

Do moving-average crossovers hold up after costs?

Evidence from 81 forex, metals and crypto markets, October 2023 to September 2026.

SixOne8 Research October 2026 10 min read

TRADES FROM	UNIVERSE	BARS	TRADES
1 Oct 2023 – 30 Sep 2026	81 markets	1h · 4h · 1d	45,129

ABSTRACT

We test a long-and-short EMA(20)/EMA(50) crossover on 81 markets and three timeframes over three years, against random entries matched on trade count, side and holding period, after the costs SixOne8 suggests for each market. The answer depends on the asset class. On forex the rule did worse than random timing on every timeframe: -1.8 pts a year at the median on 4h, with 9 of 37 pairs above their baseline. On crypto it did far better: +42 pts a year on 4h, 31 of 40 coins above baseline, across every one of the 24 parameter settings we tested. But on 1-hour bars crypto costs took all of it, leaving a median net return of -6.0% a year. No single market passed a correction for testing 243 combinations at once; the crypto result is one effect shared by coins that move together, measured over one strong cycle.

KEY FINDINGS

-1.8 pts Forex, 4h: median annual excess over random entries. Below the baseline on all three timeframes.	+42 pts Crypto, 4h: median annual excess. Positive at all 24 parameter settings.	-6.0% Crypto, 1h: median net return a year after 12 bps per side. Costs took the whole gross result.
---	--	--

01 Question

Moving-average crossovers are among the first rules most traders write. Buy when a fast average crosses above a slow one, sell when it crosses back. They are easy to state and easy to backtest, which also makes it easy to find a setting that looks good on one chart. We ask a narrow question across many charts at once: did a standard crossover time its trades better than chance, and did it still make money once realistic costs were paid?

Those are two different questions, and the paper keeps them apart. Costs decide whether a rule made money. A baseline of random trades decides whether its timing added anything. A rule can pass one and fail the other.

02 The rule and the baseline

The rule is the one a SixOne8 user would build: a long strategy that enters when EMA(20) crosses above EMA(50) at a bar's close, and a short strategy that enters on the opposite cross. Each position closes on the opposite cross or at a protective stop three ATRs from the entry, whichever comes first. Every market in the catalogue with history from October 2022 is included: 37 currency pairs, 4 metals and 40 crypto markets. The first year warms the averages up; trades are counted from October 2023.

For each market and timeframe we then draw 1,000 sets of random trades with the same number of trades, the same mix of long and short, and the same holding periods, placed at random without overlapping. A random set holds the same exposure to the market's trend as the rule and pays the same costs. The only thing it lacks is the rule's timing. *Excess* is the rule's annual return minus the median random set's.

$$\textit{Excess} = R(\textit{rule}) - \textit{median } R(\textit{random sets with the same trades, sides and holding periods}) \quad (1)$$


03 Results

Figure 1 follows each asset class on 4-hour bars. Crypto rose well above its baseline from the first months and stayed there. Forex drifted below zero and below its own baseline. Metals, only 4 markets, swung widely and finished a few points below their baseline.

FIGURE 1

Cumulative net return by asset class against the random-entry baseline, 4h

EMA(20)/EMA(50), long and short · equal weight within each class · percentage points, summed by month



Source: SixOne8 Research, from Twelve Data candles stored by SixOne8. Each trade's return after the app's suggested costs is booked in the month it closes; each line is the class average of those monthly sums, added up. The dashed lines are the median of 200 random-entry portfolios for the same class. Simple sums, not compounded.

The crypto line is lumpy. 67% of its final gain was booked in one month, December 2024, as positions opened during the November rally closed. A trend rule earns most of its return in a few long trends and gives some back between them, so a different three years, with fewer or shorter trends, would give a different line.

Table 1 gives the figures for every class and timeframe. The pattern holds across all three timeframes. Forex excess was -3.0 pts, -1.8 pts and -1.8 pts a year on 1h, 4h and 1d, each interval below zero. Crypto excess was +32 pts, +42 pts and +33 pts, each interval above zero.

TABLE 1 · Results by asset class and timeframe, EMA(20)/EMA(50)

Class	Bars	Markets	Trades	Gross	Net	Excess	90% interval	Above baseline
Forex	1h	37	13,169	-2.7	-5.1	-3.0	-4.7 to -1.1	10 of 37
Forex	4h	37	3,393	-1.5	-2.4	-1.8	-2.5 to -1.0	9 of 37
Forex	1d	37	608	-1.6	-1.8	-1.8	-3.9 to -0.5	10 of 37
Metals	1h	4	1,396	-1.3	-11.3	-4.5	-28.6 to +20.0	2 of 4
Metals	4h	4	371	-0.5	-2.3	-5.6	-8.3 to +11.4	1 of 4
Metals	1d	4	53	+17.8	+17.3	+0.3	-13.0 to +10.9	2 of 4
Crypto	1h	40	20,233	+34.0	-6.0	+31.5	+21.2 to +57.2	33 of 40
Crypto	4h	40	5,121	+42.6	+33.2	+42.3	+29.4 to +70.3	31 of 40
Crypto	1d	40	785	+37.0	+35.6	+33.4	+24.4 to +46.6	34 of 40
All	1h	81	34,798	+1.4	-5.1	+1.1	-0.5 to +5.4	45 of 81
All	4h	81	8,885	+1.0	-1.2	+0.2	-1.5 to +1.9	41 of 81
All	1d	81	1,446	+4.3	+3.6	+1.1	0.0 to +9.2	46 of 81

Medians across markets, percentage points a year. Gross before costs, net after the app's suggested costs. Excess: the strategy minus the median of 1,000 random-entry sets for the same market, which pay the same costs on the same number of trades, so excess is the same before and after costs. Interval: bootstrap over markets, 2,000 resamples. Coloured where it excludes zero.

Metals show why the baseline matters. On daily bars the median metal returned +17.3% a year after costs, which looks like a success. Its random baseline returned almost the same, and the excess was +0.3 pts. Metals rose over the period, and any rule that was long often enough made money. The crossover's timing added nothing to that.


04 Across markets

Pooling all 81 markets gives a median 4h excess of +0.2 pts, a number that describes no market at all. Figure 2 shows why: a cluster of currency pairs just below zero and a long tail of crypto markets far above it.

FIGURE 2

Annual excess return over the baseline, one bar per range of markets, 4h

81 markets · percentage points a year · the last bar holds every market at +120 or above



Below zero: 40 markets, 28 of them forex. At or above +120: 7 markets, all crypto. The median hides two different populations, which is why the paper reports each asset class separately.


Figure 3 shows the same thing one market at a time, with each rule beside its own baseline. Read down the colours of the names. A handful of coins sit at the very top, where the rule did worst of all; below them the currency pairs gather, short of their baselines by a few points; and most of the coins fill the bottom, clear of theirs by a wide margin. Crypto holds both the best and the worst results, which is what a large effect in a volatile class looks like.

FIGURE 3

Every market's rule against its own random baseline, 4h

Annual return before costs · filled dot the crossover rule, hollow dot the median random set · sorted by the gap

■ Rule above its baseline ■ Rule below its baseline



Market names are coloured by class: forex violet, metals teal, crypto in black. A > marks a value beyond the axis (ZEC/USD). The baselines sit in a narrow band near zero; the rules spread out from them, down for most currency pairs and up for most coins.

We then asked whether any single market was distinguishable from its own baseline once we allow for having tested 243 market-and-timeframe combinations. Ranking each against its 1,000 random sets and controlling the false-discovery rate at 10% (Benjamini–Hochberg), 0 passed. One market's three years of trades is too few to tell timing from luck. The crypto result rests on the class as a whole, and because coins move together, 40 crypto markets are much less than 40 independent confirmations.


05 Parameter stability

A result that appears at one setting and vanishes at its neighbours is usually noise. We repeated the test for 24 pairs of fast and slow averages. Figures 4 and 5 show the median excess for crypto and for forex on 4-hour bars.

FIGURE 4


Crypto: median excess across the parameter grid, 4h

Fast EMA by slow EMA · percentage points a year · blank where fast $\geq$ slow



All 24 settings beat the baseline at the median, from +22.6 to +78.5 points. The default, 20/50, sits in the middle of the range, not at the top.

FIGURE 5

Forex: median excess across the parameter grid, 4hFast EMA by slow EMA · percentage points a year · blank where fast $\geq$ slow

All 24 settings fell short of the baseline at the median, from -3.4 to -1.0 points. No region of the grid escapes it.

Neither picture has an isolated peak. Every crypto setting was above its baseline and every forex setting below. The headline numbers use the 20/50 default chosen before the study, not the strongest cell, which on crypto would have looked considerably better.


06 Sensitivity to costs

Faster bars mean more trades, and every trade pays twice. On 1-hour crypto the rule traded about 169 times a year per market. At the 12 basis points per side SixOne8 suggests for crypto, costs came to about 40 points a year. That is more than the gross result of +34, leaving -6.0% net. Figure 6 shows where each timeframe crosses zero as costs rise.

FIGURE 6

Crypto: net return against cost per side, by timeframe

Median across 40 crypto markets · percentage points a year · EMA(20)/EMA(50)



Costs are charged on both sides of every trade. The 1h line crosses zero at about 10 bps per side, below the 12 bps the app suggests for crypto (a typical taker fee plus spread). The slower timeframes trade too rarely for costs in this range to matter.

Forex costs are far lower, 1.5 pips or about one basis point per side on most pairs, and still the rule lost money on every timeframe, because it lost before costs. Table 2 sets out what each version could carry.

TABLE 2 · What costs could each version carry?

Class	Bars	Trades a year	Default cost	Break-even cost	Net positive
Forex	1h	120	1.1	none	3 of 37
Forex	4h	31	1.1	none	8 of 37
Forex	1d	5	1.1	none	13 of 37
Metals	1h	116	5.0	none	2 of 4
Metals	4h	30	5.0	none	1 of 4
Metals	1d	5	5.0	190.9	3 of 4
Crypto	1h	169	12.0	10.2	19 of 40
Crypto	4h	43	12.0	52.4	29 of 40
Crypto	1d	6	12.0	278.3	35 of 40

Medians across markets. Costs in basis points per side. Break-even: the cost per side at which the median market's gross return falls to zero; "none" where the median market lost money before costs.

07 Methodology

Every figure in this paper comes from the settings below. Where we made a choice, we state it.

DATA	1-hour candles built from 1-minute Twelve Data candles stored by SixOne8, 6 October 2022 to 30 September 2026. 4-hour and daily bars are folded from them on the app's own grid: New York 17:00 for forex and metals, UTC for crypto.
UNIVERSE	Every forex, metals and crypto market in the SixOne8 catalogue with history from October 2022: 37 currency pairs, 4 metals and 40 crypto markets. No market was removed after the fact.
WINDOW	Trades from 1 October 2023 to 30 September 2026. The year before warms the averages up, so every setting starts from fully formed values.
ENTRY	EMA(fast) crosses above EMA(slow) at a bar's close for a long, below for a short; the app's crossover operator, with SMA-seeded EMAs as the app computes them. Filled at the next bar's open. Default 20 and 50, chosen before the study.
EXIT	The opposite cross, filled at the next bar's open, or a stop $3 \times \text{ATR}(14)$ from the entry, filled at the stop or at the open of a bar that gaps through it. No target and no time limit.
COSTS	SixOne8's suggested spread per side for each market: 1.5 pips on currency pairs, 2–8 bps on metals, 12 bps on crypto. Sensitivity from 0 to 20 bps in Figure 6.
BASELINE	For each market and timeframe, 1,000 random sets (200 for the parameter grid) with the rule's trade count, sides and holding periods, shuffled and placed without overlap.

STATISTICS

Per-market rank against its baseline, Benjamini–Hochberg at 10% across all 243 combinations. Intervals for medians from 2,000 bootstrap resamples of markets.

CHECKED AGAINST

The SixOne8 backtest engine, running the same two strategies on EUR/AUD, BTC/USD and XAU/USD (4h, same window). BTC/USD gave the same 118 trades. On EUR/AUD 82 of the engine's 84 trades matched; the other two entered at a Sunday reopen four minutes later. On gold all 89 of the engine's trades matched, and this simulation took 11 more at the reopen after gold's daily break, where a SixOne8 backtest currently takes no trade.

$$Net = \Sigma \text{ trade returns} - 2 \times \text{cost per side} \times \text{trades, per year} \quad (2)$$

08 Limitations

Three years is short, and for crypto it was one strong cycle: long trends up and sharp falls, which is the market a trend rule suits. A different three years could look different. The crypto universe is the coins still in our catalogue with four years of history, so coins that collapsed or were delisted before today are missing, and their absence flatters any long-biased result. Costs are a fixed spread; real costs widen in fast markets, when a crossover is most likely to trade.

The study describes how the rule behaved. It says nothing about whether it will behave the same way next.

09 Test it yourself

RUN IT YOURSELF

Build this rule in SixOne8 and test it on your markets

Every condition in this paper is one you can build in the app. Test it on the markets you follow, with your own broker's costs in place of our defaults.

1. Create a long strategy with one condition: EMA, 4h, fast 20 crosses above slow 50.
2. Backtest it with a stop of $3 \times \text{ATR}(14)$ and an exit condition: EMA, 4h, fast 20 crosses below slow 50.
3. Repeat with a short strategy and the crosses reversed, then try other timeframes and costs.
4. Add the strategy to a watchlist if you want an alert when the next crossover happens.

CITE THIS PAPER

SixOne8 Research (2026). Do moving-average crossovers hold up after costs? Research Paper No. 2, October 2026.

Historical simulation and measurement only. Results describe how prices and rules behaved on past data under the stated assumptions and do not indicate future results. SixOne8 does not provide investment advice. Trading carries risk of loss.