



# Simulating Financial Markets at the Atomic Level

September 2021

Time to read: 13 minutes


We demystify the inner workings of an exchange by working with a simulator which replicates the matching engine. The simulator allows market participants to interact with order flows instead of just working with static historical observations. The ability to simulate such interactions with the order flow of a market also enables participants to quantify the expected impact of a specific order and thereby develop execution algorithms which minimise such impact costs.

For institutional investor, qualified investor and investment professional use only. Not for retail public distribution.

## Authors



**James Powrie, PhD**  
Head of Execution  
Research, Man Group



**Zihao Zhang, PhD**  
Postdoctoral Researcher,  
Oxford-Man Institute



**Stefan Zohren, PhD**  
Senior Quant, Man Group  
Central Trading; Faculty  
Member, Oxford-Man  
Institute

## Introduction

It's fair to say that trading has become unrecognisable over the last 50 years, moving rapidly from face-to-face pit trading to a world where electronic trading is ubiquitous. But what is sometimes forgotten is the pace at which industry has continued to evolve during the last decade. With ever-more sophisticated technological advances and efficient data management, investors face increased competition. As such, great efforts have been made to generate returns by exploiting state-of-art algorithms, reducing latency and obtaining access to superior information.

Data has always been an essential driver of trading. It can be broadly categorised into price data, fundamental data and alternative data (such as texts and images).

High-frequency market data is the primary source of price data, and the detailed records of order flow and price dynamics provide a granular description of short-term supply and demand. Such so-called microstructure data has been studied extensively by researchers at the Oxford-Man Institute, including works by Zhang, Zohren, and Roberts (2019) and Zhang and Zohren (2021) that adopt deep learning algorithms to predict high-frequency price movements.

But understanding how new data and technological changes have affected trading and order flows is easier said than done. Face-to-face pit trading was at least easy to conceptualise; after electronic trading became the norm, market dynamics were still relatively comprehensible. But after the data deluge, it is sometimes hard to know where to begin when discussing the inner workings of modern exchanges.

To remedy this, we give a brief description of microstructure data and demystify the inner workings of an exchange. We do so by working with a simulator which replicates the matching engine of an exchange. The simulator further allows us to interact with order flows instead of just working with static historical observations. The ability to simulate such interactions with the order flow of a market enables us to quantify the expected impact of a specific order and thereby develop execution algorithms which minimise such impact costs.

Apart from replicating the matching logic, the simulator can inspire a range of interesting applications that are specific to high-frequency market data. We demonstrate a use case that investigates the price impact of different order sizes. The simulator can also be deployed to other application domains including trade execution, passive fills for fast trading strategies and general price discovery.

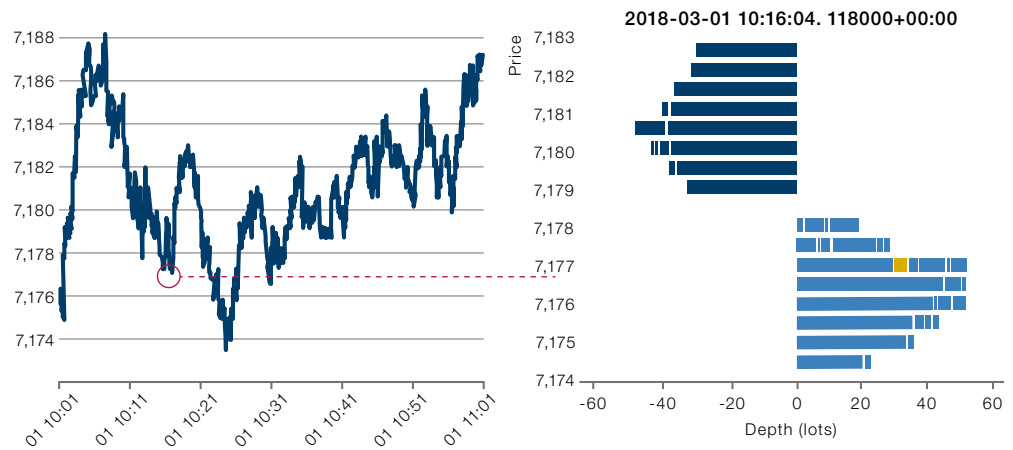
## Price Discovery Under the Microscope – Supply and Demand at the Microstructure Level

Most of us are used to standard price series of financial instruments. However, if we zoom in at any given moment in time, as shown in Figure 1, we observe that the price is just a manifestation of the available liquidity at the exchange at that instant. This is represented by the limit order book ('LOB') which is essentially a record of all outstanding intentions to trade. Traders can place various types of orders and a limit order allows traders to specify the price and quantity of interest. A limit order only executes if the market price reaches the specified price, otherwise it will stay alive in the order book until expiry or cancellation. On the contrary, a market order is matched immediately but with a higher cost, as it potentially executes against multiple levels in the order book. It is often used when immediate execution is required.



We demystify the inner workings of an exchange by working with a simulator which replicates the matching engine of an exchange.”

**Figure 1. Price Series at 1-Minute Frequency (Left) and LOB Corresponding to a Given Instance of Time (Right)**



Source: Man AHL. For illustrative purposes only.

Note: Levels in the LOB are made up of resting orders reflecting the supply and demand at that moment in time. The LOB snapshot (right) corresponds to the mid-price shown at the single point of time highlighted (left). Note that in the LOB, by convention, depth is shown positive for buy orders and negative for sell orders.

A LOB is one of the common ways to visualise high-frequency microstructure data. It has two sides: ask and bid, representing all orders to sell and buy, respectively. A LOB is sorted into different price levels based on the submitted orders. The first level of the ask (bid) is known as best ask (bid) because this is the best price level from which a buy (sell) order can be matched – in other words, the best ask (bid) is the lowest (highest) price visible on the orderbook at which a participant is willing to sell (buy).

Overall, a LOB reflects a snapshot of the current market environment and indicates the short-term demand and supply. The LOB data, up to a given depth of the LOB, is sometimes also referred to as level 2 (or simply L2) data among practitioners. L2 data provides more depth than the commonly used L1 data, which only contains the top-of-book (prices and sizes of best ask and bid), as well as recorded trades.

## The Inner Workings of an Exchange



Market by order data provides a sequence of order instructions from all market participants, each of which describes the action of a specific trader at a given time point.”

In this section, we illustrate the order lifecycle and demonstrate how an order book is maintained in an exchange. When an order is instructed by a trader, this message will flow through several intermediaries including exchanges, banks, brokers and clearing firms. Exchanges provide access to the ensemble of such messages as a real-time data feed, and the sequence of messages allows for the reconstruction of the LOB.

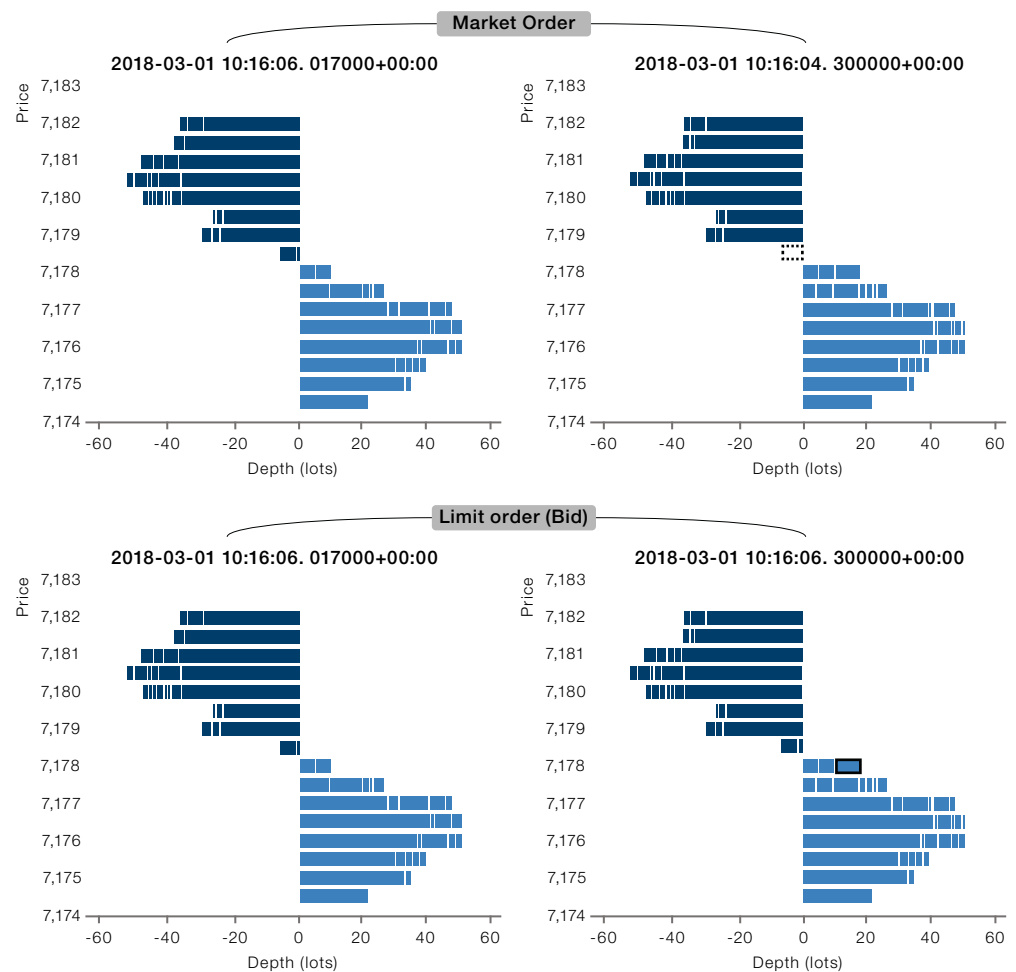
The full-resolution collection of these messages, which is termed market by order (‘MBO’) data, provides one of the most granular sources of microstructure information. Essentially, MBO data provides a sequence of order instructions from all market participants, each of which describes the action of a specific trader at a given time point. The core components of MBO data<sup>1</sup> consist of: time stamps recorded when an instruction is given; a unique and anonymous ID for order identification; and actions that indicate whether the trader is buying, selling or cancelling an existing order. Note that MBO data is sometimes also referred to as level 3 (or L3) data. By replaying L3 data through an exchange simulator, we can reconstruct the L2 data up to any price level – this is the internal state of the matching engine – as well as infer and record trades as contained in the L1 data. As such, L1 and L2 (LOB) data can be viewed as data derived from L3 (MBO) data.

1. The Oxford-Man Institute work of Zhang, Lim, and Zohren (2021) applies deep learning algorithms to MBO data to make predictions. Interested readers are referred to their work for detailed results.

MBO data also shows if an incoming order is a market or a limit order. If a market order is instructed to buy a certain number of shares, it will be matched immediately on the ask side of the LOB, as shown in Figure 2 (top chart). On the contrary, a trader can submit a limit order with specified price and volume on the bid side, as shown in the bottom chart in Figure 2.

At any given time, a market order is more expensive than a limit order because of the spread cross (i.e. the difference between the best ask and bid), paying a premium for immediate matching. A limit order is potentially cheaper, but it requires passive filling since participants must wait until the market reaches the specified price level. During this time, there is a risk of the price moving away. As a result, there is a time uncertainty for filling limit orders. Indeed, this is one of the reasons why a microstructure data simulator is needed, since it allows the study of trading behaviours for high-frequency strategies that require passive filling.

Figure 2. An Illustration of How MBO Data Updates a LOB



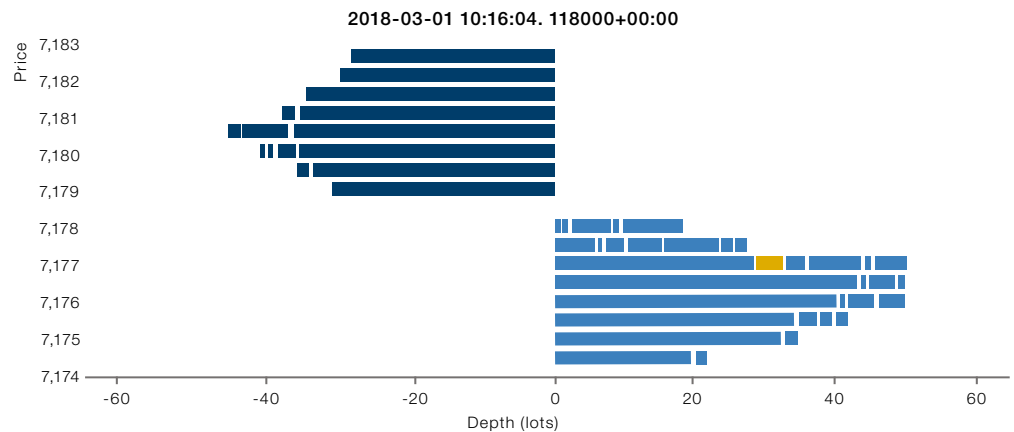
Source: Man AHL. For illustrative purposes only.

Top: A market order of "buy 6 shares" that results in immediate matching at price 7,178.5; Bottom: A limit order "buy 6 Shares at price 7,178" that waits on the LOB for passive filling.

## Simulating the Exchange's Matching Engine

In this section, we introduce a microstructure data simulator by explaining matching rules and presenting an example for passive filling. Figure 3 shows a snapshot of an animation from a simulator that presents how a typical LOB evolves throughout time.<sup>2</sup> A LOB updates whenever a message of MBO data comes in. Note that the time interval is irregular in physical clock time. In general, we expect more updates at the open and close of markets because of larger trading volumes and higher volatility.

Figure 3. A LOB Through a Market Simulator



Source: Man AHL. For illustrative purposes only.  
Note: A given limit order is marked in yellow.



In general, the most common priority mechanism is price-time, which means that for active orders, priority is given to those with the best price, and ties are broken by selecting the order with the earliest submission time.”

Note that several active orders can have the same price at a given time. LOBs employ a priority system for matching active orders at each individual price level. In general, the most common priority mechanism is price-time, which means that for active orders, priority is given to those with the best price, and ties are broken by selecting the order with the earliest submission time. Another priority mechanism is pro-rata, which is popular in some futures markets. Under this mechanism, each relevant active order receives a share of the matched quantity proportional to the fraction of the depth available that it represents at that price. There are also other, less common matching mechanisms.

Our simulation is essentially a replay of the historical messages through an artificial matching engine. To illustrate this, Figure 3 shows the evolution of a LOB and demonstrates the life cycle of a specific order (marked yellow) from placement to settlement.

Being able to simulate markets to this degree of granularity can be useful in several important applications in execution and fast trading, including:

- Determining fills of passive orders;
- Analysing price impact;
- Building, testing and refining execution algorithms.

### Analysis of Price Impact Through Microstructure Simulation

Another use case of the simulator can be illustrated by focusing on the analysis of price impact. Price impact refers to the effect the participant's own trading has on the price. It is difficult to study in live trading as there is no way of knowing what would have happened had a trade not been placed.

Empirically, we can only observe what happens on average on a large number of orders, where, in absence of trading, the price would move up or down equiprobably. In instances where a trade is conducted, we observe a subtle upward price trend in

2. To see the full animation, please visit: [www.man.com/sites/public/mangroup/content/research/lob\\_animation.gif](http://www.man.com/sites/public/mangroup/content/research/lob_animation.gif)




The powerful aspect of having a market simulator is that it effectively enables the re-running of different realities where our interventions with the market either do, or do not, occur.”

buy cases and a downward trend in sell cases. The powerful aspect of having a market simulator is that it effectively enables the re-running of different realities where our interventions with the market either do, or do not, occur.

Note that there are generally two types of price impact. Firstly, there is mechanical or direct price impact. This happens, for example, when a market order to buy is placed. This removes the best ask and, as a result, both the best ask price and the mid-price change. This type of price impact is automatically captured by the simulator as it is a direct consequence of the matching logic. Secondly, there is indirect price impact, which is also referred to as signalling impact. This is where one participant’s actions influence others. In the above example, even if not using a market order, a large buy limit order resting on the best bid would indicate a strong desire to buy, and even though no trading has occurred could encourage sellers on the best ask to cancel their orders and resubmit them at a higher price. Indirect price impact is harder to model, but even simple probabilistic models mimicking the reaction of other traders on a microstructure level can reproduce complex empirical price impact formulas at the aggregate level.

Figure 4 shows the analysis of price impact for varying order sizes obtained using simulation experiments where we replay the past over-and-over again, interacting with the market in different ways. As expected, we observe that the larger the order size, the higher the price impact. More precisely, we see that the functional relationship between order size and impact is well described by a square root. This is in line with a well-known empirical finding, also known as the square-root law, which states that price impact increases with the square root of order size. Interestingly, this empirical formula is not explicitly used. Rather, mechanical price impact and simple probabilistic logic on an order-by-order level naturally give rise to this behaviour in aggregate.

Figure 4. Analysis of Price Impact for Different Order Sizes



Source: Man AHL. For illustrative purposes only.

Note: When plotting the data on a log-log plot, the well-known square root law of price impact emerges from first principles.



In the absence of a simulator, research is often guided by a combination of intuition, and trial and error. However, learning by trial and error can be a slow and expensive process.”

---

The above results demonstrate the potential cost from executing large orders as less favourable prices are obtained. As a result, execution algorithms play an important role for large investment firms as they split larger orders into sequences of smaller size orders, avoiding severe price impacts and thereby saving costs. The study of such execution algorithms requires an interaction with order books, so the simulator facilitates this type of research by allowing trading behaviours and transaction costs for different execution algorithms to be examined. In the absence of a simulator, research is often guided by a combination of intuition, and trial and error. In the complex world of market microstructure, however, insights that seem intuitively appealing can often turn out to be flawed. Learning by trial and error can also be a slow and expensive process.

Further, the simulator allows the use of agent-based models for general price discovery by explicitly modelling price dynamics. The simulator can be further utilised as an environment for Reinforcement Learning algorithms, directly minimising execution costs through an agent interacting with the environment.

## Conclusion

We introduce high-frequency microstructure data and describe the role of limit order books ('LOBs') in price formation. Specifically, we demonstrate how a LOB evolves by processing incoming order placements and cancellations, and discuss the order lifecycle within an exchange. We do so by introducing a simulator which replays the various matching rules of an exchange. A use case is presented to demonstrate how the simulator can be used to study price impact for different order sizes, an analysis which is vital for studying execution algorithms. Beyond execution, other important applications include the simulation of passive fills in faster trading strategies.

## References

- Zhang, Zihao, Bryan Lim, and Stefan Zohren. "Deep Learning for Market by Order Data." arXiv:2102.08811 (2021). [Full paper](#).
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. "DeepLOB: Deep convolutional neural networks for limit order books." IEEE Transactions on Signal Processing 67 (11): 3001–3012. arXiv:1808.03668 (2019). [Full paper](#).
- Zhang, Zihao, and Stefan Zohren. "Multi-Horizon Forecasting for Limit Order Books: Novel Deep Learning Approaches and Hardware Acceleration using Intelligent Processing Units" arXiv: 2105.10430 (2021). [Full paper](#).

## Authors

---

### James Powrie, PhD

Head of Execution Research, Man Group



James Powrie is Head of Execution Research at Man Group. James is responsible for all aspects of algorithmic execution, optimal flow routing and trading analytics. Prior to taking on his current role in 2018, James has held multiple positions across Man Group, including Quantitative Analyst and Head of Systems Development having first joined the Man AHL team in 1992. He was also instrumental in the design of Man AHL's Futures Execution algorithm. James holds a PhD in Computer Science and a BSc in Electronics with Computer Science from the University of Wales.

---

### Zihao Zhang, PhD

Postdoctoral Researcher, Oxford-Man Institute



Zihao Zhang is a postdoctoral researcher at the Oxford-Man Institute and Machine Learning Research Group at the University of Oxford. Zihao's research focuses on quantitative finance with a special emphasis on applying deep learning models to financial time series modelling. Zihao's current projects include portfolio optimisation and reinforcement learning. Zihao holds a PhD and MSc in Applied Statistics from the University of Oxford and a BSc in Economics and Statistics from University College London.

---

### Stefan Zohren, PhD

Senior Quant, Man Group Central Trading; Faculty Member, Oxford-Man Institute



Stefan Zohren is a Senior Quant at Man Group's Central Trading division, where he is responsible for execution signals research across the range of asset classes traded by the company. Much of his research leverages modern machine learning techniques and involves close collaboration with Man Group's investment engines. Stefan has worked with Man Group's Central Trading division since 2018, first as an External Scientific Advisor, and then transitioning to his current role in early 2020. Stefan also serves as a Faculty Member of the Oxford-Man Institute at the University of Oxford, where his research is focused on machine learning in finance, including deep learning, reinforcement learning, network and NLP approaches, as well as early use cases of quantum computing. Before that he worked as a Quantitative Strategist at two leading high-frequency trading firms. Stefan holds a PhD in Mathematical Physics from Imperial College.



## Simulated Hypothetical Results

Simulated Hypothetical Results are calculated in hindsight, invariably show positive rates of return, and are subject to various modeling assumptions, statistical variances and interpretational differences. No representation is made as to the reasonableness or accuracy of the calculations or assumptions made or that all assumptions used in achieving the results have been utilized equally or appropriately, or that other assumptions should not have been used or would have been more accurate or representative. Changes in the assumptions would have a material impact on the Hypothetical Results and other statistical information based on the Hypothetical Results.

The Hypothetical Results have other inherent limitations, some of which are described below. They do not involve financial risk or reflect actual trading by an Investment Product, and therefore do not reflect the impact that economic and market factors, including concentration, lack of liquidity or market disruptions, regulatory (including tax) and other conditions then in existence may have on investment decisions for an Investment Product. In addition, the ability to withstand losses or to adhere to a particular trading program in spite of trading losses are material points which can also adversely affect actual trading results. Since trades have not actually been executed, Hypothetical Results may have under or over compensated for the impact, if any, of certain market factors. There are frequently sharp differences between the Hypothetical Results and the actual results of an Investment Product. No assurance can be given that market, economic or other factors may not cause the Investment Manager to make modifications to the strategies over time. There also may be a material difference between the amount of an Investment Product's assets at any time and the amount of the assets assumed in the Hypothetical Results, which difference may have an impact on the management of an Investment Product. Hypothetical Results should not be relied on, and the results presented in no way reflect skill of the investment manager. A decision to invest in an Investment Product should not be based on the Hypothetical Results.

No representation is made that an Investment Product's performance would have been the same as the Hypothetical Results had an Investment Product been in existence during such time or that such investment strategy will be maintained substantially the same in the future; the Investment Manager may choose to implement changes to the strategies, make different investments or have an Investment Product invest in other investments not reflected in the Hypothetical Results or vice versa. To the extent there are any material differences between the Investment Manager's management of an Investment Product and the investment strategy as reflected in the Hypothetical Results, the Hypothetical Results will no longer be as representative and their illustration value will decrease substantially. No representation is made that an Investment Product will or is likely to achieve its objectives or results comparable to those shown, including the Hypothetical Results, or will make any profit or will be able to avoid incurring substantial losses. Past performance is not indicative of future results and simulated results in no way reflect upon the manager's skill or ability.

## Important Information

This information is communicated and/or distributed by the relevant Man entity identified below (collectively the 'Company') subject to the following conditions and restriction in their respective jurisdictions.

Opinions expressed are those of the author and may not be shared by all personnel of Man Group plc ('Man'). These opinions are subject to change without notice, are for information purposes only and do not constitute an offer or invitation to make an investment in any financial instrument or in any product to which the Company and/or its affiliates provides investment advisory or any other financial services. Any organisations, financial instrument or products described in this material are mentioned for reference purposes only which should not be considered a recommendation for their purchase or sale. Neither the Company nor the authors shall be liable to any person for any action taken on the basis of the information provided. Some statements contained in this material concerning goals, strategies, outlook or other non-historical matters may be forward-looking statements and are based on current indicators and expectations. These forward-looking statements speak only as of the date on which they are made, and the Company undertakes no obligation to update or revise any forward-looking statements. These forward-looking statements are subject to risks and uncertainties that may cause actual results to differ materially from those contained in the statements. The Company and/or its affiliates may or may not have a position in any financial instrument mentioned and may or may not be actively trading in any such securities. This material is proprietary information of the Company and its affiliates and may not be reproduced or otherwise disseminated in whole or in part without prior written consent from the Company. The Company believes the content to be accurate. However accuracy is not warranted or guaranteed. The Company does not assume any liability in the case of incorrectly reported or incomplete information. Unless stated otherwise all information is provided by the Company. Past performance is not indicative of future results.

Unless stated otherwise this information is communicated by the relevant entity listed below.

**Australia:** To the extent this material is distributed in Australia it is communicated by Man Investments Australia Limited ABN 47 002 747 480 AFSL 240581, which is regulated by the Australian Securities & Investments Commission (ASIC). This information has been prepared without taking into account anyone's objectives, financial situation or needs.

**Austria/Germany/Liechtenstein:** To the extent this material is distributed in Austria, Germany and/or Liechtenstein it is communicated by Man (Europe) AG, which is authorised and regulated by the Liechtenstein Financial Market Authority (FMA). Man (Europe) AG is registered in the Principality of Liechtenstein no. FL-0002.420.371-2. Man (Europe) AG is an associated participant in the investor compensation scheme, which is operated by the Deposit Guarantee and Investor Compensation Foundation PCC (FL-0002.039.614-1) and corresponds with EU law. Further information is available on the Foundation's website under [www.eas-liechtenstein.li](http://www.eas-liechtenstein.li). This material is of a promotional nature.

**European Economic Area:** Unless indicated otherwise this material is communicated in the European Economic Area by Man Asset Management (Ireland) Limited ('MAMIL') which is registered in Ireland under company number 250493 and has its registered office at 70 Sir John Rogerson's Quay, Grand Canal Dock, Dublin 2, Ireland. MAMIL is authorised and regulated by the Central Bank of Ireland under number C22513.

**Hong Kong SAR:** To the extent this material is distributed in Hong Kong SAR, this material is communicated by Man Investments (Hong Kong) Limited and has not been reviewed by the Securities and Futures Commission in Hong Kong. This material can only be communicated to intermediaries, and professional clients who are within one of the professional investors exemptions contained in the Securities and Futures Ordinance and must not be relied upon by any other person(s).

**Japan:** To the extent this material is distributed in Japan it is communicated by Man Group Japan Limited, Financial Instruments Business Operator, Director of Kanto Local Finance Bureau (Financial instruments firms) No. 624 for the purpose of providing information on investment strategies, investment services, etc. provided by Man Group, and is not a disclosure document based on laws and regulations. This material can only be communicated only to professional investors (i.e. specific investors or institutional investors as defined under Financial Instruments Exchange Law) who may have sufficient knowledge and experience of related risks.

**Switzerland:** To the extent this material is distributed in Switzerland, this material is communicated by Man Investments AG, which is regulated by the Swiss Financial Market Authority FINMA.

**United Kingdom:** Unless indicated otherwise this material is communicated in the United Kingdom by Man Solutions Limited ('MSL') which is an investment company as defined in section 833 of the Companies Act 2006. MSL is registered in England and Wales under number 3385362 and has its registered office at Riverbank House, 2 Swan Lane, London, EC4R 3AD, United Kingdom. MSL is authorised and regulated by the UK Financial Conduct Authority (the 'FCA') under number 185637.

**United States:** To the extent this material is distributed in the United States, it is communicated and distributed by Man Investments, Inc. ('Man Investments'). Man Investments is registered as a broker-dealer with the SEC and is a member of the Financial Industry Regulatory Authority ('FINRA'). Man Investments is also a member of the Securities Investor Protection Corporation ('SIPC'). Man Investments is a wholly owned subsidiary of Man Group plc. The registration and memberships described above in no way imply a certain level of skill or expertise or that the SEC, FINRA or the SIPC have endorsed Man Investments. Man Investments, 452 Fifth Avenue, 27th fl., New York, NY 10018.

This material is proprietary information and may not be reproduced or otherwise disseminated in whole or in part without prior written consent. Any data services and information available from public sources used in the creation of this material are believed to be reliable. However accuracy is not warranted or guaranteed. ©Man 2021.

MKT002218/NS/GL/W