority among M-ELOs and M-ELO+CBs in a dynamic environment. That is, it ensures that no M-ELO or M-ELO+CB with an unexpired Holding Period at the time of a Change Event will end up becoming eligible to execute later than a M-ELO entered after the Change Event which has a shorter Holding Period applicable to it.

If at any time, the System detects extraordinary instability in a symbol, then the System will activate a "stability protection mechanism" to provide an extra layer of protection to M-ELO and M-ELO users from the heightened risks of adverse selection that exists during such periods of instability.[17] The stability protection mechanism will

---

17 For purposes of this Rule, the System determines that "extraordinary instability" for a symbol exists through observations it makes following every change in the NBBO for that symbol that occurs during the trading day. When the NBBO changes, the System looks back at the prior three seconds of trading and measures the difference between the highest and the lowest NBBO midpoint values that occurred during that period, and then it compares that measurement to a threshold value for the symbol. The System concludes that extraordinary instability exists for a symbol if the measurement exceeds the threshold value.

The threshold value for a symbol, in turn, is the difference between the highest and the lowest NBBO midpoint values for the symbol that, if applied to its trading override the prevailing Holding Periods for M-ELOs and M-ELO+CBs in a symbol experiencing extraordinary instability and immediately increase the duration of those Holding Periods to 12 milliseconds for a period of 750 milliseconds. The System may activate the stability protection mechanism even between Change Events. The System will evaluate, at each NBBO update, whether market conditions remain extraordinarily unstable and, if so, it will restart the 750 millisecond Stability Protected Period and maintain the 12 millisecond Holding Period until conditions stabilize. Once the System determines that market conditions have stabilized (i.e., all measurements for the symbol are at or below the threshold value throughout the duration of the prevailing Stability Protected Period), the System will revert the duration of the Holding Periods to that which prevailed as of the Change Event that occurred immediately prior to the activation of the stability protection mechanism or, if the stability protection mechanism was active when a Change Event occurred, to the duration selected at the immediately preceding Change Event. The System will then proceed to reevaluate the duration of the Holding Periods as per the regular schedule of Change Events.

The following is an illustration of the operation of the stability protection mechanism. At 11:10:04 AM, the prevailing Holding Period for M-ELOs in symbol XYZ is 1.5 milliseconds. At the same time, the NBBO for symbol XYZ updates. The System looks back at the prior three seconds of trading in symbol XYZ and finds that during that period, the highest observed NBBO midpoint was $10.05, and the lowest was $10.00, such that the difference between these two values is a range of $0.05. The

activity during the prior trading day, would have caused the System to deem trading in the symbol to be extraordinarily unstable for as close to one percent of that day as possible.

System then looks back at trading behavior for symbol XYZ during the immediately preceding trading day. In doing so, the System calculates the value of the threshold that would have caused the symbol to be deemed extraordinarily unstable for one percent of the trading day; the System determines that this threshold value is a range of $0.03. The System then compares the $0.03 threshold to its measurement of the prior three seconds of NBBO changes ($0.05), and concludes that over these past three seconds, the symbol is extraordinarily unstable. Accordingly, the System activates the stability protection mechanism and the Holding Period for M-ELOs in symbol XYZ immediately increases to 12 milliseconds for a period of 750 milliseconds. However, 5 milliseconds after the Stability Protection Period commences, the NBBO updates again, thus prompting the System to repeat its assessment of the stability of the symbol in light of the update. This reassessment reveals that the symbol remains unstable, such that a new Stability Protection Period of 750 milliseconds begins at that time (overriding the pre-existing Period). Over the course of this new Stability Protection Period, the NBBO shifts two more times, but each of the ensuing reassessments indicate that the NBBO ranges for the symbol have fallen below the $0.03 threshold. The Stability Protection Period elapses 750 milliseconds after it began with the symbol remaining stable. Thus, the Holding Period reverts to 1.5 milliseconds.

If the Exchange halts trading in a symbol, then upon resumption of trading, any new M-ELO or M-ELO+CB in that symbol and any pending M-ELO or M-ELO+CB in that symbol with an unexpired Holding Period will be subject to a new 12 milliseconds Holding Period (running from the time when trading resumes) until the next scheduled Change Event, at which point the System may determine to adjust that Holding Period to a duration within the range applicable under normal market conditions.[18] If, however, the System determines that extraordinary instability in the symbol exists, it may instead determine to activate the stability protection mechanism and maintain the duration of the Holding Period at 12 milliseconds for another 750 milliseconds. This design will help to ensure that M-ELOs and M-ELO+CBs receive added protection coming out of halt conditions.[19]

The Exchange notes that same dynamic process described above will also apply to and govern the time periods during which Midpoint Orders on the Continuous Book must rest before they will become eligible to interact with M-ELO+CBs (provided that participants have opted for their Midpoint Orders to interact with M-ELO+CBs). Thus, the same Holding Period duration that the System sets for a M-ELO+CB in a symbol during Regular Market Hours will also be the length of time that a Midpoint Order must rest on the Continuous Book must rest before it may interact with a M-ELO+CB.

---

18 Prior to commencement of a new 12 millisecond Holding Period for a new or pending M-ELO or M-ELO+CB following a Halt, the System will first determine whether the M-ELO or M-ELO+CB is or remains eligible for execution. That is, the Holding Period will commence only if, upon commencement of trading following the Halt, the midpoint price for the Order is within the limit set by the participant. If not, the System will hold the Order until the midpoint falls within the limit set by the participant, at which time the 12 millisecond Holding Period will commence.

19 Also as a safeguard, the System will apply a default Holding Period of 12 milliseconds to a M-ELO or M-ELO+CB if ever it fails to receive a signal during a Change Event as to whether the System should adjust or maintain the duration of the prevailing Holding Period. The System will continue to apply the default 12 millisecond Holding Period until the next Change Event where the signal is restored and the System is able to act dynamically again.

Apart from these impacts of Dynamic Holding Periods, M-ELOs and M-ELO+CBs will continue to behave as they do now in all respects, and as set forth in Rules 4702(b)(14) and (15).

It is important to note that within the parameters discussed herein and in the White Paper, the Exchange will continue to re-train Dynamic M-ELO and M-ELO+CB regularly so that the model will continue to learn from and act upon the basis of new data, and further improve its performance over time. However, the Exchange will not modify the underlying structure of Dynamic M-ELO and M-ELO+CB without first obtaining the Commission's approval to do so, including modifications to the conditions under which the model will adjust the duration of Holding Periods, the frequency with which the model my adjust the Holding Periods, and the range of Holding Period durations available to M-ELOs and M-ELO+CBs.[20]

Implementation

The Exchange intends to make the proposed change effective for M-ELOs and M-ELO+CBs in the Second or Third Quarter of 2023, but that time frame is subject to change. The Exchange will publish a Trader Alert in advance of making the proposed change effective.

---

20 In addition to the proposed changes described above, the Exchange proposes to delete an extraneous reference in Rule 4702(b)(15) to M-ELO+CB being eligible to execute against a Midpoint Order on the Continuous Book if the Continuous Book order has the "Midpoint" Trade Now Attribute enabled. In a prior filing, the Exchange folded the concept of "Midpoint Trade Now" into the general "Trade Now" Attribute. See Securities Exchange Act Release No. 34-92180 (June 15, 2021), 86 FR 33420 (June 24, 2021)(SR-NASDAQ-2021-044).

2. Statutory Basis

The Exchange believes that its proposal is consistent with Section 6(b) of the Act,[21] in general, and furthers the objectives of Section 6(b)(5) of the Act,[22] in particular, in that it is designed to promote just and equitable principles of trade, to remove impediments to and perfect the mechanism of a free and open market and a national market system, and, in general to protect investors and the public interest, by allowing for more widespread use of M-ELOs and M-ELO+CBs.

When the Commission approved the M-ELO and the M-ELO+CB, it determined that these Order Types are consistent with the Act because they "could create additional and more efficient trading opportunities on the Exchange for investors with longer investment time horizons, including institutional investors, and could provide these investors with an ability to limit the information leakage and the market impact that could result from their orders."[23] Nothing about the Exchange's proposal should cause the Commission to revisit or rethink this determination. Indeed, the proposal will not alter the fundamental design of these Order Types, the manner in which they operate, or their effects.

Even with Dynamic M-ELO Holding Periods, M-ELOs and M-ELO+CBs will continue to provide their users with protection against information leakage and adverse

---

21 15 U.S.C. 78f(b).

22 15 U.S.C. 78f(b)(5).

23 M-ELO Approval Order, supra 83 FR at 10938–39; M-ELO+CB Approval Order, supra, 84 FR at 48980.

selection – and they will do so at levels which are substantially undiminished from that which they provide now.[24]

At the same time, however, the proposal will benefit market participants and investors by reducing the opportunity costs of utilizing M-ELOs and M-ELO+CBs. The proposal, in other words, will re-calibrate the lengths of the Holding Periods so that M-ELOs and M-ELO+CBs will operate in the "Goldilocks" zone – their Holding Periods will not be so short as to render them unable to provide meaningful protections against information leakage and adverse selection, but the Holding Periods also will not be too long so as to cause participants and investors to miss out on favorable execution opportunities. Nasdaq believes the proposal will render M-ELOs and M-ELO+CBs more useful and attractive to market participants and investors, and this increased utility and attractiveness, in turn, will spur an increase in M-ELO and M-ELO+CB use cases on the Exchange, both from new and existing users of M-ELOs and M-ELO+CBs. Ultimately, the proposal should enhance market quality by increasing opportunities for midpoint executions on the Exchange.

The Exchange notes that use of Dynamic M-ELOs and M-ELO+CBs remains voluntary for all market participants. Accordingly, if any market participant feels that the dynamic Holding Periods are still too long or too short or because competing venues offer more attractive delay mechanisms, then the participants are free to pursue other trading strategies or utilize other trading venues. They need not utilize Dynamic M-ELOs or M-ELO+CBs.

---

24 See note 6, supra.

Finally, the Exchange notes that it will continue to conduct real-time surveillance to monitor the use of M-ELOs and M-ELO+CBs to ensure that such usage remains appropriately tied to the intent of the Order Types. If, as a result of such surveillance, the Exchange determines that the Dynamic M-ELO Holding Periods do not serve their intended purposes, or adversely impact market quality, then the Exchange will seek to make further re-calibrations.

B. Self-Regulatory Organization's Statement on Burden on Competition

The Exchange does not believe that the proposed rule change will impose any burden on competition not necessary or appropriate in furtherance of the purposes of the Act. To the contrary, the Exchange believes that this proposal will promote the competitiveness of the Exchange by rendering its M-ELO and M-ELO+CB Order Types more attractive to participants.

The Exchange adopted the M-ELO and M-ELO+CB as pro-competitive measures intended to increase participation on the Exchange by allowing certain market participants that may currently be underserved on regulated exchanges to compete based on elements other than speed. The proposed change continues to achieve this purpose. With Dynamic M-ELO Holding Periods, both M-ELOs and M-ELO+CBs will afford their users with a level of protection from information leakage and adverse selection that is better from what is achievable at present.[25] At the same time, the Dynamic Holding Periods will increase opportunities to interact with other like-minded investors with longer time horizons while also lowering the opportunity costs for participants that utilize M-ELOs and M-ELO+CBs, particularly for securities that trade within the "Goldilocks"

---

25 See White Paper, supra.

zone. In sum, the proposed changes will not burden competition, but instead may promote competition for liquidity in M-ELOs and M-ELO+CBs by broadening the circumstances in which market participants may find such Orders to be useful. With the proposed changes, market participants will be more likely to determine that the benefits of entering M-ELOs and M-ELO+CBs outweigh the risks of doing so.

The proposed change will not place a burden on competition among market venues, as any market may adopt an order type that operates similarly to a M-ELO or a M-ELO+CB with Dynamic M-ELO Holding Periods.

C. Self-Regulatory Organization's Statement on Comments on the Proposed Rule Change Received from Members, Participants, or Others

No written comments were either solicited or received.

III. Date of Effectiveness of the Proposed Rule Change and Timing for Commission Action

Within 45 days of the date of publication of this notice in the Federal Register or within such longer period (i) as the Commission may designate up to 90 days of such date if it finds such longer period to be appropriate and publishes its reasons for so finding or (ii) as to which the Exchange consents, the Commission shall: (a) by order approve or disapprove such proposed rule change, or (b) institute proceedings to determine whether the proposed rule change should be disapproved.

IV. Solicitation of Comments

Interested persons are invited to submit written data, views, and arguments concerning the foregoing, including whether the proposed rule change is consistent with the Act. Comments may be submitted by any of the following methods:

Electronic comments:

- Use the Commission's Internet comment form (http://www.sec.gov/rules/sro.shtml); or
- Send an e-mail to rule-comments@sec.gov. Please include File Number SR-NASDAQ-2022-079 on the subject line.

Paper comments:

- Send paper comments in triplicate to Secretary, Securities and Exchange Commission, 100 F Street, NE, Washington, DC 20549-1090.

All submissions should refer to File Number SR-NASDAQ-2022-079. This file number should be included on the subject line if e-mail is used. To help the Commission process and review your comments more efficiently, please use only one method. The Commission will post all comments on the Commission's Internet Web site (http://www.sec.gov/rules/sro.shtml).

Copies of the submission, all subsequent amendments, all written statements with respect to the proposed rule change that are filed with the Commission, and all written communications relating to the proposed rule change between the Commission and any person, other than those that may be withheld from the public in accordance with the provisions of 5 U.S.C. 552, will be available for website viewing and printing in the Commission's Public Reference Room, 100 F Street, NE, Washington, DC 20549, on official business days between the hours of 10:00 a.m. and 3:00 p.m. Copies of the filing also will be available for inspection and copying at the principal office of the Exchange. All comments received will be posted without change; the Commission does not edit personal identifying information from submissions. You should submit only information that you wish to make available publicly.

All submissions should refer to File Number SR-NASDAQ-2022-079 and should be submitted on or before [insert date 21 days from publication in the Federal Register].

For the Commission, by the Division of Trading and Markets, pursuant to delegated authority.[26]

J. Matthew DeLesDernier
Assistant Secretary

[26] 17 CFR 200.30-3(a)(12).

EXHIBIT 3

Applying Artificial Intelligence & Reinforcement Learning Methods Towards Improving Execution Outcomes

Diana Kafkes,* Josep Puig Ruiz*, Drew Rooks*, Douglas Hamilton*, and Michael O'Rourke*

Nasdaq Stock Exchange

October 10, 2022

1 Introduction

Capital markets are dynamic and always evolving. This constant evolution provides opportunities to enhance timer based solutions to improve execution outcomes and further mitigate adverse selection and volatility. This paper proposes the introduction of a dynamic model-based timer that can be applied to the Nasdaq Midpoint Extended Life Order (M-ELO).

The current M-ELO uses a holding period— a brief waiting timer applied to the queue on both sides— statically set at 10ms (10^{-3} s). This holding period achieves relatively favorable markout when trading institutional loads with respect to the continuous book, achieving the intended purpose of matching like-minded long-term investors. However, this occurs at some expense of fill rate given the differences in the dynamics of individual securities. Our research shows that the interplay between fill rate and markout is more nuanced than a simple trade-off, as previous literature indicates. Here we demonstrate that achieving both higher fill rate and lower markout

*diana.kafkes@nasdaq.com, josep.ruiz@nasdaq.com, drew.rooks@nasdaq.com, douglas.hamilton@nasdaq.com, michael.orourke@nasdaq.com

is possible through leveraging Artificial Intelligence (AI) towards achieving dynamically-improved timer execution.

Our Dynamic M-ELO system[1] leverages an AI control scheme known as reinforcement learning to evaluate the duration of the holding period timer based on local market conditions. Our research shows that applying a dynamic timer to MELO achieves an increase in fill rate of 20.3% and a simultaneous decrease in markout of 11.4% compared to the current static holding period. We believe that these results are indicative of the many improvements AI-enhanced mechanisms can bring to capital markets. Here we detail the development and results of our proposed timer-update model and advance its adoption as a step towards the future of dynamic market order types.

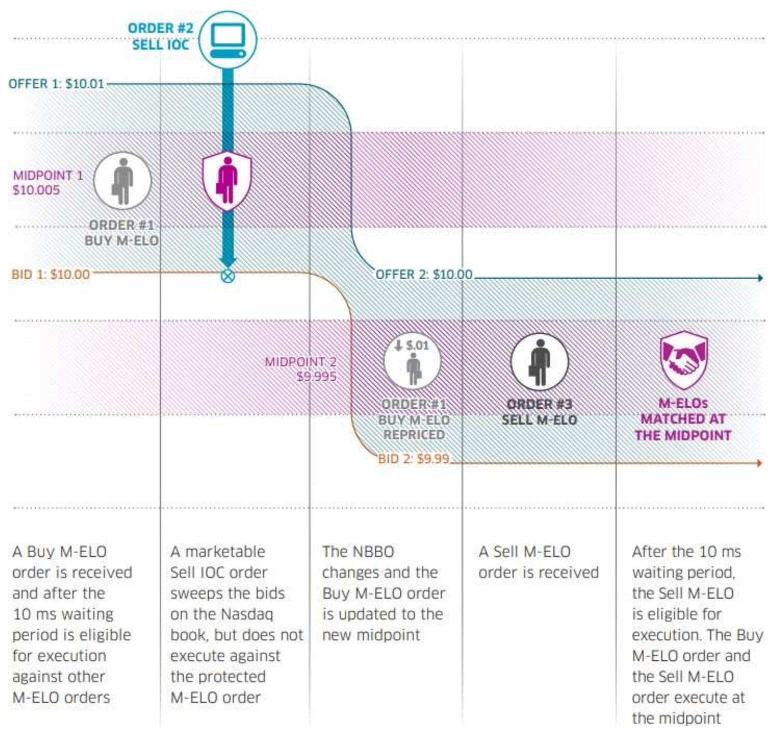

Figure 1: An example of a M-ELO trade.

[1]Our Dynamic M-ELO system is patent pending; we have filed several patent applications covering various aspects of it.

2 Assessing the State of the Current Static Holding Period

First, we present some details on how the current holding period used with M-ELO[2] functions in order to provide context for the improvement gains our system is able to make on the original.

The current static holding period market is designed to match like-minded longer-term investors on a broker-neutral exchange. When a M-ELO order is entered, it is held for a period of time before it becomes executable (Figure 1). This period of time depends both on the static holding period timer and a M-ELO order eligibility condition.

2.1 M-ELO Eligibility to Trade

Incoming M-ELO orders in the current implementation are assigned a static waiting timer of 10ms. This timer was updated from its earlier value of 500ms in May 2020.

This timer is only activated by a specific midpoint-crossing eligibility condition: the NBBO (National Best Bid and Offer) of the security or fund must have favorably crossed the price of the entered order. For buy orders, this means that the orders become eligible if the midpoint of the NBBO for that security is either at or lower than the bid price of the M-ELO order. Similarly, for sell orders, the midpoint of the NBBO must be either at or higher than the ask price of the M-ELO order in order to be eligible. As soon as an order becomes eligible by this criteria, it begins its waiting period of 10ms.

It is important to note that the waiting timer is not set off immediately when an order is entered, and to distinguish the two distinct waiting periods an order faces in the M-ELO queue: one before it is deemed eligible based on a midpoint-crossing condition and one that depends on the holding period timer assigned, in this case 10ms.

[2]Please note that the proposed system can be applied to both Nasdaq Midpoint Extended Life Order (M-ELO) and Nasdaq Midpoint Extended Life Order Plus Continuous Book (M-ELO+CB). For the sake of brevity in notation, and unless otherwise noted, we will henceforth refer to both M-ELO and M-ELO+CB as simply "M-ELO".

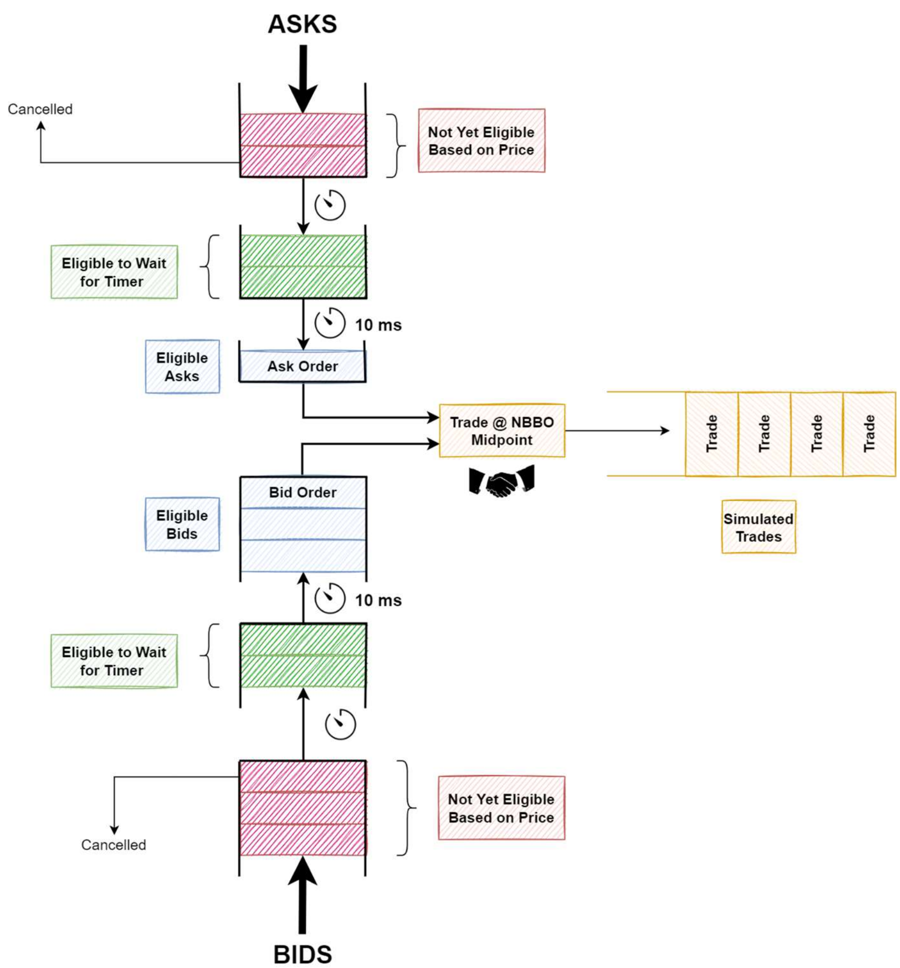

Figure 2: How Orders Become Eligible To Trade on M-ELO

2.2 Fill Rate-Markout Trade-off

While all orders are treated equally by the assignment of the same timer in the simplest notion of the word, certain trade-offs arise from the static timer's interactions with non-static market conditions that affect orders unequally.

2.2.1 Defining Fill Rate and Markout

In order to address these inefficiencies, we first must define two metrics that give us some notion for the underlying liquidity and execution quality of the market: fill rate and markout, respectively [**5**, **6**, **7**]. We will refer to these metrics throughout this White Paper in order to evaluate the performance of the current M-ELO market as well as our proposed adjustment to it that achieves the best of both worlds: increased fill rates and decreased markout. We define fill rate for a period of time and markout by trade as follows:

$$\text{FR}_{\text{period}} = \frac{\text{Shares Traded}}{\text{Total Incoming Shares}} \tag{1}$$

$$\text{MO}_{\delta,\text{bps,trade}} = \begin{cases} MO_{\delta,\text{bps,trade Buy}} = \left| \frac{-100*100(M_t - M_{t+\delta})}{M_{t+\delta}} \right| \\ MO_{\delta,\text{bps,trade Sell}} = \left| \frac{100*100(M_t - M_{t+\delta})}{M_{t+\delta}} \right| \end{cases} \tag{2}$$

where M_t is the midpoint of the NBBO at time t and δ_t denotes a specific time horizon.

Note here that the markout is always positive, i.e. it is defined as the absolute value of the price (dis)improvement for each side and counts as a penalty for any price movement within the time horizon. The negative sign inside the absolute value operator for the "Buy" markout component is not strictly needed, but we've decided to leave it in place to highlight the fact that generally price improvement for one side of a trade is met with price disimprovement for the other side of the same trade. Regardless, the goal is to minimize price deviation in *any* direction, hence the absolute value operator.

It's also worth noting that, while we considered many different time horizons in our research, here we will focus on discussing the specific case with δ_t = 1 second (1s), i.e. the markout expressed in basis points (bps) one second after trading. Different time horizons yield similar results.

Lastly, note also that we are using a simplified aggregate shares formula for fill rate instead of fulfillment × hit rate. Using this alternative definition also yields similar results.

2.2.2 Static Timer vs. Non-Static Market Conditions

We propose here a solution that leverages AI to provide a bespoke intraday symbol-specific recommendation for the holding period at any given time, based on the specific conditions for that symbol. We find that by allowing for the holding period to change for each specific security and every 30 second interval we can further enhance the favorable outcomes that M-ELO participants already see.

We present a few simple examples to show how M-ELO currently operates and why it is effective. After that, we will demonstrate as well the additional value of having a changing holding period under certain situations.

We present first two examples that showcase the current behavior of M-ELO, and why it works well in providing favorable execution for like-minded parties, through Figures 3a and 3b. Two M-ELO orders, A and B, are placed, wait for the holding period of 10ms to elapse, and then become eligible to trade at the midpoint of the NBBO.

- Figure 3a depicts a situation in which the market conditions are very stable throughout. Under such circumstances, M-ELO allows for execution to occur, matching the two like-minded parties, and the outcome is satisfactory to both sides.
- Similarly, Figure 3b shows a situation in which market conditions get momentarily unstable, and price fluctuates significantly. Since the order that arrived later needs to wait as well for the holding period to elapse, the execution does not occur during this unstable moment, but rather occurs at a slightly later time, once the new price level is established. As a result, the later execution is favorable to both sides.

Now, we highlight how a changing holding period can further increase the performance of M-ELO. Order A is placed, waits for 10ms, and is now ready to trade, but there is no liquidity on the other side yet, so it rests. Moments later, a sell M-ELO order, Order B, is placed too. The conditions at this point in time are very stable for this symbol: if Order B was to trade immediately with Order A, the execution outcome would be favorable for both sides.

However, Order B needs to wait as well for the holding period to elapse. We explore two possible situations based on the duration of the holding period:

- With the static timer of 10ms, Order B needs to wait for 10ms. At some point, while waiting for the 10ms to elapse, market conditions worsen. Once the timer has elapsed, Order B becomes eligible to trade. Consequently, a trade occurs between Orders A & B at the midpoint of the NBBO. Due to the recent instability, the execution outcome is less favorable. This situation is depicted in Figure 4a.

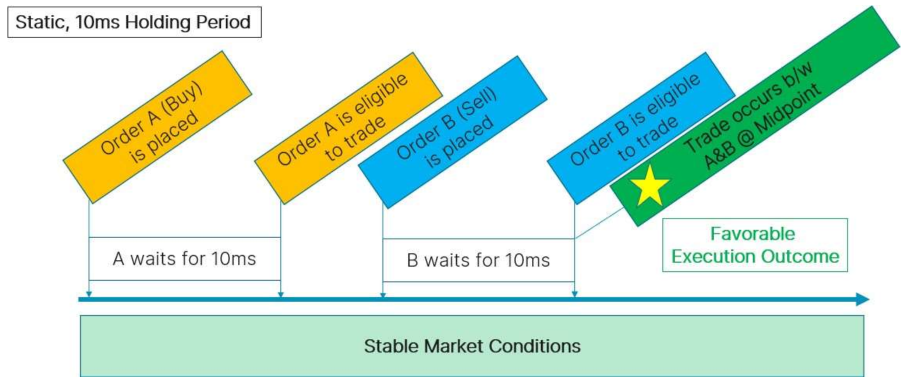

(a) M-ELO in normal market conditions.

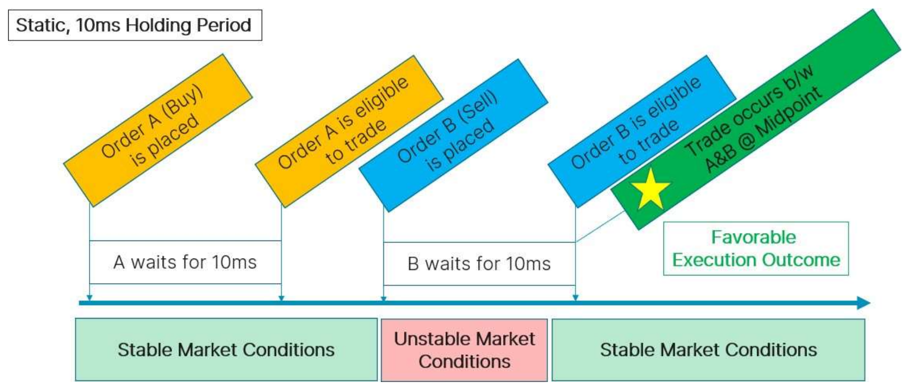

(b) M-ELO in unstable market conditions.

Figure 3: Example showing how current static M-ELO operates and how it effectively provides favorable execution outcomes that would have otherwise occurred in both normal and unstable market conditions.

- In contrast, assume a shorter dynamic timer, for example 2ms, and a stable market. The shorter dynamic allows A & B to trade sooner, which is a more favorable outcome for both parties during a stable period in the market. It is indeed a current market reality that execution algorithms brokers provide to asset managers and other institutions often need to rebalance liquidity across venues which may cause orders that could potentially match to miss each other due to a timer that is longer than it needs to be to provide the right level of protection. Fewer missed executions increases the efficacy of those broker algorithms and therefore improves investor experience in M-ELO. Indeed, there are situations in which a changing holding period provides desirable outcomes, highlighting the need for the proposed solution. This situation is depicted in Figure 4b.

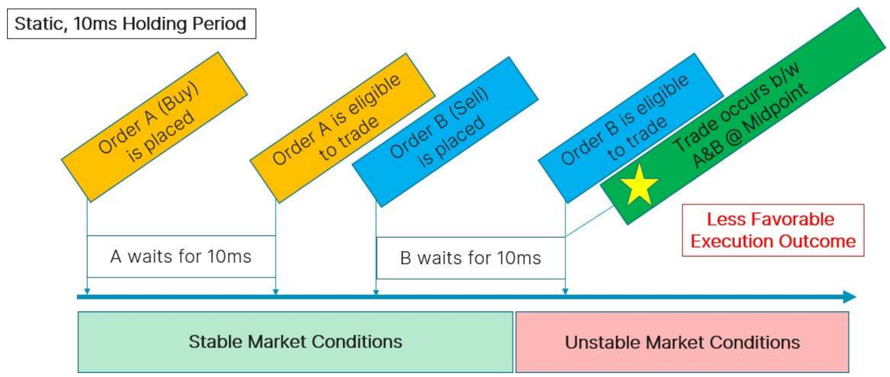

(a) Static timer might result in less favorable outcomes under specific circumstances.

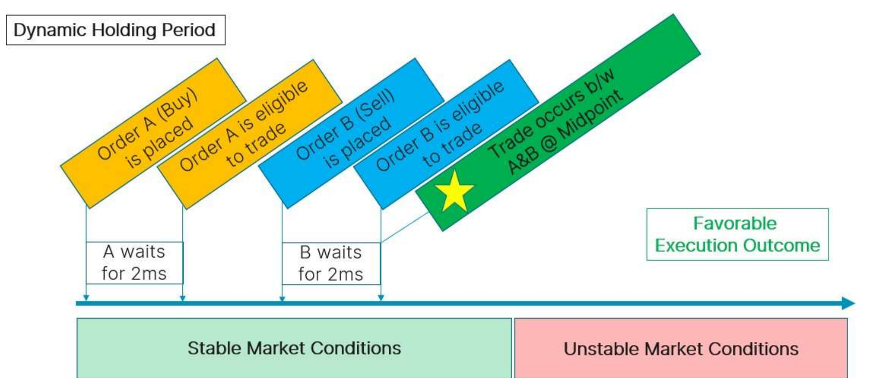

(b) Shorter timer might result in more favorable outcomes under specific circumstances.

Figure 4: Example where a static timer might result in a less favorable execution outcome, showcasing the value of dynamically reducing the timer under specific circumstances.

Lastly, for the rare situation in which market conditions significantly deteriorate and a longer holding period might be desirable to shield against volatility, the stability protection mechanism will be momentarily activated, increasing the holding period to a higher value for a short amount of time. This mechanism, which is described in more depth in Section 4.3, is depicted in Figure 5.

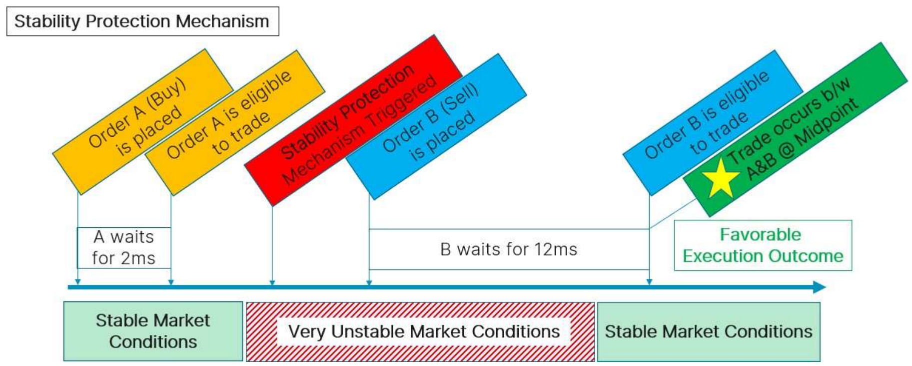

Figure 5: The stability protection mechanism acts as a safeguard against undesired execution outcomes driven by periods of high, unusual volatility. More details on this can be found in Section 4.3

3 Building an AI-Based Timer-Control System

Dynamic M-ELO is an AI-based control system that learns from watching market behavior and can dynamically adjust the holding period timer to improve client outcomes with respect to both liquidity and execution quality. The model evaluates changes to the timer between the range of .25 to 2.5ms by .25ms increments every 30 seconds (780 times per day) and includes stability protection enhancement (discussed at length in Section 4.3) that briefly overrides the timer to 12ms in the case of high volatility.

Here we discuss the methodology involved in the creation of Dynamic M-ELO, including an overview of reinforcement learning, how we built and validated a surrogate environment for our training process, and insight into our model development.

3.1 Reinforcement Learning

We framed our dynamic timer control solution as a reinforcement learning problem. We did this because we needed a way for our model to inter- act with and receive feedback from the trading system based on each timer chosen. Other machine learning approaches would have involved simulating a large or even intractably-large—depending on the frequency of timer updates chosen— number of timer-paths across our chosen learning period. This made reinforcement learning the clear choice.

Reinforcement learning is an AI paradigm in which a model ("agent") is trained to take the most optimal actions in an environment. This is achieved by the model taking random actions at first and receiving feedback— both positive and negative— in the form of a reward that is then used to tailor the model's future approach [**3**]. Overtime, the agent gradually switches from exploring its environment through random actions to exploiting what it has learned through these earlier actions [**3**]. The training of the model in this way involves a loop (Figure 6) that borrows heavily from the psychology of how we learn to walk or ride a bike— with all the falls, cuts, and scrapes included on the way to eventually achieving balance.

In a generic single iteration known as an episode: the agent takes an action and receives feedback from the environment in the form of both a reward and the next state from which it can take another action. Each time around the loop generates a (state, action, reward, next state) "experience" which is stored in what is known as the agent's memory buffer and accessed

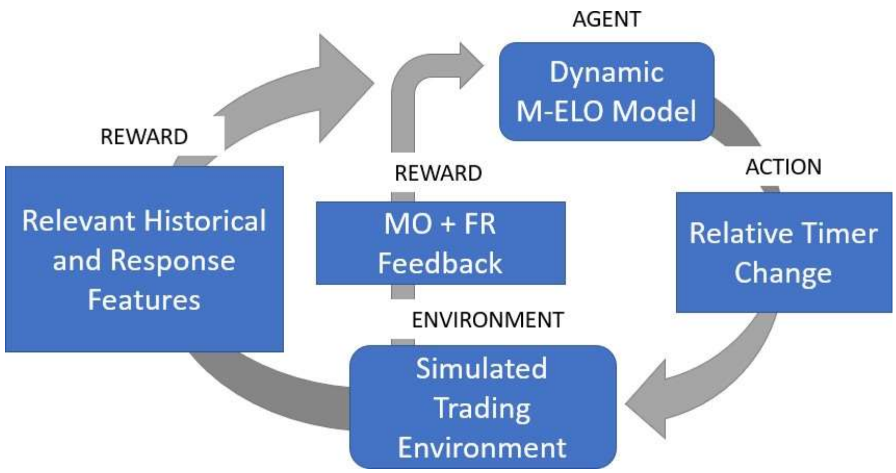

Figure 6: The Reinforcement Learning framework used to train Dynamic M-ELO.

for training [**3**].

Likewise, in our Dynamic M-ELO training process, we have a model (agent) select a relative timer change (action) that then impacts a simulated trading environment which provides feedback in the form of the next state and reward. In this case, the next state provided is the relevant historical and response features returned from the environment and the reward is a share-weighted linear combination of fill rate and markout, the proxies for liquidity and execution quality we discussed before in Equations 1 and 2:

$$\frac{\lambda}{q_i} \sum_{i=1}^{T} q_i(\mathrm{MO}_{\delta,\mathrm{agent,bps}} - \mathrm{MO}_{\delta,\mathrm{synthetic@10,bps}}) + (1-\lambda)(\mathrm{FR}_{\mathrm{agent}} - \mathrm{FR}_{\mathrm{sim@10}}) \tag{3}$$

where q_i is the number of shares traded in trade i, T is the total number of trades in a period, and λ is a weighting design factor. This can be interpreted as the marginal advantage the Dynamic M-ELO agent is able to achieve with respect to the 10ms M-ELO as represented in our simulated environment. For a comprehensive explanation of how the simulated environment at 10ms corresponds with historical M-ELO data, please see Section 5.1.

3.2 Developing a Simulation of M-ELO with a Non-Static Holding Period

In order to train our agent, we first had to build a simulated M-ELO trading environment. Without it, the Dynamic M-ELO agent would not have had a safe place to interact with and learn from without causing undue damage. This is often done in reinforcement learning to ensure that the actual environment is not disturbed. Within this simulated trading environment, the agent is able to change the duration of the holding period timer and receive positive and negative feedback relating how the environment's internal state reacts to this timer change.

The simulation built mostly consists of the M-ELO order eligibility logic detailed in Section 2.1 wrapped around a matching engine that ingests historical and slightly modified data. This matching engine accepts orders, cancellations, and replacements and has the basic capability of matching buy/sell orders of the same symbol at the same price.

The reason why slightly modified data is needed as an input to the matching engine is that the simulation is built to show what the trading environment *would have been* like at timers other than 10ms. This timer change importantly affects user's cancellation behavior.

3.2.1 Simulating Cancellations

We defined a way to statistically model appropriate cancellation behavior at timers other than 10ms based on historical past behavior. For each order without historical cancellation after the 10ms holding period, we sample from a binomial distribution to figure out whether or not the order would have a chance of cancelling. This binomial distribution was fit using the likelihood of that user cancelling for each specific symbol at a given time of day. If the user in question did not historically cancel that symbol at that time of day, we used that user's overall probability of cancelling for that symbol in general.

If the results of the binomial draw were favorable for cancellation, we then sampled the time interval between when the order was placed and when the cancellation would occur from an exponential distribution. This exponential distribution's parameters were fit from past user historical cancellation behavior for that symbol at that time of day. Similarly to the likelihood of cancelling calculation, missing values were filled in first by the user for that symbol, and then for that user in general.

3.2.2 Validating our Simulation

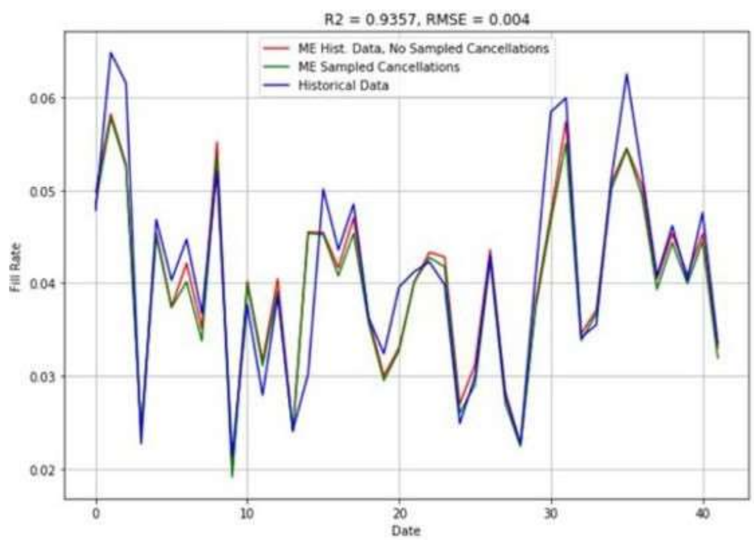

Figure 7: Validation of simulation against historical with and without simulated cancellation behavior. Results presented are a subset.

We ran our simulation forward at the 10ms timer to validate it. We received high RMSE values (averaging above .95, and all greater than .88) when plotting our resulting trades for specific symbols in comparison with the historical. The overall trend is that our simulation tends to undershoot the historical trades by a small amount. Performing this study at 10ms verified that we could replicate historical M-ELO activity with high fidelity and gave us high confidence that we could use this simulation to provide a reasonable estimate for the trading environment at different timers. Therefore, we could use the simulation as a training ground for our agent.

4 Model Development

Our timer-set agent is a deep neural network with multiple layers and over 35,000 parameters. It was trained following the DDQN (Double Deep Q-Network) training paradigm, where the agent's state-action value (Q-value) is approximated using a deep neural network [**2**].

4.1 Understanding DDQN

The Q-value provides some notion of the long-term value of a discrete sum of the agent's discounted actions [**1**]. To make this more concrete, the agent's actions will update the state it is in, sometimes in favorable ways and sometimes in nonfavorable ways. The favorable updates will increase the Q-value, and the nonfavorable actions will decrease the Q-value. These incremental updates are summed together for a discrete number of times around the RL Loop (Figure 6) with a nearsighted discount factor that makes Q-value updates that happened earlier in the chain decrease in impact as time goes by. Through many iterations of the training loop, the Q-value can be optimized.

The DDQN algorithm is a variant of the DQN algorithm, and involves initializing two identical models— a main and target network— and updating the target model's weights less frequently than the main model's [**2**]. While the main model goes around the loop and collects batches of (state, action, reward, next state) experiences in its memory buffer, the target model's parameter updates happen less frequently than the main model's for the sake of stability. This means that the target model's parameters cannot become over-adapted to one set of experiences, rendering it more stable than the main network which adapts itself to each set of sampled experiences. This lag is advantageous as the main model's parameters are periodically updated to be a weighted average of the main model and the more stable target model, ensuring it retains a middle ground of adaptation.

4.2 The Training Process

We trained our agent on data derived from our simulation operating between January 1, 2022 and April 1, 2022 (Q1)— a time period which we assessed to be representative in both duration and market behavior, including a range of relevant market volatility conditions. From this period of time, we used 380 symbols that represent a subset of the 6257 symbols that are actively traded with M-ELO. This subset covers 67% of the current M-ELO volume and includes both tick- and nontick-constrained stocks.

Each trading day that the agent experiences consists of 380 (number of symbols) symbol-days. Each symbol-day is discretized into 30 second periods that our agent steps through, meaning there are (number of trading days) $*$ (number of tickers) $*$ ($780 \times \frac{\text{30 sec periods}}{\text{day}}$) episodes— times around the RL Loop— in our training process.

We restrict our agent to take 5 possible actions each 30 second period: it is able to change the timer by [-.5, -.25, 0, +.25, +.5] ms. It does this by evaluating the past 30 seconds of market conditions which are output from our simulation. This includes historically-based information such as features derived from the one- and five-day moving averaged M-ELO historical orders and trades, NBBO historical moments, the continuous book; as well as timer-dependent information such as the resultant fill rate, markout, and metrics derived from agent-made trades and simulated M-ELO activity.

There are 142 features in total, and they make up the state and next state part of the (state, action, reward, next state) experiences that the agent samples at the end of each calendar day to train from. For more technical details regarding the training process, including the pseudo-code, please see Appendix 7.

4.2.1 Motivating Allowable Timer Values

The optimal range of timers was chosen through substantial experimentation. First, we researched the lowest value the timer could take while retaining similar markout to those achieved by the original M-ELO during stable periods. We then found a timer that could effectively mitigate markout across periods of mild instability. For extremely unstable periods (which will be properly defined in Section 4.3), we chose 12ms as our stability protection timer since we found it to be the lowest value that achieves substantial gains over the current M-ELO during high volatility events.

Likewise, the 30 second update cadence was chosen by considering the difficulty of learning at different frequencies and our system's latency requirements. We assert that the theoretical best-possible timer would update with every incoming order. However, training a reinforcement learning agent on every order would be a difficult problem to generalize, resulting in a brittle model. Additionally, despite recent advances in fast electronics, this would be difficult to implement in real-time where nanosecond-latency is required. We

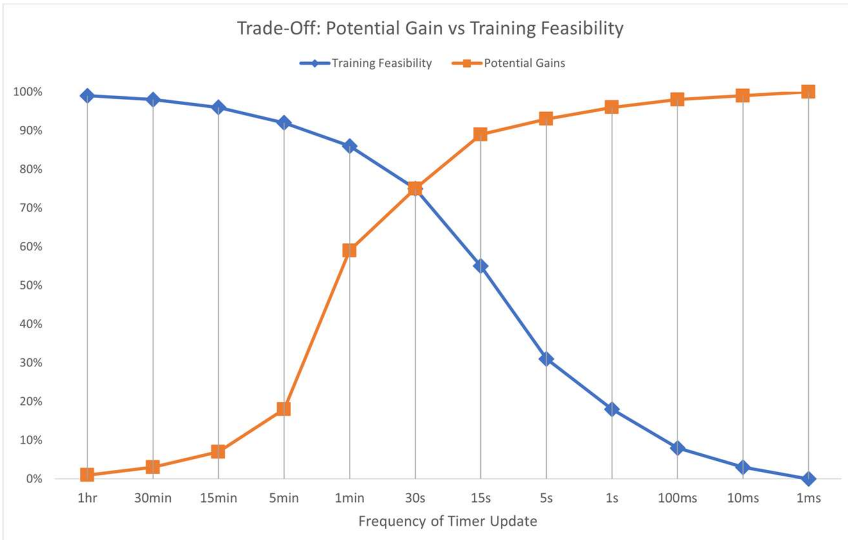

(a) Trade-off between training feasibility and potential upside gains, as a function of the frequency of timer update.

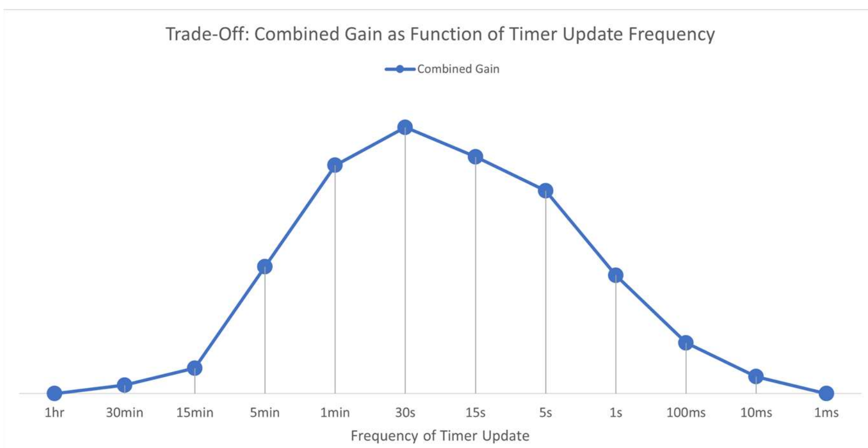

(b) Optimal frequency of timer update, based on the trade-off above.

Figure 8: Choosing the frequency of timer update.

investigated achievable update cadences and found that the point of marginal returns was between 15 and 30 second updates, and decided to go with 30 seconds to give our agent the best chance of learning generalizable actions (Figure 8).

4.2.2 Model Maintenance in Production

We can expect M-ELO environment market dynamics, like any other capital market dynamics, to evolve over time. These sorts of changes account for what is known in statistics as distribution shift. In general, models are not guaranteed to perform well when tested on data that has shifted from the training set.

Given this issue, we implement a form of transfer learning known as retraining to combat the natural performance degradation associated with distribution shift in deep learning. This retraining leverages the idea that training and testing the model on more recent market conditions will yield better results. After all, market conditions of a given day tend to be generally more similar to the the day/week before than they are to previous months.

Retraining involves freezing the weights in the first few layers of our network after a certain point, and allowing the other weights to train on more recent data. Here we commence retraining upon reaching an indicated retrain start date in the simulated environment.

We implement two retraining independent schedules: one for daily retraining and another one for weekly retraining. Both retraining schedules involve two sequential stages. In the first stage, all 35,000+ parameters of the model are trained until the specific retrain start date, in this case January 20, 2022. Then, in the second stage, some of the initial layers of the model are frozen and the end ones are allowed to update based on the most recent data, generating many additional models that are tied to being tested for specific dates. This process is depicted in Figure 9

In the case of the daily retraining, the unfrozen parameters are updated at the end of each subsequent day and saved to be tested on the next day. Likewise, in the case of the weekly retraining, the training process updates the unfrozen parameters throughout days which constitute the subsequent trading week. This model is saved and then tested on the next week. These schedules proceed iteratively until all days or weeks being tested are covered by retrained models.

Beyond the freezing of parameters to allow part of the model to adapt to

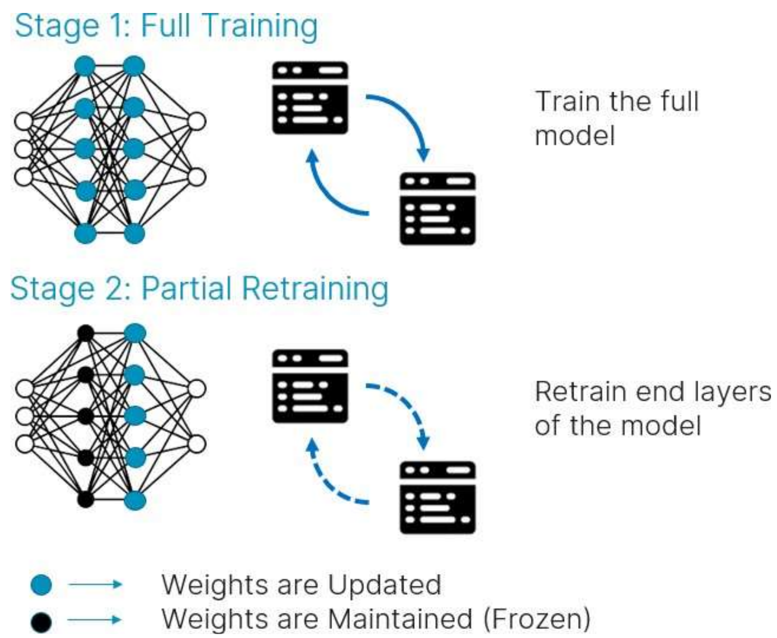

Figure 9: Visual representation of the proposed training and retraining schedule, using a simplified neural network with three input neurons, two hidden layers, and two outputs.

more recent temporally-local conditions, we also recommend the full retraining of the model on a monthly to quarterly basis to ensure it keeps up with longer-scale larger distribution shift.

4.3 Stability Protection Mechanism

Finally, as a system safeguard against volatility, we added a stability-protection mechanism to our timer-set model. In early development, we noted a significant number of one-off adverse markout events during some 30 second periods. This motivated us to design such a mechanism reactive to changes in the NBBO price.

Often these volatility events happen at a much faster frequency than the available timer change cadence, so we developed a method capable of reacting to these events at a more appropriate scale on the order of seconds. We verified that this mechanism improves M-ELO execution quality and

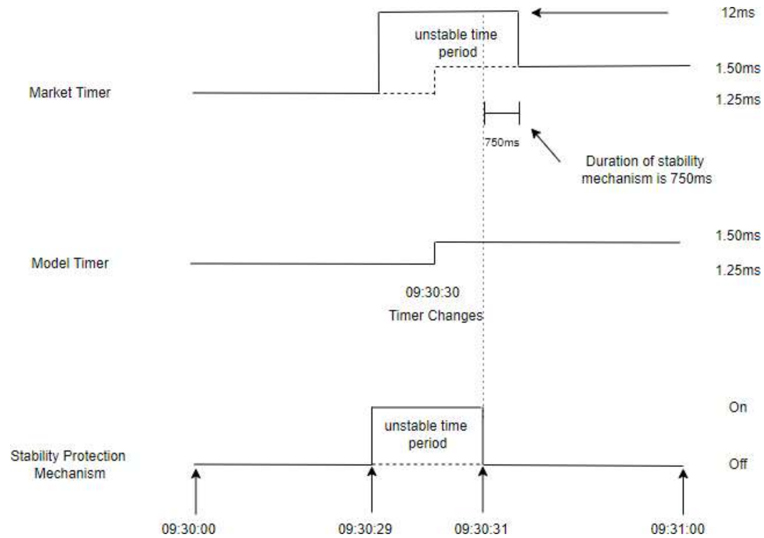

Figure 10: The market timer is set by combining the model timer and the signal from the stability protection mechanism.

performance overall, and protects participants in unstable periods.

Periods which trigger our instability condition (see below) cause the timer to be set to 12ms for 750ms, thereby mitigating negative impact from unfavorable execution. It is important to note that this 750ms duration is independent of the model setting the timer every 30 second period. The market timer is determined by combining the model timer and the stability protection mechanism signal (Figures 11a and 11b). An example of how the stability protection mechanism can improve execution outcomes can be found in Figure 5, from Section 2.2.2

4.3.1 Setting a Stability Threshold

We deem the NBBO unstable if we detect unusually high activity based on the range of NBBO values recorded in the last three seconds. We do this by comparing the range to a threshold for a given ticker and 30 second time interval. The threshold is calculated based on a bisection-quantile method using the previous days' aggregated values (see pseudo-code in 7).

This calculation aims to achieve a certain percent of each day that is

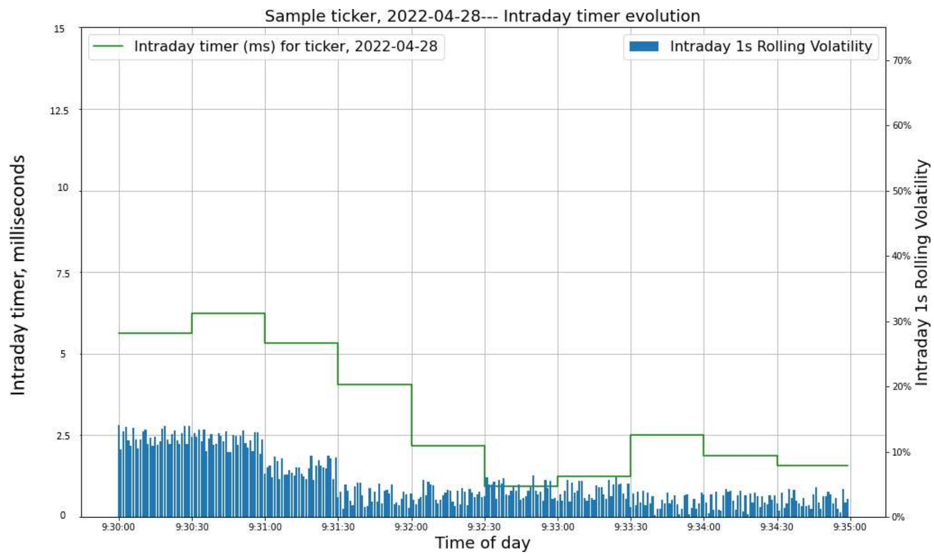

(a) Timer movement given stable NBBO prices.

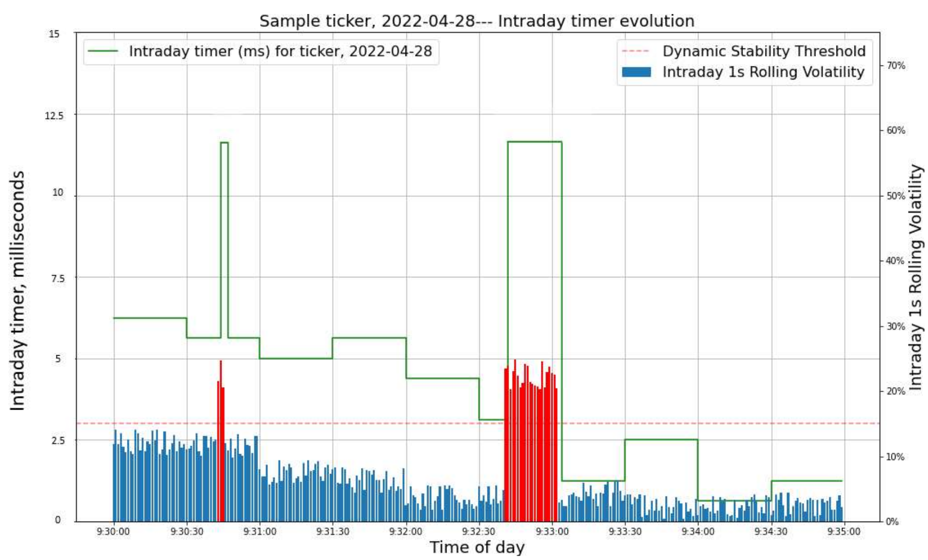

(b) Timer movement given unstable NBBO prices.

Figure 11: Demonstration of stability protection mechanism.

considered unstable. Currently, we aim to consider 1.0% of the day unstable, i.e. to have on average 1.0% of time in between the open at 9:30 AM and close at 4:00 PM be considered unstable. This can also be interpreted as flagging only the top 1.0% most unstable regions to trigger the 12ms condition. This level of protection was chosen through substantial experimentation.

While it is possible to envision a more sophisticated way of implementing a protection mechanism based on the NBBO midpoint changing, we have found this one to be sufficiently performant. Planned future further exploration of different thresholds includes moving to an intraday moving average or some sort of varying timeline-weighted threshold.

5 Behavior of Dynamic M-ELO on the Market

Our Dynamic M-ELO timer demonstrates an average combined volume-weighted improvement of 31.7% over the current static 10ms. This includes a 20.3% increase in fill rate and a 11.4% decrease in markout when compared to the current M-ELO and tested over a quarter-long period from January 1, 2022 to April 1, 2022. This corresponds with our retraining period: every date that is tested uses a model that was strictly trained on previous days' information. For more details, please revisit Section 4.2.2.

During this period, the average number of timer changes per day across all symbols was 140 out of a possible 780 (see Figure 16). This means the model decided to dynamically change the timer about 18% of the time, leaving it at the previous set value about three fifths of the time.

Here we explain how we measure performance gains over the current M-ELO timer, review our aggregate results, and present a firm-level analysis that confirms Dynamic M-ELO does not suffer from systematic-bias. Additionally, we thoroughly assess the interpretability of our model's decisions through a feature explainability study.

5.1 How are Gains Measured?

Our results represent an average combined volume-weighted improvement of 31.1% when compared to the static 10ms M-ELO *as represented in our simulation*. We assert the need to compare against the 10ms M-ELO simulation rather than the historical for two important reasons.

First, as mentioned in Section 3.2, our simulation M-ELO trading environment tends to slightly undershoot trades when compared to historical values, so in order to provide an homogeneous metric for comparison, we measure its performance relative to the simulated baseline. Erring on the side of prudence, we note that, as a direct consequence, this methodology provides a rather conservative estimate of performance, when compared to how it would be if using actual historical values as benchmark.

Second, the agent-set timer often results in orders trading that never could have traded at the static value of 10ms (e.g., by getting cancelled before the 10ms have elapsed, but after a hypothetical shorter holding period). This introduces a source of bias when comparing the aggregated markout the agent achieves in different time periods against the historical aggregated markout. If there are more trades, which is often the case, there is more aggregate markout rendering this comparison unreasonable. Instead, we compare the agent-achieved markout with the synthetic markout from shifting the agent-achieved trades to have occurred at the static 10ms holding period.

Given our baseline as outlined above, the respective improvement gains are calculated as follows:

$$\text{FR Improvement (as a \%)} = \frac{\text{FR}_{agent} - \text{FR}_{sim@10}}{\text{FR}_{sim@10}} \tag{4}$$

$$\text{MO Improvement (as a \%)} = \frac{\text{MO}_{agent} - \text{MO}_{synthetic@10}}{\max(|\text{MO}_{synthetic@10}|, |\text{MO}_{agent}|)} \tag{5}$$

and added together to give the total improvement. Note that the markout improvement calculation differs from a simple percent. We revised this formula for metrics which can vary between positive and negative values. We found it to provide more robust and conservative comparisons.

Regardless, as can be inferred from Equations 4 and 5, the gains for both metrics are measured following the usual standard convention, i.e., a positive value corresponds with an improvement with respect to the baseline.

5.2 Aggregate Backtest Results

To provide a better idea of the distribution of markout, fill rate, and combined improvement across all symbols, we have broken down our aggregate +20.3% fill rate, +11.4% markout gain result into quartiles as shown in the box-and-whisker plot in Figure 12 and in the descriptive statistics in Table 1. The box-and-whisker plot clearly shows that the combined results are right-skewed, since there are more positive than negative outliers in both fill rate and markout improvement. This means that, beyond the net positive average across symbols, the positive combined improvement happens more frequently than negative loss per-symbol.

Additionally, we benchmarked our model's performance against both static and random timers. This was done to assure that the positive increase in fill rate and decrease in markout can be attributed to the model and not to merely lowering the timer or having it switch randomly between allowable values.

Likewise, for the static timers, we see that the Dynamic M-ELO model outperforms all but the lowest allowable static timer in terms of combined

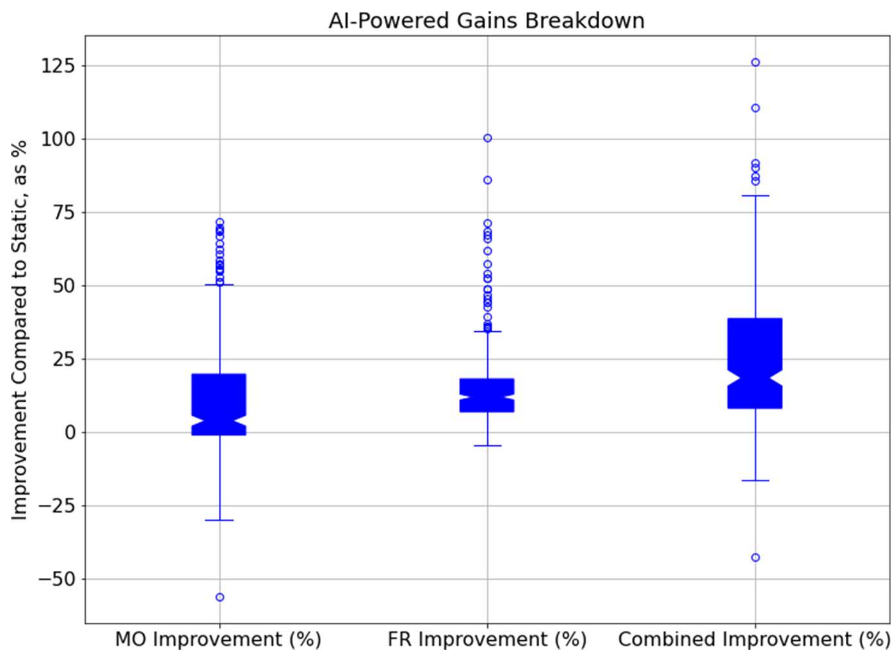

Figure 12: Distribution of fill rate, markout, and combined improvement across tickers. Note that markout and fill rate are treated as independent for their descriptive statistics, so the both column is not just a sum of the other two. For descriptive statistics, see Table 1.

gain, losing only to the .25ms timer by fill rate. However, given the fact that .25ms is the minimum allowed timer value, this actually provides the upperbound for the agent's possible fill rate gain. In terms of markout, Dynamic M-ELO is a clear winner, yielding a substantial net positive markout reduction compared to a static .25ms timer.

5.2.1 Firm-Level Analysis

In order to ensure that our results do not provide an unfair advantage to any specific firm, but rather distribute the gains equitably, we tried to identify patterns and trends that could potentially signify a systematic bias towards specific firms.

	MO	FR	Both
Avg	11.4%	20.3%	31.7%
25^{th} Percentile	-1.1%	6.8%	8.0%
50^{th} Percentile	3.8%	11.8%	18.4%
75^{th} Percentile	19.6%	18.0%	38.6%

Table 1: Distribution of fill rate, markout, and combined improvement across tickers. Note that markout and fill rate are treated as independent for their descriptive statistics, so the both column is not just a sum of the other two. See Figure 12 for a box and whisker plot representation.

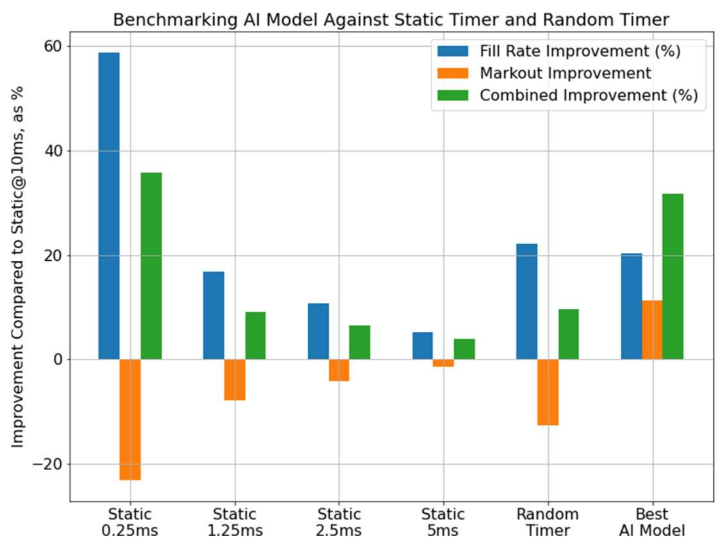

Figure 13: Dynamic M-ELO results compared with relevant static timers and random switching between allowable timer values. All the metrics in this chart are measured using the static version of M-ELO as baseline.

To this end, we performed a detailed analysis on the distribution of AI-driven gains to assess how they particularly impacted each firm that actively participates in M-ELO, regardless of volume. This yielded satisfactory results: no overall trend emerged, confirming that Dynamic M-ELO will not result in systematic-biased execution towards any one firm.

5.3 Explainability Study

Practical deployment of AI systems is not possible without substantial investigation into how they make their decisions. Indeed, garnering trust for the outcome of the system is wholly dependent on an understanding of how the model works. Here we address any concerns arising from explainability in two ways. First, we present a philosophical discussion of how AI decisions differ from human decisions, and how and why this can be beneficial. Then, we hone in on our model and provide an in-depth analysis of how specific inputs map onto our model's outputted timer change recommendations as well as how the complex interactions between inputs impact the timer duration value.

5.3.1 AI Decision Making

Humans make decisions *ex ante*, meaning they are based on our expectations or predictions of the future. Unfortunately, our predictions of the future are rarely correct and are often clouded by a myriad of obscure, confounding factors. We are limited in three main ways: information overload, biases, and decision fatigue. Information overload is the psychological concept that describes how we cannot leverage all available information due to our limited ability to process it; cognitive and emotional biases covers any deviations from strict rationality, which happen more often then we care to admit; and decision fatigue describes how our fallible brains tire after too much use.

AI systems, on the other hand, do not suffer from information overload or decision fatigue. As for cognitive biases— the likes of hindsight bias— and emotional biases— loss/regret aversion and overconfidence — they can be said to not fall prey to these specific examples as well. A broader discussion of how AI systems can be biased by their training sets, especially when their training sets involve demographic information, is out of the scope of this paper.

Furthermore, AI systems make decisions *a priori*, i.e., strictly based on what they have observed from the past. Throughout the training process, AI systems learn a knowledge base from which to make their decisions. This means that, when trained for one specific task, they can become knowledge domain experts and execute on that task much more efficiently than humans ever could aim to.

Our Dynamic M-ELO model is one such system. It builds its knowledge base to control the duration of the hold period timer, and thereby becomes the equivalent of a subject matter expert timer setting system that can be tasked to deal with the complexity of incoming information and the cadence of decision making without fatigue— unlike a human operator. To the extent that our model is temporally biased to its retraining set, this is by design: our model performs better by adapting to local conditions. The data it ingests is reliable— coming straight from the Nasdaq core internal messaging system.

5.3.2 Feature Sensitivity and Feature Interaction Study

Here we make a case for the explainability of our model's decisions given this reliable data. As stated in Section 3.2, our model ingests 142 features that capture market dynamic information about the continuous book and recent M-ELO activity, as well as other information associated with how the timer impacts the simulation environment. We performed both a localized feature sensitivity analysis and a global feature interaction study to motivate our model's timer-set decisions.

Our feature sensitivity analysis involved varying each feature's value and seeing how changing it (while holding all other features constant) directly impacted the trained model's timer duration decision. The results of this study indicate that, out of the 142 features, 27 have been found to be directly correlated with higher (slower) timers, while 25 have been found to be inversely correlated. For the sake of interpretability, we share some of these below. For a more comprehensive analysis of our feature explainability study, please see Table in Section 7.1 (Appendix 7.1).

Features which, as they increase, tend to contribute toward slowing (increasing) the holding period include: increases in the standard deviation of NBBO prices, number of unique firms placing sell orders on M-ELO, and the volume-weighted average NBBO spread. These factors are all associated with higher volatility of the underlying, and Dynamic M-ELO rightly reacts by increasing the timer to try to combat the possibility of momentarily high markout trades.

On the other hand, features which, as they increase, tend to contribute to speeding up (reducing) the holding period include an increase in the median and maximum number of shares per trade and the number of resting M-ELO

(a) (b)

Figure 14: Example 30 second time periods resulting in action 1.0 (a), mapping onto a timer decrease of .25ms, and 4.0 (b), mapping onto a timer increase of .5ms.

bids left. These factors are associated with high fill rate, relatively stable periods, and Dynamic M-ELO reacts by decreasing the timer to try to allow as little friction as possible when trading.

Furthermore, we can analyze the interactions between these features and how they contribute to the model's decisions. For example, an increase in the standard deviation of M-ELO buy-side signed markout and a decrease in the median and maximum number of shares per trade was found to be strongly correlated with an increase in timer. Likewise, a decrease in the average NBBO midpoint and an increase in the NBBO price standard deviation, corresponding with a rapid negative change in price, correlate with an increase in timer.

More complicated interactions can be seen in the Shapley value study below, which uses principles from game theory to calculate the marginal contribution of each feature to a model's output [**4**]. Here we present the ordered top seven influential features that, when combined with all of the other 142 features for a given 30 second time period, give our model's output. It is important to note however, that combining the sum total of the relative impact (Shapley) values yields the maximum Q-value from our model's output and not the actual action itself. This is because the model outputs the Q-values of each action for a given 30 second period, which are argmaxxed to give the model's prescribed relative timer change action. Two examples can be seen in Figure 14b.

6 Conclusion

We have proposed a dynamic-timer modification to the Nasdaq Midpoint Extended Life Order (M-ELO) holding period market that, based on our simulation, achieves an increase in fill rate of 20.3% and a decrease in markout of 11.4% compared to the current static 10ms M-ELO timer.

Our Dynamic Timing system leverages AI to evaluate and determine the duration of the M-ELO holding period in relation to everchanging local market conditions. In doing so, it is able to achieve combinations of fill rate and markout results that were previously unreachable by static timers (see Figure 15). We assert that Dynamic M-ELO has shifted the fill rate-markout Pareto frontier towards more favorable trading execution for all parties involved, representing a meaningful innovation.

These overwhelmingly positive results suggest further opportunity to improve market quality with conditionally-attuned products and controls. In an evolving world with more access to data and computational resources than ever before, innovations like Dynamic M-ELO are not only possible and effective, but also represent opportunities to move towards more dynamic market solutions.

7 Appendix

7.1 Further Analysis

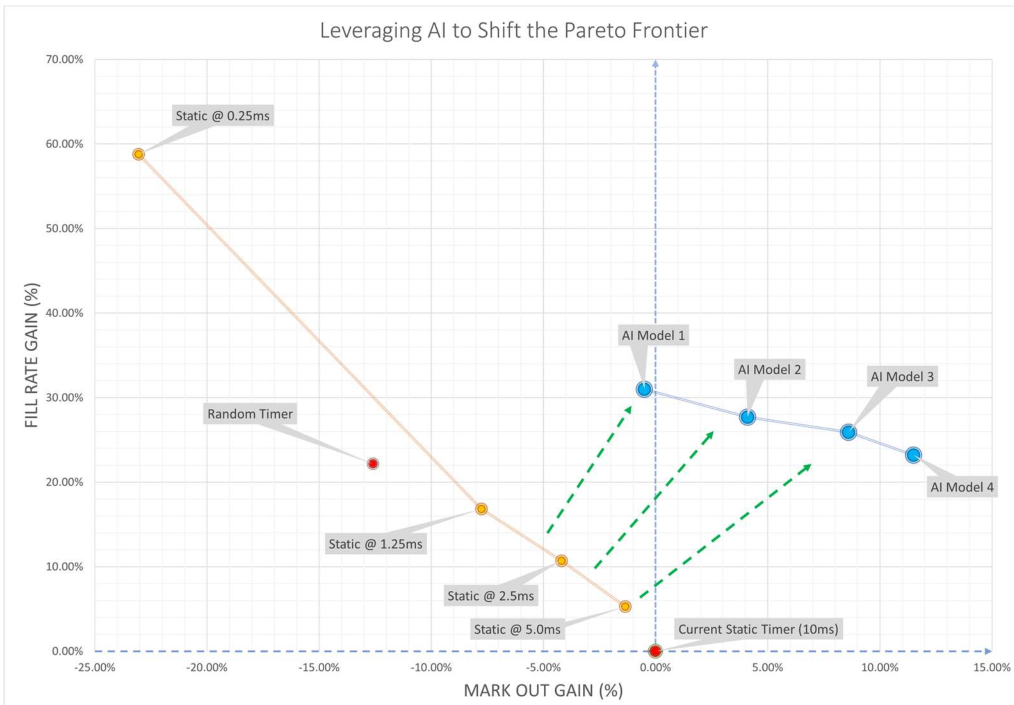

Figure 15: Demonstrated ability of AI to shift Pareto frontier of allowable fill rate and markout values. In particular, the AI models, indicated in blue, significantly outperform an hypothetical static version of M-ELO with timers lower than 10ms, indicated in yellow. The green dashed arrows signify the simultaneous gains driven by the AI system.

Table 2: More comprehensive overview of local feature explainability study.

Relationship with Holding Timer	**Feature**
Strongly Directly Correlated	NBBO price stdev # firms placing M-ELO sell orders Volume-weighted avg NBBO spread
Somewhat Directly Correlated	Max NBBO spread NBBO price level skewness & kurtosis # of resting M-ELO asks M-ELO sell trade qty stdev # unique firms on M-ELO M-ELO buy signed markout, stdev and max Short-term kurtosis of NBBO trade qty
Somewhat Inversely Correlated	# of resting M-ELO bids Shares/trade average Short-term NBBO trade qty skewness NBBO midpoint average # unique firms placing M-ELO buys Proportion of buys in incoming M-ELO orders
Strongly Inversely Correlated	# of shares/trade median and max Timer (tendency towards lower timers when possible)

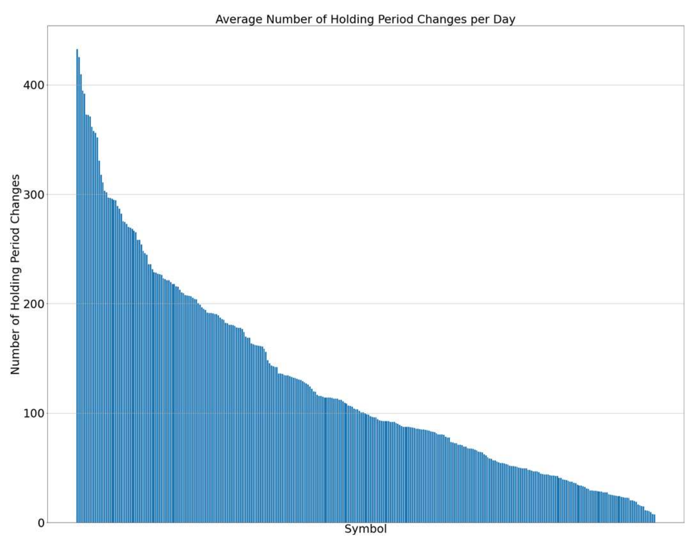

Figure 16: Average timer changes per day by symbol. The maximum number of changes per day is 780, corresponding to the amount of 30 second periods in a trading day.

7.2 Algorithms

7.2.1 D3QN: Modified DDQN Algorithm

Algorithm 1 D3QN: Dynamic M-ELO-DDQN

```
for date d in {date, end date} do
  – – – – – – – – – Begin Agent Experience Collection – – – – – – – – –
    for symbol s in {symbols} do
        for iteration i in 780 × (30 sec periods / 1 day) do
            Gather and store (state, action, reward, next state) experiences
        end for
    end for
  – – – – – – – – – End Agent Experience Collection – – – – – – – – –
  – – – – – – – – – – – Begin Agent Training Loop – – – – – – – – – – –
    for m in number of symbols do
        Sample a batch of (state, action, reward, next state) experiences
        for experience e in {batch} do
            Generate target Q-values
        end for
        Calculate loss L
        Update main model parameters via gradient descent
    end for
    Update target weights to be evolving weighted linear combination of
  main model and target model weights
  – – – – – – – – – – – End Agent Training Loop – – – – – – – – – – –
end for
```

The training process we implemented for Dynamic M-ELO actually differs from the vanilla DDQN algorithm in several ways [**2**]. First, instead of adding all episodes to the agent's memory, we only add specific episodes which have a non-zero reward, i.e. that return a non-zero fill rate and/or markout value. Additionally, instead of using a single memory buffer, we use a multi-buffer, separating experiences by each component of the reward function to ensure balanced training. Finally, instead of sampling from the buffer at the end of each episode only once, we sample from the buffer a number of times corresponding to the product of the number of ticker times and the number of steps in each training iteration (at the end of every symbol-day).

For further details, please see annotated pseudo-code above, which is broken up into two distinct phases— agent experience collection and agent training loop— for ease of understanding.

7.2.2 Stability Coverage Algorithm

Algorithm 2 Quantile-Bisection Stability Coverage Method

```
for symbol s in {symbols} do
    Decide on optimal coverage c*
    Initialize naive price range lower bound l and upper bound u
    Compute midpoint m between l and u
    Calculate coverage c for m
    while (l, u) do not offer optimal coverage c* do
        if c < c* then
            Replace the value of u with m
        else
            Replace the value of l with m
        end if
        Update m to be the midpoint of the new l and u
        Update c to be the coverage of the new midpoint m
    end while
    Return m as optimal threshold
end for
```

References

[1] Mnih, V., Kavukcuoglu, K., Silver, D. et al. *Human-level control through deep reinforcement learning*. Nature 518, 529-533 (2015). 10.1038/nature14236.

[2] Hasselt, H. *Deep Reinforcement Learning with Q-Learning*. Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16) 23, 2094-2100 (2016). 10.5555/3016100.3016191

[3] Sutton, R.S. and Barto, A.G. *Reinforcement Learning: An Introduction*. MIT Press, 2018. ISBN: 9780262193986.

[4] Lundberg, S.M. and Lee, S. *A Unified Approach to Interpreting Model Predictions*. Proceedings of the 31st Neural Information Processing Systems Conferece (NeurIPS-17).

[5] Lehalle, C. and Laruelle, S. *Market Microstructure in Practice*. World Scientific Publishing Co. Pte. Ltd., 2018. ISBN: 9789813231122.

[6] Hasbrouck, J. *Empirical Market Microstructure*. Oxford University Press, 2007. ISBN: 9780195301649.

[7] O'Hara, M. *Market Microstructure Theory*. Blackwell Publishing, 1995. ISBN: 9781557864437.

EXHIBIT 5

Deleted text is [bracketed]. New text is underlined.

The Nasdaq Stock Market LLC Rules

* * * * *

Equity Rules

* * * * *

Equity 4: Equity Trading Rules

...

4702. Order Types

(a) No change.

(b) Except where stated otherwise, the following Order Types are available to all Participants:

(1) – (13) No change.

(14) (A) A "Midpoint Extended Life Order" is an Order Type with a Non-Display Order Attribute that is priced at the midpoint between the NBBO and that will not be eligible to execute until a minimum time period [of 10 milliseconds]("Holding Period") has passed after acceptance of the Order by the System. Eligible Midpoint Extended Life Orders may only execute against other eligible Midpoint Extended Life Orders and M-ELO+CB Orders. Buy (sell) Midpoint Extended Life Orders will be ranked in time order at the midpoint among other Buy (Sell) Midpoint Extended Life Orders and buy (sell) MELO+ CB Orders. A Midpoint Extended Life Order may be cancelled at any time. If a Midpoint Extended Life Order is modified by a member (other than to decrease the size of the Order or to modify the marking of a sell Order as long, short, or short exempt) during the Holding Period, the System will restart the Holding Period. If a Midpoint Extended Life Order is modified by a member (other than to decrease the size of the Order or to modify the marking of a sell Order as long, short, or short exempt) after it is eligible to execute, the Order will have to satisfy a new Holding Period to become eligible to execute.

At the commencement of Market Hours, the initial Holding Period for a Midpoint Extended Life Order in a particular symbol will be 1.25 milliseconds. However, the System may, based upon its proprietary assessment of market conditions for that symbol, decide to vary the default Holding Period of all Midpoint Extended Life Orders in a symbol in increments of 0.25 or 0.50 milliseconds, and within an overall range of between 0.25 and 2.50 milliseconds during normal market conditions (described below), beginning 30 seconds after Market Hours commences and then at 30 second intervals thereafter until Market Hours conclude (“Change Events”). Whenever a Midpoint Extended Life Order in a symbol enters the Exchange Book during Market Hours, it will adopt the then-prevailing Holding Period that the System chose for all Midpoint Extended Life Orders in that symbol as of the immediately preceding Change Event. The Holding Period for the Midpoint Extended Life Order will not be eligible for modification until the next scheduled Change Event occurs for all Midpoint Extended Life Orders in that symbol, at which point all pending Midpoint Extended Life Orders in a symbol with unexpired Holding Periods will adopt the modifications that the System then makes to the Holding Period (retroactive to the time of acceptance of the Midpoint Extended Life Orders).

Notwithstanding the above, whenever the System determines that market conditions for a symbol have become extraordinarily unstable (including in between Change Events), the System will activate a stability protection mechanism. The stability protection mechanism will override the prevailing Holding Period for a Midpoint Extended Life Order in a symbol experiencing extraordinary instability and immediately increase the duration of the Holding Period to 12 milliseconds for a period of at least 750 milliseconds (the "Stability Protected Period"). The System may activate the stability protection mechanism even between Change Events. The System will evaluate, at each NBBO update, whether market conditions remain extraordinarily unstable and, if so, it will restart the 750ms Stability Protected Period and maintain the 12ms Holding Period until conditions stabilize. Once the System determines that market conditions have stabilized (i.e., all measurements for the symbol are at or below the threshold value throughout the duration of the prevailing Stability Protected Period), the System will revert the duration of the Holding Periods to that which prevailed as of the Change Event that occurred immediately prior to the activation of the stability protection mechanism or, if the stability protection mechanism was active when a Change Event occurred, to the duration selected at the immediately preceding Change Event. The System will then proceed to reevaluate the duration of the Holding Periods as per the regular schedule of Change Events. For purposes of this Rule, the System determines that "extraordinary instability" for a symbol exists through observations it makes following every change in the NBBO for that symbol that occurs during the trading day. When the NBBO changes, the System looks back at the prior three seconds of trading and measures the difference between the highest and the lowest NBBO midpoint values that occurred during that period, and then it compares that measurement to a threshold value for the symbol. The System concludes that extraordinary instability exists for a symbol if the measurement exceeds the threshold value. The threshold value for a symbol, in turn, is the difference between the highest and the lowest NBBO midpoint values for the symbol that, if applied to its trading activity during the prior trading day, would have caused the System to deem trading in the symbol to be extraordinarily unstable for as close to one percent of that day as possible.

If a limit price is assigned to a Midpoint Extended Life Order, the Order will be: (1) eligible for execution in time priority if upon acceptance of the Order by the System, the midpoint price is within the limit set by the participant; or (2) held until the midpoint falls within the limit set by the participant at which time the Holding Period will commence and thereafter the System will make the Order eligible for execution in time priority. For example, if the Best Bid was $11 and the Best Offer was $11.06, the price of the Midpoint Extended Life Order would be $11.03. If a participant enters a Midpoint Extended Life Order to buy with a limit of $11.02, the Holding Period would not begin until the midpoint price reached $11.02. If a Midpoint Extended Life Order has met the Holding Period requirement but the midpoint is no longer within its limit, it will nonetheless be ranked in time priority among other Midpoint Extended Life Orders and M-ELO+CBs if the NBBO later moves such that it is within the Order's limit price. Midpoint Extended Life Orders will not execute if there is a resting non-displayed Order priced more aggressively than the midpoint between the NBBO, and will be held for execution until the resting non-displayed Order is no longer on the Nasdaq Book or the midpoint of the NBBO matches the price of the resting non-displayed Order.

Midpoint Extended Life Orders in existence at the time a halt is initiated will be ineligible to execute and held by the System until trading has resumed and the NBBO has been received by Nasdaq. Upon resumption of trading in a halted symbol, any new Midpoint Extended Life Order in that symbol and any pending Midpoint Extended Life Order in that symbol with an unexpired Holding Period will be subject to a 12 milliseconds Holding Period (running from the time when trading resumes) until the next scheduled Change Event, at which point the System may determine to adjust that Holding Period to a duration within the range applicable under normal market conditions. If, however, the System determines that extraordinary instability in the symbol exists, it may instead determine to activate the stability protection mechanism and maintain the duration of the Holding Period at 12 milliseconds for another 750 milliseconds. Prior to commencement of a new 12 millisecond Holding Period for a new or pending M-ELO or M-ELO+CB following a Halt, the System will first determine whether the M-ELO or M-ELO+CB is or remains eligible for execution. That is, the Holding Period will commence only if, upon commencement of trading following the Halt, the midpoint price for the Order is within the limit set by the participant. If not, the System will hold the Order until the midpoint falls within the limit set by the participant, at which time the 12 millisecond Holding Period will commence.

Nasdaq will publish on Nasdaqtrader.com weekly aggregated number of shares and transactions of Midpoint Extended Life Orders executed on Nasdaq by security. The weekly aggregated data would be published with a delay of two weeks for NMS stocks in Tier 1 of the NMS Plan to Address Extraordinary Market Volatility, and four weeks for all other NMS stocks. Nasdaq will also publish on Nasdaqtrader.com monthly aggregated block-sized trading statistics of total shares and total transactions of Midpoint Extended Life Orders executed on Nasdaq. A transaction would be considered "block-sized" if it meets any of the following categories of criteria: (1) 10,000 or more shares; (2) $200,000 or more in value; (3) 10,000 or more shares and $200,000 or more in value; (4) 2,000 to 9,999 shares; (5) $100,000 to $199,999 in value; or (6) 2,000 to 9,999 shares and $100,000 to $199,999 in value. For each of these categories, Nasdaq will publish monthly transaction count and share executed volume information. The data will be published no earlier than one month following the end of the month for which trading was aggregated.

(B) The following Order Attributes may be assigned to a Midpoint Extended Life Order:

• Minimum Quantity.

• Size.

• Time-in-Force. Regardless of the Time-in-Force entered, a Midpoint Extended Life Order may not be active outside of Market Hours. A Midpoint Extended Life Order entered during Pre-Market Hours will be held by the System in time priority until Market Hours. Midpoint Extended Life Orders entered during Post-Market Hours will not be accepted by the System. A Midpoint Extended Life Order remaining unexecuted after 4:00 p.m. ET will be cancelled by the System. If a Midpoint Extended Life Order is entered with a Time-In-Force of IOC, it will execute against available eligible resting interest immediately upon the expiration of the Holding Period; if no such resting interest is available, or shares of the Order remain unexecuted after executing against eligible resting interest, then the System will automatically cancel the Order or the remaining shares of the Order, as applicable; if the Order is ineligible to begin the Holding Period upon entry, then the System will cancel it immediately.

• Non-Display. All Midpoint Extended Life Orders are Non-Displayed.

(15) A "Midpoint Extended Life Order Plus Continuous Book" or "M-ELO+CB" is an Order Type that has all of the characteristics and attributes of a Midpoint Extended Life Order, as set forth above in subparagraph (14), except as follows:

• A M-ELO+CB that satisfies the Holding Period shall be eligible to execute (at the midpoint of the NBBO) against other eligible M-ELO+CBs, eligible Midpoint Extended Life Orders, and as described below, Non-Displayed Orders with Midpoint Pegging and Midpoint Peg Post-Only Orders (collectively, "Midpoint Orders") resting on the Exchange's Continuous Book. A M-ELO+CB shall be eligible to execute against a Midpoint Order if: (i) the Midpoint Order has the [Midpoint]Trade Now Attribute enabled; (ii) no other order is resting on the Continuous Book that has a more aggressive price than the current midpoint of the NBBO; (iii) the Midpoint Order has rested on the Exchange's Continuous Book for a minimum of [10 milliseconds]<u>the duration of the Holding Period then applicable to a M-ELO+CB in that symbol </u>after the NBBO midpoint falls within the limit set by the participant; and (iv) the Midpoint Order satisfies any minimum quantity requirement of the M-ELO+CB. A buy (sell) M-ELO+CB will be ranked in time order at the midpoint among other buy (sell) M-ELO+CBs, buy (sell)

Midpoint Extended Life Orders, and buy (sell) Midpoint Orders, as of the time when such Orders become eligible to execute.

- QIX is not available for the entry of a M-ELO+CB.

- Nasdaq will include M-ELO+CB executions in the statistical information it publishes on Nasdaqtrader.com for M-ELOs.

* * * * *