



TO CATCH A FALLING KNIFE

What Happens After Stocks Fall Sharply?

Iain Law, Aakash Shetty, Bramantio Utomo and
Srinivasan Vaidyanathan

Executive Summary

The adage “never catch a falling knife” warns against buying stocks in free fall, yet decades of long-term return reversal literature suggests that severely depressed stocks may outperform and turn out to be wise purchases.

This study tests that tension using U.S. and global equity data spanning 1985-2024, examining 3,919 U.S. stocks and 15 additional international markets, together comprising over 12,700 securities and more than 19,000 instances of a stock falling 60% or more from its prior peak.

Analyzing these fallen stocks on 1-, 3-, and 5-year benchmark-adjusted performance from their troughs within the United States, the falling knives tended to outperform their respective benchmarks at a rate of more than 50%.

This positive finding, however, reflects an attribute of a look-ahead bias where investors in real-time would not have been able to tell where that trough occurred. We addressed this problem by standardizing the outperformance measurement from the date when the stocks fell 60%. Using this new standardization, we found the stocks did not outperform the benchmark more than 50% of the time.

This result motivates us to understand the reason behind *why* certain stocks rebounded and others did not. Using a sequence of increasingly rigorous logistic regression and gradient-boosting models, we found that momentum variables were the most important factors to determine if a stock would rebound in the United States.

In summary:

- a slow price decline was generally a better sign than a sudden collapse
- a stock that had already suffered an extremely deep decline was less likely to recover
- the path of the crash mattered more than the stock’s fundamentals, including valuation, at the time

The Authors

Iain Law, Aakash Shetty, Bramantio Utomo and Srinivasan Vaidyanathan graduated with Master of Quantitative Finance degrees from UC San Diego’s Rady School in June 2026.

As part of their “Capstone Class” during the Summer Quarter prior to graduation, they revisited a study on “falling knives” published by the Brandes Institute in 2004—extended the period under review and added new angles of analysis.

Their methods and conclusions are reflected in this report. The Brandes Center thanks them for their extensive and detailed work.

At a group level, developed-market recovery rates increased modestly from the 1-year to the 5-year horizon, whereas emerging-market rates declined. But this pattern was not universal, and the developed-versus-emerging result should therefore be interpreted as an aggregate tendency rather than a rule that applies to every constituent market.

Taken together, the results offer a qualified rebuttal to the “falling knife” adage as buying indiscriminately at the point of severe distress was not a winning strategy on a risk-adjusted basis.

At the same time, our results show the crash itself may contain enough information to meaningfully separate future winners from losers. Investors may catch knives, but they must do it selectively.

1. Introduction

In this study, we challenge the investing wisdom of never catching a falling knife by exploring what happens when investors actually choose to ignore this suggestion and purchase stocks that have fallen 60% or more from their peak. We used data from 1985 to 2024 within the United States and global equity markets and build on work previously done by the Brandes Institute, which investigated stock and bond prices worldwide that had fallen sharply.

The primary objective of this project is to examine whether stocks that have fallen sharply subsequently outperform their corresponding benchmark indexes over the following 1-, 3-, and 5-year periods. Specifically, we investigate why some stocks

that experience substantial drawdowns of at least -60% in a 12-month window recover and rebound strongly, while others continue to decline. This rebounding phenomenon, known as long-term return reversal, aligns with principles of contrarian investing and stands in contrast to the momentum effect.

Our research also seeks to identify the characteristics of stocks that successfully recovered following severe drawdowns and compare them with those that failed to rebound. In particular, we investigate whether common factors such as industry/sector, company fundamentals, and macroeconomic conditions influenced the likelihood and speed of recovery. We also examine whether the rebound dynamics differed across markets, specifically between U.S. and non-U.S. equities. Furthermore, we analyze whether stocks with higher market betas were more likely to recover after large drawdowns, given their greater sensitivity to overall market movements. By identifying the factors associated with successful recoveries, this study aims to provide insights into the potential drivers of long-term stock reversals and contribute to the broader understanding of contrarian investment strategies.

2. Literature Review

2.1. Long-Term Return Reversal/Contrarian Investing

The theoretical framework for a falling-knife strategy dates back to De Bondt and Thaler’s (1985) findings, which, based on CRSP (Center for Research in Security Prices) monthly return data, show that stock prices are consistent with an overreaction hypothesis rooted in the fundamental psychological idea

that individuals tend to overreact to unexpected and dramatic news in the media. They found the portfolios of significant losers in the past substantially outperformed portfolios of significant winners in the past, holding consistent over multi-year holding periods; this is also known as long-term return reversal. This fundamental overreaction notion drives the falling-knife hypothesis: since investors systematically overreact to negative news, stocks that have fallen dramatically should, on average, be relatively underpriced and should make a mean-reverting recovery.

De Bondt and Thaler's ideas of reversal, however, are consistently contested across the investment literature. The reversal effect is in direct contrast to Jegadeesh and Titman's (1993) work that demonstrates buying past winners and selling past losers creates significant positive returns over 3- to 12-month holding periods, also known as a momentum effect. These contradictions are generally resolved on a timeframe horizon level: momentum (Jegadeesh and Titman) dominates at a shorter horizon (3-12 months), while reversal (De Bondt and Thaler) dominates at a longer timeframe (3-5 years). This distinction is directly relevant in our own holding period research design, leading us to analyze returns over 1-, 3-, and 5-year horizons.

2.2 Falling-Knife/Distress-Recovery Literature

Our study is framed around research from the Brandes Institute (BI), the former research arm of Brandes Investment Partners. The BI's 2004 study, "Falling Knives Around the World," defined a falling knife as a stock that has fallen 60% or more within a

12-month period. The BI tested whether buying stocks after the decline generated excess returns relative to the benchmark. Furthermore, another 2007 study by the Brandes Institute, "Fixed Income Falling Knives Phase Two: Examining the Relationship Between Issuer-Specific Bond and Equity Returns," further examined issues related to fixed income securities. It studied whether issuer-specific bond price declines hold similar predictive implications for future equity performance, testing whether the falling-knife phenomenon holds true in credit markets. This 2007 study extended the work done in the original 2004 study across different asset classes and offers distinct evidence for the falling-knife effect outside of equities.

The most significant findings from the Brandes Institute work are that while there was an average outperformance of falling knives, the median result indeed underperformed the benchmarks. In essence, a small number of issues resulted in extremely significant outperformance that skewed the distribution. This leads us to believe that when considering post-drawdown returns, skewness is a significant factor, as a small subset of stocks may drive a large proportion of the recovery.

In contrast to the BI work, a study by the Flossbach von Storch Research Institute used a different approach when defining a falling knife, one where a stock suffers a price decline of 15% or more within five trading days rather than a 12-month drawdown period. They studied S&P 500 stocks over a 20-year window that retained an inflation-adjusted market capitalization of about \$7.5 billion pre-crash.

Falling knife literature has not converged on a singular methodological definition of a falling knife in terms of its drawdown threshold or window length.

In more recent years, empirical work has been done to narrow and stress-test the falling-knife effect, rather than confirming its existence. A 2026 study by Kusumahadi and Madyakusumawati (2026) used Indonesian equity data from 2010-2020 and identified 194 individual falling-knife observations and found that return volatility has a negative impact on stock recovery probability. However, the magnitude of price decline had no statistical effect on recovery probability. Importantly, only 22.7% of falling knives in the sample consistently generated positive returns over a 3-year holding period.

This fundamentally reinforces the skewness pattern as noted in Brandes' work and suggests that the average outperformance, but poor median outcome, may be generalized beyond U.S. markets. Such findings led us to analyze a broad range of developed and emerging countries across the globe.

...while there was an average outperformance of falling knives, the median result underperformed the benchmarks. In essence, a small number of issues resulted in extremely significant outperformance that skewed the distribution.

3. Hypotheses

Based on prior empirical research on long-term return reversals and contrarian investing, we expect that stocks experiencing substantial price declines will, on average, outperform selected benchmark indexes over a 1-, 3-, and 5-year holding period. However, we also expect the distribution of post-drawdown returns to be highly skewed, with only a relatively small subset of losing stocks generating exceptional recoveries, while many continue to underperform or experience further declines.

We anticipate that the likelihood and magnitude of a stock's potential recovery will depend on several firm-specific characteristics. In particular, companies with stronger financial fundamentals, more attractive valuations, smaller market capitalizations, and healthier balance sheets are expected to have a higher probability of recovering following large drawdowns. We also expect industry or sector-specific factors to influence recovery patterns, as certain sectors may be better positioned to benefit from economic or market cycles than others. In addition to firm-level characteristics, we expect macroeconomic conditions to play a significant role in determining recovery outcomes.

Specifically, we hypothesize that stocks are more likely to rebound during periods of economic expansion and market recovery than during recessions or periods of heightened uncertainty. Furthermore, we expect differences in recovery behavior across geographic markets with U.S. and non-U.S. equities potentially exhibiting distinct rebound characteristics due to differences in

market structure, investor behavior, and economic conditions. Finally, we expect stocks with higher market betas, reflecting greater sensitivity to overall market movements, to recover more rapidly during market upturns than lower-beta stocks. Collectively, these findings are expected to provide a deeper understanding of the factors that distinguish stocks that successfully recovered from severe drawdowns from those that continued to underperform, thereby offering valuable insights into the effectiveness and limitations of contrarian investment strategies.

4. Data & Variables

4.1 Defining Falling Knife Variables

As previously noted, we define a falling-knife stock as one experiencing a drawdown of 60% or more within a trailing 12-month window. For each security, we compute a trailing 12-month running maximum of its adjusted price (13 month-end observations, inclusive of the current month) and calculate the drawdown at month t as the ratio of price to that running peak, minus one. A security is flagged whenever this drawdown reaches 60% or worse, and a stock must have at least six months of prior price history to be eligible for flagging.

Due to the reference peak being a trailing-window maximum instead of a fixed historical high, subsequent price recovery does not retroactively erase a flag. Moreover, a slow, sustained decline that never drops 60% within 12 months is excluded. Consecutive flagged months are grouped into a single episode; a stock must remain clear of the 60% threshold for three

consecutive months before a new breach is counted as a separate episode. Each episode is summarized by four reference points: pre-crash peak; breach date (the first month the 60% line is crossed, serving as event time 0); trough; and recovery date (the first month the price returns to the pre-crash peak).

Because firms that head toward bankruptcy often stop reporting prices before a full decline is reflected, a pure price screen understates incidences of failure. Therefore, we supplement the price screen with the CRSP delisting file, treating a delisting with code 400–599 (liquidation, or dropped for deficiency) or a delisting return at or below 60% as a crash event. The delisting return is folded into the final price so that a total loss registers as approximately 100% rather than being dropped as a missing observation.

Our data is extracted via the Wharton Research Data Services (WRDS) API. We collect monthly adjusted closing prices for U.S. and globally listed stocks from January 1985 to December 2024. U.S. stocks are sourced from the Center for Research in Security Prices (CRSP) while global stocks are sourced from Compustat Global.

We define a falling-knife stock as one experiencing a drawdown of 60% or more within a trailing 12-month window. We collect monthly adjusted closing prices for U.S. and globally listed stocks from January 1985 to December 2024.

4.2 U.S. Sample

The U.S. stock universe is constructed from the CRSP Monthly Stock File (crsp.msfl), joined with the CRSP Names File (crsp.msfnames) and restricted to ordinary common shares (share codes 10 and 11) and to the delisting file (crsp.msdelist). Prices are adjusted for splits and distributions using the CRSP price adjustment factor. Securities are identified by PERMNO (CRSP's permanent security identifier) and by ticker. The initial price screen across all common shares identifies 38,207 falling-knife episodes across 15,607 unique firms between January 1985 and December 2024. To focus on economically important equity recoveries, two filters are applied that reduce our overall universe of securities for study:

1. *Data Sufficiency*: Stocks are required to have adequate prior fundamental and price history in WRDS master data files.
2. *Micro-Cap Exclusion*: Micro-cap stocks are excluded due to severe illiquidity, wide bid-ask spreads, and extreme idiosyncratic noise, which distorts recovery signaling.

Market capitalization size tiers are determined annually using the CRSP Market Indexes methodology. At the end of June each year, the entire CRSP common-share universe is ranked by market capitalization, computed as: *Price x Shares Outstanding*

Size boundaries are set based on cumulative market capitalization percentiles, as shown below in Table 1.

Table 1 | CRSP Market-Capitalization Size Classification Thresholds

Large-Cap	Top 0%-70% of cumulative market cap
Mid-Cap	Next 70%-85%
Small-Cap	Next 85%-98%
Micro-Cap	Bottom 2%

Following the Fama–French convention, June classification in year t governs the sample from July of year t through June of year $t+1$, and each crash episode is assigned the size tier of the June preceding its pre-crash peak. In addition, due to the June cutoffs being retained in the panel, we maintain a point-in-time monthly tier to capture dynamic size drift that reclassifies each firm-month against the respective year's cutoffs.

The U.S. panel ends in December 2024, the final month released in the CRSP legacy flat-file

format used for this extraction. Compustat Global coverage extends through December 2025; we truncate the global sample to December 2024 to maintain a common window across markets.

After applying size filters and history requirements, the finalized U.S. falling-knife sample consists of 3,919 stocks (3,256 small-cap, 443 mid-cap, and 220 large-cap firms) experiencing 9,341 crash episodes. The event-study analysis utilizes the first breach per firm, while the panel analysis retains all crash episodes. The resulting master panel contains 738,991 firm-months spanning January 1985 to December 2024.

Each firm-month in the U.S. master panel includes 54 variables categorized into six main groups:

1. *Identifiers*: PERMNO, GVKEY, ticker, and calendar month-end.
2. *Price & Market Metrics*: Adjusted price, raw CRSP price, shares outstanding, and market capitalization.
3. *Size & Industry Classifications*: Formation year, annual June market cap-size cutoffs, monthly point-in-time size tier, GICS sector code, GICS industry description, and localized industry tags.
4. *Market Risk*: Rolling 60-month market beta, estimated with a minimum 12-month window against the CRSP value-weighted total return index, accompanied by observation counts.
5. *Delisting & Distress Indicators*: Delisting date, CRSP delisting reason code, delisting return, coarse delisting reason, and a 4-way outcome flag (Active, M&A/Private, Other Delisting, Failure). Failure collects CRSP codes 400–499

(liquidation) and 500–599 (dropped for deficiency/bankruptcy). In survival models, M&A/private exits censor an episode, whereas failures represent competing risks.

6. *Accounting Ratios & Macro Metrics*: 27

accounting variables sourced from the WRDS Financial Ratios Suite (wrdsapps.firm_ratio) and Compustat Annual, spanning valuation, profitability, financial soundness, efficiency, and real EPS growth. Accounting variables are winsorized at the 1st and 99th percentiles, with high-missingness variables tracked via binary missingness indicators. Finally, monthly macro indicators from the Federal Reserve Economic Data (CPI-U, Industrial Production, and Effective Federal Funds Rate in levels and Year-over-Year changes) are merged into the panel.

4.3 *Global Sample*

Global equity data is drawn from Compustat Global (comp.g_secd and comp.g_security) restricted to common shares at month-end. Because Compustat Global records prices in local currency, the adjusted price series, given a total-return price series, is computed as:

$$\text{Adjusted Price Series} = \frac{\text{Daily Close}}{\text{AdjFactor}} \times \text{TotalReturnFactor}$$

Note: Returns for all countries in this study were calculated in local currencies.

Securities are identified via the GVKEY-IID issue pair to evaluate cross-listed securities independently. Each record retains ISIN, SEDOL, and exchange codes.

An unconstrained global screen across all non-U.S. markets identifies 129,075 crash episodes across 52,188 security issues (44,684 firms) in 117 different countries. From the global universe, we select 15 primary, non-U.S. markets. Selection is based on end-of-June 2024 market capitalization converted to USD using World Bank official exchange rates. Primary issues are retained to prevent double-counting cross-listed entities. Markets were chosen across the Americas, APAC, and EMEA regions to balance MSCI developed and emerging market classifications.

Along with the United States, the final dataset covers 15 key national markets (five developed and 10 emerging). The non-U.S. countries included within our analysis are shown in Appendix A, along with their market capitalization weights. These countries were chosen as a representative sample of their respective areas and market classification, which we divide according to Morgan Stanley Capital International (MSCI). This includes Asia-Pacific (APAC), Europe, Middle East, and Africa (EMEA), and the Americas.

These 16 countries in total were chosen to equally represent all regions of the globe and maintain a balanced number of developed and emerging market classifications while consisting of a high number of drawdowns, large market capitalizations, available securities' data, and a comparable index benchmark.

Market capitalization figures in Appendix A are calculated as of end-of-June 2024 and reflect the aggregate market value of the securities included in our study sample from that specific country, not that country's total equity market. For each country, we sum the June 2024 market capitalization of every stock in our sample belonging to that country, giving a single, total listed market value per country. This is a study-sample aggregate, not an independent measure of the country's broader public equity market.

Sample weight is each country's share of this study-sample total, equivalent to each country's aggregate market capitalization divided by the sum of aggregate market capitalization across all countries in the sample. Therefore, it reflects each country's relative representation within our sample, not its share of true global market capitalization.

Size tiers for non-U.S. equities are constructed using the same cumulative

Along with the United States, the final dataset covers 15 key national markets (five developed and 10 emerging). The non-U.S. countries included within our analysis are shown in Appendix A, along with their market capitalization weights.

market capitalization bands applied within each country and formation year. Ranking within the country-year ensures that local currency values are properly relative and currency exchange factors cancel out within local distribution rankings. Country-years with fewer than 10 listed companies were left unclassified and dropped along with micro-cap companies.

After applying country filters and removing micro-caps, the non-U.S. Compustat Global panel contains 13,601 security issues across 12,288 companies, comprising 22,456 crash episodes across 2,985,112 issue-months. The global sample contains 9,901 small-cap, 2,203 mid-cap, and 1,497 large-cap issues, classified by size at first crash date.

Pooling the U.S. and international equity datasets yields a combined multi-market analysis set of 33,615 falling knife episodes across 16 countries, categorized as 26,225 small-cap, 4,614 mid-cap, and 2,776 large-cap events.

While the global panel tracks price, size, sector, and delisting/failure indicators (derived from Compustat deletion reason codes), fundamental accounting ratio coverage is confined to the U.S. dataset due to WRDS Financial Ratios Suite availability constraints. As a result, we decided to reconstruct the ratios that we retrieved from the WRDS Financial Ratios Suite for the various countries apart from the United States.

We approached this task by extracting the quarterly fundamentals data from the Compustat data vendor located on WRDS. Utilizing the line items in the dataset, we

calculated ratios using relevant formulas. For example, to calculate the gross profit margin, we divided the trailing 12-month gross profit by the trailing 12-month revenue. We replicated this 19 times for 19 of the 21 ratios. The two ratios we decided to omit were the R&D spending/sales ratio and the Shiller PE Ratio. Reconstructing the R&D spending/sales ratio was not possible due to the differences in how GAAP and IFRS standards recommend reporting R&D expenditures. This discrepancy resulted in no R&D expenditure data available on WRDS. Calculating the Shiller PE Ratio was also too time-consuming, as we would have needed to take an average of inflation-adjusted earnings for 10 years prior to the crash. Due to these two challenges, we did not use those two variables in our global analyses.

Among the 19 variables we were able to construct, there were varying degrees of data availability. For example, our lowest availability was around 20% of the total data we needed, but the majority were higher, around 50-60%. We accounted for sparsity in the data for our models that were a part of our analysis. We initially thought of using similar line items as proxies for our ratio construction; however, we decided against this method because we would be comparing two different metrics across countries.

For example, although using income before extraordinary items instead of net income could have increased data availability for the net profit margin calculation, it would not measure the same underlying variable. Despite the similarity between these line items, relying on proxies would prevent us from consistently analyzing the relationships between identical variables across countries.

Similar to our United States study, we are exploring the effect of macroeconomic factors on the recovery of fallen stocks. We use the same variables as the United States study. Initially, we attempted to retrieve the data through a single source, as we believed that it would be standardized in a method similar to what we attempted for the fundamental data for international markets. However, after exploring various datasets, we couldn't get the specific data we wanted.

We were always missing one of the four variables for a monthly comparison. In addition, if we were able to pull data for all four variables, we ended up with annual data instead of monthly. So, macroeconomic data were collected from a combination of national statistical agencies, central banks, and international databases. However, because the agencies reported GDP growth differently, some of our GDP growth rate values were month-over-month, and some were quarter-over-quarter.

Even though there was a difference, we believe both could still be used in the models because they both measure economic output *changes* over standardized time frames, even though quarter-over-quarter captures short-term momentum and year-over-year smooths out seasonal noise. In addition, this information could be used for other markets that we didn't cover in this study, as we used both types of data in our models.

5. Methodology

5.1 Sample Construction

The analysis began with the complete dataset of drawdowns across all countries. To ensure consistency in market structure, benchmark selection, and data availability, the sample was initially restricted to U.S.-listed stocks. The initial screening of removing the micro-cap stocks reduced the number of unique stocks from 15,607 to 3,919. See Exhibit 1.

Exhibit 1 | From Episodes to Comparable Stocks



Source: WRDS, Jan. 1985 to Dec. 2024

The drawdown episodes span the period from 1985 to 2024. Consequently, the number of observations decreased as the investment horizon increased. The first reason for this reduction is right-censoring: episodes occurring toward the end of the sample period (for example, a drawdown in 2022 did not have sufficient subsequent data to evaluate a complete five-year recovery period). The second reason is the historical availability of the benchmark indices. We will share additional details on this issue in the next section. Following these adjustments, the final number of benchmark-comparable observations consisted of 3,072 episodes for the 1-year analysis, 2,883 episodes for the 3-year analysis, and 2,771 episodes for the 5-year analysis.

5.2 Return & Benchmark Methodology


For each stock in the final sample, post-drawdown performance was evaluated by calculating the annualized geometric returns over 1-year, 3-year, and 5-year investment horizons following the date on which the stock's price reached its trough. These return horizons were selected to capture both the short-term and longer-term recovery patterns following severe price declines. To assess whether these returns represented abnormal performance, each stock's return was compared against an appropriate benchmark index based on its size classification before the drawdown. Traditionally, a stock's sensitivity to the market (beta) is used to measure abnormal returns. We chose not to use this because the beta is likely not stable (pre-crash vs. post-crash). A post-crash beta reflects a leverage effect from the crash itself instead of a change in the firm's business risk. At the same time, a pre-crash beta doesn't reflect the firm's post-crash risk profile at all.

Specifically, large-cap stocks were benchmarked against the S&P 500 Index, mid-cap stocks against the S&P MidCap 400 Index, and small-cap stocks against the S&P SmallCap 600 Index. (In the last section of this report, we alluded to an issue with historical availability of the indices. To clarify, the S&P MidCap 400 Index was introduced on 19 June 1991, while the S&P SmallCap 600 Index was introduced on 28 October 1994. As a result, early drawdown episodes involving mid-cap and small-cap stocks could not be benchmarked for periods before the introduction of these indices and were therefore excluded from the benchmark comparison.) Using size-specific benchmarks provides a more appropriate comparison than a single market index, as it accounts for the well-documented size premium and allows stock performance to be evaluated relative to firms with similar market capitalizations. Following the sample selection process, the final dataset comprised 3,919 stocks, of which 264 were classified as large-cap stocks, 509 as mid-cap stocks, and 3,146 as small-cap stocks. See Exhibit 2 on the next page. The analysis then evaluated the post-drawdown performance of each stock relative to its corresponding benchmark over a given investment horizon, as stated above. For each horizon, the annualized return of the stock was compared with the annualized return of its benchmark over the same period. Stocks whose returns exceeded those of the benchmark were classified as outperformers, whereas those with lower returns than the index were classified as underperformers. See Exhibit 3 (also on the next page), where the numbers next to each bar reflect the recovery rates.

5.3 Post-Drawdown Recovery Rate Relative to Benchmarks

Having established the criteria for classifying stocks as outperformers or underperformers, Exhibit 3 presents the overall proportion of stocks outperforming their respective benchmarks across the three holding periods. The results provide initial support for a contrarian investment strategy in the period immediately following a severe drawdown. Over the 1-year horizon, about 62% of comparable stocks outperformed their benchmarks, indicating that stocks experiencing declines greater than 60% exhibited relatively strong performance during the initial recovery period.


Exhibit 2 | Sample by Market-Cap Tier



Source: WRDS, Jan. 1985 to Dec. 2024

Exhibit 3 | Post-Drawdown Relative Recovery Rate by Holding Periods


Share of Stocks Recovering Relative to Benchmark



Source: WRDS, Jan. 1985 to Dec. 2024

However, this effect weakens as the investment horizon increases. The recovery rates declined to 53.0% over the 3-year horizon and further to 47.4% over the 5-year horizon. One possible explanation is a base effect following the substantial first-year rebound, which raises the starting point for subsequent returns and makes continued outperformance more difficult. The declining recovery rate may also indicate that the momentum associated with the initial recovery gradually dissipated over time, potentially giving way to longer-term return reversal. Overall, the results suggest that the benefits associated with a contrarian strategy following severe drawdowns were strongest over the shorter investment horizon and became progressively weaker over time. Exhibit 4 further examines the results by separating the sample according to market-capitalization category and extends the time horizon to three and five years, as well.

Exhibit 4 | 1-, 3- and 5-Year Post-Drawdown Recovery Rate by Market Cap



Source: WRDS, Jan. 1985 to Dec. 2024

A clear pattern emerges across the three size groups, with the likelihood of recovery increasing with firm size. Approximately 70% of large-cap stocks outperformed their respective benchmark during the first year following the drawdown vs. 65% of mid-cap stocks and 60% of small-cap stocks. One possible explanation is that larger firms may possess stronger financial resources, more established business models, and greater access to capital, which could improve their ability to withstand periods of financial or market stress. Nevertheless, these factors represent potential explanations rather than conclusions directly established by the analysis. The overall 1-year relative recovery rate remained closer to 62% because the sample is heavily concentrated in small-cap stocks, which constituted the majority of the observed drawdown episodes and exhibited the lowest relative recovery rate among the three size categories.

Extending the analysis to the 3-year horizon shows the overall recovery rate relative to the market declined to 53.0%, compared with approximately 62% over the 1-year period. The results also indicate a notable change in the relationship between firm size and post-drawdown performance.

Large-cap stocks continued to exhibit the highest relative recovery rate at approximately 70%, whereas the relative recovery rates for both mid-cap and small-cap stocks declined to approximately 51%.

By the 5-year horizon, the initial post-drawdown advantage weakened considerably, with the overall proportion of stocks outperforming their respective benchmarks falling to 47.4%. This marks a notable reversal from the shorter holding periods, as the majority of stocks in the sample no longer outperformed their benchmarks over the longer term.

Despite this overall decline, the differences across market-cap categories remain pronounced. Large-cap stocks recorded a higher recovery rate of approximately 65%, compared with 48% for mid-cap stocks and 45% for small-cap stocks. Large-cap firms therefore remain the only category in which a clear majority of stocks continued to have higher recovery rates relative to their benchmarks over the 5-year period.

In contrast, both mid- and small-cap stocks fell below the 50% threshold, indicating that the post-drawdown relative recovery rates observed at shorter horizons were less persistent among smaller firms.

Viewed alongside the 1- and 3-year findings, these results highlight a consistent deterioration in relative performance as the holding period increases, particularly among mid- and small-cap stocks.


The evidence therefore suggests that the contrarian effect following severe drawdowns was primarily concentrated in the earlier stages of the recovery. However, more large-cap stocks showed greater persistence in a higher recovery rate relative to the market over longer horizons.

Exhibit 5 on the next page provides quantitative support for the pattern observed in the preceding analysis by comparing the median returns of the stocks with those of their respective benchmarks across the 1-, 3-, and 5-year horizons.

The median is particularly useful in this context because it is less sensitive to extreme observations than the mean, providing a more representative measure of the typical stock in the sample.

By the 5-year horizon, the initial post-drawdown advantage weakened considerably, with the overall proportion of stocks outperforming their respective benchmarks falling to 47.4%. This marks a notable reversal from the shorter holding periods....

Exhibit 5 | Median Returns: Stocks vs. the Benchmark



Source: WRDS, Jan. 1985 to Dec. 2024

At the 1-year horizon, the median stock return was 45% compared with 21% for the benchmark, representing a substantial 2400 basis-point return differential. However, this advantage diminished considerably over time.


At the 3-year horizon, the annualized median stock return fell to 16%, compared with 13% for the benchmark, reducing the differential to only 300 basis points. By the 5-year horizon, the relationship actually reversed: stocks generated an annualized median return of 11%, compared with 12% for the benchmark, producing underperformance of 100 basis points.

Overall, Exhibit 5 reinforces the earlier findings by suggesting that the strongest relative performance occurred in the period immediately following the drawdown, while the magnitude of this outperformance progressively declined and eventually disappeared over longer investment horizons.

Exhibit 5 reinforces the earlier findings by suggesting that the strongest relative performance occurred in the period immediately following the drawdown, while the magnitude of this outperformance progressively declined and eventually disappeared over longer investment horizons.


The three charts in Exhibit 6 (A, B and C) extend the analysis in Exhibit 5 by examining the distribution of stock-level benchmark-adjusted returns across the 1-, 3-, and 5-year horizons. Going back to Exhibit 3, at the 1-year horizon, 61.9% of stocks generated positive excess return. In addition, as shown in Exhibit 6A, the median excess return was 23.0% over the 1-year period. This indicates that a clear majority of stocks had a higher recovery rate relative to the market during the first year following the drawdown. This result was also consistent with Exhibit 5, where the median stock return of 45% exceeded the corresponding median benchmark return of 21% by 2400 basis points.

Exhibit 6A | 1-Year Excess Return (Stock-Benchmark) Distributions [Not Risk-Adjusted]



Source: WRDS, Jan. 1985 to Dec. 2024

Exhibit 6B | 3-Year Excess Return (Stock-Benchmark) Distributions [Not Risk-Adjusted]



Source: WRDS, Jan. 1985 to Dec. 2024


The 2400-basis point difference in Exhibit 5 and the 23.0% median excess return in Exhibit 6 are closely related but are not mathematically identical. Exhibit 5 calculates the difference between two separate medians, whereas Exhibit 6 first calculates each stock's return relative to its corresponding benchmark and then takes the median of these individual excess returns.

The shape of the 1-year distribution (first chart A in Exhibit 6) provides additional insight into these results. The distribution is substantially right-skewed, with most observations concentrated closer to the center of the distribution but with a long right tail generated by stocks experiencing exceptionally strong recoveries relative to their benchmarks. This asymmetry suggests considerable cross-sectional variation in the magnitude of post-drawdown recovery: although many stocks had a higher recovery rate during the first year, a smaller group experienced particularly large positive excess

return. The wide range of the 1-year distribution is also noteworthy, with the displayed x-axis extending from approximately -112% to 734%. (The -112% lower bound does not imply that a stock itself lost more than 100% of its value. Excess return represents the difference between the stock's return and its benchmark return rather than the stock's standalone return. Consequently, excess return can fall below -100% when a stock experiences a substantial loss while its benchmark generates a positive return. For example, a stock return of -80% combined with a benchmark return of +32% produces an excess return of -112 percentage points.) The figure also excludes a small number of observations at both extremes from the displayed range for readability, preventing extreme recoveries or losses from dominating the histogram scale.

The shape of the distribution changes noticeably as the investment horizon increases. At the 3-year horizon, the

Exhibit 6C | 5-Year Excess Return (Stock-Benchmark) Distributions [Not Risk-Adjusted]



Source: WRDS, Jan. 1985 to Dec. 2024

the distribution becomes substantially more symmetric and approximately bell-shaped, more closely resembling a normal distribution than the strongly skewed 1-year distribution. At the same time, the proportion of stocks generating positive excess return fell to 53.0%, while the median annualized excess return declined sharply to 2.5%. By the 5-year horizon, the distribution is also approximately bell-shaped and appears more evenly distributed around zero, while only 47.4% of stocks generated positive excess return and the median annualized excess return turned negative at -1.3%. The increasing symmetry of the 3- and 5-year distributions, together with their movement toward zero excess return, suggests that the extreme dispersion and strong positive skew observed immediately following severe drawdowns became less pronounced when performance was evaluated over longer horizons.

Taken together, Exhibits 5 and 6 suggest that the relative performance advantage following severe drawdowns was primarily a shorter-term phenomenon. The typical stock substantially outperformed its benchmark during the first year, but the 1-year distribution also demonstrated that this period was characterized by considerable heterogeneity and a relatively small number of exceptionally strong recoveries. As the horizon extended to three and five years, the distributions became more symmetric, median excess return moved toward and eventually *below* zero, and the proportion of stocks with a higher relative recovery rate fell from 61.9% to 53.0% and then 47.4%. Thus, the initial post-drawdown advantage not only declined in magnitude, but also became progressively less evident across the broader cross-section of stocks.

5.4 Crisis-Period Analysis: Methodology

A stock's recovery following a severe drawdown may depend heavily on the environment surrounding the crash. A stock that falls 60% during a systemic market crisis may recover largely because the broader market rebounds with it, while a stock that falls 60% in relative isolation, with no accompanying market-wide stress, must recover on the strength of its own fundamentals or firm-specific catalysts. Treating every drawdown as economically equivalent risks attributing recovery patterns to drawdown severity alone when they may instead be driven, in whole or in part, by the surrounding market environment.

To examine whether post-drawdown recovery patterns differ between systemic market crises and more idiosyncratic stock-level declines, the full sample was separated into crisis and non-crisis observations. Crisis periods were identified using the months associated with the most severe drawdowns in the relevant market-cap benchmarks: the S&P 500 for large-cap stocks, S&P 400 for mid-cap stocks, and S&P 600 for small-cap stocks. Specifically, a month is classified as a crisis month if the benchmark's month-end close is at least 20% below its highest month-end close over the trailing 12 months. The 20% threshold corresponds to the conventional definition of a bear market, while the rolling 12-month reference peak ensures that the windows capture episodes of active market stress. Note that a month may be classified as a crisis month despite a positive realized return, provided the index remains more than 20% below its recent peak.

Importantly, the classification was performed using the stock's trough date rather than its pre-crash peak or subsequent recovery date. A stock was classified as a crisis observation when the trough of its first drawdown exceeding 60% occurred *within* one of the designated crisis windows corresponding to its pre-crash market-cap tier. The windows were made to be month-inclusive to correspond with the monthly frequency of the price data.

The crisis windows used in the analysis were:

- Large-cap: Oct 1987–Jan 1988; Mar–May 1988; Aug 1988; Mar 2001; Jul–Oct 2001; Jul–Oct 2002; Dec 2002–Mar 2003; Sep 2008–Aug 2009; Mar 2020; Jun 2022; Sep 2022.
- Mid-cap: Aug 1998; Sep–Oct 2002; Dec 2002–Mar 2003; Oct 2008–Jul 2009; Sep 2011; Mar–Apr 2020; Jun 2022; Sep 2022.
- Small/micro-cap: Aug–Oct 1998; Feb–Mar 1999; Jul 2002–Apr 2003; Oct 2008–Aug 2009; Sep 2011; Dec 2018; Mar–May 2020; Sep 2022.

This tier-specific approach is important because market stress does not necessarily affect all capitalization segments simultaneously. This reflects a lead-lag effect across market-cap groups, whereby stress may become evident in one segment of the market before subsequently affecting another. For example, during some market-wide downturns, large-cap stocks may experience severe declines earlier, with the effects subsequently spreading to mid- and small-cap stocks. As a result, the different capitalization segments may reach their most

severe drawdowns at different points in time. This does not imply that large-cap stocks always lead small-cap stocks, but rather that the timing of market stress can differ systematically across capitalization groups. Applying a single crisis calendar to all stocks could therefore incorrectly classify observations as crisis or non-crisis events. Separate S&P 500, S&P 400, and S&P 600 stress windows were consequently used to account for these differences in the timing of market stress.

The August–October 1998 small-cap crisis window provides a useful illustration of how the exclusion procedure operates. Because the underlying price panel is monthly, the August–October window corresponds to the three, month-end trading observations of August 31, September 30, and October 31, 1998. A small- or micro-cap stock was classified as a crisis observation if the trough of its first drawdown exceeding 60% occurred on one of these dates. Approximately 195 small/micro-cap stocks met this condition: around 97 troughed in August, 68 in September, and 30 in October. These observations covered a broad range of drawdown severities, from stocks declining by approximately 60% to others losing more than 95% from their pre-crash peaks.

The unusually large concentration of troughs during this short period is important. The August–October 1998 period coincided with substantial market stress associated with the Russian currency devaluation and the Long-Term Capital Management crisis, particularly among smaller companies. Consequently, these 195 observations should not necessarily be interpreted as 195 completely independent, firm-specific crashes. Instead,

many stocks were exposed to a common systemic shock and reached their troughs at approximately the same time. Leaving this cluster within the non-crisis sample could therefore allow a single market-wide episode, together with the subsequent broad market recovery, to exert disproportionate influence on the estimated post-drawdown recovery statistics.

The same procedure was applied to every crisis window and capitalization tier. A stock was tested only against the crisis calendar associated with its own pre-crash size classification. For example, the August–October 1998 small-cap window affects only small- and micro-cap stocks; a large-cap stock reaching its trough during the same months would not be removed unless that date also fell within a separately identified S&P 500 crisis window. This prevents crisis periods specific to one capitalization segment from automatically being imposed on the entire market.

Across all capitalization tiers and crisis windows, applying this procedure reduces the sample from 3,919 stocks in the full dataset to 2,815 observations in the non-crisis subset. Therefore, 1,104 observations are associated with the identified systemic crisis periods and are excluded when constructing the non-crisis sample.

Separating the observations in this manner allows the subsequent analysis to investigate whether the strong recovery patterns observed after severe drawdowns were primarily associated with systemic market crises or with more idiosyncratic stock-level declines. In particular, removing crisis-associated observations helps prevent

synchronized market crashes and their subsequent broad recoveries from dominating the overall results. Comparing the full sample with the non-crisis subset then shows how the inclusion of crisis-associated observations influences the estimated recovery patterns.


It is important, however, not to interpret the difference between the full-sample and non-crisis statistics as the crisis-period statistic itself. The full sample contains both crisis and non-crisis observations, meaning that subtracting or comparing its results with those of the non-crisis subset identifies the influence of including crisis observations on the aggregate results, rather than directly estimating the performance of the crisis group. A direct comparison of crisis versus non-crisis recovery would instead require the 1,104 excluded observations to be analyzed explicitly as a separate crisis subsample. While we share some results later in this report from our analysis of crisis-specific periods, due to the broad scope of our research and time constraints, we didn't explore these periods in great detail. This offers a potential field for further exploration in the future.

This crisis and non-crisis decomposition was not extended to the 15 international markets analyzed later in this report. Several of those markets have considerably thinner sample sizes than the United States, and splitting an already-thin sample further into crisis and non-crisis subsets would leave some markets with too few observations to support reliable analysis.

5.5 Crisis-Period Results

The results in Exhibit 7 reflect the full and *non*-crisis decomposition applied to the U.S. sample, as outlined above in Section 5.4:

Exhibit 7 | Crisis Periods' Impact on Relative Recovery Rate



Source: WRDS, Jan. 1985 to Dec. 2024

Exhibit 7 shows the effects of removing crisis-associated drawdowns. The recovery rate relative to their benchmarks fell from 61.9% to 55.4% at the 1-year mark, from 53.0% to 46.1% at three years, and from 47.4% to 40.1% at five years. This is intuitive because returns were measured from the drawdown trough. During systemic crises, stocks with greater market sensitivity can fall disproportionately during the sell-off and subsequently rebound more strongly as market conditions recover. Measuring returns from the trough therefore captures this post-crisis recovery effect. In contrast, a 60%+ drawdown occurring *outside* a broader market crisis is more likely to contain a firm-specific component and such stocks may not benefit from the same market-wide recovery. We believe this helps explain why the non-crisis relative recovery rate fell below 50% at three and five years. As a robustness check, we calculated CAPM-based risk-adjusted returns to determine whether stocks exceeded the returns implied by their exposure to market risk. Expected returns were estimated using the local compounded policy rate, the corresponding benchmark-index return, and each stock's beta. Beta was estimated from monthly stock and market excess returns over the 36 months preceding the 60% drawdown, with a minimum requirement of 24 observations. When insufficient data were available, beta was set to 1.0. CAPM alpha was then calculated as the difference between the stock's realized return and its CAPM-implied expected return. Stocks with positive alpha were classified as having a higher recovery rate relative to the market, and the resulting rates were compared with the benchmark-adjusted results.

This similarity is consistent with the benchmark-matching methodology used throughout the analysis where stocks were matched to indices according to their pre-crash market-cap classification. Overall, the 1-year results shown in Exhibit 8 suggest that the market environment surrounding a severe drawdown was an important determinant of subsequent recovery, rather than drawdown severity alone.


Exhibit 8 provides a breakdown of the 1-year post-drawdown relative recovery rates by market-cap category after removing observations created during crises. The impact of removing crisis periods was minimal for large-cap stocks, with the outperformance rate declining by only 30 basis points. This suggests that large-cap recovery patterns were relatively similar across crisis and non-crisis environments. One possible explanation is that large firms tend to have stronger balance sheets, greater

access to capital and more diversified operations, while investors may also retain greater confidence in established companies during periods of market stress. In some systemic crises, large and systemically important firms may additionally benefit from policy interventions or broader measures aimed at stabilizing financial markets.

Therefore, rather than interpreting this strictly as a “too big to fail” effect, the result may reflect the greater resilience and financing capacity of large-cap firms.

The crisis effect was considerably stronger among mid- and small/micro-cap stocks. Mid-cap stocks showed the largest difference, with the relative recovery rate falling by 780 basis points when crisis observations are removed, compared with a 670-basis point decline for small-cap stocks. See Exhibit 9 on the next page.


Exhibit 8 | Crisis Period Impact by Cap Segment: 1-Year Post-Drawdown Relative Recovery Rate



Source: WRDS, Jan. 1985 to Dec. 2024

The particularly large effect among mid-caps may reflect their position between the two extremes: they may have greater exposure to cyclical and market-wide conditions than large firms, while still possessing sufficient financial strength to survive a systemic downturn and participate strongly in the subsequent recovery. Small and micro-cap firms may experience even greater declines during crises, but weaker balance sheets, financing constraints and a greater risk of permanent impairment may prevent some from participating fully in the rebound. This could explain why the crisis-related recovery effect was strongest among mid-caps rather than increasing monotonically as firm size decreases.

Exhibit 9 | Crisis Period Impact by Cap Segment: 3-Year Post Drawdown Relative Recovery Rate




Source: WRDS, Jan. 1985 to Dec. 2024

Exhibit 9 presents the 3-year relative recovery rate differential by market-cap category and reveals a notable divergence between large-cap stocks and the smaller capitalization groups. For large-cap stocks, removing crisis observations increases the relative recovery rate by 400 basis points, reversing the pattern observed at the 1-year horizon. One possible explanation is that large-cap stocks generally have greater analyst and institutional coverage, resulting in more information being incorporated into prices.

Over a 3-year period, the underlying trajectory of the business may also become considerably clearer, allowing the market to distinguish between firms experiencing temporary distress and those facing more persistent fundamental problems. As a result, large-cap stocks suffering severe, non-crisis drawdowns may have greater scope for selective long-term recovery, while the initial post-crisis rebound becomes less important over time.

In contrast, the crisis effect remained negative for both mid- and small/micro-cap stocks, with relative recovery rates falling by 700 and 760 basis points, respectively, when crisis observations were removed. The slightly larger effect among small/micro-caps may indicate that these firms remain more dependent on a broader recovery in economic conditions, financing availability and investor risk appetite. Overall, Exhibit 10 suggests that the influence of systemic crises on recovery varies with both firm size and investment horizon: it faded and reversed for large caps over three years, while remaining substantial for mid- and small/micro-cap stocks.

Exhibit 10 | Crisis Period Impact by Cap Segment: 5-Year Post Drawdown Relative Recovery Rate




Source: WRDS, Jan. 1985 to Dec. 2024

Exhibit 10 shows that at the 5-year horizon, the pattern begins to resemble the 1-year results again, with the crisis effect becoming negative across the market-cap groups. In particular, the large-cap differential, which was positive at three years, returns to negative territory. One possible explanation is a reversal of the intermediate 3-year effect as the longer horizon allows more information about the underlying businesses to become available. Over five years, much of the uncertainty surrounding whether a severely distressed company can successfully recover is likely to have been resolved, providing clearer evidence as to whether the original recovery thesis was justified.

The crisis effect also becomes more negative for mid-caps than for small/micro-caps. This may reflect differences in the timing and persistence of recovery across capitalization groups, with mid-cap firms potentially benefiting more from the sustained normalization in economic conditions following systemic crises. Overall, the 5-year results suggest that the relationship between crisis exposure and subsequent relative recovery rates was not linear across time or


firm size, with the relative effects changing as short-term rebounds, intermediate momentum, and longer-term fundamental information were progressively incorporated into prices. Previously, we looked at returns and recovery rates in all periods and non-crisis periods. Here, we sharpen our focus a bit more and compare median returns for falling knives and their benchmarks in non-crisis and crisis periods. Exhibit 11 illustrates the findings. In each of the two charts, we calibrated the Y-axes identically to help illustrate the performance disparity.

Exhibit 11A | Median Return: Stock vs. Benchmark During Non-Crisis Periods



Source: WRDS, Jan. 1985 to Dec. 2024


Exhibit 11B | Median Return: Stock vs. Benchmark During Crisis Periods



Source: WRDS, Jan. 1985 to Dec. 2024


Isolating the returns during crisis periods highlights the stronger recoveries falling knives made across the 1-, 3-, and 5-year periods. The 3- and 5-year annualized returns were measured from the trough. In addition, we calculated median returns across the market-cap spectrum over 1-, 3- and 5-year horizons during non-crisis and crisis periods. Not surprisingly, but as shown in Exhibit 12, median returns were higher following crisis periods, especially for mid-cap stocks. Again, in the two charts for Exhibit 12, we calibrated the Y-axes identically to help show the disparity.

Exhibit 12A | Median Return by Market Cap During Non-Crisis Periods



Source: WRDS, Jan. 1985 to Dec. 2024

Exhibit 12B | Median Return by Market Cap During Crisis Periods



Source: WRDS, Jan. 1985 to Dec. 2024

Overall, the crisis analysis shows that the market environment surrounding a severe drawdown played an important role in subsequent recovery. Crisis-associated stocks generally exhibited stronger post-trough relative recovery rates and returns, consistent with the rebound following systemic market declines, while the strength and persistence of this effect varied across market-cap groups and investment horizons. These findings suggest that drawdown severity alone is insufficient to explain recovery; both market conditions and firm size matter. Next, we applied modeling techniques to see whether we could identify specific traits among falling knives that might help predict future recoveries.

6. Can Models Help Forecast Future Returns for Falling Knives?

(Editor's Note: With the aid of ChatGPT and human editing, this section has been shortened considerably vs. what the team shared in its final report for its Capstone class. If readers are interested in more details, please contact the Executive Director of The Brandes Center: rpschmidt@ucsd.edu)

We investigate two, practical approaches to the critical question of trying to predict future returns among falling knives. First, we wondered whether recovery would depend on traditional financial information, such as valuation, profitability, debt, or company size. Second, we thought aspects of the decline itself, such as how quickly and how severely a stock fell, might contain more useful information. To address these hypotheses, we use two statistical tools: logistic regression and a machine-learning method called XGBoost. Logistic regression estimates

how different factors affect the probability of an outcome. In this case, the outcome is whether a stock recovered over a 1-, 3-, or 5-year period. XGBoost provides a flexible method that can detect nonlinear relationships and interactions between fundamental variables.


We tested many possible predictors, including price-to-earnings ratios, price-to-book ratios, profitability, debt, market capitalization, inflation, interest rates, industrial production, and two measures that described each crash episode. Using each approach, our analysis covers two related samples: 3,919 individual companies and 9,341 separate crash episodes. The biggest issue in the early models was the point at which returns were measured. At first, we measured performance from the stock's eventual lowest price or trough. But this creates a significant problem. An investor cannot know in real time when a stock reaches that trough. Buying at the exact low makes the investment look much more attractive than it would have been in reality.

With this in mind, we changed the starting point for return measurement. Instead of pretending an investor bought at the eventual bottom, we measured returns from the first date when the stock crossed the 60% decline threshold. This represents a point at which an investor could actually have recognized that the stock had suffered a major crash.

This adjustment changed the results substantially. Ex-post (trough-anchored) relative recovery rates of 61.9%, 53.0%, and 47.4% at 1, 3, and 5 years fall to 42.1%, 39.6%, and 35.2% once returns are measured from the first observable crossing. See Exhibit 13 on the next page.

The more practical, real-time results show that buying stocks that have fallen sharply was not automatically a profitable strategy.

Exhibit 13 | “Real-Time” Performance Reframed Recovery Rate: Ex-Post vs. Real-Time Anchor



Source: WRDS, Jan. 1985 to Dec. 2024

6.1 Model Refinement

After our preliminary findings, we sought to improve the models in several ways. First, we tried to account for many financial variables being closely related. For example, return on assets, profit margins, and return on equity may all reflect a company’s profitability. When predictors overlap too much, it becomes difficult to determine which one is actually important. We used a statistical test called the variance inflation factor to remove excessive overlap.

Second, we dealt with class imbalance. More stocks eventually outperformed than underperformed, so a model could appear accurate simply by predicting the more common result every time. To account for this, we used balanced models and evaluated them with measures such as AUC or the “area under the curve.” In essence, an AUC of 0.50 reflects a random guess. Higher values suggest a greater ability to distinguish between eventual winners and losers. Third, we used out-of-sample testing. The model was trained on earlier data and tested on later data that it had not seen. This approach provides a more realistic way to assess a model vs. testing it on the same information used to build it.

Finally, we introduced a walk-forward approach. The model was repeatedly trained on the past and then tested on the next period. This helped reveal whether the relationships remained useful across different market environments.

These corrections generally made the results weaker, but also more credible. The findings that survived these corrections were more likely to reflect a genuine pattern rather than a statistical accident.

6.2 The Models' Results

We found the strongest predictors were not traditional valuation ratios. They were two measures of the crash's shape. The first was "speed in months" or how long it took the stock to fall 60% from its previous high. Stocks that declined slowly were more likely to outperform later than stocks that collapsed quickly.

The second was "running drawdown" or how far the stock had already fallen by the time it was identified as having crossed the 60% threshold. Stocks that had fallen even further beyond that threshold by the detection date were less likely to outperform. Put simply:

- A slow decline was generally a better sign than a sudden collapse.
- A stock that had already suffered an extremely deep decline was less likely to recover.
- The path of the crash mattered more than the stock's fundamentals, including valuation, at the time.

Traditional financial information was not completely useless. Profitability measures sometimes showed a relationship with future performance when considered individually. However, once the crash-shape variables were included, most of the additional information from valuation and profitability became small or inconsistent. This suggests

that a company's financial ratios may partly reflect information already captured by the way its stock price has fallen.

6.3 Logistic Regression vs. XGBoost

We also compared logistic regression results with XGBoost. XGBoost was modestly better overall, especially for longer time horizons. At one year, the two methods were usually close. At three years, XGBoost was ahead by roughly 0.01 to 0.04 in AUC. At five years, its advantage increased to approximately 0.05 to 0.07.

Those differences were not large. This suggests that the underlying relationships are mostly straightforward rather than highly complicated.

A carefully designed logistic regression model captures most of the useful information, while XGBoost provides an incremental improvement by recognizing some nonlinear relationships.

The final models produced AUC scores of roughly 0.65 to 0.70. That is meaningfully better than random guessing, but it is far from a perfect prediction. In essence, the models can rank stocks somewhat better than chance, but they cannot reliably identify every future winner.

The path of the crash mattered more than the stock's fundamentals, including valuation, at the time.

6.4 Important Limitations

The results should not be interpreted as a ready-to-use trading strategy. Another important limitation? Approximately 80% of the sample consisted of relatively illiquid, small-cap stocks. These stocks may be difficult or expensive to trade. The analysis also did not include transaction costs, bid-ask spreads, market impact, or the practical difficulty of buying and selling during a crash.

In addition, in some market periods, the model performed worse than random guessing. Results by industry were also noisy because some industries had relatively few observations. This means that the broad crash-shape findings are more reliable than claims about a particular industry.

The main conclusion is therefore modest: recovery after a major stock decline is partly predictable, but only to a limited degree. The most consistent information comes from how quickly and how severely the stock fell—not from simply finding a stock that looks cheap. These results are perhaps best understood as a ranking exercise. They help identify which stocks that fell sharply appear more likely to recover than others, but the findings do not prove that buying those stocks would produce attractive profits after considering real-world trading costs and risks.

The main conclusion is therefore modest: recovery after a major stock decline is partly predictable, but only to a limited degree.

7. Results: Falling Knives in Global Markets

In this section, we revisit our modeling techniques and build upon the results for the U.S. market shared earlier. We extend our analysis to 15 global markets, including 12,728 securities and 19,496 crash episodes, using an identical real-time detector but a different, more rigorous risk adjustment. Rather than the U.S. study's tiered S&P benchmark, every market here measures relative recovery against a CAPM-expected return. Because of this methodological difference, the U.S. figures throughout this section are shown as a reference baseline and are not directly comparable.

Overall, our key finding points to no single explanation for stock recoveries around the world. In the United States, as described in the last section, technical information about the crash itself, especially how large the decline was and how quickly it occurred, was especially important. Outside the United States, however, the key return driver tended to vary by country. Technical factors were most important in several markets, economic conditions were most important in others, and, still in others, company fundamentals dominated.

The results also suggest that developed and emerging markets tend to behave differently over longer investment periods. Developed markets generally showed slightly better relative recovery rates over five years than over one year. Emerging markets showed the opposite pattern: their relative recovery rates were stronger over one year, but tended to decline over longer periods.

7.1 Prediction Model Methodology for Global Markets

We used a statistical prediction model to evaluate 33 traits. These “predictors” were divided into three broad groups:

- Technical factors: Information about the stock’s recent price behavior, especially the size of its decline and how long the decline lasted.
- Macroeconomic factors: Information about the broader economy, including interest rates, inflation, and industrial production.
- Fundamental factors: Information about the company itself, such as its size, profitability, debt, valuation, and other financial ratios.

Similar to our analysis for U.S. falling knives, we used walk-forward testing, training the model using older data and then testing it on later data to simulate how an investor would actually apply the approach in real time.

We excluded five countries (Brazil, Egypt, Mexico, Switzerland and Peru) from the predictor-classification analysis because we believed their samples were too small or unreliable to deliver dependable conclusions. We list all of the countries included in our analysis in the Appendix.

7.2 Classification of Predictors and Results

After excluding the five unreliable markets, we classified the remaining markets according to the broad category of predictors that appeared most important.

Technical Markets

Technical factors dominated in the United States, China, India, Indonesia, the United Kingdom, and Turkey. The United States was the clearest example. Technical variables accounted for approximately 93% of the model’s measured importance. This supports the earlier U.S. finding that the size and speed of a stock’s decline were much more useful in predicting future performance than most other variables.

China, India, Indonesia, and the United Kingdom were also classified as technical markets, although their results were less extreme. For example, technical factors accounted for about 68% of the importance in China and 59% in India. This does not necessarily mean that company financial statements or the economy do not matter in these countries. Rather, it means that, when all the variables competed in the model, recent price behavior contained the strongest overall predictive signal.

Macroeconomic Markets

Macroeconomic factors dominated in Hong Kong and Taiwan. In both markets, the most important individual variable was growth in industrial production. Industrial production is a broad measure of how much a country’s

The United States was the clearest example. Technical variables accounted for approximately 93% of the model’s measured importance.

factories and businesses are producing. The relationship was negative and statistically significant across the 1-, 3-, and 5-year periods. In plain language, changes in industrial production appeared to be consistently related to the future recovery of falling knives in these markets.

This result suggests that stock recoveries in Hong Kong and Taiwan may be especially sensitive to the health of the broader economy. Investors in these markets may need to pay closer attention to economic growth and production trends than investors in markets where price behavior or company fundamentals are more important.

Fundamental Markets

Fundamental factors dominated in South Africa, Australia, and Japan. Fundamental information includes measures such as company size, profitability, leverage, valuation, and other financial characteristics. The results for South Africa were particularly strong: fundamentals represented approximately 94% of the model's importance. However, this result was based on only one usable sample, so it should be treated cautiously.

Australia was classified as fundamental-based, but only narrowly. Fundamental factors accounted for 45.2% of the importance, compared with 42.2% for macroeconomic factors. This is better understood as a mixed result rather than clear evidence that fundamentals alone explain recoveries.

Japan also showed mixed results. Technical, fundamental, and macroeconomic factors

each accounted for roughly one-third of the model's importance. Although fundamentals were labeled the dominant category, the difference was so small that it would be misleading to describe Japan as purely fundamental-based.

Developed Markets vs. Emerging Markets

We also compared developed and emerging markets, as the two markets displayed different patterns when adjusted by what we call "sample weight." As defined earlier in this report, this is each country's share of this study's sample total, equivalent to each country's aggregate market capitalization divided by the sum of aggregate market capitalization across all countries in the sample.

Therefore, it reflects each country's relative representation within our sample, not its share of true global market capitalization.

The recovery rates among developed market stocks, relative to benchmark and adjusted for sample weight were:

- 38.7% after one year
- 39.0% after three years
- 40.7% after five years

Emerging markets had higher short-term recovery rates, but lower long-term rates:

- 42.6% after one year
- 37.6% after three years
- 31.1% after five years

Exhibit 14 on the next page shows these results. Note: we truncated the Y-Axis to better illustrate the regions' disparity.

Thus, emerging markets appeared more likely to recover during the first year after a crash, but that advantage faded over time. Developed markets started somewhat lower, but improved modestly over the longer periods.

While Exhibit 14 only illustrates *what* the results reveal, we share a few ideas as to *why* we see these findings. First, crashes in emerging markets may be caused by sudden shocks, such as currency crises, political events, capital flight, or panic among investors. These events can produce a sharp decline followed by a relatively quick rebound. However, the underlying problems may not be fully resolved, making longer-term performance less dependable.

Second, emerging markets are often more exposed to currency devaluations, capital controls, and inflation shocks. These risks can accumulate over a 5-year holding period.


A stock may recover quickly from an initial shock, but still face continuing economic and currency headwinds. Developed-market crashes may more often result from slower economic cycles. These recoveries can take longer, but once economic fundamentals improve, the recovery may be more durable.

However, the United States was an important exception. Unlike many other developed markets, its recovery rate declined over time:

- 42.1% after one year
- 39.6% after three years
- 35.2% after five years

The United States therefore looked more like the emerging markets than like other developed markets. Because the U.S. sample had the strongest data reliability at every horizon, we believe the decline may be a real feature of the U.S. market rather than simply a statistical accident.

Exhibit 14 | Developed vs. Emerging Market Recovery Rate Relative to Benchmark



Source: WRDS, Jan. 1985 to Dec. 2024

Which individual variables mattered most?

When evaluating individual variables across all 15 countries, the most frequently significant variable was the Federal funds rate, a measure of short-term U.S. interest rates. It was significant in 10 of the 15 countries and had a negative relationship in eight of those cases. Industrial production growth was significant in nine countries. Inflation was significant in eight. Market capitalization, running drawdown, and speed of the decline were each significant in seven countries.

This result may seem inconsistent with the finding that technical factors were the most common dominant category. However, these two angles of analysis ask different questions. The individual-variable analysis asks: “In how many countries was this particular variable statistically significant?”

The category analysis asks: “Which broad group of variables carried the most predictive weight inside the full model?” Interest rates were the most frequently significant individual variable, but macroeconomic factors rarely dominated an entire market. By contrast, technical variables were individually significant in fewer countries, but technical factors were the dominant category in the largest number of markets.

Emerging markets tended to show stronger short-term recoveries but weaker long-term relative recovery rates.

Overall, our analysis in global markets yields four main conclusions:

1. International falling knife recovery rates were broadly similar to the U.S. results, even though the international analysis used a different risk adjustment.
2. Reliability matters. Only 10 of the 15 international markets had enough dependable data to support strong conclusions about the predictors of recovery. Results from smaller samples should not be treated as equally convincing.
3. The U.S. finding that crash size and crash speed dominate other variables does not apply everywhere. Non-U.S. markets respond differently depending on their economic structure, financial system, and the kinds of shocks they experience. Still, technical factors were the most common dominant category overall.
4. Finally, the length of the investment horizon matters. Emerging markets tended to show stronger short-term recoveries but weaker long-term relative recovery rates. Developed markets generally showed the reverse pattern, although the United States and Australia did not follow this recovery pattern.

We believe the broader lesson is that investors should be cautious about applying a single recovery strategy to every country. A model based mainly on recent price declines may work well in the United States and several other markets, but economic conditions or company fundamentals may matter more elsewhere.

In addition, apparent patterns must be judged in light of sample size and statistical reliability. A result that looks powerful may be based on too little information to support a confident investment decision.

8. Conclusion

This study set out to test whether the popular caution against "catching a falling knife" holds up empirically, and the answer is more conditional than either side of that debate typically allows.

Our initial result showed that falling knife stocks contained higher recovery rates than their respective size benchmark by 62%, 53%, and 47% over the 1-, 3-, and 5- year holding periods, respectively. This result, however, is an artifact of measuring returns from the eventual trough, a point no investor can identify in real time.

Once returns were re-anchored to the first observable moment a stock crossed into severe distress, the contrarian strategy no longer reliably beat its benchmark at any horizon tested. We also found that a large portion of the returns were driven by a small number of stocks that vastly outperformed their respective benchmark.

That said, the crash itself was not uninformative. How deep a stock had already fallen and how long it took to get there consistently outweighed conventional valuation and profitability fundamentals in predicting which distressed stocks went on to recover. This holds true among most sectors; however, we noticed three industries were propelled by different factors. Financial firms'

returns tended to be driven by speed of the fall. Technology businesses' returns were determined by macroeconomic factors and healthcare companies' returns were shaped by valuation metrics.

This pattern, however, was hyper-specific to the United States. Extending the analysis globally, the general magnitude of underperformance replicated closely across markets, but the underlying structure did not. Macroeconomic conditions carried more weight across countries than crash-shape variables did, and the direction of the crash-depth effect itself reversed in several major markets.

Developed and emerging markets also diverged in how recovery odds evolved, with the U.S. standing out as an exception to its own category's pattern. Overall, these results offer a qualified rebuttal to the falling-knife adage: buying indiscriminately at the point of severe distress was not a winning strategy on a risk-adjusted and benchmark-relative basis, but the crash carried enough information to likely separate future winners from losers.

We thought that broader market downturns would allow us to determine specific windows which would allow for a greater number of winners. While this was true, the benefits we

...buying indiscriminately at the point of severe distress was not a winning strategy on a risk-adjusted and benchmark-relative basis

found in terms of recovery rates and strong median returns for stocks vs. their benchmarks tended to fade more quickly vs. non-crisis periods during our 5-year post-fall study period.

Several questions raised by this project remain open for future research. What economic mechanism explains why crash depth predicts recovery in opposite directions across markets? Is it a genuine reflection of differing institutional or informational environments?

Does the developed-emerging divergence in recovery trajectories reflect a real structural difference, or is it substantially driven by a small number of large markets? This analysis is also explanatory rather than a tradable strategy, since it does not yet account for transaction costs or sample illiquidity, which is a natural next step in testing whether the signal survives realistic implementation frictions.

Finally, several international markets currently lack enough data for reliable modeling, and revisiting them as more history accumulates may clarify whether their unreliability is a matter of scale or a more structural matter. We leave these questions and others to future researchers.

Appendix

Non-U.S. Market Capitalization, Classification, and Sample Weights by Country

Country	Region	Market Classification	Market Cap (USD trillions)	Sample Weight
China	APAC	Emerging	8.29	22.67%
Japan	APAC	Developed	6.36	17.39%
India	APAC	Emerging	4.92	13.45%
Hong Kong	APAC	Developed	3.60	9.83%
United Kingdom	EMEA	Developed	2.68	7.32%
Taiwan	APAC	Emerging	2.49	6.80%
Australia	APAC	Developed	1.60	4.37%

Appendix, continued . . .

Non-U.S. Market Capitalization, Classification, and Sample Weights by Country

Country	Region	Market Classification	Market Cap (USD trillions)	Sample Weight
Switzerland	EMEA	Developed	1.58	4.31%
Indonesia	APAC	Emerging	0.67	1.82%
Mexico	Americas	Emerging	0.44	1.20%
Turkey	EMEA	Emerging	0.41	1.13%
South Africa	EMEA	Emerging	0.41	1.12%
Brazil	Americas	Emerging	0.34	0.93%
Peru	Americas	Emerging	0.08	0.23%
Egypt	EMEA	Emerging	0.04	0.11%

Bibliography

Brandes Institute. (2004). "Falling Knives Around the World." Brandes Investment Partners.

Brandes Institute. (2007). "Fixed Income Falling Knives Phase Two: Examining the Relationship Between Issuer-Specific Bond and Equity Returns." (Brandes Institute Research Paper No. 2007-08). Brandes Investment Partners.

Choi, H.-S., & Jayaraman, N. (2009). "Is Reversal of Large Stock-Price Declines Caused by Overreaction or Information Asymmetry: Evidence from Stock and Option Markets." *Journal of Futures Markets*, 29(4), 348–376. <https://doi.org/10.1002/fut.20360>

De Bondt, W. F. M., & Thaler, R. H. (1985). "Does the Stock Market Overreact?" *The Journal of Finance*, 40(3), 793–805. <https://doi.org/10.1111/j.1540-6261.1985.tb05004.x>

Flossbach von Storch Research Institute. "Should You Catch a Falling Knife?" FvS Research Institute. <https://www.flossbachvonstorch-researchinstitute.com/en/comments/detail/should-you-catch-a-falling-knife>

Jegadeesh, N., & Titman, S. (1993). "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance*, 48(1), 65–91. <https://doi.org/10.1111/j.1540-6261.1993.tb04702.x>

Kusumahadi, T. A., & Madyakusumawati, S. (2026). "Do Volatility and Price Decline Matter? Evidence from Falling Knives in Indonesia Using Logistic Regression." *Jurnal Akuntansi*, 20(1), 52–78. <https://doi.org/10.25170/jak.v20i1.7688>

The Brandes Center thanks Dr. Michael Melvin and Prof. Chris Girard at UC San Diego's Rady School of Management for sharing their guidance and insights with the team on this project.

Disclosures

This document is for general information and educational purposes only, and must not be considered investment advice or a recommendation that the reader is to engage in, or refrain from taking, a particular investment-related course of action. Any such advice or recommendation must be tailored to your situation and objectives. You should consult all available information, investment, legal, tax and accounting professionals, before making or executing any investment strategy. You must exercise your own independent judgment when making any investment decision.

Past performance is not a guarantee of future results.

No investment strategy can assure a profit or protect against loss.

Diversification does not assure a profit or protect against a loss in a declining market.

All information contained in this document is provided “as is,” without any representations or warranties of any kind. We disclaim all express and implied warranties including those with respect to accuracy, completeness, timeliness, or fitness for a particular purpose. We assume no responsibility for any losses, whether direct, indirect, special or consequential, which arise out of the use of this presentation.

All investments involve risk. There can be no guarantee that the strategies, tactics, and methods discussed in this document will be successful.

It should not be assumed that any security transactions, holdings or sector discussed were or will be profitable, or that the investment recommendations or decisions we make in the future will be profitable or will equal the investment performance discussed herein.

Stocks of small companies usually experience more volatility than mid and large sized companies.

Data contained in this document may be obtained from a variety of sources and may be subject to change. We disclaim any and all liability for such data, including without limitation, any express or implied representations or warranties for information or errors contained in, or omissions from, the information. We shall not be liable for any loss or liability suffered by you resulting from the provision to you of such data or your use or reliance in any way thereon.

Unlike bonds issued or guaranteed by the U.S. government or its agencies, stocks and other bonds are not backed by the full faith and credit of the United States. Stock and bond prices will experience market fluctuations. Please note that the value of government securities and bonds in general have an inverse relationship to interest rates. Bonds carry the risk of default, or the risk that an issuer will be unable to make income or principal payment. There is no assurance that private guarantors or insurers will meet their obligations. The credit quality of fixed income investments is no guarantee of their safety or stability.

Disclosures-2

International and emerging markets investing is subject to certain risks such as currency fluctuation and social and political changes; such risks may result in greater price volatility. Strategies discussed herein are subject to change at any time by the managers in their discretion due to market conditions or opportunities.

Rolling periods represent a series of overlapping, smaller time periods within a single, longer-term time period. For example, over a 20-year period, there is one 20-year rolling period, eleven 10-year rolling periods, sixteen 5-year rolling periods, and so forth.

Please note that all indices are unmanaged and are not available for direct investment.

The declaration and payment of shareholder dividends are solely at the discretion of the issuer and are subject to change at any time.

The existence of a share repurchase authorization does not guarantee that any shares will be repurchased, and any such program may be amended, suspended, or discontinued at the issuer's discretion.

The MSCI data contained herein is the property of MSCI Inc. (MSCI). MSCI, its affiliates and its information providers make no warranties with respect to any such data. The MSCI data contained herein is used under license and may not be further used, distributed or disseminated without the express written consent of MSCI.

Nothing in this document should be interpreted to state or imply that past results are an indication of future performance. Investing involves substantial risk. It is highly unlikely that the past will repeat itself. Selecting an advisor, fund, or strategy based solely on past returns is a poor investment strategy.

The Brandes Center at UC San Diego's Rady School of Management was established in 2022 thanks to the philanthropic vision and support of Charles and Renée Brandes.

The Regents of the University of California and UC San Diego are not connected or affiliated with, nor do they endorse, favor, or support any product or service of Brandes Investment Partners, L.P.

4Q26

THE BRANDES CENTER

9500 Gilman Dr. #0553
La Jolla CA 92093

CONNECT WITH US

rady.ucsd.edu/brandes
brandes@rady.ucsd.edu

To receive new research and information about upcoming events,
contact us or use the QR code below:



UC San Diego

RADY SCHOOL OF MANAGEMENT
The Brandes Center