

Fixed Income TCA - Déjà Vu All Over Again

Written by Steve Glass – Co-CEO, Abel Noser Holdings

Introduction

Over the past several years, regulators, asset owners, and some managers have increasingly begun focusing on fixed income trading costs. There are several key drivers of this growing attention to best execution and its impact on fund performance:

- a. New regulatory reform, particularly in Europe, imposing stricter standards for demonstrating the achievement of “Best Execution;”
- b. Increasing sophistication and accuracy of Evaluated Pricing models for illiquid securities (which continue to improve);
- c. The electronification of fixed income trading which has increased the quantity and quality of data captured during the trade life-cycle;
- d. Increased demand for post-trade evaluation tools as the number of trading protocols (and their sophistication) have increased; and
- e. Growing focus on multi-asset class solutions and the need for a consistent review process of trading outcomes.

This begs the question, “*When will the broader fixed income marketplace start conducting Trade Cost Analysis (“TCA”) in a meaningful fashion?*” Many industry participants, including us, feel the time has come. Unfortunately, although fixed income TCA has been available for several years, we continue to see hesitation and skepticism from many market participants.

We recognize of course that given the fundamental structure and inherent OTC nature of the fixed income markets, some portions of fixed income TCA will likely never enjoy the granularity and detail of equity TCA. However, that is not to say it has no value. Indeed, in many respects, the status of fixed income TCA today is already more advanced than equity TCA as recently as the 80s and early 90s.

As the industry’s oldest TCA firm (having begun measuring equity trading costs in the mid-70s), we recall how the equity markets similarly resisted trade cost measurement. Given the natural reluctance of many traders to be measured and evaluated by metrics they’re unsure of, we’re not surprised to find that we’re encountering a number of the same concerns we heard for decades in the equity markets (albeit with a fixed income twist).

In the words of Yankee-great Yogi Berra, it feels like “*Déjà vu all over again.*” Of course, as we all know, over the ensuing years, the equity industry has evolved to the point where today, equity TCA is ubiquitous - a fiduciary best practice, in which upwards of 90% of all practitioners utilize it in one form or another.¹ It’s now time for the fixed income industry to follow suit.

¹ Greenwich Associates 2018 Market Structure and Trading Technology Study

We recognize that change is never easy, and particularly in the early-years, practitioners have to get comfortable with a few imperfections associated with fixed income TCA. However, we strongly believe that the closer managers look into these matters, the more at ease they will become. In point of fact, many of the concerns and purported rationales we’ve been hearing (for holding off on fixed income TCA), while well-intentioned, are in fact built on misconceptions, and/or flawed logic.

To be sure, as with the early days of equity TCA, fixed income TCA has important limitations that all practitioners need to be cognizant of. However, in the vast majority of circumstance, those limitations can be managed, and value derived. As the saying goes, “Perfection should never stand in the way of progress.” To this end, we feel the broad market would benefit from airing some of the more common concerns we’ve heard about fixed income TCA, *and our responses as to why those views no longer have merit.*

This article is divided into two parts. The first is a review of five of the more common concerns we’ve been hearing and our responses. The article then provides a flavor of the type of insights achievable today with fixed income TCA, and some high-level trading statistics across various fixed income sectors.² The article concludes by encouraging fiduciaries to not shy away, but rather, embrace this new tool for enhancing the investment management of institutional assets³.

Five Common Concerns (and Our Responses)

Concern #1: “Is quality market data available for only a handful of fixed income sectors (such as credit and rates)? If so, we need more comprehensive coverage of the broader fixed income marketplace before engaging in TCA.”

Response: Happily, good pricing data is no longer limited to US Treasuries and Corporate bonds. Many pricing vendors now offer comprehensive and robust global coverage. Our firm’s fixed income TCA product for example, provides coverage of Corporate and Sovereign Bonds, US Mortgage Backed Securities (Agency), CLOs and CDOs, European Mortgage and ABS, Consumer ABS, Municipal Bonds, US Treasury (“UST”), Bank Loans, Money Markets, and Credit Default Swaps.

Importantly, this coverage is not limited to just a handful of the most frequently traded securities. For example, in 2018/2019 we evaluated 80+ fixed income portfolios on behalf of our clients. In aggregate, the managers running those portfolios executed almost 200,000 trades. Those trades consisted of UST, Investment Grade (“HG”), High Yield (“HY”), Structured Products, Agencies, Municipals, and Sovereign Debt. The breakdown across sectors was as follows:

UST	19% of the trades (57% of the trade volume)
HG	26% of the trades (9% of the trade volume)
HY	16% of the trades (4% of the trade volume)
Structured	8% of the trades (15% of the trade volume)
Agency	Less than 1% of the trades (1% of the trade volume)
Municipals	10% of the trades (1% of the trade volume)

² For purposes of this article we use the term, “sectors” to denote different types of fixed income instruments such as High Grade, High Yield, US Treasuries, Agencies etc. We recognize that within the fixed income community these categories are often referred to as “asset classes.” However, we wished to avoid potential confusion by others who view all of fixed income as a single asset class.

³ *Acknowledgments:* Edward Morgan, COO of Abel Noser, Steven Amodio, Senior Developer, and Ariel Manalo, Developer played key roles in quantifying the trading cost statistics cited herein. This article could not have been completed without those efforts.

Sovereign	12% of the trades (10% of the trade volume)
Other	8% of the trades (3% of the trade volume)

Of those 200,000 trades, our TCA reports identified and priced 99.6%. Admittedly, the quality of those evaluated prices varied depending on the amount of data points available to the pricing vendor. For securities that hadn't been traded in months and had few broker-dealers offering quotes, the evaluated price was less accurate and, in worst case scenarios, was little more than an educated guess. Conversely, securities that were actively traded and had numerous dealers making markets typically had very accurate evaluated prices.

To enable industry participants to take this variability into account, evaluated pricing vendors have begun providing "confidence ratings" alongside their evaluated prices. This analytical refinement enables practitioners to assess the level of uncertainty associated with the calculations derived from the evaluated prices (whether pertaining to market values, performance, TCA, etc.).

To this end, our TCA reports provide an analysis of not just each manager's aggregate trade activity, but also the subset of trades which enjoy a high degree of certainty in their evaluated prices.⁴ While the percentage of trades deemed to have very accurate evaluated prices will, by definition, be less than the total number of trades evaluated, it is still a surprisingly high percentage – more than enough to derive value-added insights. In the 2018/2019 data set for example, while 99.6% of the trades had evaluated prices of any quality, fully 83% of the trades enjoyed sufficient market activity to be deemed "High Confidence."

Concern #2: "Are evaluated prices accurate enough to base TCA reports on it?"

Response: Quite honestly, evaluated pricing is now the market standard for estimating security values on most over-the-counter fixed income products (because, there is no continuously available liquidity, consolidated tape, nor NBBO). There are three main pricing services: Bloomberg, ICE Data Services and Markit. Bloomberg uses their BVAL pricing service to power the Bloomberg-Barclays indices (i.e. the "AGG"). ICE (formerly IDC) provides end-of-day pricing for many mutual funds. Most recently, BlackRock adopted ICE benchmarks for four of its new iShares ETFs with over \$18.9 BN in AUM⁵. Markit is the second largest index provider to global fixed income ETF issuers with more than \$80 BN in AUM⁶.

Rather than being an untested "black box," evaluated pricing is the fundamental basis upon which the intra-day and end-of-day NAV pricing for pension funds, mutual funds, ETFs, and bond indices are based. *Accordingly, evaluated prices are used to calculate the entry and exit prices for investors moving in and out of these funds, and billions of dollars per day are traded based on them!*

Indeed, the improving quality of evaluated pricing is driving growth in the ETF market by allowing market makers to price baskets of hundreds of bonds with varying levels of liquidity almost instantaneously. This has created new sources of liquidity for fixed income investors through ETF baskets and portfolio trading. Perhaps even more impressive, several firms are using evaluated prices to algorithmically make markets in

⁴ For further information on this topic see our article, "[Solving the Challenge of Fixed Income Market Data.](#)"

⁵ Blackrock Signs Bond ETF Benchmark Deal with Ice Data Services Hayley McDowell - <https://www.thetradejournal.com/blackrock-signs-bond-etf-benchmark-deal-ice-data-services/>

⁶ Markit iBoxx Indices - <https://www.markit.com/Product/File?CMSID=1741142c5c7249eab254955fe42767ac>

both the ETFs as well as the underlying bonds.

In point of fact, many of the firms who worried that evaluated prices are too inaccurate for practical applications such as TCA, are actually already "using" evaluated pricing, perhaps without even knowing it. For example, these managers are implicitly using evaluated pricing when they benchmark their portfolios to a Bloomberg, ICE or Markit index (who use evaluated prices to set their end-of-day values). Similarly, these managers are likely using evaluated prices when calculating their investment performance (both absolute returns and alpha).

This of course begs the question, "If evaluated prices are good enough to value their funds, calculate returns, and implement client fundings/redemptions, why not TCA?"

Concern #3: "The prices in your TCA report are different than the price I saw in the market at the time of the trade"; and/or, "The prices in your TCA report are different than the prices displayed by my market data provider on the same day/at the same time."

Response: One challenge for fixed income pricing is that the market is fragmented with a large number of securities that trade infrequently on appointment. As a practical matter, each trader therefore has access to a pool of liquidity which is almost certainly slightly different from their peers - in some cases better, in other cases worse.

Similarly, while each evaluated pricing service seeks to capture a market-wide executable price, they too see slightly different sets of transactions from each other. Pricing vendors may also weight inputs to their pricing models (quotes, trades, modeled prices etc.) differently. This will further result in slightly different final evaluated prices, depending on a trade's specific characteristics.

Essentially, the difference between the price a trader sees on their screens, or from their market-data provider, simply reflects the liquidity available to *that* trader/provider. While in most instances one pricing vendor's price will be very close to another vendor's price, it is unrealistic, and more importantly, unnecessary, to expect those prices to mirror every other vendors' prices exactly. This is particularly true when evaluating many trades over several quarters. In those instances, repeated poor performance relative to any of the top pricing vendors becomes progressively harder to justify.

Concern #4: "Since I traded at the displayed bid or offer, why does the TCA report say my trading costs are negative? Aren't my costs zero?"

Response: The gist of this comment reflects a misunderstanding of TCA and the definition of trading costs. When measuring trading costs in the context of portfolio values and returns, trading costs represent the loss in the overall asset value of a position as a result of a manager entering or exiting an investment position. This is NOT the same as evaluating whether a trader "did a good job."

Thought of in this way, all (or almost all) trades cost something due to the bid-ask spread and this can provide valuable information to investors. For example, it costs more to trade many emerging market corporate bonds than investment grade bonds because the bid-ask spread is larger. That is a trading cost - albeit one imposed by the structure of the market rather than the trader.

As an aside, if a trader is simply interested in knowing whether they traded at a price worse than the displayed bid or ask (depending on whether they sold or bought a bond), then we recommend measuring those trades against the near side of the market at the point of execution - this is a reference price provided by a few pricing services (including our product).

Concern #5: “Your TCA report indicates that some of my trades had positive trading costs. How can trading costs be positive?”

Response: As with Concern #4, this type of comment misconstrues the concept of trading costs and meaningful TCA, particularly in the context of understanding the impact trading has on performance. As stated above, trading costs measure the impact on the clients’ portfolio value incurred as a result of a trader building or unwinding a position. In most cases, a trade measured against the mid-price immediately prior to the trade will result in a negative trading cost (due to the bid-ask spread). However, under certain circumstances, there may be instances where an executed price represents a positive trading gain. These include the following:

1. If the client wishes to measure a trade against a reference point other than the prior-mid, this introduces the risk of market movements into the trade cost calculation. That is a legitimate measurement and provides valuable insight into the impact of such “market risk” (and how well a manager/broker “managed” that risk).

Put another way, if a trader decides to wait until the end of the day before executing a trade, evaluating that transaction vs. the immediately preceding mid, effectively gives that trader a “free pass” on whether his/her decision to wait all day before executing was a good decision. Nor does it give the trader credit for making a good decision if markets moved in a favorable direction throughout the day. It is only by using an earlier reference point (e.g. the market open, arrival price etc.) that these insights can be discerned.

2. Given the over-the-counter nature of fixed income, one manager (or dealer) may have access to liquidity at a price not available to the wider market. These trades may therefore look better than the prevailing market price and reflect the advantage in utilizing that particular manager/dealer.
3. If a client is providing trades without timestamps, trades may be measured against the market open price, resulting in a cost estimate that again includes market movements. This is no different than equity TCA that uses custodial data. While certainly less granular than time-stamped analysis, this level of analysis is not without value – particularly when evaluating trends/patterns in large number of trades executed over longer time horizons.

As shown above, practitioners can take comfort that the fixed income marketplace has evolved sufficiently for fiduciaries to begin incorporating TCA into the investment process. This brings us to a second question of whether such analysis has any practical value. We believe the answer is yes.

Clearly, new regulatory mandates such as those required by MiFID II⁷ and PRIIPS⁸, require firms transacting in fixed income securities to measure trading costs and demonstrate best execution. However, aside from

⁷ Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU, OJ L 173/349.

⁸ Packaged Retail and Insurance-based Investment Products (PRIIPs) - Regulation (EU) No 1286/2014

compliance with these new regulatory obligations, we believe fixed income TCA has also reached a point where it can offer useful insights that substantively assist the investment process. The remaining portion of this article describes some of those practical applications engendered by fixed income TCA.

What Are We Seeing?

Similar to equity TCA, our firm maintains a database of fixed income transactions. As a firm dedicated to assisting our clients in better understanding the impact associated with trading costs, we systematically monitor and data-mine our trade universe. This ongoing exercise, in turn, enables us to help clients put their execution costs into context (i.e. given the trade characteristics and prevailing market conditions).

For example, most industry participants recognize that on-the-run Treasuries will be cheaper to trade than Emerging Market or High Yield bonds. But how much more expensive? More to the point, what are the expected costs to trade an odd lot, versus a round lot, versus a trade of extremely large size? Similarly, how different are the expected costs to trade two securities within the same sector, but with different characteristics (e.g. a short duration BBB security versus a longer duration AA security)? What types of trades are better suited for electronic trading versus “voice” transactions, and what are the potential cost savings/penalties with making the right/wrong choice?

This is where having a sense of what different types of securities can (or should) cost, provides useful insights. Depending on the circumstances, these types of insights can have applications for: buy-side trading desks evaluating themselves or their counterparties, compliance officers monitoring outliers, portfolio managers fine-tuning their portfolio construction views, asset owners evaluating their managers’ trading efficiency, ECNs demonstrating best execution, and broker-dealers making efficient markets.

Following below are four practical applications that provide a flavor of how TCA can generate value-added insights.

Application #1: Provide boundaries and expectations for reasonable costs - across sectors

Not surprisingly, on average, depending upon each security’s duration, rating, and size, the trading costs observed within different sectors varied tremendously.⁹ For example, trades in high grade securities incurred costs ranging from less than -1 bp to -50 bp, with a mean cost of approximately -10 bp. In contrast, the average trading costs for high yield securities ranged from around -1 bp to almost -70 bp, with a mean cost of -17 bp. Municipal bonds were similarly expensive to trade, ranging between -1 bp and over -100 bp, with a mean of -14 bp. Not surprisingly, US Treasuries were the cheapest sector to trade, exhibiting costs of less than -1 bp to about -28 bp, with a mean cost of less than -2 bp. However, non-US Sovereign securities are almost as cheap, ranging from less than -1 bp to upwards of -70 bp, and a mean cost of -11 bp (with a median cost of -2 bp, primarily driven by developed market European Sovereigns).

⁹ For purposes of the data discussed herein, trading costs are calculated pursuant to the following protocols: Trades are measured as the difference between execution price and mid-market evaluated price at the time of the trade, costs are calculated in bps of market value traded, and when timestamps are not available, costs are measured against the market opening price and/or mid of the bid/ask spread over the entire day.

Application #2: Understanding the expected costs to trade across different venues

With the explosion of electronic trading platforms and capabilities, a critical trading decision for every desk is whether to execute a transaction “electronically” or by “voice.” Neither option will be the optimal choice for every trade. Further, even when trading electronically, a buy-side trader must choose between an electronic trading platform (that facilitates client-to-client transactions) versus using an electronic market-maker.

Currently, depending on the sector, some studies have estimated anywhere from approximately 15% (high yield) to over 60% (mortgage backed securities) of fixed income trades are executed electronically, and it’s generally understood that when used in the right circumstances, electronic trading incurs lower trading costs. The key is to understand what types of trades lend themselves to electronic trading. Important considerations include trade size, urgency, trading costs, operational efficiencies, and minimizing information leakage.

It goes without saying, of course, that different trading venues are better suited for different types of trades. And one job of the buy-side trader is to recognize when (and under what circumstances) each respective venue should be used. Picking the wrong venue/counterparty for the wrong trade may result in greater costs. In this regard, TCA can offer guidance.

For example, by aggregating some of the larger electronic market makers (e.g. Citadel, Jane Street, Millennium, Sumridge etc) it’s possible to begin getting a sense of the average trading costs incurred with electronic broker-dealers. Similarly, by excluding those electronic market maker trades from the transactions we see being executed across the larger electronic trading venues (that facilitate open trading for both dealer-to-client and client-to-client trading) such as Bloomberg, LiquidNet, MarketAxess, TradeWeb etc., it’s possible to also see the average costs incurred on electronic trading venues by non-market maker dealers and/or buy-side counterparties. Viewing the remaining transactions that were not executed across electronic trading platforms further enables rough comparisons versus traditional “voice” executions.

In this regard, on high grade credit trades we currently see average trading costs of around -1 bp incurred on transactions executed with electronic market makers (often large UST transactions), -3 bp incurred on trades executed across electronic trading platforms with other counterparties, and -7 bp on trades executed via “voice.” Of course, these costs are aggregate numbers across multiple sectors, and individual trades in all three groups can often range upwards of -20 bp or more (particularly on transactions that do not lend themselves to that particular approach). Deeper analysis of these transactions (e.g. by sector, trade size, maturity etc.) can thereby assist practitioners in navigating these types of decisions.

Application #3: Understand the liquidity cost of trading on-the-run versus off-the-run UST

While on-the-run Treasuries can be traded at virtually no cost, the duration, size, and even the side of a UST trade (i.e. Buy/Sell) can all impact the costs incurred. In this regard, given the many reasons why managers traffic in USTs (e.g. taking a view on interest rate moves, engaging in relative value trades, buying or selling Treasuries for collateral purposes, and repo trades) and the large volumes often involved in those transactions, even a few extra basis points in cost can translate into meaningful cost penalties. Foreknowledge of such potential costs can thereby facilitate more informed decision-making.

For example, selling small amounts of even long-duration USTs can generally be executed at approximately the same cost (within 1 bp) as purchasing those securities. However, selling larger amounts of long-duration USTs can often cost upwards of -5 bp more than buying the same securities.

Application #4: The cost of duration

As alluded to above, at a macro-economic level, the search for yield often entails shortening, or lengthening, the aggregate duration of a portfolio. In making this investment decision, it may therefore be helpful to have a quantitative understanding of the cost “penalty” associated with achieving that goal.

It also goes without saying that the size of this cost penalty varies depending upon the sector involved. For example, Table 1 below shows the average trading cost incurred to trade different maturities across various sectors (while keeping other characteristics constant).

Table 1

Fixed Income Sector	Average Cost to Trade Shorter-Duration Securities	Average Cost to Trade Longer- Duration Securities
UST	Less than -1 bp	Up to -10 bp
Non-US Gov. Sovereigns	Less than -1 bp to -3 bp	Up to -21 bp
High Grade	-3 bp to -9 bp	Up to -26 bp
High Yield	-9 bp to -20 bp	Up to -25 bp
Municipal	-13 bp to -14 bp	Up to -31 bp

Conclusion

As with equities, the ultimate purpose of fixed income TCA should ideally be two-fold:

- Address fiduciary and compliance oversight responsibilities; and
- Better inform the investment process.

Particularly with respect to helping maximize investment performance, fixed income TCA can be utilized to help validate/challenge trading tactics and strategies, evaluate the execution quality of counterparties and venues, demonstrate best execution, and even potentially provide useful inputs into portfolio construction models.

Fixed income TCA, as with equity and foreign exchange TCA, is an ever-evolving science. Future iterations will inevitably expand into the pre-trade phase. However, for the reasons discussed in this article, we believe that, fixed income practitioners who care can now use TCA in its current state to immediately derive a competitive edge for their investment process.

"It's déjà vu all over again."

Yogi Berra describing Mickey Mantle and Roger Maris repeatedly hitting back-to-back home runs in the Yankees' seasons in the early 1960s

[LEGAL DISCLAIMER](#)