NORTHFIELD | modeling financial markets worldwide

Fixed Income Liquidity Management

Dan diBartolomeo
Webinar February 2021

Introduction

- We have recently received numerous inquiries on fixed income liquidity analysis.
- This presentation illustrates how liquidity and transaction costs can be explicitly estimated for fixed income securities
- We will provide a “liquidity policy” framework allowing liquidity considerations can be incorporated into portfolio risk assessment.
- We will also present a new econometric approach to the estimating the liquidity needs of mutual funds subject to investor withdrawals.
- The final part of the presentation will illustrate a “liquidity tilted” portfolio replication.

The Time Bomb Problem

- During the financial crisis years of 2007-2009, much of the declines and volatility experienced by global fixed income markets purportedly had to do with liquidity.

- Since then, regulators such as the US SEC and the various aspects of MIFID II (and ESMA) in Europe have begun to require that asset managers of open-end funds and ETFs carry out analyses of their liquidity risk. In addition, regulators desire trading practices that do not unfairly shift the cost burden of large liquidations to remaining investors from those investors withdrawing.

- Various industry approaches are being taken to analyze liquidity risk. Unfortunately, we find that in all but a few cases the analytical approaches are unsound. These flawed analyses give the impression that market liquidity to transact securities is far greater than actually it is.

The Global Financial Crisis and Beyond

- The impact of the liquidity effects of interest began with the "hedge fund meltdown" of August 2007 and the destabilization of money markets triggered by the failure of Lehman Brothers. The near-failure of numerous other financial institutions contributed further to the misery, to which central banks responded with unprecedented injections of massive liquidity into financial markets.
- Since then, the equities world has been subjected to lots of discussions on "crowding" of strategies and factors. Interest rates have gone to zero or even negative in many countries (and may remain so as an outcome of the COVID pandemic).
- The rapid growth of ETFs makes the current problem worse, as many ETFs are traded with high liquidity but without regard to the fact that many of the underlying securities may not be equally liquid. There have already been significant dislocations of this sort as was the case in December 2018.

ETFs and ETNs

- ETFs and related ETNs have grown exponentially in recent years. The routine operations of these securities imply liquidity as both an exchange traded assets and through the potential for the creation/destruction of units.
- There is also the structural possibility of *in-kind contributions and withdrawals of securities.* Most ETFs and ETNs are based on indices.
  - Some exotic ETFs (levered, inverse, VIX related) have daily rebalancing rules required by the prospectus. The rebalancing of the underlying portfolios is predictable to a material extent, putting a bound on “normal” trading. Some funds also use prearranged “in kind” flows of securities to reduce realization of taxable capital gains.
  - Debt funded market makers and hedge funds have been providing liquidity for rebalances in return for transaction spreads, *particularly those securities with low trading volume.*

The Overly Simplistic Answer

- Encouraged by regulators, many asset managers now analyze liquidity by considering how long it would take them to liquidate their positions, assuming a given rate of trade participation and current (i.e. usually normal) market conditions.
  - For example, how long would it take an asset manager to liquidate the positions in a $10 billion mutual fund, assuming their trades could be 20% of typical trading volume in the relevant securities. Such a measure is calculable only for securities where the trading volume is relatively consistent across time.
  - Typically, the result is a matter of a relatively small number of days, *which seems ok.*
  - What if our $10 billion fund was part of a $1 Trillion fund family? Why would all investors in a given fund want to liquidate, but investors in the same manager's other funds would not? If a crisis or scandal hit a firm, it is likely all funds would suffer withdrawals.

Fire in a Theater

- In a global market where an asset class might be valued at tens of trillions, how would a $10B fund persistently manage to be one fifth of the total transaction volume? Put another way, a maximum of five funds can possibly be 20% participants in one side of the trading volume, when many thousands exist.

- If we repeat the analysis at a participation rate comparable to the fund's participation in global AUM, the results would show that most funds would require months or years to liquidate.

- A much more likely scenario is that either an entire firm is impacted by a "run on the bank" or that a macro event (e.g. the August 2007 hedge fund crisis) impacts entire markets. Now the simple analysis is even more implausible. *If there is a fire in a crowded theatre, and everyone stampedes to the exits it is not possible for everyone to be the fastest runner.*

Liquidity Aware Risk Policies

- Several years ago we proposed that funds should calculate their risk metrics such as volatility, tracking error, or Value at Risk in a liquidity-adjusted fashion.
- As a portfolio gets larger and larger in value, it becomes more cumbersome to run in the same way that a large ship is not as agile as a small boat. The market impact of larger trades increases the cost of a large liquidation so the effective risk level of a portfolio is a function of portfolio size.
- All traditional portfolio theory is based on the assumption of *infinite liquidity* so risk is a function only of portfolio weights.
- *In the real world, trading costs are an increasing function of position size so both portfolio size and weights matter in risk calculations.* A good theoretical discussion of this point can be found in Acerbi and Scandolo (2008).

Basic Liquidity Aware Risk

- We can formulate a liquidity policy as: we have to be able to liquidate X% of the portfolio in N trading days. Given our models of cost and the size of our positions, we can estimate the cost of liquidation during normal or crisis conditions. To adjust our portfolio risk estimate for liquidity, we convert our portfolio volatility estimate to parametric Value at Risk for the length of time specified in our liquidity policy then add the expected cost of fulfilling the liquidation to VaR. Finally, we convert the adjusted VaR value back to the equivalent volatility (or other risk metrics as needed).
- The estimated portfolio volatility is 7.5% per year and we will consider a 3 standard deviation VaR (covers 99.8% of normal distribution). The percentage % Parametric VaR = 4.98 = $[7.5 * 3 * (10/252)^{.5}]$ and we will assume that our liquidation cost will average 1% of the portfolio value. This cost estimate is estimated under our expected distribution of market conditions, as well as the size of positions to be liquidated and allowable time interval. Therefore the % Parametric VaR with Cost = 5.98 = [4.98 + 1]. The "effective" portfolio volatility = 10.06 = $[5.98 * / 3 * (252/10)^{.5}]$.
- Conditional on our liquidity policy the effective risk of this portfolio has more than doubled, and the position level risk contributions of large illiquid positions will have increased even more so.

Econometric Estimates of Possible Redemptions.

- If we believe that markets are efficient (as at least passive managers must), then investors are rational actors. Let's assume we have a bond fund with a 5% annualized volatility.

- Such a fund would be the optimal level of aggressiveness for an investor whose consumption liabilities (e.g. present value of future pension payments) is 85% of the value of their assets, so "net worth" is 15%.

- The 85% value is derived from the Discretionary Wealth Hypothesis (Wilcox, *Journal of Portfolio Management*, 2003). It is conceptually similar to "Kelly betting" (Kelly, 1956).

- An intuitive explanation can be found in Estimating an Investor's Volatility/Return Tradeoff: The Answer is Always Six (northinfo.com).

Redemption Forecasting: Example

- Consider a case where a fund with a 5% annual volatility has a 6% decline in one month (about a 4 standard deviation event).
  - The portfolio value goes from 100% of the prior value to 94% of the prior value, and "net worth" goes from 15% to 9% (94-85).
  - To ensure maintenance of solvency, the efficient investors would shift *40% of their investment over the month* out of the fund and into a risk-free investment similar to "constant proportion portfolio insurance".
- This construct is probably a "worst case" scenario for single funds
  - Many investors are active and will perceive the decline in the value of portfolio assets as a "buying" opportunity.
  - Market conditions that generate a 6% loss in a bond fund might cause even bigger losses in equity funds, triggering money to flow into bond funds.
  - Larger investors are very aware of transaction costs. Since costs are not zero this acts as a friction reducing actual redemptions.
  - Cirri and Tufano (JoF, 1988) suggests investors underreact to poor returns
  - Effects could be larger across a fund family or fund peer group

Bond Trading Costs are Tractable

- For fixed income securities, trading is still largely opaque and decentralized through a network of securities dealers that are usually part of investment banks.

- Since the Global Financial Crisis, data transparency has been improving through data aggregation sites such as TRACE (USA) and ZEN(UK).
  - A daily summary of TRACE data is hosted by FINRA and Morningstar, TRACE Market Aggregate Statistics (morningstar.com)

- Some custody banks also provide aggregated market information based on the transactions of their clientele, A framework for analyzing Liquidity risk in negotiated markets for Fixed Income securities (northinfo.com)

- As such, alternative methods are required to explicitly create an ex-ante estimate of costs associated with a specific bond trade of a particular size to be executed within a particular time interval.

Northfield Transaction Cost Framework

- Northfield uses a functional form for predicting trading costs that captures the empirical feature that costs rise rapidly when trades get *too big per unit time* for a market to absorb.

$$M = A + BS + CS^{.5} + D * MAX[S-L,0]^2$$

M = percentage transaction cost

S = proposed number of units to trade

L = the number of units where routine liquidity is exhausted

A = agency costs (broker fees, normal bid-asked spread)

B, C, D = constants per security (to be estimated)

Trading Cost Assumptions

- B, C, D and L assume a known time horizon (e.g. one day) for trade completion and can be rescaled for different trading intervals. All inputs can be restated as percentages of expected trading volume (as is customary for equity trades).
  - When trades get larger than S > L, then the price impact of trading imbalances gets very large, very fast. As previously noted, traders avoid doing such large trades over short time intervals.
  - In the GFC period, costs got large because both S values went up (lots of people trying to sell) and L went down (market makers withdrew from activity due to margin calls, particularly for exotic fixed income securities derisively referred to as "toxic waste").
  - The first response from central banks was to flood markets with money to bring L values back up. The second response was to start buying securities so that larger values of S (selling trades) could be accommodated without creating imbalances.

Lighting the Fuse to the Bomb

- It should be noted that many liquidity analyses being carried out for regulatory reporting assume participation rates that would routinely be sufficient to put markets into imbalances such that $S > L$, triggering potentially very large price impacts.

- In the event of some kind of market wide imbalance, there is little doubt that $S > L$ and large price impacts would arise
  - Recalling the October 1987 crash in global equities

- The question is whether asset managers would make matters worse by trying to transact at impractical participation rates previously reported to regulators.

Investor Fairness

- Assuming large price impacts would arise from a liquidity event, there is a fairness issue between investors seeking to liquidate and investors choosing to remain.

- Once the explicit cost of liquidation are considered, asset managers might be incented to sell off more liquid assets to minimize costs, but leaving the remaining investors with less liquid securities that could be even worse in the event of a second liquidity event.
  - The incentive comes from wishing to preserve AUM fees and minimize the perceived impact on investment performance.

Approaches to Fairness

- One sensible approach is to force the post liquidation portfolio to have a maximum amount of tracking error relative to whatever the portfolio was before the sell-off.
  - As long as this pre-set limit is not breached, managers would have the flexibility to select which assets to sell to minimize costs and would have clear reasoning for spreading trades over multiple days.

- Some regulators have introduced the concept of “swing pricing” where security prices are adjusted for the impact of expected withdrawals, but such processes are guesses at best.

An Illustration of Estimating Bond Liquidity

- At least for sovereign bonds in major currencies, the government bond markets are extremely deep with very transparent bid-asked spreads.
- That liquidity is further supplemented by exchange traded derivative markets offering futures and options.
- When we observe material illiquidity in bond markets it arises in issues where there is high uncertainty in the future cash flow stream such as high yield debt (i.e. "junk" bonds) or structured products where there is significant risk of unfavorable prepayments. See Lo, Getmansky and Makarov (2003).
- As described in The Volatility of Financial Assets Behaving Badly (northinfo.com), we will observe that all illiquid portfolios will have positive serial correlation in their reported returns. For some extremely illiquid assets (e.g. large commercial real estate) the autocorrelation in reported returns can be as high as 90%. A less extreme but statistically significant degree of autocorrelation is observed in bond yield spreads.

The No Arbitrage Argument

- Once we have estimated the degree of autocorrelation in spreads, we can use a simple "no arbitrage" argument to estimate trading costs. The presence of autocorrelation means that investors have a material degree of foreknowledge of *what will happen to security prices tomorrow*.

- In the absence of transaction costs, this foreknowledge would allow for the proverbial "free lunch" in terms of easily obtained trading profits. In a competitive market such easy profits should not be present. If we can estimate the autocorrelation and volatility of yield spreads (provided by our risk models) we can infer the trading costs that must be present to eliminate arbitrage from the credit spread data on any cohort of bonds.

- It should be noted that while this example uses index level data, our internal analysis can be more granular to the individual bond (as illustrated in the subsequent discussion):

US$ Example

- We will first consider the time series of the credit yield spread for the ICE/Bank of America US$ AAA bond index.
- Using the last five years of daily trading data, we observe that the average daily change in the spread is approximately zero, with a standard deviation of 2.6 basis points.
- The first order autocorrelation of the series is .42 leading to an arbitrage profit of 1.1 basis points per day per unit of effective bond duration in the absence of transaction costs.
- We confirm this positive serial correlation by noting if the series was random, we would expect to see an equal number of cases where a change in the spread was followed the next day by a change that was the same or opposite in sign.
  - For this series, adjacent changes of the same sign occurred in 285 instances, while adjacent changes of the opposite sign occurred 185 times.

Euro Denominated Example

- Applying the same analysis to the ICE bank of America Euro High Yield index leads to a quite different result. While the average daily change in the spread is approximately zero, the volatility in the spread is much higher at 8.2 basis points.

- The first order autocorrelation is .34 leading to a plausible arbitrage profit of 2.8 basis points per day per unit of effective bond duration in the absence of transaction costs.

- Again an examination of "runs" shows repetition of same sign changes 641 events compared to 372 events in which the sign of changes was different from one day to the next.
  - The T statistics on both serial correlation estimates indicate an extremely high level of statistical significance.

A Couple Refinements

- To do a better job in establishing ex-ante estimates of transaction costs, we can make several improvements.
  - Adjust the standard deviation realized for the observed autocorrelation. For the ICE/BAML high yield example above, the effective daily volatility increases by 43%, leading to a historic spread volatility value of around 11.8 basis points per day.
  - Alternatively, we could adjust the historic daily volatility to account for skew and kurtosis in the historical distribution of spread changes using the method of Cornish and Fisher (1938). This procedure requires additional parameters: whether the investor is trying to buy or sell (leading the counterparty to be potentially long or short credit volatility) and the size of the position relative to typical trading volumes.

Further Enhancements

- We can use ex-ante estimates of credit spread volatility rather than historic values. Such risk estimates at the index, cohort, or individual bond level are available in the "near horizon" version of the Northfield risk models. In these systems, credit risk aspects of fixed income securities are updated daily via automated analysis of financial news text[1].
- We can add data from a Northfield methodology called "Optimal Deal Flow" from Belev and Gold (*Real Estate Finance*, 2016) which was developed to address portfolio management of illiquid asset classes.
  – The trading of most fixed income securities also exhibits this episodic nature (maturities, new issuance, and infrequent trades). Using a modified version of the trading desk concept of "risk adjusted return on capital" (RAROC), we can estimate the impact of trades on bid/offers for transacting an illiquid asset of known position size (i.e. trades are assumed to be "all" or "nothing") as described in Optimal Deal Flow for Illiquid Assets (northinfo.com).

Relative Illiquidity: The J Score

- Another approach is to consider the relative illiquidity between portfolios without estimating transaction costs explicitly.
- Replicating this index as a passive portfolio is impractical for large investors because of illiquidity for many of the constituent bonds issues.
  - To address this task we created a proprietary metric for relative liquidity of various specific bonds which we will call the "J" score. The process of estimating J first classifies all bonds into groups by sector, and rating level. If the mandated index was of a global nature, the groups would be further delineated by geographic region of issue (e.g. Asia).
  - From our internal data we calculate the "option adjusted spread" (OAS) for each bond for each day on which that bond was known to have traded (i.e. we know the price is reliable). Higher average OAS spreads within a cohort are associated with less liquid bonds, as investors demand more yield in compensation for less liquidity. A detailed rationale for this hypothesis is presented in a Northfield research paper[2].

J Score Framework

- For each trading day, we transform the bond-level OAS values *within each cohort* into a unit normal distribution (i.e. mean zero, standard deviation one). The distributions of OAS levels comparable across time.
  - We now have a time series of OAS Z-scores (on irregular dates) for each bond and can calculate the time series standard deviation and standard error of mean. If the dispersion measure of a standardized OAS series is high, the pricing of the bond relative to similar bonds is erratic, which is typically caused by larger imbalances between supply and demand.
  - The dispersion measures will reflect the number of observations from which the mean was calculated. If there are relatively few trades the dispersion measures will be higher, reflecting poor sampling.
  - We can think of our illiquidity measure J as having two components. The first is the mean level of OAS and the second being the time series dispersion of bond level standardized OAS. Both measures will have higher magnitudes for illiquid bonds and lower magnitudes for more liquid securities.

Index Replication Experiment: Liquidity Tilting

- We constructed the following experiment in January of 2020 with the index constituents from the index provider. Out of the index list of approximately 600 bonds, we eliminated a few dozen where there was some data discrepancy between Northfield and the index provider.

- We broke the remaining bonds into three mutually exclusive subsets from which to form portfolios. The rationale behind this idea *is to create multiple replicating portfolios so that at least a few investors could all carry out the replication without creating competing demand for liquidity.*

- Various portfolio replication tests have been run to illustrate the ability to replicate the index. The optimization objective function is to minimize tracking error while holding the prescribed number of bonds with the lowest value of J (most liquid).

Summary Test Results: January 2020

Tracking Error		Weighting Scheme			
	N-Bond	PCT	EQWT	OPT	
Port 1	180	0.35	0.24	0.12	
Port 2	180	0.21	0.20	0.18	
Port 3	180	0.22	0.21	0.19	
Best J 180 bonds	180			0.10	
Best J 100 bonds	100			0.19	
Best J 40 bonds	40			0.23	

J Score Test Results Discussion

- The estimated absolute volatility of the portfolios and index was then around 7.6%.
- The expectation of the explanatory power (R^2) of the replicating portfolios is on the order 99.97%
- Assuming tracking error of .15% per annum, the expectation of variance drain (difference in annual arithmetic mean and geometric mean return) is approximately .00113% annually.   This is obviously small enough to be more than compensated by any measurable increase in liquidity or reduction in turnover.
- Three sets of roughly 180 bonds each were chosen at random.   The PCT portfolios are weighted by issue value, the EQWT are equally weighed and the OPT represent minimum tracking error optimal solutions (using our Fixed Income risk model as of the relevant January date).

Understanding J Score Output

- The analyses were reviewed as to the source of tracking error (interest rate risks, credit spread volatility, individual issuer idiosyncratic risks).
- Nearly all of the tracking error comes from issuer specific idiosyncratic risks. This is easy to understand by conceptually decomposing a high yield bond into two components, a riskless bond (i.e. a US Treasury) and the portion of "loss given default" (LGD).
  - The LGD portion is comparable to equity in the issuer (the value of both go to approximately zero in a bankruptcy), so the portfolio of LGD "slices" is comparable to an equity portfolio for estimation of credit risk. As bonds are very highly correlated there are an infinite number of combinations of positions that can virtually eliminate factor risk while idiosyncratic risk is minimized only when the portfolio holdings and benchmark holdings are most similar. Our review also included examination of the portfolios in terms of their similarity to the index on weight by cohort (sectors and ratings groups).
  - We also tested for various levels of numbers of portfolio securities

Conclusions

- The issue of liquidity in fixed income funds subject to investor redemption has been a material concern to regulators. Some regulators have introduced liquidity reporting schemes that we believe may identify particular funds that are currently illiquid, but may worsen the impact of illiquidity during periods of broad market distress.
- We have illustrated methods for incorporating liquidity considerations into the risk metrics for policy compliance policies.
- We have provided a new econometric method for setting the upper bound on the expected level of investor redemptions conditional on particular market events for a fund, or aggregation of many funds.
- We have illustrated our framework for estimating explicit transaction costs and provided an example of a “liquidity tilted” portfolio replication strategy.

Internal References

- diBartolomeo, Dan. “Credit Risk Assessment of Corporate Debt Using Sentiment and News”, in G. Mitra and X. Yu Editors, “Handbook of Sentiment Analysis in Finance”, Albury Books, 2016.

- diBartolomeo, Dan. “Reconciliation of Default Risk and Spread Risk”, Northfield Working Paper, 2020.