SALT FINANCIAL

RESEARCH NOTE

Risk Before Return:

Targeting Volatility with Higher Frequency Data

- **Volatility is generally easier to predict than future returns.**
- **Dynamically targeting volatility can lead to higher risk-adjusted returns vs. buy and hold.**
- **VIX can be a useful indicator but may not be ideal for targeting risk in some applications.**
- **Higher frequency data has the potential to boost performance in volatility targeting strategies.**
- **Volatility targeting compares favorably to trend-following strategies in reducing risk.**

Volatility is generally something investors seek to avoid. Most investors know increases in volatility are most often accompanied by declines in their account balances.

From 1990 through April 2020, the correlation between S&P 500 daily returns and changes in the Cboe Volatility Index (VIX) has been strongly negative—about -0.70.[1]

While volatility is undesirable, assuming some risk is required to earn a return. The name of the game is maximizing the amount of return captured per unit of risk, commonly measured by the Sharpe ratio—the average excess return (above a risk-free rate like Treasury bills) divided by its standard deviation (volatility).

Figure 1: Cboe Volatility Index (VIX), Daily

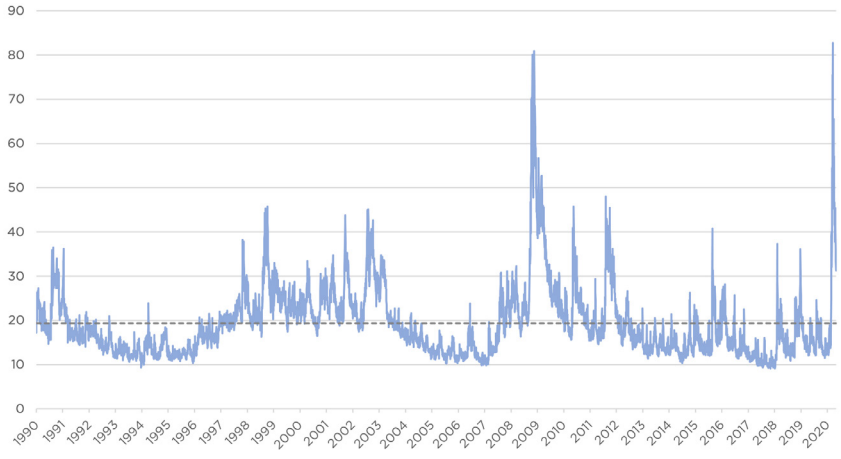

Source: Bloomberg, January 1990 – March 2020

[1] *Bloomberg data, January 1990 – March 2020*

Volatility itself has some unusual characteristics. Unlike returns that can compound for years, sharp increases in volatility eventually revert to the mean. While extreme events may cause the VIX to spike in the short-term, it tends not to stay there for long.

While predicting future returns is extremely difficult, volatility tends to "cluster" and shows more persistence, making forecasting risk a less formidable task. In statistical terms, volatility exhibits a high degree of autocorrelation–the relationship between a time series and lagged version of itself. In plain language, high volatility tends to beget more volatility and vice versa. We can use this relationship to construct a portfolio that targets return per unit of risk more effectively than a buy and hold investment.

Basics of Volatility Targeting

Volatility targeting is a relatively simple concept. The historical average volatility of the S&P 500 is about 18% per year.[2] An investor looking to target a level of 10% could dynamically adjust a stock portfolio by allocating a portion to a risk-free asset like cash to reduce volatility to the targeted level.

$$Ratio\ of\ Stock\ to\ Cash = \frac{Volatility\ Target}{Volatility\ Forecast}$$

For example, the strategy could use the historical daily volatility of the S&P 500 over the past month as a simple forecast of future volatility. If historical volatility is 20% and the target is 10%, the investor would allocate 50% of the portfolio to the S&P 500 and 50% to cash, adjusting the ratio daily as volatility changes to help match the target. The adjustment is implemented daily on a 1-day lag, allowing the portfolio to be re-allocated at the close on the day following the updated forecast.[3]

While not likely to hit the target exactly, even relatively naïve forecasts are generally able to keep portfolio volatility within a relatively tight band using this mechanism. If basic measures can work, more sophisticated forecasts of future volatility should power more accurate targeting—and better risk-adjusted returns. In this paper, we demonstrate the effectiveness of volatility targeting as a longer-term allocation strategy and incorporate higher frequency data to improve forecasting ability and risk-adjusted returns even further under a similar framework.

Comparing Results

In this analysis, we build hypothetical portfolios using the basic targeting mechanism described above with various forecasting techniques and measure their ability to impact returns. We begin with a simple buy and hold investment in the S&P 500 with some summary statistics on risk and return to use as a benchmark. The sample period is December 2003 through March 2020 for all hypothetical portfolios.

A simple buy-and-hold strategy from the end of 2003 produced solid returns but was certainly not without its rough patches. Volatility over the period was 19.4% and marked with some

Figure 2: S&P 500 Performance – Buy and Hold

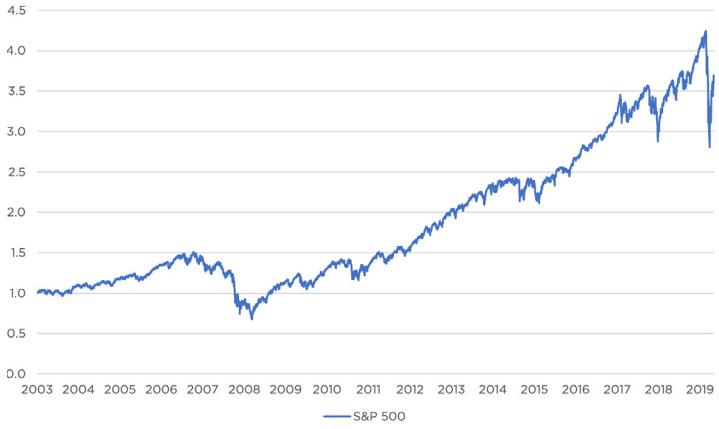

Strategy	CAGR	Avg Return	Volatility	Sharpe	Max DD
Buy and hold	8.3%	8.1%	19.4%	0.42	-55.3%

All average returns are annualized and in excess of 3-month LIBOR

[2] Bloomberg data, January 1990 – March 2020.

[3] This schedule assumes the data used as of the close on T-2 (Monday) generates a forecast that can be used to effect the changes at the close of trading on T-1 (Tuesday) to realize the returns at the new ratio on T-0 (Wednesday).

SALT FINANCIAL

Figure 3: Volatility Targeting Using 30D HV

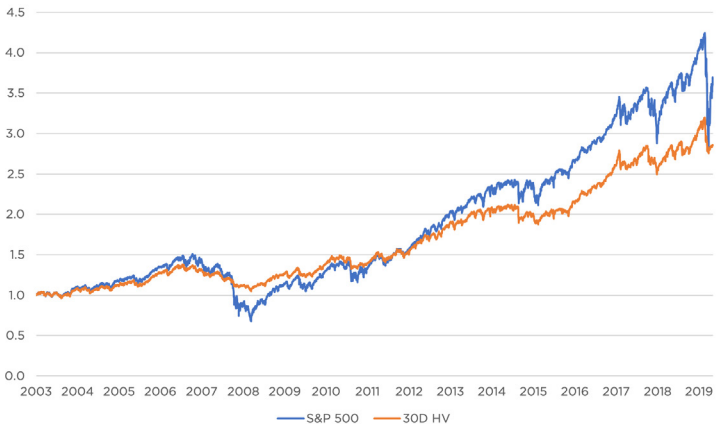

Strategy	CAGR	Avg Return	Volatility	Sharpe	Max DD
Buy and hold	8.3%	8.1%	19.4%	0.42	-55.3%
Vol targeting, 30D HV	6.6%	5.2%	10.3%	0.51	-33.9%

How effectively did the targeting mechanism match the desired level of risk? A look at rolling 1-year volatility confirms even simple targeting can consistently keep it within a tight band.[4]

The targeting strategy kept trailing volatility from ballooning during very tough periods for the overall market, never exceeding 13% on a temporary basis and hugging close to the 10% level most of the time.

Figure 5: Volatility Targeting Using VIX

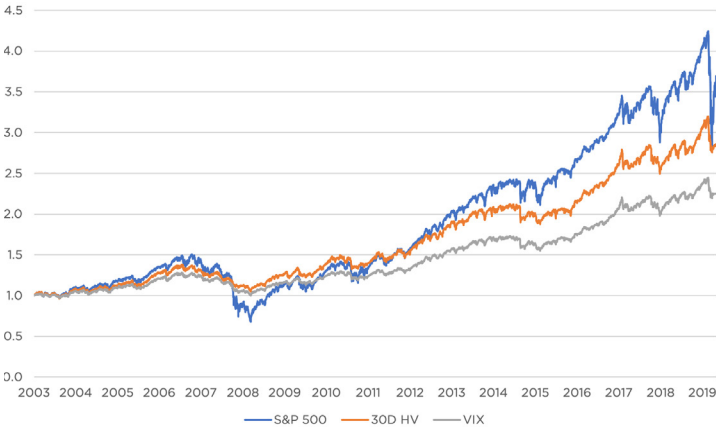

Strategy	CAGR	Avg Return	Volatility	Sharpe	Max DD
Buy and hold	8.3%	8.1%	19.4%	0.42	-55.3%
Vol targeting, 30D HV	6.6%	5.2%	10.3%	0.51	-33.9%
Vol targeting, VIX	5.1%	3.6%	8.1%	0.44	-26.0%

stomach-turning declines, including the Global Financial Crisis in 2008 and the impact of COVID-19 in 2020.

Targeting volatility at 10% using 30-day trailing historical volatility (30D HV) as the forecast with dynamic daily rebalancing can significantly reduce volatility and drawdowns.

Volatility comes close to the targeted 10% over the period with shallower drawdowns (note the dramatic outperformance in 2008). The trade-off is lower absolute returns but the strategy generated a higher Sharpe ratio, implying better risk-adjusted performance.

Figure 4: Rolling Volatility, 1-Yr Trailing, B&H & 30D HV

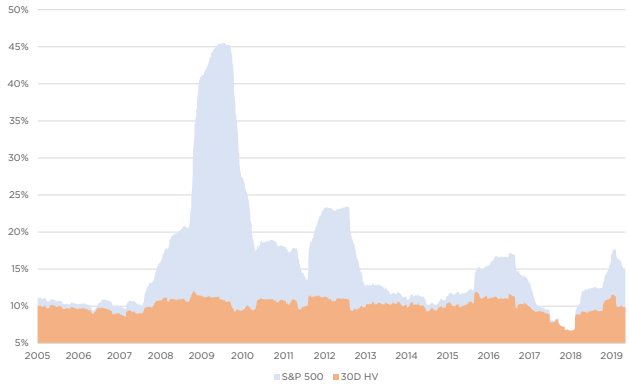

If simply using the recent past to forecast the future works well, what about using more sophisticated metrics that may be more forward-looking? For example, the VIX measures the volatility implied from prices on a series of S&P 500 put and call options, reflecting the market's expectations for volatility over the next 30 days. In theory, the recent past may not reflect the collective wisdom of market forces determining views of future volatility, making a metric like VIX a more appropriate forecast.

To the left we show the results of substituting VIX for 30D HV to target our 10% volatility portfolio.

[4] Volatility dipped well below the 10% target during a very tranquil 2017 when S&P 500 volatility dipped into the mid-single digits. While this hypothetical portfolio does not use any leverage, some volatility targeting strategies allow exposure above 100% to lever returns while volatility remains below the target.

SALT FINANCIAL

Like with 30D HV, targeting with VIX resulted in lower volatility and less severe drawdowns at the expense of lower absolute returns. However, using VIX underperformed on a couple of key measures. For one, compound returns were some 150 basis points per year lower than 30D HV. The max drawdown using VIX was lower, but the overall performance on a risk-adjusted basis was worse with a 14% lower Sharpe ratio than 30D HV.

More importantly, using VIX significantly overestimated volatility, leading the strategy to undershoot the target by almost 200 basis points. This resulted in the portfolio taking too little risk to earn enough return given the investor's preference, lagging the S&P 500 by 320 basis points per annum with only a modestly higher Sharpe ratio.

Although the "Fear Gauge" is forward-looking, implied volatility measures like VIX incorporate what is known as the volatility risk premium (VRP)—the compensation that buyers of options must pay to the sellers.[5] Selling options is like underwriting insurance. The seller collects a premium to provide the buyer with protection from adverse events (market moves). The price of the option typically implies a level of volatility that exceeds recent historical volatility, compensating the seller for assuming this risk. Without the premium, there would be little incentive for sellers to provide the insurance at all.

If recent historical volatility matches the 10% target, it is likely VIX will be priced higher (12% for example), which will result in a more conservative mix of stock to cash in targeting, underutilizing its volatility "budget" for the strategy. VIX is also a measure of the market's volatility expectation for the next 30 days, which may create a mismatch in timeframes for a strategy using dynamic daily targeting.

Adding High Frequency Data

Since volatility tends to be persistent, using recent historical data (HV) has several benefits in using for targeting. Unlike VIX or other measures driven from pricing models of the future, HV represents empirical observed results in the market. And since the VRP is not present, it does a respectable job in targeting with daily rebalancing. But it can be improved.

The wider availability of intraday market data over the last couple of decades has led to an expansion of research into using higher frequency data to drive more accurate forecasts of volatility. Known as realized volatility (RV), this measure uses intraday returns to easily compute an estimate of volatility for a single day, which can then be aggregated with observations from other days to estimate future volatility. Academic research demonstrates using intraday data can offer superior prediction accuracy in comparison to traditional estimates using only daily data.[6]

Intuitively, having more data to use for estimating a future outcome makes sense. If asked to forecast the temperature tomorrow at 3pm, hourly observations over the last three days would be a better tool to use than a snapshot once per day over the last three months. More frequent and more recent observations provide a more responsive measure to help refine the estimate.

Using higher frequency data involves some additional transformations to make the intraday values comparable to annualized volatility estimates. We begin with 15-minute returns from the very liquid SPDR S&P 500 ETF (SPY) as a proxy for the S&P 500.[7] For each trading day, we then sum the squares of these 15-minute returns (which is technically the daily variance), multiply by 252 to annualize, and then take the square root (to convert to standard deviation). Then we scale this value up by 130% to account for overnight (close to open) returns that are not included in the estimate.[8]

[5] This would also extend to other volatility instruments such as VIX futures, which normally trade in "contango" – successively higher prices for longer expiration contracts.

[6] Andersen, T., Bollerslev, T., Diebold, F., Labys, P., "Modeling and Forecasting Realized Volatility", Econometrica, 71, 529-626 (2003).

[7] Varying intervals from 5 to 30 minutes returns are used in most of the academic literature on realized volatility. Intervals smaller than 5 minutes (i.e. 1 minute or tick-by-tick) tend to introduce distortions from market microstructure noise. We find 15-minute intervals to be an appropriate length–materially better than 30 minutes but not much different than 5 minutes which involves triple the amount of data to produce.

[8] There are multiple ways to account for the impact of the overnight return on RV-based forecasts. Some methods add the squared close to open return to the daily calculation. Others dynamically scale the ratio of intraday to daily volatility to adjust towards the daily equivalents. In our research, we find the simplest method of scaling the daily value by a constant to be most effective for targeting using RV.

Figure 6: Volatility Targeting Using 2D RV

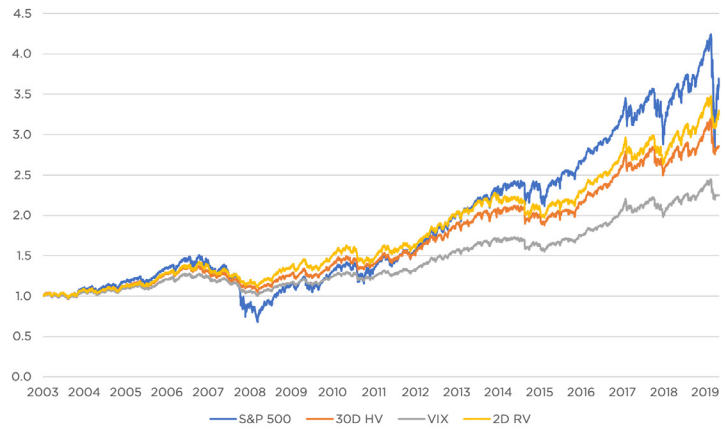

Strategy	CAGR	Avg Return	Volatility	Sharpe	Max DD
Buy and hold	8.3%	8.1%	19.4%	0.42	-55.3%
Vol targeting, 30D HV	6.6%	5.2%	10.3%	0.51	-33.9%
Vol targeting, VIX	5.1%	3.6%	8.1%	0.44	-26.0%
Vol targeting, 2D RV	7.6%	6.1%	10.0%	0.60	-30.6%

To test using RV-based forecasting in our volatility targeting framework, we devise a very simple measure calculated from average of the last two days as the forecast (2D RV).[9]

The 2D RV estimate generated 92% of the buy and hold annual return with a little over half as much volatility, leading to a 43% boost in the Sharpe ratio. Max drawdown was comparable to the other volatility-controlled strategies, falling in between the 30D HV and the more constrained VIX-driven method. Accuracy in targeting the 10% level was also superior using 2D RV compared to the other methods.

Cash: The Zero Volatility Solution

One of the reasons volatility targeting works is the stability of the asset used to dynamically adjust the portfolio: the volatility of cash is effectively zero. Switching to cash is one of the most effective means of avoiding losses in risky assets as markets tumble. However, timing the market is notoriously difficult.

Figure 7: 200D MA System vs. Volatility Targeting

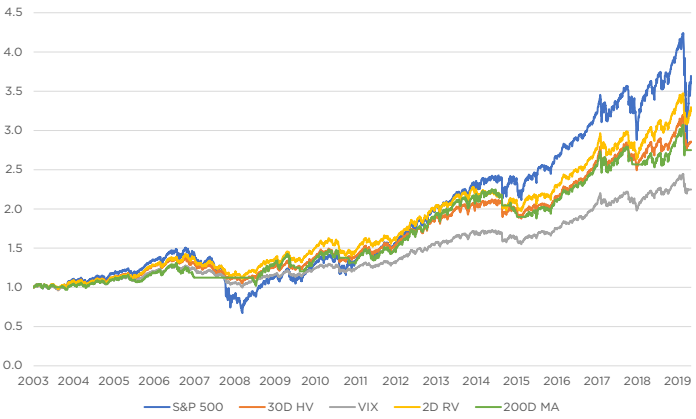

Strategy	CAGR	Avg Return	Volatility	Sharpe	Max DD
Buy and hold	8.3%	8.1%	19.4%	0.42	-55.3%
Vol targeting, 30D HV	6.6%	5.2%	10.3%	0.51	-33.9%
Vol targeting, VIX	5.1%	3.6%	8.1%	0.44	-26.0%
Vol targeting, 2D RV	7.6%	6.1%	10.0%	0.60	-30.6%
200D MA System	6.4%	5.1%	11.0%	0.46	-22.1%

To extend our analysis, we compare our hypothetical volatility-targeted strategies to another simple but still popular tool for detecting bull and bear markets: the 200-day moving average (200D MA). A market trading above its 200D MA is said to be in a bull market, with higher prices more likely, and below it a bear market when cash is preferred.

We test this premise using a very simple rule—own 100% stocks when above the 200 MA, and 100% cash when below it—using the same data, implementation assumptions, and sample period used in the volatility targeting analysis.

The basic 200D MA system does a good job of lowering volatility and generating better risk-adjusted returns than a buy and hold investment. Its ability to stay out of the way in bear markets (note the flat line during the 2008-2009 period) helps reduce its max drawdown considerably compared to the volatility-targeted strategies. But its Sharpe ratio falls short of even the naïve 30D HV strategy.

[9] The purpose of higher frequency data is to get a more responsive measure. However, ill-placed responsiveness equates to meaningless turnover and suboptimal results. A 2-day average is a simple way to retain most of the responsiveness of the leading day, while taming it slightly with the prior day.

The signal from a 200D MA system is by nature slower—it takes a significant change in price or the passage of time to cross through the moving average signal line. The volatility targeting strategy—especially the 2D RV—is more responsive to market moves. As a result, the 200D MA system tends to do well in strongly trending markets (up or down), and less so in choppy periods with more frequent selloffs followed by "v-bottom" type recoveries. The chart below illustrates the historical drawdowns over the period for the 200D MA, 2D RV, and S&P 500. The 200D MA system clearly does a bit better on avoiding the worst of the declines (2008) but can be slower to respond and lag in markets that turn around faster (2009, 2011, 2015-2016, 2020 H1). Both strategies can be effective in reducing risk, but volatility targeting—especially enhanced with realized volatility estimates—demonstrated better performance in our analysis.

Figure 8: Historical Drawdowns, 2D RV, S&P 500, 200D MA

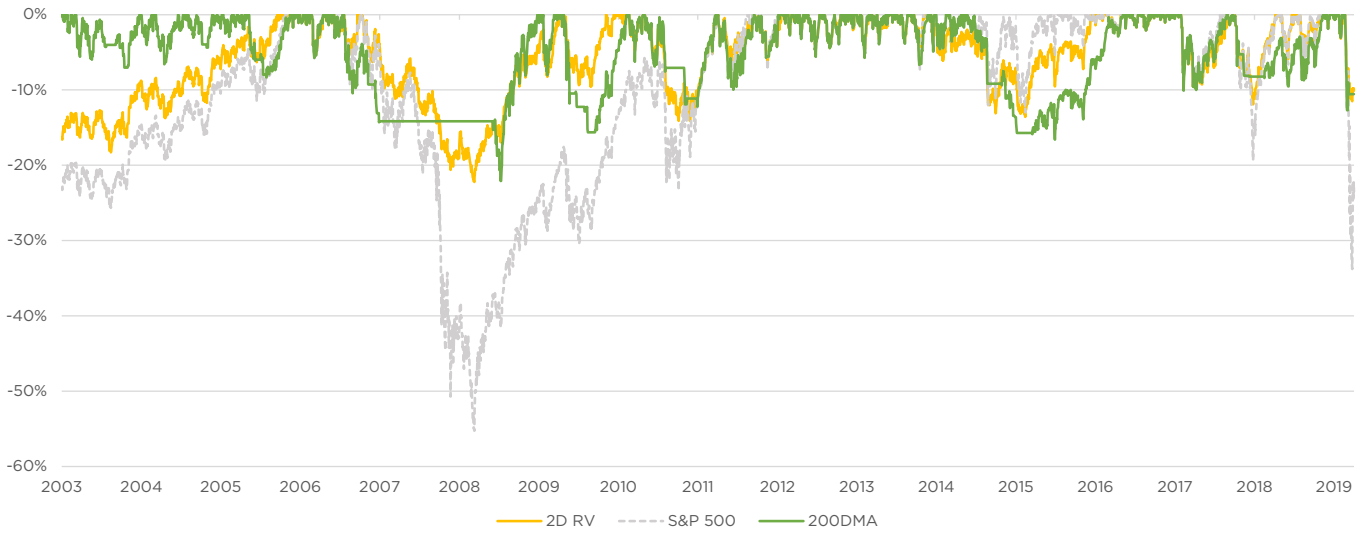

Conclusion

In the search for higher risk-adjusted returns, it is often said the best strategy is one an investor can stick with over the long term. Volatility is an unusual beast. While not a friend to investors, understanding how it behaves and harnessing its tendencies to be more predictable can help construct portfolios with the potential to outperform over a full market cycle. More effectively managing volatility can help investors keep their heads and stay on track to reach their objectives.

We demonstrate how volatility targeting, even with naïve estimates, can be effective in managing risk and enhancing risk-adjusted returns. The strategy also compares favorably to trend-following techniques using measures like the 200-day moving average. Lastly, the addition of higher frequency data to help improve responsiveness and accuracy in forecasting clearly demonstrates benefits within a volatility targeting framework.

Our very basic 2-day average of realized volatility observations performed well, but more sophisticated techniques leveraging similar underlying data have the potential to enhance performance even more. Many leading academics in research on volatility have partnered with some of the largest quantitative investment managers in the world to develop more advanced methods of forecasting.[10] While a large amount of this research is publicly available, other techniques are likely held closer to the vest to protect their intellectual property. We expect additional development of volatility forecasting methods to expand in the years to come with wider adoption of volatility-targeted strategies by more investor segments.

[10] Leading quantitative managers such as AQR and Man Group have published a body of research in collaboration with many of the academics with the most significant contributions to the literature on realized volatility since the late 1990s. See the aforementioned "Risk Everywhere: Modeling and Managing Volatility" from AQR at https://tinyurl.com/y7vant9j or "The Impact of Volatility Targeting" by Campbell Harvey and a team from Man Group, available at https://tinyurl.com/ya64qzq6.

About the Authors

Tony Barchetto, CFA is the Founder and Chief Executive Officer of Salt Financial, LLC, a leading provider of index solutions and risk analytics for investment advisors, fund sponsors, investment banks, and insurance carriers. Salt develops risk-focused index strategies built for higher precision and responsiveness to changing markets using its truBeta™ and truVol™ analytics.

Prior to founding Salt, Tony led corporate development for Bats Global Markets, helping guide the firm from its merger with rival Direct Edge to initial public offering and subsequent acquisition by Cboe Global Markets. Before Bats, Tony served in a similar capacity for global institutional trading platform Liquidnet. Earlier in his career, he served in leadership roles in trading, sales, and product management with Citigroup and Knight Capital Group.

Tony earned a BS in International Economics from Georgetown University and is a CFA Charterholder.

Ryan Poirier, ASA, CFA, FRM is Director, Index Products & Research for Salt Financial. In his role, he is responsible for the conceptualization and design of index-based strategies leveraging Salt Financials' full suite of truBeta™ and truVol™ analytics. He also leads the research efforts, focusing primarily on risk forecasting and portfolio construction.

Prior to joining Salt Financial, Ryan spent three years at S&P Dow Jones Indices as part of the Global Research & Design team, where he grew the quantitative index offerings for banks, asset managers, and insurance companies. He is a CFA Charterholder, FRM Charterholder, and an Associate of the Society of Actuaries.

Ryan earned a MS in Financial Engineering from New York University and a BS in Mathematics and Finance from SUNY Plattsburgh.

About Salt Financial

Headquartered in New York, Salt Financial is a leading provider of index-based products, risk analytics, and ETFs for the global investing community. We use higher frequency returns and modern data science to develop more sophisticated measures of volatility and market risk. Our analytics, including truBeta™ and truVol™, power index-based solutions for investment advisers and financial institutions to help target and control risk in their portfolios.

To learn more, please visit www.saltfinancial.com.

SALT FINANCIAL