

## Lessons from recent banking litigation

2021 PRINDBRF 0123 • By Samuel Weglein, Chris Feige, Hadrien Vasdeboncoeur and Ilona Mostipan, Analysis Group Inc.

---

Practitioner Insights Commentaries • April 20, 2021

---

(April 20, 2021) - Samuel Weglein, Chris Feige, Hadrien Vasdeboncoeur and Ilona Mostipan of Analysis Group Inc. consider technology used to parse vast transcripts in recent antitrust lawsuits involving benchmark interest rates and the foreign exchange market.

### Introduction

Since 2011, a wide variety of criminal and civil investigations related to benchmark interest rates (LIBOR), foreign exchange (FX) trading, and other financial products have focused a spotlight on the use of chatrooms by market-making traders.<sup>1</sup>

In these chatrooms, traders would discuss securities, post quotes, negotiate trades, share information about the market, maintain professional relationships, and even socialize. But investigators and regulators have alleged that some of these chatrooms were also used to inappropriately influence global benchmark interest rates, foreign exchange rates, and prices in other financial markets.

Follow-on litigation has been brought against these banks and traders alleging anticompetitive conduct and market manipulation, such as price-fixing and spoofing.

Electronic communications between the traders in the chatrooms, as well as in emails and on phone calls, have been cited as key pieces of evidence by regulators and plaintiffs in these investigations and litigation matters to support claims of anticompetitive conduct or market manipulation.

In these actions, the transcripts of chatroom conversations, emails, and phone calls are typically produced from each day throughout the alleged conduct period, which can span years. Given the almost-constant nature of the communications, the produced materials can include millions of pages of conversation transcripts.

Given the volume of information typically produced, it would take months or years for a person or a group of people to read the transcripts, analyze and categorize their contents, and identify useful evidence. Fortunately, such manually intensive reviews can now be largely avoided by employing data science-based approaches, running on increasingly powerful computers, to

review and analyze the vast quantities of text.

Over the past few years, such customized programmatic solutions, broadly referred to as "natural language processing" (NLP) and "machine learning" (ML) methods, have been leveraged in a number of litigations involving securities traders' chat and email transcripts.

Recent developments in NLP and ML algorithms mean that increasingly powerful and accessible tools are available to process vast quantities of text at a speed and efficiency that is potentially orders of magnitude greater than is possible with teams of manual reviewers.

These developments not only make it feasible, but also cost-effective and accurate, to apply these tools without sacrificing the quality of the analysis or compromising the reliability of conclusions.

In addition to overcoming potential time or budgetary constraints, these NLP and ML tools open up entirely new analytical possibilities. The power of these tools lies in the ability to search for patterns — not only for simple defined patterns, but also for complex linguistic patterns that do not need to be explicitly defined by the reviewer.

For example, an ML model can identify conversations that bear similarities based on the context or patterns of communication, rather than similarities based only on exact key word matches. Applying the algorithms to traders' electronic communications allows for much more complex, as well as more efficient, scrutiny of the production record than is possible with manual review alone.

This enhanced capability has led to more accurate identification of useful and relevant evidence, and has helped uncover patterns of and insights into behaviors that would likely have been very difficult and costly, if not impossible, to glean otherwise.

In this article, we have adapted disguised examples from recent litigation to illustrate the utility of ML and NLP methodologies. Although the discussion in this article focuses on traders' electronic communications, such as chats and emails, these techniques have broad potential applications to a range of text documents and data.

These could include financial statements, equity analyst reports, call transcripts, customer service support chats, job applications, insurance claims, product reviews, social media posts, and any other case in which vast sets of text documentation may contain important information.

## Creating a database of content with natural language processing tools

A worthwhile first step when dealing with large amounts of text data is to create a database that organizes the data in a structured manner. Doing so is relatively straightforward — it does not require complex NLP and ML tools — and the result is that the text information is organized and

displayed in a searchable format.

Creating a database of the information can be as simple as combining thousands of documents into one large searchable database. We can also use basic NLP tools, including scraping and parsing tools, to extract specific relevant information such as time stamps, participant names, locations, and chat or email text, based on defined patterns in the documents. Once a program is trained to 'read' documents, transcripts of chats or other documents can be efficiently read in and organized in a database.

For example, a discussion among FX traders at competing banks might take the form of the chatroom conversation shown in Figure 1. In this exchange, Traders A, B, and C are FX traders dealing in the EUR/USD currency pair at different banks.

Trader A appears to ask traders at other banks about bid-ask spreads for a particular size of trade in the EUR/USD market. The content in this chat appears in a format such that the individual components, such as date, time, time zone, trader name, trader email, and chat could be extracted into a database using a simple NLP program.

In addition, the content could then be standardized (e.g., converting the time stamp from UTC to EST), and relevant information could be extracted (e.g., identifying the trader's bank based on the email address).

**Figure 1. Capturing text information from chats between traders<sup>2</sup> into a database**

Date	Time 1	UTC	Trader A name & email	posted: How wide are you guys showing in 100 eurUSD today?
Date	Time 2	UTC	Trader B name & email	posted: consensus is 10 I think
Date	Time 3	UTC	Trader C name & email	posted: at least 10
Date	Time 4	UTC	Trader A name & email	posted: Yeah I think it's at least 10 as well

  

File	Time (UTC)	Time (EST)	Author	Content
Bates	Date and time 1	Date and time 1	Trader A name	How wide are you guys showing in 100 eurUSD today?
Bates	Date and time 2	Date and time 2	Trader B name	consensus is 10 I think
Bates	Date and time 3	Date and time 3	Trader C name	at least 10
Bates	Date and time 4	Date and time 4	Trader A name	Yeah I think it's at least 10 as well

As there are often thousands of chat transcripts, the same NLP program could be used to extract the same information into the database. This process turns thousands of text-based chat transcripts into an organized and searchable summary.

Beyond simply extracting and organizing content into a convenient searchable format, these basic applications of NLP methods for extracting text content can be used to build quantitative datasets that would otherwise be very difficult to gather.

Consider as another example Figure 2, which shows two examples of email announcements regarding the details of different securities and their offer prices. The two email extracts in Figure 2 contain similar information, but in different formats.

In a case with hundreds or thousands of such emails, the time and cost of manually extracting data for quantitative analysis of the securities pricing would be potentially prohibitive.

With a more complex NLP program than was used to analyse the chat in Figure 1, each type of document format could be "read" and select information of interest — such as the details about the securities, and prices — could be extracted, standardized, and consolidated into a database.<sup>3</sup>

**Figure 2. Capturing security pricing information from emails into a database**

*Example 1*

Amount	Tckr	Cpn	Maturity	Call	Call Offer	Comment	Cusip
Amount 1	Issuer 1	0.3	02/06/15	05/13	B/Q	Price 1	Type 1 Cusip 1

*Example 2*

```

ISSUER: Issuer 2
STRUCTURE: Type 2 and Amount 2 MM DEAL
SETTLE: 03/26/2013
MATURITY: 03/26/2018
CALL: 03/26/2014 AND THERAFTER
CPN: 1.04%
OFFER: Price 2
CUSIP: { Cusip 2 Corp <go>}
    
```

File	Date Sent	CUSIP	Issuer	Type	Amount	Coupon	Maturity	Call	Offer
Bates 1	Date 1	Cusip 1	Issuer 1	Type 1	Amount 1	0.3%	02/06/2015	02/05/2013	Price 1
Bates 2	Date 2	Cusip 2	Issuer 2	Type 2	Amount 2	1.04%	03/26/2018	03/26/2014	Price 2

### Making sense of the data

Once the relevant information from the electronic communications has been extracted into a database, the contents can be searched and organized by date, time range, participant, recipient, keyword or phrase, or even by complex language pattern.

In the example in Figure 1, where Traders A, B, and C discuss market spreads in a chat, suppose that in addition to chatting with each other in a three-way chat, all three traders were also participating in one-on-one conversations with each other, and with other market participants.

To understand the context of the chat in Figure 1, it could be important to evaluate what these traders were saying contemporaneously in those other chats. Having the data organized in a database allows us to quickly identify other chats that these traders were in at the same point in time, and quickly create a summary of all relevant discussions to understand the full context of the various discussions that were occurring at the time.

While Traders A, B, and C were chatting, they may also have been actively trading — with each

other, and with other market participants. Just as it is important to understand the full context of all conversations occurring at the time of a particular chat, it is often important to assess how the at-issue discussions align with the trading records.

Identifying all relevant conversations and trades allows for analysis of how the traders reacted to certain discussions, and potentially, whether the traders' discussions affected trading behavior. A database of the chat text allows the multiple chat conversations to be seamlessly merged with the trading records to create a single stream of content for review and analysis.

A database further lends itself to standard queries and statistical analyses of the quantitative content found in the communications, such as counts, averages, and any other comparisons over time and between instances, individuals, securities, and so forth.

For example, if one wanted to query how frequently Traders A, B, and C chatted over a years-long conduct period, how often they mentioned 'spreads', or how frequently they mentioned a particular currency before executing a trade in that currency, the database would allow for each of these analyses to be performed quickly.

In the example from Figure 2 in which a program pulled specific information about securities from emails into a database, the data could then be used for further quantitative analysis of security prices. For example, in the context of price-fixing allegations, the data could be used to analyze whether market participants priced securities similarly when announcing newly launched securities.

In the example, the program also extracted information about the emails and securities into the database, so such an analysis of pricing could also take into account changes in pricing patterns over time and distinguish between securities based on their characteristics (e.g., maturity, coupon, vintage, issuer).

As such, databases of organized chat and email content allow us to efficiently address a wide array of queries and perform complex analyses, which are both qualitative and quantitative in nature.

## Applying machine learning tools to develop insights

The analyses described above, which are mainly focused on organizing and reviewing relevant information and performing quantitative analyses, can be especially valuable and can be accomplished using relatively straightforward analytical approaches.

But more complex analyses, which might include identifying potentially relevant or helpful chat discussions among thousands of chat transcripts, or identifying patterns across larger sets of chats, may require more advanced ML techniques.

Many of the investigations and litigation matters related to traders' chatroom conduct have been focused on a handful of short chat excerpts that regulators and plaintiffs have identified as potentially problematic. Typically, such regulators or plaintiffs will focus on a subset of chats as examples which imply a broader pattern of behavior throughout the period of alleged misconduct.

These may not be random or unbiased samples and a central challenge is to assess the prevalence of these types of excerpts across the broader set of chats, which often entails identifying other similar and dissimilar chats across the vast set of chat transcripts.

An unbiased sample, or sets of other similar or dissimilar chats, can be very difficult — and very time-consuming and expensive for the client — to find by searching the database and manually reviewing the chat language.

Consider again the allegedly problematic discussion among FX traders adapted from Deutsche Bank's consent order with the State of New York.<sup>4</sup> In this exchange, Trader A appears to ask the traders at other banks about the spreads they have been quoting when transacting \$100 million of EUR/USD.

The other chat participants respond by sharing their views on the spreads in the market. The government deemed the exchange was "[suggestive] of possible coordination."<sup>5</sup>

*Trader A: How wide are you guys showing in 100 eurUSD today?*

*Trader B: consensus is 10 I think*

*Trader C: at least 10*

*Trader A: Yeah I think it's at least 10 as well*

This is just one conversation, at one point in time. Unfortunately, there is no efficient and accurate way to search for conversations similar to the one above. There are no clear keywords with which to run a search; for example, the excerpt does not explicitly mention 'spreads', which is the key topic at issue.

Other potential keywords, like 'wide', 'showing,' and 'eurUSD' would likely result in many false positives due to extensive usage in unrelated and dissimilar discussions among the traders. The discussion also does not specifically relate to a time of day or a specific market event, which could be identified using traditional search methods.

Until recently, there was no option besides manually reviewing vast sets of chat transcripts to look for other similar discussions. Even in large-scale litigation matters in which the parties have significant resources and the discovery schedule allows for it, human review of large productions is still slow, and can be prone to human error.

Chat discussions typically include novel jargon, code words, and slang, and an effective review of the transcripts requires fluency in such language. A significant investment in a manual review may still result in missed context and understanding, and a re-review on such a large scale can be even less appealing to clients.

NLP and ML methods are not a panacea for these challenges. However, an algorithm-based parsing of large sets of text documentation, like the FX chats discussed above, can significantly improve the speed and accuracy of review, and can be refined iteratively as new information and a better understanding of the conversations is developed.

At their core, these methods use computational power to determine patterns in language use, identifying conversations that are similar. Discussions of interest may consist of sequences of common words or phrases that are individually anodyne, but collectively form an identifiable pattern. NLP and ML tools can search for such patterns; a simple Boolean search that relies on finding identified keywords will be unable to identify relevant patterns.

The application of NLP and ML tools relies on the identification of chats of interest, which are then used to 'train' a model to identify passages with similar patterns and content. For example, the FX chats, identified by regulators or plaintiffs (such as the spread discussion identified above), can be put into a training group of chats, and an ML algorithm can then use those training chats to learn a rule that can identify other chats with similar patterns and contents.

The algorithm 'learns' by breaking down the training chats into their salient components, i.e., those that are distinctive of the conversation relative to the wider chat room discussions. For instance, from the example above, the individual terms 'wide' and 'eurusd', in conjunction with other important words or word groupings, may be deemed by the algorithm to be a relevant pattern that can help identify other similar chats.

A major benefit of the approach is that the key identifying features of similar chats do not need be defined ahead of time; the algorithm determines those features as part of the training process and can then surface other chats in the production that are likely to encode similar meaning.

The output of the algorithm is typically a ranking of the chats in the population of chats by similarity or relevance to the training set. Ultimately, the chats that the algorithm identifies as most similar still require human review to evaluate the discussions and identify the usefulness of them in the case.

However, by ranking the chats, the algorithm focuses that manual review on the most similar and, therefore, most-likely-to-be-relevant chats. In addition, as new information is gathered or new chats are identified by regulators or plaintiffs, the training set can be quickly updated, and the algorithm re-run.

The computerized nature of the approach means updated iterations can be completed quickly,

rather than requiring a full re-review of thousands of chat transcripts.

For example, taking the spread chat discussed above, one might want to answer a relatively simple question, such as how frequently other traders, in other chat rooms, have conversations about spreads that are similar to the example spread chat.

Using examples of spread exchanges in that chatroom that are identified as the relevant chats in a 'training set', a model can be trained to quickly review thousands of chats from other traders and rank the other chats for similarity to the training set chats.

A review of the top-ranked chats would quickly start to provide an answer to the question of whether other traders were having similar discussions in their chat rooms, or whether this was a pattern of chat discussion specific to the at-issue chat room.

Classical ML algorithms like the one described above are extremely powerful and particularly useful in situations where extensive training examples are unavailable. The new frontier in ML includes 'deep learning' techniques that have proven to be attractive alternatives in situations where large training sets are available.

Classes of algorithms within the deep learning toolkit have the ability to process text inputs word by word while keeping memories of the order of words. Algorithms of this nature can more closely mimic biological intelligence processes and have proven to be very effective at identifying relevant or similar text in even more complicated situations.

When appropriate, the modeling framework we described earlier can be augmented to train more sophisticated models and has the potential to generate more accurate results or perform more complex analyses.

More generally, common applications of these tools include:

**Quantifying incidence of alleged conduct:** Although these analyses cannot provide definitive counts of instances of specific conduct within the universe of transcripts, ML and NLP tools can provide an idea of how prevalent certain types of exchange are in a broader set of chat communications. That information can help inform arguments and further analyses that could be done, such as the appeal of reviewing other transcripts in more detail to calculate precise incidences, or identify potentially helpful examples.

**Assess uniqueness of alleged conduct:** The algorithms can be also used to determine whether similar exchanges were prevalent among market participants who are not part of the case, as in the example above. These types of analyses necessitate access to a broad range of chats beyond the chatrooms directly at issue, and can be especially potent in evaluating whether a specific chatroom was unique or similar to others.

**Identify similar but distinct evidence:** One of the most useful applications of these tools is in

efficiently identifying examples similar to the instances of misconduct identified in the litigation or investigation, but where misconduct is clearly not present.

The identification of these types of instances allows the flagged instances to be further contextualized, such as substantiating whether the interpretation of a vague excerpt might be correct or better understood in another context. They also allow the flagged instances to be evaluated for representativeness, i.e., whether the exchanges are cherry-picked for instances in which they appear to support the pattern of misconduct alleged.

There are numerous other potential uses within the context of FX, in other financial cases, or in any case with large amounts of text-based evidence.

## Conclusion

The NLP and ML tools discussed are powerful and have many applications. Of course, they also have limitations. They cannot be used to establish precise counts of discussions or the *absence* of a discussion in the way that Boolean search can definitively identify a count of instances where a particular word or phrase was used.

In addition, these are relatively new tools and not yet commonly featured in the context of finance litigation; this may be because to those who are not familiar with these methods, they appear to be 'black boxes'.

Their results can be more difficult to grasp or to demonstrate by example; therefore, they currently may be more powerful in support of developing analyses or identifying useful complementary evidence, rather than as standalone pieces of evidence or expert testimony.

However, these are powerful tools, and over time, we believe that courts will become more comfortable accepting evidence and analyses based on tools like these. The substantial gains from the efficient and speedy processing and review of text-based documents allow the pursuit of challenging analyses within those large sets of documentation that otherwise could not be done.

We expect that these tools will become more accessible and powerful, and that they will have increasingly central roles in finance and antitrust litigation, as well as in any litigation which revolves around large sets of text-based production materials.

### Notes

1 A market-making trader is a trader, typically at a large bank, who offers to buy currency from clients, or sell currency to clients, at those clients' request. A market-making trader typically does not hold currency for long periods of time with a goal of profiting on the movement of the market.

2 Chat content adapted from Deutsche Bank's consent order with the State of New York. See "Deutsche Bank Consent Order Under [New York Banking Law](#)" §§39 and 44, *In the Matter of Deutsche Bank AG and Deutsche Bank AG New York Branch*, 20 June 2018, available at <https://on.ny.gov/3wWVFSU>.

3 These examples show use cases with relatively simple formats that are particularly suited for the application of NLP and ML methods. It is important to note that if the formats of chats and other types of text are more complex and less standardized, more work would be needed to create programs to extract relevant information from the text documents. For example, if chat documents are in multiple formats, such as from different communications platforms or produced by different parties, multiple corresponding programs may be necessary. This would result in greater up-front time and effort to convert all chats for inclusion in a database. Nevertheless, in cases that involve hundreds or thousands of individual communications, the time saved by programmatically 'reading' the documents can be substantial, as computer programs generally can be written to 'teach' the computer how to extract clean, organized data from text in nonstandard or heterogeneous formats.

4 Deutsche Bank Consent Order Under [New York Banking Law](#) §§39 and 44, *In the Matter of Deutsche Bank AG and Deutsche Bank AG New York Branch*, 20 June 2018, available at <https://on.ny.gov/3wWVFSU>.

5 *Ibid*, at p. 11

By Samuel Weglein, Chris Feige, Hadrien Vasdeboncoeur and Ilona Mostipan, Analysis Group Inc.

Samuel Weglein, managing principal with Analysis Group, is a Boston-based economist who testifies and supports testifying experts in complex antitrust and securities litigation, and in international arbitration. Dr. Weglein testified recently on behalf of several large banks in litigation involving municipal bond markets, and has testified on damages in a case in the shipping industry. He has also co-led teams on several financial benchmark antitrust matters. Chris Feige is a vice president Analysis Group's London office where he specializes in the areas of finance, securities, and financial systems. In cases involving alleged market manipulation in the foreign exchange and IBOR markets, he has analyzed trade data and evaluated alleged manipulation strategies for clients. He has also developed complex valuation models including discounted cash flow models, and has analyzed asset-backed securities, collateralized debt obligations, and other securitized products in support of expert testimony in a number of bankruptcy and damages matters. Hadrien Vasdeboncoeur is a London-based manager with Analysis Group and a consulting economist specializing in securities and finance, as well as in assessing injury in antitrust and competition matters. His finance experience includes interest rate derivatives markets, the foreign exchange (FX) market, and shareholder litigation. Ilona Mostipan is a manager at Analysis Group's San Francisco office. Dr. Mostipan is a consulting economist who has worked on finance cases involving forex and bond traders, structured finance, and stock borrowing. She is experienced in applying programmatic methods to analyze large amounts of text data. Analysis Group, Inc, San Francisco.



Samuel Weglein



Chris Feige



Hadrien Vasdeboncoeur



Ilona Mostipan

---

END OF DOCUMENT

© 2021 THOMSON REUTERS. NO CLAIM TO ORIGINAL U.S. GOVERNMENT WORKS.

---

## Trending **companies**

[Catalyst Capital Advisors LLC](#)

[NuCast LLC](#)

[Prime Office Products Pennsylvania Inc.](#)

[Regency Care](#)

[Heart and Vascular Institute of North Florida](#)

# Trending law firms

Paul Hastings LLP

[Contact us](#)

[Training and support](#)

[Terms of use](#)

[Full Desktop](#)

---

1-800-WESTLAW (1-800-937-8529)

---

Westlaw Today. © 2021 Thomson Reuters

[Accessibility](#)

[Privacy](#)

