

August 11, 2023

Vanessa Countryman, Secretary
Securities and Exchange Commission
100 F Street NE
Washington DC

Re: File No. SR-NASDAQ-2022-079

Dear Ms. Countryman:

Thank you for the chance to comment. This is in response to Nasdaq's July 18, 2023 letter ("Letter"), attached to which is a draft filing ("Filing"), Nasdaq's most recent amendment. On January 31, 2023, May 2, 2023, and May 30, 2023, I submitted comments on previous versions of this proposal. Those comments are available at <https://www.sec.gov/comments/sr-nasdaq-2022-079/srnasdaq2022079.htm> and I include them here by reference. As of today the Filing hasn't posted formally as an amendment on the SEC's website or in the Federal Register, and with the September 7, 2023 deadline for SEC action (Release 34-97844) just a few weeks away it isn't clear to me - or to the public - if the Filing will post before that deadline or if the calendar can or will be extended.

The main points of my previous comments were that I believed Nasdaq's proposal for a new order type called Dynamic M-ELO was an impermissible use of discretionary methods by an exchange and directly contravened 3b-16, that Nasdaq's proposed rule text failed to define the novel, complex, and consequential operations of the order type and failed to meet exchange standards, that Nasdaq's filings failed to meet the SEC's standards for 19b-4 filings, that supplemental materials Nasdaq provided as exhibits were inaccessible and of uncertain governance, that filing texts and exhibits were inconsistent on fundamental operations, that bias was a known risk of the technology Nasdaq proposed and Nasdaq had not analyzed its test results for 6(b)(5) compliance or described how it would monitor for those concerns in production, and that a participant's own current or past behavior could unpredictably influence Nasdaq's discretionary methods, even to degrade a participant's or its counterparty's outcomes.

Compared to earlier versions, the Filing updates and clarifies some points. Other updates are flabbergasting, as I discuss below.

Over the last half century and because of Congressional, SEC, and industry initiatives, new technologies steadily advanced market transparency, determinacy, and fairness. I believe the Filing proposes new exchange technology - reinforcement machine learning, an artificial intelligence deep learning method - that will reverse transparency and determinacy and put fairness at risk. My view about Dynamic M-ELO hasn't changed. The SEC should disapprove.

Sparkling

Since January I've pointed out that while Nasdaq's white paper said its researchers looked at whether Dynamic M-ELO was biased against individual firms and believed it wasn't, Nasdaq hadn't explored whether that was true for participant types. A well known risk of reinforcement machine learning is that these technologies, as Gary Gensler and Lily Bailey carefully describe in their paper "Deep Learning and Financial Stability," can exhibit forbidden group biases. They wrote about bias in consumer finance. I believe the risk is the same here. Gensler and Bailey wrote that "Technical and regulatory approaches to the challenges of bias have yet to fully

emerge for deep learning models," and as a first step to meet these challenges I've said Nasdaq should analyze Dynamic M-ELO for participant type biases. As I wrote in my May 2, 2023 comment:

If certain participant types show distinct behavioral patterns - like order and trade sizes and counts, or order submission strategies - they might have enough of an influence on Dynamic M-ELO to move holding timers in one direction or another, even to the advantage of some participant types over others. Nasdaq says the "model is objective and designed to avoid bias and discrimination" but so far as I can tell hasn't analyzed the data to answer this question. Nasdaq's apparent lack of curiosity about this is bewildering. Do all participant types see similar fill rate and mark out results day-by-day, week-by-week, and month-by-month over the entire test set? Do some participant types see better results than others in certain periods? How will Nasdaq monitor for this kind of bias in production, and what will Nasdaq do if it turns up?

In the Filing, Nasdaq now says it explored the question and didn't find bias. But what Nasdaq said about how it looked at the question is remarkable.

Furthermore, Nasdaq performed internal tests of its AI model to detect indications of harmful bias in its performance results, and such tests concluded that no such indications exist. That is, the Exchange reviewed the impact on fill rates and mark-outs of Dynamic M-ELO, as compared to the "static" M-ELO, *for those firms that accounted for more than 95% of M-ELO activity on the Exchange during Q1 2022.*[33] The Exchange analyzed results both in an absolute and a relative sense. Testing revealed that all participants experienced at least some improvements in fill rates and mark-outs when using Dynamic M-ELO versus static M-ELO, with the volume-weighted average improvement being aligned with the results expressed in the White Paper. We detected no *material variations* that might suggest that a particular participant or category of participant (i.e., nature of firm; size of firm) benefitted from Dynamic M-ELO functionality to an extent that was *unreasonably disproportionate* to the benefits that other participants experienced. Thus, Nasdaq believes the model is objective, is designed to, and does avoid bias and discrimination. (Letter, p23-24) (Emphasis added.)

And Nasdaq wrote this:

[33] Beyond this grouping of participants, the activity levels of other individual M-ELO participants were so small as to be insignificant. In many cases, these participants entered only a handful of M-ELOs during the study period. As such, *the Exchange believes it is reasonable to exclude such participants from its analysis to avoid their data distorting the results.* (Emphasis added.)

It seems Nasdaq trimmed its data before analyzing it for bias and constrained its analysis to "those firms that accounted for more than 95% of M-ELO activity on the Exchange during Q1 2022." (I assume Nasdaq used the same data defined in the Filing as the "Training Period" for its analysis. Nasdaq doesn't say so, however.) Nasdaq doesn't describe the kind of M-ELO "activity" it filtered the data for, and specifically whether it filtered on order or trade counts or order or trade volume or some combination of two or more of these categories, or on some other factor, before removing firms from its analysis.

If Nasdaq filtered on order activity, it's worth pointing out that even 5% of order activity could be materially different from other orders and the firms sending them depending on the

characteristics of that flow. For example, Nasdaq's monthly M-ELO reports (available at <https://www.nasdaq.com/docs/M-ELO-Monthly-Report>) show that, in the months I examined, all orders of 1,000 shares or more - very large orders in today's market - are merely 5% to 7% of total M-ELO order counts, though they could be as much as half of all M-ELO order volume (my own back-of-the-envelope estimate). If Nasdaq filtered on trades, Nasdaq's monthly M-ELO reports show that even 5% of trade activity could be materially different from other trades and the firms behind them depending on the characteristics of those trades. Nasdaq's monthly M-ELO reports I examined show that trades of 201 shares or more are 5% to 7% of trade counts (identified as "Passive Executions" on the report), though they could be as much as a third of all M-ELO trade volume (my own back-of-the-envelope estimate). And if Nasdaq used the Training Period data for this analysis, remember that the Filing defines Training Period data as just "380 out of the 6,257 symbols on the M-ELO Book (accounting for approximately 67 percent of M-ELO volume)," so I wonder if Nasdaq filtered the data not once but twice from the full set of Q1 2022 M-ELO activity before Nasdaq examined it for 6(b)(5) concerns. More on this in "Hogwash," below.

I also wonder if Nasdaq culled the data to handle what Gensler and Bailey describe as one of the "inherent characteristics of deep learning. The non-linearity and hyper-dimensionality of deep learning models make them likely to be more sensitive within a certain input range, but less likely to perform well outside of that range." They go on to write "deep learning models may more frequently end up in local minima and have larger out of range tails, leading to so-called 'fat-tails' and a higher prominence of 'black swan' events," and they worried about "deep learning models' inherent challenges of robustness." (Gensler and Bailey, p24). What does this mean? It means that even if deep learning algorithms perform well on typical behavior, they may perform poorly on behavior outside of that range.

Step-by-step, how did Nasdaq select the 380 symbols for the Training Period data? And then step-by-step, how did Nasdaq choose firms to exclude for this analysis? Nasdaq hasn't said. Nasdaq does say "the Exchange believes it is reasonable to exclude such participants from its analysis *to avoid their data distorting the results*," (emphasis added) which is about as big of a gobsmacker as I've ever read in a 19b-4, and I hope caused spit-takes across the DC metro area. What a sparkling world it is when we avoid distorting the results. If everyone were free to trim data to avoid distorting the results we could joyously pretend the U.S. has never had a recession, Silicon Valley Bank is still in business, there was no Flash Crash, the American League won the All Star game this year, and I have six pack abs. I also wonder if limiting development and testing in this way goes against the spirit, if not the letter, of Reg SCI.

Hogwash

To be clear, I don't know what Nasdaq excluded and whether Nasdaq excluded firms with large orders and trades from the analysis. My point is that discarding any data could exclude activity that has qualitative or quantitative differences from the rest, and especially so if, as seems likely, that data is the Training Period data and already culled to 67% of all Q1 2022 M-ELO activity. Nasdaq says so, too - that the data is qualitatively or quantitatively different - when it says it removed those firms to avoid "distorting the results." Of course it's possible that with a close look there's evidence of bias in any so-called distorted results, or we might find that Dynamic M-ELO improves participant outcomes only on the most liquid instruments or only for odd-lots, or that Dynamic M-ELO doesn't make much difference at all - or not, we don't know. All we know is what Nasdaq told us, that it removed firms from its analysis because Nasdaq thought those firms distorted its results. In what way? Without more information about the data and whether and how Nasdaq verified its remaining data accurately reflected all Q1 2022 M-ELO

activity, I believe this analysis and Nasdaq's conclusion that "the model is objective, is designed to, and does avoid bias and discrimination" is hogwash.

(If Nasdaq used Training Period data for this analysis, trimming a dataset to 95% that's already been trimmed to 67% can suggest that the data has been repeatedly culled to exclude data that is one or more standard deviations away from an average. Assuming a normal distribution, approximately 68% of a data series is within one standard deviation of a mean and approximately 95% is within two standard deviations of a mean. Nasdaq should clarify whether that's what it did. Trimming data one or more standard deviations away from a mean is one way to remove the kinds of fat-tails and black swans Gensler and Bailey warn about.)

Because of the novel, complex, and consequential methods Nasdaq has proposed in the Filing and their known risks - risks the SEC Chairman has carefully outlined in published research - Nasdaq should show all of its work on this question. The SEC and the public shouldn't simply accept an SRO's happy report about these matters. From 201.700(b)(3)(i) to *Susquehanna* to common sense, an SRO should provide enough data in enough detail so the public and the SEC can understand, verify, and comment on the SRO's methods and substantive conclusions. In *Susquehanna*, the D.C. Circuit wrote "the SEC should have critically reviewed [the SRO's] analysis or performed its own," and quoted from *NetCoalition* where it found a "lack of support in the record" and that "The SEC had tried to rely on statements by the self-regulatory organization, but we saw 'little' supporting value in the 'self-serving views of the regulated entit[y].'" (866 F.3d at 447) Later the SEC itself wrote "The D.C. Circuit's *Susquehanna* Opinion makes clear that relying on such representations, without more, is insufficient. Rather, the Commission must critically evaluate the representations made and the conclusions drawn by [the SRO]." (Release 34-85121, p2.)

Imagine if an S-1 or 10-K took the same approach - "we exclude certain operating units to avoid distorting the results" or "our income statement eliminates a range of expenses to avoid distorting the results" or "the balance sheet drops various liabilities to avoid distorting the results." The concern isn't that Nasdaq designed Dynamic M-ELO to discriminate against participant types, it's that Dynamic M-ELO could do it anyway. If Nasdaq used Training Period data, excluding nearly 40% of the data from an analysis (95% of 67% leaves 63.65%) is no way to confirm Dynamic M-ELO is free from 6(b)(5) problems. By itself I believe this is a compelling reason to disapprove, particularly in light of Nasdaq's qualified and unsubstantiated conclusions that it detected no "*material variations*" (emphasis added) or "*unreasonably disproportionate*" (emphasis added) benefits among participant types. The SEC and the public should see all the results.

One way to proceed here is for Nasdaq to (a) train Dynamic M-ELO on a full and unfiltered set of data for Q1 2022, not just 380 of 6,257 symbols, (b) explore the full and unfiltered set of Dynamic M-ELO's results for any evidence of participant type bias, and then (c) clearly and comprehensively present all its work to the SEC and the public, "distorted" results included. Another way to proceed here is disapproval.

Exchanges, Discretion, and Certainty

The Exchange Act at 3(a)(1) defines an exchange as "a market place or facilities for bringing together purchasers and sellers of securities or for otherwise performing with respect to securities the functions commonly performed by a stock exchange as that term is generally understood." As was generally understood at the time, there's no room for a "stock exchange" to routinely intervene and set the material terms of buy and sell orders. There's also no

suggestion an exchange can be anything more than a strictly neutral forum for buyers and sellers to conduct their business together, and according to their own terms.

That neutrality has been underscored time and again by the SEC in think pieces, policy statements, and new rules. In the Market 2000 report the SEC wrote "The Commission has determined this function [an exchange] to be the provision of a trading market that is designed, whether through trading rules, operational procedures, or business incentives, to centralize trading..." which, again, defines an exchange as a neutral facility tasked simply with "the provision of a trading market" for buyers and sellers to conduct their business according to their own terms. (Market 2000, Study III - 12.) In that context, four years after Market 2000 the SEC adopted Reg ATS. Rule 3b-16(a) defines an exchange as an entity that "(1) Brings together the orders for securities of multiple buyers and sellers; and (2) Uses established, non-discretionary methods (whether by providing a trading facility or by setting rules) under which such orders interact with each other, and the buyers and sellers entering such orders agree to the terms of a trade" and again there is no suggestion that an exchange is anything but a neutral marketplace that "brings together the orders" of its participants and uses "established, non-discretionary methods" under which those buyers and sellers will conduct their business according to their own agreed terms. When Reg ATS was adopted, there wasn't even a hint an exchange would alter the material terms of a participant order based on its own undisclosed and unpredictable judgments and its own undisclosed data, as Nasdaq proposes now with Dynamic M-ELO.

The record continues. In its IEX Form 1 approval, the SEC visited the question of what "established, non-discretionary methods" means. There was some controversy at the time whether one of IEX's proposed order types was an impermissible discretionary method for an exchange. The SEC wrote "IEX has thus encoded in its rule the *totality* of the discretionary feature of its proposed discretionary peg.... the Commission does not believe that the *hardcoded* conditionality of the IEX proposed 'discretionary' peg order type provides IEX with actual discretion or the *ability to exercise individualized judgment* when executing an order." (Emphasis added.) (Release 34-78101, pp44-45.) Further, the order type was distinguished by a feature "based on a *pre-determined, objective set of conditions* that are detailed in IEX's proposed rule." (Emphasis added.) (Release 34-78101, p41.)

Merriam-Webster's defines "totality" as "the quality or state of being total" and "total" as "comprising or constituting a whole." As I pointed out in my May 2 comment, for all time in the computer age "hardcoded" has meant that programmers specifically tell a system what to invariantly do in a circumstance. Merriam-Webster's defines "individualized" as "to make individual in character" and "judgment" as "the process of forming an opinion or evaluation by discerning and comparing." It defines "pre-determined" as "to determine beforehand" and "objective" as "expressing or dealing with facts or conditions as perceived without distortion by personal feelings, prejudices, or interpretations." All to say, the Filing fails to meet the plain language used by the SEC in its IEX approvals.

Later, IEX proposed the D-Limit order type, and the SEC again visited the question of what "established, non-discretionary methods" means. In its D-Limit approval, the SEC doubled-down and wrote that the feature distinguishing the order type and a source of controversy, was allowable because it was based on a "*pre-determined, objective set of conditions that are detailed in IEX's [rules]*" and which any market participant can thus recreate on its own." (Release 34-89686, p28.)

I believe Nasdaq now proposes to flip this long and important history inside out, and maintain that a rule that says, in effect, that Dynamic M-ELO will make its own private, proprietary,

varying, and unpredictable determination of an order's material terms is an established, non-discretionary method just because the rule says that's what Nasdaq intends to do. It's as if Nasdaq proposed a rule saying it will use discretionary methods, and then with a straight face claimed the rule was an established, non-discretionary method simply because it wrote it down.

In contrast to IEX and its D-Limit, if approved by the SEC, Nasdaq's rulebook won't set out the "totality of the discretionary feature" (I believe it can't, because the totality changes week-to-week and even minute-to-minute) and it won't define the "hardcoded conditionality" of its feature (again, I believe it can't), and a market participant won't be able to "recreate on its own" what Dynamic M-ELO has done (participants can't - it's not even clear anyone will be able to, as discussed below). I believe Dynamic M-ELO will also "exercise individualized judgment" such that it can set a different time-in-force for the very same order presented in the very same market conditions on August 21 than it set on May 15, depending on the system's undisclosed individualized judgments of market conditions and participant behavior from even days or weeks in the past, characteristics of this technology Nasdaq has not disputed.

Reg ATS

I have covered Reg ATS in my previous comments and include them here by reference.

In the Filing, Nasdaq offers this from the adopting release:

A system uses established non-discretionary methods either by providing a trading facility or by setting rules governing trading among subscribers. The Commission intends for "established, non-discretionary methods" to include any methods that dictate the terms of trading among the multiple buyers and sellers entering orders into the system. Such methods include those that set procedures or priorities under which open terms of a trade may be determined. For example, traditional exchanges' rules of priority, parity, and precedence are "established non-discretionary methods," as are the trading algorithms of electronic systems. Similarly, systems that determine the trading price at some designated future date on the basis of pre-established criteria (such as the weighted average trading price for the security on the specified date in a specified market or markets) are using established, non-discretionary methods.

Without any analysis, Nasdaq simply asserts "Nothing in the Reg. ATS Adopting Release or in any of its illustrative examples suggests that Dynamic M-ELO would constitute an exercise of discretionary behavior." So, first, I'll offer the same quote and assert that *everything* in the Reg ATS Adopting Release suggests Dynamic M-ELO is an exercise of discretionary behavior.

I won't leave it there. Let's keep in mind that when this was written in 1998 exchange trading was done in fully automated systems (ECNs, Cincinnati, SuperDOT, SOES and the like), semi-automated systems (SelectNet, ITS, and others), and manually on trading floors. The Reg ATS text above ("Text") captures all these different forums with "A system" and begins by saying "A system uses established non-discretionary methods either by providing a trading facility or by setting rules governing trading." At the time and through to today, fully automated and semi-automated exchange systems were deterministic - day after day and week after week, presented with the same market data and interactions, exchange software did the same thing, invariantly. The Text then says "or by setting rules," which was intended to capture floor-based trading.

Next, "The Commission intends for 'established, non-discretionary methods' to include any methods that dictate the terms of trading among the multiple buyers and sellers entering orders into the system. Such methods include those that set procedures or priorities under which open terms of trade may be determined." And then to be more explicit about "procedures or priorities" the Text says "For example, traditional exchanges' rules of priority, parity, and precedence are 'established non-discretionary methods'" - intended to capture floor-based trading - "as are the trading algorithms of electronic systems" - intended to capture automated trading. In 1998, exchange "rules of priority, parity, and precedence" were deterministic. So far as I know, no exchange at the time allowed buyers and sellers to deviate from its rules of priority, parity, and precedence, in either their fully automated, semi-automated, or floor-based facilities, unless the exchange detailed an exception. Exceptions were deterministic too. Of course, for manual markets this was a farce. Specialists, market makers, and floor brokers deviated from the rules all the time, but, formally, exchange rulebooks "set procedures or priorities" such that priority, parity, and precedence were deterministic. For automated systems, however, this was not a farce. "[T]rading algorithms of electronic systems" were invariant and deterministic day after day, week after week.

I'm not sure why Nasdaq included the Text from here on, but Nasdaq time and again has handed me an axe and does it again, "Similarly, systems that determine the trading price at some designated future date on the basis of pre-established criteria (such as the weighted average trading price for the security on the specified date in a specified market or markets) are using established, non-discretionary methods." This shows the SEC's meaning and the meaning is squarely against Nasdaq. A weighted average trading price method was invariably deterministic in 1998 and is today. Presented with the same instant market data, the method will produce the same result, every time. Market participants know exactly how the price will be calculated and can verify the calculation themselves. What the SEC meant here in 1998 has even been plainly echoed in its IEX approvals, decades later. In 1998, especially with respect to automated systems, "pre-established criteria" and "established non-discretionary methods" meant, first and foremost, deterministic and invariant behavior, behavior which exchanges and participants could predict and verify in every detail.

It's not certain

Though there have been some rare and trivial exceptions, since the start of automated exchange trading, order type behavior has been deterministic, invariant, and immediately auditable and explainable. In that context the SEC adopted Reg ATS and it surely informed the SEC when it wrote that an exchange should operate with "established, non-discretionary methods." Presented with the same market conditions at a point in time those systems delivered the same results, over and over through time, and exchanges and participants could confirm that systems behavior followed the rules. Even with discretionary and reserve orders and other partially disclosed or undisclosed order types and facilities, exchange rulebooks described the totality of how those features worked and in such specificity that participants could either confirm they were treated correctly according to those rules or call the exchange for an accounting. Exchange support can review an order's history and deliver a full accounting, almost always on the spot. System errors were, generally, immediately apparent because participants (or the exchange) could tell whether orders were processed and executed in accordance with the exchange's rulebook.

As the industry moved on from the insanity (and inanity) of all the flavors of manual or semi-automated trading facilities over the years, for the most part exchange trading rules earned a reputation for transparency and determinacy. This wasn't accidental. It flowed from

the principle that exchanges were neutral facilities that simply provided a forum for buyers and sellers to meet and trade according to "established, non-discretionary methods."

Dynamic M-ELO departs from decades of this progress. Its behavior will not be deterministic or invariant over time, and purposefully so. Far from being hardcoded and deterministic, in its March 9, 2023 letter at page two Nasdaq said its system "trains itself... without the programmers telling the system in advance exactly what to do." As I've noted in previous letters, my Uncle Lou trains himself too. No one tells him what to do either. Because of it my Uncle Lou can never be an exchange, even if he puts his picture in a rulebook. He could be a broker though if that's what he wants. In its May 18, 2023 letter at page nine, Nasdaq bellyaches that disapproval because of 3b-16 would "create a new regulatory disparity and erect a new barrier to fair competition," but of course Nasdaq is free to abandon the many unique and exclusive privileges and revenue streams it enjoys as an exchange, turn itself into a broker, and use all the reinforcement machine learning and discretionary methods it can dream up. But I imagine Nasdaq would rather continue with those privileges and revenues as it asks the SEC for the latitudes of a broker. Or, as I wrote on May 30, 2023, to have its exchange cake and eat brokerage cake too.

If the Filing is approved, an exchange will not be limited to a neutral forum for buyers and sellers to meet according to established, non-discretionary methods. Even if an approval order constrains the scope of these technologies, over time exchanges will push against those constraints. Exchanges will file to make their own varying and uncertain judgments of current and past market conditions and their own varying and uncertain judgments of current and past participant behavior. Exchanges will file to use their own undisclosed data in those judgments and then unpredictably determine the material terms of orders under their control, based on whatever undisclosed scheme in which they expressed their objectives and in whatever undisclosed and uncertain manner their software interpreted those objectives. In this new world, participants and exchange support will only be able to guess at what behavior or market conditions determined an order's handling and execution. As Gensler and Bailey point out at page 12, no one will know what on earth those systems are doing or why.

At page 13 of their paper Gensler and Bailey write "Deep learning models' decisions and outcomes are often unexplainable," and then "explainability generally captures the notion that decisions and outcomes of a model can be explained to customers, management, and regulators." At page two of its March 9, 2023 letter Nasdaq wrote that "We note that it is not feasible to provide a precise and definitive description of how the system will react to all manner of changes to these data elements, either individually or in combination, because machine learning technology does not operate in a strictly linear fashion," I believe conceding the point for Dynamic M-ELO.

In the Filing Nasdaq writes "Dynamic M-ELO will handle and execute Orders according to published, pre-determined rules that are disclosed to the public and which provide reasonable notice of how the Order Type will behave." When a customer calls to complain that a stock moved against them immediately after a fill and asks why Nasdaq set the holding period on their order to one millisecond instead of two milliseconds, what will Nasdaq tell them? Specifically, what will Nasdaq tell the customer about how Nasdaq handled the order and set its holding period? When a customer asks for Nasdaq to verify a holding period was calculated correctly, what will Nasdaq do to verify it? If Nasdaq is confident its "published, pre-determined rules" provide reasonable notice of how Dynamic M-ELO works, this should be easy. Or when a customer calls to ask what a holding period will be for a new order and why, what will Nasdaq

say about that? Let's imagine how an exchange might handle a customer inquiry. Here's an absurdist vignette I'll call "Waiting for M-ELO":

Customer: So, uh, what's the holding period? 250 mics, right?

Nasdaq: It's not certain.

Customer: What?

Nasdaq: Nothing is certain.

Customer: What am I to say?

Nasdaq: Say, I am happy.

Customer: I am happy.

Nasdaq: We are happy.

Customer: I can't go on like this.

Nasdaq: That's what you think.

At page 26 of their paper, Gensler and Bailey write "Deep learning models are less explainable, more complex, and more dynamic than other models. Accordingly, market transparency is likely to decrease, due to both explainability and complexity challenges." For decades, new technology enhanced market transparency and determinacy. The Filing pitches technology that will reverse transparency and determinacy. For what? To scratch a techno-onanist itch? I don't even know what to make of Nasdaq's performance claims because Dynamic M-ELO was developed and tested on a subset of the data.

Data

Nasdaq's July 18, 2023 letter includes another version of the data element list. Containing varying details and enigmatic color schemes, I've described previous lists as "vague, confusing, and perfunctory" and pointed out problems I had with them, including what I believed were vague definitions, uncertain or undefined terms, and apparent disagreements with other filing materials. To date we've seen three of these lists, the first on March 9, 2023 ("Version 1"), another on May 18, 2023 ("Version 2"), and then another on July 18, 2023 ("Version 3"). Apart from some minor revisions and the odd transcription error, data names and descriptions in Version 3 are the same as those in Version 2. (Data descriptions in Version 2 were significantly different from Version 1.)

To start, the Filing states at note 14 that the list includes "an observation of the strength of the correlations that currently exist between changes to those data values and decisions the system makes," but that column went missing in Version 2 and is still missing in Version 3. And then, as an example of why I believe these lists have been vague, confusing, and perfunctory - and an example I chose at random for my previous comments on the Filing - I tried to make sense of a data field called "last_3step_incl_lookback_BUY_trade_qty_mean." In Version 1 it was described as "mean(buy # shares per trade executed in the last 90 seconds)." In Version 2 and Version 3 the description changed to "The average buyside trade quantity that came from the simulated trades made by the agent's action over the last 90 seconds, per ticker."

In my May 30, 2023 comment I said "buyside trade quantity" and other terms in the new description were mysteries to me. I wrote that I didn't know what "buyside trade quantity" meant since a buy trade quantity is always equal to the sell quantity of a trade, and how "buyside trade quantity" was distinguished from a simple "trade quantity" was unclear and confusing. The data element list now includes a brief glossary apparently intended to explain this and other mysteries. The glossary explains that a "Buyside trade attribute" is the "Evaluation of trade attributes corresponding to the perspective of the buyer of the instrument." That definition is

more than a little circular, reduces to "a buy-side trade quantity is the quantity of a buy-side trade," and seems to me like a Bronx cheer to the effort. On a trade, the trade quantity is the trade quantity. The buyer on a trade doesn't have its own "perspective" of a trade quantity. Neither does a seller. To give an idea of how strange this term is and how unusual it is even for Nasdaq, I can't find the term "buy-side trade" anywhere else on nasdaq.com or nasdaqtrader.com or for that matter on sec.gov. So again, what does "buy-side trade" mean? Is this a Gen Z thing?

I also wondered what "simulated trades" were and what "the agent's action" meant since the filing text didn't explain these terms. The glossary doesn't define "simulated trade" but points to section 3.1 of the white paper to explain "simulation." Section 3.1 of the white paper describes generally the training process and how the system will measure success, and we're left to guess that a "simulated trade" is a hypothetical trade that might come out of this process as the system trains by matching orders and executing trades in a simulated trading environment. But as for what "buy-side trade quantity" means and how it is distinguished from "trade quantity" and "sell-side trade quantity"? No idea.

Wild goose chase

What's frustrating in this wild goose chase spanning Version 1, Version 2, and Version 3 is to see again how deeply inappropriate I believe it is for the Filing - as well as the rule text itself now - to point to the white paper for details which belong in the filing and rule texts. As I wrote in my May 30, 2023 comment "it's unclear whether [the white paper] describes a research project to evaluate whether M-ELO could be improved, a development process to sharpen any improvements, or something that Nasdaq will deploy to production in the National Market System."

I find those lines blurred almost everywhere the white paper is referenced. The data element glossary is just one example. The Filing text itself is another. For example, note 18 of the Filing says "The Exchange provides further information about the algorithm and the retraining process in a White Paper attached hereto as Exhibit 3a." One question then is will retraining in production be done on only the 380 symbols and 67% of M-ELO data the white paper describes using to retrain? (See section 4.2.) Or will the system in production retrain on all M-ELO data? If Nasdaq will retrain on all M-ELO data instead of a subset of it in production, and since that's a fundamental aspect of the system and very different from what's described in the white paper, it deserves some ink somewhere. If it will retrain in production on a subset of M-ELO data, however, that deserves some ink too. As far as the data elements go, in production will all of the "simulated" fields be derived during retraining on a filtered subset of M-ELO data? The word "simulated" appears in almost 90 of the 142 data elements on the list.

We've also gone round and round in a wild goose chase about whether the system will retrain daily, weekly, monthly, or quarterly, and while the Filing and new rule text specify weekly retraining, both these texts also point to the white paper for further clarification, and daily, weekly, monthly, and quarterly cycles are all highlighted in the white paper. For such a novel, complex, and consequential initiative we should expect more care than this.

Why not

Thanks to the glossary, it seems the vast majority if not all these data elements are calculated after hours during retraining cycles. But *something's* going on during market hours. During market hours, what "market conditions" does Dynamic M-ELO use or calculate that tell it to change holding periods during a "Change Event"? That's not explicit in the data element list or anywhere else, from what I can tell. The description for "last_3step_incl_lookback_BUY_trade_qty_mean" in Version 1 was "mean(buy # shares per trade executed in the last 90 seconds)." At first I assumed that description meant the field was calculated during the trading day. In Version 2 and Version 3 the field is newly described as "The average buy-side trade quantity that came from the simulated trades made by the agent's action over the last 90 seconds, per ticker." Because of the terms "simulated" and "agent's action" in that description the glossary in Version 3 suggests to me the field is calculated during after hours weekly retraining cycles. So which, if any, of the data fields on the list does Dynamic M-ELO calculate during the trading day to assess market conditions and change holding periods? Or is that data not on the list?

Finally, I've noted twice before the field name is "last_3step_incl_lookback_BUY_trade_qty_mean" but there was no description of what "last_3step" means. There still isn't. And - why not - I'll point out again that the Filing notes "the number of resting bids left in the M-ELO and M-ELO+CB Book" influences the model, as does the white paper ("the number of resting M-ELO bids left"), but I can't find this field on any version of the data element list.

Rule Text

The Filing at last includes rule text with language attempting to define "proprietary assessment of market conditions." This new language fails. Like the Filing text, the rule text points to the white paper for "further information about the algorithm and the retraining process." The white paper describes a research project to study whether M-ELO could be improved. It does not describe something that was designed for and will be deployed into production in the National Market System and how it will be designed and deployed. Repeated references to the white paper in the Filing and rule text just raise more questions than they answer. The Filing and rule texts should describe all aspects of what Nasdaq will put into production to make its "proprietary assessment of market conditions" so participants know how the order type will work and what Nasdaq is accountable for.

The white paper begins by saying "This paper proposes the introduction of a dynamic model-based timer that can be applied to the Nasdaq Midpoint Extended Life Order (M-ELO)," which cues the reader to expect a proposal. In its opening paragraphs the white paper talks about the authors' research efforts and that the rest of the paper will "detail the development and results of our proposed timer-update model and advance its adoption as a step towards the future of dynamic market order types." Nothing in the text or tone suggests that the methods and systems described in the white paper will be put into production in the National Market System as written.

The word "production" appears just once in the white paper and only as a section heading at 4.2.2., "Model Maintenance in Production." That section then describes daily, weekly, monthly, and quarterly retraining and describes a "specific retrain start date" for its training datasets. It also describes "two sequential [retraining] stages." In the first stage "all 35,000+ parameters of the model are trained until the specific retrain start date, in this case January 20, 2022." And

then in the second stage, "some of the initial layers of the model are frozen and the end ones are allowed to update based on the most recent data, generating many additional models that are tied to being tested for specific dates." The white paper then recommends monthly or quarterly "full retraining," a function clearly distinguished from daily or weekly retraining and which the authors say will "ensure [Dynamic M-ELO] keeps up with longer-scale larger distribution shift" in the markets. And as I pointed out in "Data" the white paper says the system was trained on a filtered subset of M-ELO data, 380 names of the 6,257 names it says are "actively traded with M-ELO" and 67% of M-ELO volume, and a reader has no idea if that's how it will work in production.

Eyes Wide Open

Since the 1970s - and even earlier - Congress, the SEC, and a competitive and innovative industry embraced new technologies to improve transparency, determinacy, and fairness in U.S. stock markets. Not all new technologies and implementations helped. Not all market participants were happy about it. The policy direction though was always clear - transparency, determinacy, and fairness. The gains have been revolutionary and touched every corner of the market and every investor and listed company. The Filing is a retreat and should be disapproved.

In previous comments I listed precedents I believe approval will establish. I'll revisit those in light of the Filing, with some added commentary:

(a) That rulebook text as vague and thin as "proprietary assessment" is enough to describe an exchange's novel, complex, and eventful operations.

As discussed above in "Data" and "Rule Text," I believe the proposed rule text fails.

(b) That a mutable and highly technical "white paper" is enough to inform the public about how an exchange performs those operations.

See "White Paper" in my May 2, 2023 comment. As discussed above in "Data" and "Rule Text," the white paper doesn't describe a production implementation of the facility. Nasdaq's Filing text and rulebook reliance on the white paper raise crucial unanswered questions about how the facility will work in production.

(c) That an exchange can use undisclosed aspects of its order flow to freely and unpredictably change an order's material terms.

See discussion above in "Exchanges, Discretion, and Certainty" and my May 2, 2023 and May 30, 2023 comments.

(d) That an exchange doesn't have to hardcode rules and procedures for its order types. Rules and procedures don't have to be hardcoded at all.

See discussion above in "Exchanges, Discretion, and Certainty" and my May 2, 2023 and May 30, 2023 comments.

(e) That an exchange can use a participant's current and past behavior to freely and unpredictably change an order's material terms, even if it means worse results for that

participant relative to a collective, relative to its counterparties, or relative to what would have happened had the exchange remained neutral.

The Filing describes methods designed to improve overall average results, not trade-by-trade results. Dynamic M-ELO isn't a Lake Wobegon order type, where every fill is above average. One objective guiding the system is zero-sum - if a stock moves, one party's improved mark out comes at the cost of its counterparty's worsened mark out. Since a participant's behavior affects the statistics Dynamic M-ELO uses in its methods, that participant's behavior can determine the way in which Dynamic M-ELO alters that participant's order. How Dynamic M-ELO will alter the participant's order is unpredictable, and the participant's results could be below overall averages, worse than its counterparty's results, or worse than if the exchange hadn't used its discretion to determine the order's terms in the first place. The participant's behavior Dynamic M-ELO used to determine the order's terms could have been from the day before, the week before, or even a month or more before. It's unknowable. Dynamic M-ELO's statistics capture activity up to five days in the past, but if a retraining cycle doesn't improve results, according to the Filing "the Exchange will not deploy a retrained model if it fails to achieve performance improvements." Presumably Dynamic M-ELO could operate for weeks or even months with same model, meaning a model derived from behavior long in the past. I don't see that any of this has been refuted or disclosed. See also my May 2, 2023 and May 30, 2023 comments.

(f) That an exchange can use a participant's current and past behavior to freely and unpredictably change an order's material terms without confirming the exchange doesn't unfairly discriminate against participant types, and without monitoring for it in production.

As discussed above in "Sparkling," I believe the Filing raises more questions than it answers about bias and discrimination.

(g) That an exchange can offer de facto broker services without corresponding disclosure and best execution obligations, and without due consideration of its effect on competition.

About Dynamic M-ELO's unpredictable and indeterminate methods, I believe the Filing should go much further than it does in its disclosures. At note 22 the Filing says "During periods where the model is not undergoing retraining, the System will behave predictably from day to day, such that its decisions when presented with given set of facts and circumstances in a given security on day 1 should be the same as they would be on day 2." This is, at best, half the story. The industry would likely never imagine the half that's missing. The half that's missing is that across retraining cycles the system is unpredictable for a given set of facts and circumstances, and that reality should be disclosed in the Filing and any materials Nasdaq publishes about Dynamic M-ELO, just as a broker would almost certainly be required to disclose it. About competition, as I discussed in my May 30, 2023 comment, I believe these methods venture into discretionary broker services by an exchange, and because of an exchange's unique privileges any exchange has an unfair competitive advantage offering these kinds of services.

(h) That an exchange can use current and past counterparty behavior to change an order's material terms, affecting the "manner of execution" of a participant's order.

See "Reg ATS and IEX" in my May 2, 2023 comment.

(i) That an exchange can in fact exercise indeterminate control, judgment, and discretion over its participants' orders and retain the status and enjoy the privileges of an exchange.

Certain artificial intelligence methods are contrary to what an exchange is for and what Congress, regulators, and the industry have worked toward for the last half century. Even if the SEC believes these methods have promise, this proposal bulldozes the public interest by unraveling 50 years of steady and unambiguous policy choices toward market transparency and determinacy. Approval will set new precedents and launch a radical reinterpretation of an exchange and its role in the National Market System. At minimum, the SEC should broadly engage the industry and the public about these methods before allowing them in exchange trading systems. In the meantime, the SEC should say no to the Filing.

Sincerely,

R. T. Leuchtkafer