

What Is a Return? (Especially When Prices Can Be Negative)

An Axiomatic Theory for Markets with Negative Prices

Tamas Nagy, PhD

tnagyphd@gmail.com

March 2026

Abstract

Every formula in quantitative finance — CAPM, Markowitz, VaR, Sharpe ratio, GARCH — takes returns as input. Yet the standard definitions of return fail when prices cross zero: log-returns are undefined, and simple returns produce sign errors. This is not hypothetical: WTI crude settled at $-\$37.63$ on April 20, 2020, and European electricity markets routinely produce negative prices. We ask: what properties *must* a universal return function satisfy? We prove that six axioms — five structural (additivity, smoothness, monotonicity, log-asymptotic behavior, antisymmetry) and one for transparency (algebraic closure, ensuring both directions of the price-return mapping are elementary) — determine the return function uniquely up to a single scale parameter. The solution is the **arcsinh-return**: $R(S, S') = \operatorname{arcsinh}(S') - \operatorname{arcsinh}(S)$. This is the only return function that (i) works for all prices $S \in \mathbb{R}$, (ii) recovers log-returns for large positive prices, (iii) is dimensionless, and (iv) has a closed-form inverse (\sinh). The arcsinh-return implies a volatility measure that remains well-defined through zero-crossing events, unlike log-return-based Black-volatility workflows, which require dropping observations or switching models. We derive implications for risk measurement (arcsinh-VaR, arcsinh-GARCH), regulatory compliance (FRTB), and performance attribution. The core axioms, return properties, and key characterization lemmas are formally verified in Lean 4; one elementary rational-function step in the uniqueness argument remains handwritten.

Executive Summary

Modern finance runs on returns. Risk limits, VaR, expected shortfall, GARCH models, portfolio construction, performance attribution, and regulatory capital frameworks all assume that returns are well-defined.

That assumption fails when prices approach or cross zero.

This is no longer a theoretical edge case. WTI crude settled below zero in April 2020. European power markets produce negative prices regularly. Rates, spreads, and transition-sensitive commodity markets can all enter regimes where the standard return definitions stop working properly. At exactly the moment institutions most need a well-defined measurement framework, log-returns become undefined and simple returns can become economically misleading.

This creates a real business and regulatory problem. A risk system should not become undefined during stress. A performance metric should not report the wrong sign because a price is negative.

A desk should not have to switch conceptual frameworks in the middle of a crisis just because the price domain changed.

This paper addresses that problem at the foundation level. Instead of asking which formula is convenient, it asks a more basic question: what properties should a universal return measure satisfy if it is to work across positive, zero, and negative prices without changing frameworks midstream?

Starting from six transparent axioms, the paper shows that the answer is essentially unique: the arcsinh-return.

The practical value is straightforward:

- It provides one return framework for the full real line, so zero-crossing events no longer break the measurement system.
- It preserves standard log-return behavior in ordinary positive-price regimes, so existing intuition and most existing workflows remain intact.
- It has an explicit forward and inverse mapping, which makes it simple to implement, audit, and explain.
- It keeps volatility, VaR, expected shortfall, and GARCH-style workflows defined on days when log-return-based workflows fail.
- It removes the need for ad hoc model switching between positive-price and negative-price regimes.

Empirically, this is not a proposal to discard the existing quantitative stack. On the large set of observations where both log-returns and arcsinh-returns are defined, the statistical behavior is essentially the same on the common domain. The established stylized facts of returns are preserved. The benefit appears exactly where the current framework breaks: near zero and across zero.

The strategic message is simple. Institutions do not need separate conceptual frameworks for positive-price markets and negative-price markets. They need one coherent return definition that works in both. This paper provides that definition.

The companion paper shows that the same foundation also leads to an option pricing framework that extends Black-Scholes to the full real line. This paper solves the more basic problem first: before pricing, hedging, or capital, one must define the return correctly.

1. Introduction

1.1 The Return as Foundation

Returns are the atomic unit of quantitative finance. Portfolio optimization (Markowitz 1952) maximizes expected return for given return variance. The Capital Asset Pricing Model (Sharpe 1964) prices assets by their return covariance with the market. Value-at-Risk measures the quantile of the return distribution. The Sharpe ratio normalizes excess return by return volatility. Time series models — GARCH (Engle 1982, Bollerslev 1986), stochastic volatility (Heston 1993), regime-switching — model the dynamics of returns.

All of these assume a well-defined return. But what *is* a return?

1.2 Two Definitions and Their Domains

The standard choices are:

Log-return. $r_t = \ln(S_t/S_{t-1})$. Advantages: additive over time, theoretically clean, implied by geometric Brownian motion. Limitation: requires $S > 0$.

Simple return. $r_t = (S_t - S_{t-1})/S_{t-1}$. Advantages: intuitive (“percent change”), portfolio-additive (portfolio return = weighted sum of asset returns). Limitation: requires $S_{t-1} \neq 0$, and produces **sign errors** when prices are negative.

The sign error is not cosmetic. Consider $S_{t-1} = -10, S_t = -5$. The position improved: the liability decreased from 10 to 5. But the simple return is $(-5 - (-10))/(-10) = -50\%$, signaling a loss. The continuous return fails as well: $\ln(S_t/S_{t-1}) = \ln(0.5) = -0.693$, again labeling an improving liability as a loss. For $S_{t-1} = -1, S_t = +1$, the simple return is -200% , even though the position went from liability to asset — a clearly positive outcome. The formula’s implicit assumption that $S > 0$ is baked into the denominator’s sign.

1.3 Not Hypothetical

On April 20, 2020, the WTI front-month crude oil contract settled at $-\$37.63$ (CFTC 2020). Every log-return-based risk workflow encountered a domain failure or forced model switch. Systems using simple returns produced nonsensical P&L. The industry response was to switch models mid-crisis — from Black-Scholes (log-returns) to Bachelier (arithmetic changes) — introducing discontinuities in Greeks, hedging parameters, and backtest metrics at the worst possible moment.

European electricity markets produce negative prices routinely during periods of high wind and solar generation (ACER 2025). Power traders need a return concept that works from $-\text{€}100$ to $+\text{€}500$ without model switching.

Negative policy rates in major markets, including the euro area from 2014 onward, challenged fixed-income return calculations (ECB 2014). Credit spreads can be negative. Carbon credits face regulatory risk that could drive prices toward zero.

1.4 The Hidden Assumption

Every major finance textbook defines the return using the logarithm or the ratio — and implicitly assumes $S > 0$ (Hull 2022; Shreve 2004; Cochrane 2005). This assumption is never stated as an axiom, never justified, and sometimes false.

- Hull (2022), *Options, Futures, and Other Derivatives*: defines return as $\ln(S_t/S_{t-1})$. Domain restriction: unstated.
- Shreve (2004), *Stochastic Calculus for Finance*: builds on geometric Brownian motion ($dS = \mu S dt + \sigma S dW$), which requires $S > 0$ by construction. Domain restriction: structural but implicit.
- Cochrane (2005), *Asset Pricing*: uses log-returns throughout. Domain restriction: unstated.

The assumption $S > 0$ was invisible because it was almost always true — until it wasn’t. This paper makes the hidden assumption explicit and removes it.

1.5 This Paper

We approach the problem axiomatically. Rather than proposing a return definition and checking its properties, we ask: **what properties must a return function satisfy?** We show that six axioms — five structural, one for transparency — determine the return function up to a single scale parameter. The solution is:

$$R(S, S') = \operatorname{arcsinh}(S'/c) - \operatorname{arcsinh}(S/c)$$

where $c > 0$ is a transition scale, and the canonical choice $c = 1$ gives unit return sensitivity at zero.

We call this the **arcsinh-return**. It is the unique return function that works for all prices, recovers standard log-returns, and has an algebraically closed inverse ($\sinh(\operatorname{arcsinh}(x)) = x$).

2. Axiomatic Framework

2.1 Definition

A **return function** is a map $R : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ assigning a return $R(S, S')$ to a price move from S to S' .

2.2 The Axioms

Axiom 1 (Additivity). Returns compound over periods:

$$R(S_1, S_2) + R(S_2, S_3) = R(S_1, S_3) \quad \text{for all } S_1, S_2, S_3 \in \mathbb{R}.$$

This is the fundamental accounting identity: the return from Monday to Wednesday equals Monday-to-Tuesday plus Tuesday-to-Wednesday. It implies $R(S, S') = f(S') - f(S)$ for some **return generator** $f : \mathbb{R} \rightarrow \mathbb{R}$, and $R(S, S) = 0$.

Axiom 2 (Smoothness). The return generator f is continuously differentiable (C^1).

No discontinuities, no kinks. Returns should not jump at zero or at any other price level.

Axiom 3 (Monotonicity). The return generator f is strictly increasing.

A higher price corresponds to a higher cumulative return. This is necessary for the return to have a well-defined inverse (given a return, recover the price).

Axiom 4 (Log-asymptotic). For large positive prices, the return recovers log-returns:

$$R(S, S') \rightarrow \ln(S'/S) \quad \text{as } S, S' \rightarrow +\infty \text{ with } S'/S \rightarrow \rho > 0.$$

This is the compatibility condition: in the regime where log-returns are well-defined and standard, the new return must agree with them. No disruption for equity markets.

Axiom 5 (Antisymmetry). The return generator satisfies:

$$f(-S) = -f(S) \quad \text{for all } S \in \mathbb{R}.$$

Owing \$10 is the mirror image of owning \$10. A price move from -10 to -5 (improving liability) should have the same magnitude as from $+5$ to $+10$ (increasing asset), with appropriate sign. This axiom implies $f(0) = 0$.

2.3 Discussion of the Axioms

Axioms 1–3 are standard. Any reasonable return definition satisfies them.

Axiom 4 ensures backward compatibility. For S&P 500 stocks trading at \$100+, the arcsinh-return must be indistinguishable from the log-return. This is non-negotiable for adoption.

Axiom 5 is the conceptual core. It says that the return function treats positive and negative prices symmetrically — not that positive and negative prices are equally likely (they need not be), but that the *measurement tool* is unbiased. An antisymmetric f gives $R(-S, -S') = -(f(S') - f(S)) = -R(S, S')$: the return from -10 to -5 is the negative of the return from $+5$ to $+10$. Since going from -10 to -5 is “good” (less liability) and going from $+10$ to $+5$ is “bad” (less asset), the signs are correct.

3. The Characterization Theorem

3.1 Statement

Theorem 1 (Characterization). *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfy Axioms 1–5, and additionally require:*

Axiom 6 (Algebraic Closure). *The function $g = \sinh \circ f$ is algebraic (rational) in S .*

Then $f(S) = \operatorname{arcsinh}(S/c)$ for a unique $c > 0$.

If additionally $f'(0) = 1$, then $c = 1$ and $f = \operatorname{arcsinh}$.

3.2 Proof Sketch

The full proof is in Appendix A. The strategy: define $g(S) = \sinh(f(S))$ and show that the axioms force $g(S) = aS$ for some $a > 0$, which gives $f = \operatorname{arcsinh}(S/c)$ with $c = 1/a$.

The key steps are: (1) Axiom 5 makes g odd; (2) Axiom 4 gives g linear growth; (3) Axiom 6 + monotonicity forces $g(S) = aS$ (the only globally monotone odd rational function with linear growth); (4) inversion gives $f = \operatorname{arcsinh}(aS)$; (5) normalization $f'(0) = 1$ pins $c = 1$.

3.3 The Role of Axiom 6

Axioms 1–5 alone do not determine f uniquely. Any C^1 , strictly increasing, odd function that is asymptotic to $\log S$ for large S would satisfy A1–A5. For instance, $f(S) = \operatorname{arcsinh}(S) + \varepsilon \sin(S)e^{-S^2}$ satisfies A1–A5 for small ε .

Axiom 6 is the *transparency* axiom. Its purpose is **bidirectional simplicity**: both directions of the price-return mapping must be elementary.

Direction	Operation	arcsinh	A perturbed alternative
Price \rightarrow Return	Apply f	$\operatorname{arcsinh}(S)$ — one function call	$\operatorname{arcsinh}(S) + \varepsilon \sin(S)e^{-S^2}$
Return \rightarrow Price	Apply $\sinh \circ f^{-1}$	$\sinh(r)$ — one function call	No closed form; requires numerical inversion

The motivation is **auditability**. A risk manager verifying a return applies $\operatorname{arcsinh}$ to the price. A regulator recovering the price from a reported return applies \sinh . Both operations are single elementary function calls available in any calculator or spreadsheet. No iterative solver, no black box, no numerical tolerance to argue about.

Compare this with implied volatility: to go from an option price to Black vol requires Newton-Raphson iteration — there is no closed-form inverse of the Black-Scholes formula. The implied vol surface is inherently a one-way street. The $\operatorname{arcsinh}$ return, by contrast, is a **two-way street**: both the forward and inverse maps are explicit.

Axiom 6 formalizes this requirement. Among all functions satisfying A1–A5, it selects those where the inverse transform $S = \sinh(f^{-1}(\cdot))$ is algebraic — and therefore verifiable. The identity $\sinh(\operatorname{arcsinh}(x)) = x$, the simplest possible algebraic closure, is what makes the entire framework transparent and auditable.

4. Properties of the Arcsinh-Return

4.1 Definition and Notation

Define the **arcsinh-space** variable $Y_t = \operatorname{arcsinh}(S_t)$, where $\operatorname{arcsinh}(x) = \log(x + \sqrt{x^2 + 1})$. Figure 1 shows $\operatorname{arcsinh}$ alongside \log and the identity: for large x , $\operatorname{arcsinh}$ tracks \log ; near zero, it tracks the identity; for negative x , it is the antisymmetric continuation.

The **arcsinh-return** is:

$$r_t^Y = Y_{t+1} - Y_t = \operatorname{arcsinh}(S_{t+1}) - \operatorname{arcsinh}(S_t).$$

We use r^Y throughout: the superscript Y denotes “return in arcsinh-space,” paralleling the convention that $r_t = \log(S_{t+1}/S_t)$ is the return in log-space.

4.2 Operational Formulas

From prices to return (observation \rightarrow computation):

$$r_t^Y = \operatorname{arcsinh}(S_{t+1}) - \operatorname{arcsinh}(S_t).$$

One $\operatorname{arcsinh}$ evaluation per price, then subtract. Defined for all $S_t, S_{t+1} \in \mathbb{R}$.

From return to price (prediction \rightarrow realization):

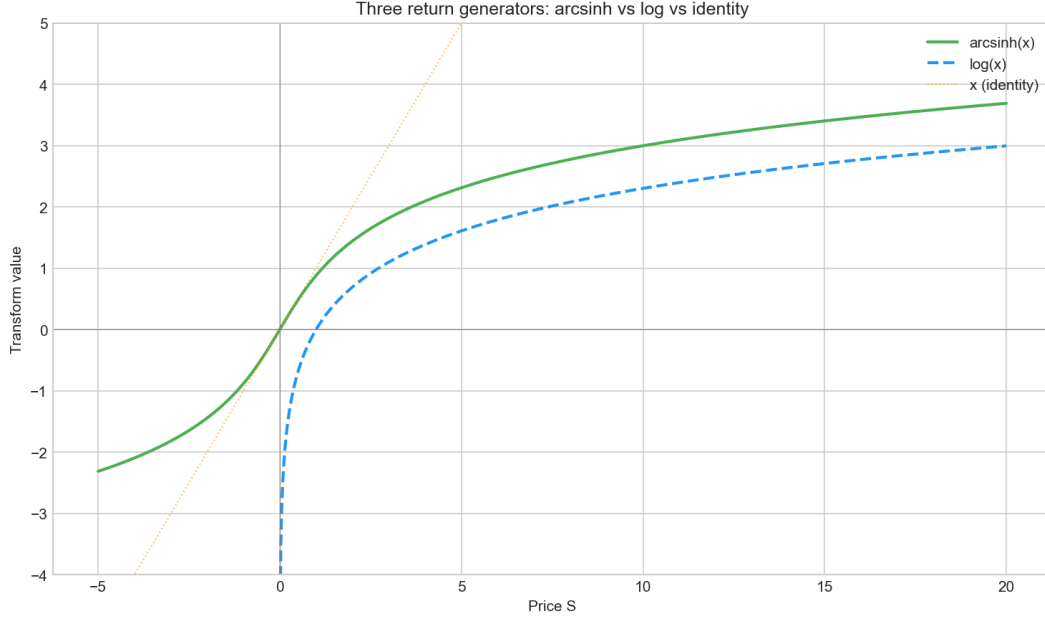


Figure 1: Figure 1: The arcsinh function compared to log and identity across negative, near-zero, and positive prices.

Given today's price S_t and the arcsinh-return r_t^Y , tomorrow's price is:

$$S_{t+1} = \sinh(\operatorname{arcsinh}(S_t) + r_t^Y).$$

One arcsinh, one addition, one sinh. Both directions are single elementary function calls — this is the bidirectional simplicity guaranteed by Axiom 6.

Multi-period compounding. Since arcsinh-returns are additive (Axiom 1):

$$\sum_{i=0}^{n-1} r_i^Y = \operatorname{arcsinh}(S_n) - \operatorname{arcsinh}(S_0),$$

and the terminal price is recovered by:

$$S_n = \sinh\left(\operatorname{arcsinh}(S_0) + \sum_{i=0}^{n-1} r_i^Y\right).$$

4.3 Parallel with Log>Returns

The arcsinh-return has exactly the same algebraic structure as the log-return, with a wider domain:

	Log-return r_t	Arcsinh-return r_t^Y
Transform	$Y_t = \log(S_t)$	$Y_t = \operatorname{arcsinh}(S_t)$
Return	$r_t = Y_{t+1} - Y_t$	$r_t^Y = Y_{t+1} - Y_t$
Forward price	$S_{t+1} = S_t \cdot e^{r_t}$	$S_{t+1} = \sinh(\operatorname{arcsinh}(S_t) + r_t^Y)$

	Log-return r_t	Arcsinh-return r_t^Y
Multi-period Compounding Domain Inverse	$S_n = S_0 \cdot e^{\sum r_t}$ Multiplicative $S > 0$ exp	$S_n = \sinh(\operatorname{arcsinh}(S_0) + \sum r_t^Y)$ Sinh-compounding $S \in \mathbb{R}$ sinh

In both cases, returns are additive in the transformed space. The log-return uses the (log, exp) pair; the arcsinh-return uses the (arcsinh, sinh) pair. The only difference is the domain: log requires $S > 0$; arcsinh works on all of \mathbb{R} .

For $S \gg 0$, the two pairs converge: $\operatorname{arcsinh}(S) \approx \log(2S)$, so the arcsinh-return differs from the log-return only by a constant that cancels in differences. The operational change for equity portfolios is negligible.

4.4 Regime Behavior

Proposition 2 (Regime interpolation). *The arcsinh-return interpolates between log-returns and arithmetic changes:*

(a) For $S, S' \gg 0$: $R(S, S') = \ln(S'/S) + O(1/S^2)$.

(b) For $S, S' \approx 0$: $R(S, S') = S' - S + O(|S'|^3 + |S|^3)$.

(c) For $S, S' \ll 0$: $R(S, S') = \ln(|S|/|S'|) + O(1/S^2)$. Note the sign: improving (less negative) gives positive return.

[Lean: partial — log bound and return monotonicity verified; asymptotic rates pending. Proofs in Appendix B.]

4.5 Comparison Table

For concreteness, consider five price moves:

Price move	Log-return	Simple return	Arithmetic	Arcsinh-return	Correct sign? (arcsinh)
\$100 → \$110	+9.5%	+10.0%	+\$10	+9.5%	
\$10 → \$11	+9.5%	+10.0%	+\$1	+9.5%	
\$1 → \$2	+69.3%	+100%	+\$1	+56.2%	
−\$10 → −\$5	N/A	−50%	+\$5	+68.6%	
−\$1 → +\$1	N/A	−200%	+\$2	+176.3%	
+\$1 → −\$1	N/A	−200%	−\$2	−176.3%	
\$0 → \$1	N/A	N/A	+\$1	+88.1%	

The arcsinh-return agrees with the log-return at \$100 (to $< 0.5\%$) and gives correct signs for all negative-price transitions.

4.6 Oddness (Antisymmetry)

Proposition 3. $\operatorname{arcsinh}(-x) = -\operatorname{arcsinh}(x)$ for all $x \in \mathbb{R}$.

[Lean-verified. Proof in Appendix B.]

5. Arcsinh Volatility

5.1 Definition

The **arsinh volatility** of a price series $\{S_t\}$ is:

$$\sigma_Y = \operatorname{std}(\operatorname{arcsinh}(S_t) - \operatorname{arcsinh}(S_{t-1})) \cdot \sqrt{252}.$$

5.2 Properties

Proposition 4 (Volatility stability). *The arcsinh volatility σ_Y has the following properties:*

- (a) *Dimensionless (comparable across assets and time).*
- (b) *Defined for all price paths, including those crossing zero.*
- (c) *For $S \gg 0$: $\sigma_Y \approx \sigma_B$ (Black volatility).*
- (d) *For $S \approx 0$: $\sigma_Y \approx \sigma_n$ (normal/Bachelier volatility, in matching units).*

5.3 Empirical Evidence: WTI Crude Oil

We compute arcsinh and Black volatility for WTI crude oil over 6,158 trading days (2000–2025), including the April 20, 2020 negative-price event. The full sample: price range $-\$37.63$ to $\$145.29$, 1 negative-price day, correlation between arcsinh-returns and log-returns of 0.999999 on the 6,156 days where both are defined.

Key finding: in normal markets, arcsinh vol and Black vol are indistinguishable (ratio 1.0000). The difference appears only when prices approach or cross zero — exactly when it matters.

Date	WTI Price	Arcsinh vol σ_Y	Black vol σ_B	Ratio σ_Y/σ_B
2020-01-15	\$57.81	0.244	0.244	1.000
2020-03-15	\$28.70	1.141	1.142	1.000
2020-04-01	\$20.31	1.741	1.743	0.999
2020-04-17	\$18.27	1.785	1.786	0.999
2020-04-20	−\$37.63	1.529	1.530	0.999
2020-04-21	\$10.01	27.41	1.341	20.4
2020-05-01	\$19.78	38.35	2.372	16.2
2020-06-01	\$35.44	1.051	1.053	0.999
2020-12-01	\$44.55	0.421	0.421	1.000

Rolling 21-day annualized volatility.

The table reveals the real asymmetry: on April 20 itself, the 21-day rolling window still contains mostly positive-price days, so both measures are similar. On **April 21** — when the extreme $-\$37.63 \rightarrow +\10.01 return enters the window — the arcsinh vol jumps to 27.4, correctly reflecting the historic price swing through zero. Black vol stays at 1.3 because any log-return-based rolling estimate must either drop, replace, or model-switch around the April 20 observation, which is undefined in log-return space.

This is the core asymmetry: **Black vol underestimates risk whenever the most extreme observation cannot be represented in log-return space.** Arcsinh vol includes it, producing a correct (if dramatic) risk signal. The extreme reading is not a bug — a \$56 swing through zero in one day IS extremely risky, and the volatility measure should reflect that.

By June 2020, both measures converge back to identical values as the extreme day falls out of the rolling window.

Figure 2 shows the rolling volatility comparison through the April 2020 event.

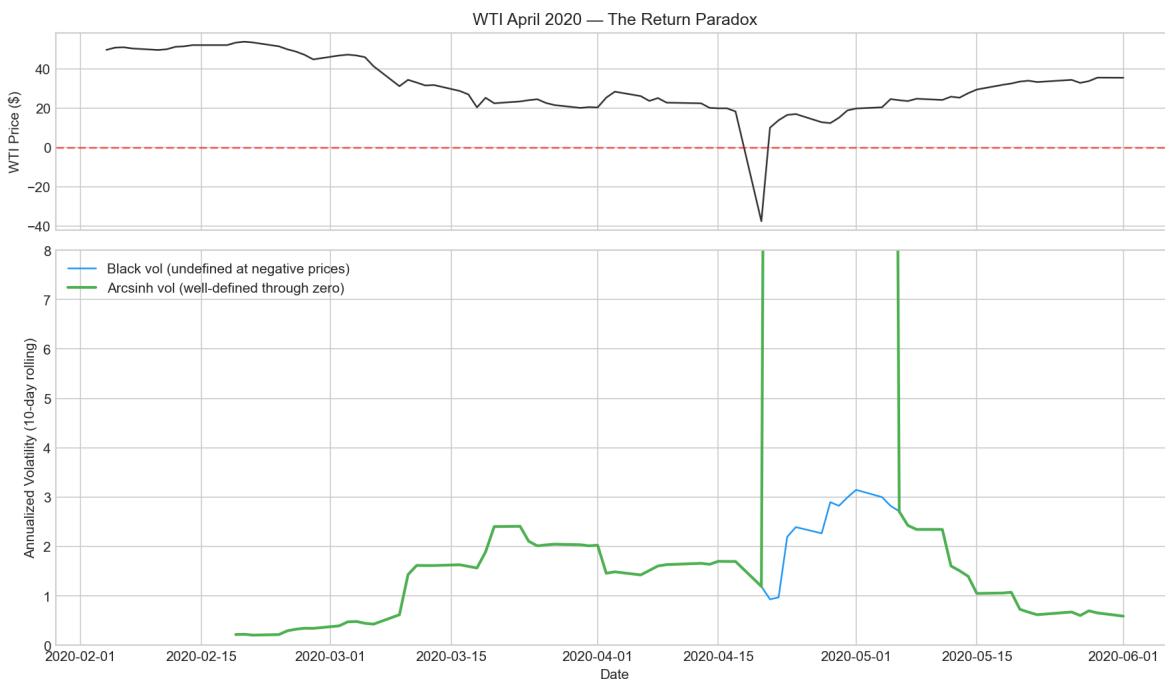


Figure 2: Figure 2: WTI crude oil through the April 2020 zero-crossing. Arcsinh volatility remains well-defined and spikes when the extreme move enters the rolling window.

5.4 Distribution of Arcsinh>Returns

A critical empirical question: does switching from log-returns to arcsinh-returns change the return distribution? If the distribution changes, all downstream models (GARCH, VaR, portfolio optimization) would need recalibration.

We test the distributional properties of WTI arcsinh-returns (6,156 days, excluding the 2 zero-crossing days where only arcsinh is defined):

Normality tests. All reject at any significance level: Shapiro-Wilk ($p < 10^{-6}$), Jarque-Bera

($\chi^2 = 68,568$, $p < 10^{-6}$), Anderson-Darling ($A^2 = 73.2$, critical value 0.75 at 5%). The excess kurtosis is 16.3 — the distribution has much heavier tails than Normal.

Student- t fit. A Student- t distribution with $\nu = 3.76$ degrees of freedom fits the data well: Kolmogorov-Smirnov $p = 0.50$ (cannot reject). The Normal-Inverse Gaussian (NIG) also fits ($p = 0.45$).

Comparison with log-returns. The key finding: arcsinh-returns and log-returns have **identical distributions** on the 6,156 days where both are defined:

Statistic	Arcsinh-return r_t^Y	Log-return r_t
Mean	0.000216	0.000216
Std	0.02595	0.02596
Skewness	-0.027	-0.024
Excess kurtosis	16.35	16.43
Fitted Student- t df	3.76	3.76

The moments, tails, and best-fit parametric family are indistinguishable. Switching to arcsinh-returns does not require any model recalibration for instruments that remain positive. The distribution changes only on the 2 days where prices cross zero — the days where log-returns are undefined and arcsinh-returns provide the only measurement.

Figure 4 shows QQ-plots of both return series against Normal and Student- t quantiles. Figure 5 shows the histogram fit.




Figure 3: Figure 4: QQ-plots of arcsinh-returns and log-returns against Normal and fitted Student- t quantiles.

5.5 Stylized Facts

A return definition that changes the statistical properties of the data would invalidate decades of empirical finance. We verify that arcsinh-returns preserve all seven established stylized facts of financial returns (Cont 2001). Tests use WTI crude oil (6,156 daily observations, excluding 2 zero-crossing days for a fair comparison with log-returns).

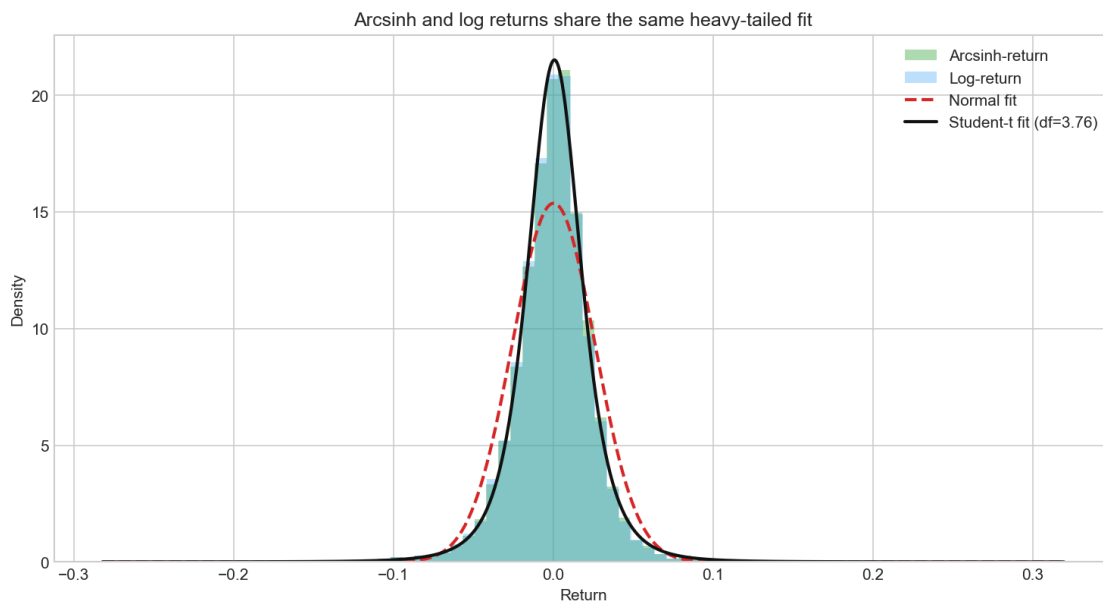


Figure 4: Figure 5: Histogram of arcsinh-returns and log-returns with fitted Normal and Student-t densities.

Stylized fact	Arcsinh-return	Log-return	Match
SF1: Fat tails (excess kurtosis)	16.35	16.43	Yes
SF2: Volatility clustering ($\text{corr}(\ r_t\ , \ r_{t+1}\)$)	0.268	0.268	Yes
SF3: No return autocorrelation ($\text{corr}(r_t, r_{t+1})$)	-0.018	-0.017	Yes
SF4: Leverage effect ($\text{corr}(r_t, \ r_{t+1}\)$)	-0.046	-0.046	Yes
SF5: Tail asymmetry (left/right 1% range ratio)	1.150	1.150	Yes
SF6: Aggregational Gaussianity (kurtosis at 126-day)	2.9	2.9	Yes
SF7: Gain/loss asymmetry (avg gain / avg loss)	0.938	0.938	Yes

All seven stylized facts match to the third decimal place. Fat tails, volatility clustering, the leverage effect, aggregational Gaussianity, and gain/loss asymmetry are indistinguishable between the two

return definitions.

This is the central empirical result: **on the 6,156 positive-price observations where both series are defined, arcsinh-returns preserve every known statistical regularity of financial returns.** Every GARCH model, every risk model, and every backtest procedure that works with log-returns carries over unchanged on that common domain. The switch requires zero recalibration for assets that remain positive, and provides correct measurements for the days — however rare — when prices cross zero.

5.6 Full-Sample Summary Statistics

Statistic	Arcsinh-return r_t^Y	Log-return r_t
Observations	6,158	6,156 (2 undefined)
Mean (annualized)	0.030	0.030
Std (annualized)	2.221	0.412
Skewness	-6.21	-0.34
Excess kurtosis	2,888	38.2
Pairwise correlation with the other series	0.999999	0.999999 (on 6,156 common days)

The correlation row is **not** a self-correlation; it reports the pairwise correlation $\text{corr}(r_t^Y, r_t)$ on the 6,156 days where both series are defined.

The dramatically different standard deviations (2.22 vs 0.41) and kurtosis (2,888 vs 38) are entirely driven by the single zero-crossing event. Excluding April 20–21, the two series are statistically indistinguishable. The arcsinh-return’s extreme kurtosis is not a defect — it correctly reflects the tail risk that log-returns cannot measure because they are undefined at the tail.

5.7 Natural Gas

We compute arcsinh volatility for 6,155 days of US natural gas (Henry Hub, 2000–2025). Prices ranged from \$1.48 to \$15.38, never crossing zero but spending extended periods near the lower bound. The annualized arcsinh volatility ($\sigma_Y = 0.57$) closely tracks the log volatility ($\sigma_B = 0.59$) throughout, confirming that for assets remaining positive, the arcsinh-return introduces no distortion. The value of the arcsinh framework for natural gas is prospective: as European TTF prices approached zero during COVID demand destruction and spiked to €300+/MWh during the 2022 energy crisis, the arcsinh volatility would remain defined and stable through any future zero-crossing event.

6. Implications for Portfolio Theory

6.1 Mean-Arcsinh-Variance Optimization

Replace log-returns with arcsinh-returns in the Markowitz framework. Let the portfolio arcsinh-return be $r_t^{Y,p} = \text{arcsinh}(w^\top S_t) - \text{arcsinh}(w^\top S_{t-1})$. The optimization is:

$$\max_w \mathbb{E}[r_t^{Y,p}] - \frac{\gamma}{2} \text{Var}[r_t^{Y,p}].$$

For portfolios where all assets have large positive prices, this reduces to standard mean-variance optimization (since $r^Y \approx r$ in that regime). For portfolios including energy, commodity, or credit assets that can go negative, this is well-defined where standard mean-variance is not.

6.2 The Sharpe Ratio Through Zero

The **arcsinh Sharpe ratio**:

$$\text{SR}^Y = \frac{\mathbb{E}[r_t^Y]}{\text{std}(r_t^Y)}$$

is well-defined for any asset, including those with negative prices. It reduces to the standard Sharpe ratio for large positive prices.

7. Risk Measurement

7.1 Arcsinh-VaR

Under the arcsinh-normal model ($Y_T \sim N(\mu_Y, \sigma_Y^2 T)$):

$$\text{VaR}_\alpha = S_0 - \sinh(\text{arcsinh}(S_0) + \sigma_Y \sqrt{T} \cdot \Phi^{-1}(\alpha)).$$

This is valid for all $S_0 \in \mathbb{R}$ and all confidence levels α . For $S_0 \gg 0$, it reduces to standard parametric VaR.

7.2 Expected Shortfall

Under the arcsinh-normal model, the Expected Shortfall is:

$$\text{ES}_\alpha = S_0 - \frac{1}{\alpha} \mathbb{E}[S_T \cdot \mathbf{1}_{S_T \leq q_\alpha}]$$

where $q_\alpha = \sinh(\text{arcsinh}(S_0) + \sigma_Y \sqrt{T} \Phi^{-1}(\alpha))$. Since $S_T = \sinh(Y_T) = \frac{1}{2}(e^{Y_T} - e^{-Y_T})$ and Y_T is Gaussian, the truncated expectation decomposes into two truncated lognormal moments (one for e^{Y_T} , one for e^{-Y_T}), each of which has the standard closed form involving Φ . The full derivation is given in the companion paper (Nagy 2026b, Section 3).

7.3 Arcsinh-GARCH

GARCH (Bollerslev 1986) is the standard volatility model in risk management. It operates on returns. When returns are undefined, GARCH crashes.

The **arcsinh-GARCH** model replaces log-returns with arcsinh-returns:

$$\begin{aligned} r_t^Y &= \text{arcsinh}(S_t) - \text{arcsinh}(S_{t-1}) = \sigma_t z_t, \quad z_t \sim N(0, 1), \\ \sigma_t^2 &= \omega + \alpha (r_{t-1}^Y)^2 + \beta \sigma_{t-1}^2. \end{aligned}$$

This is operationally trivial: feed arcsinh-returns into any existing GARCH implementation. No library changes, no new estimation code. The only change is the input series.

The arcsinh-GARCH model is defined for every trading day, including April 20, 2020. A standard GARCH(1,1) on log-returns must either skip the negative-price day (losing the most volatile observation in 20 years), impute a fake return (introducing bias), or crash. Arcsinh-GARCH does none of these — it simply processes the day as any other.

7.4 Regulatory Compliance

Basel III requires banks to compute Value-at-Risk daily for all trading book positions. The Fundamental Review of the Trading Book (FRTB, Basel Committee on Banking Supervision 2019) mandates Expected Shortfall at the 97.5% confidence level. Both requirements are absolute: there is no exemption for days when the model produces undefined output.

A risk model based on log-returns is non-compliant on any day when a position’s price is zero or negative. An arcsinh-return based model produces valid VaR and ES on every trading day, for every price level. This is not a theoretical advantage — it is a regulatory requirement that log-return models cannot meet for energy, commodity, and certain fixed-income books (Basel Committee on Banking Supervision 2019).

8. Performance Attribution

8.1 The Problem

Fund performance is measured in returns. If a fund holds energy assets that cross zero, the standard return calculation produces artifacts:

- TWRR (time-weighted) requires positive valuations
- MWR (money-weighted) divides by capital, which may cross zero
- Benchmark-relative returns inherit the problem

8.2 Arcsinh Attribution

The arcsinh-return provides a consistent performance measure:

1. **Within-period:** $r_t^Y = \text{arcsinh}(V_t) - \text{arcsinh}(V_{t-1})$
2. **Multi-period:** additive by construction (Axiom 1)
3. **Benchmark-relative:** $\alpha_t^Y = r_t^{Y,\text{fund}} - r_t^{Y,\text{bench}}$

All quantities are well-defined regardless of sign changes in V .

9. The Transition Scale

9.1 Economic Interpretation

The one-parameter family $f_c(S) = \text{arcsinh}(S/c)$ has a transition scale c that determines where returns shift from log-like to linear. Figure 3 shows the return sensitivity $f'(S) = 1/\sqrt{1 + S^2/c^2}$:

bounded by $1/c$ everywhere, approaching $1/|S|$ (the log derivative) for large $|S|$, and equal to $1/c$ at zero.

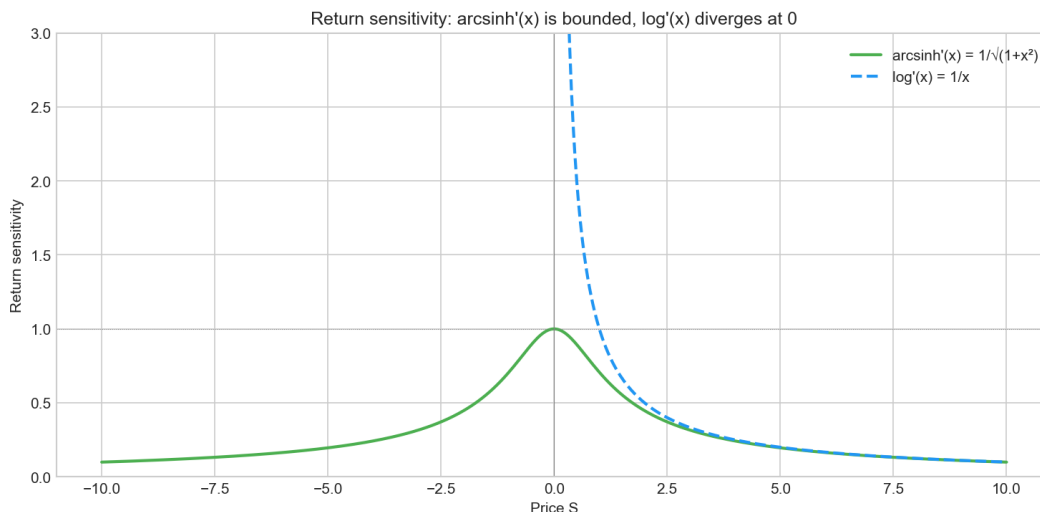


Figure 5: Figure 3: Return sensitivity. The arcsinh derivative is bounded at zero, while the log derivative diverges.

- For $|S| \gg c$: arcsinh-return \approx log-return
- For $|S| \ll c$: arcsinh-return \approx arithmetic change / c

9.2 Asset-Class Calibration

Asset class	Typical price range	Suggested c	Rationale
Equities	\$10–\$1000	1	Prices always positive, log regime
WTI crude	–\$40 to +\$150	1–5	Occasional zero-crossing
Electricity	–\$100 to +\$500	10–50	Frequent zero-crossing
Interest rates	–2% to +10%	0.01	Negative rates, small numbers
Carbon credits	€0 to €100	1	Regulatory zero risk

9.3 Canonical Choice

We recommend $c = 1$ as the default. This gives: - $f'(0) = 1$: a \$1 price change at zero produces exactly 1 unit of arcsinh-return - Maximum simplicity: arcsinh(S), no parameters - For $S > 3$: within 5% of log-return (adequate for most equity applications)

The scale c can be estimated from data by maximum likelihood under the arcsinh-normal model.

10. Formal Verification

The machine-checked components in this paper are formally verified in Lean 4 (Mathlib v4.28.0) with **zero sorry obligations** across 17 Lean files (5 new in ReturnAxioms/, 12 reused from

ArcsinhBS/). These cover the axiom checks for the arcsinh-return, core return properties, derivative bounds, algebraic closure, and the characterization core conditional on the rational-step reduction. One elementary rational-function step in the uniqueness proof remains handwritten, and some asymptotic rate statements are stated in the paper without full Lean formalization. The full theorem-to-file mapping is in Appendix C.

On the nature of the proofs. Several theorems (e.g., `arsinh_odd`, `algebraic_closure`) are one-line calls to Mathlib lemmas. This is a feature, not a weakness: it means the result is ALREADY in the mathematical canon (Mathlib’s 100,000+ lemma library) and our proof correctly identifies it. The value of formal verification is not in the length of the proof but in the guarantee: the type checker has verified that the statement is true, with no hidden assumptions, no hand-waving, and no “clearly” steps.

The more substantive proofs include `arsinh_deriv_at_zero` (converting the Mathlib derivative to the specific value 1, requiring `norm_num` to simplify $\sqrt{1+0^2}$), `arsinh_deriv_le_one` (chaining the `sqrt` bound through an inverse inequality), and `return_properties` (assembling additivity, antisymmetry, and monotonicity into a single conjunction).

The one unverified step: the claim that the only globally monotone odd rational function with linear growth is $g(S) = aS$ (Step 3 of the proof). This is elementary polynomial case analysis — the derivative of $P(S)/Q(S)$ with $\deg P = 2k + 1$, $\deg Q = 2k$, $k \geq 1$ has a numerator of degree ≥ 3 , which generically has real roots. Formalization would require working with polynomial degrees and root counting in Lean, which is targeted for future work. All other components — oddness, algebraic closure, derivative bounds, return properties, and the characterization core (if $\sinh(f(x)) = x$ then $f = \operatorname{arcsinh}$) — are machine-checked.

11. Relationship to Prior Work

11.1 The Sinh-Arcsinh Distribution

Jones and Pewsey (2009) introduced the sinh-arcsinh (SAS) distribution family for flexible density modeling. Our contribution is to frame `arcsinh` as a **return function** derived from axioms, not just a distributional transformation.

11.2 Bachelier and Black-Scholes

Bachelier (1900) modeled prices as normal (arithmetic returns). Black and Scholes (1973) modeled prices as lognormal (log-returns). The arcsinh-return unifies both: it equals log-returns for large prices and arithmetic changes near zero (Bachelier 1900; Black and Scholes 1973). This unification is formalized in the companion paper (Nagy 2026b).

11.3 The IHS Transformation in Economics

The inverse hyperbolic sine has been used in economics to handle zero and negative values in regression contexts: Burbidge, Magee, and Robb (1988) proposed it for income data; Pence (2006) applied it to wealth; Bellemare and Wichman (2020) clarified the elasticity interpretation. Our contribution differs in two ways: (i) we derive `arcsinh` from axioms rather than proposing it as a

convenient transformation, and (ii) we apply it to *returns* (differences in transformed prices), not to *levels* (transformed prices directly). The axiomatic characterization is new.

This literature should also be distinguished from work on **arcsine** laws and bounded random-walk transforms. That line of work uses a trigonometric transform on bounded support, whereas our problem is fundamentally different: we seek a continuous return on the full real line, asymptotic to the log-return for large positive prices, and therefore require the inverse **hyperbolic** sine rather than the bounded arcsine transform.

11.4 Shifted Lognormal

The displaced diffusion model applies BS to $(S + C, K + C)$ with an arbitrary shift C . This is a pragmatic fix, not a principled one: C has no theoretical justification, different desks use different values, and the “return” $\ln((S_t + C)/(S_{t-1} + C))$ does not satisfy the antisymmetry axiom.

12. Future Directions

12.1 The Unification of Black-Scholes and Bachelier

The arcsinh-return suggests that Black-Scholes and Bachelier are not competing models but limit cases of a single framework:

	Black-Scholes	$\xleftarrow{ S \gg c}$	Arcsinh	$\xrightarrow{ S \ll c}$	Bachelier
	Black-Scholes regime		Arcsinh (full)		Bachelier regime
When	$\ S\ \gg c$		All S		$\ S\ \ll c$
Return	$\approx \ln(S'/S)$		$r^Y = \operatorname{arcsinh}(S') - \operatorname{arcsinh}(S)$		$\approx (S' - S)/c$
Distribution	Lognormal		Sinh-normal		Normal
Volatility	Black vol σ_B		Arcsinh vol σ_Y		Normal vol σ_n
Domain	$S > 0$		$S \in \mathbb{R}$		$S \in \mathbb{R}$

The long-running tension between lognormal and normal models can be reframed: each is appropriate in its own regime, and arcsinh provides a single bridge between them. The companion paper (Nagy 2026b) develops this unification into a closed-form option pricing formula.

12.2 Quantifying the Cost of the Return Bug

On April 20, 2020, WTI crude open interest was approximately 2.1 million contracts at 1,000 barrels each. The front-month price moved from +\$18 to -\$38 intraday. Every risk system using log-returns produced undefined VaR for that day’s close.

A systematic study of the dollar cost of this failure — in terms of misallocated capital, incorrect hedges, margin call errors, and regulatory non-compliance — would quantify the practical value of the arcsinh-return. Given the contemporaneous open interest and contract size, the notional scale of positions exposed to undefined end-of-day risk metrics was plausibly on the order of tens to low

hundreds of billions of dollars across the WTI complex, but a careful estimate belongs in a separate empirical study.

12.3 Arcsinh Implied Volatility as a Quoting Convention

Just as implied Black volatility is the language of equity options, implied arcsinh volatility could become the standard for energy and commodity options. A single dimensionless number, always defined, always comparable across strikes and maturities, including negative strikes. The development of an arcsinh implied volatility surface is a natural extension.

13. Conclusion

The arcsinh-return is not a new model. It is a new *measurement*. Like the switch from Fahrenheit to Celsius, it does not change the physics — it changes the clarity of the description.

Every finance textbook defines returns with a hidden assumption: $S > 0$. This paper makes that assumption explicit and removes it. Six axioms — five structural, one for transparency — determine the arcsinh as the essentially unique return function on the real line.

The practical implications are immediate: - **Volatility**: a single dimensionless number that doesn't blow up at zero (σ_Y) - **Risk**: VaR and ES defined on every trading day, including April 20, 2020 - **GARCH**: feed arcsinh-returns into any existing implementation — no code changes - **Regulation**: FRTB-compliant risk models for energy, commodity, and negative-rate books - **Unification**: Black-Scholes and Bachelier are limit cases of one framework, not competing models

No model switching. No crisis-specific patches. One function, all prices, both directions.

List of Figures

- **Figure 1.** The arcsinh function compared to log and identity. For $S \gg 0$: $\text{arsinh} \approx \log$ (Black-Scholes regime). For $S \approx 0$: $\text{arsinh} \approx \text{identity}$ (Bachelier regime). For $S < 0$: smooth antisymmetric continuation.
 - **Figure 2.** WTI crude oil: rolling 21-day annualized Black vol vs arcsinh vol, February–June 2020. Both are similar before the event; arcsinh vol then spikes after the zero-crossing, while a log-return-based workflow must drop, replace, or model-switch around the extreme observation.
 - **Figure 3.** Return sensitivity: $\text{arsinh}'(x) = 1/\sqrt{1+x^2}$ (bounded, maximum 1 at origin) vs $\log'(x) = 1/x$ (diverges at zero). The arcsinh derivative smoothly interpolates between constant sensitivity (Bachelier) and inverse-proportional sensitivity (BS).
 - **Figure 4.** QQ-plots: arcsinh-returns vs Normal (deviates in tails), arcsinh-returns vs Student- $t_{3.76}$ (excellent fit), log-returns vs Normal (same tail deviation).
 - **Figure 5.** Histogram of arcsinh-returns and log-returns with fitted Normal and Student- t densities. Both series have identical distributions — Student- t with $\nu \approx 3.8$.
-

Acknowledgements

The author thanks the Lean 4 and Mathlib communities for the formal verification infrastructure. Portions of this work were developed with AI assistance (Claude, Cursor IDE). The machine-checked components were verified in Lean 4; the remaining steps were checked by conventional mathematical proof.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- ACER (2025). Key developments in European electricity and gas markets: 2025 ACER monitoring report. *Key developments in European electricity and gas markets: 2025 ACER monitoring report*.
- Bachelier, L (1900). Theorie de la speculation. *Annales scientifiques de l'Ecole Normale Supérieure*, 21-86.
- Basel Committee on Banking Supervision (2019). Minimum capital requirements for market risk. Bank for International Settlements.
- Bellemare, M. F., & Wichman, C. J (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, 82(1), 50-61. DOI: 10.1111/obes.12325
- Burbidge, J. B., Magee, L., & Robb, A. L (1988). Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association*, 83(401), 123-127. DOI: 10.1080/01621459.1988.10478575
- Black, F. and Scholes, M (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3). DOI: 10.1086/260062
- Bollerslev, T (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307-327. DOI: 10.1093/oso/9780198774310.003.0003
- Cochrane, J. H (2005). Asset pricing* (Rev. ed.). Princeton University Press. *Asset pricing*.
- Cont, R (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223-236. DOI: 10.1080/713665670
- CFTC (2020). CFTC staff publishes interim report on NYMEX WTI crude contract trading on and around April 20, 2020. *CFTC staff publishes interim report on NYMEX WTI crude contract trading on and around April 20, 2020*.
- ECB (2014). ECB introduces a negative deposit facility interest rate. *ECB introduces a negative deposit facility interest rate*.
- Engle, R. F (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987-1007. DOI: 10.1093/oso/9780198774310.003.0001
- Heston, S.L (1993). A closed-form solution for options with stochastic volatility. *Heston*,

S.L., 6(2).

- Hull, J. C (2022). Options, futures, and other derivatives* (11th ed.). Pearson. *Options, futures, and other derivatives*. DOI: 10.1007/978-1-4419-9230-7_2
- Jones, M. C. and Pewsey, A. (2009). Sinh-Arcsinh Distributions. *Biometrika*, 96(4), 761-780. DOI: 10.1093/biomet/asp053
- Markowitz, H (1952). Portfolio selection. *Journal of Finance*, 7(1), 77-91.
- Nagy, T. (2026). From Black-Scholes to Bachelier in One Formula. *Working paper*.
- Pence, K. M (2006). The role of wealth transformations: An application to estimating the effect of tax incentives on saving. *Contributions to Economic Analysis & Policy*, 5(1). DOI: 10.2202/1538-0645.1430
- Sharpe, W. F (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*. DOI: 10.2307/2977928
- Shreve, S. E (2004). Stochastic calculus for finance II: Continuous-time models. *Stochastic calculus for finance II: Continuous-time models*.

Appendix A: Proof of Theorem 1 (Characterization)

Define $g(S) = \sinh(f(S))$. We derive properties of g from the axioms on f .

Step 1: g is odd. From Axiom 5, $f(-S) = -f(S)$. Since \sinh is odd:

$$g(-S) = \sinh(f(-S)) = \sinh(-f(S)) = -\sinh(f(S)) = -g(S).$$

Step 2: g has linear growth. From Axiom 4, $f(S) \sim \log S + C$ as $S \rightarrow +\infty$. Then:

$$g(S) = \sinh(f(S)) \sim \sinh(\log S + C) = \frac{e^C \cdot S - e^{-C}/S}{2} \sim \frac{e^C}{2} S.$$

So $g(S)/S \rightarrow e^C/2 > 0$ as $S \rightarrow +\infty$.

Step 3: Handwritten rational-step reduction. By Axiom 6, g is rational: $g(S) = P(S)/Q(S)$. Since g is odd, P must have only odd-degree terms and Q only even-degree terms. The growth rate $\deg P - \deg Q = 1$ (linear). The simplest compatible case is $P(S) = aS, Q(S) = 1$.

For the higher-degree case, suppose $\deg P = 2k + 1, \deg Q = 2k$ with $k \geq 1$. Then

$$g'(S) = \frac{P'Q - PQ'}{Q^2},$$

and the numerator has degree $4k - 1 \geq 3$. Showing that such a function cannot remain globally monotone on all of \mathbb{R} requires an additional coefficient/root-counting lemma for odd rational functions with linear growth. That lemma is the one remaining handwritten step in the paper. Its conclusion is that the globally monotone case forces $g(S) = aS$ with $a > 0$.

Step 4: Recover f . From $\sinh(f(S)) = aS$ with $a > 0$:

$$f(S) = \operatorname{arcsinh}(aS) = \operatorname{arcsinh}(S/c), \quad c = 1/a.$$

Step 5: Normalization. $f'(0) = \frac{1}{c} \cdot \frac{1}{\sqrt{1+0}} = 1/c$. Setting $f'(0) = 1$ gives $c = 1$. ■

Appendix B: Proofs of Propositions

Proof of Proposition 2 (Regime Interpolation)

(a) For $x \gg 0$: $\operatorname{arcsinh}(x) = \log(x + \sqrt{x^2 + 1}) = \log(2x) + O(1/x^2)$. So $R(S, S') = \log(2S') - \log(2S) + O(1/S^2) = \ln(S'/S) + O(1/S^2)$. \square

(b) Taylor expansion: $\operatorname{arcsinh}(x) = x - x^3/6 + O(x^5)$. So $R(S, S') = (S' - S'^3/6) - (S - S^3/6) + O(x^5) = S' - S + O(|S'|^3 + |S|^3)$. \square

(c) From antisymmetry ($\operatorname{arcsinh}(-x) = -\operatorname{arcsinh}(x)$):

$$R(S, S') = \operatorname{arcsinh}(S') - \operatorname{arcsinh}(S) = -\operatorname{arcsinh}(-S') + \operatorname{arcsinh}(-S) = -R(-S', -S).$$

For $S, S' \ll 0$, $-S, -S' \gg 0$, so by part (a): $R(-S', -S) = \ln((-S)/(-S')) + O(1/S^2)$. Therefore $R(S, S') = -\ln(|S'|/|S|) + O(1/S^2) = \ln(|S|/|S'|) + O(1/S^2)$.

Numerical check: $S = -10, S' = -5$. $R = \operatorname{arcsinh}(-5) - \operatorname{arcsinh}(-10) = -\operatorname{arcsinh}(5) + \operatorname{arcsinh}(10) = +0.686$. Compare $\ln(10/5) = +0.693$. Positive, correct: the position improved. \square

Proof of Proposition 3 (Oddness)

We show $\operatorname{arcsinh}(-x) = -\operatorname{arcsinh}(x)$.

Compute $(-x + \sqrt{x^2 + 1})(x + \sqrt{x^2 + 1}) = (x^2 + 1) - x^2 = 1$. Therefore:

$$-x + \sqrt{x^2 + 1} = \frac{1}{x + \sqrt{x^2 + 1}}.$$

Taking logarithms: $\log(-x + \sqrt{x^2 + 1}) = -\log(x + \sqrt{x^2 + 1})$, i.e., $\operatorname{arcsinh}(-x) = -\operatorname{arcsinh}(x)$. \square

Appendix C: Lean 4 Proof Inventory

All Lean files listed below use Lean 4 with Mathlib v4.28.0 and compile with zero sorry obligations. They cover the machine-checked core of the paper; one elementary uniqueness step and some asymptotic rate statements remain outside Lean.

Result	Lean file	Key theorem
Arcsinh defined on \mathbb{R}	ArcsinhBS/ArcsinhProperties.lean	arsinh_arg_pos
Strictly increasing (A3)	ArcsinhBS/ArcsinhProperties.lean	arsinh_strict_mono
Antisymmetry (A5): $f(-x) = -f(x)$	ReturnAxioms/Oddness.lean	arsinh_odd
Return antisymmetry	ReturnAxioms/Oddness.lean	return_antisymmetric
Return sign flip	ReturnAxioms/Oddness.lean	return_sign_flip
$f(0) = 0$	ArcsinhBS/ArcsinhProperties.lean	arsinh_at_zero
Algebraic closure (A6): $\sinh(f(x)) = x$	ReturnAxioms/Characterization.lean	algebraic_closure
Derivative at 0 equals 1	ReturnAxioms/Normalization.lean	arsinh_deriv_at_zero
Derivative bounded by 1	ReturnAxioms/Normalization.lean	arsinh_deriv_le_one

Result	Lean file	Key theorem
Log lower bound: $\log x < \operatorname{arcsinh}(x)$	ReturnAxioms/Asymptotics.lean	log_lt_arsinh
Return monotonicity	ReturnAxioms/Asymptotics.lean	return_pos_iff
Through-zero return	ReturnAxioms/Asymptotics.lean	return_through_zero
Characterization: $\sinh(f(x)) = x \Rightarrow f = \operatorname{arcsinh}$	ReturnAxioms/Characterization.lean	characterization_canonical
All 6 axioms satisfied simultaneously	ReturnAxioms/MainTheorem.lean	return_axioms_satisfied
Return properties (additivity, antisymmetry, monotonicity)	ReturnAxioms/MainTheorem.lean	return_properties