FNCE685 AY25/26
Financial Risk Management

Does the Options Market Underreact to Firms' Left-Tail Risk?

<u>Group 2</u>

Kenneth Chan Jian Da
Olivia Tang Shu Hui
Tang En Ci Nathan
Randolph Lee Zheng Yang
Koh Thoong Yip

EXECUTIVE SUMMARY

Insurance markets are generally expected to price risk efficiently. In car insurance, higher-risk drivers pay higher premiums, ensuring that expected insurance payoffs are broadly similar across drivers. Financial markets are often assumed to operate in the same way: assets with greater downside risk should face more expensive crash insurance, implying similar expected insurance returns.

The paper "*Does the Options Market Underreact to Firms' Left-Tail Risk?" (2024)* challenges this intuition. Using US optionable stocks from 1996 to 2017, the authors find that crash insurance, proxied by bear put spreads, earns higher returns for firms with greater left-tail risk, measured by Value-at-Risk and Expected Shortfall. They attribute this anomaly to investor underreaction to volatility and the persistence of downside risk rather than compensation for systematic risk.

This paper reviews the paper's findings. First, we outline the theory and empirical evidence that the paper discusses. Second, we challenge the assumptions and evidence made by the paper. Third, we attempt to replicate the empirical results of the paper. Fourth, we discuss the potential implications of the paper for business owners and financial players.

The core of this paper is in the conduct of two "extensions" to the paper's findings: a simulated index-level crash insurance portfolio using S&P 500 data, and an analysis of long put strategies as an alternative form of crash insurance. While our finding that simulated crash insurance portfolios outperform the unhedged S&P 500 during high tail-risk periods supports the paper, our other finding that long-put strategies do not show increasing returns with left-tail risk, suggests the anomaly may be specific to bear-put spreads rather than to all crash insurance.

I. INTRODUCTION

Insurance markets are generally expected to price risk efficiently. In everyday contexts such as car insurance, individuals who are more likely to be involved in accidents are charged higher premiums, while safer drivers pay less. Although high-risk drivers pay more, the expected return on insurance coverage is broadly similar across drivers because higher accident probabilities are offset by higher premiums. In this sense, insurance pricing reflects both the likelihood and severity of losses, ensuring that riskier participants do not systematically receive better insurance deals.

Financial markets are often assumed to operate in a similar manner. Assets that are more prone to large price declines should face higher demand for downside protection, causing crash insurance to become more expensive and, in turn, yield lower expected returns. Under this logic, investors seeking to insure high-risk assets should pay a premium that fully reflects their greater exposure to crashes, just as reckless drivers pay more for car insurance. Consequently, the expected payoff from crash insurance should not increase with an asset's left-tail risk.

The paper "*Does the Options Market Underreact to Firms' Left-Tail Risk?*" (2024) explores whether such crash insurance in the financial markets is priced efficiently. In particular, they compare the cost of the "insurance premium" across assets based on their assessed riskiness. In so doing, they discover that there are supernormal returns on crash insurance for riskier assets. In simpler terms, this means that the "insurance profit" that is gained when the portfolio decreases in value is higher than expected for assets which are more likely to face crashes in their asset prices.

In technical terms, the authors found that bear put spread strategies on high left-tail-risk firms earn unexpectedly positive returns. This contradicts standard risk-pricing logic, which theorizes that bear put spread strategies return on high left-tail-risk firms should be small, given that they would be in high "demand". The authors further found that there were two key mechanisms for this phenomenon – the underreaction to price volatility and the underestimation of the persistence of downside risk.

The implications of this paper's findings are fourfold. First, for business owners, the results underscore the importance of adopting a proactive and data-driven approach to downside risk management, as market prices often underreact to rising left-tail risk, creating a false sense of security and increasing vulnerability to adverse shocks. Second, for traders, the findings suggest that deploying bear put spread strategies on firms with high left-tail risk may yield abnormally high risk-adjusted returns, as crash insurance tends to be systematically underpriced for such firms. Third, for portfolio managers, the results indicate that conventional risk management frameworks may underestimate true downside exposure, particularly when left-tail risk is persistent, implying a need for more dynamic and targeted hedging strategies. Fourth, for lenders, the evidence suggests that traditional credit assessments may underweight equity-based downside risk signals, potentially leading to mispriced credit risk and insufficient protection against borrower distress during extreme market conditions.

This paper reviews the paper "*Does the Options Market Underreact to Firms' Left-Tail Risk?*" (2024) in four parts. First, we outline the theory and empirical evidence that the paper discusses. Second, we challenge the assumptions and evidence made by the paper. Third, we attempt to replicate the empirical results of the paper on a separate paper. Fourth, we discuss the potential implications of the paper and potential areas for further research.

II. OVERVIEW OF LEFT-TAIL RISK AND RETURNS OF CRASH INSURANCE

A. LEFT-TAIL RISK

In the paper, left-tail-risk refers to the risk of experiencing extreme negative stock returns. The term left-tail-risk is derived from the distribution curve of the firm's returns over a single time period (e.g., daily), where the left-tail-risk is the risk of an extremely negative return over a single time period (i.e., a crash).

The authors use two metrics to measure a firm's left-tail-risk, or the probability of a large sharp downward drop in its prices as illustrated in **Figure 1**. First, the authors use the Value-At-Risk (VAR) metric. The VAR metric can be defined as the "maximum loss you can expect over a given time period, at a chosen market level, under normal market conditions". In simpler terms, the VAR refers to the amount an investor could lose with X% probability, provided that the distribution of the assets' returns do not change. This paper uses the VAR5 and VAR1, which correspond to the 5^{th} percentile and 1^{st} percentile of daily returns over the past 250 trading days of the assets used in the paper. A higher VAR implies a higher "left-tail-risk" in that the downward falls in asset prices, for a given probability (i.e. level of risk), is much larger.

Figure 1: VAR and ES Metric

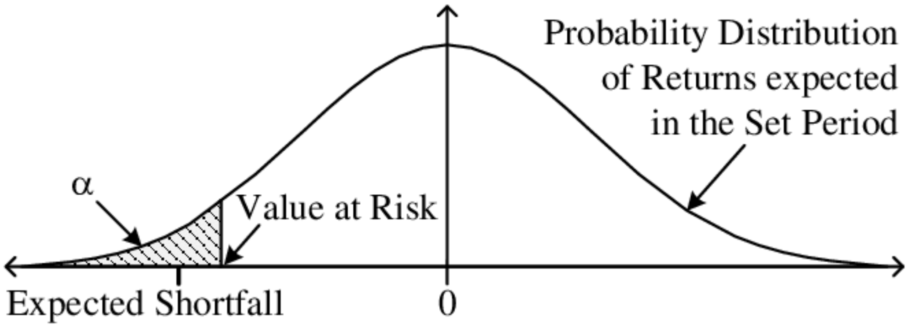

Second, the authors use the Expected Shortfall (ES) metric. This is also known as the Conditional VAR metric. In essence, the ES metric measures the losses on average in the worst X% of outcomes. While VAR shows the cutoff point for bad losses, the ES metric measures the average severity in the bad scenarios. In mathematical terms, the ES measures the area under the probability density function of returns over a single period, while the VAR measures the return at the specific probability within the same probability density function. The ES is arguably a better metric that the VAR metric because it captures how bad the worst losses actually are, and this captures catastrophic risks more closely. This paper uses the ES5 and ES1. Similarly to the VAR metric, a higher ES indicated a higher left-tail-risk.

B. CRASH INSURANCE

In this paper, crash insurance refers to insurance against sudden and drastic asset price declines. In simpler terms, it is a strategy that allows the buyer to profit when the asset price declines rapidly, thus providing insurance against price falls, as the overall return on the portfolio (including the crash insurance) would not fall as much as if there was no crash insurance.

There are many ways to obtain crash insurance. These include purchasing options, derivatives or even swaps. For example, a portfolio manager could purchase put options. These are options which give the manager the right, but not the obligation, to sell the asset at a specific strike price, within a specified period in the future. In the event that the asset price declines, the portfolio manager would be able to activate the option and sell the asset price at the specified strike price, thus reducing their losses. The portfolio manager could also directly sell the put option at a profit. Alternatively, the portfolio manager could also use volatility derivatives. Given that volatility typically spikes during crashes, the portfolio manager could purchase long volatility strategies. When volatility spikes, the portfolio manager would profit, which would offset the decline in price.

Specific to the paper, the authors measure crash insurance through bear put spread strategies. This involves buying an out of the money (OTM) put option with a relatively higher strike price and selling a deeper OTM put option with a relatively lower strike price. When combined, this ensures a positive payoff when the asset price falls, and a negative payoff when the asset price rises, thus providing a form of crash insurance as shown in **Figure 2**. Of note, this also ensures a cap in the positive and negative payoff when the asset price declines or rises respectively.

Figure 2: Payoff Structure for Bear Put Spread Strategy

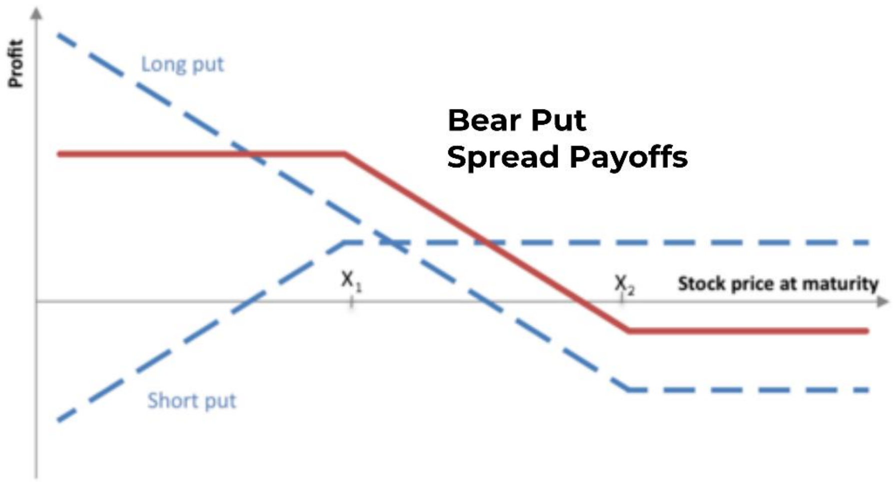

C. INTERPLAY BETWEEN LEFT-TAIL RISK AND CRASH INSURANCE

In theory, firms should have similar expected returns on their crash insurance regardless of the left-tail-risk that they have. This is because demand for crash insurance for assets with higher left-tail-risk will be higher. As such, the prices for this crash insurance will be bid up, and the returns on the insurance, when the asset price falls, would be smaller. This is analogous to car insurance, where drivers who have a higher risk of getting into an accident would pay a higher premium, and as such have a similar return on insurance to a driver who pays a smaller premium and has a smaller probability of getting into the similar accident.

In other words, there should not be higher returns for crash insurance for assets with a high amount of left-tail-risk, in the same way that there should not be a lower premium for drivers who are at a higher risk of getting into an accident. The paper, however, demonstrates that the empirical evidence contradicts this theory, by showing that firms with higher left-tail-risk do experience higher crash insurance returns.

III. DATA AND RESEARCH METHODOLOGY USED IN THE PAPER

A. DATA

The data used in the paper covers all US stocks from January 1996 to December 2017. To examine the relationship between the returns on crash insurance and left-tail-risk, the list of stocks were further reduced to only those with exchange-listed options. From there, the different data requirements were sourced from different datasets. To measure the left-tail-risk, the study used data from the Centre for Research in Security Prices (CRSP), which provided data on the stock prices, returns, market capitalisation and trading volume. To measure the returns on crash insurance, i.e., the returns on the bear put spread strategy, the study used data from OptionMetrics, which provided data on bid and ask prices for different options, the implied volatility, the open interest and volume. To ensure that their study was robust, the study also included control variables to improve comparability across firms and time periods. These were mostly firm related characteristics obtained from Compustat.

B. CONSTRUCTION OF VARIABLES

With the data, the study computed the respective variables. As elaborated above, the left-tail-risk was measured using VAR5, VAR1, ES5 and ES1. These were derived based on the distribution of returns over the past 250 days (i.e. over the past year) given that each year has 252 trading days.

The crash insurance returns were then computed first by constructing a bear put spread each month with a one month maturity, and then computing the returns on those bear put spreads. In total, the study computed 155,003 crash-insurance returns, which represented one for each stock-month combination in the sample.

Of note, there are two nuances in the construction of the bear put spread. First, the authors used midpoint prices. Because the study was back tested, it was not possible to determine whether the options would have been transacted at the bid or ask price. As such, the authors used the midpoint prices. Second, the authors used a fixed percentage to determine the appropriate strike price for the put options. This fixed percentage was the probability (based on the historical data at that time to avoid forward looking bias) that the option would finish in the money. That specific percentage was 30% for the OTM put option and 10% for the deeper OTM put option, meaning that the former had a 30% chance of ending in the money and the latter had a 10% chance of finishing in the money. This was to ensure consistency across stocks and time periods.

In addition to the variables of interest, the study also computed 19 control variables, comprising three variables to control for firm characteristics, six to control for stock return and trading characteristics, four to control for option characteristics and six to control for systematic factor exposures. This mainly ensured that the relationship between crash insurance returns and left-tail-risk was not confounded by characteristics related to the firm or the way that the asset or options were traded, or other risk factor exposures. These are shown in **Table 1**.

Table 1: List of Control Variables

Firm Characteristics	Trading Characteristics	Option Characteristics	Systematic Factor Exposures
• Size • Book-to-market • Leverage	• Momentum • Short-term reversal • Illiquidity • Idiosyncratic volatility • Skewness • Kurtosis	• Variance risk premium • Volatility-of-volatility • Risk-neutral skewness • Option demand	• βBEAR • βSTRAD • βJUMP • βVOL • βTAIL • βDOWN

C. METHODOLOGY AND KEY FINDINGS

To test the relationship between the returns on crash insurance with the firms' left tail risk, the authors used four key empirical methods. First, the authors conducted a simple univariate portfolio analysis. This involved sorting the stocks into deciles based on the different left tail risk measures (VAR and ES) and examining the spread in the returns of the bear put spread crash insurance between the highest-tail-risk groups versus the lowest-tail-risk groups. In this, the authors found that the returns of crash insurance for the highest-tail-risk groups were about 0.2% higher than those in the lowest-tail-risk groups on the VAR5 measure. The t-statistic was 4.58, which showed that it was highly statistically significant.

Second, the authors conducted a bivariate portfolio analysis. This was to safeguard against possible confounding effects in the univariate portfolio analysis, where a third variable could be causing higher left-tail-risk groups to have higher returns on crash insurance. To do this, the authors sorted the firms into the respective control variables (such as firm characteristics) first before sorting them by their left-tail-risk (similarly to the univariate portfolio analysis). In doing so, they again found that the returns on crash insurance on firms with the highest 10% of left tail risk was higher than the returns on crash insurance on firms with the lowest 10% of left tail risk. This result was statistically significant. Compared against the univariate portfolio analysis, this methodology suggested that there were no confounding variables causing the higher returns on crash insurance for firms with higher left tail risk.

The challenge with the portfolio analysis methodologies was that they only compared the top decile with the bottom decile of firms according to their left tail risk and could not demonstrate a general increase in returns on crash insurance in line with left tail risk. Hence, to do so, thirdly, the authors conducted a Fama-MacBeth regression. This was essentially a univariate and multivariate regression of crash insurance returns on the left tail risk measures (and the controls for the multivariate regression). In doing so, they found that the additional returns from an additional unit of left tail risk was about 16% (with controls) to 20% (without controls), with the t-statistic of 4.58 and 3.05 demonstrating that the result was statistically significant.

Fourth, the authors explored whether the difference in crash insurance returns could be explained by systematic factors – i.e., whether they could be explained by differences in how the firms with different left-tail-risk "benchmarked" the market. Their finding was that there was no statistically significant impact, which meant that the difference in crash insurance returns was not a result of inherent perceived "risk" of that particular asset.

To demonstrate the robustness of their results, the authors also conducted robustness tests in two areas. First, they compared the results across high volatility and low volatility periods as well across pre-crisis and post-crisis subsamples. Second, they tested different methods of constructing the crash insurance, such as using different delta thresholds, different at-the-money and out-of-the-money combinations, as well as non-delta-hedged spreads. Their findings remained consistent.

D. MECHANISMS

To explain their findings, the authors tested two hypotheses and found them to be plausible mechanisms for the empirical findings. First, the authors explored volatility underreaction. This mechanism argues that the crash insurance returns on firms with higher left-tail-risk is higher because the market underestimates the future increases in volatility after volatility begins to increase. In other words, after volatility increases, the market does not value the crash insurance sufficiently to "hedge" against future volatility increases. As such, the value of the crash insurance appreciated disproportionately. To prove this, the authors demonstrated that while firms with higher left-tail-risk today faced higher realised volatility and higher negative-return volatility in the future, the implied volatility demonstrated by the options adjusted much more slowly. This created an opportunity for abnormal returns on crash insurance.

Second, the authors explored underreaction to left-tail return momentum. This mechanism argues that when the stock price crashes, they often continue to crash. In simpler terms, if the stock price falls significantly today, it is likely to continue over the next few days. This is significant for crash insurance because investors who are aware of this phenomenon should immediately purchase crash insurance for their assets after a large fall in the asset price, and this would, in theory, result in similar returns on crash insurance for firms with different left-tail-risk. However, the authors find that the returns on crash insurance in fact increase after the stock price falls, which means that investors are "cashing in" (or in other words, undervaluing) their crash insurance for future crashes. This implies that they were underreacting to the persistence of left-tail risk.

IV. EVALUATING THE ASSUMPTIONS MADE BY THE AUTHORS

To construct and interpret their empirical results, the authors rely on several key assumptions. Although many of these assumptions align with standard practice in empirical asset pricing, each carries inherent limitations that warrant closer scrutiny. This section evaluates the major assumptions embedded in the paper and assesses the extent to which they are valid, empirically justified, and potentially restrictive.

A foundational assumption concerns the overall dataset, which consists of US optionable stocks from 1996 to 2017 and relies on asset-level analysis. While this long US sample provides depth and statistical power, it implicitly assumes that the pricing of crash insurance in this period is representative of broader market behavior. However, structural changes in options markets may limit the generalisability of the results to other time periods. It does not also guarantee that this phenomenon holds under all market conditions. Moreover, asset-level analysis abstracts from portfolio-level interactions, correlations, and hedging behavior that institutional investors typically face. As a result, the findings may not fully translate to real-world portfolio construction or to non-US markets.

A second major assumption is that historical distribution-based measures of left-tail risk, specifically VAR and ES calculated over the previous 250 trading days, accurately capture firms' true downside risk. These measures (VAR1, VAR5, ES1, ES5) are widely used in literature and allow the authors to maintain consistency with prior studies. However, it is unclear whether these results are generalizable to other forms of construction of left-tail risk. Additionally, nonparametric tail estimators are inherently noisy, particularly when applied to firms experiencing jumps, regime changes, or sparse trading activity. Given that tail events are rare, using a one-year window can produce unstable estimates that may introduce measurement error. The authors partially address this concern by demonstrating that the results are consistent across multiple alternative tail-risk metrics, thus increasing confidence in the assumption's empirical validity. Even so, the assumption remains reasonable but imperfect, and more sophisticated tail-estimation techniques could further strengthen the analysis.

Thirdly, the paper assumes that bear-put spreads are the main instrument for crash insurance. However, bear-put spreads are just but one of the ways of obtaining crash insurance. Hence, the paper's generalisability to other downside protection mechanisms may be limited. The construction of the bear-put spreads also depend on assumptions about option maturity, strike selection, liquidity, and bid–ask midpoint pricing, all of which can materially influence measured returns. Transaction costs, margin requirements, and execution frictions, especially for deep OTM options, may also further erode the practical effectiveness of bear spreads as an insurance instrument. Consequently, while bear put spreads provide a tradable and well-defined proxy for crash insurance in empirical analysis, they may not fully capture the economic properties or real-world implementation costs of alternative tail-risk hedging strategies such as dynamic option hedging, variance swaps, or portfolio-level protection.

Another important assumption is that the authors' extensive set of firm-level and option-level controls suffices to capture alternative drivers of bear put spread returns. Their models include variables such as size, book-to-market, leverage, momentum, idiosyncratic volatility, variance risk premium, risk-neutral skewness, and exposures to a battery of systematic factors. The breadth of control variables enhances confidence that omitted-variable bias is minimized.

However, no observational paper can eliminate all possible confounders. The interpretation that left-tail risk predicts crash-insurance returns independently of other firm characteristics is well supported, but the possibility of latent omitted risks cannot be entirely ruled out.

Finally, the authors assume that behavioral explanations, such as underreaction to volatility and tail-risk persistence, are superior to risk-based explanations. They justify this by showing that known factor models fail to explain the positive returns observed in high tail-risk portfolios. While persuasive, this assumption depends on the premise that the factor models they employ fully span all systematic sources of crash risk. Given that tail-risk factors are complex and potentially model-dependent, this conclusion, though reasonable, is suggestive rather than definitive.

In sum, the authors' assumptions are largely consistent with prevailing empirical-finance methodologies and are supported by extensive robustness checks. Yet several assumptions, particularly those involving hedging precision, left-tail measurement, and the completeness of factor models, remain only partially validated. While these limitations do not undermine the study's main contributions, they qualify the strength of its causal claims and point to opportunities for further refinement in future research.

V. TESTING THE ROBUSTNESS OF AUTHORS CONCLUSIONS ON DIFFERENT DATA SETS

To further ascertain the robustness of the authors' conclusions, we attempt to replicate the authors' conclusions in two different ways. These are broadly based on two of the above limitations of the study.

The first limitation of the study is that the crash insurance is determined at the asset level. In particular, the various methodologies, such as the univariate portfolio analysis, the bivariate portfolio analysis and the Fama-Macbeth regressions, are conducted at the asset level. This means that the VAR, the ES and the bear put spread returns were computed at the asset level. As such, strictly speaking, the findings in the paper should be only generalised to asset-level crash insurance, in that the returns on bear put spreads are higher for assets with higher left-tail-risk. It is perhaps true, but not necessarily true, that the same phenomenon would occur at the portfolio level. As such, the first way that we attempt to replicate the authors' conclusion is to replicate the authors' conclusion on index data, to see whether the same finding is true on index-level data, particularly on the S&P 500.

The second limitation of the study is that the crash insurance is determined mostly based on bear-put spread strategies. This is useful because it is the closest payoff structure to insurance, where you earn a fixed amount in the "bad scenario" and you lose a fixed amount in the "good scenario". There are, however, some limitations to this approach. In particular, it does not fully represent how normal investors, especially retain ones, approach managing downside risk. In such cases, buying, or going long, on a put, could represent a more natural way of hedging. In particular, one potential argument against the set-up of the bear put strategy could be that some of the positive returns stem from the trade being "undervalued" due to there being less demand, given that only professional investors are likely to take on this trade. A long put on the other hand, would likely be more "crowded" and has a lower probability of being "undervalued". As such, the second way that we attempt to replicate the authors' conclusion is to replicate the study on just long puts as crash insurance, to see whether the same finding is true in the relationship between buying a put, as well as the VAR and ES.

A. SIMULATED CRASH INSURANCE PORTFOLIO

In our first attempt to replicate the authors' empirical findings, we construct two simplified crash insurance portfolios to examine whether portfolio returns increase with left-tail risk. This exercise is intended as a conceptual replication rather than a full structural match to the original paper, allowing us to test the core intuition in a more tractable setting. We use S&P 500 index data, S&P 500 options data, and the VIX volatility index over a 10-year period. All data are obtained from publicly available sources on Kaggle.

To simplify the replication, we define crash insurance as a long position in the S&P 500 combined with a systematic put-option hedge that provides payoff during market drawdowns. Unlike the original study, which constructs delta-hedged bear spreads at the individual-stock level, our approach focuses on index-level crash protection, making the portfolio dynamics more transparent and easier to interpret. Specifically, crash insurance is implemented through OTM put options on the S&P 500 with a strike set at 85% of the prevailing index level, intended to approximate deep downside protection.

We construct two portfolios—an equal-weighted crash insurance portfolio and a volatility-weighted crash insurance portfolio—each with an initial capital of $100. Portfolio returns are simulated day by day using S&P 500 daily returns. For both portfolios, new put options are purchased monthly and held to expiry. The portfolios incur a recurring insurance premium, representing the cost of maintaining downside protection. When the S&P 500 falls below the option strike, the hedge generates a payoff that partially offsets equity losses. The resulting daily portfolio return reflects the underlying S&P 500 return adjusted for both insurance costs and option payoffs.

Table 2: Details of Portfolio Simulation

Simulation Date: 2 Jan 2015 – 19 Nov 2025 **Rebalancing Frequency**: Quarterly **Initial Hedge Ratio**: 30% **Target Strike Price**: 85% of Current S&P500 Price (To simulate deep OTM option) **Option Maturity**: 1 month **Rolling Schedule**: New OTM puts bought monthly held to expiry **Portfolio 1**: Equal Weighted Portfolio Hedge Ratio Constant 30%

Portfolio 2: Volatility Weighted Portfolio
Hedge ratio is adjusted quarterly based on VIX level. If VIX is over median, the hedge ratio will be at 50%. Otherwise, it will be at a constant 30%

As shown in the cumulative return plots in **Figure 3**, both crash insurance portfolios outperform the unhedged S&P 500 over the sample period. This suggests that systematic downside protection can meaningfully mitigate large drawdowns and improve long-term risk-adjusted performance, despite the recurring cost of insurance.

Figure 3: Cumulative Portfolio Returns Over Time

We then evaluate whether portfolio performance varies systematically with left-tail risk, measured using VAR5, VAR1, ES5, and ES1, consistent with the measures used in the original paper. To do this, we sort observations into deciles based on VAR5, VAR1, ES5, and ES1. The results as shown in **Figure 4** indicate that mean portfolio returns generally increase across higher left-tail risk deciles for both crash insurance portfolios.

In higher-risk environments, the hedged portfolios experience more stable returns and less severe downside, highlighting the protective value of crash insurance. In contrast, during low-risk periods, crash insurance portfolios underperform due to the drag imposed by insurance premiums, consistent with theoretical expectations. This pattern reinforces the notion that crash insurance primarily delivers value in adverse tail-risk states rather than in calm market conditions.

Figure 4: Mean Portfolio Returns by VAR and ES

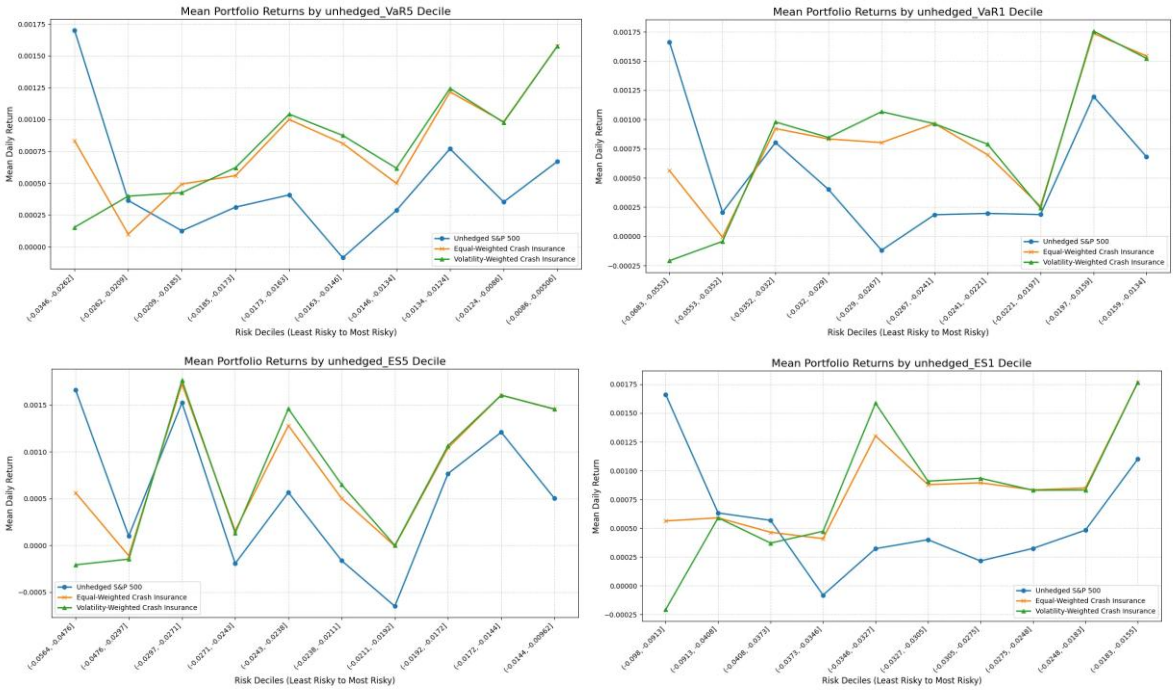

Overall, the unhedged S&P 500 exhibits a declining pattern in mean daily returns as left-tail risk increases. In contrast, both crash insurance portfolios perform significantly better in high-risk deciles, supporting the view that downside protection becomes more valuable precisely when tail risk is elevated. While this simplified replication does not fully capture the asset-level option strategies employed in the original paper, the results are directionally consistent with the authors' central finding: crash insurance delivers stronger performance in environments characterized by elevated left-tail risk.

B. LONG PUTS AS CRASH INSURANCE

In our second attempt to replicate the authors' conclusions, we study the use of long puts as crash insurance as opposed to bear-put spread as crash insurance. This is on the basis that a long put could also serve as crash insurance, albeit in a different way from a bear put spread, and on the basis that this could provide further insight into the authors' conclusions, given that a long put is less likely to be "undervalued" than a bear put spread, by virtue of the fact that it is simpler to understand and as such, more likely to be well valued, and up-bid. For this analysis, we use the US Option Level Output Data by the Wharton Research Data Services (WRDS). We also use the CRSP Daily Stock Security Data.

Due to the limited time for analysis, we simplify the analysis to focus on one single time period, to see whether the authors' conclusions still hold. In particular, we focus on a bunch of assets for which the options expire on 20 Nov 2020. This was a particularly volatile period for the financial market, given that the world was at the height of COVID-19 at that time. As such, this was a period where crash insurance would be highly valued. This is because there was always the risk of the COVID-19 situation deteriorating, more countries going into lockdown, and a sharp decline in asset prices. We did not choose the period of Mar 2020, despite that being the period where asset prices declined the most. On the other hand, we chose Nov 2020, because that was broadly one year after the COVID-19 pandemic began in Dec 2019, and the left-tail-risk calculated using the previous year's data would mostly be from the COVID-19 period, while the uncertainty in terms of whether asset prices could decline was high.

As such, we started by downloading all the option prices available within the US Options Level Output Data by WRDS in 2020. Following this, we also downloaded the CRSP Daily Stock Security Data in 2019 and 2020. From there, we obtained 12.9 million data points for options level data and 16.0 million data points for daily stock security data. From there, we further streamlined the options data by focusing only on options which expired on 20 Nov 2020. This reduced the dataset to 599,603 data points, consisting of option data from 12 Oct 2020 to 30 Oct 2020. We also computed the VAR5 and ES5 of the various securities from the daily stock security data. This was calculated based on the past year, from 14 Oct 2019 to 14 Oct 2020. This gave us 8,405 data points. From there, we further merged both data sets, using the SECID – PERMNO conversion to identify each unique ticker, with SECID corresponding to the identifier used by Options Metrics and PERMNO corresponding to the identifier used by CRSP. We also further focused on the put prices only on 12 Oct 2020 and 30 Oct 2020, with the difference being the dependent variable – i.e., the appreciation or depreciation of the put option.

In sum, this gave us a set of independent variables corresponding to the left-tail-risk, i.e. the VAR5 and ES5 based on historical data over the past year, and the dependent variable, which is the return on a put purchased on 12 Oct 2020 and sold on 30 Oct 2020, both of which has an expiration date of 20 Nov 2020. Due to a lack of time and to keep the analysis simple, we did not apply the delta hedge methodology to calculate the appropriate strike price. Instead, we used all the strike prices to calculate the returns on a put. This is not a perfect method, but it managed to keep the analysis simple. In total, we had 2,421 data points.

With this dataset, we proceeded to execute two forms of analysis. First, we did a univariate portfolio analysis, as per what the authors did in the paper. On this, we sorted the 2,421 data points based on their VAR as well as their ES. From there, we computed the average returns for a long put within the top decile of VAR and ES (i.e., assets with the highest left-tail-risk) and a long put within the bottom decile of VAR and ES (i.e., assets with the lowest left-tail-risk). The results from this analysis are as shown in **Figure 5**.

Surprisingly, the average returns for assets with a higher left-tail risk were lower (or more negative) than the returns for assets with a lower left-tail-risk. This was about 200% worse off for the top decile as compared to the bottom decile. This contradicts the findings of the paper, which should have shown that the returns for crash insurance for firms with higher left tail risk should have been higher.

Figure 5: Results of Univariate Portfolio Analysis

```
=== Test 1: Sorted by VaR_0.05 (descending) ===
Average return (top decile):    -2.475267
Average return (bottom decile): -0.274783
Difference (top - bottom):      -2.200485

=== Test 2: Sorted by ES_0.05 (descending) ===
Average return (top decile):    -2.982624
Average return (bottom decile): -0.494484
Difference (top - bottom):      -2.488140
```

Second, we did a Fama-Macbeth regression, by simply running a cross-sectional regression, with the VAR and ES variables as the independent variables, and the return on the long put as the dependent variable. The results are as shown in **Figure 6** and **Figure 7**.

Figure 6: Results of Fama-Macbeth Regression (VAR)

```
=== Regression 1: Returns on VaR_0.05 ===
                            OLS Regression Results
==============================================================================
Dep. Variable:                returns   R-squared:                       0.000
Model:                            OLS   Adj. R-squared:                 -0.000
Method:                 Least Squares   F-statistic:                    0.8226
Date:                Mon, 01 Dec 2025   Prob (F-statistic):              0.365
Time:                        21:36:57   Log-Likelihood:                -14239.
No. Observations:                2416   AIC:                         2.848e+04
Df Residuals:                    2414   BIC:                         2.849e+04
Df Model:                           1
Covariance Type:            nonrobust
==============================================================================
                 coef    std err          t      P>|t|      [0.025      0.975]
------------------------------------------------------------------------------
const         -1.5590      5.269     -0.296      0.767     -11.891       8.773
VaR_0.05     -74.8289     82.506     -0.907      0.365    -236.618      86.960
==============================================================================
```

Figure 7: Results of Fama-Macbeth Regression (ES)

```
=== Regression 2: Returns on ES_0.05 ===
                            OLS Regression Results
==============================================================================
Dep. Variable:                returns   R-squared:                       0.000
Model:                            OLS   Adj. R-squared:                  0.000
Method:                 Least Squares   F-statistic:                     1.012
Date:                Mon, 01 Dec 2025   Prob (F-statistic):              0.314
Time:                        21:36:57   Log-Likelihood:                -14239.
No. Observations:                2416   AIC:                         2.848e+04
Df Residuals:                    2414   BIC:                         2.849e+04
Df Model:                           1
Covariance Type:            nonrobust
==============================================================================
                 coef    std err          t      P>|t|      [0.025      0.975]
------------------------------------------------------------------------------
const         -2.1583      5.370     -0.402      0.688     -12.688       8.371
ES_0.05      -51.8665     51.548     -1.006      0.314    -152.949      49.216
==============================================================================
```

Overall, while not statistically significant, both regressions suggest that there is a negative correlation between left tail risk and returns on a long-put strategy. This seems to contradict the findings by the authors that assets with a higher left-tail-risk would have higher crash insurance returns.

While we have attempted to relook at the authors' findings, we were not able to establish any statistically significant contradiction to the paper, and the paper's conclusions must be taken to be sound and credible as they are supported by a comprehensive set of robustness tests. This is given that the authors validate their results across multiple left-tail risk measures (VAR and ES), alternative bear put spread constructions, different portfolio-weighting schemes, and static and daily delta-hedging methods, and have demonstrated that the positive relation between left-tail risk and future crash-insurance returns persists after controlling for an extensive list of firm characteristics, option-based variables, systematic equity factors, and systematic left-tail-risk factors. This breadth of robustness checks, spanning data, methodology, risk controls, and alternative explanations, provides strong evidence that the paper's conclusions reflect a genuine and persistent market underreaction to crash risk. Thus, the paper's findings remain highly persuasive and carry significant strategic implications to guide business owners.

VI. IMPLICATIONS OF PAPER'S FINDINGS ON BUSINESS OWNERS

In essence, the paper's findings suggest that it is prudent for business owners to adopt a proactive approach to risk management that is forward-looking and sensitive to early warning signals. There are three significant implications for business owners as follows.

A. OWNERS SHOULD NOT TRUST CALM MARKETS AND MUST ACTIVELY ACT BEFORE THE MARKET REALIZES THE RISK

First, business owners need to realise that stable markets can conceal the true state of escalating risk. The paper shows that markets consistently underprice crash insurance precisely when a firm's left-tail risk is rising, thereby creating a misleading picture where owners may erroneously interpret calm market conditions as an indication that all is well.

For businesses, it is inherently flawed to rely on market signals to assess exposure to left-tail risks. Warning signs, such as deteriorating fundamentals, widening uncertainty and persistent negative trends, often accumulate well before investors recognize their severity. The result is a perilous window of time in which a firm's actual vulnerability is far greater than what market-based measures imply.

Since crash insurance is underpriced during this period, inaction can be costly. Owners who fail to hedge early effectively gamble against the statistical tendency of markets to underreact, leaving themselves exposed when risk finally re-prices. Conversely, those who obtain crash insurance while it is still underpriced find themselves better positioned not only to survive adverse shocks but also to capitalize on competitors who were unprepared. Therefore, crash insurance becomes not only a defensive measure, but also an opportunity for strategic advantage.

B. THE PERSISTENCE OF LEFT-TAIL RISK DEMANDS PROACTIVE MANAGEMENT

Second, business owners must proactively manage their left-tail risk given the persistence of such risk. As illustrated in the paper, 77% of firms in the highest left-tail risk decile remain in the top two deciles even a year later. The persistence of such risk matters as it signals that negative signals rarely resolve quickly, but are indicative of an ongoing structural vulnerability.

For business owners, the persistence of left-tail risk transforms what might appear to be "temporary noise" into critical information. Small warning signs, such as a stock price nearing its 52-week low, increasing volatility or sustained margin compression, are often early indicators of deeper and enduring fragility. Owners who ignore these signals or attribute them to short-term fluctuations underestimate how long problems can endure and how costly delayed action can become.

The persistence of left-tail risks also affects external perceptions. Investors reward firms that demonstrate strong tail-risk management with greater confidence, which leads to lower financing costs and improved access to capital. Conversely, persistent tail risk, even if not immediately obvious to the market, ultimately erodes trust and raises borrowing costs when conditions worsen. Thus, proactive management of left-tail risk strengthens both operational resilience and financial credibility.

C. BUSINESS OWNERS SHOULD USE CRASH INSURANCE PROACTIVELY, NOT REACTIVELY, AS IT IS THE MOST VALUABLE WHEN THE MARKET IS STILL UNDERPRICING IT

Third, business owners should use crash insurance proactively because it is most valuable precisely when risk is not yet visible and the market is still underpricing downside protection. The paper's findings suggest that the highest returns to crash insurance occur during periods when investors are complacent, i.e., when sentiment is optimistic, volatility appears low, and warning signs have not yet been incorporated into market prices. For business owners, this means that waiting until risk becomes obvious is financially costly. Once the market recognizes the mounting risk, the price of protection surges, leaving late-responding firms with higher hedging costs and less effective risk coverage.

As crash insurance yields the highest returns when markets overlook or underestimate future risk, the practical implication for business owners is that hedging is most valuable when it feels least necessary. Therefore, owners should treat hedging as a core component of long-run risk management. The evidence in the paper shows that downside protection performs best when purchased before sentiment shifts, before volatility returns, and before the market wakes up to real left-tail vulnerability. This counterintuitive timing, hedging in calm periods rather than during stress, is exactly what creates the opportunity.

Under this proactive mindset, crash insurance evolves from a defensive safeguard into an offensive strategic instrument. Firms that lock in protection early are not only shielding themselves from losses but also positioning themselves to operate from a position of strength when downturns emerge. Such firms can maintain investment plans, avoid fire-sale behavior, negotiate better financing terms, and even acquire distressed competitors whose lack of preparation leaves them exposed. In this way, crash insurance converts market underreaction into a competitive advantage. When firms secure protection early, they effectively "lock in" the benefits of market underreaction. By the time risk becomes visible and hedging becomes expensive, proactive firms have already insulated themselves.

In sum, proactive crash-risk management not only preserves value during crises but enhances a firm's ability to thrive when competitors are weakened, making it a powerful component of strategic planning rather than a marginal financial tactic.

VII. PROPOSED LEFT-TAIL RISK ASSESSMENT FRAMEWORK

For business owners to proactively manage left-tail risk, they need an organized approach to identify, understand, and mitigate catastrophic exposures. We have developed a four-stage left-tail risk assessment framework to help owners recognize vulnerabilities early, act before the market catches up to reality, and use crash insurance proactively as a strategic advantage instead of a last-minute defense.

A. STAGE 1: PERFORM LEFT-TAIL RISK STRESS TESTING

The first stage in addressing markets' consistency in underreacting firms' left-tail risk is for businesses to identify their own weak points long before they become visible externally. A comprehensive stress-testing process allows firms to assess both operational and financial exposures that could generate left-tail outcomes.

Operational tail risks may include supply chain fragility, excessive regulatory exposure, reliance on single points of failure, or customer-concentration risk. Financial tail risks involve leverage levels, liquidity buffers, and the ability to refinance short-term obligations under stress. Firms should also monitor market-based signals such as industry volatility, competitor tail events, and abrupt sentiment changes that may foreshadow broader instability.

At this stage, business owners can take concrete actions to stress test left-tail risk, such as running worst-case cash-flow scenarios, stress-testing liquidity and debt structures, modeling substantial revenue shocks to assess the business's resilience, and evaluating operational vulnerabilities that could amplify an unexpected downturn. By simulating severe but plausible scenarios, owners gain a realistic sense of whether their business could withstand a catastrophic event. This early diagnostic work provides the foundation for effective crash insurance and strategic preparation.

B. STAGE 2: IDENTIFY SIGNS OF MARKET UNDERREACTION

The second stage is for business owners to assess if current market conditions are masking the firm's true vulnerabilities. This requires assessing whether the firm is in a phase where bad news travels slowly or where investors and counterparties are overly optimistic.

Business owners can consider key questions like: (1) is the business currently benefiting from hype cycles, rapid growth, or unusually strong investor enthusiasm that may suppress perceived risk? (2) Has internal risk quietly increased, perhaps due to rising leverage, new operational exposures, or an untested product launch, without corresponding market awareness? (3) Are favorable external conditions, such as cheap credit or strong industry sentiment, concealing deeper fragility inside the business model?

If the answer to any of the above questions is yes, the firm is likely operating in an environment where risk is real but underpriced. This mismatch between true exposure and market perception creates a window of opportunity, where crash insurance strategies can be deployed when they are still inexpensive and effective. Recognizing market underreaction early enables businesses to act before rising volatility causes protection costs to spike.

C. STAGE 3: IDENTIFY SIGNS OF MARKET UNDERREACTION

Once risks are identified, the third stage is for business owners to determine whether they can mitigate them internally, hedge them financially, or transfer them externally.

Internal mitigation focuses on reducing the underlying vulnerability arising from the firm's business model. This may involve diversifying suppliers to reduce dependency, broadening the customer base to avoid concentration risk, or redesigning the operational architecture to eliminate fragile nodes. For risks originating in the business model itself, mitigation is often the most direct and durable form of protection.

Hedging is essential when financial exposures can be insulated using market instruments. As downside protection is typically cheaper than it should be for high-left-tail-risk firms, owners have a unique advantage as they can hedge using underpriced crash insurance such as put options or bear spreads to gain cost-effective protection against severe downturns. Proactive hedging ensures resilience well before the market corrects its underreaction.

Finally, risk transfer involves shifting tail exposures to other third parties through contractual mechanisms such as indemnity clauses and force majeure provisions or outsourcing high-volatility operations. This provides structural insulation for risks that are difficult or costly to mitigate internally.

D. STAGE 4: MONITOR "TAIL PERSISTENCE" AND EARLY WARNING SIGNALS

Finally, to address the high persistency of left-tail risk, business owners should monitor "tail persistence" by implementing ongoing monitoring systems such as monthly dashboards that track internal metrics such as cash burn rate, short-term debt proportions, and customer concentration, alongside external indicators such as industry volatility, credit spreads, and commodity or FX movements.

Business owners should also go beyond simple monitoring and establish automatic response triggers tied to measurable "tail scores" or a crash-severity index. These triggers function as pre-committed action points in that when specific indicators push the tail score above defined thresholds, the business automatically shifts into a higher defensive posture. This may include activating hedging programs, tightening working-capital controls, pausing discretionary spending, or accelerating contingency plans. Pre-defined triggers reduce the influence of emotion and managerial delay during early stages of a downturn, ensuring the firm responds while protection is still inexpensive and before conditions worsen. By embedding these automated thresholds into their risk dashboard, owners create a disciplined, systematic framework that converts early warning signals into timely and value-preserving action.

VIII. OTHER IMPLICATIONS

Beyond implications for business owners, the findings are also relevant for financial practitioners in at least the following three ways.

First, for traders, the findings suggest that purchasing bear put spread strategies on firms with high left-tail risk may generate abnormally high returns relative to the risk assumed. If crash insurance is systematically underpriced for high tail-risk assets, traders who specialise in options and tail-risk strategies can exploit this mispricing by selectively targeting securities with elevated downside risk. Importantly, this does not imply unconditional profitability; bear spreads still carry negative carry during normal market conditions. However, by conditioning trades on firm-specific left-tail risk, traders may improve the risk–return profile of crash insurance strategies and enhance portfolio returns during market stress periods.

Second, for portfolio managers, the results imply that conventional risk management frameworks may underestimate the true downside risk embedded in portfolios, particularly during periods of elevated tail risk. Standard volatility-based measures and diversification assumptions may fail to capture persistent left-tail exposure, leading to insufficient hedging. The findings suggest that portfolio managers should consider scaling hedging intensity dynamically in response to measures of tail risk rather than relying on static or volatility-only hedging rules. Increasing the level of downside protection for assets or portfolios with elevated left-tail risk may improve drawdown control and risk-adjusted performance, even at the cost of higher insurance premiums.

Third, for lenders, the findings suggest that traditional credit risk assessments may underweight information contained in equity market tail-risk measures. Persistent left-tail risk in a firm's stock price can signal heightened vulnerability to extreme downside events, which may not be fully reflected in standard leverage ratios, credit ratings, or historical volatility measures. If lenders underreact to such signals, they may misprice credit risk or extend financing on terms that inadequately compensate for potential losses during stress periods. Incorporating equity-based downside risk metrics into credit evaluation and covenant design could therefore improve risk pricing and help lenders better anticipate deterioration in borrowers' financial conditions during adverse market environments.

IX. CONCLUSION

This report set out to evaluate whether crash insurance in financial markets is priced in a manner consistent with standard risk-pricing intuition by reviewing, critiquing, and extending the findings of "*Does the Options Market Underreact to Firms' Left-Tail Risk?*" (2024). Beyond summarising the original paper, we examined its core assumptions and tested the robustness of its conclusions using alternative data and implementations.

First, we reviewed the paper's empirical framework, including its measurement of left-tail risk using VAR and ES, its construction of crash insurance through bear put spread strategies, and its use of portfolio sorts and Fama–MacBeth regressions. We then evaluated key assumptions related to the dataset at asset-level, the construction of left-tail risk, and the generalisability of bear put spreads as a proxy for crash insurance.

To address these limitations, we conducted two extensions. The first involved constructing simulated index-level crash insurance portfolios using S&P 500 data to test whether the relationship between tail risk and crash-insurance performance holds at the portfolio level. The results show that systematic downside protection performs relatively better during periods of elevated tail risk, broadly aligning with the paper's intuition. The second extension examined long put strategies as an alternative form of crash insurance during a high-volatility period. In contrast, this analysis does not show increasing insurance returns with higher left-tail risk.

Regardless, the paper's conclusions can be taken as credible within the specific context of bear put spread strategies. For practitioners, the results highlight the importance of understanding how different hedging instruments embed downside risk, rather than assuming uniform pricing across strategies. For business owners, portfolio managers, and lenders, the findings reinforce the value of proactive, data-driven tail-risk management, while cautioning against overreliance on volatility-based signals or single hedging approaches. Overall, the evidence points to a nuanced view of crash-risk pricing – one that depends critically on implementation choices and warrants further investigation.

Access to Datasets and Notebooks:
https://drive.google.com/drive/folders/1yKUuRvzfL2us_OwOFM59M53Swi9Elyr5?usp=sharing