

Copy trading, without the bullshit.

Who is worth copying, how the pipes actually work, and whether the right portfolio is one trader, fifteen, or fifteen hundred.

The answer: do not wire Miroscope to one leaderboard winner. Do not blindly copy 1,500 accounts. Start with a deliberately different basket of **8 to 15 strategies**, paper-only, with risk normalization, correlation caps, and a kill switch.

The mechanism works. The shortcut does not.

Copy platforms can faithfully relay another account's orders. That does not prove the source trader has a repeatable edge, and it does not guarantee the follower receives the same result.

ONE TRADER

No

One hidden leverage change, one strategy drift, or one bad regime can wreck the whole account.

8 TO 15 STRATEGIES

Best test

Enough breadth to reduce idiosyncratic risk, still small enough to audit every trade and control overlap.

1,500 TRADERS

Noise

Looks diversified, but usually recreates market beta, duplicates crowded trades, multiplies costs, and becomes unauditable.

Hard truth: I could not verify a single current “best trader” whose identity, full closed-trade history, fees, leverage, and follower execution were all independently auditable. Public leaderboards are candidate lists, not proof.


What I would build

Read-only
leader feedsNormalize
risk and costsPaper-copy
90+ daysPromote only
if gates pass

This keeps Miroscope aligned with its current demo-only rule. No real-money switch belongs in this project until the copied basket passes its own frozen, forward test.

What “copy” actually copies

The leader’s trade is a signal. The follower’s result is a new execution, translated through different capital, leverage, broker rules, spread, liquidity, and latency.



Tracking error is unavoidable. cTrader warns that leverage mismatch can cause insufficient margin or earlier stop-out, and that equity stops may fill below the requested level because of volatility, gaps, liquidity, or technical factors.^[3]

The four numbers that matter

Layer	Measure	Why it matters
Signal	Leader entry/exit timestamp	Separates strategy timing from follower delay.
Translation	Target risk, leverage, size	Raw lots are meaningless across unequal accounts.
Execution	Fill price, rejected orders, slippage	Leader return can be uncopyable in practice.
Outcome	Net P&L after every fee	The only result the follower owns.

Where the usable pipes are

Platform	What is copied	Control / access	Research verdict
eToro CopyTrader	Popular Investors' activity	\$200 minimum per investor; up to 100 at once. ^[1]	Best consumer UX Useful for native testing, weak as a raw research feed.
MQL5 Signals	MT4/MT5 provider orders	Automated, cross-broker, paid fixed subscription possible. ^[2]	Best MT5 fit Closest to Miroscope, but one signal per account and broker drift matter.
cTrader Copy	Live or demo strategies	Dedicated copy account; equity stop; performance, management, and volume fees may apply. ^[3]	Best instrumentation Clear follower-level stats and risk controls.
Darwinex	Risk-normalized DARWIN indices	Risk engine targets 3.25% to 6.5% monthly VaR; 20% performance fee is documented. ^{[4][5]}	Best risk wrapper Cleaner comparison than raw leaderboard ROI.
ZuluTrade	Third-party Leaders	Broker-connected social copying; ranking and protection tools are platform-controlled. ^[6]	Needs export audit Broad marketplace, but headline rank is not enough.
Collective2	Strategies into a brokerage account	Stocks, options, futures, forex; results shown may be hypothetical. ^[7]	Model risk Useful breadth, but live executability must be re-measured.
NAGA	Trader positions	Public social-copy discovery and ranking. ^[8]	Discovery only Treat profiles as leads, not validated track records.

US access is not uniform. Product availability, instrument set, and regulatory entity differ by country and account. Verify the actual legal entity and permissions before treating any platform as implementable.

Platform terms, fees, and access can change. Figures above are from official help or product material observed 1 October 2026, not a promise of current account eligibility.

A winner is easy to display. Skill is hard to prove.

The leaderboard answers: “Who looks best in this window?”
It does not answer: “Whose edge survives the next window in my account?”

What the screen rewards

- High recent ROI
- Leverage and concentration
- Survival long enough to appear
- Strategies with smooth closed wins and hidden open losses
- Popularity and social proof

What an investor needs

- Complete closed and open trade history
- Return after leverage normalization
- Maximum equity drawdown
- Follower fills after costs
- Stable behavior out of sample

Research warning: an eToro network study found that follower count followed a power-law pattern and ranking did not correlate with performance among the 300 most-followed users it examined.^[9]

Behavior warning: a 2024 Scientific Reports paper found social-trading followers took more risk, reacted strongly when signal providers increased risk, and obtained more negative returns in the studied setting.^[10]

Why I am not naming a “best” individual

- Public pages are rolling snapshots. The trader can change leverage, stop publishing, or alter strategy immediately after selection.
- Some records are model or hypothetical results, not a follower’s executable history.
- A promotional profile is not an independent audit.
- The identity that matters is the strategy process, not a screen name.

This is not a claim that no skilled trader exists. It is a claim that the evidence gathered here does not justify presenting one person as a durable, copyable winner.

The trader has to survive the boring questions first

	Gate	Minimum evidence	Automatic rejection
01	History	At least 24 months and 200 closed trades, spanning more than one regime	Short promotional window
02	Completeness	Open positions, deposits, withdrawals, and closed trades visible	Equity or open losses hidden
03	Risk	Leverage, margin use, worst trade, max equity drawdown	Leverage spikes or no hard loss bound
04	Payoff	Average win, average loss, expectancy, profit factor	Very high win rate with catastrophic left tail
05	Concentration	Top trade, top symbol, and top month share of P&L	Edge depends on one event or one market
06	Copyability	Follower fills, slippage, rejected orders, minimum balance	Leader result cannot be reproduced
07	Incentives	Fee method and provider compensation disclosed	Paid mainly for volume without controls
08	Persistence	Frozen first-half selection succeeds in unseen second half	Only in-sample ranking works
09	Behavior	Stable instruments, holding time, leverage, and sizing	Style changes after attracting followers
10	Operations	Feed uptime, timestamps, broker, account type, execution logs	Unexplained gaps or account replacement

Score risk before return. A 40% return with a 60% drawdown is not “better” than a 12% return with a 7% drawdown. Raw ROI alone is not a ranking system.

The numeric thresholds are proposed Miroscope research rules, not industry guarantees. Freeze them before examining candidate results to avoid moving the goalposts.

One, fifteen, or fifteen hundred?

1 SOURCE

Concentrated bet

- Simple to monitor
- Clean attribution
- Maximum manager risk
- One failure can dominate

15 SOURCES

Controlled ensemble

- Diversifies real differences
- Still auditable
- Supports strategy-level caps
- Correlation can be measured

1,500 SOURCES

Synthetic crowd

- Operationally huge
- Many duplicate exposures
- Slow signals get crowded
- Fixed fees can explode

The math people miss

NOMINAL DIVERSIFICATION

15 names

Fifteen traders all long US tech or all selling volatility are one position wearing fifteen hats.

EFFECTIVE DIVERSIFICATION

Risk clusters

Count independent return drivers after correlation, common symbols, direction, holding period, and tail behavior.

1,500 is not “even better.” Equal-weighting a giant public crowd tends to dilute any real skill, import market exposure, amplify common trades, and make due diligence impossible. More feeds improve breadth only when the added feed contributes an independent, net-positive edge.

Decision rule

Start at eight. Add a strategy only if it reduces predicted portfolio risk or adds a distinct positive expectancy after fees. Stop near fifteen unless the next source demonstrably improves the portfolio out of sample.

Build a basket of behaviors, not a list of celebrities

Target distinct sleeves

- 1 **Trend**
Longer holding, directional persistence
- 2 **Mean reversion**
Shorter holding, different payoff shape
- 3 **Breakout**
Convex profile, accepts frequent small losses
- 4 **Macro / swing**
Lower frequency, event-aware
- 5 **Market-neutral**
If execution and borrow are truly copyable

RISK BUDGET, NOT CAPITAL SPLIT

Weight by expected loss

Equal dollars do not mean equal risk. Volatility-scale every sleeve, then cap any one strategy and any one correlated cluster.

- Single strategy: max 15% of portfolio risk
- Single symbol/theme: max 25%
- Single broker/platform: max 35%
- Unallocated reserve: 20% to 30%

Daily ensemble logic

When signals agree	When signals conflict	When risk spikes
Combine only up to the symbol cap. Three correlated longs do not get three full allocations.	Net cautiously only if the strategies were validated independently. Otherwise preserve attribution in separate sleeves.	Scale gross risk down. Do not let a leader's sudden leverage hike force the follower to match it.

Do not average raw orders. Convert each signal to a target risk contribution first. Then aggregate. This prevents a high-leverage trader from quietly owning the portfolio.

Use native execution. Build our own measurement.

ROUTE A · RECOMMENDED FIRST

Platform-native paper copy

- Open a demo or virtual allocation where permitted.
- Let the platform perform its own sizing and order routing.
- Export leader and follower trades every day.
- Measure the copy gap trade by trade.

Why: lowest engineering risk and closest to the user experience we are evaluating.

ROUTE B · RESEARCH LAYER

API or terminal observer

- Read positions, deals, quotes, and account events.
- Normalize symbols, timestamps, currency, and risk.
- Write signals into a shadow ledger.
- Never infer a missing close.

Why: lets Miroscope compare providers on one honest basis.

ROUTE C · LATER

Independent copier

Subscribe to authorized feeds, calculate target positions, and execute through a broker API or MT5 EA. This creates legal, operational, and execution responsibilities.

ROUTE D · REJECT

Screen scraping + live money

Fragile, delayed, possibly against platform terms, and impossible to reconcile reliably. A leaderboard page is not an execution feed.

Platform hooks worth testing

System	Native mechanism	Use in the study
MQL5 / MT5	Terminal subscription and automatic copying ^[2]	Primary demo execution lane for Miroscope compatibility.
cTrader	Dedicated copy account plus official Open API ^{[3][11]}	Best clean-room benchmark for follower statistics and event ingestion.
eToro	Native CopyTrader allocation ^[1]	Consumer behavior benchmark, not first automation target.
Darwinex	Risk-managed investable indices ^[4]	Benchmark the value of provider-independent risk normalization.

Gross leader return is not your net follower return

LEADER RECORD

Gross

- May ignore subscription fee
- May not include follower spread
- May reflect a better fill
- May include unrealized P&L

FOLLOWER ACCOUNT

Net

- Subscription or management fee
- Performance fee
- Spread and commission
- Slippage, rejects, financing

Fee structures change behavior

Fee	Provider incentive	Portfolio effect
Fixed subscription	Attract and retain subscribers	Small accounts suffer high fixed drag.
Performance fee	Generate profit above a high-water mark	Closer alignment, but can reward tail risk near reset points.
Management fee	Grow allocated capital	Paid even during flat or losing periods.
Volume fee / rebate	Trade more	Direct conflict if activity, not quality, drives pay.

Regulators see the conflict. ESMA specifically highlights remuneration structures, inducements, copied-trader qualifications, costs, and suitability as supervisory issues.^[12]

cTrader documents provider-set performance fees up to 40%, management fees up to 10%, and volume fees up to USD 10 per million copied, depending on the strategy.^[3] Darwinex documents a 20% performance fee, split 15% to the provider and 5% to Darwinex.^[5]

Fee ceilings are platform rules, not the fee every strategy charges. Always capture the actual strategy fee at subscription time.

How a beautiful curve lies

Martingale camouflage

Many tiny wins, occasional enormous loss. Screen for rising size after losses, long open losers, and average loss far above average win.

Leverage masquerading as skill

Raw ROI can rank the most aggressive account first. De-leverage returns before comparing providers.

Survivorship bias

Dead accounts vanish while survivors remain visible. A current leaderboard is selected on being alive today.

Strategy drift

A provider builds a safe record, attracts capital, then changes symbols, frequency, or leverage. Monitor behavior, not just P&L.

Crowding and slippage

Popular short-term trades become least copyable when many followers chase the same thin liquidity after the leader.

Fake venue or operator

Guaranteed returns, crypto-only deposits, withdrawal pressure, and no verifiable authorization are stop signs.

Regulatory line: the FCA says automatic copy trading can be portfolio management, bringing authorization, suitability, conduct, and reporting obligations.^[13]

Fraud is not theoretical. The CFTC obtained a \$1.7 billion judgment tied to Mirror Trading International and described it as a fraudulent multilevel-marketing scheme involving bitcoin.^[14]

Immediate kill triggers

Feed silence during open risk, unexplained account replacement, leverage beyond the approved band, missing open positions, withdrawal problems, altered trade history, or follower drawdown exceeding the frozen limit.

A 90-day paper test is the beginning, not the proof

DAYS 0 TO 30

Plumbing

- Confirm every signal arrives
- Reconcile open and closed positions
- Measure latency and fill gap
- Reject broken symbol mappings

DAYS 31 TO 60

Behavior

- Measure style drift
- Cluster correlations
- Recompute net expectancy
- Stress common-loss days

DAYS 61 TO 90+

Forward gate

- Freeze weights and rules
- No retrospective substitutions
- Track follower, not leader, P&L
- Promote nothing by instinct

Pass / fail dashboard

Gate	Pass condition	Fail condition
Data integrity	100% of trades reconciled or explicitly explained	Missing events or silent backfill
Copy gap	Net follower expectancy remains positive	Costs erase the leader's edge
Drawdown	Inside the frozen sleeve and portfolio limits	Limit breach, even if later recovered
Diversification	No one strategy or cluster dominates loss	Nominal names collapse into one bet
Stability	Leverage, symbols, and holding time stay in band	Unapproved drift
Out-of-sample	Positive net expectancy after freeze	Only backtest or selection window works

Do not “fix” the test midstream. If a trader fails, record the failure. Replacing losers with new recent winners creates a self-refreshing survivorship machine that can never be falsified.

Ninety days may be too short for low-frequency strategies. The trade-count gate matters more than the calendar. The minimum should be set before testing and extended when evidence is sparse.

Build the microscope first. Leave the trigger disconnected.

Phase 1 · Observe

- Ingest read-only feeds from MT5 and one independent platform.
- Store every signal, quote, fill simulation, fee, and rejection.
- Create immutable provider IDs so renamed accounts do not reset history.
- Keep raw events and normalized trades separate.

Phase 2 · Score

- Run the ten selection gates.
- Split every record into selection and unseen evaluation windows.
- Cluster by returns, symbols, direction, holding time, and tail losses.
- Publish a reason for every accept and rejection.

Phase 3 · Shadow ensemble

- Choose 8 to 15 distinct strategies.
- Size by risk contribution, not follower count or raw ROI.
- Cap strategy, symbol, cluster, and platform exposure.
- Compare ensemble fills with each native platform.

Phase 4 · Decision

- Require frozen forward gates.
- Review every loss-limit breach.
- Keep live trading OFF unless Michael explicitly authorizes it after the evidence passes.
- If no candidate passes, the correct portfolio is cash.

The proposed product is not “copy the best trader.”

It is an evidence engine that proves whether any trader is copyable, then combines only independent survivors under one risk budget.

Final answer

One is too fragile. Fifteen is a sensible research ceiling. Fifteen hundred is a data universe, not a portfolio. Use the 1,500 to discover candidates. Use 8 to 15 to test an ensemble. Use zero real dollars until the follower results, not the leader screenshots, earn the right.

Evidence used

1. **eToro Help, CopyTrader.** Minimum allocation and 100-investor limit. https://help.etoro.com/s/article/What-is-CopyTrader?language=en_GB
2. **MetaTrader 5 Help, Trading Signals.** Automated copying, cross-broker use, statistics, subscription model. <https://www.metatrader5.com/en/terminal/help/signals>
3. **cTrader Copy Help.** Account types, fees, leverage mismatch, equity stop behavior, follower statistics. <https://help.ctrader.com/ctrader-copy/investing-in-strategies/>
4. **Darwinex Help, Risk Engine.** Risk normalization and VaR range. <https://help.darwinex.com/risk-manager>
5. **Darwinex Help, Performance Fees.** 20% fee split and high-water-mark method. <https://help.darwinex.com/performance-fees-darwin-provider>
6. **ZuluTrade, Social Trading.** Official product description. <https://www.zulustrade.com/social-trading?ib=197>
7. **Collective2 product description, secondary source.** AutoTrade and hypothetical-result caveat. <https://slashdot.org/software/comparison/Collective-2-vs-UpTrader.io/>
8. **NAGA, Copy Trading.** Official product page. https://naga.com/en/copy-trading?_branch_match_id=1052767529234529539
9. **Pan et al., “Decoding Social Influence and the Wisdom of the Crowd.”** Follower distribution and rank-performance relationship. <http://alumni.media.mit.edu/~panwei/pub/socialcom12.pdf>
10. **Pelster et al., Scientific Reports 14, 1819 (2024).** Social comparison, risk taking, and returns. <https://www.nature.com/articles/s41598-023-49648-3>
11. **cTrader Open API.** Official programmatic-access documentation. <https://help.ctrader.com/open-api/>
12. **ESMA35-42-1428 (30 March 2023).** Supervisory expectations for copy trading. https://www.esma.europa.eu/sites/default/files/2023-03/ESMA35-42-1428_Supervisory_Briefing_on_Copy_Trading.pdf
13. **UK Financial Conduct Authority, Copy Trading.** Portfolio-management classification. <https://www.fca.org.uk/firms/copy-trading>
14. **US Commodity Futures Trading Commission, Mirror Trading International case.** Fraud judgment and scheme description. <https://www.cftc.gov/PressRoom/PressReleases/8854-24>

Scope limit: This was a public-source study, not a logged-in account audit. I did not verify private statements, actual follower fills, current US eligibility, or a complete live leaderboard export. That is why no individual was certified as “best.”

Observed 1 October 2026. Platform terms and product availability can change. This document is research and system-design guidance, not individualized investment advice.