

# Trading Strategies and Algorithms

A practical tour of the major strategy families, the execution algorithms that implement them, and the risk controls that keep them alive.

## 1. Two separate questions

Every time an order is placed, two independent questions must be answered. The **strategy** decides *what* to buy or sell and *when*: the signal that justifies the trade, the size, the stop-loss. The **execution algorithm** then decides *how* the chosen order is split, routed, and placed in the market to minimise transaction costs and information leakage.

Confusing the two is a common mistake. A brilliant idea can be destroyed by poor execution — and a flawless execution algorithm cannot rescue a bad signal. The rest of this document treats them as distinct layers.

## 2. Strategy taxonomy

Most systematic strategies fall into one of five families. They differ in the market phenomenon they try to capture, their typical holding period, and the kind of risk they take on.

Family	What it exploits	Typical horizon	Main risk
Trend-following / momentum	Persistence of returns — winners keep winning over weeks to a year.	1 month – 1 year	Sharp reversals; whipsaw in sideways markets.
Mean reversion	Temporary mispricings revert to a statistical anchor (moving average, fair value).	Minutes – weeks	The trend stays and the position bleeds; 'picking up pennies'.
Carry	Earning the spread between a high-yielding and a low-yielding asset (FX carry, futures roll, credit).	Weeks – years	Sudden risk-off events that collapse the spread.
Factor / smart-beta	Long-run premia from characteristics such as value, size, quality, low-vol, momentum.	Months – years	Long drawdown periods; factor crowding and fading.
Market-making / arbitrage	The bid-ask spread, or tiny price differences across venues/instruments.	Milliseconds – minutes	Adverse selection, inventory risk, technology failure.

## 3. Trend-following and momentum

The oldest systematic family. The empirical observation, documented across equities, futures, FX, and crypto for decades, is that assets that have gone up recently tend to keep going up for a while — and vice versa. The signal is direction-persistence.

Classic implementations rank a universe of instruments by their recent returns and hold the top decile long (optionally the bottom decile short). The 12-1 momentum factor — the last 12 months' return skipping the most recent month — is the textbook formulation for cross-sectional equity momentum. Time-series momentum instead trades each asset against itself: go long if the asset's own trailing return is positive, short if negative.

- **Signal:** past return over a lookback window (typically 3 – 12 months), optionally volatility-scaled or smoothed with a moving average.
- **Entry:** cross-sectional rank crosses a threshold, or the signal flips sign.
- **Exit:** trailing stop in units of ATR (2 – 3x is common), or the signal flips back.
- **Sizing:** inverse-volatility — smaller positions in noisy assets, larger in quiet ones, so that each position contributes the same risk budget.

Trend-following works because investors are slow to react to new information (underreaction in the short term) and because risk-off episodes produce persistent trends. Its Achilles' heel is *whipsaw*: in range-bound, choppy markets every entry is quickly reversed by a counter-move, and realized Sharpe drops to zero for long stretches. Robust trend systems accept this and survive by having enough breadth across asset classes and by sizing small.

## ATR-based exits

The exit rule mentioned above — *trailing stop in units of ATR (2 – 3x is common)* — deserves unpacking, because it is the workhorse risk-management primitive in trend systems. ATR (Average True Range) measures the typical daily price swing of an instrument and lets the stop adapt to each asset's own volatility instead of a one-size-fits-all percentage.

## How ATR is computed

**True Range (TR)** for each bar is the largest of three distances, chosen so that overnight gaps are not missed by a plain *high – low*:

$$TR = \max(\text{high} - \text{low}, |\text{high} - \text{prev\_close}|, |\text{low} - \text{prev\_close}|)$$

**ATR** is then a smoothed average of TR over  $N$  periods (typically  $N = 14$ ). Two conventions are in use:

- **Simple ATR:** arithmetic mean of the last  $N$  TRs. Easy, but recomputes from scratch each bar.
- **Wilder's ATR** (the original used by most platforms) — exponential smoothing with  $\alpha = 1/N$ :

$$ATR\_today = ((N - 1) \times ATR\_yesterday + TR\_today) / N$$

The first ATR value seeds with the simple mean of the initial  $N$  TRs; subsequent bars update incrementally. ATR is expressed in **price units** (dollars), so it scales naturally with the instrument — a \$500 stock will have a much larger ATR than a \$20 stock without any difference in volatility. To compare across instruments use **ATR / price** (ATR%).

## What 2 – 3x ATR means in practice

The ATR multiple  $k$  defines how far the price is allowed to move against the position before the trade is closed. Two common variants:

- **Initial stop:**  $\text{stop} = \text{entry} - k \times \text{ATR}$  (long) or  $\text{entry} + k \times \text{ATR}$  (short). Example: long entry at \$100 with  $\text{ATR} = \$2$  and  $k = 3$  places the stop at \$94.
- **Trailing / Chandelier stop:**  $\text{stop} = \text{max\_close\_since\_entry} - k \times \text{ATR}$ . The stop ratchets upward as new highs are made but never moves down — locking in profit while letting the trend run.

The 2 – 3x range is the classic risk/whipsaw trade-off:

- **$k = 2x$**  — tight stop. Smaller per-trade losses, but you get knocked out by routine noise; better suited to short-term/swing systems.
- **$k = 3x$**  — looser stop. Lets normal volatility pass through, captures bigger trends, but the realized loss when a trade does fail is larger; standard for medium-term trend followers.

ATR also drives **position sizing** in the same family of systems:  $\text{size} = (\text{risk budget}) / (k \times \text{ATR})$ . With a fixed risk budget per trade (e.g. 0.5% of equity) and a  $k \times \text{ATR}$  stop, a noisier instrument receives a smaller position automatically — equalising risk across the book.

## 4. Mean reversion

The mirror image of trend-following. The assumption is that prices oscillate around a statistical centre of gravity — a moving average, a volume-weighted average, a cointegrated pair's long-run ratio — and deviations get reabsorbed.

The canonical retail example is Bollinger Bands: when the close sits more than two standard deviations below its 20-day moving average, buy; close when it returns to the middle band. More sophisticated implementations use Ornstein–Uhlenbeck fits, Kalman filters to re-estimate fair value online, or z-score baskets across many correlated names.

- **Holding period** is shorter than trend: often intraday to a few days.
- **Hit rate** is high (60 – 70% winners), but wins are small and losses can be enormous if the mean shifts. The classic risk description is *picking up pennies in front of a steamroller*.
- **Works best** in liquid, low-volatility regimes and range-bound markets — precisely when trend-following struggles. The two families are natural diversifiers.

A dangerous mistake is applying mean reversion to an asset that has entered a structural trend (earnings miss, regime change, regulatory shock). The model keeps signalling 'buy the dip' while the dip deepens. Regime detection filters (ADX, realized vol, drawdown depth) are commonly used to disable the strategy during such periods.

## 5. Statistical arbitrage

Mean reversion applied to *relative* rather than absolute prices. Find two (or more) assets whose prices have moved together for a long time — their spread is **cointegrated** — and trade the spread when it deviates from its historical mean.

The most famous form is **pairs trading**. Ranking thousands of pairs by historical correlation and cointegration, the strategy goes long the under-performer and short the out-performer when their

spread widens beyond, say, two standard deviations, and unwinds when it contracts.

$$\text{Spread z-score} = (\text{spread} - \mu) / \sigma, \text{ entry } |z| > 2, \text{ exit } |z| < 0.5$$

The beauty of pairs trading is that it is *dollar-neutral and beta-neutral*: market-wide moves cancel, leaving only the idiosyncratic convergence. The ugliness is that cointegration relationships are not stable: a merger, a spin-off, or a regulatory change can destroy them overnight, at which point the spread keeps widening and the trade becomes a pure directional loss. Modern stat-arb books manage this with *hundreds* of small pairs, strict half-life cut-offs, and automatic stop-outs when cointegration tests fail.

## 6. The VIX — trading volatility and predicting SPX

The **CBOE Volatility Index (VIX)**, often called the *fear gauge*, measures the market's expectation of S&P 500 volatility over the next 30 days. It is computed in real time from the implied volatilities of a wide strip of SPX options across many strikes and two nearest expirations, then annualised. A VIX of 20 means the options market is pricing approximately  $\pm 20\%$  annualised volatility — or roughly  $\pm 5.8\%$  for the next 30 days. The long-run average is near 19–20; readings below 12 signal complacency, readings above 30 signal stress, and readings above 40 are associated with outright panic (2008 financial crisis: 80+, March 2020: 85+).

### How VIX is traded

VIX itself is a calculation, not an asset — you cannot buy or sell spot VIX directly. Exposure is obtained through three instruments:

- **VIX Futures** — monthly contracts listed on CBOE, settled in cash at expiry against the SOQ (Special Opening Quotation) of VIX. Because the market almost always expects volatility to revert to its long-run average, the futures curve is typically in *contango* (front-month futures > spot VIX). A trader long the front-month future must roll it into the next month before expiry, and that roll is almost always done at a higher price — a structural headwind that erodes long positions over time. The mirror image is true for shorts: contango is a tail-wind for short-VIX traders.
- **VIX ETPs (ETFs and ETNs)** — these products track a rolling exposure to VIX futures, not to spot VIX. The distinction matters enormously: during calm periods the spot-to-futures gap (contango drag) decays the ETP value even when spot VIX is flat. Main products: VXX and VIXY (1× long, front-month roll), UVXY (1.5× leveraged long), and SVXY (−0.5× inverse, i.e. short volatility). Long ETPs have historically lost most of their value over any multi-year window; inverse ETPs tend to compound positively in calm regimes but suffer catastrophic drawdowns in vol spikes.
- **VIX Options** — options on VIX futures (not on spot VIX), cash-settled. Used to build structured exposure: long call spreads as cheap tail-hedge insurance, or short puts to collect premium from the volatility risk premium without the unlimited downside of an uncovered short.

### VIX ETP mechanics in depth

Since VIX cannot be held directly, ETPs must roll futures contracts monthly — selling the expiring front-month contract and buying the next-month. In contango (the normal state), next-month futures trade *above* the front-month: the ETP systematically *sells low and buys high* on every roll. This is the **contango drag**: even if spot VIX never moves, the ETP loses value on each roll cycle — often 5–10% per month in deep contango environments.

Ticker	Type	Exposure	Characteristic
VXX, VIXY	Long	1× front-month roll	Steady decay in calm markets; provides crisis payoff when VIX spikes sharply.
UVXY	Long leveraged	1.5× (was 2× pre-2018)	Decays faster than VXX; leverage was cut after Feb 2018 Volmageddon.

<b>SVXY</b>	Inverse	-0.5x (was -1x pre-2018)	Collects contango as income; halved to -0.5x after near-wipeout in Feb 2018.
-------------	---------	--------------------------	--

**Long ETP decay.** VXX has lost roughly 99.9% of its value since its 2009 launch (through multiple reverse splits). You are paying contango drag every trading day for insurance you may never collect on. The only scenario where a long VXX position pays off sufficiently to offset the carry cost is a fast, large vol spike — and even then the futures-spot basis means you capture only a fraction of the VIX move.

**Inverse ETP danger.** SVXY collects contango as steady income in calm regimes — a reliable carry trade. But VIX can spike 5–10x in days: the Feb 2018 'Volmageddon' event wiped out -1x inverse ETPs in a single session; March 2020 produced a similar drawdown. The -0.5x redesign reduces but does not eliminate this tail risk. Pre-committing exits (covering when VIX crosses 25, or when term structure inverts) is non-negotiable.

The core tension: long VIX ETPs are expensive insurance that bleeds in calm markets and pays in crises; inverse ETPs earn steady carry that occasionally blows up catastrophically. Neither is a viable long-term hold in isolation — they are tactical instruments for short-term hedging, speculation, or mean-reversion strategies with strict risk controls.

The key practical rule: *VIX ETPs track futures, not spot VIX.* A VIX spike from 15 to 35 will produce only a fraction of that move in VXX because the front-month future was already trading above 20 before the spike. Never use VXX to hedge an equity portfolio without accounting for the contango drag and the futures-spot basis.

## VIX as a predictor of SPX direction

VIX and the S&P 500 have a strong contemporaneous inverse relationship (correlation roughly -0.7 to -0.8 on daily returns): when equity markets sell off sharply, demand for put-option protection drives implied volatility up, and VIX rises. The causal direction runs both ways — equity price moves drive VIX, and VIX-regime signals can be used to anticipate SPX behaviour:

## Extreme VIX readings as contrarian signals

When markets crash fast, investors panic and rush to buy put options for protection. That buying pressure drives implied volatility — and therefore VIX — sharply higher, often *above the actual volatility of the move itself*. Options markets reflect maximum fear, not a rational forecast: participants price in scenarios that rarely materialise (a 2008-style depression, a complete market breakdown). VIX above 35–40 historically marks the *tail end* of fear, not the beginning of further catastrophe. The overshoot unwinds in two steps:

- **VIX spikes:** put buyers drive implied vol above the actual market move — implied vol > realized vol.
- **VIX mean-reverts:** the fear premium deflates, lowering required risk premia and lifting equity prices as VIX converges back toward its long-run average (~19–20).

Historically, buying SPX within 5% of a VIX peak above 35 has produced strongly positive forward returns. The 3- and 6-month horizons are the most reliable; the 1-month window is noisier because equities can keep sliding slightly while VIX is already declining from its peak.

Event	VIX Peak	6m SPX return	Note
March 2020 (COVID)	~85	+60%	Liquidity panic; Fed backstop accelerated the recovery. Signal worked.
Dec 2018 (rate fears)	~36	+25%	Fed pivot ended the sell-off. Signal worked.
Feb 2018 (Volmageddon)	~50	+15%	Vol-targeting unwind, no recession. Signal worked.
Oct 2008 (GFC)	~80	Negative	Structural/solvency crisis — banks insolvent. Signal failed; bottom took 18 months.

**When the signal fails.** The exception is a *structural* crisis — bank insolvency, credit market freeze — where the underlying damage is severe enough to keep equities falling after VIX peaks. The signal is most reliable in *liquidity-driven* panics where fear outpaces the fundamental damage. This distinction (liquidity shock vs. solvency shock) is the key filter before acting on an extreme VIX reading.

- **VIX term structure as a regime filter.** When the VIX futures curve is in contango (front-month < 2nd-month, i.e. the market expects a calm near term), SPX tends to trend positively. When the curve inverts into *backwardation* (front-month > back-month), it signals acute near-term stress; trend strategies perform poorly and mean reversion systems face elevated tail risk. The VIX3M / VIX ratio is commonly used: values below 1 indicate backwardation.
- **VIX level as a regime classifier.** Many systematic traders split their universe into three regimes by VIX level: low-vol (<15) favours carry and mean-reversion; medium (<25) is broadly friendly to trend and factor strategies; high (>25) triggers defensive overlays or outright position cuts. The specific thresholds are calibrated by backtesting; the logic is that in high-vol regimes correlations spike, diversification collapses, and the normal rules break down.
- **Intraday VIX for short-term signals.** In intraday strategies, a sharp VIX move of more than 5% in 30 minutes is associated with dislocated order flow; some practitioners use this as a momentary 'do not trade' flag to pause systematic algos until the move stabilises.

### The VIX-based algorithm: harvesting the volatility risk premium

The app implements this as the **Volatility Risk Premium** (short VIX) strategy under the Carry family. The theoretical foundation is simple: implied volatility on equity index options consistently trades *above* the subsequent realized volatility — historically by 3–4 vol points on average. This gap exists because end-investors pay a structural insurance premium for crash protection via puts. Selling that protection harvests the premium.

- **Entry condition:** the VIX futures term structure is in contango (front-month future trades above spot VIX, or equivalently VIX3M/VIX > 1.0) AND realized 30-day SPX vol is below current VIX — confirming the premium is real, not a compensation for actual elevated risk. The position is initiated by shorting front-month VIX futures (or selling ATM SPX straddles) with notional sized to a defined volatility budget.

- **Roll yield:** in contango, each monthly roll captures the futures-to-spot convergence. As expiry approaches, the front-month future must converge downward to the spot VIX settlement; short sellers collect this difference as P&L on each roll.
- **Exit / stop conditions:** the short is covered (closed) when VIX spikes through a regime threshold (typically  $VIX > 25$  from a calm baseline below 18), when the term structure flips into backwardation (front-month  $>$  back-month), or on a hard 'max-loss' stop expressed as a multiple of initial premium received. Discipline on the exit is non-negotiable: the Feb 2018 'Volmageddon' event wiped out the XIV ETP in a single session because exits were not pre-committed.
- **Tail-hedge overlay:** institutional implementations spend approximately 10% of premium income on far-out-of-the-money SPX puts or low-strike VIX calls. This converts the payoff from 'insurance-writer' (unlimited loss, bounded gain) to a defined-risk structure that survives a 2008-style vol explosion.
- **Sizing:** because the strategy is negatively skewed, position size is set conservatively — typically 5–15% of total portfolio risk budget. Running it alongside trend-following (which benefits from the same high-vol regimes that hurt short-VIX) is the canonical diversifier.

Short-VIX trades look like a steady income stream in backtests and live trading during calm regimes — until they don't. A full understanding of VIX mechanics (contango drag, futures basis, term structure inversion) is a prerequisite before allocating real capital.

## 7. Carry

Carry strategies earn a return simply by *holding* an instrument, independent of price moves: the coupon on a bond, the dividend on a stock, the interest-rate differential in an FX pair, the contango/backwardation in a futures curve. The bet is that the return from holding will exceed the risk premium priced in.

- **FX carry** — borrow in a low-rate currency, lend in a high-rate one. Profits unless the high-rate currency depreciates enough to offset the rate differential. Famous for sudden 'carry crashes' in risk-off events.
- **Commodity futures carry** — go long markets in backwardation, short markets in contango; the 'roll return' is the carry.
- **Equity carry** — dividend-yield strategies, or the volatility-risk-premium (short VIX) trade, which earns the gap between implied and realized variance.

Carry trades look smooth for long stretches and then suffer catastrophic shocks when the priced-in risk finally materialises. Their payoff profile resembles writing insurance: small positive drift most of the time, occasional large losses.

## 8. Factor investing and smart beta

Decades of academic work have documented **factors** — systematic characteristics of securities that have historically earned a return premium not explained by market beta alone. Smart-beta or factor-investing strategies tilt portfolios toward these characteristics.

- **Value** — cheap stocks (low P/E, low P/B) beat expensive ones on average. Fama & French (1992).

- **Size** — small-cap stocks beat large-cap on average, with higher volatility.
- **Momentum** — see section 3. Also a formal factor in multi-factor models.
- **Quality** — profitable, low-debt, stable-earnings firms outperform their peers.
- **Low volatility** — the counter-intuitive finding that low-vol stocks earn the same or higher returns than high-vol, with a much better Sharpe. Known as the *low-vol anomaly*.
- **Investment** — firms that invest conservatively (low asset growth) outperform firms that expand rapidly.

Factor strategies are run as long-only tilts (most retail smart-beta ETFs), long-short market-neutral books (many quant hedge funds), or multi-factor composites that blend several premia. Risk is concentrated in *factor drawdowns*: a given factor can underperform for 5 – 10 years (as value did between 2007 and 2020). Diversification across factors is therefore essential, and even then there are correlated bad periods like the 'quant quake' of August 2007.

## 9. Market-making and high-frequency strategies

A market-maker posts resting limit orders on both sides of the book, earning the bid-ask spread each time they're crossed. The economics are tiny per trade (fractions of a basis point) but enormous in aggregate at market-making scale.

The central challenge is **adverse selection**: a smart trader only lifts your ask when they know the price is about to go up, leaving you short at the worst time. Modern electronic market-makers mitigate this with (a) ultra-fast cancellation when flow turns informed, (b) dynamic spread widening as inventory imbalance grows, and (c) sophisticated fair-value models from many correlated signals (other venues, index futures, news feeds).

- **Latency arbitrage** — exploit microsecond lags between correlated venues.
- **Index arbitrage** — ETF versus basket, or futures versus cash index.
- **Triangular arbitrage** — inconsistent quotes across three FX pairs.
- **Rebate capture** — post liquidity on exchanges that pay maker rebates, collecting the rebate as pure edge when the order fills.

These strategies are infeasible for retail participants: they require co-located servers, direct exchange feeds, FPGA-accelerated decision loops, and total monthly volumes measured in billions. Mentioning them here is context — as a non-HF trader, you are on the other side of their trades by default, which is a cost to budget for.

## 10. Execution algorithms

Once the strategy has decided to trade, the execution algorithm decides how. Dropping a large order into the market in one shot moves the price against you (**market impact**) and signals your intent to predatory traders (**information leakage**). Execution algorithms break the parent order into child orders placed across time and venues to minimise both costs.

Algorithm	Rule	When to use
-----------	------	-------------

<b>TWAP</b>	Slice the parent order into equal chunks across the target window (e.g. 60 slices over an hour, one per minute).	Low-urgency orders; when you have no view on intraday volume.
<b>VWAP</b>	Slice the order to match the historical intraday <b>volume profile</b> — more participation when volume is naturally high (open, close), less in the quiet mid-day.	Benchmarked executions: large institutional orders where beating VWAP is the performance metric.
<b>POV / Participation</b>	Trade at a fixed percentage of actual volume as it happens (e.g. 10%). Speeds up in busy markets, slows down in quiet ones.	When you need to complete within a market session but don't want to move price aggressively.
<b>Implementation Shortfall</b>	Optimize the trade-off between market impact (pay more to finish now) and opportunity cost (pay more if the price drifts while you wait). Almgren–Chriss is the classic model.	Alpha-driven orders where delay has a cost because the signal decays.
<b>Iceberg</b>	Display only a small slice of the true order size in the book; replenish as fills come in.	Hiding large intent in a visible order book.
<b>Liquidity seeking</b>	Smart order router that hunts for block liquidity across lit exchanges, dark pools, and conditional venues, using minimum fill-size constraints to avoid pinging.	Large blocks where adverse selection is the main concern.
<b>Peg</b>	Limit order that floats with the market (pegged to bid, mid, or ask) and re-prices as the market moves.	Market-making or liquidity provision; also as a passive version of a limit order.

The right choice depends on **signal urgency**. A momentum signal that is minutes old is worthless — use aggressive algorithms (implementation shortfall, liquidity-seeking). A factor rebalance over a week is insensitive to intraday timing — use VWAP or POV and focus on minimising impact.

## 11. Risk management building blocks

- **Position sizing.** Size each trade so its loss at a predefined stop is a small, fixed fraction of equity (0.25 – 1% is typical). Formally:  $\text{size} = (\text{risk budget}) / (k \times \text{ATR})$ .
- **Stop losses.** Pre-committed exit levels remove the discretion to turn a small loss into a big one. Trailing stops ( $k \times \text{ATR}$  below the high) ratchet higher with the trend without re-arming on every bar.
- **Volatility targeting.** Scale total portfolio exposure inversely to realized volatility so that the book runs at a constant risk level instead of a constant notional.
- **Exposure limits.** Caps on gross leverage, net market exposure, single-name weight, sector weight, country weight. Simple, boring, and the difference between surviving a bad month and blowing up.
- **Correlation and factor monitoring.** Today's 'diversified' portfolio can be tomorrow's concentrated factor bet if correlations spike. Recompute factor exposures and pairwise

correlations frequently, not just at inception.

- **Kill switches.** Automatic flatten-all triggers on extreme drawdown, stale market data, or broken broker connectivity. Human reflexes are too slow.

## 12. Backtesting and the path to real money

A backtest shows what a strategy *would have* earned on historical data. It is the cheapest way to evaluate ideas — and the easiest way to fool yourself. The industry-wide failure mode is **overfitting**: parameters tuned so well to the past that they have no predictive power on new data.

- **Look-ahead bias** — accidentally using information in the decision at time  $t$  that wasn't actually available until  $t + k$ . Classic cause: joining today's market-cap weights into a 20-year-old price series.
- **Survivorship bias** — testing only on securities that still exist today. The universe you would have traded in 2005 included Enron and Lehman; a clean backtest must include them.
- **Data snooping** — running thousands of variants and reporting the best one. Statistically, even pure noise produces a Sharpe  $> 2$  if you try enough combinations. Adjust for the number of trials (e.g. Bonferroni, deflated Sharpe ratio).
- **Ignoring transaction costs** — every trade pays commission, spread, slippage, and for larger orders, market impact. A strategy with a 1% expected edge per trade and realistic 50 bps round-trip costs still nets 0.5%. Small edges die in real execution.
- **Regime blindness** — a strategy that works on 2010 – 2020 data was fit to the longest equity bull market ever seen. Test it over 2000 – 2003, 2008, 2022 separately before trusting it.
- **Parameter instability** — if a tiny change in a lookback window breaks the backtest, the strategy has no real edge. Walk-forward and parameter-sensitivity analysis are essential sanity checks.

The practical path from idea to capital looks like: generate the hypothesis → in-sample backtest on a subset of history → out-of-sample validation on untouched history → paper-trade in live market for weeks or months → deploy with a fraction of target size → scale up only after live Sharpe stabilises near backtested expectations. Each stage kills more ideas than it promotes; this is by design.

## 13. Combining strategies

No single strategy works in every market. Trend-following loses money in sideways markets; mean reversion loses in strong trends; factor strategies suffer multi-year drawdowns. Combining weakly-correlated strategies reduces total risk without proportionally reducing return — the formal statement of diversification.

Portfolio Sharpe  $\approx \sqrt{N}$  × per-strategy Sharpe (uncorrelated strategies, equal risk)

Four uncorrelated strategies each with Sharpe 0.5 combine to an aggregate Sharpe of 1.0. In practice correlations are never zero and the compounding is smaller, but the principle holds: *strategy diversification is the only free lunch a systematic trader gets*. Allocating risk equally across strategies (rather than capital) is usually a better starting point than concentrating on the single best-looking backtest.

## 14. Ethics and regulation

All of the above is legal when done transparently. Practices that cross the line include **spoofing** (placing and immediately cancelling large orders to mislead other participants), **layering** (multiple

spoof orders at different levels), **wash trading** (self-matched trades to inflate apparent volume), and **front-running** (trading ahead of a client order you know about). Regulators — SEC, CFTC, ESMA — actively surveil for these patterns; penalties range from disgorgement to criminal prosecution. A systematic strategy must not only have positive expected return, it must have it for reasons a regulator would approve of.

## 15. Further reading

- Ernest Chan — *Algorithmic Trading: Winning Strategies and Their Rationale*. Practical stat-arb and mean-reversion with code.
- Marcos López de Prado — *Advances in Financial Machine Learning*. How to apply ML without overfitting; the definitive treatment of backtest validation.
- Robert Kissell — *The Science of Algorithmic Trading and Portfolio Management*. Deep dive on execution algorithms, Almgren–Chriss, market impact models.
- Andrew Lo & Jasmina Hasanhodzic — *The Heretics of Finance*. Interviews with practitioner technicians; useful antidote to purely academic efficient-markets literature.
- Antti Ilmanen — *Expected Returns*. Comprehensive survey of every major return premium with historical evidence.

---

Financial Analysis — internal reference document on trading strategies and algorithms.