Required fields are shown with yellow backgrounds and asterisks.

Page 1 of * 63

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

File No. * SR 2024 - * 112
Amendment No. (req. for Amendments *)

Filing by NYSE Arca, Inc.

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 *).

Proposal to amend Rule 7.31-E(h)(3) to adopt the Selective Midpoint Order

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 * Le-Anh | Last Name * Bui

Title * Senior Counsel, NYSE Group Inc.

E-mail * Le-Anh.Bui@ice.com

Telephone * (202) 661-8953 | Fax (212) 656-8101

Signature

Pursuant to the requirements of the Securities Exchange of 1934, NYSE Arca, Inc.
has duly caused this filing to be signed on its behalf by the undersigned thereunto duly authorized.

Date 12/18/2024

By Martha Redding
(Name *)

(Title *) Corporate Secretary

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.

Martha Redding
Digitally signed by Martha Redding
Date: 2024.12.18 11:19:31 -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

19b-4 - NYSE Arca Selective Midpoint

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

Ex. 1 NYSE Arca Selective Midpoint C

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

Ex. 3 - Selective Midpoint Indicator Wh

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

Ex. 5 - NYSE Arca Selective Midpoint

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) Pursuant to the provisions of Section 19(b)(1) of the Securities Exchange Act of 1934 (the "Act")[1] and Rule 19b-4 thereunder,[2] NYSE Arca, Inc. ("NYSE Arca" or the "Exchange") proposes to amend Rule 7.31-E to adopt the Selective Midpoint Order.

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

(b) The Exchange does not believe that the proposed rule change will have any direct effect, or any significant indirect effect, on any other Exchange rule in effect at the time of this filing.

(c) Not applicable.

2. Procedures of the Self-Regulatory Organization

Senior management has approved the proposed rule change pursuant to authority delegated to it by the Board of the Exchange. No further action is required under the Exchange's governing documents. Therefore, the Exchange's internal procedures with respect to the proposed change are complete.

The person on the Exchange staff prepared to respond to questions and comments on the proposed rule change is:

Le-Anh Bui
Senior Counsel
NYSE Group, Inc.
(202) 661-8953

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 Rule 7.31-E(h)(3) to decommission the Discretionary Pegged Order ("DPO") and introduce the Selective Midpoint Order ("SeMi Order"). The SeMi Order will be similar to the DPO in that it is a discretionary order type, but will, unlike the current DPO, provide price protection during periods of market instability based on input from a gradient-boosting machine learning model.

---

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

[2] 17 CFR 240.19b-4.

Background

The Exchange currently offers the DPO, which is a non-displayed order to buy (sell) that is pegged to the same side of the PBBO and assigned a working price equal to the lower (higher) of the midpoint of the PBBO (the "Midpoint Price") or the limit price of the order.[3] Any untraded shares of such order are assigned a working price equal to the lower (higher) of PBB (PBO) or the order's limit price, which is automatically adjusted in response to changes to the PBB (PBO) for buy (sell) orders up (down) to the order's limit price. A DPO will exercise the least amount of discretion necessary from its working price to its discretionary price (defined as the lower (higher) of the Midpoint Price or the limit price of the order) to trade with contra-side interest.

Prior to November 2022, the DPO would not exercise discretion if the PBBO was determined to be unstable via a "quote instability calculation" that assessed the probability of a change to the PBB or PBO.[4] The Exchange used the quote instability calculation along with real-time relative quoting activity of protected quotations to assess the probability of an imminent change to the PBBO (the "quote instability factor"). When the quoting activity met certain predefined criteria and the quote instability factor calculated was greater than the Exchange's predefined quote instability threshold, the Exchange treated the quote as unstable and restricted DPOs from exercising discretion. In November 2022, the Exchange amended Rule 7.31-E(h)(3) to eliminate the quote stability calculation, allowing DPOs to exercise discretion even during potential periods of quote instability.[5]

Proposed Rule Change

The Exchange proposes to modify Rule 7.31-E(h)(3) to replace the DPO with the SeMi Order. The SeMi Order will share the same basic attributes as the DPO. Like the DPO, the SeMi Order would be a non-displayed order to buy (sell) that is pegged to the same side of the PBBO and assigned a working price equal to the

---

3 See Rule 7.31-E(h)(3). As defined in NYSE Arca Rule 1.1, "PBBO" means the Best Protected Bid and the Best Protected Offer. Rule 1.1 also defines "PBB" as the highest Protected Bid and "PBO" as the lowest Protected Offer.

4 See Securities Exchange Act Release No. 95154 (June 24, 2022), 87 FR 39134 (June 30, 2022) (SR-NYSEArca-2022-13) (Notice of Filing of Amendment No. 2 and Order Granting Accelerated Approval of a Proposed Rule Change, as Modified by Amendment No. 2, To Amend Rule 7.31-E(h)(3) Relating to Discretionary Pegged Orders).

5 See Securities Exchange Act Release No. 96322 (November 15, 2022), 87 FR 69376 (November 18, 2022) (SR-NYSEARCA-2022-76) (Notice of Filing and Immediate Effectiveness of Proposed Rule Change to Amend Rule 7.31-E). The Exchange resumed offering the DPO in November 2022, after previously filing to temporarily suspend its use in August 2022. See Securities Exchange Act Release No. 95584 (August 23, 2022), 87 FR 52826 (August 29, 2022) (SR-NYSEARCA-2022-54) (Notice of Filing and Immediate Effectiveness of Proposed Rule Change To Amend Rule 7.31-E).

lower (higher) of the Midpoint Price or the limit price of the order. Any untraded shares of a SeMi Order would be assigned a working price equal to lower (higher) of the PBB (PBO) or the order's limit price and automatically adjusted in response to changes to the PBB (PBO) for buy (sell) orders up (down) to the order's limit price. In order to trade with contra-side orders on the NYSE Arca Book, a SeMi Order to buy (sell) would exercise the least amount of price discretion necessary from its working price to its discretionary price, which is defined as the lower (higher) of the Midpoint Price or the SeMi Order's limit price.

SeMi Orders would not be displayed, must be designated Day, and would be eligible to be designated for the Core Trading Session only. SeMi Orders designated for the Early Trading Session or Late Trading Session would be rejected.

When exercising discretion, SeMi Orders (like DPOs today) would maintain their time priority at their working price as Priority 3 - Non-Display Orders and are prioritized behind Priority 3 - Non-Display Orders with a working price equal to the discretionary price of a SeMi Order at the time of execution. If multiple SeMi Orders are exercising price discretion during the same book processing action, they would maintain their relative time priority at the discretionary price.

Accordingly, the Exchange proposes to replace existing references to the "Discretionary Pegged Order" in Rule 7.31-E(h)(3) and subparagraphs (A) and (B) thereunder with references to the "Selective Midpoint Order."[6]

The Exchange proposes to add new Rule 7.31-E(h)(3)(C) to provide that the SeMi Order, as proposed, would exercise price discretion to its discretionary price, except during periods of quote instability as identified by the Selective Midpoint Indicator ("SMI") (as discussed in further detail below). If the SMI determines the PBBO for a particular security to be an unstable quote, both an arriving and resting SeMi Order would wait for a PBBO that is stable before the order's working price is adjusted and the order becomes eligible to trade. In other words, a SeMi Order would be ineligible to trade when the SMI identifies unstable market conditions and would remain in that state until the SMI determines that market conditions have stabilized, thereby preventing potentially undesirable executions during volatile or unstable market conditions. Whereas the DPO previously relied on one static logistical regression model to forecast market instability and prevented DPOs in any symbol from exercising discretion to trade when the model anticipated an unstable market, the SeMi Order, as proposed, would rely on the SMI to predict market instability using a symbol-specific gradient-boosting machine learning model and would protect SeMi Orders from trading when the SMI predicts quote instability for a given symbol.

---

[6] The Exchange also proposes conforming changes in Rules 7.18-E(b)(1), 7.18-E(c)(1) and (5), 7.31-E(i)(3), and 7.34-E(c)(1)(A) to replace references to "Discretionary Pegged Orders" with references to "Selective Midpoint Orders."

The Exchange also proposes new Rule 7.31-E(h)(3)(D) to provide for SeMi Orders to be optionally designated as Liquidity Providing. Proposed Rule 7.31-E(h)(3)(D)(i) would provide that a SeMi Order designated as Liquidity Providing will only execute on arrival against resting orders that include a Non-Display Remove Modifier and are priced within the Liquidity Providing SeMi Order's discretionary range.[7]

Proposed Rule 7.31-E(h)(3)(D)(ii) would provide that if a resting contra-side order that does not include a Non-Display Remove Modifier is priced within an arriving Liquidity Providing SeMi Order's discretionary range, the Liquidity Providing SeMi Order will be placed on the NYSE Arca Book, and its discretionary range will be adjusted to equal the resting price of a non-displayed contra-side order or to one MPV less aggressive than the resting price of a displayed contra-side order.[8]

Proposed Rule 7.31-E(h)(3)(D)(iii) provides that a Liquidity Providing SeMi Order resting on the NYSE Arca Book will not trade with an arriving contra-side order that cannot remove liquidity.[9] Once such arriving contra-side order is placed on the NYSE Arca Book, the discretionary range of the Liquidity Providing SeMi Order will be adjusted to equal the resting price of a non-displayed contra-side order or to one MPV less aggressive than the resting price of a displayed contra-side order.

---

7 The Exchange also proposes related conforming changes to Rules 7.31-E(d)(2)(B), 7.31-E(d)(3)(F), and 7.31-E(e)(1)(C) to provide that Non-Displayed Limit Orders, MPL Orders, and Non-Routable Limit Orders designated with a Non-Display Remove Modifier will trade as takers against Liquidity Providing SeMi Orders.

8 The Exchange notes that allowing Liquidity Providing SeMi Orders to trade with resting orders with a Non-Display Remove Modifier, as well as adjusting the discretionary range of such orders, would be consistent with how similar discretionary order types function on other equities exchanges. See, e.g., Cboe EDGX Exchange, Rule 11.8(g)(5) ("[Midpoint Discretionary Orders ("MDOs")] that are not entered with a [Quote Depletion Protection ("QDP")] instruction, as defined in Rule 11.8(g)(10), will only act as the liquidity provider. MDOs entered with a QDP instruction will instead be allowed to remove liquidity, by default, unless the User chooses to require that the MDO only act as a liquidity provider. If the instructions included on an MDO do not permit the MDO to remove liquidity, it will only execute on entry against resting orders that include a Super Aggressive instruction priced at the MDO's pegged price if the MDO also contains a Displayed instruction, and against resting orders that include an NDS instruction priced either at the MDO's pegged price or within its discretionary range. If a resting contra-side order that does not include [a Non-Displayed Swap] instruction is priced within the discretionary range of an incoming MDO that is not permitted to remove liquidity, the incoming MDO will be placed on the EDGX Book and its discretionary range will be shortened to equal the limit price of the resting contra-side order. Likewise, where an incoming order with a Post Only instruction does not remove liquidity on entry pursuant to Rule 11.6(n)(4) against a resting MDO, the discretionary range of the resting MDO will be shortened to equal the limit price of the incoming contra-side order with a Post Only instruction.").

9 The Exchange notes that this proposed handling is also consistent with the handling of similar discretionary order types by other equities exchanges. See id.

Proposed Rule 7.31-E(h)(3)(D)(iv) would provide that, once resting on the NYSE Arca Book, the discretionary range of a Liquidity Providing SeMi Order will be adjusted based on resting contra-side interest as described in proposed subparagraphs (ii) and (iii) of this Rule when its working price Changes. In addition, proposed Rule 7.31-E(h)(3)(D)(iv)(a) and (b) provide that a Liquidity Providing SeMi Order to buy (sell) will not be eligible to trade at a price equal to or above (below) any sell (buy) orders that are displayed and that have a working price equal to or below (above) the working price of such Liquidity Providing SeMi Order, or at a price above (below) any sell (buy) orders that are not displayed and that have a working price below (above) the working price of such Liquidity Providing SeMi Order.

Finally, the Exchange proposes to renumber current Rule 7.31-E(h)(3)(C) as 7.31-E(h)(3)(E). The Exchange proposes that, as with DPOs, if the PBBO is locked or crossed, SeMi Orders would wait for a PBBO that is not locked or crossed before the working price is adjusted and the order becomes eligible to trade. Accordingly, the Exchange proposes to substitute "Selective Midpoint Order" for "Discretionary Pegged Order" in the text of proposed Rule 7.31-E(h)(3)(E).

Selective Midpoint Indicator

As described in further detail in the white paper attached as Exhibit 3 to this proposed rule change,[10] the SMI was developed using a decision tree model, which is a supervised learning algorithm used for classification and regression tasks. A decision tree model is trained sequentially, and each successive model seeks to improve on errors in the previous model by focusing on accurately predicting where the previous model performs poorly.

The Exchange determined to use a decision tree model in developing the SMI in large part because of the transparency it offers to its designers. The output of a decision tree is a hierarchy of questions that allows a user to follow the model's decision-making process, assess the importance of a given feature to the model's output, and examine the reasons underlying a specific output from the model. The Exchange believed it was important to build a model that could be interpreted and understood in this way to allow for the evaluation of, among other things, the relationships between market shifts, feature selection, and feature weightings, as well as to be able to assess overall model performance.

Market Instability Assessment

The Exchange set out to develop a model that would allow the SMI to predict market instability. For purposes of the model, the Exchange defined "instability"

---

[10] See Exhibit 3, Selective Midpoint Indicator (SMI): A Gradient Boosted Signal Enabling Stable Order Executions (the "White Paper").

at a high level as relatively large price moves during a relatively short time frame. The Exchange chose PBBO updates as the fundamental data points for the model based on the evenly distributed nature of PBBO updates throughout the trading day and the granular level of information such updates offer. The model's objective is to identify windows where changes to the PBBO are unstable, in order to predict unstable markets in real time and prevent SeMi Order trading accordingly.

The Exchange established the concept of a "price jump" to further understand and categorize periods of instability. A price jump is defined as a PBBO mid-price change of a pre-defined percentage of a symbol's spread (referred to as the spread threshold ***X***), in either direction, within a configurable time interval (referred to as the time horizon ***G***). The Exchange sought to identify price jumps following a PBBO update, whether positive or negative. A price jump is marked at a given point in time if, looking back to the start of the time horizon, the mid-price was at least the spread threshold in difference from the current mid-price. This approach identifies discrete price jumps over configurable time intervals, thereby offering the Exchange flexibility to adjust the spread threshold and/or time horizon parameters according to symbol-specific dynamics.[11]

To build the SMI as a continuous signal of market instability, the Exchange next applied the price jump definition to delineate continuous periods of market instability. At each PBBO update (at index time *i*), the Exchange determined whether there had been a price jump. If there was no price jump, the market would be considered stable. If there was a price jump, the Exchange continued to look for price jumps within the time horizon ***G***, until no price jump occurred. At that point, the Exchange would mark an unstable time window starting at the PBBO update that occurred at *i*-1 or 50 microseconds prior to the PBBO update at time *i*, whichever is closest to *i*, and ending at the last price jump identified.[12]

Finally, the Exchange employed an additional parameter for a minimum time increment ***g*** to establish that unstable windows identified using the spread threshold and time horizon parameters are indicative of persistent price instability. If multiple price jumps occur within a small timeframe, such that the distance between the first and last price jumps is less than ***g***, the Exchange would not mark the window as unstable. The application of the minimum time increment parameter allows the Exchange to more accurately identify true periods of instability by filtering out temporary price jumps that are quickly followed by a reversion to the price prior to the observed price jump.[13]

---

11 Figure 1 in Section 4.1 of the White Paper provides a detailed example of how a price jump is calculated.

12 Figure 2 in Section 4.2 of the White Paper provides an example of the identification of unstable windows.

13 Figure 3 in Section 4.2 of the White Paper demonstrates the operation of the minimum time increment parameter.

In developing the SMI to facilitate the SeMi Order's ability to provide protection against potentially unfavorable executions, the Exchange also wanted to be able to differentiate between quote instability on the bid side and ask side. For example, the Exchange would not want to prevent executions of buy orders based on upwards instability of the mid-price of the PBBO because executions under those conditions would likely be favorable. Accordingly, if the mid-price of the PBBO is higher at the end of an unstable window, the Exchange would only mark PBBO updates during such window as unstable for the ask side (and vice versa for the bid side).

Model Development

To develop and train the models underlying the SMI, the Exchange used data from the NYSE Arca Book from August 29, 2024 through October 22, 2024 for a set of 500 symbols selected to reflect a representative sample of the U.S. equity markets.[14] The Exchange created evaluation data points for each PBBO update that include the closest book depth state at the time of that update. This merged data set of PBBO updates and book depth data has the benefit of the information from both types of data without overinflating the size of the data set. All features and parameters used by the models are calculated based on this merged data set.

The Exchange identified a set of features based on NYSE Arca Book data—e.g., book depth information, PBBO updates, number of IOC orders—that would contribute to the models. The Exchange selected 83 unique features that could be considered for incorporation into the model.[15] These features were identified based on an iterative process, and features were selected based on their ability to explain unstable periods identified by the application of the parameters discussed above and, going forward, to predict market instability on a real-time basis so that the SMI can effectively protect SeMi Orders from potentially unfavorable executions. In production, the Exchange proposes to select a subset of these 83 features to be weighted in the model's assessment of market instability (rather than having the model utilize all available features), because focusing on the features that have been identified as most important in predicting market instability for a given symbol would both optimize prediction accuracy and processing speed.[16]

---

14 The full symbol list is included in Appendix B of the White Paper. Symbols were chosen based on criteria including absolute price level, spread in dollars, spread in basis points, and liquidity (daily ADV). Because stable periods generally far outnumber unstable periods in the U.S. equity markets, the Exchange used under-sampling methods where appropriate to reduce the number of stable data points in the data set and randomly shuffled data before training.

15 The full list of features is included in Appendix A to the White Paper. Appendix A also identifies the subset of features that were selected for the model training described in the White Paper.

16 Once the SMI is implemented in production, the full list of features currently calculated in real-time and available for evaluation for inclusion in the SMI models will be published daily on the Exchange's website.

Symbol-Specific Models

The Exchange proposes that the SMI will rely on symbol-specific models to leverage the ability of the models to incorporate different features and weightings to respond to individual symbols' unique profiles (e.g., the features most likely to accurately predict instability for a given symbol). To explore how symbol-specific models should be distributed, the Exchange ranked each of the 500 symbols in the representative data set according to the total number of unstable data points. The Exchange trained a model for each of the 500 symbols and used the model trained for the symbol SPY as a default "market model."

The Exchange found that the SMI performed better with a symbol-specific model for more active symbols (i.e., those with more unstable data points) than on the market model (i.e., testing showed sharp declines in precision and increases in overlocking behavior using the market model), whereas the SMI with a market model performed well for less active symbols (i.e., high recall with relatively small loss of precision and minimally more overlocking). The Exchange concluded that, for the SMI to provide optimal information and protection to SeMi Orders, more active symbols would benefit significantly from symbol-specific models, while less active symbols (which have fewer unstable data points to inform a symbol-specific model) could successfully default to a market model.

The Exchange believes that models tailored to individual symbols' specific characteristics would provide for better performance by the SMI and thus enhanced price protection by the SeMi Order. Accordingly, in production, the Exchange anticipates that at least 200 (and up to approximately 1,000) symbols that trade on the Exchange will have an individualized model that incorporates features that have been specifically identified for predicting market stability for that symbol. The remaining symbols that trade on the Exchange would use the market model, which would apply the same features and weightings for all symbols.[17] Each day, the Exchange will identify the 1,000 symbols with most unstable data points and evaluate those symbols to determine whether a symbol-specific model or market model would yield better performance. The Exchange will publish on its website a list of the symbols that have an individual model to provide transparency to market participants regarding the operation of SeMi Orders.

Performance Metrics

The Exchange strategically prioritized developing the models to have a high rate of recall, which was intended to maximize the models' ability to accurately capture unstable PBBO updates while accepting that the models might identify

[17] The Exchange proposes that a new symbol will operate on the market model until the Exchange has gathered at least three days' worth of data to be able to train a symbol-specific model and determine whether it outperforms the market model for that symbol.

more periods of instability than would exist in realistic market conditions. The Exchange focused on three metrics: (1) recall, or the model's ability to accurately identify true unstable data points; (2) precision, or the model's ability to identify only true unstable data points (i.e., to not misidentify stable data points as unstable); and (3) overlocking, or the model's ability to minimize the amount of time (measured in seconds) that the model incorrectly predicts unstable market conditions.[18] The Exchange intends for the models to maximize recall and precision, while minimizing overlocking.

To evaluate the performance of the models, the Exchange selected initial baseline values for each the three parameters ***X***, ***G***, and ***g*** to maximize these performance metrics. To define an unstable PBBO period, the Exchange selected a minimum time increment ***g*** of 100 microseconds; spread threshold ***X*** of 25%; and time horizon ***G*** of one millisecond.[19]

In final performance testing of the model, the aggregate results demonstrate that the model achieves an average recall rate or 90% and average precision of 30%, with overlocking occurring for an additional 3.8 seconds on average.[20]

Production Integration

The models' compact size and average prediction speed of approximately two microseconds allows for seamless integration of the model's prediction process into the NYSE Pillar trading platform ("Pillar") on which the Exchange operates.[21] The Exchange believes that it has designed the SMI to produce an output rapidly enough to keep up with real-time trading, without overburdening Exchange systems or otherwise impacting current system performance.

Pillar will have access to full real-time trading data and will continuously maintain required features for the model, including PBBO updates, order entries, order cancellations, and book depth information. The models are invoked as soon as an evaluation trigger (e.g., a PBBO update) is received.[22] The evaluation

---

18 Recall, precision, and overlocking are discussed in more detail in Section 6.1 of the White Paper.

19 The analysis the Exchange performed to arrive at these parameter values is discussed in more detail in Section 6.2 of the White Paper. The Exchange expects that the parameter values may change over time to ensure proper calibration. The Exchange anticipates implementing the SMI in production with these parameter values, but will continue to analyze data and train the models until the date of implementation and may update these values to the extent that its analysis suggests that different values would improve performance.

20 Additional discussion of the Exchange's performance testing of the models appears in Section 6.4 of the White Paper.

21 Section 7 of the White Paper discusses the Exchange's analysis of the model's prediction time in more detail, as well as the integration of the model into the Pillar platform (see Figure 9).

22 In addition to PBBO updates, Pillar will respond to a timer-based evaluation trigger. The timer-based trigger is intended to ensure that the model remains updated when the NYSE Arca Book

process concludes by sending a message the Pillar matching engine to indicate the beginning and end of an unstable period, which would inform whether SeMi Orders are eligible to trade. For example, when Pillar receives a SeMi Order, the SMI will indicate whether the market is stable or unstable. If the market is stable, the SeMi Order will be allowed to post to the NYSE Arca Book and trade. If the market is unstable, the order will be prevented from trading until the SMI next predicts that the market is stable. For as long as the SMI predicts that the market will be unstable, a SeMi Order will remain ineligible to trade and will not execute until the SMI evaluates the market as stable.

On a given trading day, the SMI models will use the feature weights determined from the previous night's training, and the features will be calculated using real-time intraday data. The list of symbols with individualized models will be dynamically constructed daily and published before the start of trading each day.

The model will not change intra-day. The Exchange proposes to continue to retrain the model within the parameters described in this filing and the White Paper daily, outside of the Core Trading Session (on days when trading takes place on the Exchange). Retraining will be based on the last three trading days' worth of historical data. Retraining may result in changes to the features used by the model and/or the weighting of such features. The values assigned to the three parameters $\boldsymbol{X}$, $\boldsymbol{G}$, and $\boldsymbol{g}$ will not be adjusted as a result of regular model training but may be updated periodically based on the Exchange's analysis of overall model performance. Retraining is a standard and accepted process in the use of machine learning models like the ones underlying the SMI. The retraining process is not intended to result in significant or unexpected changes to the performance of the SMI or the behavior of the SeMi Order. Rather, retraining would help ensure that the SMI continues to perform well in dynamic circumstances, by allowing the models to learn from and incorporate more recent data points and would facilitate improved model performance over time. The Exchange also notes that retraining would build on the models' existing state (i.e., existing data inputs and knowledge base) and would not alter the model's objectives; retraining would result in new behaviors only to the extent that the model had not previously encountered a given scenario, and even then, any new behavior would be consistent with the model's objectives. If the Exchange determines that a retrained model would not be as successful as an existing model in achieving its objectives based on the metrics defined above, the Exchange will not implement the retrained model in production.

The Exchange will file a subsequent proposed rule change if it seeks to modify the underlying structure of the models underlying the SMI, such as the parameters $\boldsymbol{X}$, $\boldsymbol{G}$, and $\boldsymbol{g}$ used to label unstable windows or new features that could be incorporated into the models, but will not seek Commission approval prior to

changes in the absence of a PBBO update (such as when non-displayed liquidity is added to the NYSE Arca Book).

retraining the models to adjust the weighting of features that have been disclosed as potential inputs for the models or modifications to the value of any of the three identified parameters. The Exchange will also retain copies of each historical iteration of the models as part of its books and records and will make such records available to the Commission upon request. The Exchange will also publish a Trader Update in advance of implementing a retrained version of the SMI models when the Exchange has a reasonable belief that the retrained version(s) would yield results that differ materially from the prior version(s).

*****

Because of the technology changes associated with this proposed rule change, the Exchange will announce the implementation of this change by Trader Update. Subject to approval of this rule filing, the Exchange is prepared to implement the proposed rule change in 2025.

(b) Statutory Basis

The proposed rule change is consistent with Section 6(b) of the Act,[23] in general, and furthers the objectives of Section 6(b)(5),[24] in particular, because it is designed to prevent fraudulent and manipulative acts and practices, to promote just and equitable principles of trade, to foster cooperation and coordination with persons engaged in facilitating transactions in securities, 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.

The Exchange believes that the proposed change to eliminate the DPO and introduce the SeMi Order would remove impediments to, and perfect the mechanism of, a free and open market and a national market system, as well as protect investors and the public interest, by continuing to provide market participants with the benefits of an order type that can exercise discretion to trade with contra-side interest. The SeMi Order will operate in a substantially similar manner to the existing DPO, with the benefit of the SMI to provide price protection to SeMi Orders during periods of market instability. The Exchange also believes that the proposed rule change would remove impediments to, and perfect the mechanism of, a free and open market and a national market system and protect investors and the public interest because the SMI would provide improved functionality as compared to the regression model previously used by the DPO to predict market instability, which used static parameters for all symbols. Specifically, the Exchange believes that the SMI would provide enhanced price protection for SeMi Orders because its fast, lightweight, and transparent models can be seamlessly integrated into Pillar to predict future microsecond-level market stability on a symbol-specific basis. The Exchange

23 15 U.S.C. 78f(b).

24 15 U.S.C. 78f(b)(5).

believes that SeMi Orders, as proposed, would perfect the mechanism of, a free and open market and a national market system and protect investors and the public interest by relying on the SMI to restrict SeMi Orders from trading during times of predicted high market volatility, thereby avoiding potentially undesirable executions and increasing the potential for price improvement for such orders at the cost of slightly reduced fill rates.

The Exchange further believes that the proposed change to allow SeMi Orders to be designated as Liquidity Providing (an option that was not previously available to DPOs) would remove impediments to, and perfect the mechanism of, a free and open market and a national market system and protect investors and the public interest because it would afford increased flexibility to users of the order type.

The Exchange believes that the proposed change would remove impediments to, and perfect the mechanism of, a free and open market and a national market system and promote just and equitable principles of trade because the SeMi Order and SMI will operate within strict, well-defined, and transparent parameters. Although the SMI models will undergo daily retraining (outside of market hours), such retraining will aim to improve the performance of the SMI in achieving its stated objectives; retraining is not intended to alter the basic design parameters, features, or objectives of the models without prior Commission approval.[25] Moreover, the Exchange will not deploy a retrained model if it fails to achieve performance improvements based on the metrics described above. As noted above, a list of all features that may be incorporated in the models will be publicly available, and the Exchange will publish on its website daily the full list of features used for real-time calculation and available for inclusion in the SMI models. The Exchange will also retain each historical iteration of models employed by the SMI as part of its books and records and make such information available to the Commission upon request. The Exchange will also publish a Trader Update in advance of implementing a retrained version of an SMI model when the Exchange has a reasonable belief that the retrained version(s) would yield results that differ materially from the prior version(s).

The Exchange notes that neither the SMI nor the SeMi Order are designed or intended to further the performance of any participant or any category of participant over others. The Exchange believes the models underlying the SMI are objective and designed to avoid bias and discrimination. Use of the SeMi Order (like use of the DPO) remains voluntary for all market participants. Accordingly, if any market participant feels that the SeMi Order does not meet their needs, they are free to pursue other trading strategies.

---

[25] As discussed above, the Exchange will not seek Commission approval prior to allowing the models, as part of its retraining process, to vary the weighting of the features it uses. The Exchange believes this is appropriate because such variance will only occur to the extent that it will improve a model's performance with respect to pre-defined objectives.

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 that is not necessary or appropriate in furtherance of the purposes of the Act. The Exchange believes that the proposed change would promote competition by offering market participants the optional use of an order type designed to protect against potentially undesirable executions by preventing trading during periods of market instability as identified by the SMI.

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

The Exchange has neither solicited nor received written comments on the proposed rule change.

6. Extension of Time Period for Commission Action

Not applicable.

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

Exhibit 1 – Form of Notice of Proposed Rule Change for Federal Register

Exhibit 3 – White Paper

Exhibit 5 – Text of Proposed Rule Change

EXHIBIT 1

SECURITIES AND EXCHANGE COMMISSION
(Release No. 34- ; File No. SR-NYSEARCA-2024-112)

[Date]

Self-Regulatory Organizations; NYSE Arca, Inc.; Notice of Filing of Proposed Rule Change to amend Rule 7.31-E

Pursuant to Section 19(b)(1)[1] of the Securities Exchange Act of 1934 ("Act")[2] and Rule 19b-4 thereunder,[3] notice is hereby given that, on December 18, 2024, NYSE Arca, Inc. ("NYSE Arca" or the "Exchange") filed with the Securities and Exchange Commission (the "Commission") the proposed rule change as described in Items I, II, and III below, which Items have been prepared by the self-regulatory organization. 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 Rule 7.31-E to adopt the Selective Midpoint Order. The proposed rule change is available on the Exchange's website at www.nyse.com, 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 self-regulatory organization 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 those statements may be examined at the

---

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

2 15 U.S.C. 78a.

3 17 CFR 240.19b-4.

places specified in Item IV below. The Exchange has prepared summaries, set forth in sections A, B, and C below, of the most significant parts of such statements.

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

1. Purpose

The Exchange proposes to amend Rule 7.31-E(h)(3) to decommission the Discretionary Pegged Order ("DPO") and introduce the Selective Midpoint Order ("SeMi Order"). The SeMi Order will be similar to the DPO in that it is a discretionary order type, but will, unlike the current DPO, provide price protection during periods of market instability based on input from a gradient-boosting machine learning model.

Background

The Exchange currently offers the DPO, which is a non-displayed order to buy (sell) that is pegged to the same side of the PBBO and assigned a working price equal to the lower (higher) of the midpoint of the PBBO (the "Midpoint Price") or the limit price of the order.[4] Any untraded shares of such order are assigned a working price equal to the lower (higher) of PBB (PBO) or the order's limit price, which is automatically adjusted in response to changes to the PBB (PBO) for buy (sell) orders up (down) to the order's limit price. A DPO will exercise the least amount of discretion necessary from its working price to its discretionary price (defined as the lower (higher) of the Midpoint Price or the limit price of the order) to trade with contra-side interest.

Prior to November 2022, the DPO would not exercise discretion if the PBBO was determined to be unstable via a "quote instability calculation" that assessed the probability of a

[4] See Rule 7.31-E(h)(3). As defined in NYSE Arca Rule 1.1, "PBBO" means the Best Protected Bid and the Best Protected Offer. Rule 1.1 also defines "PBB" as the highest Protected Bid and "PBO" as the lowest Protected Offer.

change to the PBB or PBO.[5] The Exchange used the quote instability calculation along with real-time relative quoting activity of protected quotations to assess the probability of an imminent change to the PBBO (the "quote instability factor"). When the quoting activity met certain predefined criteria and the quote instability factor calculated was greater than the Exchange's predefined quote instability threshold, the Exchange treated the quote as unstable and restricted DPOs from exercising discretion. In November 2022, the Exchange amended Rule 7.31-E(h)(3) to eliminate the quote stability calculation, allowing DPOs to exercise discretion even during potential periods of quote instability.[6]

Proposed Rule Change

The Exchange proposes to modify Rule 7.31-E(h)(3) to replace the DPO with the SeMi Order. The SeMi Order will share the same basic attributes as the DPO. Like the DPO, the SeMi Order would be a non-displayed order to buy (sell) that is pegged to the same side of the PBBO and assigned a working price equal to the lower (higher) of the Midpoint Price or the limit price of the order. Any untraded shares of a SeMi Order would be assigned a working price equal to lower (higher) of the PBB (PBO) or the order's limit price and automatically adjusted in response to changes to the PBB (PBO) for buy (sell) orders up (down) to the order's limit price. In order to trade with contra-side orders on the NYSE Arca Book, a SeMi Order to buy (sell) would exercise the least amount of price discretion necessary from its working price to its

---

5 See Securities Exchange Act Release No. 95154 (June 24, 2022), 87 FR 39134 (June 30, 2022) (SR-NYSEArca-2022-13) (Notice of Filing of Amendment No. 2 and Order Granting Accelerated Approval of a Proposed Rule Change, as Modified by Amendment No. 2, To Amend Rule 7.31-E(h)(3) Relating to Discretionary Pegged Orders).

6 See Securities Exchange Act Release No. 96322 (November 15, 2022), 87 FR 69376 (November 18, 2022) (SR-NYSEARCA-2022-76) (Notice of Filing and Immediate Effectiveness of Proposed Rule Change to Amend Rule 7.31-E). The Exchange resumed offering the DPO in November 2022, after previously filing to temporarily suspend its use in August 2022. See Securities Exchange Act Release No. 95584 (August 23, 2022), 87 FR 52826 (August 29, 2022) (SR-NYSEARCA-2022-54) (Notice of Filing and Immediate Effectiveness of Proposed Rule Change To Amend Rule 7.31-E).

discretionary price, which is defined as the lower (higher) of the Midpoint Price or the SeMi Order's limit price.

SeMi Orders would not be displayed, must be designated Day, and would be eligible to be designated for the Core Trading Session only. SeMi Orders designated for the Early Trading Session or Late Trading Session would be rejected.

When exercising discretion, SeMi Orders (like DPOs today) would maintain their time priority at their working price as Priority 3 - Non-Display Orders and are prioritized behind Priority 3 - Non-Display Orders with a working price equal to the discretionary price of a SeMi Order at the time of execution. If multiple SeMi Orders are exercising price discretion during the same book processing action, they would maintain their relative time priority at the discretionary price.

Accordingly, the Exchange proposes to replace existing references to the "Discretionary Pegged Order" in Rule 7.31-E(h)(3) and subparagraphs (A) and (B) thereunder with references to the "Selective Midpoint Order."[7]

The Exchange proposes to add new Rule 7.31-E(h)(3)(C) to provide that the SeMi Order, as proposed, would exercise price discretion to its discretionary price, except during periods of quote instability as identified by the Selective Midpoint Indicator ("SMI") (as discussed in further detail below). If the SMI determines the PBBO for a particular security to be an unstable quote, both an arriving and resting SeMi Order would wait for a PBBO that is stable before the order's working price is adjusted and the order becomes eligible to trade. In other words, a SeMi Order would be ineligible to trade when the SMI identifies unstable market conditions and would

---

[7] The Exchange also proposes conforming changes in Rules 7.18-E(b)(1), 7.18-E(c)(1) and (5), 7.31-E(i)(3), and 7.34-E(c)(1)(A) to replace references to "Discretionary Pegged Orders" with references to "Selective Midpoint Orders."

remain in that state until the SMI determines that market conditions have stabilized, thereby preventing potentially undesirable executions during volatile or unstable market conditions. Whereas the DPO previously relied on one static logistical regression model to forecast market instability and prevented DPOs in any symbol from exercising discretion to trade when the model anticipated an unstable market, the SeMi Order, as proposed, would rely on the SMI to predict market instability using a symbol-specific gradient-boosting machine learning model and would protect SeMi Orders from trading when the SMI predicts quote instability for a given symbol.

The Exchange also proposes new Rule 7.31-E(h)(3)(D) to provide for SeMi Orders to be optionally designated as Liquidity Providing. Proposed Rule 7.31-E(h)(3)(D)(i) would provide that a SeMi Order designated as Liquidity Providing will only execute on arrival against resting orders that include a Non-Display Remove Modifier and are priced within the Liquidity Providing SeMi Order's discretionary range.[8]

Proposed Rule 7.31-E(h)(3)(D)(ii) would provide that if a resting contra-side order that does not include a Non-Display Remove Modifier is priced within an arriving Liquidity Providing SeMi Order's discretionary range, the Liquidity Providing SeMi Order will be placed on the NYSE Arca Book, and its discretionary range will be adjusted to equal the resting price of a non-displayed contra-side order or to one MPV less aggressive than the resting price of a displayed contra-side order.[9]

---

[8] The Exchange also proposes related conforming changes to Rules 7.31-E(d)(2)(B), 7.31-E(d)(3)(F), and 7.31-E(e)(1)(C) to provide that Non-Displayed Limit Orders, MPL Orders, and Non-Routable Limit Orders designated with a Non-Display Remove Modifier will trade as takers against Liquidity Providing SeMi Orders.

[9] The Exchange notes that allowing Liquidity Providing SeMi Orders to trade with resting orders with a Non-Display Remove Modifier, as well as adjusting the discretionary range of such orders, would be consistent with how similar discretionary order types function on other equities exchanges. See, e.g., Cboe EDGX Exchange, Rule 11.8(g)(5) ("[Midpoint Discretionary Orders ("MDOs")] that are not entered with a [Quote Depletion Protection ("QDP")] instruction, as defined in Rule 11.8(g)(10), will only act as the

Proposed Rule 7.31-E(h)(3)(D)(iii) provides that a Liquidity Providing SeMi Order resting on the NYSE Arca Book will not trade with an arriving contra-side order that cannot remove liquidity.[10] Once such arriving contra-side order is placed on the NYSE Arca Book, the discretionary range of the Liquidity Providing SeMi Order will be adjusted to equal the resting price of a non-displayed contra-side order or to one MPV less aggressive than the resting price of a displayed contra-side order.

Proposed Rule 7.31-E(h)(3)(D)(iv) would provide that, once resting on the NYSE Arca Book, the discretionary range of a Liquidity Providing SeMi Order will be adjusted based on resting contra-side interest as described in proposed subparagraphs (ii) and (iii) of this Rule when its working price Changes. In addition, proposed Rule 7.31-E(h)(3)(D)(iv)(a) and (b) provide that a Liquidity Providing SeMi Order to buy (sell) will not be eligible to trade at a price equal to or above (below) any sell (buy) orders that are displayed and that have a working price equal to or below (above) the working price of such Liquidity Providing SeMi Order, or at a price above (below) any sell (buy) orders that are not displayed and that have a working price below (above) the working price of such Liquidity Providing SeMi Order.

---

liquidity provider. MDOs entered with a QDP instruction will instead be allowed to remove liquidity, by default, unless the User chooses to require that the MDO only act as a liquidity provider. If the instructions included on an MDO do not permit the MDO to remove liquidity, it will only execute on entry against resting orders that include a Super Aggressive instruction priced at the MDO's pegged price if the MDO also contains a Displayed instruction, and against resting orders that include an NDS instruction priced either at the MDO's pegged price or within its discretionary range. If a resting contra-side order that does not include [a Non-Displayed Swap] instruction is priced within the discretionary range of an incoming MDO that is not permitted to remove liquidity, the incoming MDO will be placed on the EDGX Book and its discretionary range will be shortened to equal the limit price of the resting contra-side order. Likewise, where an incoming order with a Post Only instruction does not remove liquidity on entry pursuant to Rule 11.6(n)(4) against a resting MDO, the discretionary range of the resting MDO will be shortened to equal the limit price of the incoming contra-side order with a Post Only instruction.").

[10] The Exchange notes that this proposed handling is also consistent with the handling of similar discretionary order types by other equities exchanges. See id.

Finally, the Exchange proposes to renumber current Rule 7.31-E(h)(3)(C) as 7.31-E(h)(3)(E). The Exchange proposes that, as with DPOs, if the PBBO is locked or crossed, SeMi Orders would wait for a PBBO that is not locked or crossed before the working price is adjusted and the order becomes eligible to trade. Accordingly, the Exchange proposes to substitute "Selective Midpoint Order" for "Discretionary Pegged Order" in the text of proposed Rule 7.31-E(h)(3)(E).

Selective Midpoint Indicator

As described in further detail in the white paper attached as Exhibit 3 to this proposed rule change,[11] the SMI was developed using a decision tree model, which is a supervised learning algorithm used for classification and regression tasks. A decision tree model is trained sequentially, and each successive model seeks to improve on errors in the previous model by focusing on accurately predicting where the previous model performs poorly.

The Exchange determined to use a decision tree model in developing the SMI in large part because of the transparency it offers to its designers. The output of a decision tree is a hierarchy of questions that allows a user to follow the model's decision-making process, assess the importance of a given feature to the model's output, and examine the reasons underlying a specific output from the model. The Exchange believed it was important to build a model that could be interpreted and understood in this way to allow for the evaluation of, among other things, the relationships between market shifts, feature selection, and feature weightings, as well as to be able to assess overall model performance.

---

[11] See Exhibit 3, Selective Midpoint Indicator (SMI): A Gradient Boosted Signal Enabling Stable Order Executions (the "White Paper").

Market Instability Assessment

The Exchange set out to develop a model that would allow the SMI to predict market instability. For purposes of the model, the Exchange defined "instability" at a high level as relatively large price moves during a relatively short time frame. The Exchange chose PBBO updates as the fundamental data points for the model based on the evenly distributed nature of PBBO updates throughout the trading day and the granular level of information such updates offer. The model's objective is to identify windows where changes to the PBBO are unstable, in order to predict unstable markets in real time and prevent SeMi Order trading accordingly.

The Exchange established the concept of a "price jump" to further understand and categorize periods of instability. A price jump is defined as a PBBO mid-price change of a pre-defined percentage of a symbol's spread (referred to as the spread threshold ***X***), in either direction, within a configurable time interval (referred to as the time horizon ***G***). The Exchange sought to identify price jumps following a PBBO update, whether positive or negative. A price jump is marked at a given point in time if, looking back to the start of the time horizon, the mid-price was at least the spread threshold in difference from the current mid-price. This approach identifies discrete price jumps over configurable time intervals, thereby offering the Exchange flexibility to adjust the spread threshold and/or time horizon parameters according to symbol-specific dynamics.[12]

To build the SMI as a continuous signal of market instability, the Exchange next applied the price jump definition to delineate continuous periods of market instability. At each PBBO update (at index time *i*), the Exchange determined whether there had been a price jump. If there was no price jump, the market would be considered stable. If there was a price jump, the

[12] Figure 1 in Section 4.1 of the White Paper provides a detailed example of how a price jump is calculated.

Exchange continued to look for price jumps within the time horizon ***G***, until no price jump occurred. At that point, the Exchange would mark an unstable time window starting at the PBBO update that occurred at *i*-1 or 50 microseconds prior to the PBBO update at time *i*, whichever is closest to *i*, and ending at the last price jump identified.[13]

Finally, the Exchange employed an additional parameter for a minimum time increment ***g*** to establish that unstable windows identified using the spread threshold and time horizon parameters are indicative of persistent price instability. If multiple price jumps occur within a small timeframe, such that the distance between the first and last price jumps is less than ***g***, the Exchange would not mark the window as unstable. The application of the minimum time increment parameter allows the Exchange to more accurately identify true periods of instability by filtering out temporary price jumps that are quickly followed by a reversion to the price prior to the observed price jump.[14]

In developing the SMI to facilitate the SeMi Order's ability to provide protection against potentially unfavorable executions, the Exchange also wanted to be able to differentiate between quote instability on the bid side and ask side. For example, the Exchange would not want to prevent executions of buy orders based on upwards instability of the mid-price of the PBBO because executions under those conditions would likely be favorable. Accordingly, if the mid-price of the PBBO is higher at the end of an unstable window, the Exchange would only mark PBBO updates during such window as unstable for the ask side (and vice versa for the bid side).

Model Development

To develop and train the models underlying the SMI, the Exchange used data from the

---

13 Figure 2 in Section 4.2 of the White Paper provides an example of the identification of unstable windows.

14 Figure 3 in Section 4.2 of the White Paper demonstrates the operation of the minimum time increment parameter.

NYSE Arca Book from August 29, 2024 through October 22, 2024 for a set of 500 symbols selected to reflect a representative sample of the U.S. equity markets.[15] The Exchange created evaluation data points for each PBBO update that include the closest book depth state at the time of that update. This merged data set of PBBO updates and book depth data has the benefit of the information from both types of data without overinflating the size of the data set. All features and parameters used by the models are calculated based on this merged data set.

The Exchange identified a set of features based on NYSE Arca Book data—e.g., book depth information, PBBO updates, number of IOC orders—that would contribute to the models. The Exchange selected 83 unique features that could be considered for incorporation into the model.[16] These features were identified based on an iterative process, and features were selected based on their ability to explain unstable periods identified by the application of the parameters discussed above and, going forward, to predict market instability on a real-time basis so that the SMI can effectively protect SeMi Orders from potentially unfavorable executions. In production, the Exchange proposes to select a subset of these 83 features to be weighted in the model's assessment of market instability (rather than having the model utilize all available features), because focusing on the features that have been identified as most important in predicting market instability for a given symbol would both optimize prediction accuracy and processing speed.[17]

---

15 The full symbol list is included in Appendix B of the White Paper. Symbols were chosen based on criteria including absolute price level, spread in dollars, spread in basis points, and liquidity (daily ADV). Because stable periods generally far outnumber unstable periods in the U.S. equity markets, the Exchange used under-sampling methods where appropriate to reduce the number of stable data points in the data set and randomly shuffled data before training.

16 The full list of features is included in Appendix A to the White Paper. Appendix A also identifies the subset of features that were selected for the model training described in the White Paper.

17 Once the SMI is implemented in production, the full list of features currently calculated in real-time and available for evaluation for inclusion in the SMI models will be published daily on the Exchange's website.

Symbol-Specific Models

The Exchange proposes that the SMI will rely on symbol-specific models to leverage the ability of the models to incorporate different features and weightings to respond to individual symbols' unique profiles (e.g., the features most likely to accurately predict instability for a given symbol). To explore how symbol-specific models should be distributed, the Exchange ranked each of the 500 symbols in the representative data set according to the total number of unstable data points. The Exchange trained a model for each of the 500 symbols and used the model trained for the symbol SPY as a default "market model."

The Exchange found that the SMI performed better with a symbol-specific model for more active symbols (i.e., those with more unstable data points) than on the market model (i.e., testing showed sharp declines in precision and increases in overlocking behavior using the market model), whereas the SMI with a market model performed well for less active symbols (i.e., high recall with relatively small loss of precision and minimally more overlocking). The Exchange concluded that, for the SMI to provide optimal information and protection to SeMi Orders, more active symbols would benefit significantly from symbol-specific models, while less active symbols (which have fewer unstable data points to inform a symbol-specific model) could successfully default to a market model.

The Exchange believes that models tailored to individual symbols' specific characteristics would provide for better performance by the SMI and thus enhanced price protection by the SeMi Order. Accordingly, in production, the Exchange anticipates that at least 200 (and up to approximately 1,000) symbols that trade on the Exchange will have an individualized model that incorporates features that have been specifically identified for predicting market stability for that symbol. The remaining symbols that trade on the Exchange

would use the market model, which would apply the same features and weightings for all symbols.[18] Each day, the Exchange will identify the 1,000 symbols with most unstable data points and evaluate those symbols to determine whether a symbol-specific model or market model would yield better performance. The Exchange will publish on its website a list of the symbols that have an individual model to provide transparency to market participants regarding the operation of SeMi Orders.

Performance Metrics

The Exchange strategically prioritized developing the models to have a high rate of recall, which was intended to maximize the models' ability to accurately capture unstable PBBO updates while accepting that the models might identify more periods of instability than would exist in realistic market conditions. The Exchange focused on three metrics: (1) recall, or the model's ability to accurately identify true unstable data points; (2) precision, or the model's ability to identify only true unstable data points (i.e., to not misidentify stable data points as unstable); and (3) overlocking, or the model's ability to minimize the amount of time (measured in seconds) that the model incorrectly predicts unstable market conditions.[19] The Exchange intends for the models to maximize recall and precision, while minimizing overlocking.

To evaluate the performance of the models, the Exchange selected initial baseline values for each the three parameters ***X***, ***G***, and ***g*** to maximize these performance metrics. To define an unstable PBBO period, the Exchange selected a minimum time increment ***g*** of 100 microseconds; spread threshold ***X*** of 25%; and time horizon ***G*** of one millisecond.[20]

---

[18] The Exchange proposes that a new symbol will operate on the market model until the Exchange has gathered at least three days' worth of data to be able to train a symbol-specific model and determine whether it outperforms the market model for that symbol.

[19] Recall, precision, and overlocking are discussed in more detail in Section 6.1 of the White Paper.

[20] The analysis the Exchange performed to arrive at these parameter values is discussed in more detail in Section 6.2 of the White Paper. The Exchange expects that the parameter values may change over time to

In final performance testing of the model, the aggregate results demonstrate that the model achieves an average recall rate or 90% and average precision of 30%, with overlocking occurring for an additional 3.8 seconds on average.[21]

Production Integration

The models' compact size and average prediction speed of approximately two microseconds allows for seamless integration of the model's prediction process into the NYSE Pillar trading platform ("Pillar") on which the Exchange operates.[22] The Exchange believes that it has designed the SMI to produce an output rapidly enough to keep up with real-time trading, without overburdening Exchange systems or otherwise impacting current system performance.

Pillar will have access to full real-time trading data and will continuously maintain required features for the model, including PBBO updates, order entries, order cancellations, and book depth information. The models are invoked as soon as an evaluation trigger (e.g., a PBBO update) is received.[23] The evaluation process concludes by sending a message the Pillar matching engine to indicate the beginning and end of an unstable period, which would inform whether SeMi Orders are eligible to trade. For example, when Pillar receives a SeMi Order, the SMI will indicate whether the market is stable or unstable. If the market is stable, the SeMi Order will be allowed to post to the NYSE Arca Book and trade. If the market is unstable, the

---

ensure proper calibration. The Exchange anticipates implementing the SMI in production with these parameter values, but will continue to analyze data and train the models until the date of implementation and may update these values to the extent that its analysis suggests that different values would improve performance.

21 Additional discussion of the Exchange's performance testing of the models appears in Section 6.4 of the White Paper.

22 Section 7 of the White Paper discusses the Exchange's analysis of the model's prediction time in more detail, as well as the integration of the model into the Pillar platform (see Figure 9).

23 In addition to PBBO updates, Pillar will respond to a timer-based evaluation trigger. The timer-based trigger is intended to ensure that the model remains updated when the NYSE Arca Book changes in the absence of a PBBO update (such as when non-displayed liquidity is added to the NYSE Arca Book).

order will be prevented from trading until the SMI next predicts that the market is stable. For as long as the SMI predicts that the market will be unstable, a SeMi Order will remain ineligible to trade and will not execute until the SMI evaluates the market as stable.

On a given trading day, the SMI models will use the feature weights determined from the previous night's training, and the features will be calculated using real-time intraday data. The list of symbols with individualized models will be dynamically constructed daily and published before the start of trading each day.

The model will not change intra-day. The Exchange proposes to continue to retrain the model within the parameters described in this filing and the White Paper daily, outside of the Core Trading Session (on days when trading takes place on the Exchange). Retraining will be based on the last three trading days' worth of historical data. Retraining may result in changes to the features used by the model and/or the weighting of such features. The values assigned to the three parameters $\boldsymbol{X}$, $\boldsymbol{G}$, and $\boldsymbol{g}$ will not be adjusted as a result of regular model training but may be updated periodically based on the Exchange's analysis of overall model performance. Retraining is a standard and accepted process in the use of machine learning models like the ones underlying the SMI. The retraining process is not intended to result in significant or unexpected changes to the performance of the SMI or the behavior of the SeMi Order. Rather, retraining would help ensure that the SMI continues to perform well in dynamic circumstances, by allowing the models to learn from and incorporate more recent data points and would facilitate improved model performance over time. The Exchange also notes that retraining would build on the models' existing state (i.e., existing data inputs and knowledge base) and would not alter the model's objectives; retraining would result in new behaviors only to the extent that the model had not previously encountered a given scenario, and even then, any new behavior would be

consistent with the model's objectives. If the Exchange determines that a retrained model would not be as successful as an existing model in achieving its objectives based on the metrics defined above, the Exchange will not implement the retrained model in production.

The Exchange will file a subsequent proposed rule change if it seeks to modify the underlying structure of the models underlying the SMI, such as the parameters ***X***, ***G***, and ***g*** used to label unstable windows or new features that could be incorporated into the models, but will not seek Commission approval prior to retraining the models to adjust the weighting of features that have been disclosed as potential inputs for the models or modifications to the value of any of the three identified parameters. The Exchange will also retain copies of each historical iteration of the models as part of its books and records and will make such records available to the Commission upon request. The Exchange will also publish a Trader Update in advance of implementing a retrained version of the SMI models when the Exchange has a reasonable belief that the retrained version(s) would yield results that differ materially from the prior version(s).

*****

Because of the technology changes associated with this proposed rule change, the Exchange will announce the implementation of this change by Trader Update. Subject to approval of this rule filing, the Exchange is prepared to implement the proposed rule change in 2025.

2. Statutory Basis

The proposed rule change is consistent with Section 6(b) of the Act,[24] in general, and furthers the objectives of Section 6(b)(5),[25] in particular, because it is designed to prevent

[24] 15 U.S.C. 78f(b).

[25] 15 U.S.C. 78f(b)(5).

fraudulent and manipulative acts and practices, to promote just and equitable principles of trade, to foster cooperation and coordination with persons engaged in facilitating transactions in securities, 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.

The Exchange believes that the proposed change to eliminate the DPO and introduce the SeMi Order would remove impediments to, and perfect the mechanism of, a free and open market and a national market system, as well as protect investors and the public interest, by continuing to provide market participants with the benefits of an order type that can exercise discretion to trade with contra-side interest. The SeMi Order will operate in a substantially similar manner to the existing DPO, with the benefit of the SMI to provide price protection to SeMi Orders during periods of market instability. The Exchange also believes that the proposed rule change would remove impediments to, and perfect the mechanism of, a free and open market and a national market system and protect investors and the public interest because the SMI would provide improved functionality as compared to the regression model previously used by the DPO to predict market instability, which used static parameters for all symbols. Specifically, the Exchange believes that the SMI would provide enhanced price protection for SeMi Orders because its fast, lightweight, and transparent models can be seamlessly integrated into Pillar to predict future microsecond-level market stability on a symbol-specific basis. The Exchange believes that SeMi Orders, as proposed, would perfect the mechanism of, a free and open market and a national market system and protect investors and the public interest by relying on the SMI to restrict SeMi Orders from trading during times of predicted high market volatility, thereby avoiding potentially undesirable executions and increasing the potential for price improvement for such orders at the cost of slightly reduced fill rates.

The Exchange further believes that the proposed change to allow SeMi Orders to be designated as Liquidity Providing (an option that was not previously available to DPOs) would remove impediments to, and perfect the mechanism of, a free and open market and a national market system and protect investors and the public interest because it would afford increased flexibility to users of the order type.

The Exchange believes that the proposed change would remove impediments to, and perfect the mechanism of, a free and open market and a national market system and promote just and equitable principles of trade because the SeMi Order and SMI will operate within strict, well-defined, and transparent parameters. Although the SMI models will undergo daily retraining (outside of market hours), such retraining will aim to improve the performance of the SMI in achieving its stated objectives; retraining is not intended to alter the basic design parameters, features, or objectives of the models without prior Commission approval.[26] Moreover, the Exchange will not deploy a retrained model if it fails to achieve performance improvements based on the metrics described above. As noted above, a list of all features that may be incorporated in the models will be publicly available, and the Exchange will publish on its website daily the full list of features used for real-time calculation and available for inclusion in the SMI models. The Exchange will also retain each historical iteration of models employed by the SMI as part of its books and records and make such information available to the Commission upon request. The Exchange will also publish a Trader Update in advance of implementing a retrained version of an SMI model when the Exchange has a reasonable belief that the retrained version(s) would yield results that differ materially from the prior version(s).

---

[26] As discussed above, the Exchange will not seek Commission approval prior to allowing the models, as part of its retraining process, to vary the weighting of the features it uses. The Exchange believes this is appropriate because such variance will only occur to the extent that it will improve a model's performance with respect to pre-defined objectives.

The Exchange notes that neither the SMI nor the SeMi Order are designed or intended to further the performance of any participant or any category of participant over others. The Exchange believes the models underlying the SMI are objective and designed to avoid bias and discrimination. Use of the SeMi Order (like use of the DPO) remains voluntary for all market participants. Accordingly, if any market participant feels that the SeMi Order does not meet their needs, they are free to pursue other trading strategies.

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 that is not necessary or appropriate in furtherance of the purposes of the Act. The Exchange believes that the proposed change would promote competition by offering market participants the optional use of an order type designed to protect against potentially undesirable executions by preventing trading during periods of market instability as identified by the SMI.

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

No written comments were solicited or received with respect to the proposed rule change.

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 up to 90 days (i) as the Commission may designate if it finds such longer period to be appropriate and publishes its reasons for so finding or (ii) as to which the self-regulatory organization consents, the Commission will:

(A) by order approve or disapprove the 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 (https://www.sec.gov/rules/sro.shtml); or
- Send an email to rule-comments@sec.gov. Please include file number SR-NYSEARCA-2024-112 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-NYSEARCA-2024-112. This file number should be included on the subject line if email 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 website (https://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 a.m. and 3 p.m. Copies of the filing also will be available for inspection and copying at the principal office

of the Exchange. Do not include personal identifiable information in submissions; you should submit only information that you wish to make available publicly. We may redact in part or withhold entirely from publication submitted material that is obscene or subject to copyright protection. All submissions should refer to file number SR-NYSEARCA-2024-112 and should be submitted on or before [INSERT DATE 21 DAYS AFTER DATE OF PUBLICATION IN THE *FEDERAL REGISTER*].

For the Commission, by the Division of Trading and Markets, pursuant to delegated authority.[27]

Sherry R. Haywood,

Assistant Secretary.

[27] 17 CFR 200.30-3(a)(12).

EXHIBIT 3

Selective Midpoint Indicator (SMI): A Gradient Boosted Signal Enabling Stable Order Executions

Hanyu Zhang, Shreejith Billenahalli Lokegowda, Stefanos Bazinas, Sophia Lichoulas

New York Stock Exchange

November 25, 2024

1. Introduction

Equity securities, particularly highly liquid stocks, exhibit frequent price fluctuations even within just a few microseconds. While numerous studies analyze market dynamics at levels of multiple milliseconds, few consider microsecond granularity. Furthermore, prevailing research prioritizes alpha-seeking strategies over market stability and execution quality. To extend the current literature, this paper presents a novel decision tree-based model capable of forecasting market stability for a given symbol at the microsecond level.

Our Selective Midpoint Indicator ("SMI") provides an ultra-fast, lightweight, and transparent model which can be seamlessly integrated into the NYSE Pillar trading system and greatly improves the performance of an order type leveraging input from the SMI to prevent executions when the SMI predicts market instability. The SMI achieves recall rates as high as 90% for liquid symbols while minimizing the duration of the instability windows during a trading day, thus ensuring fill rates are minimally impacted.

2. Current State and Motivation

NYSE Arca currently offers a discretionary order type that market participants can leverage to provide liquidity and interact with incoming orders in a more flexible way than traditional limit orders. A Discretionary Pegged Order (as defined in NYSE Arca Rule 7.31-E(h)(3)) is a non-displayed order type that pegs to the same side of the Protected Best Bid and Offer ("PBBO") - the PBB for buy orders and PBO for sell orders - with discretion to trade up to the midpoint of the PBBO if there is available liquidity on the opposite side. Despite the flexibility that this order

type currently offers, it does not currently provide any price protection during highly volatile market conditions. Such periods of volatility often occur rapidly and last for less than 1ms, thereby offering an inherent advantage to market participants with the lowest latency profiles.

Such order types have become more prevalent in recent years and are already available on some U.S. equity exchanges. Existing offerings all share a major limitation - a timer-based approach. [1] Upon order acknowledgment, orders are protected (either by limiting the order's discretionary range or its working price) for a time interval determined at the time of order acknowledgment. Current offerings differ from each other in (1) the way the time interval is determined (static vs. dynamic) and (2) the way orders are handled during unstable periods. However, all of the existing offerings currently use a pre-determined time interval, meaning that none of the existing indicators react to information and data generated between order acknowledgement and the expiration of the timer. This pitfall can lead to missed executions simply because the timer is not making real-time predictions, and orders are protected for a non-dynamic and pre-determined time interval.

Our new SMI seeks to provide an improved experience for market participants by altogether eliminating the need for a timer. The SMI employs a gradient boosting machine learning model in evaluating real-time data from NYSE Arca's order book to predict future microsecond-level market stability for a given symbol. Selective Midpoint orders ("SeMi Orders"), as proposed in SR-NYSEARCA-2024-112, will leverage the SMI and will be locked during times when the SMI predicts quote instability for a given symbol, significantly improving execution quality for these orders and improving markouts at the cost of minimally reduced fill rates.

3. Methodology

3.1 Understanding Decision Trees

We modeled our stability prediction problem as a decision tree problem. Decision trees are supervised learning algorithms used for classification and regression tasks. They fit our problem well since we are using structured data. The algorithm splits the data into subsets based on the value of input features. The splitting forms a branching decision process and continues until a stopping criterion is met (e.g., a maximum depth or minimum number of samples per leaf).

---

[1] CBOE's Quote Depletion Protection (QDP), Nasdaq's Midpoint Extended Life Order (M-ELO) and Dynamic M-ELO and IEX's D-Peg all employ machine learning approaches to provide some level of price protection for incoming orders.

Essentially, a decision tree creates a sequential series of questions about the input data that returns an output label.

Gradient boosting is an ensemble learning technique. It involves two kinds of models: (1) weak learners, which are usually decision trees, and (2) a strong learner formed by combining multiple weak learners. Weak learners are trained sequentially, and each successive model works to improve the error from the previous model by focusing on accurately predicting cases where the first model does poorly. The final prediction in a classification problem uses majority voting of all weak models.

3.2 Decision Tree vs. Neural Networks

Decision trees are an industry standard when working with tabular data. While they are an obvious model to consider, we will explain why they are ultimately better suited than Deep Learning for our use case.

The main advantage of decision trees is explainability. The output of a decision tree is a hierarchy of specific questions such that users can easily follow the decision-making process. One can assess the importance of a feature or even examine what caused a specific model output. In other use cases, explainability would be a nice-to-have, but in our case it is critical; ensuring the best possible execution for clients is of crucial importance. Markets change in unexpected ways over time. We thus require our model to be flexible and interpretable in cases of shifting market dynamics, regime changes, or regulatory updates. An interpretable model allows us to evaluate the relationships between such market shifts and feature selection and weightings, as well as overall model performance.

The primary requirements of our model are best-in-class performance and memory efficiency. With over 80 features under consideration, we must select a model that scales well with number of features. It would be impossible to integrate a neural network model into the NYSE Pillar trading system while at the same time maintaining a sufficiently low inference time.

4. Market Stability Model

Before developing the model, we must first define “instability.” At a high level, we want our definition to capture relatively large price moves during a relatively short time frame. We consider symbols independently and recognize that market conditions change over time. We parameterize our ground truth labeling process to closely represent market instability for various symbols and changing conditions.

For clarity, we introduce an intermediate label termed "price jump" and then use that term to define "unstable," and "side-aware unstable." In the included examples, we assume there is one PBBO update per microsecond for simplicity.

4.1 Price Jump

A “price jump” is defined as a case where the PBBO mid-price moves by a pre-defined percentage of a symbol’s spread (the spread threshold ***X***) in either direction within a configurable time interval (the time horizon ***G***). If the move is positive (negative) - i.e. the mid-price increases (decreases) - then the price jump is classified as positive (negative) and is assigned a value of 1 (-1). Otherwise, the price jump value is 0.

Otherwise stated, a price jump is labeled 1 or -1 at a given point in time if, looking back to the start of the time horizon, the mid-price was at least the spread threshold in difference from the current mid-price. This approach of identifying discrete price jumps over configurable time intervals offers us the flexibility to adjust either of the parameters (spread threshold and/or time horizon) according to symbol-specific dynamics. Note that a price jump can be labeled in real time since the definition looks backwards in time.

Figure 1

Price jumps over time

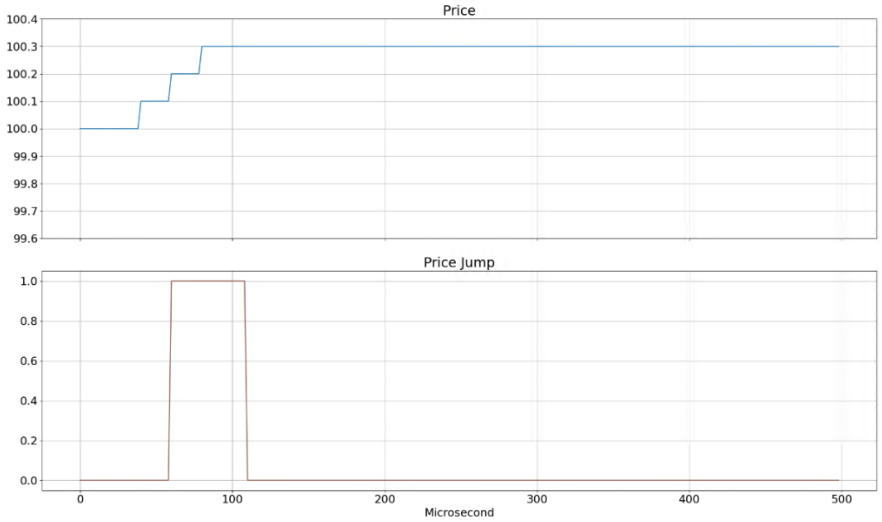

Figure 1 shows an example of how price jumps are calculated, for a sample symbol with a spread threshold of 20%, current spread of $1.00, and a time horizon

of 50us. The PBBO mid-price is $100 from 0us to 40us, and it increases by $0.10 every 20us until 80us for a total of three $0.01 price moves. At 60us, the PBBO mid-price has now increased by at least 20% of the $1.00 spread ($0.20) over the preceding 50us. Thus, all PBBO updates from 60us (the first threshold breach) to 130us (the last threshold breach plus the time horizon) are labeled with a positive price jump of 1.

4.2 Unstable Quote

Price jumps label discrete PBBO updates. Since the SMI is a continuous signal, we need to apply our price jump definition to further define continuous periods of market instability. We use the following process to define market stability.

For each PBBO update (index *i*), we check whether the conditions for a price jump are satisfied (independent of direction). If so, we then look for subsequent price jumps within our time horizon ***G***. If one can be found, we iterate this process until no price jump can be identified and classify an unstable time window starting at the *i-1* PBBO update or 50us prior to the *i* PBBO update (whichever is closest to *i*) and ending at the last price jump.

Figure 2

Unstable window (time horizon of 400us)

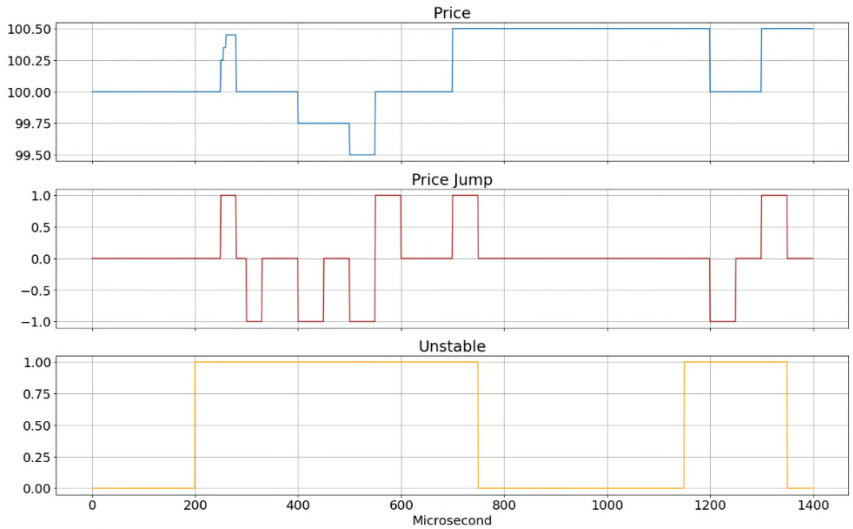

To ensure that the unstable windows we are capturing are persistent, we employ a third parameter ***g***. If multiple price jumps are too close to each other and

fall into one small time window such that the distance between the first and last price jump is less than g (e.g., g = 100us), then we do not mark this window as unstable. This is a pivotal step in the process, as it allows us to isolate and correctly identify persistent price moves and filter out temporary ones that quickly reverse back to the original price. Figure 3 shows how the configuration of parameter g filters out temporary price fluctuations.

Figure 3

Stable window with price jumps

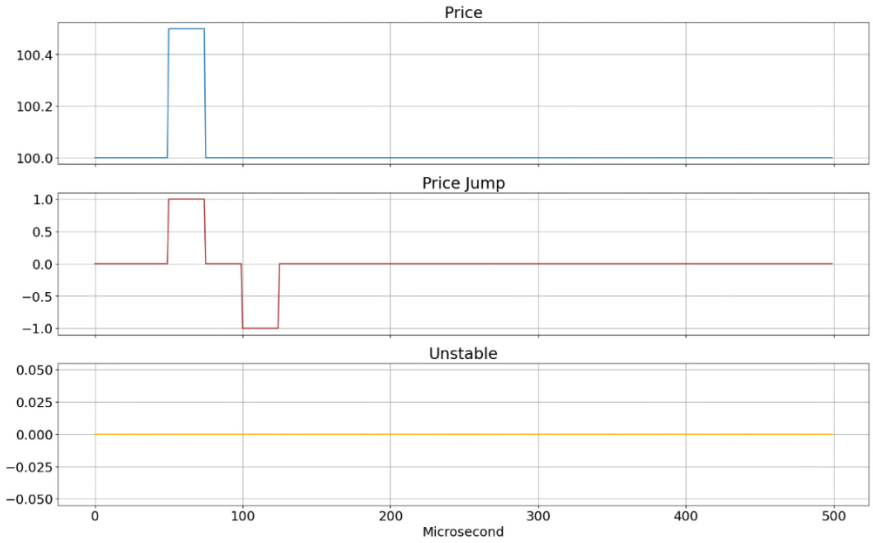

4.3 Side-Aware Unstable Quote

In the development of the SMI, it is important to differentiate between quote instability on each side of the quote, i.e., separately for the bid and ask side. This is of crucial importance when considering its implementation in future order type behavior, since the aim of SeMi Orders is to provide protection against adverse movements in a stock's price.

In this context, we would want to avoid preventing executions of buy (sell) orders based on upwards (downwards) instability of the mid-price of the PBBO, since these would be considered favorable conditions for order execution. In other words, if the mid-price of the PBBO is higher (lower) at the end of an unstable window, we only mark PBBO updates in this window as unstable for the ask (bid) side of the model (Figure 4). Side-aware instability is therefore less restrictive than the instability indicator described in section 4.2 and leads to separate models for each side of the PBBO, a bid model and an ask model.

Figure 4

Side-aware unstable window

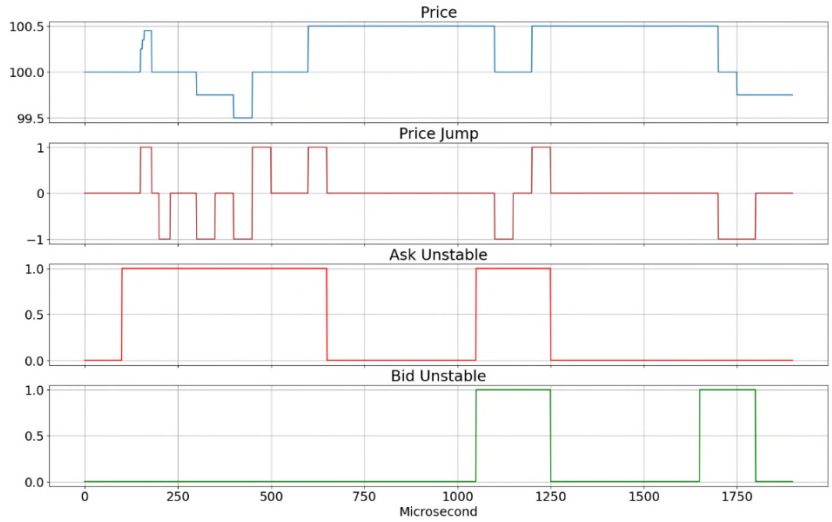

4.4 Final Model Parameters

As outlined in sections 4.1-4.3, the three parameters that need to be accurately tuned and used in our model evaluation are the:

1. Spread threshold ***X***, used in identifying and labelling price jumps.
2. Time horizon ***G***, used in identifying and labelling unstable time windows.
3. Minimum time increment ***g***, used to filter for persistent price changes during unstable time windows.

5. Model Development

5.1 Feature Selection

We carefully choose features to avoid any unnecessary model input. The features come from both the NYSE Arca order book and market data, and data points include but are not limited to book depth information, PBBO updates, number of immediate-or-cancel orders/shares and price aggressiveness of such orders. Appendix A contains a complete list of the 83 features under consideration as inputs in the SMI model. This list was developed through an iterative process of evaluating individual feature significance.

5.2 Data Gathering & Preprocessing

As part of the training process, we use data from the NYSE Arca order book for the period 8/29/24-10/22/24 and consider a set of 500 symbols chosen to reflect a representative sample of the US equity markets (full symbol list in Appendix B). Symbols are chosen based on multiple criteria, including but not limited to absolute price level, spread (in dollars), spread (in basis points), and liquidity (daily ADV) to ensure a representative sample across all criteria categories is evaluated.

PBBO updates are labeled according to Section 4.3 as 1s or 0s for each side of the PBBO. Due to the nature of U.S. equity markets, stable datapoints usually largely outnumber unstable datapoints. As previously argued, unstable market conditions tend to last for very short time intervals but tend to represent the majority of unfavorable executions on a given day. To account for this data imbalance and to avoid a biased model output, we employ under-sampling methods when appropriate to reduce the number of observations in the majority (stable) class and randomly shuffle the data before training.

5.3 Training Process

Each model is retrained daily using data from the preceding three days. Evaluation data points are created for each PBBO update, rather than each book depth update. Book depth information has much more granularity than PBBO updates: two adjacent data points in the first several levels of book depth could be the same if the change happens on a deeper level. To reduce the size of the dataset and increase the informational value-add of each data point, we create new data points upon each PBBO update and merge the closest book depth state as of the time of each PBBO event. We calculate all features and labels based on this merged dataset.

6. Results

6.1 Performance Metrics

The objective of our framework is to correctly predict and capture as many unstable PBBO updates as we can, while at the same time limiting the number of times we incorrectly predict a stable PBBO update as unstable. Our emphasis lies more in achieving greater recall rather than solely pursuing an elevated overall F1

score. This strategic shift reflects our willingness to accept instances where the model identifies more data points as unstable than indicated by the ground truth, provided that the number of false unstable instances (and aggregate time associated with those instances) remains proportionally insignificant compared to the total number of PBBO updates (or total seconds) within a given day. As a result, our focus is on:

(1) Maximizing model recall, i.e., making sure we are capturing the maximum number of truly unstable data points.

$$Recall = \frac{True\ Unstable}{True\ Unstable + False\ Stable}$$

(2) Maximizing model precision, i.e., making sure we mislabel as few truly stable data points as possible in our predictions.

$$Precision = \frac{True\ Unstable}{True\ Unstable + False\ Unstable}$$

(3) Minimizing overlocking, i.e., minimizing the amount of time (in seconds) that the model predicts unstable market conditions over the ground truth.

$$Overlocking = \frac{Model\ predicted\ unstable\ time\ (sec)}{Actual\ unstable\ time\ (sec)}$$

We evaluate model performance holistically across metrics. Price-driven performance supersedes fill rates when evaluating price protection mechanisms. To improve price-based markouts, recall is our North Star metric whereas precision and overlocking serve as guardrail metrics. So long as recall is strong and overlocking is low, we are willing to accept a precision baseline of at least 20% during the model evaluation process, which is reasonable given the imbalanced nature of our dataset.

6.2 Parameter Calibration

In order to evaluate the performance of our model, it is important to first calibrate our three parameters (see Section 4.4). In the below sections, we evaluate the stability and performance of a baseline model for different levels of each parameter to arrive at the final parameter values. Given the dynamic nature of U.S. equity markets and the diversity of symbols it is composed of (e.g., price, spread,

liquidity), these parameter values are likely to change over time to ensure proper calibration.

6.2.1 Minimum Time Increment (*g*)

Table 1 below shows the distribution of time between successive PBBO updates for 15 of the most liquid U.S. equity symbols as well as the entirety of the U.S. equity markets for the period July-October 2024. For these 15 symbols, 10% of all PBBO updates happen within at most 64us from one another. This result holds for the aggregate of U.S.-listed equity securities, with 10% of all PBBO updates occurring within 46-55us during the four months. This further illustrates our earlier point about the ever-increasing speed of price updates in today's markets and highlights the need to ensure proper identification of market instability.

Some PBBO updates reflect a real change in the price of a security, while some simply capture a reversion in price when a security was temporarily dislocated. We do not attempt to decide what the "true" price of a security is, but it stands to reason that quick corrections of dislocations would often have relatively short times between PBBO updates.

Table 1

Distribution of time (microseconds) between successive PBBO updates

	Jul-24				Aug-24				Sep-24				Oct-24			
	P10	P25	P40	P50	P10	P25	P40	P50	P10	P25	P40	P50	P10	P25	P40	P50
AAPL	25	79	292	695	29	101	386	910	33	117	433	916	37	144	509	1,095
AMD	27	99	411	960	36	139	508	993	49	211	623	1,241	46	180	540	1,064
BAC	13	43	108	271	13	43	106	235	13	46	117	265	16	55	146	362
F	9	35	75	131	5	32	66	108	9	40	85	149	7	39	86	151
NIO	9	32	69	114	8	30	60	99	9	31	64	107	6	30	68	117
NVDA	27	92	341	801	30	105	387	840	32	118	440	1,000	30	95	388	937
PFE	13	43	93	183	13	42	91	175	10	38	81	148	11	39	86	160
QQQ	64	324	1,332	3,597	55	234	739	1,857	53	204	635	1,473	59	258	890	2,494
SNAP	16	46	114	293	13	40	96	205	11	38	87	167	12	39	86	167
SOXL	35	110	770	3,955	28	73	396	2,216	25	63	293	1,386	24	60	222	984
SOXS	18	60	169	447	28	91	482	2,872	22	66	200	617	17	53	140	349
SPY	57	241	833	2,488	57	242	725	1,802	58	237	714	1,747	57	243	786	2,191
T	14	42	93	186	11	39	84	155	7	37	82	144	9	38	83	149
TQQQ	22	57	257	2,600	26	72	357	2,676	23	62	221	1,193	20	50	111	394
TSLA	23	80	272	589	25	102	348	774	41	141	531	1,125	36	122	453	1,047
All	46	282	1,927	15,613	53	321	1,964	12,709	55	319	1,796	11,880	53	277	1,330	10,811

To identify the typical length of such reversion events, we analyze all PBBO updates where the mid-price of the PBBO is different within 25us but reverts to the original price within an additional 175us and extract the time interval elapsed until the price reversion. Table 2 shows the distribution of these intervals for the same set of U.S.-listed equity securities and period as Table 1. 75% of all such price reversions occur within at most 100us from the original PBBO update. Stated

otherwise, it only takes a total of 100us for a price to change and revert to the original price for 75% of all such cases. This value of 100us provides a very accurate and reasonable starting estimate for our minimum time increment ***g***.

Table 2

Distribution of time (microseconds) until price reversion for PBBO changes occurring within 50us

	Jul-24			Aug-24			Sep-24			Oct-24		
	P25	P50	P75	P25	P50	P75	P25	P50	P75	P25	P50	P75
AAPL	45	70	101	43	69	100	47	72	101	50	76	109
AMD	55	83	108	52	81	107	56	83	115	54	83	115
BAC	45	73	112	47	75	116	46	74	117	48	77	119
F	50	79	120	51	82	129	51	80	122	54	83	126
NIO	50	79	123	54	83	126	51	78	118	50	79	120
NVDA	48	76	107	48	74	107	47	72	106	46	71	106
PFE	48	76	117	50	77	119	50	78	121	51	78	121
QQQ	51	80	119	52	80	120	53	80	117	51	79	119
SNAP	44	70	107	47	76	117	47	74	115	48	77	117
SOXL	48	73	107	46	72	109	45	70	108	46	75	112
SOXS	45	74	113	46	73	111	47	76	116	45	75	117
SPY	49	74	111	52	79	118	53	81	119	50	77	117
T	48	77	121	50	79	125	50	81	126	52	82	129
TQQQ	50	81	100	45	74	102	47	76	100	53	82	109
TSLA	57	85	114	55	89	138	52	81	111	46	74	108
All	49	78	108	49	78	112	50	78	113	50	78	115

6.2.2 Spread Threshold (*X*)

The spread threshold determines the minimum change in the mid-price of the PBBO that we aim to predict. Figure 5 below shows the relationship between (1) recall and precision and (2) overlocking and precision for various levels of spread thresholds and time horizons, with each point in the plot representing the median values for a given day across the 500 symbols under consideration.

As shown in the scatterplot, a threshold value of 25% greatly outperforms parameter values of $0.01 and 50% in terms of all the metrics under evaluation: maximizing recall, maximizing precision, and minimizing overlocking. In particular, spread threshold values of $0.01 or 50% only marginally increase recall at the cost of multifold decrease in precision and multifold increase in overlocking. Additionally, a spread threshold value of 25% also outperforms the $0.005 parameter value by providing a recall and precision boost at only a slightly higher overlocking ratio. Since over 70% of U.S.-listed equity securities have daily average spreads of greater than $0.02 - meaning that 25% of the spread for these symbols would equal at least $0.05 - we choose the spread threshold level of 25% which maximizes the performance metrics analyzed.

Figure 5

Performance metrics for different levels of model spread threshold

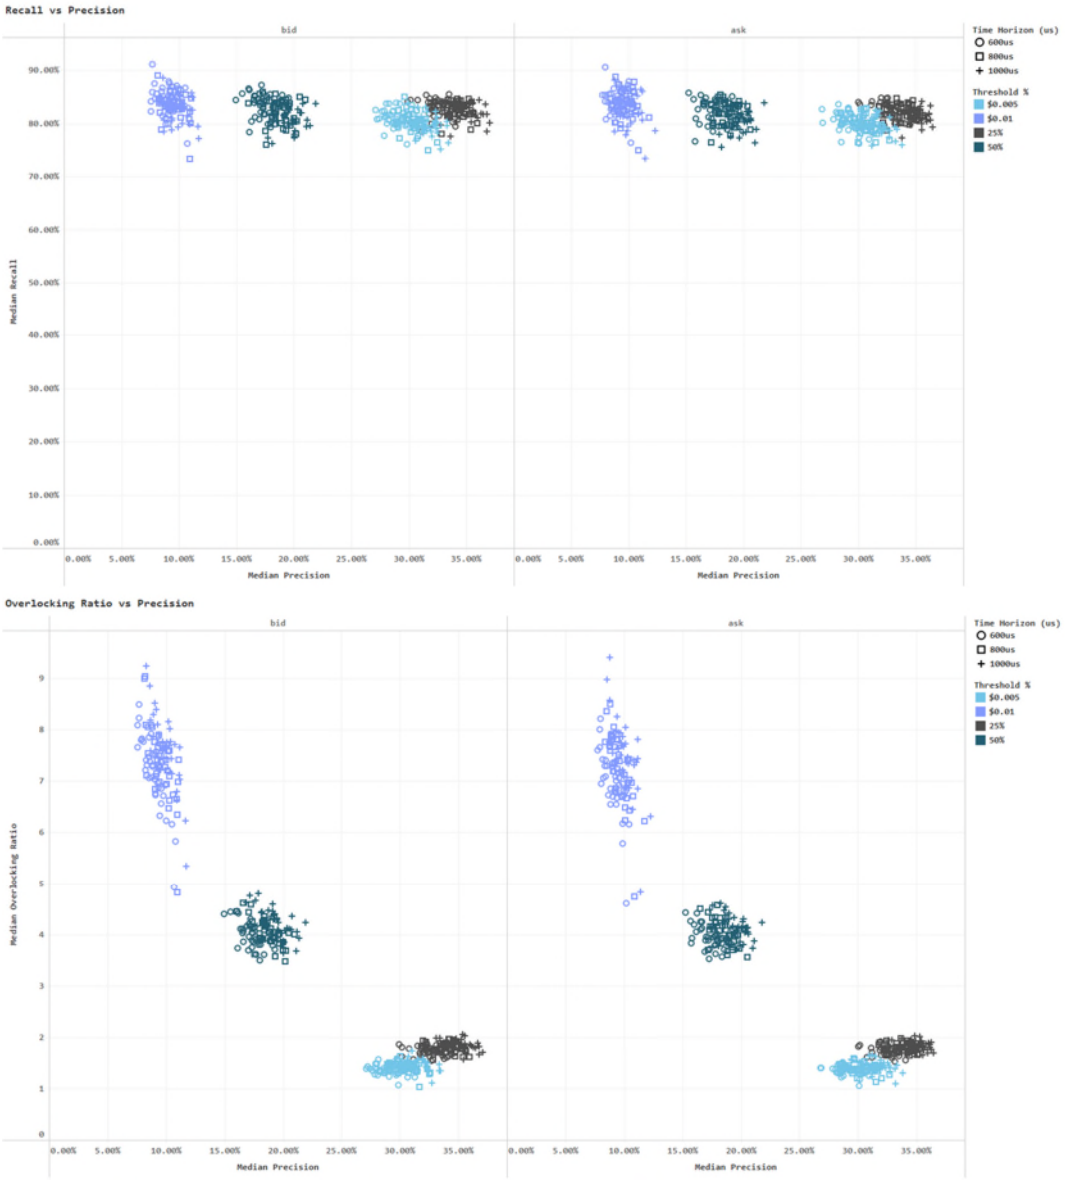

6.2.3 Time Horizon (*G*)

Our selection of time horizon will greatly affect our model performance as shown in Table 3. For both aforementioned levels of spread threshold ($0.005 and 25%) for the 500 symbols in our universe, increasing our time horizon from 600us to 1ms greatly increases our model's precision, with very minimal impact on the model's recall and overlocking performance.

Table 3

Performance metrics for different levels of model time horizon

		Bid			Ask		
		Median Recall	Median Precision	Median Overlocking Ratio	Median Recall	Median Precision	Median Overlocking Ratio
$0.01	**600us**	80.91%	28.81%	1.355	80.33%	28.95%	1.349
$0.01	**800us**	80.40%	30.45%	1.407	80.07%	30.50%	1.385
$0.01	**1000us**	80.09%	31.54%	1.455	79.69%	31.70%	1.43
25%	**600us**	83.08%	32.32%	1.735	82.59%	32.41%	1.724
25%	**800us**	82.68%	34.04%	1.799	82.23%	34.02%	1.786
25%	**1000us**	82.24%	35.03%	1.865	81.78%	35.06%	1.849

6.3 Symbol-Specific vs. Market-Based Models

With over ten thousand securities traded listed and traded in the U.S. equity markets, and NYSE Arca trading over 90% of them daily, deciding how to scale our model across this vast universe of symbols is crucial. Even though maintaining a one-model-per-symbol solution could potentially be efficient enough to be incorporated into a trading system such as that of NYSE Arca, such an approach misses the more important question. What is the value-add of creating and maintaining a one-model-per-symbol framework?

To answer this question, we rank our universe of 500 symbols according to their total number of unstable data points that we were able to label for our training purposes. We then compare the performance of our universe of symbols on each symbol's symbol-specific trained model against its performance on the SPY trained model. In this context, we are using the SPY model as a proxy for a "market model" that could potentially be used for symbols with insufficient number of training data points. Figure 6 plots the difference of each symbol's SPY-model performance and symbol-specific performance. More active/volatile symbols (with more unstable datapoints) are color-coded in red, while less active/volatile symbols (with less unstable datapoints) are color-coded in green.

Symbols with fewer unstable data points perform much better on the SPY-trained model; recall is significantly higher, at the cost of a relatively small loss of precision and minimally more overlocking. In sharp contrast, more active/volatile symbols with more unstable data points to train each symbol's symbol-specific model perform noticeably worse on the SPY model; small gains in recall are paired with sharp declines in precision (as much as -30%) and large increase in overlocking behavior. As a result, our framework will need to employ a symbol-specific model for more active symbols, while symbols with less activity will need to default to a 'market' model, i.e. the SPY symbol-specific model.

Figure 6

Performance metrics for symbol-specific models vs SPY-model

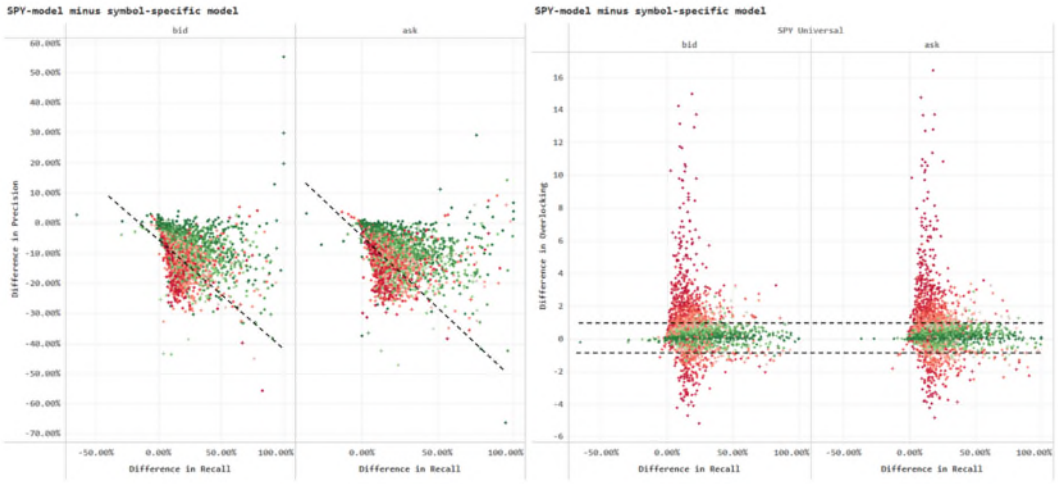

6.4 Final Model Performance

To evaluate the performance of our final model, we combine the results and conclusions from Sections 6.1-6.3. We choose 100us as our minimum time increment (***g***), 1ms as our time horizon (***G***), and 25% as our spread threshold (***X***). We split our symbol universe into two separate groups – group 1 consisting of the symbols in the 20th percentile of average daily unstable datapoints during the evaluation period (a total of 100 symbols), and group 2 consisting of the remaining 80% of symbols (a total of 400 symbols). For symbols in group 1, we use each symbol's symbol-specific trained model. For symbols in group 2, we use the model trained on SPY data (the market model).

Table 4 shows the aggregated performance results of this combination of models. In aggregate, our model achieves an average recall rate of 90% and average precision of 30%, while overlocking for an additional 3.8 seconds on average. As discussed in Section 6.1, maximizing the model's recall rate and minimizing the overlocking window are of primary significance when evaluating the performance of our model. The model precision of 30% is well above the baseline of 20% and warranted to achieve high levels of recall without incurring substantial overlocking.

Table 4

Final Model Results

	Recall	Precision	Overlocking (seconds)	Overlocking Ratio
Bid	90.49%	30.18%	3.977	2.863
Ask	89.70%	29.95%	3.621	2.953
Aggregate	90.10%	30.06%	3.799	2.908

Symbol-specific model for 100 symbols with the highest number of unstable datapoints & Market-model for the rest 400 symbols
Time horizon: 1000us (1ms)
Spread threshold: 25%

7. Model In Production Environment

7.1 Prediction Time

The prediction speed of our model is critical for its implementation within a production trading environment. The NYSE Pillar matching engines on average match orders within a few microseconds. If the model's process required more time than the matching time, locked symbols could fail to unlock quickly enough to execute against incoming orders, potentially causing missed opportunities and low fill rates. The model's current implementation is both efficient and fast. The model output is produced fast enough for real time trading, does not place an undue burden on the trading system, and does not degrade performance.

To achieve our fast prediction speeds, we maintain a compact model size, achieved through the incorporation of early stopping mechanisms during the training phase to limit the number of trees in the model. Additionally, we compile the model into a C static library. This strategic approach yields a substantial reduction in average prediction time – from 10ms to a mere 0.5us – when using one model to predict continuously with an Intel Xeon Platinum 8259CL CPU.

In practice, we have multiple models running sequentially in one process, one for each symbol (according to process described in section 6.4). Figure 7 illustrates the average prediction time when multiple such models are running within a single process. We selected one model at a time using a round-robin approach and did not utilize cache space when the combined size of all models exceeded the cache size, underscoring the importance of maintaining small model sizes. We additionally employ the existing symbol sequence of real production orders within NYSE Pillar to assess speed, resulting in an average prediction time of approximately 2us. Such efficiency allows seamless integration of the model prediction process into the NYSE Pillar trading system.

Figure 7

Average model prediction time by number of models utilized

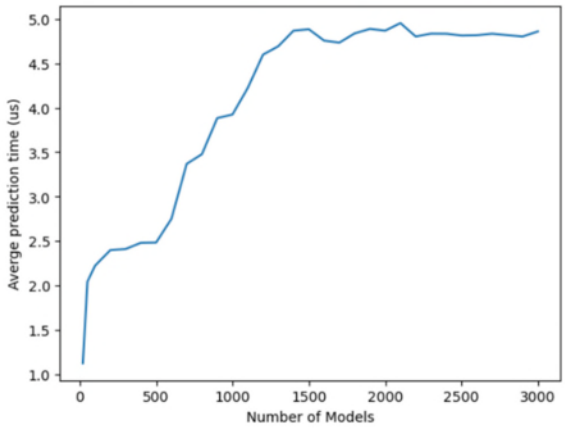

7.2 Daily model maintenance

The daily model maintenance includes data collection, data preprocessing, training, model compiling, and shared library compiling. The first three steps are the same as described in Section 5; we collect historical exchange data, preprocess them on multiple instances on cloud or on-prem servers, and train the models. Furthermore, as discussed in Section 5, we proceed to compile all models into C static libraries. Subsequently, we consolidate these static libraries into a single shared library featuring just one function entry point. This function accepts a symbol as one of its parameters, facilitating the invocation of the respective model associated with the symbol.

7.3 Intraday prediction: event driven pipeline

An internal trading system process which has access to the full real time trading data continuously maintains all the required features for this model, including but not limited to PBBO updates, order entries, order cancellations and book depth information. This process invokes the model for inference as soon as an evaluation trigger is received. Evaluation triggers primarily include PBBO updates as received by the NYSE Pillar matching engine. Since the model incorporates PBBO-agnostic features (e.g., individual order entries/cancellations), the PBBO trigger is complemented by a timer-based trigger. The inclusion of a timer-based trigger ensures that the model remains up-to-date when NYSE Arca's order book changes without an associated PBBO update, such as when non-displayed liquidity is added to or removed from the order book. This end-to-end process concludes by sending a

message to the NYSE Pillar matching engine to indicate the beginning and end of an unstable period (Figure 9).

Figure 9

Model integration with NYSE Pillar trading system (simplified)

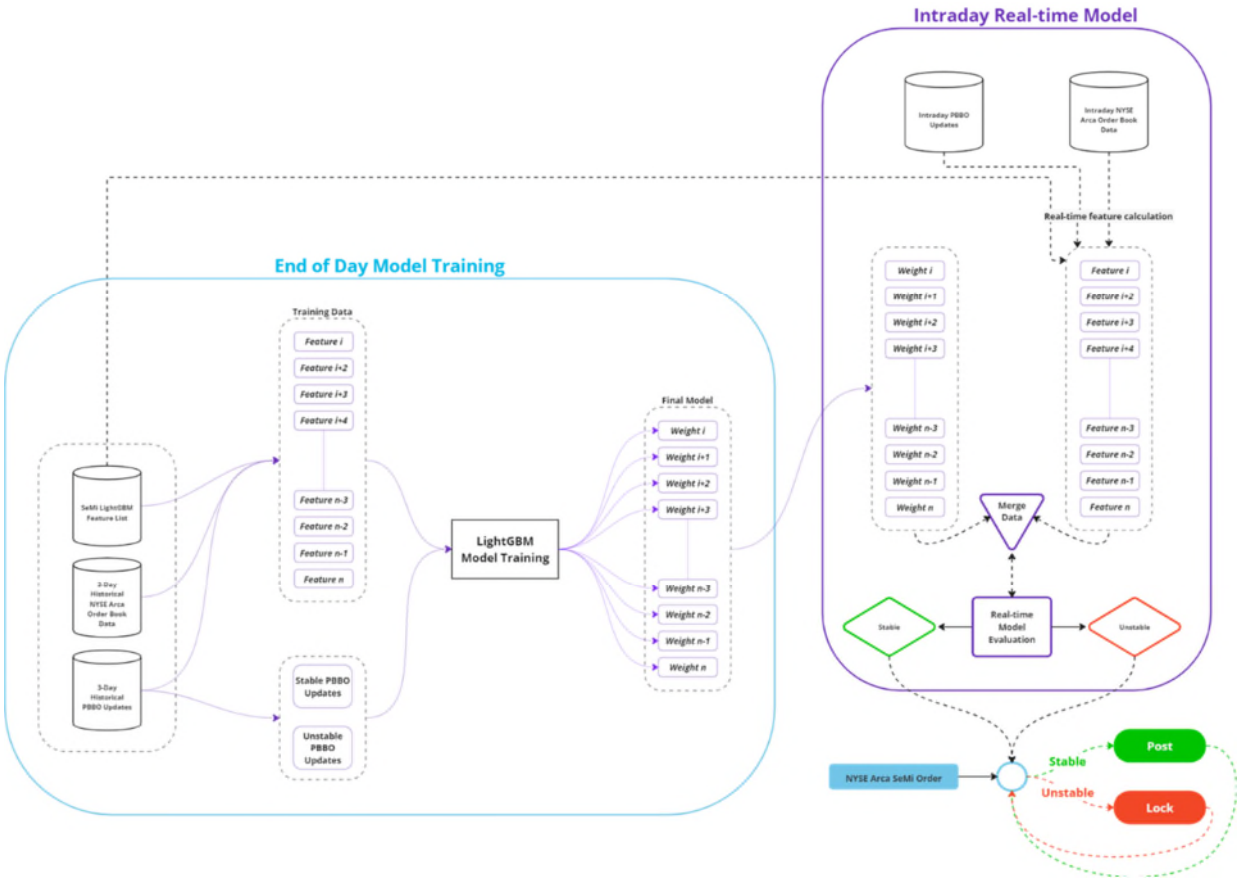

8. Conclusion

We have proposed a new machine learning indicator, the SMI, to predict market instability at any given point in time for a given security. Our solution is not just a machine learning model, but a machine learning system. The SMI uses a gradient boosted decision tree model to predict market instability and provides a very fast and explainable model. The speed of prediction of the model allows us to provide a near continuous inference pipeline to move from the traditional timer paradigm to a real-time solution.

The SMI uses data from the consolidated feed as well as NYSE Arca book data to provide these predictions. Liquid symbols with enough training data points

retrieve predictions from each symbol's symbol-specific SMI model, while less active symbols use a market SMI model for inference. The SMI achieves recall rates of at least 90% for the majority of liquid symbols while minimizing the overlocking time windows within a given day.

NYSE Arca's SeMi Orders will leverage this signal to provide an improved trading experience, especially for low latency market participants. When the SMI signals an unstable market, a SeMi Order will be prevented from executing against liquidity on the opposite side of the market until the SMI indicator predicts stable market conditions. Once the SMI signals a stable market, a SeMi Order is allowed to execute discretion up to the midpoint of the PBBO to trade. SeMi Orders avoid undesirable order matching during times of high market volatility, increasing the potential for price improvement for such orders at the cost of slightly reduced fill rates. Our strong performance results and new technology capabilities suggest further opportunities to improve client outcomes using real-time signals. We seek to continue to improve market quality, especially when benefits can be shared among multiple market participants.

Appendix A - Model Feature List

Feature Name	Explanation
ord_buy_qty_10_ms[3]	Total ordered shares on buy side in the last 10 ms
ord_buy_qty_20_ms	Total ordered shares on buy side in the last 20 ms
ord_buy_qty_50_ms	Total ordered shares on buy side in the last 50 ms
ord_buy_qty_100_ms	Total ordered shares on buy side in the last 100 ms
ord_sell_qty_10_ms[3]	Total ordered shares on sell side in the last 10 ms
ord_sell_qty_20_ms	Total ordered shares on sell side in the last 20 ms
ord_sell_qty_50_ms	Total ordered shares on sell side in the last 50 ms
ord_sell_qty_100_ms	Total ordered shares on sell side in the last 100 ms
n_ioc_ord_10ms[3]	Number of IOC distinct orders submitted in the last 10 ms
n_ioc_ord_20ms	Number of IOC distinct orders submitted in the last 20 ms
n_ioc_ord_50ms	Number of IOC distinct orders submitted in the last 50 ms
n_ioc_ord_100ms	Number of IOC distinct orders submitted in the last 100 ms
n_ioc_share_10ms	Number of IOC shares in the last 10 ms
n_ioc_share_20ms	Number of IOC shares in the last 20 ms
n_ioc_share_50ms	Number of IOC shares in the last 50 ms
n_ioc_share_100ms	Number of IOC shares in the last 100 ms
pbbo_updates_1ms[3]	Number of PBBO updates in last 1ms
bid_bps_0[1,3]	Price on book depth level 0 on bid side
bid_bps_1[1,3]	Price on book depth level 1 on bid side
bid_bps_2[1,3]	Price on book depth level 2 on bid side
ask_bps_0[1,3]	Price on book depth level 0 on ask side
ask_bps_1[1,3]	Price on book depth level 1 on ask side
ask_bps_2[1,3]	Price on book depth level 2 on ask side
bid_vol_0[3]	Volume on book depth level 0 on bid side
bid_vol_1[3]	Volume on book depth level 1 on bid side
bid_vol_2[3]	Volume on book depth level 2 on bid side
ask_vol_0[3]	Volume on book depth level 0 on ask side
ask_vol_1[3]	Volume on book depth level 1 on ask side
ask_vol_2[3]	Volume on book depth level 2 on ask side
price_jump_1ms[3]	Sum of price jumps (with direction) within a 1ms sliding window prior to the current PBBO update
price_jump_abs_1ms[3]	Number of price jumps within a 1ms sliding window prior to the current PBBO update
bid_momentum_100us[2,3]	All bid side momentum which is calculated within a 100us sliding window prior to the current PBBO update
bid_momentum_500us[2,3]	All bid side momentum which is calculated within a 500us sliding window prior to the current PBBO update
bid_momentum_1ms[2,3]	All bid side momentum which is calculated within a 1ms sliding window prior to the current PBBO update
bid_momentum_2ms[2,3]	All bid side momentum which is calculated within a 2ms sliding window prior to the current PBBO update
bid_momentum_5ms[2,3]	All bid side momentum which is calculated within a 5ms sliding window prior to the current PBBO update
bid_momentum_10ms[2,3]	All bid side momentum which is calculated within a 10ms sliding window prior to the current PBBO update
ask_momentum_100us[2,3]	All ask side momentum which is calculated within a 100us sliding window prior to the current PBBO update
ask_momentum_500us[2,3]	All ask side momentum which is calculated within a 500us sliding window prior to the current PBBO update
ask_momentum_1ms[2,3]	All ask side momentum which is calculated within a 1ms sliding window prior to the current PBBO update
ask_momentum_2ms[2,3]	All ask side momentum which is calculated within a 2ms sliding window prior to the current PBBO update
ask_momentum_5ms[2,3]	All ask side momentum which is calculated within a 5ms sliding window prior to the current PBBO update
ask_momentum_10ms[2,3]	All ask side momentum which is calculated within a 10ms sliding window prior to the current PBBO update
bid_level_change_1ms[3]	Sum of level shifts (with direction) on bid side in the book within 1ms sliding window prior to the current PBBO update
bid_level_change_abs_1ms[3]	Number of level shifts on bid side in the book within 1ms sliding window prior to the current PBBO update
ask_level_change_1ms[3]	Sum of level shifts (with direction) on ask side in the book within 1ms sliding window prior to the current PBBO update
ask_level_change_abs_1ms[3]	Number of level shifts on ask side in the book within 1ms sliding window prior to the current PBBO update
hour[3]	Hour of current PBBO update time
minute[3]	Minute of current PBBO update time
second[3]	Second of current PBBO update time
ord_buy_n_10_ms	Total orders on buy side in the last 10ms
ord_buy_n_20_ms	Total orders on buy side in the last 20ms
ord_buy_n_50_ms	Total orders on buy side in the last 50ms
ord_buy_n_100_ms	Total orders on buy side in the last 100ms
ord_sell_n_10_ms	Total orders on sell side in the last 10ms
ord_sell_n_20_ms	Total orders on sell side in the last 20ms

Feature Name	Explanation
ord_sell_n_50_ms	Total orders on sell side in the last 50ms
ord_sell_n_100_ms	Total orders on sell side in the last 100ms
cxl_buy_10ms[3]	Cancelled shares on buy side in the last 10ms
cxl_buy_20ms	Cancelled shares on buy side in the last 20ms
cxl_buy_50ms	Cancelled shares on buy side in the last 50ms
cxl_buy_100ms	Cancelled shares on buy side in the last 100ms
cxl_sell_10ms[3]	Cancelled shares on sell side in the last 10ms
cxl_sell_20ms	Cancelled shares on sell side in the last 20ms
cxl_sell_50ms	Cancelled shares on sell side in the last 50ms
cxl_sell_100ms	Cancelled shares on sell side in the last 100ms
exec_share_10ms	Executed share in the last 10ms
exec_share_20ms	Executed share in the last 20ms
exec_share_50ms	Executed share in the last 50ms
exec_share_100ms	Executed share in the last 100ms
exec_order_10ms	Executed order in the last 10ms
exec_order_20ms	Executed order in the last 20ms
exec_order_50ms	Executed order in the last 50ms
exec_order_100ms	Executed order in the last 100ms
ioc_exec_share_10ms	Number of IOC shares executed in the last 10 ms.
ioc_exec_share_20ms	Number of IOC shares executed in the last 20 ms.
ioc_exec_share_50ms	Number of IOC shares executed in the last 50 ms.
ioc_exec_share_100ms	Number of IOC shares executed in the last 100 ms.
ioc_exec_ord_10ms	Number of IOC distinct orders executed in the last 10 ms.
ioc_exec_ord_20ms	Number of IOC distinct orders executed in the last 20 ms.
ioc_exec_ord_50ms	Number of IOC distinct orders executed in the last 50 ms.
ioc_exec_ord_100ms	Number of IOC distinct orders executed in the last 100 ms.

[1] Price on book depth level X is normalized by following formula:

- Bid side:

$$D_{bid,\,Lvl=l} = \frac{P_{mid} - P_{bid,Lvl=l}}{P_{mid}} * 10^6$$

- Ask side:

$$\mathrm{D}_{\mathrm{ask,\,Lvl=l}} = \frac{P_{ask,Lvl=l} - P_{mid}}{P_{\mathrm{mid}}} * 10^6$$

[2] Momentum of volume change in book levels within a time window. This feature consists of book depth information in a continuous period and shows the trend of volume and price movement. For example, if we see volume increasing aggressively on the bid side and decreasing on the ask side, then the price will be more likely to change upwards. After we calculate momentum on each PBBO update, we add all momentum within a 0.1/0.5/1/2/5ms sliding window prior to each PBBO update as additional data points.

[3] Subset of features with high significance during evaluation period and subsequently used in model performance testing.

Appendix B - Test Symbol Universe

NVDA	ARKB	IYR	CMCSA	DAL	PINS	VONG	IYE	SCHD	BKAG
SPY	PDD	C	S	BMY	U	LUMN	PTON	EFA	BTF
QQQ	PLTR	IYW	KSS	NOVA	MTUL	FNDA	UTEN	GOLD	CALF
TQQQ	BABA	NVDY	BAC	CCL	NCLH	IP	PM	D	NOBL
SOXL	ASTS	UCO	WDC	TMV	AAPB	CL	MTCH	VST	FHLC
IWM	USO	AVLV	IREN	MS	SO	URA	UITB	YINN	AIQ
QQQM	IJR	TLH	MDLZ	VTV	AAXJ	Z	DOW	MSFU	VNQI
AAPL	AAPU	FCX	IWN	GME	COLB	CM	DUK	QQQJ	TPR
GOOGL	DDM	VNQ	SPXU	LQD	BSX	UPS	MTUM	ZIM	AOR
AMD	IWD	CLSK	VTI	UPST	MODL	CONL	ADMA	COIN	SDS
IVV	TSM	TLT	XLU	SMCI	TCOM	GNOM	META	VFH	SLV
QLD	SVXY	WFC	CVX	FTNT	BBWI	IWX	DBND	PSTG	DFUV
SPXL	SCHG	RIVN	APDN	TSLL	TSLY	ALLY	CGGO	VCLT	XRAY
VOO	BITU	WMT	VYM	AA	RTX	BHP	EW	FETH	UNL
SOXS	UBER	RDFN	GAP	GDXD	MNST	AVIG	SQQQ	KDP	W
VXX	UWM	KO	TFC	PICB	IOVA	ETHE	GUSH	ULTY	NDAQ
GOOG	GDXJ	SPYG	COP	ONON	EQT	BITO	ROKU	UNG	AMLP
UPRO	QID	SPMO	MMIN	MAGS	APA	FMDE	IGM	MOS	CARR
TNA	VTWO	MSFT	MCHP	SPHB	VZ	HIMS	RETL	GIS	QGRW
TSLA	SBUX	ON	ASST	PG	LCDS	EFIV	VIXY	DIG	MSTR
XLK	TSLZ	QCOM	VOXX	FELG	SU	OVV	BBY	CNQ	EXC
UDOW	JDST	JD	XLB	SPYI	FNDX	EWT	JCI	LYFT	AVTR
AMZN	IWFL	CHWY	SLB	MDT	TSLR	XP	USSG	WTAI	WPM
DIA	INTC	UAL	EBAY	RIO	TRFM	EWY	RIOT	X	CHX
FBTC	IBIT	SOXX	AMDY	ZM	IWP	XLF	DLN	MOAT	TGT
SPYU	ITOT	SQ	CFG	USPX	LUNR	PFE	FLEX	IAT	MMM
XLY	BITX	XHB	FITB	QQQG	SPTM	TECK	ORCL	BAX	ARKF
XLI	SARK	CVS	EWW	IWR	SNPE	FNGO	ENVX	BRRR	BILL
NVDL	JNUG	BERZ	BBAX	XEL	CZR	LRGF	VFC	SE	MMLG
UVXY	FNGD	AMDL	FFIE	EWZ	RKT	JGRO	GDS	CLF	NWS
MU	CMG	AVGO	CPRT	TJX	AI	SPSM	BZ	DPST	EMLP
NVDS	DIS	LI	OXY	AEP	GLW	JNJ	KHC	PENN	IUSG
GLD	DKNG	TSLT	HOOD	QUAL	MO	FCOM	BLV	RPG	RXRX
GBTC	QQXT	XME	TSDD	MRK	BK	ABT	SCHB	FTV	RY
TMF	NKE	CSCO	GM	SNOW	AEO	ITB	AR	BOX	KEY
MRVL	XRT	HPQ	AGQ	MGM	NUAG	HAIN	MNTS	AAOI	FDD

NVDX	ACWI	VT	DVN	VRAX	JPM	SPUS	BNRG	SDVY	DJT
RSP	SHOP	UFIV	CCJ	LUV	KR	IJH	SCHV	ESGU	CAG
TSLQ	GILD	PST	DYNF	DBB	BNS	LQDW	HPE	NVS	BE
SSO	VIXM	CONY	HUT	VPL	FBL	JQUA	SYF	CLS	STNE
AFRM	NEE	GDX	SRTY	TZA	STM	TSLS	TECL	CPB	IYG
XLV	IGV	AMZU	PSQ	LVS	O	ZION	CRM	FAST	XLRE
SMH	NUGT	RBLX	XLC	DFSV	IYH	CRH	DFAR	ERX	XRTX
XLE	BITB	DELL	WOLF	SPLG	GSK	CSX	WBA	OARK	PCAR
KRE	RUN	AZN	FL	SOXQ	HAL	DG	JWN	APH	IHDG
XOM	XOP	URTY	BKR	VRT	ETHA	AGGY	ARKG	CLSM	TXN
PYPL	IVW	TD	NVDU	CTSH	CORZ	TOST	NBSM	ICF	FE
XBI	FNGS	NEM	XLP	BILI	NVAX	EWJ	PAAS	SPYX	OMFL
SVIX	ARKK	SCHW	DGRW	NU	SW	EVRG	AEM	BOIL	KWEB
MARA	KBE	USB	PEP	SHEL	BN	WMB	MET	HIBS	DLTR

EXHIBIT 5

Additions: Underlined
Deletions: [Bracketed]

Rules of NYSE Arca, Inc.

* * * * *

Rule 7-E EQUITIES TRADING

* * * * *

Section 1. General Provisions

* * * * *

Rule 7.18-E. Halts

* * * * *

(b) The NYSE Arca Marketplace will not conduct a Trading Halt Auction in a UTP Security and will process new and existing orders in a UTP Security during a UTP Regulatory Halt as follows:

(1) cancel any unexecuted portion of Market Orders, Non-Displayed Limit Orders, MPL Orders, Tracking Orders, Market Pegged Orders, [Discretionary Pegged Orders] Selective Midpoint Orders, RPI Orders, and orders not eligible to trade in the current trading session on the NYSE Arca Book;

* * * * *

(c) The NYSE Arca Marketplace will process new and existing orders in securities listed on the Exchange during a halt, suspension or pause as follows:

(1) cancel any unexecuted portion of Non-Displayed Limit Orders, MPL Orders, Tracking Orders, Market Pegged Orders, [Discretionary Pegged Orders] Selective Midpoint Orders, and RPI Orders;

(2) maintain any unexecuted quantity of Market Orders;

(3) re-price all other resting orders in the NYSE Arca Book to their limit price;

(4) accept and process all cancellations;

(5) reject incoming Limit Orders designated IOC, Non-Displayed Limit Orders, MPL Orders, Tracking Orders, Market Pegged Orders, [Discretionary Pegged Orders] Selective Midpoint Orders, RPI Orders, and Retail Orders; and

(6) accept all other incoming order instructions until the Auction Processing Period for the Trading Halt Auction, at which point, Rule 7.35-E(g) will govern the entry of incoming orders and order instructions.

* * * * *

Section 3. NYSE Arca Marketplace

* * * * *

Rule 7.31-E. Orders and Modifiers

* * * * *

(d) Orders with a Conditional or Undisplayed Price and/or Size

* * * * *

(2) Non-Displayed Limit Order. A Limit Order that is not displayed and does not route. A Non-Displayed Limit Order is ranked Priority 3 - Non-Display Orders. A Non-Displayed Limit Order must be designated Day, is valid for any trading session, and does not participate in any auctions.

(A) The working price of a Non-Displayed Limit Order will be adjusted both on arrival and when resting on the NYSE Arca Book based on the limit price of the order. If the limit price of a Non-Displayed Limit Order to buy (sell) is at or below (above) the PBO (PBB), it will have a working price equal to the limit price. If the limit price of a Non-Displayed Limit Order to buy (sell) is above (below) the PBO (PBB), it will have a working price equal to the PBO (PBB).

(B) A Non-Displayed Limit Order may be designated with a Non-Display Remove Modifier. If so designated, a Non-Displayed Limit Order to buy (sell) will trade as the liquidity-taking order with an Aggressing ALO Order or MPL-ALO Order to sell (buy) that has a working price equal to the working price of the Non-Displayed Limit Order, or with a Liquidity Providing Selective Midpoint Order to sell (buy) as specified in Rule 7.31-E(h)(3)(D) and the subparagraphs thereunder.

(3) Mid-Point Liquidity Order ("MPL Order"). A Limit Order to buy (sell) that is not displayed and does not route, with a working price at the lower (higher) of the midpoint of the PBBO or its limit price. An MPL Order is ranked Priority 3- Non- Display Orders. MPL Orders are valid for any session and do not participate in any auctions.

* * * * *

(F) MPL Orders designated Day may be designated with a Non-Display Remove Modifier. If so designated, an MPL Order to buy (sell) will trade as the liquidity-taking order with an Aggressing ALO Order or MPL-ALO Order to sell (buy) that has a working price equal to the working price of the MPL Order, or with a Liquidity Providing Selective Midpoint Order to sell (buy) as specified in Rule 7.31-E(h)(3)(D) and the subparagraphs thereunder.

* * * * *

(e) Orders with Instructions Not to Route

(1) Non-Routable Limit Order. A Limit Order that does not route.

* * * * *

(C) Non-Routable Limit Order may be designated with a Non- Display Remove Modifier. If so designated, a Non-Routable Limit Order to buy (sell) with a working price, but not display price, equal to the working price of an Aggressing ALO Order or MPL-ALO Order to sell (buy) or within the discretionary range of a Liquidity Providing Selective Midpoint Order to sell (buy) will trade as the liquidity taker against such order.

* * * * *

(h) Pegged Orders. A Limit Order that does not route with a working price that is pegged to a dynamic reference price. If the designated reference price is higher (lower) than the limit price of a Pegged Order to buy (sell), the working price will be the limit price of the order.

* * * * *

(3) [Discretionary Pegged Order] Selective Midpoint Order. A Pegged Order to buy (sell) that upon entry to the NYSE Arca Marketplace is assigned a working price equal to the lower (higher) of the midpoint of the PBBO ("Midpoint Price") or the limit price of the order. Any untraded shares of such order are assigned a working price equal to the lower (higher) of PBB (PBO) or the order's limit price and is automatically adjusted in response to changes to the PBB (PBO) for buy (sell) orders up (down) to the order's limit price. In order to trade with contra-side orders on the NYSE Arca Book, a [Discretionary Pegged Order] Selective Midpoint Order to buy (sell) will exercise the least amount of price discretion necessary from its working price to its discretionary price (defined as the lower (higher) of the Midpoint Price or the [Discretionary Pegged Order's] Selective Midpoint Order's limit price).

(A) [Discretionary Pegged Orders] Selective Midpoint Orders are not displayed, must be designated Day, and are eligible to be designated for the Core Trading Session only.

[Discretionary Pegged Orders] Selective Midpoint Orders that include a designation for the Early Trading Session or Late Trading Session will be rejected.

(B) When exercising discretion, [Discretionary Pegged Orders] Selective Midpoint Orders maintain their time priority at their working price as Priority 3 - Non-Display Orders and are prioritized behind Priority 3 - Non-Display Orders with a working price equal to the discretionary price of a [Discretionary Pegged Order] Selective Midpoint Order at the time of execution. If multiple [Discretionary Pegged Orders] Selective Midpoint Orders are exercising price discretion during the same book processing action, they maintain their relative time priority at the discretionary price.

(C) A Selective Midpoint Order is eligible to exercise price discretion to its discretionary price, except during periods of quote instability as identified by the Selective Midpoint Indicator ("SMI"). If the SMI determines the PBBO for a particular security to be an unstable quote, both an arriving and resting Selective Midpoint Order will wait for a PBBO that is stable before the order's working price is adjusted and the order becomes eligible to trade.

The SMI reflects a proprietary assessment of market conditions conducted using supervised learning models developed by the Exchange. The models rely on a combination of SIP data and NYSE Arca Book data and incorporate up to 83 features to predict periods of market instability, with the objectives of maximizing recall and precision while minimizing overlocking. The Exchange will retrain the models underlying the SMI daily, outside of the Core Trading Session, based on data from the preceding three days and will deploy updated version(s) of the model(s) if retraining indicates that such updated version(s) would improve performance in achieving the above objectives. More information on the SMI is available in the white paper attached as Exhibit 3 to SR-NYSEARCA-2024-112, which is available on the Exchange's website.

(D) A Selective Midpoint Order may be designated as Liquidity Providing ("Liquidity Providing Selective Midpoint Order").

(i) A Liquidity Providing Selective Midpoint Order will only execute on arrival against resting orders that include a Non-Display Remove Modifier and are priced within the Liquidity Providing Selective Midpoint Order's discretionary range.

(ii) If a resting contra-side order that does not include a Non-Display Remove Modifier is priced within an arriving Liquidity Providing Selective Midpoint Order's discretionary range, the Liquidity Providing Selective Midpoint Order will be placed on the NYSE Arca Book, and its discretionary range will be adjusted to equal the resting price of a non-displayed contra-side order or to one MPV less aggressive than the resting price of a displayed contra-side order.

(iii) A Liquidity Providing Selective Midpoint Order will not trade with an arriving contra-side order that cannot remove liquidity. Once the arriving contra-side order is placed on the NYSE Arca Book, the discretionary range of the Liquidity Providing Selective Midpoint Order will be adjusted to equal the resting price of a non-displayed contra-side order or to one MPV less aggressive than the resting price of a displayed contra-side order.

(iv) Once resting on the NYSE Arca Book, the discretionary range of a Liquidity Providing Selective Midpoint Order will be adjusted based on resting contra-side interest as described in subparagraphs (ii) and (iii) when its working price changes. A resting Liquidity Providing Selective Midpoint Order to buy (sell) will not be eligible to trade:

(a) At a price equal to or above (below) any sell (buy) orders that are displayed and that have a working price equal to or below (above) the working price of such Liquidity Providing Selective Midpoint Order; or

(b) At a price above (below) any sell (buy) orders that are not displayed and that have a working price below (above) the working price of such Liquidity Providing Selective Midpoint Order.

[(C)] (E) If the PBBO is locked or crossed, both an arriving and resting [Discretionary Pegged Order] Selective Midpoint Order will wait for a PBBO that is not locked or crossed before the working price is adjusted and the order becomes eligible to trade.

(i) Additional Order Instructions and Modifiers

* * * * *

(3) Minimum Trade Size ("MTS") Modifier. A Limit IOC Order, Non-Displayed Limit Order, MPL Order, Tracking Order, [Discretionary Pegged Order] Selective Midpoint Order, or non-displayed ALO Order may be designated with an MTS Modifier.

* * * * *

Rule 7.34-E. Trading Sessions

* * * * *

(c) *Orders Permitted in Each Session.*

(1) Early Trading Session. Unless otherwise specified in paragraphs (c)(1)(A) - (F), orders and modifiers defined in Rule 7.31-E that are designated for the Early Trading Session are eligible to participate in the Early Trading Session.

(A) Market Orders and Pegged Orders are not eligible to participate in the Early Trading Session. Market Orders and Pegged Orders that include a designation for the Early Trading Session will be rejected. Market Pegged Orders and [Discretionary Pegged Orders] Selective Midpoint Orders, regardless of the session designated for the order, may not be entered before or during the Early Trading Session and will be rejected.

* * * * *