
CONSIDERATIONS FOR PUT PURCHASES

Ideas, outline, and figures by Nikhil Jaisinghani, Placid Sound Capital

Drafting by Claude

July 2026

CONSIDERATIONS FOR PUT PURCHASES

AI Disclosure

The ideas, graphs, and thinking presented in this paper were developed by the author. The author created a detailed outline of the paper, with key messages and points for each paragraph. The author then used Claude to convert the outline to a paper using the following prompt: “Please read the outline in the file [filename] and write a paper in an academic style based on the outline.” The author’s review of the Claude generated paper found the paper to represent the author’s ideas as provided in the outline, and no other ideas were introduced by Claude.

Abstract

Exchange-traded index put options are, on average, priced above their risk-neutral fair value, reflecting a persistent premium collected by option sellers in exchange for providing downside protection. Yet the popularity of tail-risk strategies pioneered by managers such as Universa Investments suggests that a disciplined approach to put buying can be profitable under the right conditions. This paper examines three empirical regularities in the pricing and behavior of out-of-the-money index put options that bear on this question. First, the term structure of option time decay deviates systematically from the pattern predicted by a constant-volatility option pricing model, with short-dated options embedding richer implied volatility than longer-dated options. Second, the distribution of multi-month returns in the underlying exhibits a pronounced left tail that is largely absent from the distribution of single-month returns, a pattern more consistent with an asymmetric, leverage-driven relationship between price declines and subsequent volatility

than with simple directional momentum. Third, the price of out-of-the-money puts responds asymmetrically to moves in the underlying, rising more in percentage terms following a decline than they fall following an equivalent advance. Taken together, these findings suggest that put buyers may improve their expected returns by favoring longer-dated contracts and by actively managing positions rather than holding them to expiration, while put sellers may be better compensated concentrating at the shorter end of the maturity curve.

Keywords: options, implied volatility, put options, theta decay, tail risk, leverage effect, volatility term structure

1. Introduction

A substantial body of empirical work has found that exchange-traded put options are, on average, priced above their risk-neutral value: the average payoff of a put at expiration tends to fall short of the premium paid for it. This finding is often interpreted as compensation collected by option sellers for bearing tail and volatility risk that buyers are willing to pay to offload. If this premium is persistent, a naive buy-and-hold strategy in out-of-the-money puts should be expected to lose money over time, on average, even if it occasionally pays off spectacularly during a market decline.

Despite this headwind, put buying, and tail-risk hedging more broadly, has become increasingly fashionable among institutional and retail investors alike. Tail-risk funds such as Universa Investments have popularized the idea that a systematic allocation to out-of-the-money puts can be a valuable portfolio complement, not because the puts themselves are cheap, but because of how and when they are used. This raises a natural question for any investor considering a put-buying program: if puts are overpriced on average, how can an investor purchase them in a manner that produces a positive expected return, rather than simply paying a slow, steady insurance premium to the market?

This paper identifies several empirical trends in the pricing and behavior of out-of-the-money puts that are relevant to this question; we do not offer a complete trading strategy but instead insights which others may be able to use to construct a trading strategy. We examine the term structure of time decay across different maturities, the relationship between the shape of the return distribution of the underlying and the pricing of volatility, and the asymmetric response of put prices to moves in the underlying. Each of these trends carries practical implications for how a put buyer, or a put seller, might structure a position to improve its expected return.

2. Literature Review

[This section is intentionally left incomplete pending a targeted literature search. Candidate references identified to date relate to the leverage effect in equity volatility, including Black (1976) and French, Schwert, and Stambaugh (1987). Additional references on the pricing of variance risk, the term structure of implied volatility, and the empirical performance of put-buying and put-selling strategies will be incorporated once identified.]

3. Greater Fool Theory

The greater fool theory holds that an investor can profit from purchasing an overpriced asset provided that the asset can later be sold to another buyer, the “greater fool,” at a still higher price. As it applies to puts, this theory suggests that a put may be overpriced relative to the true probability that it expires in the money, and yet its owner may still profit by selling the put once other market participants are willing to pay an even greater premium for the same exposure. This is a materially different proposition from buying a put with the expectation of holding it to expiration and collecting an intrinsic payoff.

In practice, successful put buyers rarely hold their positions to expiration. Instead, they roll positions forward and wait for episodes in which the price of the put they hold spikes, whether because the underlying has declined, because implied volatility has risen, or both, and sell into that strength. Many successful put-buying programs close positions before the put has acquired any meaningful intrinsic value at all, monetizing the rise in extrinsic value rather than waiting for the option to settle.

This suggests a practical approach to put buying: initiate positions in out-of-the-money puts when implied volatility is low and be prepared to exit when the underlying declines and

market-priced volatility rises. In that scenario, two distinct forces push the price of the put upward simultaneously, the decline in the underlying itself and the accompanying spike in implied volatility, and a put buyer who is positioned to capture both effects, rather than waiting for the put to expire in the money, is better placed to realize a positive return net of the premium originally paid.

4. Theoretical Time Decay for Out-of-the-Money Options

It is a common misconception that time decay is gradual far from expiration and accelerates as expiration approaches. This characterization is accurate for at-the-money options only. It is not, however, a general property of options away from the money.

For sufficiently out-of-the-money puts, theory implies the opposite pattern: holding all else equal, these options should lose much of their value further from expiration, approaching expiration almost valueless. This has a direct practical implication for put sellers, many of whom focus on options with zero or near-zero days to expiration on the assumption that time decay is greatest on the day of expiration itself. For out-of-the-money contracts, this assumption does not hold, and the theoretical decay curve implies that options can be profitably sold well before expiration.




Figure 1: Theoretical decay of puts with different levels of out-of-the-moneyness.

5. Empirical Time Decay for Out-of-the-Money Options

However, the empirical time decay curves differ from the theoretical curves. While deep out of the money puts should lose money quickly far from expiration, the data show that the actual time decay is more gradual.




Figure 2: The empirical theta curve for options with strike prices 20% out of the money.




Figure 3: The empirical theta curve for options with strike prices 10% out of the money.

Note on Figures 2 and 3: Red markers denote observations priced during March 2020, when the pandemic's impact on stock prices, the uncertainty of the government's policy response, and the economic impact of the virus were all at their greatest. These are clear outliers but illustrate what can happen when a high degree of uncertainty, particularly over the short term, is priced into options. The

green lines represent a hypothetical linear rate of time decay and are included purely as a visual reference; they are not intended to represent the trend in the data.

Figures 2 and 3 present the price of out-of-the-money puts at various times to maturity relative to the price of a nine-month put, for strikes 20 percent and 10 percent out of the money, respectively. Each data point compares the prices of puts with one month, two months, three months, and so on to expiration against the nine-month put, where all of the puts being compared were priced on the same day and differ only in their expiration date. This same-day, cross-sectional comparison was repeated across many pricing dates to construct the empirical curves shown.

The resulting curves run counter to the pattern implied by the theoretical, constant-volatility decay curve presented in Figure 1. The direction of the deviation is consistent: shorter-dated puts are persistently priced richer, relative to the nine-month benchmark, than the theoretical curve would predict. Because the theoretical curve holds implied volatility fixed across all maturities while the empirical curves reflect whatever volatility the market actually prices into each maturity, the deviation between the two is attributable to the term structure of implied volatility rather than to time decay in the strict sense: the market does not price a single, flat level of volatility across the maturity curve; in fact, given not a single daily cross-section of the out of the money puts declined at the theoretical rate, we can conclude that in all cases the market priced more volatility into near-term contracts than into longer-dated ones.

Two, not mutually exclusive, explanations are worth considering for why the market persistently prices more risk into near-term options than into longer-dated ones. One possibility is behavioral: investors may simply be more attentive to, and more willing to pay for protection

against, risks that could materialize in the immediate future than risks of comparable magnitude further out in time. A second, more mechanical possibility is that the market structurally prices longer-dated options to imply lower volatility, independent of any change in the near-term outlook, so that the observed term structure reflects a persistent risk premium concentrated at the short end of the curve rather than a purely time-varying response to near-term news.

6. Tail Risk

The observation from the last section, that options have historically priced greater volatility into puts closer to expiration, begs the question of whether stocks exhibit greater tail risk over shorter periods than longer periods. In fact, relying on simple statistical theory, it would make sense that if there was a tail in the one-month return distribution that this tail would be smoothed out in the six-month distribution. Assuming the six-month return of the S&P 500 was simply the aggregate of six independent samples from the one-month return distribution, the Central Limit Theorem tells us that the tail would all but disappear in the 6-month return distribution. If so, it could be reasonable that the market exhibited greater concern for risk in shorter duration options.

However, historical returns do not support this idea. The one-month return distribution of the S&P 500 closely resembles a normal distribution. To the extent a tail exists at the one-month horizon, it appears to be small and not extreme. The six-month return distribution, by contrast, exhibits a clear left tail. If six-month returns were instead constructed by randomly sampling six one-month returns from the one-month distribution and compounding them, no comparable tail would emerge; the fact that one does indicates that monthly returns are not independent, identically distributed samples from a single, stable distribution.




Figure 4: Rolling six-month returns for SPY, January 2005 to May 2026.




Figure 5: Rolling one-month returns for SPY, January 2005 to May 2026.

Figures 4 and 5: The distribution of six-month returns shows a visible elevated left tail near the -32 percent mark. The distribution of one-month returns shows no comparably clear elevated tail risk. That a

tail emerges only over longer holding periods implies that monthly returns are not independent of one another, at least with respect to declines.

One candidate explanation is the leverage effect, whereby the volatility of a given month's return is influenced by the sign and magnitude of the prior month's return: a decline tends to be followed by a period of elevated volatility to a greater degree than an equivalent advance. This is distinct from momentum in the ordinary sense of the term. It does not require that a negative month makes a subsequent negative month more likely in direction; rather, it implies that the distribution from which the next month's return is effectively drawn widens following a decline, which is sufficient on its own to produce a fatter left tail once returns are compounded over multiple months. Figure 6 presents evidence consistent with this mechanism: a decline of more than five percent in month N-1 is associated with materially greater dispersion of returns in month N.




Figure 6: A negative return, particularly a decline greater than 5 percent, in month N-1 tends to be associated with greater dispersion of returns in month N.

This dependency structure has a direct implication for options positioning. A negative month creates an opportunity for put sellers as well as for put buyers who are positioned correctly: following a decline, the leverage effect causes implied volatility to rise, and the put itself moves closer to at-the-money, both of which increase its price. A seller who is willing to take on this risk after the volatility repricing has occurred, rather than before, may be better compensated for it; a buyer holding the position into the decline benefits from the same repricing.

7. Put Prices Rise Disproportionately During Stock Price Declines

Equity indices tend to rise in value more often than they decline, which creates a persistent headwind for anyone systematically buying puts. Several characteristics of the option chain, however, can partially offset this headwind and can be used to a put purchaser's advantage.

Options chains exhibit a degree of convexity: on a static option chain, out-of-the-money puts rise more in value when the market declines five percent than they lose when the market gains five percent. The magnitude of this asymmetry depends on where along the chain a given strike sits; the ratio of upside on a decline to downside on an advance varies meaningfully by strike placement.



Figure 7: The impact of a change in the price of the underlying on the price of a ten percent out-of-the-money SPY put, controlling for other variables. Prices are for puts expiring March 19, 2027, as priced on June 17, 2026. The graph illustrates the convexity in the payoff, with larger gains when the underlying declines than losses when the underlying rises by an equivalent amount.

A second, related factor compounds this effect. When the price of the underlying declines by five percent, implied volatility itself can spike, consistent with the leverage effect discussed above: volatility tends to rise more following a decline in the price of an asset than following a comparable advance, and the market prices this asymmetry into option prices, resulting in both a higher put price and a higher implied volatility. This spike in implied volatility compounds the increase in the put's value already driven by the decline in the underlying, amplifying the return of a put position during a sudden price decline. This effect is not permanent: the implied

volatility spike tends to fade once the downward trend ends or the catalyst for the decline is resolved, implying that the timing of the exit is important for put purchasers to generate returns.

As a result, buying a well-placed out-of-the-money put when implied volatility is at a low to normal level offers the potential for larger gains when the underlying declines than losses when the underlying rises. Because implied volatility is already low, it is unlikely to compress much further on a five percent rise in the underlying, but it may rise significantly on a five percent decline. Identifying puts that benefit from both sources of asymmetry, payoff convexity and the volatility response to a decline, can help identify contracts that are priced attractively enough to be worth owning.

8. Conclusion

Time decay in out-of-the-money puts does not follow the pattern implied by a constant-volatility theoretical benchmark. This deviation is attributable to the market pricing higher implied volatility into short-dated contracts than into longer-dated ones. Yet it is over longer durations that tail risk in the underlying actually materializes, which on its face might suggest that the market should price greater, not lesser, risk into longer-dated contracts.

The richness priced into short-dated puts appears to reflect compensation for immediate, event-driven risk, while the tail risk that emerges over multi-month horizons reflects the propagation of elevated volatility following a decline, a phenomenon that compounds over time rather than one that is fully priced in at the outset. Practically, this suggests that put buyers may be better served by longer-duration puts, which are priced with comparatively less risk premium while still capturing the tail risk that tends to build over multi-month horizons, whereas put

sellers may be better compensated selling shorter-duration puts, which command a richer premium relative to the risk realized over their shorter life.

Considering the greater fool framework discussed above, these findings suggest that a purely static, buy-and-hold approach to put purchasing is unlikely to be optimal. A put buyer who actively manages a position, favoring longer-dated contracts and being prepared to exit once a decline and its accompanying volatility spike have repriced the position favorably, is better positioned to capture a positive expected return than one who purchases a put with the intention of holding it to expiration.