

Thin rallies, heavy sell-offs: A Fragility Ratio for when price and volume tell conflicting stories about stock returns

Florin Cepraga

investing.principles.com@gmail.com

<https://substack.com/@investingprinciples2048>

September 22, 2026

ORCID: 0000-0002-1220-1403

Abstract. To better understand what sits under the hood of returns, a conceptualization of fragility is needed first. This white paper posits that fragility - and not risk itself - comes from two market regimes: higher prices on lower volumes (hollow rally), and lower prices on higher volumes (deep breakdown). Taken together, these regimes can be informative about how the returns are obtained. A Fragility Ratio is proposed here, and it prints how much return an asset delivered per unit of fragility, narrowed down to these two regimes of volatility. Next, a Fragility Cost of Return is introduced, and it prints how much fragility an investor had to endure to obtain 1% of return. These insights are built upon the ratios proposed by Sharpe and Sortino, and the novelty of this paper lies in asking a narrower research question: What does an asset's fragility stemming from price and volume diverging from each other reveal on its own and in relation to the returns investors actually earn? To make these theoretical concepts visible and easy to use, the [Fragility Ratio indicator](#) was created using the code made available via Claude AI, and it is publicly available on TradingView for financial analysts, investors - institutional or retail - and journalists to use in their reports and articles.

keywords: #fragility #ratio #rally #sell #stock #return #trading #indicator

Introduction

Understanding how returns are obtained and under what circumstances has put many great minds through decades-long series of framing and reframing the philosophical and methodological constructs in an effort to explain why an investment might be favorable in itself or against other types of investments.

William F. Sharpe changed how investors look at returns when he put into a single, tractable expression the concepts of expected return, risk-free rate, and volatility, in what would later become known as the Sharpe ratio (1966).

Sharpe gave the financial world what it needed: structure and coherence by looking at excess returns in a way that can be measured in a very practical and quick manner across different asset classes. It took some 15 years before another researcher, Frank Sortino, brought forward new insights by arguing that investors should not be so much concerned about the upside volatility of prices, but should rather focus on what could go wrong; hence the focus on downside volatility. This ratio and the calculations for it have been around since the early eighties (1980-1981).

Understanding the perspective investors adopt when they look at returns has been narrowed down by introducing the volatility of volumes into the equation. By adjusting the Standard Deviation (SD) of prices while controlling for the SD of volumes, other bodies of research proceeded to investigate how much information asymmetry (Ceprega, 2025), as reflected by price and volume variations, an investor had endured to earn 1% of excess return.

Building on these insights, this white paper dives deeper into price and volume volatility, and aims to conceptualize fragility by looking at how the returns of liquid assets are obtained by taking into account only two divergent regimes: higher prices on lower volumes (hollow rally), and lower prices on higher volumes (deep breakdown). More specifically, this paper puts forward the following research question:

RQ: What does an asset's fragility stemming from price and volume diverging from each other reveal on its own and in relation to the returns investors actually earn?

The paper opens by briefly presenting the Sharpe ratio (1966), the Sortino ratio (1981), the information asymmetry (Ceprega, 2025), and then introduces the Fragility Ratio and the Fragility Cost of Returns. Examples of the calculations are provided in a table. Furthermore, the [Fragility Ratio indicator](#) is openly accessible on TradingView, while the code was made available via Claude AI (Anthropic, 2026). More on how to customize the indicator can be found in the Results section. For ease of interpretation, the appendix includes a screenshot of a symbol accompanied by an explanation of the output produced by the [Fragility Ratio indicator](#).

2. Theoretical Framework

2.1. The Sharpe Ratio

The Sharpe ratio is defined as:

$$S = \left(\frac{R_p - R_f}{\sigma_p} \right)$$

where

R_p is the portfolio's return,

R_f is the risk-free rate of return, and

σ_p is the portfolio's standard deviation

The numerator represents the reward in the form of excess return earned above a risk-free baseline, while the denominator represents variability. The higher the rate, the higher the returns.

The biggest weakness of the Sharpe ratio is that it makes no distinction between the types of volatility, treating upside or downside volatility just the same.

2.2 The Sortino ratio

Frank Sortino argued in 1980-1981 that investors should be more concerned with the downside deviation when assessing the presence of risk.

The Sortino ratio is defined as:

$$S = \frac{(R - T)}{TDD}$$

Where: R = the average period return

T = the target rate of return also known as MAR (minimum acceptable return)

TDD = the target downside deviation

While the Sortino ratio can be regarded as a modified Sharpe ratio, it has the merit of comparing performance, especially when returns are skewed (Sortino & Price, 1994). But Sortino did more than this. Sortino brought forward a new perspective: while standard deviation is a measure of dispersion around the mean, meaning taking into account both data above and below its mean, the downside deviation measures the dispersion of data below a mean or a specified threshold. This small nuance made fund managers, researchers, and investors look at the returns from a new perspective when it comes to measuring the presence of risk.

2.3. The Information Asymmetry Ratio (IAR)

Pushing the thought process further about the presence of risk, arguing that prices are not born in a vacuum, and that the total variability in prices and volumes can reflect information asymmetry (Slim & Dahmene, 2016), a new perspective was put forward by measuring how much information asymmetry, as reflected by price and volume variations, an investor must endure to earn 1% of excess return (Ceprega, 2025). The Information Asymmetry Ratio (IAR) is presented below: the numerator includes the price and volume volatility and their correlation, and the denominator reflects the excess return.

$$\text{IAR} = \frac{\sqrt{\sigma_r^2 + \sigma_v^2 + 2\rho_{r,v}\sigma_r\sigma_v}}{E[R_p - R_f]}$$

In sum, the Sharpe ratio measures reward-to-variability regardless of direction, the Sortino ratio measures excess return per downside volatility, and the Information Asymmetry Ratio (IAR) measures fragility among the positions taken by the market participants, while incorporating the co-movement between prices and volumes.

2.4. The gap in measuring returns: the conceptualization of fragility

There is a gap that is unexplored throughout these concepts, although it might transpire through all of them. How fragile are these returns? Or more precisely where does fragility even come from? Before diving into this, for clarity purposes, the author of this paper does not use *risk* and *variability / volatility* interchangeably. Price and / or volume *fluctuations / variability / volatility* can signal the presence of risk, but they do not have to be risk itself. This line of thought has been deeply elaborated before by Howard Marks (How to think about risk with Howard Marks, 2024).

So if fragility is not risk but signals the presence of risk, and if it is not the fluctuation / variability / volatility itself, *what is fragility to begin with?* Across disciplines, fragility describes a system's susceptibility to disturbances, going from the general dictionary definition "the quality or state of being easily broken or destroyed" (Merriam-Webster), to natural ecosystems defined as the degree of change in an environment following disturbances (Nilsson & Grelsson, 1995), to financial fragility that an economy can exhibit resulting in reduced returns and diminished welfare (Lagunoff & Schreft, 2001). The literature on financial fragility is extensive, and the purpose of this paper is not to look into the reasons, causes or consequences of this phenomenon but rather - similar to an extensive real-time market survey administered to all participants in the market/sector/company - to provide a snapshot of the state a company is in based on the results of the decisions made in the

trading pool. In other words, in what conditions of fragility have other market participants left the company as a result of their trades?

To answer this, it is important to understand where fragility comes from in the first place. The most important task is to identify the gap between what a system looks like and what it can withstand because fragility lives in this gap, not in volatility per se. Markets communicate information through price, but price is not supported in a vacuum. Therefore, the information communicated by the market participants is realized through the *combination* of price and volume. The combinations where price and volume agree or are in sync with each other are less prone to endangering the returns of a liquid asset. When looking at Graphic 1 (*Price and volume relation*), higher prices on higher volumes (top-left) indicate a rally with real participation; lower prices on lower volumes (bottom-left) indicate a decline with little participation.

The central argument of this paper is that the greatest threat to returns is less likely to come from the left-side of the quadrant (convergent regimes) and more likely to come from rallies unsupported by volume, which can be seen as thin market optimism, and sell-offs accompanied by heavy volume, which can be described as panic-driven liquidation. The real threat to returns that is more likely to expose the fragility of a construction hides, therefore, in the divergent regimes of the quadrant, visually represented at the right-hand side of the graphic. Higher prices on lower volumes (top-right) indicate a rally with little participation (let us label or frame this as a *hollow rally*); lower prices on higher volumes (bottom-right) indicate a decline with heavy participation (let us frame this as a *deep breakdown*).

Therefore, when looking at Graphic 1 (*Price and volume relation*), the author focuses on the right-hand side of the graphic and the story revealed by the divergent regimes.

Graphic 1

Price and volume relation in convergent (left-side) and divergent (right-side) regimes



Based on this line of thinking, the paper aims to answer the overarching research question:

RQ: What does an asset's fragility stemming from price and volume diverging from each other reveal on its own (RQ3) and in relation to the returns investors actually earn (RQ1, RQ2)?

by breaking it down into the following research sub-questions:

RQ 1: Among all the volatility this asset exhibited, how well was the investor compensated for the specific slice of it that came from price and volume telling conflicting stories?

RQ 2: Among all the volatility this asset exhibited, how much fragility coming from price and volume telling conflicting stories did an investor have to endure to obtain 1% of (excess) return?

RQ 3: Among all the volatility this asset exhibited, how much fragility coming from price and volume telling conflicting stories did this asset carry regardless of what it returned?

2.4.1. The Fragility Ratio (FR)

In other words, the question that needs to be answered with a mathematical formula is the one asking: *Among all the volatility this asset exhibited, how well was the investor*

compensated for the specific slice of it that came from price and volume telling conflicting stories?

While this is a reward-per-variability type of question - structurally, it resembles what Sharpe asks (i.e. return per unit of total variability) and what Sortino asks (i.e. return per unit of downside volatility) - the focus on the fragility ratio is narrowed down to a specific kind of volatility: the variability generated only on days where price movement and volume movement diverged from each other (i.e. a real up-move that volume did not confirm, or a real down-move that volume did confirm). It deliberately ignores every day where prices and volumes agreed or where nothing eventful happened, leaving that kind of variability left for Sharpe or Sortino to capture. A higher Fragility Ratio (FR) means the asset was earning good compensation specifically for tolerating divergent variability, not just general variability or downside variability.

$$FR = \frac{\bar{R}_p - R_f}{\sigma_F}$$

The entire novelty of the FR is in the denominator, since the numerator is the same as in Sharpe or Sortino. Looking at the numerator, the fragility volatility σ_F is defined as follows:

$$\sigma_F = \sqrt{\frac{1}{T} \sum_{t=1}^T \left[1_A(t) \cdot R_t^2 \cdot \min(1 + \max(-z_t, 0), \kappa) + 1_B(t) \cdot R_t^2 \cdot (1 + \max(z_t, 0)) \right]}$$

Looking at:

$$\frac{1}{T} \sum_{t=1}^T [\cdot]$$

this represents the average over the *entire sample*, not just the flagged days. This is a deliberate choice because most periods contribute exactly 0 since they are neither Regime A nor Regime B, so σ_F is diluted by however many clean timeframes sit alongside the fragile ones. This choice was made because it lets σ_F answer how much fragility an asset carries *as a whole, and* not how bad the fragile days were in isolation. It is also the reason why this ratio's magnitude is not directly comparable to Sharpe or Sortino.

$$1_A(t), 1_B(t)$$

this represents the *regime gates* that make this a *divergence detector of variability*. Ordinary up-days, ordinary down-days, and quiet noise all get gated to zero; instead, only the days where prices and volumes genuinely disagree in one of the two specified ways are allowed to contribute. This regime gates are also a distinction from Sharpe or Sortino.

$$R_t^2$$

this represents the returns squared on any of the flagged days. Squaring makes the plus or minus sign irrelevant since a fragile up-move and a fragile down-move both add to the fragility of the construction. Also, this means larger divergent moves are penalized disproportionately as compared to small ones. This is in line with the standard variance convention in Sharpe and Sortino, where big surprises matter more than small ones.

$$z_t = \frac{V_t - \bar{V}_t}{\sigma_{V,t}}$$

this represents how unusual the volume was in the asset's own terms, with one important distinction: it measures volume in standard deviation units instead of a raw multiple of the mean because the paper argues that “unusual” should be asset-specific. So instead of arbitrarily applying a general yardstick, the z-score is better suited to reflect each asset's own volume behaviour because a stock whose volume generally prints wild swings needs a much bigger deviation as compared to a stock that prints more stable volumes.

$$\min(1 + \max(-z_t, 0), \kappa)$$

this represents the Regima A (hollow rally) weight. $\max(-z_t, 0)$ isolates only the negative side of z_t (volume below its own average) and turns it into a positive penalty that grows the weaker the volume gets to define what a *hollow rally* is: *the further the volume falls short, the less trustworthy the gain*.

The $\min(\cdot, \kappa)$ cap exists purely for theoretical numerical safety so that one near-zero-volume print does not get to dominate the entire fragility estimate. Based on the data for the assets presented as an example in this paper, this cap had no influence on the data, so treating it as a guardrail rather than a constraint is a better way to look at it.

$$1 + \max(z_t, 0)$$

this represents the Regime B (Confirmed Breakdown) weight, and it mirrors the logic on the other side: $\max(z_t, 0)$ isolates the positive deviations for the volumes above the average, and it gets to define what a *confirmed breakdown* is: *the weight grows the more confirmed the sell-off is*.

The confirmed breakdown is deliberately left uncapped as compared to the hollow rally because, economically speaking, more conviction behind a decline is unambiguously worse than a thin rally, therefore, there is no reason to cap how much a heavily-traded breakdown should weigh.

$$\varphi_{up}, \varphi_{down}$$

this represents the minimum move size, and is embedded in the two regimes 1A and 1B. The minimum move size was chosen so that a near-zero return sitting on the "wrong" side of zero (pure noise) does not get flagged just because it is technically positive or negative.

2.4.2. The Fragility Cost of Return (FCR)

In other words, the question that needs to be answered with a mathematical formula is the one asking: *Among all the volatility this asset exhibited, how much fragility coming from price and volume telling conflicting stories did an investor have to endure to obtain 1% of (excess) return?*

This is a cost-per-outcome question as compared to a reward-per-variability question for FR, not a reward-per-risk question. For the Fragility Ratio, a higher number is better, while for the Fragility Cost of Return, lower is better. The nuance here lies in that the FCR is asking what "price" in terms of fragility had to be paid to get the actual return. A noteworthy assumption here is that the returns are expected to be positive.

The mathematical formula is the same for FR but inverted:

$$= \frac{1}{FR} \times 1\%$$

Or

$$= \frac{\sigma_F}{\bar{R}_p - R_f} \times 1\%$$

To ensure clarity and uniformity, how the results can be reported and interpreted is presented separately in the following section.

2.4.3. Fragility volatility σ_F annualized

In other words, the question that needs to be answered with a mathematical formula is the one asking: *Among all the volatility this asset exhibited, how much fragility coming from price and volume telling conflicting stories did this asset carry regardless of what it returned?*

The formula is:

$$\sigma_{F,\text{annualized}} = \sigma_F \times \sqrt{N}$$

where N is the number of time periods, in the examples from this paper this is 252 trading days.

The σ_F in daily units is defined as:

$$\sigma_F = \sqrt{\sigma_F^2}$$

where

$$\sigma_F^2 = \frac{1}{T} \sum_{t=1}^T \text{contribution}_t$$

The full steps along with the formulas on how to calculate all the elements from this paper can be found in [the Excel file available here](#). Should you not be able to download the file properly, please inquire for a copy of it, and it will be sent to you.

3. Results

3.1. Results based on a real-world case study

To better see how these elements reflect in real-world examples, see Table 1 below (*The Fragility Ratio and the Fragility Cost of Return*) for some examples of how these ratios apply.

Table 1

The Fragility Ratio and the Fragility Cost of Return: NVDA and AAPL - the two largest public companies in the world by market cap, PLTR - a “flat” stock, NKE - a “falling” stock

Metric	NVDA	AAPL	PLTR	NKE
Actual total return	26.12%	41.30%	-0.41%	-50,75%
Fragility volatility σ_F annualized	32,60%	25,15%	49.92%	49,85%
Fragility Ratio (annualized)	0.93	1.51	0.35	-1.29
Fragility Cost of Return* (per 1% of return)	1.25%	0.61	_*	_*
Regime A days (Hollow Rally)	45	22	40	31

Regime B days (Confirmed Breakdown)	24	17	27	38
---	----	----	----	----

*Note: Calculations made for NVDA, AAPL, PLTR, and NKE from September 18, 2025, to September 18, 2026; NVDA and AAPL selected because they are the largest companies by market cap, PLTR is described as a “flat” stock assuming “flat” involves a +/-0.5% range around 0; NKE is described as a “falling” stock assuming “falling” involves substantial declines. All calculations assume 0% risk-free rate, price movements are set at +/- 0.5%, and volume movements at +/- 0.5 SD; For Fragility Ratio, higher is better; For Fragility Cost of Return, lower is better. *FCR is only applicable to positive returns*

3.2. Reporting the results

To make the reporting meaningful and informative, the author of this paper encourages analysts to report not only the percentage but also, in brackets, the number of hollow rally days and confirmed breakdown days.

In general terms, this could look like:

Name of the symbol AND name of the function in italics AND number of trading windows* AND the first letter of the trading period** (d for days, m for month, etc.) AND, in brackets, the number of hollow rallies first, and secondly the number of deep breakdowns) AND = Result AND add *p* (price) and *SD* (standard deviation) in italics, followed by the = if different from the default $p=\pm 0.5\%$, $SD=\pm 0.5$

* 252d is assumed by default; to decluster reporting, only use if another timeframe is reported

** (d for days, m for month, etc.)

In the examples provided in this paper, this could look like this:

NVDA:

$$FR(45,24) = 0.93$$

$$FCR(45,24) = 1.25\%$$

AAPL:

$$FR(22,17) = 1.51$$

$$FCR(22,17) = 0.61\%$$

PLTR:

$$FR(40,27) = 0.35$$

$$FCR(40,27) = - \text{ (n/a; not reported)}$$

NKE:

$$FR(31,38) = -1.29$$

$$FCR(31,38) = - \text{ (n/a; not reported)}$$

3.3. Interpreting the results

Fragility Ratio (FR)

For every unit of fragility-specific volatility NVDA carried in one year from September 18, 2025, to September 18, 2026, volatility coming from 45 hollow rallies and 24 deep breakdowns, the stock delivered about 0.93 units of (excess*) return, meaning the compensation / the return was *roughly proportional* (1:1) to the fragility undergone.

For every unit of fragility-specific volatility AAPL carried in one year from September 18, 2025, to September 18, 2026, volatility coming from 22 hollow rallies and 17 deep breakdowns, the stock delivered about 1.51 units of (excess*) return, meaning the compensation / the return was *strongly above par* (1:1) to the fragility undergone.

For every unit of fragility-specific volatility PLTR carried in one year from September 18, 2025, to September 18, 2026, volatility coming from 40 hollow rallies and 27 deep breakdowns, the stock delivered about 0.35 units of (excess*) return, meaning the compensation / the return was *moderately above par* (1:1) to the fragility undergone.

For every unit of fragility-specific volatility NKE carried in one year from September 18, 2025, to September 18, 2026, volatility coming from 31 hollow rallies and 38 deep breakdowns, the stock delivered about -1.29 units of (excess*) return, meaning the compensation / the return was *substantially negative below par* (1:1) to the fragility undergone.

*Note here that since in the calculation $R_f = 0\%$, this is essentially return per unit of fragility. For the interpretation of the Fragility Ratio, a higher score is better. Also, as a suggestion: the effect size up to ± 0.2 is described as slightly/roughly, up to ± 0.5 moderate, up to ± 0.8 strong, and what exceeds this can be described as substantial.

Fragility Cost of Return (FCR)

To earn each 1 percentage point of the 26.12% NVDA actually returned over the year, an investor had to tolerate about 1.25% of annualized fragility-specific volatility. Framed differently, the fragility “cost” was roughly \$1.25 of divergence-variability exposure per \$1 of return earned, a cost moderately above par (1:1).

To earn each 1 percentage point of the 41.30% AAPL actually returned over the year, an investor had to tolerate about 0.61% of annualized fragility-specific volatility. In other words,

the fragility “cost” was roughly \$0.61 of divergence-variability exposure per \$1 of return earned, a cost moderately below par (1:1).

*Note here that the direction of interpretation is the opposite of the Fragility Ratio: for FCR, lower is better here, since this is a cost figure, not a reward-per-variability figure. Also, for PLTR or NKE, the FCR is not useful to report since the two stocks report negative returns. As suggested above: the effect size up to +/-0.2 is described as slightly/roughly, up to +/-0.5 moderate, up to +/-0.8 strong, and what exceeds this can be described as substantial.

3.4. Comparing the results

But what if an analyst wants to compare two or more companies? And when is the σ_F annualized important to report separately? When looking at these two examples, one might expect to read something along these lines in a report (for brevity reasons, suggestions are made only for two companies):

AAPL looks meaningfully less fragile than NVDA with a higher Fragility Ratio (1.51 vs 0.93, better reward per unit of divergence), a lower Fragility Cost of Return (0.61% vs 1.25%, roughly half the fragility endured per 1% of return), and fewer deep breakdown days in particular (17 vs 24), meaning AAPL’s declines occurred less often when validated by heavy volume.

AAPL also delivered a materially higher raw return (41% vs NVDA’s 26%), and *when analyzing the fragility volatility annualized*, AAPL’s σ_F annualized (25.15%) is lower than NVDA’s (32.60%) in absolute terms as well. *Therefore, support was found for AAPL’s better Fragility Cost of Return*, reflecting genuinely less fragility coming from divergent price-volume behavior.

When comparing FR:

NVDA earned its returns at a Fragility Ratio of 0.93, versus AAPL’s 1.51, meaning AAPL delivered significantly more return per unit of fragility (price-volume divergence) than NVDA did.

When comparing FCR:

NVDA’s investors endured a Fragility Cost of Return of 1.25%, versus AAP’s 0.61%, meaning NVDA’s investors had to tolerate roughly twice as much fragility-specific volatility to earn each 1% of the return they received.

3.5. Visualizing the results: The Pine Script code adaptation for the Fragility Ratio indicator

Based on this general formula, please note that the author of this paper has instructed Claude AI (Anthropic, 2026) to generate a Pine Script to create the [Fragility Ratio indicator](#), which can be publicly used on TradingView. A screenshot of how the [Fragility Ratio indicator](#) looks in practice can be consulted in the Appendix.

Some of the parameters were set from the get-go as defaults to remove some noise but still avoid overfitting the data. Most notably, the [Fragility Ratio indicator](#) takes into account price movements above or below 0.5% and volume movement above or below 0.5 standard deviations. Let us call these inputs the *stress threshold*.

The two regime conditions written with the specific thresholds plugged in can be identified here:

$$\begin{aligned} \mathbb{1}_A(t) &= \begin{cases} 1 & \text{if } R_t > 0.5\% \text{ and } z_t < -0.5 \\ 0 & \text{otherwise} \end{cases} \\ &\quad \quad \quad \\ \mathbb{1}_B(t) &= \begin{cases} 1 & \text{if } R_t < -0.5\% \text{ and } z_t > +0.5 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

A bare sign test ($R_t > 0$) would flag a +0.01% tick as a rally, which should probably be seen as noise. Requiring the move to clear $\pm 0.5\%$ before counting as a candidate ensures the regime is measuring a more realistic directional move instead of statistical noise. The same reasoning applies to the volume movement: rather than flagging any below-average or above-average volume day, the volume movement has to sit at least 0.5 standard deviations away from its own average before it can be considered meaningful.

Please note that the [Fragility Ratio indicator](#) published on TradingView can be customized on several layers: the % movement of the prices can be changed from the default +0.5% or -0.5%, the SD movement of the volumes can be changed from the default of +0.5 SD or -0.5 SD, the lookback period can be changed from 20 days, the rolling back period can be changed from the default 60 days, the periods per year can be changed from 252 days, and the risk-free rate can be changed from 0%. More on these elements, in the limitations and recommendations for future research.

4. Conclusion and discussion

From the intersection of communication science and capital markets, the main purpose of this white paper is to build on existing knowledge and ask a new research question about what is happening under the hood of how returns are generated. This is done by putting forward a *conceptualization of fragility*, and the paper argues that *the fragility of returns of*

liquid assets lies in price and volume telling conflicting stories. It is precisely in this gap between these conflicting stories, defined as *hollow rallies* (higher prices on lower volumes) and *deep breakdowns* (lower prices on higher volumes), where fragility lies. By asking “What does an asset’s fragility stemming from price and volume diverging from each other reveal on its own (RQ3) and in relation to the returns investors actually earn (RQ1, RQ2)?” this piece of research is an invitation to analysts, shareholders, investors (institutional or retail), journalists, and other stakeholders participating in the capital markets to look at the returns through a new lens. In other words, it puts into words and numbers what the market participants intuitively ask.

Questions like “how well am I compensated for being invested in this asset?” (answered by the Fragility Ratio) or “what is the cost of the fragility I have to endure to get 1% of return?” (answered by the Fragility Cost of Return) or simply “how much conflict lies within this asset?” (answered by the fragility volatility annualized) now have an answer that is easy to grasp and visualize.

The Fragility Ratio (FR) is useful mainly because it answers a question that investors, from institutional investors to retail investors, other shareholders, analysts, or journalists ask intuitively. They want to know how well or not they have been compensated for tolerating the conflicting stories the other investors told about the same company. Secondly, the FR does not ask anyone to learn a new formula or language because the idea of return per unit of variability has been present for decades thanks to William Sharpe and Frank Sortino, and entire generations of investors have learned how to ask how well they were paid for putting up with a specific investment. Thirdly, the FR is closer to a mental image than to a formula since people do not need a background in finance to picture a rally where the support gets thinner or a sell-off where the crowd rushes towards the exit.

The Fragility Cost of Return (FCR) is useful when it comes to translating an abstract variability-adjusted ratio into a concrete statement about the realized outcome by showing what a specific return costs in fragility terms. It is therefore an intuitive number to hand to someone asking about how clean or how fragile a specific but positive return was. More on why FCR does not work with negative returns can be found in the limitations.

The number of hollow rallies and confirmed breakdowns is useful for extra insights into the divergent regimes. However, aggregating them into one number could reveal a deceiving picture. In the examples from the paper, NVDA, PLTR, and NKE show a similar aggregated number of hollow rallies and confirmed breakdowns (NVDA 69 days, PLTR 67 days, NKE 69 days). The author argues that it is not only useless to cluster the two regimes’ days together into one number, but it also clutters judgement.

The fragility volatility annualized is useful to report, and it should be noted that it should not be interpreted as a standalone metric because two companies with similar variability can reveal different stories. As the examples showed, both PLTR and NKE posted an annualized fragility volatility of around 50%, however they tell different stories.

Limitations

By building on previous knowledge about standardizing how returns are obtained when accounting for *all variability in prices* (Sharpe, 1966), *downside variability in prices* (Sortino,

1981), and moving on to incorporating *all variability of prices and volumes* (Cepraga, 2025), this paper aims to understand the fragility of returns when looking through the lens of price-volume divergence. This does not come without limitations.

First, one of the limitations is conceptual. It stems from saying that fragility lies in the space generated by conflicting stories (higher prices on lower volumes, and lower prices on higher volumes). Critics argue that fragility does not have to be directional, and that the returns can be fragile even when obtained in a regime of lower prices on lower volumes. For example, a stock caught in a downward momentum can display such a behaviour for a long period of time. The counterargument is that a quiet decline such as lower prices on lower volumes can be better ascribed to sellers running out of conviction, which - time depending on the horizon, the company not going bankrupt or not being delisted - tends to precede stabilization or reversal. In addition, the examples given throughout this paper emphasize that the Fragility Ratio (FR) has a place in distinguishing between strong growth candidates, such as NVDA or AAPL, a near-flat stock such as PLTR, or a falling stock such as NKE.

The Fragility Cost of Return (FCR), on the other hand, is only useful when the returns are positive, as theorized in the theoretical framework. It becomes irrelevant to infer anything about two companies displaying a negative FCR. Although the Excel file allows the calculation for PLTR and NKE and displays negative percentages, the negative print is not helpful for interpretation. Therefore, FCR can be used as a lens to look at positive returns.

Another limitation is that the metric in itself, whether the Fragility Ratio or the Fragility Cost of Return, is not meant to substitute the Sharpe or the Sortino ratios. Also, neither FR nor FCR should be used in isolation from other metrics, nor do they make any inferences about the quality of a business (management, products, services, earnings, etc.). FR and FCR are designed to offer a lens to look at the returns by understanding how other investors who have traded the shares of the company have left the business at a certain moment, and in doing so, they take into account the conflicting story told by the mismatch of prices and volumes.

And finally, a limiting factor comes from the examples covered in this paper. All the companies presented here are large companies when measured by market cap, all these companies are listed in the US, and the US is a developed and highly liquid market. Understanding how fragility works in less liquid markets is therefore an important lens when looking at the returns of the stocks from other Developed or Emerging Markets.

Future research

A potential avenue for future research comes in testing or backtesting some of the parameters used in the calculations. For reasons of convenience, the paper assumes a price deviation of $\pm 0.5\%$, and a volume deviation of ± 0.5 SD. Other researchers might suggest leaving these thresholds at 0 or 0.25 or 0.75 or 1 or any other magnitude, and they might identify solid reasons for updating this based on the power of their tests. To ensure the uniformity of results, these parameters should always be published if values other than the default ones are used.

The same applies to the Fragility Ratio indicator itself. The default trailing volume window is set to 20, and the lookback window is set to 60, based on the simplification that 20

represents roughly a month of trading days and 60 roughly a quarter. Other researchers are encouraged to test different ranges of window/lookback periods, and they might find that other pairs are better suited for a wider range of scenarios.

Overall, this paper represents an invitation to look under the hood of the returns posted by publicly traded liquid assets, and takes the intuitive questions that investors ask to a level where these questions can be answered with a number while bringing to life the mental image told by the story of conflicting sides of prices and volumes.

Disclaimer

The author retains shares in some of the companies presented as examples in this paper. Note that the examples do not constitute investment advice. Their mention was for illustrative purposes only.

Conflict of interest

I have no conflicts of interest related to this research. I have not received financial support, consulting fees, or other benefits from organizations that may influence the interpretation or conclusions of this work.

References

Anthropic. (2026). *Claude*. [Large language model]. <https://claude.ai>

Cepraga, F. (2025). Revisiting the Sharpe ratio: Incorporating price and volume volatility into modern reward-adjusted performance to gauge information asymmetry - Proposing a new survey method to measure the collective judgment of market participants. SSRN: <https://ssrn.com/abstract=5775105> DOI: <http://dx.doi.org/10.2139/ssrn.5775105>

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431-1451.

How to think about risk with Howard Marks. (2024, September 12). Oaktree Capital. <https://www.oaktreecapital.com/insights/insight-video/education/how-to-think-about-risk-with-howard-marks>


Lagunoff, R., & Schreft, S. L. (2001). A model of financial fragility. *Journal of Economic Theory*, 99(1-2), 220-264.

Nilsson, C., & Grelsson, G. (1995). The fragility of ecosystems: a review. *Journal of Applied Ecology*, 677-692.

Sharpe, W. F. (1966). Mutual Fund Performance. *The Journal of Business*, 39(1), 119–138. <http://www.jstor.org/stable/2351741>

Sortino, F. A., & Price, L. N. (1994). Downside risk framework. *The Journal of Investing*, 3(3), 59-64.

Appendix: Screenshot of the Fragility Ratio indicator published on TradingView with the code made available via Claude AI



Note: The [Fragility Ratio indicator](#) from TradingView displays here the data for NVDA. In the upper part of the screen, the graph shows Regime A (hollow rally) and Regime B (deep breakdowns). They are displayed in the last 2 rows of the table presented in this screenshot under Fragility Ratio (z-score).

The components of the Fragility Ratio indicator are also visually represented in a table under Fragility Ratio (z-score). The Fragility Ratio itself is displayed in purple and is represented visually by the line that has a lot of ups and downs: the number displayed in the right part of the screen (2.19) is the same number as the first number displayed in the table presented in this screenshot. The second purple number and the filled purple body reading 0.3226 make the visual representation of the fragility volatility annualized (the second number presented in the table as a percentage, namely 32.26%).

As indicated in the body of this paper, please note the default settings of the Fragility Ratio indicator published on TradingView: the % movement of the prices can be changed from the default +0.5% or -0.5%, the SD movement of the volumes can be changed from the default of +0.5 SD or -0.5 SD, the lookback period can be changed from 20 days, the rolling back period can be changed from the default 60 days, the periods per year can be changed from 252 days, and the risk-free rate can be changed from 0%. If you make changes to the default settings, and to ensure uniformity in comparisons, please describe them in your articles or reports.

