OPTIONS MARKET MAKER PROTECTIONS:
A BEST-IN-CLASS APPROACH

Optiver supports open discussion and debate on all market structure topics that would lead to an improvement of the market.

To give feedback on this paper – or to discuss any other market structure topic – reach out to the Optiver Corporate Strategy team at CHI_CorporateStrategy@optiver.us

When it comes to mass quoting, some market maker protections (MMPs) work better than others. Despite the variety of protections out there, we feel that a small core provide sufficient protection and lend themselves to sound administration in practice. They consist of volume- and Greek-based dimensions of protection, with limit settings scoped at the badge- and product-level. Our hope is that the recommendations in this paper help guide exchanges in the design and evolution of their MMP infrastructure.

Mass quoting, an important part of any well-functioning market, requires market makers to post liquidity in a vast range of instruments. The benefit for market participants is that they see liquidity in far more products than any one market maker would ever be willing to trade at a given time. But providing this level of liquidity across a significant number of products, strikes and tenors/expiries exposes the market maker to enormous risks.

To mitigate these risks, exchanges provide market maker protections (MMPs). These help market makers cap their instantaneous risks well below the sum of the aggregate liquidity they are quoting, and without them there would be less displayed liquidity. MMPs are built directly into the exchange's matching engine, as pre-emptive controls won't work without real-time, accurate information.

Exchanges offer various metric-based MMPs, some of which we find more effective than others (see graphic). Some target various "**dimensions**" – or different kinds of exposure or trading scenarios – that might exceed the liquidity provider's risk appetite. Exchanges also define the "**scope**," or the segment of quotes to which these protections apply, differently. Finally, there are subtly different procedures among exchanges for what happens in the event of a "**limit breach**," or when an MMP counter exceeds its threshold, as well as the procedure for checking and managing limits.

Our experience as a global market-maker suggests that volume- and Greek-based dimensions of protection are the most effective and relevant in today's options markets, with limit settings scoped at the badge[1]- and product-level. In this paper,

---

[1] The word "badge" is used here to mean an identifier used by an exchange to identify the market maker submitting quotes.

we present a high level overview of the various protections and procedures, along with our recommendations for operational changes to increase flexibility.

METRICS: DIMENSIONS OF PROTECTION

Surveying a number of options trading venues, the following broad dimensions of protection emerge. These metrics are a proxy for the underlying financial risks that market makers face, offering fine-grained, mechanical control over how their quotes are permitted to be traded against.

- **Execution events** use the number of discrete trading events to which quotes might be subject to:
    - Total quote hits
    - Total first-time quote hits
    - Total complete quote fills
- **Volume** in either absolute or relative terms:
    - Cumulative volume traded
    - Percentage of posted volume traded
- **Notional value** bounds the cumulative notional value tradeable, effectively allowing deep in-the-money options to tally differently from out-of-the-money options
- **Greeks** directly account for the different risks to which market makers are exposed by taking option "moneyness" into account:
    - Cumulative deltas traded
    - Cumulative vegas traded
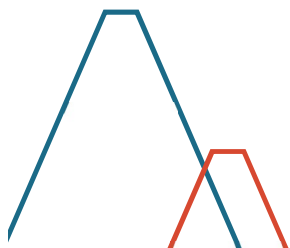
- **MMP breaches** – when a limit is breached – are a second-order measure to bound the aggregate number of protection breaches that a firm is permitted to incur across all of its market-making activities

Market makers administer their protections by specifying threshold values, such as limits, for whatever metrics are defined by the exchange or clearing member within the various dimensions of protection that their MMPs encompass. These threshold values are for an exchange-defined time interval. For example, a market maker might specify a maximum value for cumulative volume that is allowed to trade within a 100-millisecond time window.

It is tempting to cover as many MMP dimensions as possible. This way, a large number of potentially dangerous scenarios can be covered. However, having to manage too many limits can itself constitute a risk. The chance of administrative error, as well as the operational burden to both market makers and exchanges, goes up exponentially as complexity increases. In the end, attempting to cover too many dimensions becomes a technical risk in and of itself, undermining the very goals of these protections.

Based on our experience as a market maker on a variety of exchanges, we find that just two MMP dimensions, volume and Greeks, are sufficient for robust risk management of options mass quoting. Together, these dimensions provide protection for a wide range of scenarios and are easy to calculate and to configure.

SEGMENTATION: SCOPE OF PROTECTION

In order to achieve effective protection, exchanges apply the various MMP metrics over different segments, or scopes, of a market maker's quotes. Scope of protection tends to fall into three buckets:

- **Session**:
  - MMP limit settings are confined to the scope of a single quote-streaming session, meaning an individual TCP connection to the exchange
  - All quotes submitted via a session are subject to the limits associated with that session
- **Badge:**
  - MMP settings are associated with all quoting sessions for a specific market maker badge[2] to stream quotes
  - If a single badge is used by multiple quoting sessions, MMP protections for that badge cover all the quotes submitted via those separate sessions
- **Product:**

---

[2] Terminology is not consistent across exchanges. Terms such as "MP ID," "MM ID" and "EFID" are used by various exchanges to mean essentially the same thing. A single market-making firm will typically use more than one badge to stream quotes. We make the assumption here that a badge is a session-level concept (a session is associated with a single badge), but a given badge might be used on multiple sessions.
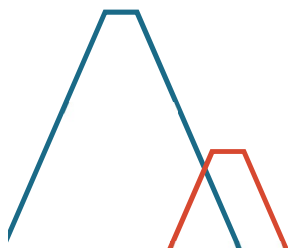
o MMP settings are associated with all quotes for instruments falling within an exchange-defined product.[3] Like those scoped at the badge level, protections can span sessions. However, they do not necessarily encompass all quotes on those sessions because each session might be used to stream quotes for more than one product

o The product-level scoping idea can cover more than one product, yielding protections at a more aggregate level. This is useful for market makers who post liquidity in options series spanning highly correlated products. An ability to specify distinct MMP settings at the aggregate level allows the market maker to cap exposures at levels less than the sum of the constituent product-level settings.

In practice, we find that session-level scoping is unnecessarily complicated because it introduces technical risk without significantly reducing market risk. That is why we recommend that exchanges support both badge and product level scopes for risk management while eliminating session scoping, or at least providing the ability to opt in and out of it.

Two scopes of protection still provide market makers with a choice in how they specify threshold values for the MMP metrics. Badge-level scoping allows the market maker to specify limits in a manner that suits specific strategies. Product-level scoping provides a firmwide perspective on the fundamental market risks for the various products being quoted, regardless of the associated badge. In practice, market makers will likely want to set an MMP threshold value for a given product that is lower than threshold values for the badge on which that product is quoted.

Likewise, with two scopes of protection, there are two ways an MMP breach can occur. When a product-level counter exceeds its threshold value, the breach occurs at the product level. This event would then impact all quotes submitted for that product, irrespective of the associated badge. Likewise, if it is a badge-level counter that exceeds its threshold, then the breach is confined to the associated badge. In that instance it will affect all quotes associated with that badge, irrespective of product.

LIMIT BREACH

A limit breach occurs when one of the MMP counters exceeds its threshold value. All quotes falling within that counter's scope of protection, and which are still resting on the book, must be prohibited from further execution.

In principle, market makers would like MMP counter values tallied and checked against threshold values between each and every matching event. That way, a breach

---

[3] Again, terminology is not consistent in this area. In this document, the term "product" can be loosely understood to mean an exchange-defined collection of (options) instruments. For example, a product might comprise all options series for a given underlying contract.
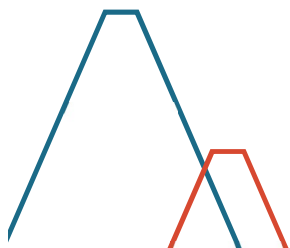
can be detected between successive matching events, allowing for preemptive action to be taken to protect the remaining liquidity. The standard way in which to do this is to remove (cancel) all remaining quotes from the book. Exchanges are generally consistent in this respect. However, their implementation differs, often in very subtle ways.

There are a number of nuances in how counter values are tallied. For example, some exchanges trigger a breach as soon as a counter reaches its threshold value. Others do so only when the threshold is exceeded. Further subtleties can arise in the context of pro-rata matching rules and implied matching in complex order books. All of these nuances make it hard for market makers to determine their real exposures and complicate administration of threshold settings. To a market maker, the ideal would be for a threshold value to define the maximum permissible value for a given counter.

From the market maker's perspective, the following sequence of events is optimal:

1. All quotes falling within the breaching limit's scope of protection are canceled from the book, and this occurs before any remaining matching events are permitted to occur. Any quote messages received after a breach but before a reset are rejected.
2. An explicit "MMP Breach" message is sent to all quoting sessions impacted by the event. This single message should embed at least the following information:
   a. Which MMP limit (i.e. which counter) was breached
   b. Which scope is impacted; and
   c. The number of quotes that have been removed from the book as a result
3. Upon receipt of the "MMP Breach" message, the market maker should be able to safely assume that all quotes within the associated scope of protection have been removed from the book. That is, there is no need to explicitly stream quote-canceled messages for all the affected quotes, an action that needlessly consumes matching engine computing resources and network bandwidth

To resume normal quoting activity after a breach, market makers should have to send an explicit "MMP Reset" message to the matching engine for all quoting sessions impacted by the breach.[4] Quote updates issued in the interim should be rejected for the affected sessions, and the reason for the rejections should be explicitly communicated back to the market maker. This requirement forces market makers to acknowledge the breach and gives them complete control as to when they re-engage with the market.

---

[4] Thus, the breach-in-effect state is a property of a quoting session.
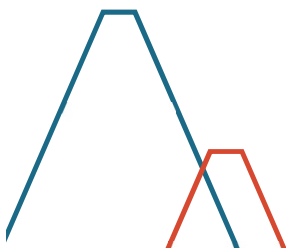
OPERATIONAL CONSIDERATIONS

The administrative and operational realities of MMP controls are as important as the technical and functional details.

SETTING LIMIT VALUES

For an optimal operational setup of MMPs, market makers need the ability to adjust settings and limits in real-time. Some exchanges permit this to be done only manually, either via a web-based portal or through their clearing firm. Others support it directly and electronically via a market maker's quoting sessions. A direct, electronic capability to set limits is preferable because it automates MMP administration, reducing the chances of human error.

Ideally, changes to MMP settings happen instantaneously. This allows market makers to benefit from being able to adapt quickly to changing market conditions. Furthermore, quick turnaround on limit changes allows market makers to correct setting misconfigurations as soon as they are detected. This would also allow greater flexibility and more dynamic risk management by firms providing liquidity, while keeping exchanges and clearing firms in control of the total potential exposure. This market structure change should improve liquidity levels during volatile markets.

AUDITING LIMIT VALUES

While support for electronic setting of limit values is useful, the ability to check and confirm current settings electronically is critically important. Market makers need to do this in order to ensure that the threshold values they have established at the exchange match the firm's intentions. In other words, support for electronic real-time check and confirm of current settings facilitates robust, mechanical auditing of an MMP's setup.

OPTIONALITY

An exchange without MMPs is not an attractive venue at which to post liquidity. MMP settings should therefore be mandatory on all badges and for all products quoted by market makers. Further, exchanges should not permit market makers to rely on default values for any limits – all thresholds should be explicitly specified by the market maker. This should be a prerequisite to streaming any quotes.

CONCLUSION

We base our views on the preferred characteristics of MMPs on our perspective as a professional options market-making firm, with experience in managing protection controls across a range of options quoting venues.

The concepts of dimension and scope of protection are the metrics best suited to model – and to bound – market makers' risks as well as the granularity at which they are able to associate their quotes with specific limit settings. We have distilled the
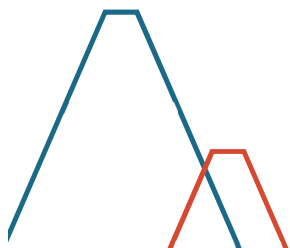
set of alternatives down to a small core which, we believe, provide sufficient protection and which lend themselves to sound administration in practice. That core consists of volume- and Greek-based dimensions of protection, with limit settings scoped at the badge- and product-level.

We've also weighed in on more pragmatic topics related to the operational integrity of the MMP setup. Chief among these is an ability to programmatically check and confirm current limit settings, as well as support for a mechanical audit of the current production setup. Lastly we recommend a level of flexibility and dynamic updating below the overall MMP cap of the mass-quoting firms.

Our hope is that these core principles and basic requirements will help guide exchanges in the design and evolution of their MMP infrastructure. With a robust set of MMPs, market-making firms will have greater confidence in their ability to provide a maximum level of liquidity when mass quoting, a benefit to the entire marketplace.

Have feedback? Reach out to the Optiver Corporate Strategy team at CHI_CorporateStrategy@optiver.us

ABOUT OPTIVER
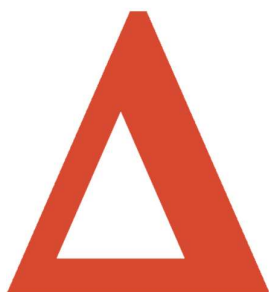
Optiver is a leading global electronic market maker, headquartered in Amsterdam, with over 1400 employees working from offices in Amsterdam, Austin, Chicago, Hong Kong, London, Shanghai, Singapore, Sydney and Taipei. Through pricing, execution and risk management, we provide liquidity to financial markets using our own capital, at our own risk, trading a wide range of products: listed derivatives, cash equities, ETFs, bonds and foreign currencies. Our independence allows us to improve the markets through pioneering trading strategies and sophisticated technology.