

Turning art into science

Pricing new issues on intuition and market feel is ancient history. Artificial intelligence and algorithms are setting the market price for credit, using factors and correlations humans can guess but not follow. Is AI the latest black box risk that will bring illiquid credit markets low or could it make them more efficient?

By: Peter Lee

All around the securities markets, technologists are trying to turn what was once an art into a science. The new game is to recreate using technology what the great traders, investors and bankers of the past used to carry in their heads – a strong sense of who owned what, who wanted to buy, sell or issue what and how all that pent-up desire for action mapped against prevailing market conditions.

Whether at banks, fintechs, trading platforms or technology companies, the new masters of the markets want to replace intangible market feel, intuition and creativity with machine learning; to ditch anecdotes and rely instead on data – the bigger, the better.

Artificial, not human, intelligence increasingly drives the bond and credit markets just as it does equity – and so drives the cost at which governments, companies and people borrow money.

Everyone in finance has to adapt. When Euromoney sits down with Ken Jacobs, chief executive of Lazard, we are expecting to talk about the great takeover battles of history. But even here, the talk is all about AI moving the markets through which the firm advises its clients how to tread.

In February, Lazard announced the expansion of Lazard Asset Management's alternative investments platform with the addition of a new team, headed by Shaunak Khire and Trevor Mottl, that uses explainable AI (XAI).

Before joining Lazard, Khire co-founded an AI company. Jacobs hasn't talked about this much and so Euromoney offers him the opportunity and an even greater challenge: can he explain AI to us?

"What is AI? To me, it is a series of tools to improve your ability to predict an outcome," he tells Euromoney. "In financial markets, that outcome might be the performance of a stock or a basket over a set period relative to the S&P500 index."

Jacobs says: "We used to try to predict outcomes based on data in the public domain: cash flows, growth

rates, earnings multiples, historic trading ranges. We could only work on what was available to us. But now AI is taking huge amounts of data and analyzing relationships that we cannot otherwise see, teaching itself in the process, and presenting an outcome – a prediction that may lead to actions that are based on factors we may not have considered or can even see today."

Jacobs is not the only senior banker Euromoney has talked to recently who hears echoes in what is happening from the wave of market activity first unleashed by the advent of computing power to analyze financial data back in the 1980s.

Vikram Pandit, founder of Orogen and former chief executive of Citigroup, says that the advent of cloud computing today is a reminder of how the marginal cost of computing capacity and communications lurched down towards zero 20 years ago and led to the transformation of the cash equities business by algos and high-frequency trading.

But self-learning machines are new.

"There is a reasonable probability that many of the buy and sell decisions being made in financial markets today are based on factors that algorithmic traders identify but are unknown to the users of these capabilities," says Jacobs.

"Explainable AI is the next step of identifying the invisible factors that the black box uses to make decisions. Ultimately, the question will be – is the algorithm better alone or better with a human?"

This all comes as the bond markets continue adjusting to a post-crisis structure characterized by constrained dealer balance sheets relative to market size and customer needs, as well as a regulatory drive for more transparency that continually butts against the absence of a price tape and a central limit order book.

Concerted central bank buying eased that adjustment, but markets have struggled to cope with its unwind: witness the disorder evident at the end of 2018. And the

shrinking of the single-name credit default swap market since the crisis has made it harder for banks to hedge big risk positions taken on to accommodate customer flow.

As one of the biggest bond trading platforms, MarketAxess has seen this change from an inside seat. It allied with BlackRock's Aladdin trading network back in 2013 as it promoted so-called open trading or all-to-all trading, in which asset managers trade not just with dealers responding to their request for quotes (RFQs) but also with each other, sometimes even providing liquidity to dealers themselves instead of the other way around.

"Historically, investors were reliant on dealers for

"The next step is to automate that provision, which requires a systematic way to ensure that pricing is consistent within a market context," he says.

And that takes us to Composite+, MarketAxess's pricing algorithm for corporate bonds, which its people call CP+ for short.

This seeks to deliver a pre-trade reference for investors that will predict the clearing price for a bond by taking historic reported price data and combining it with real-time unstructured data – such as indications of interest being streamed by dealers and RFQ responses sent by liquidity providers via the MarketAxess trading platform. It also extracts any other shreds of price information relevant to a particular bond that may be embedded within other sources such as CDS index prices.

"There is a huge universe of bond securities," says David Krein, head of research at MarketAxess. "So even if the particular bond you would like to deal in may not have traded for several days – which means pricing engines bound by hard rules to only take account of observed prices will not be much help – indications of interest and the information content derived by machine learning from other instruments can provide an accurate pre-trade reference level."

Others are trying to do similar things, offering structured and unstructured data to the self-teaching machines, including Overbond. There are obvious challenges: the infrequency of trading in many seasoned credit bonds being just one.



"AI is taking huge amounts of data and analyzing relationships that we cannot otherwise see, teaching itself in the process and presenting an outcome – a prediction that may lead to actions that are based on factors we may not have considered or can even see today"

Ken Jacobs, Lazard

price discovery," says Gareth Coltman, head of product management for Europe and Asia at MarketAxess. "But with the availability of more data through Trace [Trade Reporting & Compliance Engine] in the US and Trax [the regulatory reporting and trade matching service] in Europe, investors can now do much more to independently price the liquidity that they themselves are increasingly providing to the market.

The fact is that there is no one price for any bond at any moment. Dealers will quote and transact at different levels for different sized lots in the same security – tighter for regular institutional size trades of \$2 million, wider for larger riskier trades above \$5 million and for smaller ones of \$150,000 that maybe look more trouble than they are worth.

They may deal at different prices for different clients depending on those clients' importance to the bank.

And investors themselves will have different price expectations given different priorities for any trade. They may accept some price slippage if their aim is just to get risk into or out of their portfolios as quickly as possible. Or they may instead take their time and optimize both spread – sometimes dealing at the mid-price between bid and offer, sometimes even letting orders rest and capturing spread themselves – and volume weighted average price.

"We have taken a very big-data approach with CP+," says Krein. "It looks at every data point in the market and uses it as a potential input to the price of each individual security. Obviously when trying to price a bond, it helps to look at the most recent trades in that exact same security; then at other bonds from the same issuer; then at similarly

rated issuers in the same sector. The further you go from the original security, the less relevant the information content becomes, but it is never zero. And even when the most proximate signals are missing, CP+ can learn to triangulate the less predictable features successfully.”

The CP+ algorithm provides a two-way price for over 24,000 bonds. It recalibrates the model for pricing each of these bonds every night. The prediction engine continually learns by incorporating new historic price data and then collating incremental data through the trading day from actual transactions, indications of interest and other sources.

It updates the two-way price for each security every 15 to 60 seconds: so it is not quite real time but as close as, especially for a bond that may not have traded for two weeks.

“The algorithm is self-learning,” says Krein. “We focus the calibration on institutional size dealer-to-client trades, but remember we put our predicted price out there pre-trade – to provide a reference before we can know which clients will transact in what size and with what priorities.”

He makes an astonishing sounding claim: “We continually compare our pricing predictions against every actual trade we then see come through Trace and Trax and find that the average difference is zero.”

Euromoney is aware that the “average difference” may hide some important variations, but still that sounds pretty impressive to us.

Not everyone, however, is convinced about the power of AI to derive information hidden inside the iTraxx Crossover to predict the likely next price for an infrequently traded, unusual maturity bond from a little followed and infrequent corporate issuer.

Euromoney calls in at Barclays to talk about the strong start to the year in debt capital markets, coming after the horrible secondary market selloff and near closure at the end of 2018.

“We are deriving historic pricing from years of central bank-supported markets with very constrained volatility,” says Lee Cumbes, managing director and head of public sector EMEA at Barclays. “I can tell you, as a practitioner, that in the fourth quarter of last year, predicting the price for your next deal off the last few could have led you to disaster. Investors who had previously been saying: ‘We need yield, we particularly like this name,’ were suddenly saying: ‘Look, we’re just done. We’re not buying anything.’”

That hasn’t persisted for long.

Marco Baldini, head of European and Japan bond syndicate at Barclays, says: “Fast forward to the start of this year and markets are open as never before and order books are quickly reaching record size.

“But it looks to me like issuance windows are going to snap open and shut in unpredictable ways, and that what we went through last November and December might



“Both among banks and asset managers we’re seeing a desire to be much more data driven”

Usman Khan, Algomi

return even though issuers, investors and bankers would all much rather the market wasn’t so stop-start. It’s a new primary market volatility, partly driven by lack of secondary liquidity, perhaps by automated stop-loss orders.”

Cumbes picks up the thread: “Even if the stops don’t last too long, managing across those episodes is helped by experienced people picking up suddenly changing tones from conversations with key decision makers. I’m not convinced that artificial intelligence can do that.”

It is a debate that will run for years, pitting the bankers promoting the art of the deal against the data scientists pitching big data and machine learning.

Krein at MarketAxess sees the value of CP+ very differently.

“If markets are very volatile and a client is seeking to price an esoteric, rarely traded instrument, that’s when it is extremely useful,” he says. “It is a super-powerful tool that analyzes historic data in the context of how the markets are behaving right now.”

It is important to remember that CP+ is not a trade execution algorithm. The advantage of a credible predicted price is that it provides a baseline against which to judge



“We’re not selling our technology as a standalone package. We seek to provide IT in combination with our market expertise and execution capability”

Vittorio Nuti, Deutsche Bank

responses to RFQs. That baseline can also be embedded into buy-side order management systems to monitor trade execution desks’ search for best execution. It gives investors a benchmark for review of responses to RFQs and enables subsequent automation of more trading.

“We can measure the average delta, or difference, between the CP+ price in typical \$2 million lots and the additional cost for executing larger blocks of say \$5 million to \$20 million for bonds that may be less liquid or that may be on special,” says Krein.

That helps investors quantify some of the trade-offs around dealing in size and speed. Investors need to think whether it may be better to do large trades with a single dealer at a wide spread or split them and work them, which may reduce friction cost but could have a price impact.

What is Krein seeing now in the credit markets?

“In the US investment grade market, the share of total market volume done in blocks seems to have come down slightly over the last 24 months, while the frequency of very large blocks is also lower than historical levels. And for anything done under block size, so below \$5 million, I don’t get the impression there is much risk of information leakage moving the price.”

Platforms like MarketAxess and banks like ING with Katana Lens are each seeking to provide technology tools

for price discovery to the buy side.

Are they going to come increasingly into competition?

It’s complicated.

“Banks, dealers, investors and alternative liquidity providers are all customers of CP+,” says Coltman. “Those with a sufficient technology capacity will seek to consume data from a variety of sources. Our unique selling point is that we have very large sets of data given our position as a leading trading platform and from Trax.”

One of the questions for banks as they seek to sell

technology to clients, including to smaller banks, is the importance customers attach to the independence of key technology providers.

You don’t really want to buy from a direct competitor. How confident should you be taking software from – or plugging into the risk management and price discovery systems of – a bank that wants to sell you something those software packages may claim is cheap and good value or buy from you what they say is expensive and over-valued?

The asset management industry provides some lessons. BlackRock developed its Aladdin trading network for its own trade execution desks first and then marketed it to other institutional investors that compete with it in the core business of managing money. BlackRock struggled to make much headway with these competitors until it folded its network into the MarketAxess all-to-all bond trading platform.

Similarly, Alliance Bernstein developed Alfa, a fixed income data aggregation and market surveillance tool, for its own dealing desk but then sold it in May 2017 to Algomi, the creator of Honeycomb.

Algomi Alfa provides cross-market information on liquidity and trade intent in order to give the buy side a real-time view of the whole bond market, including government bonds, investment grade, high yield, emerging market, municipal debt and structured credit.

Many of the world’s largest asset managers, particularly in the US, now use it, whereas some may have been reluctant to take it from Alliance Bernstein.

“Both among banks and asset managers we’re seeing a desire to be much more data driven,” says Usman Khan, co-founder and chief technology officer of Algomi. “Banks have spent the past five years digitizing their internal stores of unstructured data – not just transaction history but also prior indications of interest and RFQs received from investors that may not have led to trades – and the question now is how to make best use of that.

“Algomi Alfa helps the buy side assimilate data from multiple trading venues together with their own data, such as indicated prices from the banks, prices they have previously paid and how new prices compare. This

combination vastly enhances the data lake investors can access through APIs [application programme interfaces],” he says. “The precursor to deciding where to trade is having sufficient normalized data to identify the venue with the best price and also the best cost of trading.”

Algomi Alfa is not an execution platform. Banks have remained committed to market making in fixed income, turning over smaller dedicated balance sheets with much greater velocity rather than warehousing risk. The direction of travel is towards a more automated market – albeit never as automated as equity – with sell-side and buy-side algorithms that talk to each other and smart order routers that supply useful transaction cost analysis to asset managers’ customers and to regulators.

“Asset managers can learn from how their own execution desks’ algos have traded with various venues and start to distinguish liquidity and maybe put \$5 million of a \$20 million block through one venue, \$5 million through another and perhaps \$10 million through a third or a single dealer,” says Khan. “Buy-side algos can even automate negotiation of trades within certain tolerances.”

In the bond markets it is still not entirely clear how much of the over-the-counter market remains dark, with big banks crossing trades internally for different investors. But at the insistence of regulators, the market is lighting up, with at least post-trade reporting of blocks and, in Europe, even a requirement to disclose quotes received in response to an indication of interest.

“The top asset managers in the US love Algomi Alfa for the data lake and the market colour – and there is strong demand in Europe,” says Khan. “In Europe, we’d like to integrate with the trading venues so that you can click a button on Algomi Alfa that opens up a trade on one of the venues.”

“The banks that used to prioritize their single-dealer platforms are realizing that aggregators like Algomi are where they now need to provide their content and are good venues to meet the buy side.”

Fintechs of all types and size continue to challenge

the banks, although in secret, Euromoney suspects, the biggest banks are not too alarmed.

We talk to Samir Assaf, chief executive of global banking and markets at HSBC.

“Technology is just a tool to bridge issuers and investors,” he says. “Every five or 10 years there’s a cycle of innovation, which is a challenge for banks – and newcomers say the banks are finished. We saw it with the internet 20 years ago and we’re seeing it with fintech today. But I’m not worried about fintechs. We will deal with them as clients and partners.”

He is worried about something else – and it is not AI and machine learning.



“If markets are very volatile and a client is seeking to price an esoteric, rarely traded instrument, that’s when CP+ is extremely useful. It is a super-powerful tool that analyzes historic data in the context of how the markets are behaving right now”

David Krein, MarketAxess

“The difference today is the big established tech companies running in parallel to the big banks,” he says. “The whole world is a client of Amazon and Facebook. The banks’ clients are all on their platforms: that is a potential challenge we banks have not seen before.”

“But it is a challenge for the big tech companies as well. What do they want to do in banking? What might they be allowed to do? If they want to compete, they may face constraints. And what impact might it have on their stock prices. Amazon’s market cap is huge, more than five times that of HSBC, but its profit is a fifth of ours while its price-to-earnings ratio is a multiple of ours.”

So, perhaps, that takes us back for one last crack at

the question: can banks provide technology to customers? They are spending enough on it, after all.

Inside fixed income, currencies and commodities divisions, the path to bank-developed in-house trading technology being provided to customers has been trodden before, by FX dealers.

It is five years since Barclays, a technology pioneer through

its Barx system, developed Gator, an aggregator tool for its own FX dealers to test the ability to successfully execute trades in an increasingly fragmented FX market with 50 or more electronic communication networks all displaying prices, which often proved to be phantom liquidity.

Behind the advertised prices were often the same few

do large volumes immediately at zero spread, but only very rarely is that actually possible.

“We are developing apps that give a real-time quantitative view of the market, like volumes or impacts. We aim to exemplify through our pre-trade tools the trade-offs customers might consider on speed, price and spread. Those tools empower customers to decide when they should execute a trade and how they should execute it.”

Its position as a leading market maker in FX has given Deutsche the building blocks on which to base this technology.

“The quality of data is the crucial problem,” says Nuti. “Given our strong history in FX, we have created market-making engines to provide prices to our clients, enabling us to perfect our data quality and modelling, with which other banks and technology vendors may not be so well equipped.”

And the bank is not just a data warehouse, it is also a centre of IT development.

“Some of the tech we use in FX originated on the rates side,” says Nuti. “We have a central IT support staff that ensures consistency in tech development and enables re-use.”

And how does Deutsche get paid for providing technology first developed for internal users to clients of the bank? After all, this technology, such as its new FX market-colour app within Autobahn, takes a lot of development and quantitative analysis.

“Our return is in volume,” says Nuti, “while in the algo business we charge a fee. Clients need to have transparency on what they are paying for and how they are paying for it.”

So, what of the role of banks as pure technology providers?

“I think some banks may be tempted to try what Goldman did with Redi,” says Nuti.

This was the equity execution management system Goldman Sachs first acquired from Spear, Leeds & Kellogg in 2000 and then built up over the next 13 years into a leading single-dealer platform, across which hedge funds and other investors accessed its order execution algorithms.

But the tide turned against single-dealer platforms and Goldman first sounded out other banks to take stakes and then spun Redi off in 2013. Thomson Reuters bought it in 2017.

“We’re not selling our technology as a standalone package,” says Nuti. “We seek to provide IT in combination with our market expertise and execution capability.”

For now, banks’ provision of financial services and associated enabling technology remains bundled. But regulators are pushing them to unbundle, market structures are changing and the banks’ own search for new revenues in recompense for their hefty tech investments may yet take them in surprising directions.

EWI



“Managing across episodes is helped by experienced people picking up suddenly changing tones from conversations with key decision makers. I’m not convinced that artificial intelligence can do that”

Lee Cumbes, Barclays

liquidity providers and they could withdraw on last look. By constantly testing each venue with large and small trades Barclays became confident in its own traders’ ability to work orders across multiple venues, splitting them to achieve best price. It then, after considerable internal debate, chose to provide the same tool to its customers.

FX has developed closer to equities than bonds. Voice trading gave way to tech firms 20 years ago to build an array of complex market making engines.

“The FX market is now so quant-driven that it is often hard to decipher trends in a digestible format,” says Vittorio Nuti, head of algo and segregated execution at Deutsche Bank. “Every end customer wants to be able to