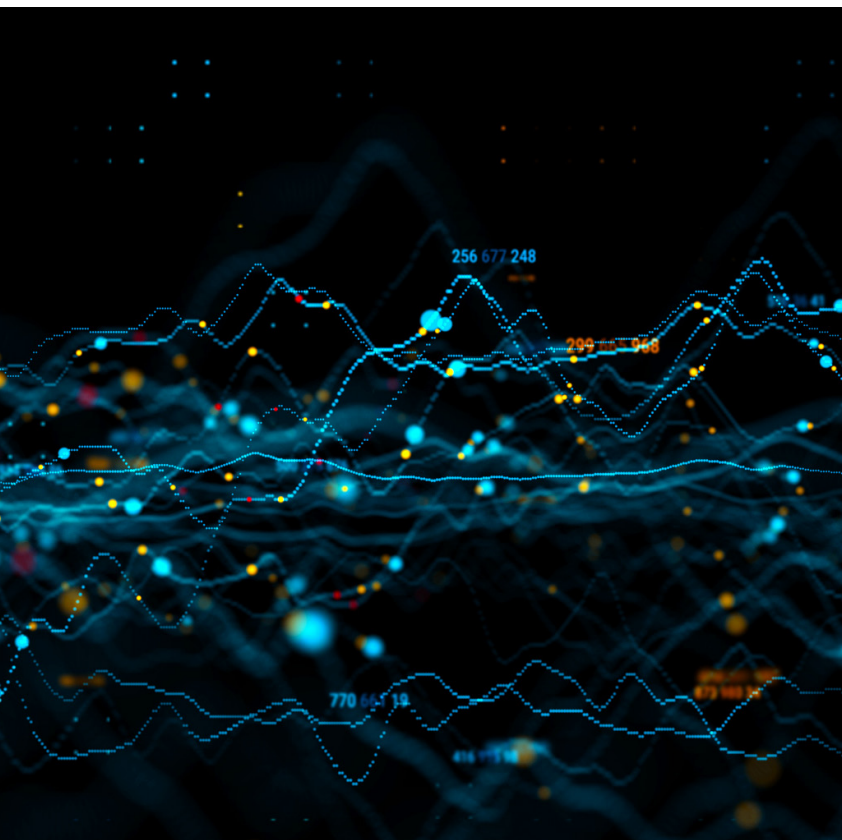


Composite+™

Algorithmic Pricing for the European Government Bond Market



Julien Alexandre

Global Head of Research

Chisom Amalunweze

Senior Research Analyst

Michael Wehbe

Research Analyst

Introduction

Composite+ (CP+) is our AI-powered algorithmic pricing engine for global fixed income securities leveraging our proprietary Trax® data and the MarketAxess trading platform. It is designed to support a variety of trading functions including pre-trade price discovery, auto-execution, crossing, liquidity provision, and transaction cost analysis.

By calibrating to dealer-to-client transactions, CP+ brings pricing continuity to the European Government Bond market. It produces an unbiased, two-sided price for more than 600 European Government Bonds.

In this paper, we will review the CP+ methodology, characteristics, and applications with a focus on European Government Bonds, diving deeper into the design and potential roles for CP+ in traders' workflows.



Key differentiators

AI-powered pricing

A robust machine learning algorithm calibrated to the traded market.

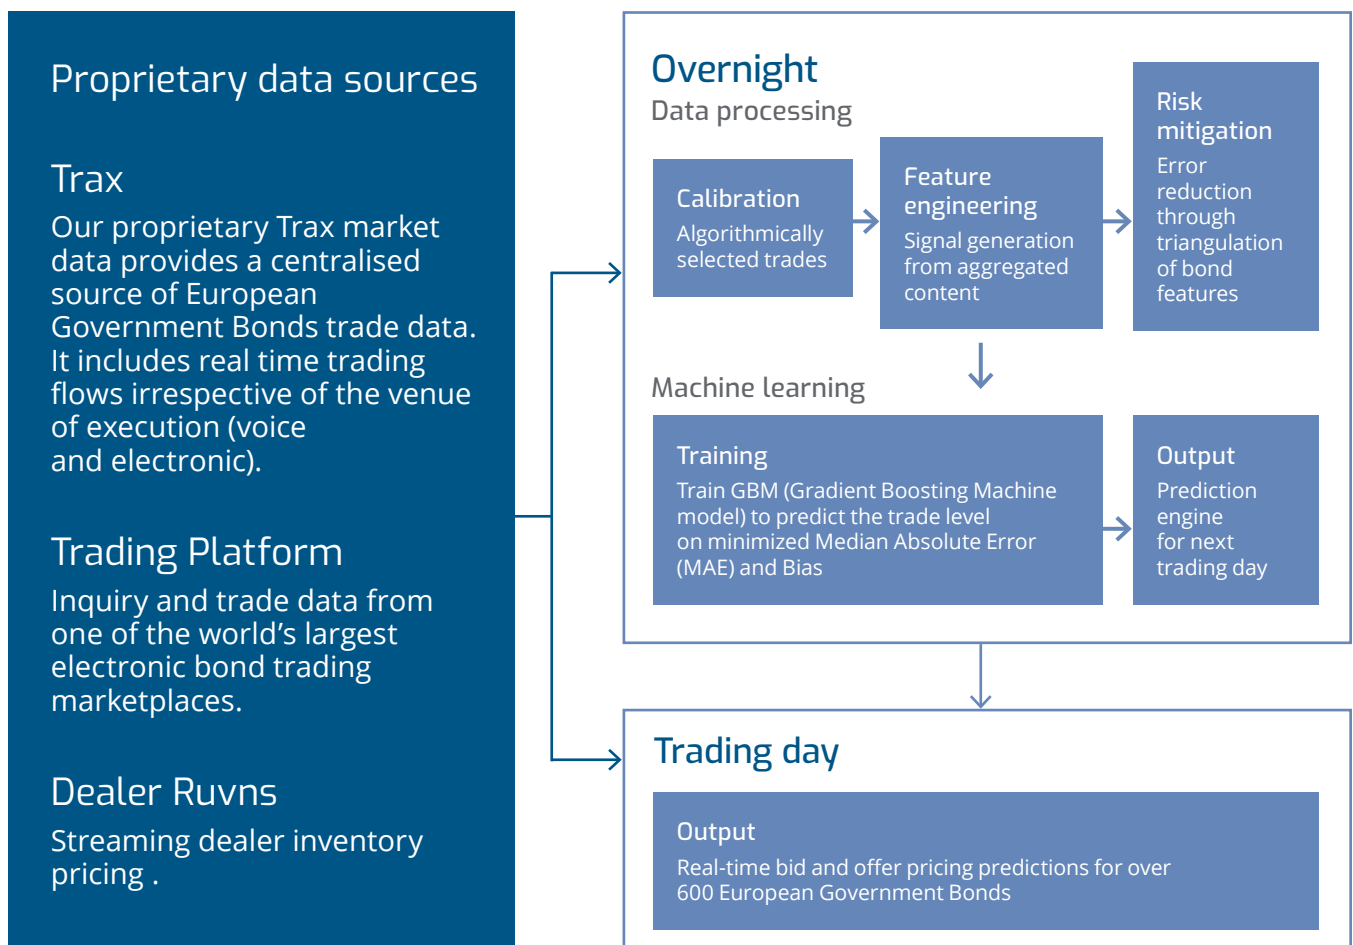
Trax® European trade data

Unique and centralised data source of cleansed fixed income trades reported in real time. It includes dealer-to-client and dealer-to-dealer flow irrespective of the venue of execution—be it voice or electronic platform.

MarketAxess global trading platform

Data sourced from the MarketAxess trading system, which connects over **1,800 global institutions**.

Data and construction



Construction

Composite+ works in two steps:

- a. The prediction engine is trained daily using historical data.
- b. The engine generates its predictions in real-time while incorporating all available intraday information. All predictions are unique combinations of time, bond, and direction.

Calibration

The prediction engine rests on a supervised machine learning algorithm. The supervisory variable is the price of all qualifying reported dealer-to-client Trax trades. Qualifications help narrow the subset of Trax trades to those that best represent dealer risk pricing for institutional clients in a competitive market. For example, we limit consideration to institutionally sized trades \$1mm or larger, and exclude seemingly non-principal risk trades—e.g., paired transactions for which there is a corresponding dealer-to-client trade of the same size at roughly the same time in the opposite direction. These conditions are essential because they define the calibration of the algorithm.

CP+ at time t predicts the likeliest trade price if there was a qualifying dealer-to-client trade at this same time t .

With this goal and design, CP+ brings pricing continuity to the dealer-to-client segment of the market.

Prediction

The prediction engine employs unique sources of bond trading data: earlier Trax trades, indicative bond price data streamed by dealers, and RFQ responses sent by liquidity providers via the MarketAxess trading platform.

We begin by creating a set of features, best defined as explanatory signals that aggregate the data from the past and the present. Each feature has a different predictive power that will depend on the vicinity to the prediction.

The most proximate features include indicative levels streamed by dealers and the most recent Trax traded prices for the bond of interest.

Additional subsets of features include progressively relaxed constraints. For example, features variously consider activity on the opposite side and in the interdealer market, and activity in other bonds with matching attributes. By aggregating a broad and diverse set of data sources, including proprietary Trax and MarketAxess trading platform activity, it predicts subsequent trade levels with minimal error.

“Overall, the algorithm consumes more than 200 features since it can make good use of even low information content factors to improve its prediction. The process also calls an outlier detection mechanism to identify and exclude off-market levels.”

DAVID KREIN, Global Head of Research and Data

Even when the most proximate signals are missing, the algorithm will learn to triangulate the less predictive features successfully. This approach leverages all the available information to increase coverage without impacting quality.

Overall, the algorithm consumes more than 200 features since it can make good use of even low information content factors to improve its prediction. The process also calls an outlier detection mechanism to identify and exclude off-market levels.

The features are fed into a tree-based machine learning algorithm called Gradient Boosting Method (GBM). GBM provides better accuracy than other algorithms like deep learning and neural networks since it is better at identifying the signal within noisy bond market data.

Coverage

Currency	Country	Avg Number of Bonds	% of Trax prints
Euro	Austria	27	100%
Euro	Belgium	35	96%
Euro	Finland	21	100%
Euro	France	85	99%
Euro	Germany	73	98%
Euro	Greece	15	100%
Euro	Ireland	17	100%
Euro	Italy	119	100%
Euro	Luxembourg	5	100%
Euro	Netherlands	30	100%
Euro	Portugal	18	100%
Euro	Spain	48	99%
GBP	United Kingdom	86	98%
Swiss Franc	Switzerland	20	100%
Norwegian Krone	Norway	7	100%
Swedish Krona	Sweden	5	100%
Danish Krone	Denmark	6	100%

February 2021 Data | Source: MarketAxess

Benchmarking

The CP+ algorithm—designed to predict, with minimal error, the price of qualified dealer-to-client Trax trades that are \$1mm or larger—can easily be tested out-of-sample with historical Trax data.

Specifically, we benchmark Trax traded prices above \$1mm to the CP+ prediction at the time of the trade.

Further, we can benchmark those same Trax traded prices to the Trax dataset itself, using the most recent, or last, Trax traded price on the same side and with the same constraints listed above.

First, we define the delta (Δ):

$\Delta = \text{Price}_{\text{trade}} - \text{Price}_{\text{CP+}/\text{last trade}}$ if Dealer Sells

$\Delta = \text{Price}_{\text{CP+}/\text{last trade}} - \text{Price}_{\text{trade}}$ if Dealer Buys

For example, if a dealer sells a bond at €99.78, while CP+ offer is at €99.81, then the Δ for this trade is negative three cents; if a dealer buys a bond at €99.78 and CP+ bid is at €99.81, then the Δ for this trade is positive three cents.

The error (Median Absolute Error or MAE) is defined as the median of all absolute distances (Δ).



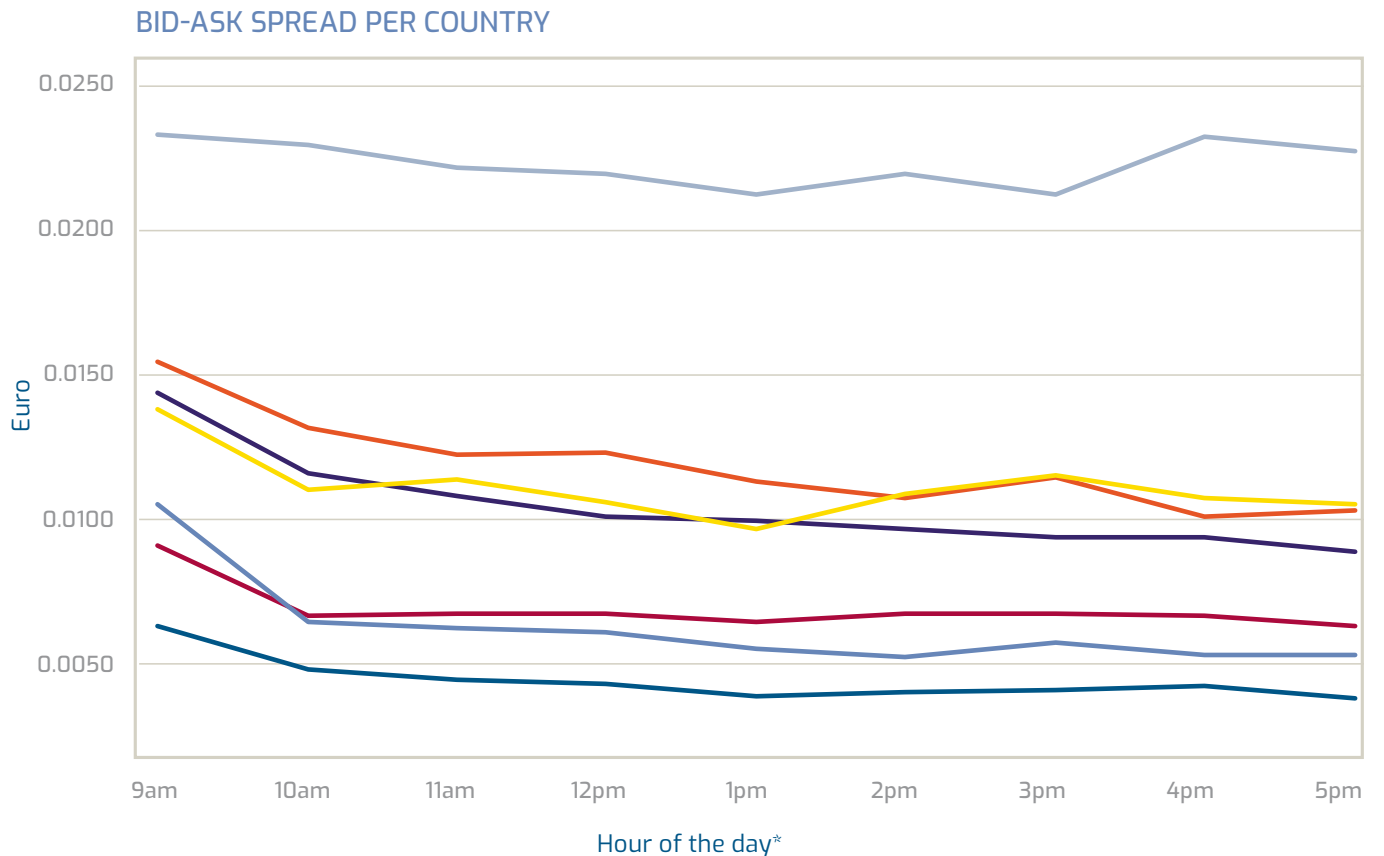
Benchmark performance

Currency	Aggregate	Trade Count	Composite+ Median Absolute Error (MAE, cents)	Last Price Median Absolute Error (MAE, cents)
Euro		52,833	2.25	4.65
Side				
Euro	Dealer Buy	27,039	2.20	4.50
Euro	Dealer Sell	25,794	2.30	4.80
Size Bucket				
Euro	\$1-10mm	31,833	2.50	5.20
Euro	>\$10mm	21,000	1.90	3.80
Country				
Euro	Austria	572	3.30	13.60
Euro	Belgium	3,838	5.10	13.00
Euro	Finland	646	3.10	12.00
Euro	France	8,173	2.60	7.00
Euro	Germany	12,343	1.80	3.80
Euro	Greece	797	5.60	10.00
Euro	Ireland	386	2.50	11.80
Euro	Italy	17,175	1.50	2.40
Euro	Luxembourg	<250	3.70	16.70
Euro	Netherlands	1,440	2.30	6.70
Euro	Portugal	1,015	4.30	8.30
Euro	Spain	6,428	3.60	6.00
GBP	United Kingdom	19,445	5.10	13.20
Swiss Franc	Switzerland	<250	5.20	10.00
Norwegian Krone	Norway	<250	4.20	16.50
Swedish Krona	Sweden	<250	4.60	12.50
Danish Krone	Denmark	<250	5.70	17.80

February 2021 Data | Source: MarketAxess

Applications

Intraday bid-ask spread



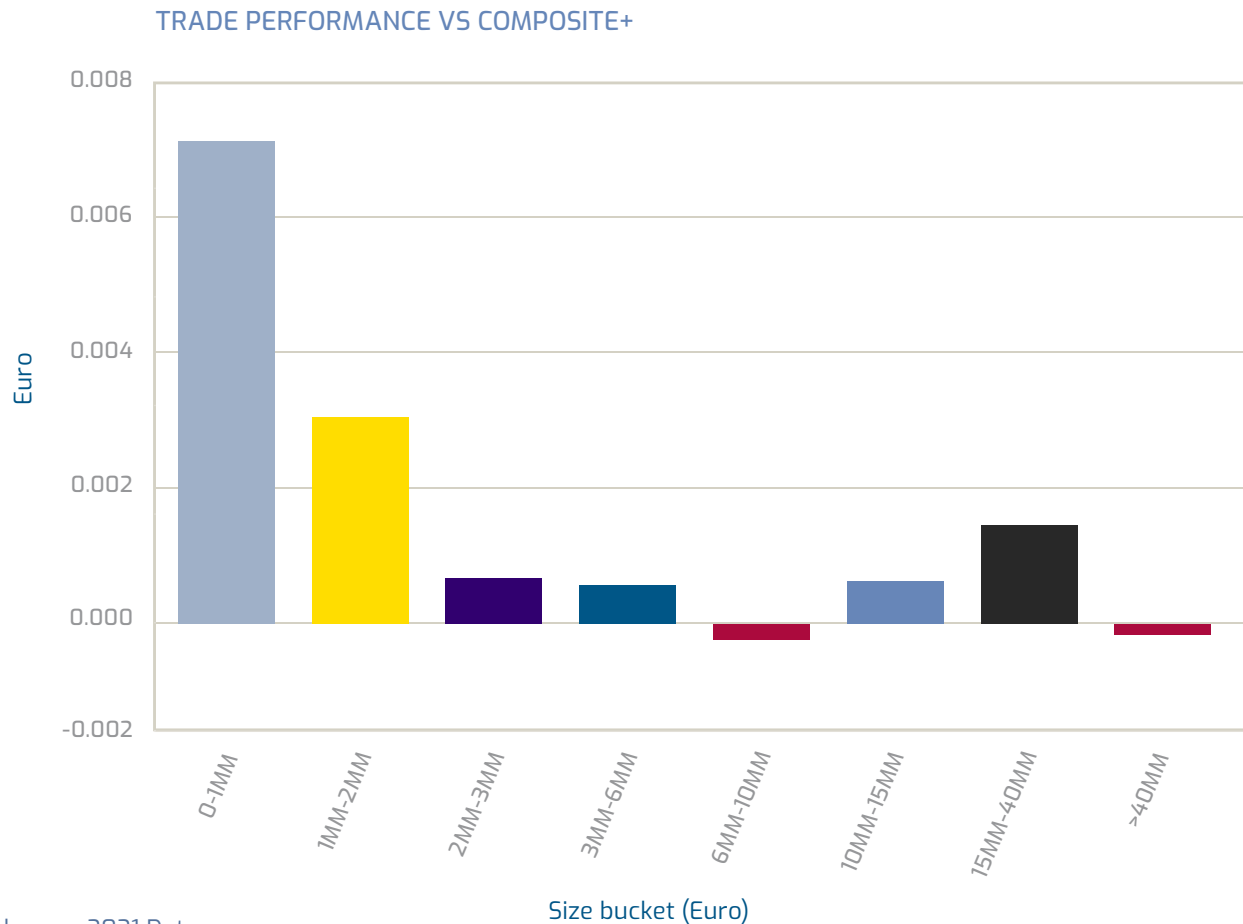
February 2021 Data
Source: MarketAxess
*London local time

— Austria
— Belgium
— Spain
— Netherlands
— Italy
— France
— Germany

The graph shows the hourly bid-ask spreads for Euro Denominated Government Bonds in February 2021 measured in EUR terms. We see that the hourly bid-ask decreases throughout the day as the trading information increases.

This is consistent with expectations. In early trading, there is little information available to the market, and by extension the CP+ model reflects that risk with wider bid-ask spreads. As the day goes by, more traded prices and RFQ responses come into play and we observe the bid-ask spreads tighten accordingly.

Transaction cost analysis



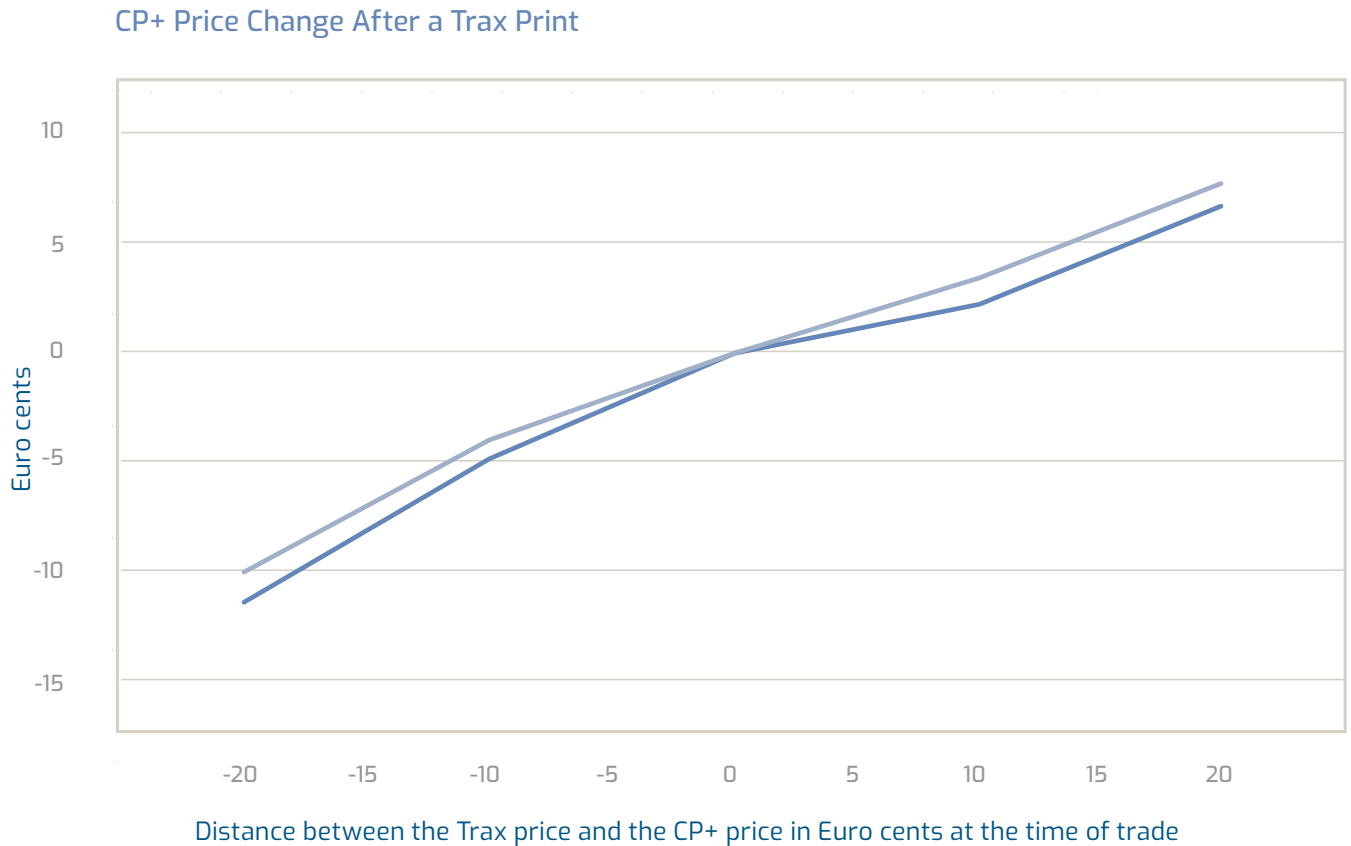
February 2021 Data
Source: MarketAxess

The graph shows the transaction cost for Euro Denominated Government Bonds (EUR) by size bucket. Transaction cost is the distance between the trade price and the CP+ price on the corresponding side. A negative value on the y-axis means that the trade was inside (more competitive than) CP+. A positive value means that it was outside (less competitive than) CP+. Each size bucket's value in the chart is the median of all observations. For example, if a dealer or liquidity maker sells a bond at €99.78, while CP+ Offer is at €99.81, then the dealer's transaction cost is negative three cents if a

dealer buys a bond at €99.78 and CP+ Bid is at €99.81, then the dealer's transaction cost is positive three cents.

The graph demonstrates that trades under \$1mm are priced wider than CP+ while trades above \$1mm are inline with CP+. This confirms the assumption that institutional-sized trades start at \$1mm and that calibrating the algorithm on these trades only, reduces the bias while including the maximum number of realized samples.

Price impact of Trax trade on Composite+(1)



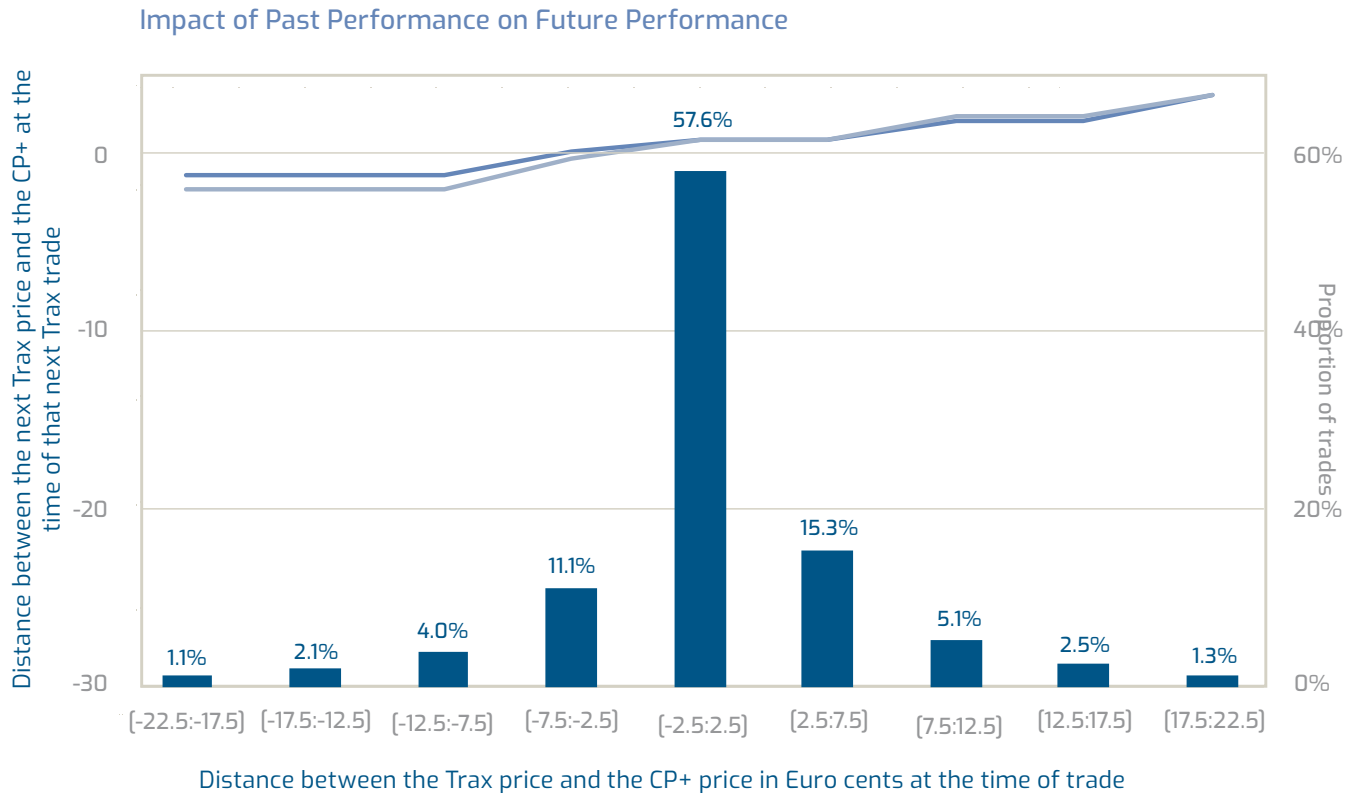
February 2021 Data
Source: MarketAxess

— Dealer Buy
— Dealer Sell

The graph above represents the movement of CP+ after a trade hits Trax. The x-axis is the distance between the Trax price and the CP+ price at the time of the trade (i.e., generated slightly before the trade occurred). On the y-axis lays the price impact, defined as the median difference between the CP+ 15 minutes after the trade and the CP+ at the time of the trade.

When a trade hits Trax and deviates from CP+ by Δ , the impact on CP+ is less than Δ . The algorithm still weighs the other features when adjusting its price. That's especially important since bonds don't trade continuously; therefore other data points are better at capturing price movement. In other words, even though the most recent Trax trade price is one of the most important features that CP+ uses, a single trade can't be the only source that sets the price of a bond. Every other meaningful data point, even those less proximate to the prediction, will have an influence on CP+.

Price impact of Trax trade on Composite+(2)



The graph above evaluates how optimal the CP+ price adjustment is after a Trax trade occurs. It uses the same x-axis as the previous graph, i.e., the distance between the Trax price and the CP+ price at the time of the trade. The y-axis shows the distance between the next Trax trade price on the same bond as the x-axis and the CP+ at the time of that next Trax trade. The quasi-flat line shows that the CP+ error exhibits

almost no autocorrelation. Whether a Trax trade was very close to CP+ or not, CP+ adjusted optimally. It did not overreact, nor did it underreact. It was able to evaluate very precisely how much to adjust in each situation. The bar chart gives the reader context around distribution of CP+ error. More than 57% of trades happen at a price less than 2.5 cents away from CP+.

Conclusion

- By leveraging the latest AI techniques, CP+ functions to transform proprietary and public data points into an unbiased two-way reference price for more than 600 individual European Government Bonds.
 - By aggregating a broad and diverse set of data sources, including Trax, it predicts subsequent trade levels with minimal error.
 - CP+'s calibration on institutionally sized risk-based trades gives market participants a robust, precise, meaningful, and consistent price.
-

For more information, please contact research@marketaxess.com



Now you're in the market™

©2023 MarketAxess Holding Inc. (the “Company”). MarketAxess Corporation is a member of FINRA and SIPC. MarketAxess Canada Company is a member of the New SRO. MarketAxess Europe Limited (04017160), MarketAxess Capital Limited (09777893), and MarketAxess Post-Trade Limited (01917944) are incorporated in England and Wales, and are each authorised and regulated by the Financial Conduct Authority. MarketAxess NL B.V. (69592888) and MarketAxess Post-Trade B.V. (69597774) are incorporated in the Netherlands, and are each licensed by the Netherlands Authority for the Financial Markets. MarketAxess Singapore Pte. Limited (201702319R) is incorporated in Singapore and recognised by the Monetary Authority of Singapore. MarketAxess is a trading name of the aforementioned entities. To the fullest extent permitted by applicable law, all warranties and representations are disclaimed. All information presented herein is considered to be accurate at the time of writing, provided “as is” with no express or implied representations or warranties—including without limitation in relation to fitness for purpose, accuracy, or completeness of contents. Whilst sources of information are believed to be reliable, MarketAxess does not independently verify information obtained from third party sources. Any information presented is not intended to represent an offer or solicitation of any financial instrument. For the purposes of clients accessing MarketAxess services from EU27 countries, this information is provided by MarketAxess NL B.V. and /or MarketAxess Post-Trade B.V., as appropriate. Proprietary and confidential information of the Company—all forms of copying prohibited. Further information can be found on our website at [marketaxess.com](https://www.marketaxess.com).

01/25/2023