Best Practices For Automated Trading Risk Controls And System Safeguards

July 2024

TABLE OF CONTENTS

Introduction

The technology of trading has experienced several bursts of innovation in the past 40 years. From open outcry trading on the floors of exchanges, to the "point and click" screen-based trading of the 1990s, to the development of automated trading systems in the 2000s. In parallel with this trend, participants in the exchange-traded markets have continuously enhanced their processes and tools for mitigating the risks of automated trading.

Risk controls related to automated trading, developed through more than a decade of collaboration among market participants, exchanges and regulators, are key components of stability for the industry and, when they were used, have proven effective through various volatility events.

As technologies have evolved, so too have risk management practices. Accordingly, earlier this year FIA undertook a review of these risk controls, as we have done periodically for the past 15 years. This paper outlines current practices across a broad range of subjects including pre-trade risk management, exchange volatility control mechanisms, post-trade analysis and testing, among others. These practices are designed to ensure trading systems operate safely, regardless of the technology used. Thus, the practices outlined in this paper apply to both existing technologies and evolving ones, such as artificial intelligence (AI).

FIA has a long history of working with market participants on industry standards and best practices. We strongly support innovation, and we believe innovation in trading has the best chance of success when it occurs within a framework that promotes responsible use and appropriate risk management. This paper consolidates and updates our previous work, including our most recent paper on Exchange Volatility Control Mechanisms, focusing on the tools and protocols used to mitigate the risks of today's technologically advanced futures markets.

1 Pre-Trade Controls

Localized pre-trade risk controls, not credit controls, should be the primary tools used to prevent inadvertent market activity due to unauthorized access, system failures and errors. These controls can be implemented at various points in the execution order flow—at the trader, at the broker or at the exchange itself. Such localized controls use various approaches and act on a very granular level. Similar pre-trade controls may often exist at multiple points within the order flow and are used to mitigate risk from different viewpoints. For example, a maximum size order would be employed by:

- A trader, to prevent it from submitting an order to the market that is larger than its risk tolerance.
- A broker, to block a customer from submitting an order larger than the limits previously determined as part of a risk review exercise.
- An exchange or other trading venue at a product level, as a secondary control to avoid inappropriately large orders that could affect price discovery within a market.

Without prescribing specific risk control implementations, a degree of standardization may be achieved across market participants, regardless of their trading strategy, by implementing risk controls at the exchange. By doing so, participants are required to pass each of their orders through the exchange's risk controls prior to submitting for execution. This provides a baseline of risk controls within the marketplace regardless of the type of access used or the type of market participant. Transparency on exchange-level rules and policies is an important consideration. Public disclosure of rules or policies on exchange websites or through other public forums is strongly encouraged. If policies are changed, prompt dissemination to market participants should occur.

This section will list different types of pre-trade controls and make recommendations for deploying them within the order flow and by whom. Regardless of where, or who sets the controls, configuring them correctly and at the appropriate level is essential. The implementation and configuration of the controls set forth below should be reviewed on a regular basis and updated as necessary. When possible, in addition to manual methodologies, the provision of an Application Programming Interface (API) for limit setting and maintenance is beneficial.

1.1 Maximum Order Size

Maximum order size sets the maximum quantity allowed to be submitted per

order, commonly referred to as “fat-finger” limits. Errors may be prevented by rejecting the order in the case of a limit breach.

This risk control should be applied when a new order is submitted or an existing order is modified. Requiring each order to pass pre-trade order size checks can facilitate the entry of all orders into the market within parameters that protect the natural price discovery process from aberrant and accidental behavior, such as generating unintentionally large orders.

An appropriate maximum order size should be applied across different types of instruments appropriate to typical activity. For example, different maximum order sizes may be applied to orders on futures, options or spreads on the same underlying instrument. These limits may be adjusted for each instrument and each trading venue.

Systems should prevent orders from being placed in cases where no order size limits have been set for an instrument.

Such limits may be set at the trading application level. Depending on the type of market access, the broker providing access also should set limits within their own trading infrastructure for indirect access participants or using an exchange-provided tool for direct access participants, based on a review of the risk limits appropriate for the market participant.

Additionally, exchanges should apply a similar control to all participants at a product level to prevent accidental disruption of the market caused by orders that exceed the liquidity of the marketplace. However, care should be taken that order size levels set at the exchange level are not too restrictive. Exchange mandated limits should be publicly available, enabling market participants to take them into consideration when designing and configuring their systems.

1.2 Maximum Intraday Position

A maximum intraday position is the maximum long or short position that can be taken within a given system intraday. Errors may be prevented by rejecting the order in the case of a limit breach. Warnings may be employed when the limit is close to being breached.

When a new order is submitted or an existing order is modified, both current positions and working orders should be evaluated to determine whether a breach of the limit could occur. It is important to include working orders such that limits would not be breached if that order is filled, even though it may not be immediately executable.

Such position limits may be considered simple pre-trade risk limits as opposed to credit limits since an accurate picture of start-of-day positions is difficult to derive in a timely fashion across multiple execution channels. It is important to note that not all systems can use this type of limit. This may work for automated trading operations that leverage more than one trading system, over-the-counter trading or floor-based trading. Where implemented, it should be considered a "speed bump" to prevent accidental overtrading and, as such, should be employed with appropriate post-trade risk controls (see Section 4).

Maximum intraday position limits are generally set by instrument, by individual trader, group of traders and/or for the whole firm. They may be set within the trading application or a separate risk management system that oversees activity across multiple trading strategies.

Depending on the type of market access, the broker providing access may also set limits within their own trading infrastructure for indirect access participants or use an exchange-provided tool for direct access participants.

A process should be established for setting the initial limits and for making necessary amendments over time. Authorized staff independent of trading activities should manage the process whenever possible to avoid conflicts.

1.3 Price Tolerance

A price tolerance limit is the maximum amount an individual order's limit price may deviate from a reference price, such as the instrument's current market price, and is typically applied on orders generated from an automated trading system before the order is sent to the exchange. Errors may be prevented by rejecting orders with limit prices placed outside the acceptable range. Price tolerance checks should be applied when a new order is submitted or when an existing order is modified. Requiring each order or amendment to pass price tolerance checks makes it more likely that all orders entered into the market fall within parameters that protect the natural price discovery process from aberrant and accidental behavior, such as generating orders unintentionally far away from the current market price.

Price tolerance limits should be set at the trading application level. Depending on the type of market access, the broker entity providing access may also set limits within their own trading infrastructure for indirect access depending on the requirements of the participant and/or the exchange.

1.4 Cancel-On-Disconnect

Cancel-On-Disconnect (COD) is a service provided by exchanges that monitors for

a loss of connectivity between a participant's trading session and the exchange's trading platform. If a loss of connection is detected, COD initiates a best-effort attempt to cancel all resting orders for the disconnected session. COD provides participants with the safeguard of knowing that all working orders are cancelled at the exchange if the market participant loses its connection to the exchange.

COD functionality at the exchange should be optional, allowing market participants to decide whether COD mitigates risk by cancelling orders in the event of a disconnection or adds to risk in such a situation.

It should be at the discretion of the exchange, i.e., the entity responsible for triggering COD functionality, to define what disconnection means. For example, it might involve detection of a network-level error or even loss of application-level heartbeats. What matters is that the exchange triggers COD when it has determined that a trading session has suffered an unexpected disconnect.

In terms of which orders ought to be cancelled upon disconnect, it should be considered that many market participants maintain multiple trading sessions (i.e., connections) to an exchange, and order cancellation should be done at the granularity of an individual session so that all orders originating from the disconnecting session should be cancelled and those originating from other sessions should remain working on the exchange.

It is important to note that it is increasingly common for brokers to also employ COD for their connections to the exchange. This allows the broker to manage their risk across customers in the event of a loss of connection. As with a market participant, the broker also should decide whether using COD in the event of an issue mitigates risk or increases risk, and they should advise customers accordingly.

Brokers also provide the ability for customers to route orders to an exchange— or multiple exchanges— through the broker's infrastructure. The broker should advise its customers whether it is possible to pass through cancel requests to the exchange in the event of an unexpected disconnection by the customer from the broker's infrastructure. At present this is typically unsupported, and the customer would need to contact the broker to manually cancel any working orders.

1.5 Kill Switches

A kill switch is a control that, when activated, immediately disables all trading activity for a particular participant or group of participants, typically preventing the ability to enter new orders and cancelling all working orders. It also may allow for risk-reducing orders while preventing risk-increasing orders. This can be considered an effective safeguard against situations such as an automated

trader breaching limits defined by a broker, or erroneous trading activity that may be caused by an automated trading system malfunction or the generation of unintended orders released into the market.

Activation of a kill switch is based on a decision that such action protects market integrity or the financial integrity of the counterparties involved. Such a control may provide exchanges, brokers and market participants with an immediate and effective way to remove or reduce risk. The conditions under which a kill switch may be used by an exchange or a broker should be clearly communicated to their counterparties.

However, kill switches offer just one of many different types of risk controls that comprise an effective suite of risk controls and only invoked based on a qualitative decision taken as a last resort when other actions have failed or may not be feasible. In an environment where adequate pre-trade risk controls are implemented at all appropriate levels, i.e., the market participant, broker and exchange, a kill switch may ultimately be considered redundant.

Market participants are encouraged to build their own kill switch functionality into their trading applications, and where possible to implement it on a sufficiently granular level to identify individual trading systems. Such functionality may be separate from the trading application itself and can be operated both by the trader and by the person responsible for risk. When made available, this functionality should serve in addition to and as a final backstop for the pre-trade risk functionality outlined above.

A broker may want to implement kill switch functionality for both its direct and indirect automated trader customers, although typically the revocation of customer trading access takes place through the broker's pre-trade risk controls or using exchange-provided tools for direct access customers. Such a control should be granular enough to identify individual customers and/or trading systems as appropriate.

Where a broker must rely on an exchange-provided risk management control, for example for a kill switch for a direct access participant, the exchange control should operate at a suitable level to control only that customer's order flow and should not be shared across customers. It is important to note that exchange risk management tools vary in implementation based on how the exchange identifies trading sessions or operator IDs. Where an automated trader also has access to the exchange-provided control, the automated trader should not be able to override a kill switch invoked by the broker, effectively ensuring that the trader cannot bypass the broker's decision.

Where an exchange provides such a control, the exchange should have a registration process and entitlement system that requires market participants and brokers to specify which staff have authorization to use the functionality. The system itself should provide explicit warnings informing the authorized users of the consequences of activating the kill switch.

1.6 Exchange-Provided Order Management

The ability to manage orders independently from the automated trading system is an important risk mitigation device. Errors may be prevented by cancelling any working orders if there is a system failure.

Exchanges should provide an independent mechanism for viewing and cancelling working orders for a given session or user. Such functionality should be independent from the trading access that might be subject to disconnection or disruption and may be used in conjunction with COD functionality (see Section 1.4), or in cases where COD is not provided. Such exchange-provided order cancellation and COD ought not be viewed as alternative approaches; they are often complementary. Alternative order cancellation channels also allow a firm to proactively pull orders on behalf of trading sessions that they have themselves deemed in error.

It is important to note that major network failures at the client may impact alternative order management channels. At that point, the only mechanical means by which orders can be removed is the exchange's COD capability.

Brokers providing exchange access to market participants also may have access to the same alternative order management tools. In the event of a major system failure, authorized personnel at the broker may use this tool to confirm that orders have been cancelled and/or initiate the cancellation on behalf of the market participant.

2 Exchange Provided Volatility Control Mechanisms (VCMs)

The objective of volatility controls is not to prevent volatility altogether; volatility typically does not weaken the integrity of markets, nor does it lessen investor protection. Volatility is a natural byproduct of price discovery and efficient markets, as new information is incorporated into prices. Volatility controls "should be designed to avoid market disruptions without unduly interfering with that market's price discovery function"[1] and "to preserve the efficient and orderly functioning of financial markets."[2]

Volatility controls should be designed to mitigate the impact of extreme short-term volatility driven by events, such as operational errors or temporary supply and demand imbalances, but without impeding the ability of an instrument to reach a new equilibrium price level based on healthy market dynamics. In a volatility event, exchanges should focus on keeping markets trading. Trading interruptions should be used sparingly.

Exchanges typically design and deploy several different, complementary controls working in tandem to ensure markets operate in an efficient and orderly manner across all trading hours. Because there are different triggering events that may impact volatility on electronic trading platforms (such as fat-finger order entry quantity and pricing errors, message frequency aberrations and micro- and macro-economic events), there is no panacea. No single control can accomplish the goal of preserving orderly markets, and no one control is in all cases more effective than another.

In addition to designing VCMs to mitigate the impact of different types of triggering events, exchanges should regularly test and review these controls. This involves regularly assessing the efficacy of the controls, including a review of thresholds or parameters, the frequency of a control being triggered and whether the controls meet their purpose, as well as incorporating feedback from market participants or other stakeholders.

During stressed markets, exchanges may potentially amend the price band parameters which trigger a VCM, and they should regularly review and validate the criteria used to determine whether markets are stressed. Stressed market conditions are characterized by significant short-term changes in price and volume and can be caused by a variety of occurrences including economic releases. The conditions and criteria that must be met to determine whether a market is

1 17 C.F.R. Appendix B to Part 38
2 MiFID

"stressed" should be transparent and readily available to market participants, enabling them to independently calculate and anticipate the triggering of these stressed market condition price and volatility controls.

Information pertaining to VCMs should be made available to market participants in a transparent manner. Exchanges should publish information on how the controls operate and the thresholds that would result in the controls being triggered. These criteria should be publicly disclosed and replicable. Access to these controls and their settings should be made available to brokers and traders via an API, enabling efficient, integrated and more automated monitoring and management of these controls within these firms' existing tools and processes. This should be in addition to GUI-based interfaces suitable for manual operations. Many exchanges have webpages dedicated to volatility controls that include detailed descriptions of the controls, technical specifications, explanatory videos and frequently asked questions. Exchanges should also notify participants when a VCM is triggered.

2.1 Exchange Dynamic Price Collar

The purpose of price bands or price collars on orders is to prevent erroneous orders, such as limit bids at prices well above the market or limit offers at prices well below the market, from entering the market and resulting in trades at aberrant prices. With price bands, each order submitted to an exchange's trading platform is subject to a price validation control. Orders with prices outside of a pre-defined band, collar or validation range will be held or rejected by the exchange.

Price bands can vary from product to product. They should be dynamic and regularly recalculated to allow price discovery to continue whenever possible. Price bands are generally based on the reference price plus or minus a fixed or variable increment, whichever is deemed to be suitable for the market and/or asset class. The size of the fixed increment, or the relative percentage of price range, should be reviewed regularly and calibrated based on historical intraday price statistics as well as current market conditions.

Price bands typically allow a wide range of prices to be traded above or below a particular product's reference price. In certain situations, price bands may become stale (too narrow) and/or no longer reflective of the current market; this is more likely to happen in illiquid products such as options or in stressed market conditions. In this situation, some exchanges may modify the price bands to permit a wider band of prices to be traded using theoretical prices or another interpolated mechanism if no last trade price is available. While the width of price bands need not necessarily be symmetrical, if one side of a price band is adjusted then the other side of the band should be adjusted by an equivalent amount (i.e., a "sliding

band"). Exchanges should conduct regular price band reviews to ensure they are recalculating bands with sufficient frequency.

When price bands are in use, they should be active across all trading hours. However, the size of the band may vary during different trading sessions.

2.2 Daily Price Limits

Daily price limits, where implemented, represent the maximum price range permitted for each contract during a prescribed time interval (i.e., daily or trading session). Both the daily price limits and the actions taken when limits are reached can vary from product to product. Some markets may have a hard daily price limit (i.e., prices cannot trade above or below the limit) and other markets may temporarily interrupt trading (short or long duration market pause) until the price limit can be expanded to a higher/lower interval. Not all contracts have daily price limits.

Generally, daily price limits are calibrated or determined by the exchange using historical price statistics (such as averages or quantiles) and the specifics of the contract. Daily price limits apply to up and down price movements. They may not necessarily be symmetric such that a limit down value can be different from a limit up value. Daily price limits should apply across all trading hours, including overnight trading, although the size of the limits applicable to different trading sessions may vary.

To allow for continued price discovery in volatile markets, contracts with hard daily price limits may also have provisions for an "expanded daily limit." An expanded limit goes into effect when one or more contract months settle at an upper or lower price limit. The expanded daily limit remains in effect until one or more of the contract months settle below a high limit or above a low limit, at which time the contract reverts to the original hard daily price limit.

2.3 Mechanisms to Interrupt Continuous Trading

Short duration market pauses are one of the VCMs used to prevent market prices from moving "too far, too fast" by placing limits on the amount a market can move within a preset time period. This control usually analyses price moves within a very short timeframe, typically no more than a few seconds. If the price move within this timeframe exceeds predefined parameters, such as the fixed price band value, the market is paused or held, again, typically for no more than a few seconds.

Long duration market pauses (commonly referred to as circuit breakers or dynamic circuit breakers) are a VCM that provides another layer of protection from extreme volatility events. Except for the timing or duration component, circuit

breakers function similarly to short duration market pauses. While circuit breaker functionality varies, they typically assess market moves over a longer period (e.g., price moves over a rolling one-hour period or over the course of a trading session) and involve longer market pauses than short duration market pauses (e.g., a few minutes versus a few seconds). The price moves required to trigger a circuit breaker event are also usually wider than price moves for short duration market pauses (e.g., 5%+ of previous settlement price).

Short and long duration market pause triggering parameters are configurable by the exchange and can be tailored for each contract or group of contracts. These parameters should be transparent and readily accessible to market participants.

When a short or long duration market pause is triggered, in addition to pausing the market, there are a variety of potential additional parameters that an exchange can set to determine more specifically how the pause functions. Some examples include:

- Apply a pause where no orders are allowed.
- Apply trading constraints over a defined period where orders and matching beyond a specific price level are restricted.
- Apply a pre-open state where orders are submitted or cancelled but no matching takes place.

Interrupting continuous trading by transitioning to an auction state is another VCM an exchange can use if prices are outside a pre-defined range. During the call phase of the auction, orders can be submitted to the order book without order matching and an indicative price can be displayed. This decelerates trading, allowing market participants sufficient time to react to new information and for the market to attract more liquidity.

Exchanges have different procedures for resuming continuous trading after the activation of an interruption mechanism. In instances where a market pause is implemented, markets transition to a pre-open state where the exchange rebuilds its order book by collecting orders from market participants in advance of resuming order matching at a specified time. In markets that have moved to an auction state, the call phase can establish a re-opening price based on a trade where buyers on the bid and sellers on the ask are matched.

Regardless of the specific operating model, exchanges should give due consideration to the design of the trade resumption mechanism to foster robust price formation, enhance market transparency and strengthen confidence. Although sometimes necessary, interruptions to continuous trading can interfere

with the price discovery function. As such, an appropriate degree of care and deliberation should be applied to the design of the procedures for resumption of trading. Exchanges should incorporate features that facilitate greater participation and incentivize liquidity to boost confidence in price formation during periods of excessive volatility, such as providing indicative prices or order balance information which could attract the liquidity needed to reestablish a healthy market.

3 Other Tools and Controls

In addition to the pre-trade risk controls and exchange-provided volatility control mechanisms described in the previous sections, additional tools and functionalities exist to further enhance market stability and mitigate unintended consequences. This section delves into a range of mechanisms employed by market participants, brokers and exchanges to safeguard market integrity and prevent inadvertent errors.

3.1 Market Data Responsibility

A market data reasonability check is a tool designed to control whether the data used to generate orders by an automated trading system is within acceptable boundaries. Errors may be prevented by having the existence of aberrant market data escalated to the supervisor or support team of the trading system. Orders generated because of this data should be cancelled or rejected prior to submission to the market.

Trading and risk management systems should have such checks on incoming market data as well as on values generated using the market data. For example, automated trading systems should have controls that validate whether actionable data is reasonable based on a variety of factors that may include the time since the last update was received, previous price, bid/offer spread or deviation from an average price. If there appears to be a deviation from what is expected, then an alert should be provided flagging market data may be stale, and any orders should be blocked while the deviation is investigated.

Both exchanges and commercial data providers should make efforts to disseminate accurate data. This is especially important at times of high volatility when the possibility of incorrectly disseminated data (including a data outage) could affect the ability of participants to manage their risk.

3.2 Repeated Automated Execution Limits

A repeated automated execution limit is the maximum number of times a

strategy or identical order is filled and then re-enters the market without human intervention. After a configurable number of repeated executions, the strategy should be disabled until an authorized person re-enables it.

Due to the dependency on the type of strategy and on market conditions, these controls should be set by the automated trader and not by the broker or the exchange. The appropriate limit will vary depending on the strategy in use and should be configured accordingly.

While it is the responsibility of the automated trader to detect incorrectly generated repeated executions, exchanges and brokers may also detect suspected repeated executions through their regular monitoring of market activity. They should attempt to contact the automated trading system operator regarding any action deemed appropriate to protect market integrity.

3.3 Exchange Message Programs

Exchanges are in the best position to monitor a market participant's messaging practices to help safeguard the integrity of the market and the exchange platform.

The exchange should be responsible for setting messaging measures for each instrument based on many factors, including the capacity and performance of its network and matching engine, the matching algorithm and the unique characteristics of the financial instrument, particularly as it relates to liquidity. Messages can include orders, cancellations, modifications and notifications of execution. Messaging measures should not be dynamic because market participants need to know what is expected of them. Details of exchange messaging programs should be transparent and publicly available.

It is reasonable for messaging programs for designated liquidity providers to be different from those for other market participants. Designated liquidity providers are often required to quote two-sided markets in many products simultaneously, and an overly restrictive limit will inhibit their ability to perform their duties and properly manage the risk associated with those duties.

One example of a messaging program looks at a participant's order-to-trade ratio, which compares the number of orders submitted to the executed quantity. For each instrument or instrument group, the acceptable threshold ratios are set by the exchange and publicly documented. It is important to note that this type of analysis is not done in real-time but after a trading session is complete. Also, it may not be possible to set meaningful order-to-trade ratios for newer products or products that trade infrequently.

3.4 Message Throttles

Message throttles are controls designed to prevent excessive messaging which could disrupt, slow down or impede normal market activity.

There are a wide variety of approaches to message throttles that can be applied at various points in the order flow. Exchanges can establish controls at their gateways that monitor message rates and send warnings or reject messages from a participant when certain rates of messages per second are sustained. Such controls reduce inadvertent market activity by preventing high message rates that can stress the infrastructure at the automated trader, broker or exchange. High message rates generated by one automated trader can introduce slowness for other market participants (i.e., more messages than a market data feed can consume or disseminate) and can increase message processing risk as participants cannot process the latest messages because they are still processing earlier ones. To avoid introducing undue risk into the marketplace, a message throttle should never reject an order cancellation request due to breached message rate limits. Exchanges that implement message throttles should publish their limits.

Exchange-based message throttles may be supplemented by message rate limits at the market participant or broker level. If an automated trader chooses to implement their own message rate limits, the limits must be flexible to address the market participant's unique and diverse risk management requirements. Brokers may choose to implement such controls to minimize disruption to execution services due to abnormal activity from a customer. They would also prevent any knock-on effect to post-trade services caused by capacity constraints from processing abnormal rates of pre-trade activity for indirect market participants. Brokers should be transparent to their customers regarding the reason for the additional control and the maximum message rate that can be supported by the broker.

Finally, care should be exercised regarding any attempt to mandate the use of message throttles, as it would be difficult to devise universal controls that are appropriate for all trading strategies and all financial instruments. Implementation of such controls should not be limited to specific types of market participants as this is likely to distort fair and even access to the market and may ultimately impact market integrity. For example, controls and their associated costs applied exclusively to designated liquidity providers or market makers could potentially discourage them from performing the critically important role of providing liquidity.

3.5 Self-Match Prevention

Self-match prevention is a functionality designed to prevent a market participant

from inadvertently trading with itself. To consider such functionality, it is necessary to distinguish between three types of self-match trades that could occur on an exchange:

- **Wash Trades**: Intentional self-matches created with the aim of distorting or manipulating the market, generally prohibited by rules and regulations.
- **Bona Fide and Allowable Self-Match Trades:** Buy and sell orders for accounts with common beneficial ownership that are independently initiated for legitimate and separate business purposes by independent decision makers and which coincidentally cross with each other in the competitive market.
- **Inadvertent Self-Matches that Occur on More than an Incidental Basis:** Orders submitted by the same person, automated trading system, trading team or related groups of traders are matched despite best efforts to avoid self-matching.

Market participants should have policies and procedures that prohibit wash trades and other forms of undesirable self-match trades. A variety of tools may be used to prevent inadvertent self-matches.

It is important to note that due to the diversity of trading operations and strategies, there is not a one-size-fits-all solution to self-match prevention. For example, a market participant that predominately acts as a liquidity provider may not want its resting quotes to prevent new hedge orders from being accepted for execution by the exchange. Similarly, a market participant that rests large limit orders for extended periods of time may not want those orders to be cancelled because of submitting a new, aggressing order (or market order) to the exchange.

Exchanges should offer participants a selection of self-match tools to allow market participants to tailor self-match prevention to their individual needs by offering various options (e.g. cancel resting, cancel new, cancel both and decrement order size) and various levels of granularity (e.g. firm level, group level, trader ID level, customer account level and strategy level). However, providing flexibility can increase the complexity of implementation, and an appropriate balance between flexibility and complexity should be found. It should be noted that certain levels may be combined or offered in conjunction with another level.

An important benefit of the exchange providing self-match prevention is consistency across market participants in terms of available functionality and cost impact. However, given the different requirements from market participants and different implementations at the exchange level, such controls should remain optional, and a decision by a market participant not to implement available functionality because it does not suit their business structure should not be construed as intent to bypass responsibilities regarding self-matches.

4 Post-Trade Analysis

A combination of post-trade controls, monitoring and data collection should be used in conjunction with pre-trade controls to watch for potential credit events or unintended trading. These post-trade functions will vary depending on the size and complexity of the market participant and the variety of asset classes being traded.

4.1 Drop Copy Reconciliation

Drop copy is a report that details a participant's execution activity on a trading venue and is generated as close to real-time as possible. Drop copy feeds differ from cleared trade feeds in that they (a) may contain additional information to aid a participant's risk management, such as order state changes, modifications, rejections and cancellations, and (b) are generated at the point of execution, rather than when the trade has been cleared. Currently, the contents and method of delivery for drop copy feeds vary by trading venue. All participants may use drop copies for real-time trade reconciliation, including both market participants and brokers. This reconciliation process typically compares the information provided by a drop copy in real-time with the trade notifications received from production trading sessions. This comparison process allows firms to reconcile their electronic trading activity with an independent source of exchange-provided trade notifications. In the event of a discrepancy, the responsible party may take action immediately to address trading risk, determine the cause of the discrepancy and resolve any issues.

Market participants may also supplement their risk management process by using drop-copy functionality to consolidate multiple trading session reports into a single data feed. This consolidated data feed may then be used by operational staff to more efficiently monitor a participant's overall trading activity.

Drop copies should be available for all trading venues and products whenever technologically practicable. Exchanges should seek consistency in the format of drop-copy reports to assist in consolidation across exchanges. Trade reports and other information provided by drop copy should be disseminated to the market participant in real-time or as near real-time as technologically and operationally practicable. Updates provided by drop copy, or any other order and trade report, should include any necessary information required to identify the order described in the update and interpret the changes to that order. Additional details may be provided to increase the utility of the order and trade report.

Both exchange drop copy and feeds from the broker clearing the trades offer good sources of information for reconciling a market participant's own system.

A frequent reconciliation process where the firm balances its trading systems to drop copy or clearing information can serve as an early warning for potential problems and can help mitigate risks due to errors or malfunctions.

4.2 Post-Trade Credit Controls

Brokers should establish post-trade credit limits that are appropriate for the market participant's capital base, clearing arrangements, trading style, experience and risk tolerance. Credit limits should be determined by a broker's assessment of their customer's assets and history, and they should be monitored across the customer's entire portfolio. Monitoring of customer credit limits includes their utilization of margin on positions carried by the broker, executed through the broker, those "given in" from other executing brokers, as well as the collateral posted in customers' accounts at the broker clearing the trades.

It is important to distinguish between pre-trade risk controls and credit controls. Pre-trade risk controls are designed to prevent trading activity from creating market disruptions, i.e., what is acceptable in terms of order size, number of orders and other checks outlined in Section 1. Credit controls are designed to prevent a credit event and are calculated on a post-trade basis.

Credit controls are a key feature of how a broker manages its exposure to its customers through the different types of market activity in which they participate. As such, these credit controls need to be employed on a post-trade basis due to the diversity of information required to accurately calculate exposure where market participants have the ability to use multiple systems and/or multiple brokers to access the market. In such circumstances both market participants and the broker clearing the trades may use drop copies and clearing system trade feeds on a near real-time basis and should also maintain this data for historical review.

Both market participants and brokers may set daily position and/or loss limits by account as a form of credit control. These limits should be monitored and alerts generated at appropriate thresholds so a discussion can occur between the broker and the market participant to decide what action should be taken to mitigate the relevant risk before the limit is breached. Such post-trade controls have both a quantitative and qualitative nature, and judgment should always be exercised before invoking certain controls, such as a kill switch (see Section 1.5).

4.3 Exchange Error Trade Policies

In addition to pre-trade controls, exchanges typically also provide other forms of recourse to market participants where trades occur at erroneous price levels, including exchange error trade policies. Exchange rulebooks typically provide

the exchange with the authority to review and adjust or cancel trades that are executed at clearly erroneous prices within a defined time limit.

Exchange rules will describe the exchange's authority to adjust or cancel trades where, in its absolute and sole discretion, it believes a price adjustment or cancellation is necessary to mitigate market disrupting events caused by a participant's error. These error trade policies should be designed to balance market participants' need for trade certainty with the adverse effects of trades being executed at prices inconsistent with prevailing market conditions.

The goal of any error trade policy should be to promote a marketplace where all trades stand as executed. Where this is not possible, a price adjustment is the preferred action. Adjusting the price of an erroneous trade is typically less disruptive than a cancellation, as an adjustment only impacts the value of the erroneous trade and not the overall position or any ensuing trading decisions made by the impacted market participants. However, it is recognized that there may be exceptional circumstances when erroneous trades must be cancelled to retain market integrity.

5 Testing

Testing pertains to the work done by organizations to confirm that their trading systems and environments function as designed and within acceptable parameters. Organizations should have a process for testing trading software and infrastructure before any new system, new versions or updates are released to the production environment.[3] While market participants may utilize a wide range of testing techniques, conformance testing is one type that should be consistently implemented.

5.1 Exchange-Based Conformance Testing

Exchange-based conformance testing is a type of testing that typically follows a script of tests designed and administered by an exchange to confirm that market participants' systems interact with an exchange's systems properly. By administering and performing such tests, exchanges can confirm that each market participant system exhibits a baseline level of functionality that has been deemed necessary for maintaining orderly markets.

- Exchanges should provide market participants with an environment that sufficiently mimics the production trading environment for conformance testing.
- Exchanges should require that conformance testing is performed whenever a market participant wishes to deploy a new exchange-facing software interface to the production environment. Follow-up conformance tests should be performed when material changes have been made to previously approved exchange-facing software interfaces.
- Market participants are responsible for initiating conformance testing whenever necessary.
- Exchanges should provide an appropriate series of tests for market participants to perform in conjunction with conformance testing. If a market participant's system does not perform certain functions described within a conformance testing script, the exchange may grant a waiver for the associated tests.
- Exchanges should provide documentation to market participants to confirm the successful completion of conformance testing.

3 Please refer to the FIA Guide to the Development and Operation of Automated Trading Systems for a comprehensive set of testing methodologies.

About FIA

FIA is the leading global trade organization for the futures, options and centrally cleared derivatives markets, with offices in Brussels, London, Singapore and Washington, D.C.

FIA's mission is to:

- *support open, transparent and competitive markets,*
- *protect and enhance the integrity of the financial system, and*
- *promote high standards of professional conduct.*

As the leading global trade association for the futures, options and centrally cleared derivatives markets, FIA represents all sectors of the industry, including clearing firms, exchanges, clearing houses, trading firms and commodities specialists from about 50 countries, as well as technology vendors, law firms, and other industry service providers.

BRUSSELS
Office 502
Square de Meeûs 37
1000 Brussels, Belgium
+32 2.791.7571

SINGAPORE
One Raffles Quay North Tower
Level 49
Singapore 048583
Tel +65 6622.5781

LONDON
Level 28, One Canada Square
Canary Wharf
London E14 5AB
Tel +44 (0)20.7929.0081

WASHINGTON
2001 K Street NW
Suite 725, North Tower
Washington, DC 20006
Tel +1 202.466.5460

FIA.org