Relationship Between Trend-Following and Options

This Version: November 6, 2024

Vineer Bhansali is the Founder and CIO at LongTail Alpha, LLC

vb@longtailalpha.com

Jeremie Holdom[1] is an Economic Research Strategist at LongTail Alpha, LLC

jh@longtailalpha.com

LongTail Alpha, LLC

500 Newport Center Drive, Suite 820

Newport Beach, CA 92660

[1] Corresponding Author: Email: jh@longtailalpha.com

Introduction

In this note, we will discuss the deep relationship between two different risk-mitigation strategies that have become important in the evolution of diversification strategies, namely "Trend-Following" and "Options Based Hedging". Both these strategies, as is well-established empirically, aim to, and have delivered positive performance when risk assets, especially equity markets, are doing badly. However, at first blush, the two strategies look very different. Trend-following primarily aims to provide diversification by following systematic rules for entering linear instruments, such as futures and swaps across a wide set of assets. On the other hand, the use of options-based risk-mitigation is based on the inherent non-linearity contractually built into options contracts. The actions of a trend-follower, by extension, depend on the realized volatility of the underlying instruments, whereas the actions of an options-based strategy are primarily driven by the implied volatility of the options contracts.

Trend-following strategies, which are characterized by actively increasing exposure in rising markets and decreasing exposure in falling ones, embody the principle of "buy high and sell even higher". Options, on the other hand, provide investors with the right to buy or sell assets at pre-defined strike prices. Call options provide the right to buy an underlying asset within a specified period, while put options grant the right to sell. As the price of the underlying asset increases, the value and delta (exposure to the underlying) of a call option rises and the value of a put option falls. Conversely, when the underlying price decreases, the put option's value (and negative exposure to the underlying) increases and the call option's value decreases. Thus, both the call and put options, at face value, pro-cyclically increase the absolute value of the delta, or their exposure to the underlying. While trend-following does it through an active increase or decrease in exposure, options do the same contractually. Options have a fixed expiration date, whereas trend-following does not have an expiration date, but requires periodic rebalancing to maintain exposures.

The question which has previously been addressed by many researchers, and which we revisit in this paper, is whether the apparent similarities (and some differences) between trend-following and options can be reconciled at a deeper level and used for practical portfolio construction. If made rigorous, this reconciliation has the potential to add value to both strategies. In other words, can a trend-following algorithm be improved by evaluating the optionality embedded in the algorithm and thus supplementing "classic" trend-following by using explicit options? Going in the other direction, can a purely options-based strategy be improved by using systematic timing rules such as those used in trend-following?

It is our belief that this is indeed the case. This paper explores the intrinsic and deep relationship between trend-following and options, highlighting how the dynamic adjustments in exposure mirror the payoff structures of certain option positions. It then follows with suggestions for practical implementation of the relationships for mutual improvement in both trend-following and options-based strategies. As a preview of our main motivation for writing this paper, note that the inclusion of options in trend-following has the potential to mitigate the Achilles Heel of trend-following, i.e. sharp reversal risks. On the other side of the coin, by using trend-following signals, options-based tail risk hedging strategies may be made more cost efficient.

Conceptual Observations

By taking long or short positions in the underlying asset equal to the option's delta ($\Delta = \partial O/\partial S$), one can locally replicate the profit and loss of a call or put option. This delta replication strategy involves buying the underlying asset in rising markets and selling in falling markets for both call and put options, thus adhering to the trend-following principle described above. One could think of delta replication as the activity of a market maker who is net short the payoff of an option, and thus has to trade in the underlying to remain hedged to leading order.

When evaluating a trend-following program that can take such replicating long and short positions, it is instructive to compare it to a combination of call and put options, such as a straddle, so that the terminal payoff is symmetrical. The strike price of the straddle is analogous to the threshold for switching from long to short, or vice versa, and can be taken as a moving average of prices or other similar filter which determines when the program shifts between long and short positions. The magnitude of exposure around this threshold is analogous to the option's delta, representing sensitivity to changes in the underlying asset's price. The gamma, indicating the rate of change of delta, reflects how quickly the program adjusts its exposure over time. Both the trend-following strategy and the options positions are affected by the impact of volatility. However, a key distinction is that options have exposure to implied volatility expectations: the option's premium incorporates the market's forecast of future volatility, allowing it to benefit immediately from anticipated volatility increases. In contrast, the trend-following program gains from longer-term volatility only as it unfolds over time through its trading activities. In other words, to be analogous to an option, the trend-follower has to construct active exposures to realized volatility through specific systematic choices in the replication algorithm.

Empirical Connections

Observing that trend-following strategies exhibit convex payoffs is well-established. Fung and Hsieh (2001) connected these strategies to options by demonstrating that the performance of trend followers can be replicated through delta replication of straddles and lookback straddles on the underlying assets, thus explaining the convexity of such strategies. While the standard straddle is familiar to most investors, in a lookback straddle the investor has the option of “looking-back” over a fixed window and benefiting from the highest price observed in the lookback window for a lookback put, and the lowest price observed in the lookback window for a lookback call. While both standard straddles and lookback straddles enable investors to profit from substantial price expansion in either direction, subtle differences between them are particularly relevant to trend-following strategies. The key distinction lies in how their delta exposures change relative to movements in the underlying market, a topic which we will discuss in greater detail in the following section. Exhibit 1, adapted from Fung and Hsieh (2001), illustrates these differences: the blue line represents the delta exposure of a lookback straddle struck initially at a one month expiration, and the gray line depicts that of a standard straddle, with the underlying market price shown in purple on the right-hand axis.

Exhibit 1: Comparing Delta Exposures of Lookback and Standard Straddles

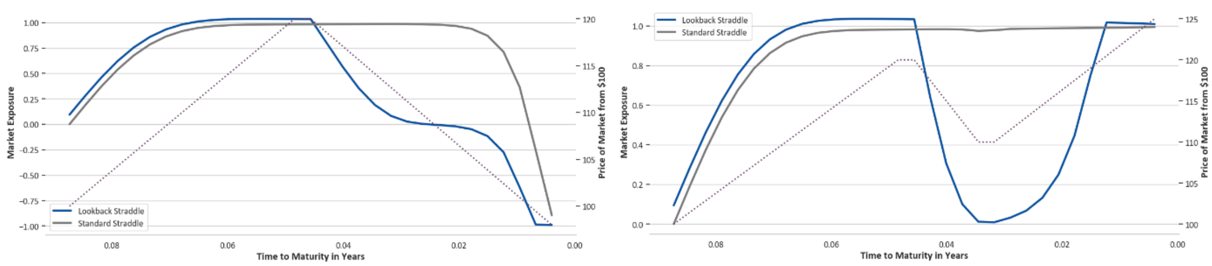

In the first example, the standard straddle and the lookback straddle have similar deltas for the increase in price, but as the price falls, the look-back straddle starts to flip its exposure much quicker than the standard straddle. Similarly, in the second example in Exhibit 1, the lookback straddle again flips its exposures more dynamically than a standard straddle as the underlying first declines and then recovers.

Based on historical performance, trend-followers operating across multiple asset classes have demonstrated a convex payoff relative to equity markets, serving as effective diversifiers during critical periods such as the Global Financial Crisis (GFC) and the 2022 inflationary selloff. This is shown in Exhibit 2 below, which displays a quadratic fit of monthly returns for two trend-following indices versus the MSCI ACWI index.

Exhibit 2: Quadratic Fit of Trend-Following Indices vs Global Equities (Monthly Returns)

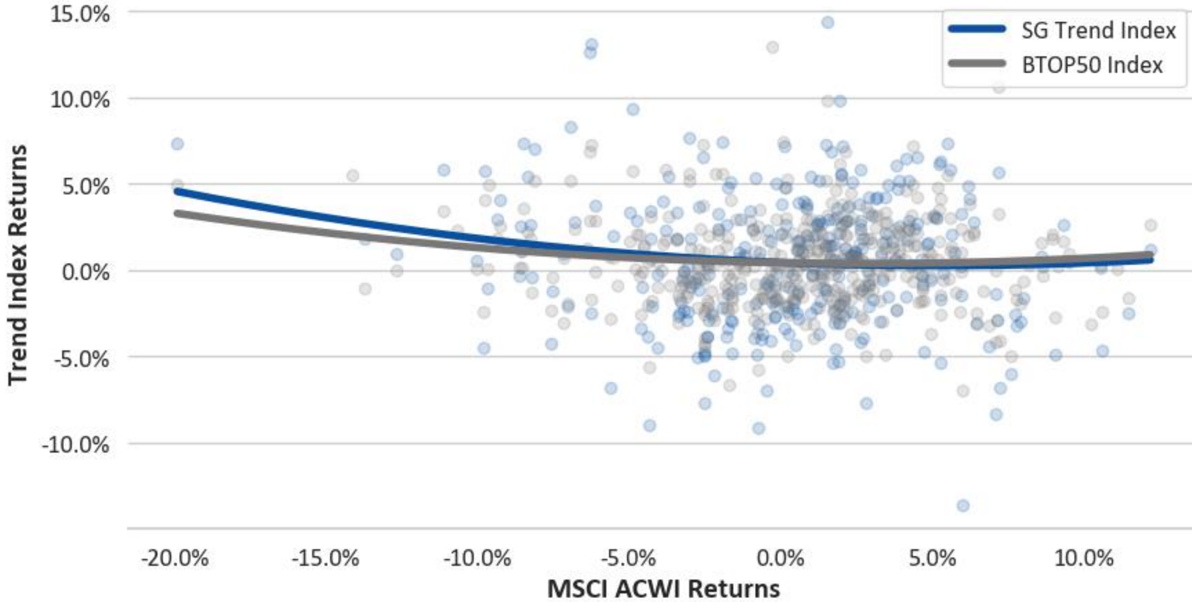

This favorable payoff profile has led to the inclusion of trend-following strategies in risk mitigation and diversification allocations within institutional portfolios. Exhibit 3 below shows how a 50/50 allocation to global equities and the BTOP index was able to meaningfully increase CAGR, reduce the beta to global equities, and reduce the portfolio drawdown. Zooming in on the most recent two decades, the 50/50 portfolio of global equities and BTOP was not able to keep up with the CAGR of a 100% allocation to equities but was still able to deliver superior risk-adjusted returns. Over the same period, a 60/40 portfolio of global equities and the CBOE long-volatility index was able to deliver similar risk-adjusted returns, showing the benefits of options-based long-volatility strategies.

Exhibit 3: Portfolio Statistics of Including Trend-Following to Global Equities

Dec 1987 - Jun 2024	MSCI ACWI	50/50 ACWI + BTOP	Dec 2004 - Jun 2024	MSCI ACWI	50/50 ACWI + BTOP	60/40 ACWI + CBOE Long Vol
CAGR	5.87%	6.32%	CAGR	5.46%	4.80%	5.11%
Max Drawdown	-56.23%	-27.86%	Max Drawdown	-56.23%	-27.86%	-26.80%
Calmar Ratio	0.10	0.23	Calmar Ratio	0.10	0.17	0.19
Monthly Vol (ann.)	15.20%	8.54%	Monthly Vol (ann.)	15.83%	8.59%	8.50%
Monthly Sharpe	0.45	0.76	Monthly Sharpe	0.42	0.59	0.63
Beta Global Equities	1.00	0.48	Beta Global Equities	1.00	0.50	0.51

A few questions immediately arise: (1) trend-following uses multiple asset classes, and primarily linear instruments like futures and swaps. How is it then possible to offer crisis protection against equity markets? (2) How important are the details of a particular trend-following algorithm to offer the benefits of diversification and convexity against equity market selloffs? (3) In what sense is trend-following like an option i.e. what are the "greeks" analogous to the "delta", "gamma", "vega", "theta" etc.? A consequence of the ability to answer these questions is that an investor can then optimally combine the greeks from both options and the trend-following algorithm to create a portfolio that is potentially better than just using trend-following or options alone.

In the following section, we aim to delve into the underlying mechanisms that confer convexity to trend-following strategies and explore how running a trend-following rule on an individual security can be effectively linked to long option portfolios to answer these questions.

Theoretical Connections

Assessing the performance of a fund or strategy by drawing analogies to options markets has long been a significant topic in finance. Merton (1981) developed a model demonstrating that a market timer—an investor who can predict, with some error, whether equity markets will outperform the risk-free rate—achieves a payoff profile equivalent to holding a straddle. The market timer invests in equities when equities are expected to outperform the risk-free rate and shorts them when they are expected to underperform. Merton showed that the equilibrium management fee such a market timer should charge is equal to the price of the straddle. This establishes the theoretical link between options markets and strategies that profit from both upward and downward market movements, akin to trend-following strategies.

The link between profiting from price expansion and trend-following strategies becomes evident when examining the payoff profiles of straddles across different maturities. Denote $\phi_t(S_t, T) = C_t(S_t, T) + P_t(S_t, T)$ as the price of a straddle at time t expiring at T, where C_t, P_t represent the prices of the ATM call and put components and S_t is the underlying spot. Define the payoff "range" of the straddle as the difference between the price of the straddle if we shocked the spot to expiration minus its initial entry costs.

$$R(S_T, T) = \phi_0(S_T, T) - \phi_0(S_0, T)$$

We can consider the payoff range of the straddles as analogous to a desired payoff profile of a delta replication strategy. Exhibit 4 shows the payoff ranges for straddles of various expirations as a function of the underlying terminal price. The zero days-to-expiration ("0DTE") payoff range could be achieved by using a binary trend signal of +1 if the market moves above the current strike, and -1 if the market moves below the strike, and thus represents the target profit and loss of a simple trend-program that goes long or short around the initial strike price. Varying the maturity of the straddle changes the convexity of its payoff range. We can see that despite the subtle differences between the straddles, all profits increase significantly when the underlying market moves substantially in either direction.

Exhibit 4: Straddle Payoff Ranges

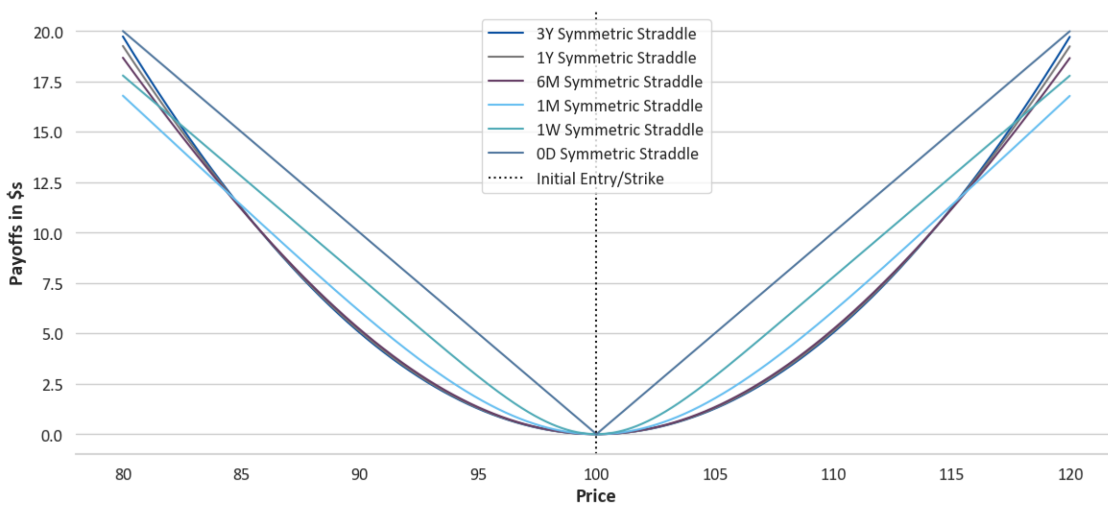

To make the straddle payoffs symmetrical on either side, we use a variation of the Black-Scholes option pricing formula, namely the Bachelier model, which assumes the price of the underlying market is normally distributed as opposed to log-normally distributed. Using the Black-Scholes model produces a very similar illustration but creates a tilt in the payoffs from the log-normal assumption. Specifically, the straddle payoff range is equal to $R(S,T) = \phi(S,T) - \phi(100,T) = C(S,T) + P(S,T) - C(100,T) - P(100,T) = \left[(S-100)N(d) + \sigma\sqrt{T}n(d)\right] + \left[(100-S)N(-d) + \sigma\sqrt{T}n(d)\right] - 2\sigma\sqrt{T}n(d)$ where S is the underlying price, T is the maturity of the put and call options in years, N, n are the normal CDF and PDF, and $d = (S-100)/(\sigma\sqrt{T})$ with σ being the volatility of the price of the underlying asset. Payoffs are normalized so that they all have the same delta exposure at a +-10% move, i.e., $\phi(S)/|\phi'(110)| = \phi(S)/|\Delta(110)| = 1$.

Delta replication involves adjusting positions in the underlying asset by buying or selling amounts equal to the delta of the combined options in the straddle, effectively mirroring the straddle's exposure. Consider a strategy that replicates one of the payoff ranges in Exhibit 4. Each time step (i.e. hour, day, etc.), the strategy determines the quantity to buy or sell based on the delta of the straddle at $t = 0$, namely $\Delta_t = \partial\phi_0(S_t, T)/\partial S_t$. This means that at each time step, the delta references the straddle struck at inception, keeping the time until expiration fixed at T and in a sense preventing the payoff range being replicated to season. The profit and loss of this replication strategy is approximately equal to the payoff range minus the sum of the gamma at each time step multiplied by the "quadratic variation" of the underlying asset's price.[2]

$$
\begin{aligned}
PNL_t^R &= \sum_{s=1}^{t} \Delta_{s-1}(S_s - S_{s-1}) \\
&\approx [\phi_0(S_t, T) - \phi_0(S_0, T)] - \frac{1}{2}\sum_{s=1}^{t} \Gamma_{s-1}(S_s - S_{s-1})^2 \\
&= \underbrace{R(S_t, T)}_{Payoff\ Range} - \underbrace{\frac{1}{2}\sum_{s=1}^{t} \Gamma_{s-1}(S_s - S_{s-1})^2}_{Quadratic\ Variation\ Costs}
\end{aligned}
$$

Quadratic variation represents the cumulative effect of small but frequent fluctuations that render a process stochastic. This term arises due to the slippage, the inherent cost of replicating any payoff in a stochastic or random process. Consequently, these replication strategies and by extension trend-following resemble being long on longer-term price expansion but short shorter-term volatility, reflecting the imperfect replication of the target payoff range (see Bouchaud (2017)).

In contrast to the delta replication, consider the profit from purchasing the straddle, which requires an upfront payment of a premium. This includes a premium decay from the options getting closer to their expiration, and by paying this time decay, the straddle can obtain the payoff range without incurring the costs from delta replication. This highlights the fundamental difference between delta replication and purchasing a straddle: while delta replicating attempts to mimic the option payoff dynamically (incurring slippage costs along the way), purchasing the straddle involves a cost of time decay without the need for continuous adjustments. Thus, we can see how the greek "theta" of the option is related to the unavoidable cost from whipsaws and transactions costs in the replicating strategy. Moreover, since the price of the straddle is determined by its implied volatility at any given time, the straddle can benefit from changes in volatility expectations on a marked-to-market basis while the delta replication strategy must continue to trade until expiration to realize those expectations, assuming they are indeed correct. An under-evaluation of the long term realized volatility would then penalize the trend-following strategy and benefit the explicit options-based strategy.

---

[2] This relationship can be demonstrated by applying Itô's formula to the replication process, assuming that the underlying asset S_t is a stochastic Itô process. Rearranging Itô's formula for a function f that depends only on S_t, we get $\int_0^t f'(S_s)dS_s = [f(S_t) - f(S_0)] - \frac{1}{2}\int_0^t f''(S_s)[S]_s$, where $[S]_s$ represents the quadratic variation of the underlying asset. Substituting $\Delta_t = f'(S_t)$ and $\Gamma_t = f''(S_t)$ into the previous formula gives us the continuous analogue of the discrete formula. When the payoff function f is convex ($\Gamma_t > 0$), the trading strategy incurs a cost due to the quadratic variation of the process. Conversely, if the payoff is concave, one may profit by trading in a mean-reverting manner, effectively earning from the quadratic variation.

$$\phi_t(S_t,T) - \phi_0(S_0,T) = [\phi_t(S_t,T) - \phi_0(S_t,T)] + R(S_t,T)$$
$$= \underbrace{[\phi_t(S_t,T) - \phi_0(S_t,T)]}_{Premium\ Decay} + \underbrace{PNL_t^R + \frac{1}{2}\sum_{s=1}^{t} \Gamma_{s-1}(S_s - S_{s-1})^2}_{Delta\ Replication\ Without\ Costs}$$

One can change the delta replication strategy to use the delta of the straddle at t, namely $\Delta_t = \partial\phi_t(S_t,T)/\partial S_t$. The difference between this method and the one that references the delta of the option straddle at inception when $t = 0$ is that the payoff profile being replicated is equal to the profit from purchasing the option straddle. This is because the delta used in the replication is being updated over time as the option straddle matures. Unlike the option straddle however, the delta replication does not lose any money simply from the passage of time, and as a result is able to achieve the option straddle profit without incurring the time decay $\theta_t < 0$. But it still pays the slippage due to quadratic variation, although slightly different from the previous replication strategy since the gamma is based on the straddle at t.[3]

$$PNL_t^S = \sum_{s=1}^{t} \Delta_{s-1}(S_s - S_{s-1})$$
$$\approx \underbrace{[\phi_t(S_t,T) - \phi_0(S_0,T)] - \sum_{s=1}^{t} \theta_s}_{Straddle\ PNL\ Without\ Time\ Decay} - \underbrace{\frac{1}{2}\sum_{s=1}^{t} \Gamma_{s-1}(S_s - S_{s-1})^2}_{Quadratic\ Variation\ Costs}$$

So far, we have considered option straddles where the strike remains fixed throughout the lifetime of the trade. Trend-following typically has an entry rule, such as an exponential moving average that determines whether to go long or short. We can now connect this concept of using a dynamic strike for a straddle as opposed to a fixed strike for a straddle.

[3] This relationship can be demonstrated by using Itô's formula for a function f that depends on S_t and time t. From $\int_0^t \frac{\partial f}{\partial S}(S_s)dS_s = [f(S_t) - f(S_0)] - \int_0^t \frac{\partial f}{\partial t} ds - \frac{1}{2}\int_0^t \frac{\partial^2 f}{\partial S^2}[S]_s$, we substitute $\Delta_s = \frac{\partial f}{\partial S}(S_s), \Gamma_s = \frac{\partial^2 f}{\partial S^2}, \theta_s = \frac{\partial f}{\partial t}$ to get the continuous analogue of the discrete formula.

Fung and Hsieh (2001) argued and showed that a trend-following program could be approximated using lookback straddles. A lookback straddle combines a lookback put, which pays the difference between the maximum over the horizon and the terminal spot price $S_{max} - S_T$, with a lookback call, which pays the difference between the terminal spot price and the minimum over the horizon $S_T - S_{min}$. Therefore, the lookback straddle combination is a single derivative contract with terminal payoff equal to $S_{max} - S_{min}$. The general idea behind employing a lookback straddle to represent trend-following strategies is that it allows the holder to profit from a large spread between the maximum and minimum over the time horizon, which can emerge due to increasing longer-term realized volatility or price expansion.

Goldman et. al. (1979) developed a closed-form expression for the price of lookback put and call options and illustrated how lookback options can be replicated. The general replication strategy for a lookback straddle is as follows: begin with two straddles each struck at the same initial strike equal to spot. Whenever a price hits a new maximum, sell one of the straddles and restrike it at the new maximum. Similarly, whenever a price hits a new minimum, sell the other straddle and restrike it at the new minimum. This replication argument shows how a lookback straddle, as proposed by Fung and Hsieh (2001) to replicate trend-following programs, is continuously adjusting straddles whenever prices reach new extremes. In the same fashion, referencing a straddle that is continuously restruck to the filter, such as the moving average, of a trend program is conceptually very similar to the ideas behind using a lookback straddle. In both cases the static replication is replaced by a path-dependent, dynamic replication strategy.

We will make one final connection between trend-following strategies and options by comparing the length of trend signals to the tenor or expiry of an option straddle being delta-replicated. Consider three simple market models that can generate significant price expansion relative to the initial price: constant positive drift, positive autocorrelation, and time-varying drift. Correspondingly, we examine three trend signals of varying lengths: a short-term 10-day moving average (MA), a medium-term 50-day MA, and a long-term 200-day MA. The trend positions are binary, taking values of +1 or -1 depending on whether the closing price is above or below the moving average.

We report the trading statistics of these strategies under the specified market dynamics. Longer-term signals tend to profit from drift, while shorter-term signals capitalize on autocorrelation. A short-term trend signal adjusts its exposure rapidly, which in option terminology corresponds to a high gamma—the rate at which delta changes with respect to the underlying price. Shorter-dated options exhibit higher gamma and more rapidly changing delta exposures, making them analogous to short-term trend strategies. Conversely, longer-dated options have relatively lower gamma and are better suited for capturing drift, similar to long-term trend strategies. Drift that varies over time is a generalization of the constant positive drift model, and we can see that a longer-term signal unsurprisingly does better than the long-only portfolio.

Exhibit 5: Trend Signal Lengths Profit from Different Market Dynamics

Constant Positive Drift

$$r_t = \mu + \epsilon_t$$

	Annualized Return	Volatility	Information Ratio
MA 10	0.41%	20.02%	0.02
MA 50	1.11%	20.02%	0.06
MA 200	2.25%	20.01%	0.11
Long Only	9.99%	20.01%	0.50

Positive Autocorrelation

$$r_t = \mu + \alpha r_{t-1} + \epsilon_t$$

	Annualized Return	Volatility	Information Ratio
MA 10	7.26%	20.01%	0.36
MA 50	4.17%	20.01%	0.21
MA 200	3.76%	20.01%	0.19
Long Only	10.02%	20.01%	0.50

Time-Varying Drift

$$r_t = \mu_t + \epsilon_t$$

	Annualized Return	Volatility	Information Ratio
MA 10	2.65%	20.04%	0.13
MA 50	5.84%	20.03%	0.29
MA 200	10.37%	20.02%	0.52
Long Only	2.03%	20.01%	0.10

Results are created by running the various trading strategies on simulated price data, where the log- returns of the price series are generated as specified by the respective model. The constant positive drift model assumes an annualized return of 10%, and an annualized volatility of 20%. The positive autocorrelation model assumes an autoregressive coefficient of $\alpha = 0.05$. *The time-varying drift model uses the methodology described in* (Giraitis 2014) *to create a time-varying autoregressive coefficient but instead applied to the drift* μ_t.

Below, we present trading statistics for three straddle delta replication strategies with different tenors. Each day, these strategies trade the number of shares in the underlier required to match the delta of a straddle consisting of a long ATM put and ATM call that is re-struck once the options expire. The behavior of longer-dated options parallels that of long-term trend strategies, while shorter-dated options mimic the characteristics of short-term trend strategies.

Exhibit 6: Option Maturities Profit from Different Market Dynamics

Constant Positive Drift

$$r_t = \mu + \epsilon_t$$

	Annualized Return	Volatility	Information Ratio
Straddle 1W	0.26%	9.99%	0.03
Straddle 1M	0.54%	11.05%	0.05
Straddle 1Y	1.97%	12.26%	0.15
Long Only	10.00%	20.01%	0.50

Positive Autocorrelation

$$r_t = \mu + \alpha r_{t-1} + \epsilon_t$$

	Annualized Return	Volatility	Information Ratio
Straddle 1W	4.72%	10.12%	0.47
Straddle 1M	3.07%	11.25%	0.27
Straddle 1Y	2.66%	12.44%	0.20
Long Only	9.94%	20.00%	0.50

Time-Varying Drift

$$r_t = \mu_t + \epsilon_t$$

	Annualized Return	Volatility	Information Ratio
Straddle 1W	0.86%	10.02%	0.08
Straddle 1M	1.92%	11.13%	0.17
Straddle 1Y	6.79%	12.96%	0.51
Long Only	2.06%	20.01%	0.10

While substantial further research is needed to answer all three questions posed above, we can now see indications of the linkages. (1) Even though trend-following involves multiple assets, while equity options only involve equity puts and calls, the reason trend-following can diversify against equity risk is because the volatility expansion for a diversified pool of assets captures the changes in risk perception in equities. This seems to be the case empirically, since all asset returns tend to become correlated in the limit where there are large, systemic shocks as in the case of equity market crises. (2) The details of the trend-following algorithm do matter. For trend-following to act like options, the exposures of the program have to be systematically increased and decreased. (3) The "greeks" of options can be mapped into other, more natural features of trend-following. For instance, the option gamma, or rate of change of the delta, can be thought of as the specific rule for increasing the exposure of the trend-following program. While there is no "theta" or time-decay in trend-following, the cost of the replication now is transmuted into the transactions costs and whipsaws that are a natural consequence of volatility.

Use Cases

Armed with both theoretical and empirical insights on how trend-following is related to options, we now turn to practical applications. Applying the option framework to trend-following design choices provides valuable insight into optimizing these strategies for better performance and risk management. This generalized optionality framework that connects an explicit options based strategy with an implicit option like strategy is useful for studying risk mitigation strategies more broadly.

1. Periodic Rebalancing and Monetization Rules

Trend-following programs typically apply some weighting scheme around a threshold entry rule or signal. As previously discussed, we can express this strategy as replicating a constant maturity straddle where the strike is dynamically being updated based on the threshold filter. Instead of replicating a constant maturity straddle, one can choose to replicate a straddle that seasons over time (this is the difference in delta methodologies that were used to derive PNL_t^R and PNL_t^S above), effectively adding monetization rules that reduce exposure based on a calendar interval. Exhibit 7 below illustrates the tradeoff between these approaches: clearly in a strong trending market, it is better to maintain maximum exposure, but this idea of monetizing periodically can be beneficial in range-bound or mean-reverting periods. Thus, effective option monetization (see our paper "Monetization Matters") translates naturally to profit taking in a trend-following program. One conclusion of our monetization research was that some monetization of options is better than no monetization. Therefore, it follows that some profit-taking of trend is better than no profit-taking. Otherwise, the trend reverses and most, if not all of the profits, are given back.

Exhibit 7: Fixed Strike vs Calendar Restrike Trend-Following Examples

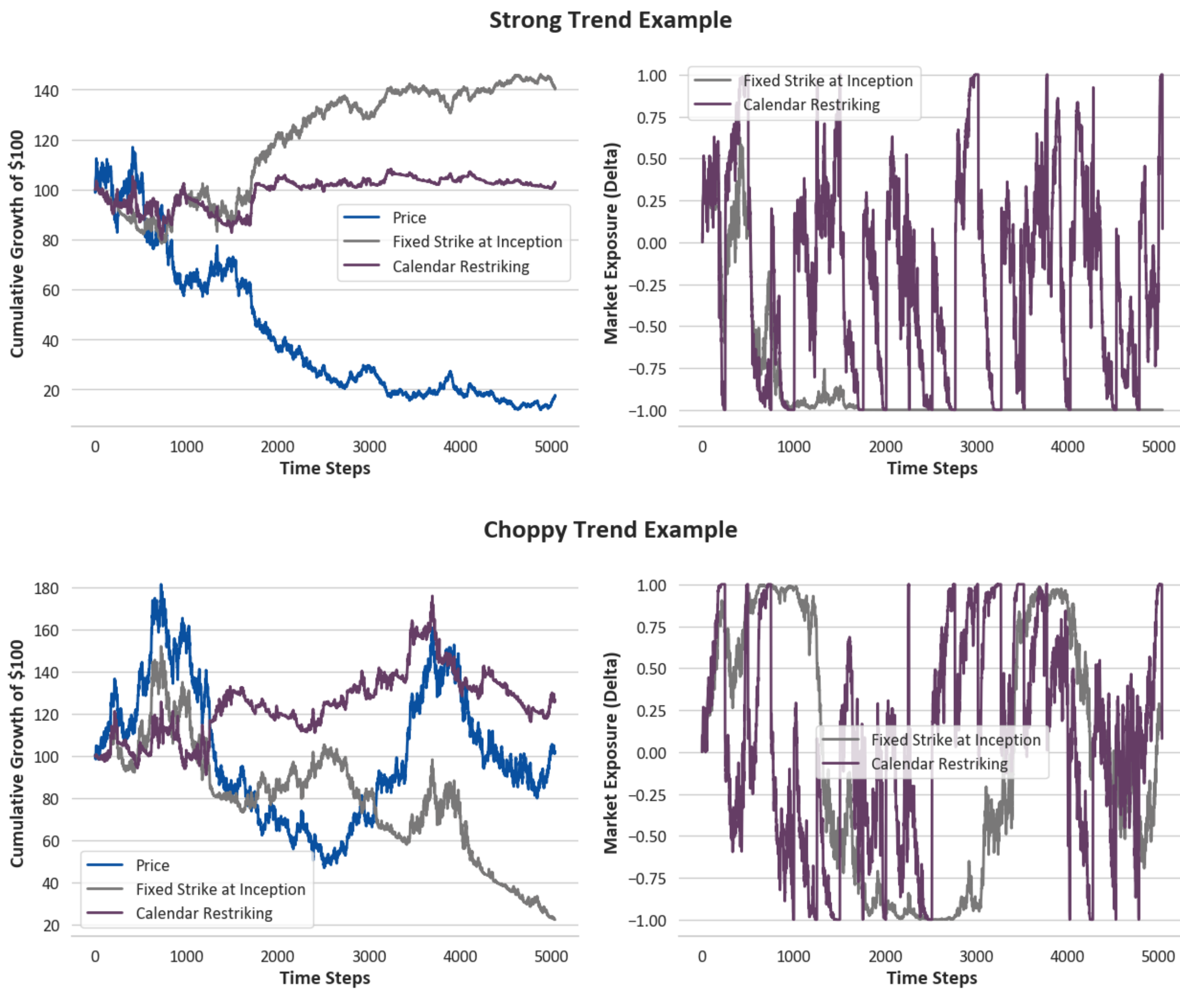

2. Signal Confirmation and Strangles

Our focus thus far in this paper has been on connecting trend-following to option straddles, but varying the strikes of the put and call options can create different market exposure profiles. Exhibit 8 below shows how strangles allow one to take a view that trends are going to appear only after a certain move. Though the differences between strangles and straddles are most pertinent for shorter maturities. In the trend-following context switching from straddles to strangles, which means spreading the strikes, would be akin to creating entry thresholds, i.e do not take a position until the position is confirmed by a large enough move in the underlying. Of course, in the option context the cost of spreading the strikes is a slower increase in convexity. In the trend-following context, the analogue is the cost of patience and confirmation. The longer one waits, the more sure one is that the trend is persistent, which can prove to be beneficial, but at the same time, there is a cost, i.e. the initial potential profits are not captured.

Exhibit 8: Delta Market Exposures for Straddles vs Strangles

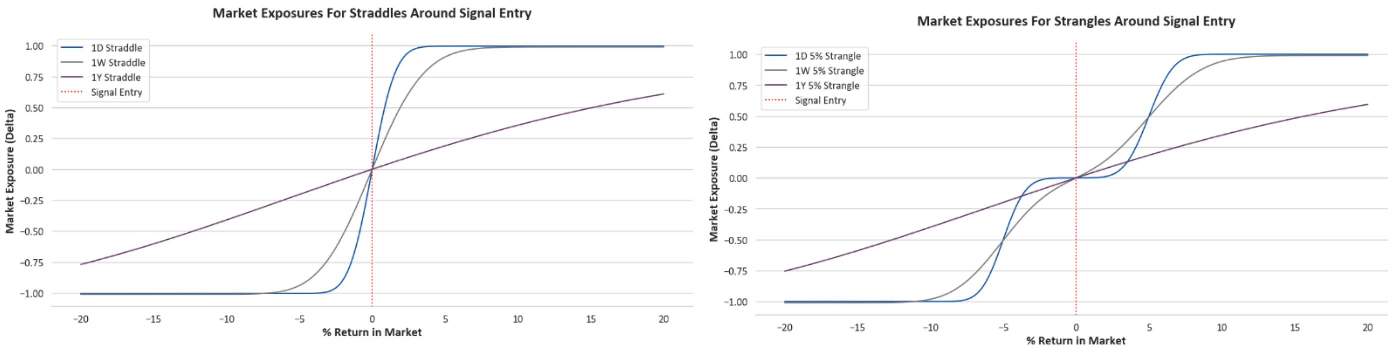

3. Stop Losses and Explicit Options

Many trend programs enter positions with pre-defined exit levels, effectively acting as stop-losses. Kaminski and Lo (2014) showed how stop-losses always decrease a strategy's expected return if prices follow a random walk, but if returns are positively autocorrelated, then stop-losses can add value to the strategy. It is also useful to consider how stop losses can be replicated using options to analyze additional tradeoffs one is making when using stops. Unlike a simple option, where the trade exits when the option expires, a stop loss will exit at a certain trigger level. This payoff profile is akin to exotic knock-out options, where in the case of a down-and-out call option, it expires worthless if a certain barrier threshold is hit at or below its strike $B \leq K$.

Let us first consider the case where $B = K$ <= S, corresponding to the left-side image in Exhibit 9. If the price remains above the strike, we are long the underlier, but if we hit the barrier, we contractually have no exposure. This is identical to a stop-loss order except that the stop-loss takes the jump risk of the price gapping through the barrier. When $B < K$, as in the right-side image, Carr (1998) showed that a down-and-out call is approximately equal to being long a $C(K)$ and short a further OTM put option $P(B^2K^{-1})$ (i.e. long a C100 and short a P90 when the barrier is 95). In both cases, we see that stop-losses are like being long a call option but short either jump risk at the barrier or a tail put below the barrier. As it relates to trend-following strategies, stop-losses suffer if the quadratic variation or short-term volatility of the underlying market is expected to rise at the barrier level.

The analogy here is clear: if one is to use stop-losses to manage downside risk, then it behooves them to consider what the cost of the stop-loss is in terms of potential sacrificed profits in case the underlier reverses course after the stop-loss is triggered. If the cost of the pure option replication is low enough, then it would be more efficient to replace the stop-loss with an option. When short-term volatility is lower than long-term volatility, and a trend program is designed to benefit from long term volatility increase, then using short term options might be a better approach than to exit a position and re-enter it later, assuming low short-term option volatility. This is a fundamental and key reason why we think exploring the possible role of options in mitigating reversal risk in trend-following is key, which is the subject of the next section.

Exhibit 9: Stop Losses are Similar to Knockout Options

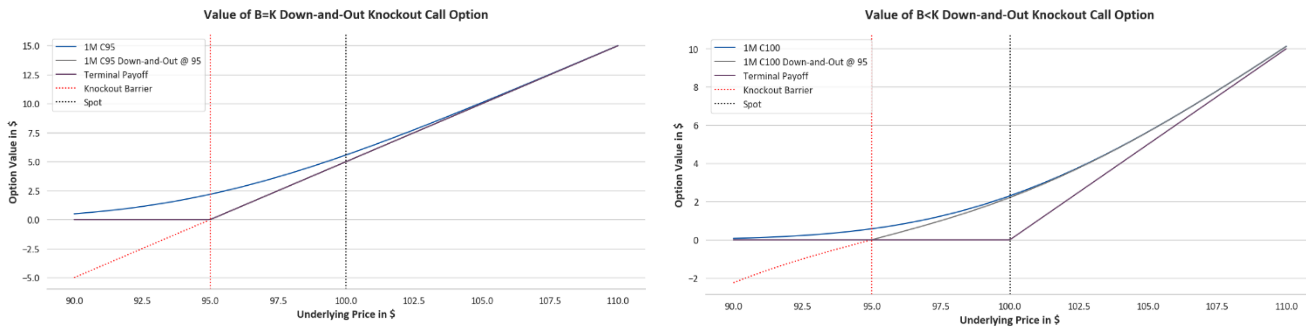

4. Managing Trend Reversal Risks With Options

In the previous section, we discussed how predefined exits in trend-following strategies, such as stop-losses, are analogous to knock-out options. These exits inherently expose the strategy to jump risks or an effective short tail put below that level. This exposure explains why rapid and sharp reversals against common trend-following positions can lead to significant losses - a phenomenon we refer to as reversal risks. Indeed, reversal risks are a "feature, not a bug" of trend-following. For trend-following to work, there has to be a reason why. To us, the reason is that trends can reverse without warning. But not all hope is lost. By sacrificing a small amount of potential gain, an investor can insure against the unknowable reversal by hedging the reversal through options. Indeed, for readers who have followed our discussion so far, it is also possible to hedge short term reversal risks by exiting based on short term momentum reversal signals. But the risk of a premature exit is that it can result in being out of the trend when the trend re-reverses and gains strength.

To mitigate these hidden reversal or tail risks, incorporating options can be effective. The choice and blending of options to mitigate reversal risks depend on several key factors. First, understanding the nature of potential reversals is crucial. If reversals are expected to be rapid or driven by strong autocorrelation, shorter-dated options with higher gamma are preferable, as they allow for swift adjustments in delta exposure. Conversely, if reversals are anticipated to be gradual and possibly accompanied by increases in future volatility expectations or sudden jumps, options with higher vega become more suitable due to their sensitivity to changes in implied volatility. Second, the cost of the options – specifically, the implied volatility premium over realized volatility at the time the hedge is initiated – plays a significant role because utilizing options entails an unavoidable time decay. Finally, market liquidity is an important consideration. In situations where market liquidity is low, the costs associated with quadratic variation in a delta replication or trend-following strategy can escalate significantly, making options a more attractive alternative.

Thus, one implementation of this framework is to manage trend with a longer-term expansion of volatility in mind, i.e. replicate the algorithm with a longer-term lookback straddle as the reference replicating option but manage the reversal risks with shorter term options as hedges. As long as the volatility expands over time due to the specific details of the underlying replication algorithm, this allows an investor to stay with their positions. In other words, while one waits for long term volatility expansion to pay off in the trend program, the investor is not taken out of their positions with short term reversals because the relatively inexpensive short-term options provide the insurance to hold the long-term positions. More details on this will be forthcoming in a future paper.

Conclusion and Future Work

In this paper, we have shown that trend-following strategies can be viewed as option replication strategies aiming to profit from longer-term price and dispersion expansions. Delta replication incurs costs due to quadratic variation, leading to slippage as it attempts to replicate a convex payoff. Purchasing an outright straddle involves paying time decay to achieve the same payoff and can benefit from implied volatility expansions during the trade. But the deep relationship between the two approaches shows how incorporating options into trend-following programs can assist an investor in managing reversal risks, in managing the tradeoffs and contributions from various option greek analogues, and creating better trend-following portfolios. Moving in the other direction, using trend-following signals can assist investors in potentially creating more efficient options-based hedging portfolios.

Another important direction for future exploration is the relationship between option moneyness and the skewness and higher moments of the underlying return distribution. The shape of this distribution significantly impacts the optimal choice of strike price or moneyness in option-based strategies. Allocating the same premium to further out-of-the-money options yields a more convex payoff but decreases the probability of finishing in-the-money. Conversely, in-the-money options provide higher initial delta exposure but offer less convexity if the trend intensifies substantially. Therefore, the ex-post return distribution should guide the preference for option moneyness, making the skewness of the distribution a crucial consideration.

References

Bhansali, Vineer, and Chang, Linda, and Holdom, Jeremie, and Johnson, Matthew. 2024. "Option Total Return and Active Option Portfolio Management." *To Appear in the Journal of Portfolio Management*.

Bhansali, Vineer, Linda Chang, Jeremie Holdom, and Marcy Rappaport. 2020. "Monetization Matters: Active Tail Risk Management and the Great Virus Crisis." *The Journal of Portfolio Management* 47, no. 1: 16-28.

Black, Fisher, and Scholes, Myron. 1973. "The Pricing of Options and Corporate Liabilities." *The Journal of Political Economy* 81 (3) 637-654.

Bouchaud, J. P., T. L. Dau, C. Deremble, Y. Lempière, T. T. Nguyen, and M. Potters. 2017. "Tail protection for long investors: Convexity at work." *J Inv Strat* no. 1 61-84.

Carr, Peter, and Ellis, Katrina, and Gupta, Vishal. 1998. "Static Hedging of Exotic Options." *The Journal of Finance* 53 (3) 1165-1190.

Choi, Jaehyuk, Minsuk Kwak, Chyng Wen Tee, and Yumeng Wang. 2022. "A Black–Scholes user's guide to the Bachelier model." *Journal of Futures Markets* 42, no. 5 (2022): 959-980.

Fung, W., & Hsieh, D. A. 2001. "The risk in hedge fund strategies: Theory and evidence from trend followers." *The review of financial studies* 14(2) 313-341.

Gatheral, Jim. 2001. "Lecture 4: Barrier Options." *Case Studies in Financial Modelling Course Notes*.

Giraitis, L., Kapetanios, G. and Yates, T. 2014. "Inference on stochastic time-varying coefficient models." *Journal of Econometrics* 179(1), pp.46-65.

Goldman, Barry, and Sosin, Howard, and Gatto, Mary. 1979. "Path Dependent Options: "Buy at the Low, Sell at the High"." *Journal of Finance* 34 (5) 1111-1127.

Kaminski, Kathryn M., and Andrew W. Lo. 2014. "When do stop-loss rules stop losses?" *Journal of Financial Markets* 18 234-254.

Lo, Andrew and Wang, Jiang. 1995. "Implementing Option Pricing Models When Asset Returns Are Predictable." *The Journal of Finance* 50(1).

Merton, Robert C. 1981. "On market timing and investment performance. I. An equilibrium theory of value for market forecasts." *Journal of business* 363-406.

IMPORTANT DISCLOSURES

The authors of and contributors to this paper are members of LongTail Alpha, LLC. Any opinions or views expressed herein are solely those of the authors and contributors, and do not necessarily reflect the opinions or views of LongTail Alpha, LLC or any of its affiliates. You should not treat any opinion expressed herein as investment advice or as a recommendation to make an investment in any particular investment strategy or investment product.

The data and information contained herein is not intended to predict the performance of any investment strategy based on market conditions. There can be no assurance that actual outcomes will match the assumptions or that actual returns will match any cumulative performance presented. The information contained herein is subject to change, and LongTail Alpha, LLC assumes no obligation to update the information. This is not an official statement and should not be relied upon as such. Several processes, assumptions and data sources were used to create the information provided. It is possible that different methodologies may have resulted in different outcomes. This data and information may not reflect the effect of material economic and market factors.

LongTail Alpha, LLC ("LongTail") is registered with the Securities & Exchange Commission as a registered investment adviser. LongTail Alpha is also registered with the Commodity Futures Trading Commission as a CTA and CPO and as a member of the National Futures Association. Registration does not imply a certain level of skill or training. This paper is furnished on a confidential basis and is not for redistribution or public use. The data and information presented are for informational purposes only and LongTail does not make representations as to the completeness or accuracy of any information contained herein. The information contained herein should be treated in a confidential manner and may not be transmitted, reproduced or used in whole or in part for any other purpose, nor may it be disclosed without the prior written consent of LongTail. All investing involves risk of loss, including the possible loss of all amounts invested. This document is not intended as and does not constitute an offer to sell any securities to any person or as a solicitation of any offer to purchase any securities, nor is it legal, tax, accounting or investment advice.

This document should not be viewed as a current or past recommendation to invest or to adopt any investment strategy discussed herein. The financial information and data contained in this report represents unaudited financial information and is subject to future adjustment and revision.

The performance shown was prepared by LongTail and has not been compiled, reviewed, or audited by an independent accountant. The results are based on internal books and records and are subject to adjustment following year-end audit. The strategy's returns are shown, in each case, at the end of the period indicated. The results are based on the periods as a whole, but results for individual months or quarters within each period will vary and will be more or less favorable than the average. The performance shown reflects investment of limited funds for a limited period and does not reflect performance in different economic or market cycles. Investors may not experience returns, if any, comparable to those shown. Past performance is not necessarily indicative of future results.

Certain of the exhibits included in this paper are examples for illustrative purposes only and are presented through hypothetical scenarios with hypothetical returns. Hypothetical performance results have many inherent limitations, some of which, but not all, are described herein. No representation is being made that any strategy will or is likely to achieve profits or losses similar to those shown herein. In fact, there are frequently sharp differences between hypothetical performance results and the actual results subsequently realized by any particular investment strategy. One of the limitations of hypothetical performance results is that they are generally prepared with the benefit of hindsight. In addition, hypothetical trading does not involve financial risk, and no hypothetical trading record can completely account for the impact of financial risk in actual trading. For example, the ability to withstand losses or adhere to a particular trading program in spite of trading losses are material points which can adversely affect actual trading results. The hypothetical performance results contained herein represent the application of certain strategies as currently in effect and there can be no assurance that the strategies will remain the same in the future or that an application of the current strategies in the future will produce similar results because the relevant market and economic conditions that prevailed during the hypothetical performance period will not necessarily recur. There are numerous other factors related to the markets in general or to the implementation of any specific trading program which cannot be fully accounted for in the preparation of hypothetical performance results, all of which can adversely affect actual trading results. Discounting factors may be applied to reduce suspected anomalies. Hypothetical performance results are presented for illustrative purposes only and should not be relied upon in making an investment decision.

This information is provided to you on the understanding that, as a sophisticated investor, you understand and accept the inherent limitations of the data presented, and you will not rely on it in making any investment decision. No representation is being made that any of the strategies will or are likely to achieve returns similar to any of those included. The financial information and data contained in this document represent unaudited financial information and is subject to future adjustment and revision.