Replication and Offset Methodology

A high-level overview for research and engineering teams

Replicurve

September 2026

Abstract

Research teams routinely approximate return streams they cannot hold directly—factor ETFs, fund sleeves, restricted single names, or cross-listed proxies—using liquid instruments in a candidate basket. Static factor regressions and spreadsheet models break when correlations shift, universes grow, or rebalance cadence must keep pace with production research pipelines.

Replicurve is a portfolio replication API. It accepts aligned daily return series you already maintain and returns time-varying long/short weights plus tracking metrics on each run. This document explains *why* dynamic replication matters, *what* the engine computes at a conceptual level, and *how* outputs map to typical research workflows—with summary statistics and charts from live engine runs. Proprietary estimation details remain server-side.

1 Who this document is for

Portfolio engineers, quantitative researchers, and platform teams at systematic funds and prop shops who:

- maintain in-house replication or proxy logic (Excel, linear regressions, ad-hoc Python);
- evaluate whether an API can replace bespoke maintenance without sacrificing tracking quality;
- need descriptive methodology before integrating `POST /v1/replication` or `POST /v1/offset`.

Replicurve is infrastructure, not investment advice. Examples below are **illustrations** only; tickers and parameters are not recommendations.

2 The research problem

2.1 What teams are trying to compute

Given a *target* return series r_t^T and a set of *candidate* return series $\{r_t^i\}_{i=1}^N$, the objective is to find portfolio weights $w_{i,t}$ such that the replicating portfolio return

$$r_t^P = \sum_{i=1}^{N} w_{i,t-1}\, r_t^i \tag{1}$$

tracks r_t^T over an evaluation window. In *offset* mode the same machinery targets $-r_t^T$: the portfolio is constructed to neutralize exposure to the target stream rather than replicate it.

Outputs are used in backtests, risk dashboards, and execution research—wherever a liquid proxy or hedge basket is needed. Replicurve returns weights and diagnostics; downstream systems decide how outputs are consumed.

2.2 Why static models fall short

Many in-house workflows estimate fixed betas once (or refresh them monthly) via ordinary least squares on a long history:

$$r_t^T \approx \alpha + \sum_i \beta_i \, r_t^i + \varepsilon_t. \tag{2}$$

That works until it does not:

- **Regime change.** Factor loadings drift; a static basket lags the target after volatility spikes or sector rotations.
- **Universe scale.** Screening 500+ names by hand or in a brittle script does not scale to weekly research batches.
- **Key-person risk.** One senior quant maintains the spreadsheet; turnover or absence stalls the pipeline.
- **Incomplete diagnostics.** Without correlation, tracking error, and slippage on every run, teams cannot compare targets or rebalance policies consistently.

Dynamic replication updates weights on a schedule (typically weekly) using only information available up to each rebalance date—the same constraint production research faces.

3 Replicurve approach (conceptual)

The engine runs server-side from aligned return inputs. You supply daily simple returns (not prices, not log returns). Each basket series must share the target calendar.

At a high level, every replication run follows four stages:

1. **Alignment and validation.** Dates are intersected; missing observations are handled consistently before any estimation begins. Payload validation (`/v1/validate`) checks structure before a full run.

2. **Candidate screening.** From the full basket universe, the engine ranks names by recent co-movement with the target and retains a workable subset. This keeps estimation stable when N is large (e.g. S&P 500 or Nasdaq-100 constituents).

3. **Dynamic weight estimation.** Weights are time-varying and estimated with an adaptive, state-space-style procedure that responds to changing correlations while penalizing unnecessary turnover. Long, short, and long/short constructions are supported; gross exposure is normalized according to internal policy. *Specific regularization, selection rules, and filter parameters are proprietary and not exposed on the public API.*

4. **Scheduled rebalance with lag.** Weights update on a configurable cadence (`daily` or `weekly`). A one-day implementation lag separates signal formation from the return period applied—matching realistic research assumptions.

The public request surface exposes only `mode`, `rebalance_freq`, and `max_allocation` (cap on names with material weight). Fixed engine parameters (e.g. lookback length, internal thresholds) are versioned server-side so methodology improvements do not require client changes.

4 Replication and offset modes

Replication is the usual entry point: proxy an ETF, fund sleeve, or synthetic series with a liquid basket (preset universe or custom tickers on Professional plans).

Offset applies when the research question is exposure neutralization—for example a concentrated single name or sector ETF whose return stream must be hedged in simulation without unwinding the underlying holding.

Mode	Endpoint	Intended output
Replication	`POST /v1/replication`	Portfolio return tracks target return; weights approximate the target stream.
Offset	`POST /v1/offset`	Portfolio return offsets target exposure; same request body, sign flipped internally.

Table 1: API modes share one request schema.

5 Illustration scenarios

The empirical section below reproduces every live example on `https://replicurve.com/examples`. Narrative groupings:

- **Factor ETFs** — MTUM, VLUE, USMV, QUAL replicated with Nasdaq-100 or S&P 500 baskets.
- **UCITS proxy** — VFMF target with `ucits_etfs` basket.
- **Liquidity proxy** — IJR small-cap target with S&P 500 basket.
- **Single-name offset** — AAPL, GOOG (Nasdaq-100) and JPM (S&P 500).
- **ETF sleeve offset** — EFA with US listed ETF basket.

6 Empirical results

All runs below use weekly rebalance, `max_allocation=10`, and market data from 2019-06-01 onward (first usable rebalance dates vary by listing history). Statistics are computed on aligned weekly portfolio vs. target returns returned by the engine. Tickers and pairings are **illustrations only**.

6.1 Summary statistics

Target	Mode	Basket	ρ	TE	R^2	Hit	Slip.
VFMF	Repl.	UCITS ETFs	0.712	12.8%	50.8%	73.8%	12.27%
MTUM	Repl.	Nasdaq-100	0.765	14.2%	58.5%	80.5%	7.25%
VLUE	Repl.	S&P 500	0.820	10.4%	67.3%	79.9%	-50.20%
USMV	Repl.	S&P 500	0.848	6.7%	72.0%	81.5%	16.31%
QUAL	Repl.	S&P 500	0.843	9.0%	71.0%	79.6%	-12.19%
IJR	Repl.	S&P 500	0.795	12.9%	63.2%	78.9%	34.23%
GOOG	Offset	Nasdaq-100	0.502	32.9%	25.2%	72.8%	-4.52%
AAPL	Offset	Nasdaq-100	0.463	28.0%	21.4%	70.6%	7.06%
JPM	Offset	S&P 500	0.827	15.0%	68.4%	88.2%	1.13%
EFA	Offset	S&P 500	0.758	11.3%	57.4%	75.4%	0.14%

Table 2: Tracking metrics for all live site examples (ρ = correlation, TE = annualized tracking error, Hit = hit rate, Slip. = cumulative slippage).

6.2 Replication charts

6.3 Offset charts

7 Tracking and quality metrics

Every successful run returns a `metrics` object alongside weight history. Table 2 and the chart captions summarize the fields most teams review first:

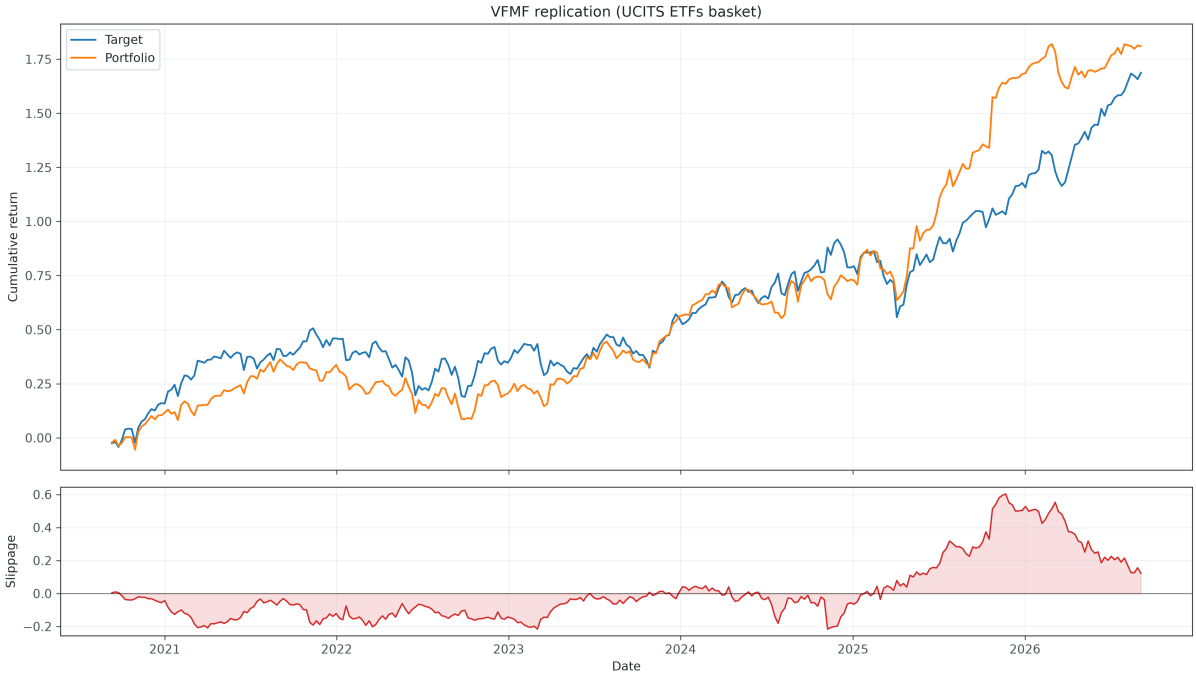

Figure 1: US multifactor via UCITS ETFs — target VFMF, UCITS ETFs basket, weekly rebalance. Correlation 0.712, tracking error 12.8%, R^2 50.8%, hit rate 73.8%, slippage 12.27% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

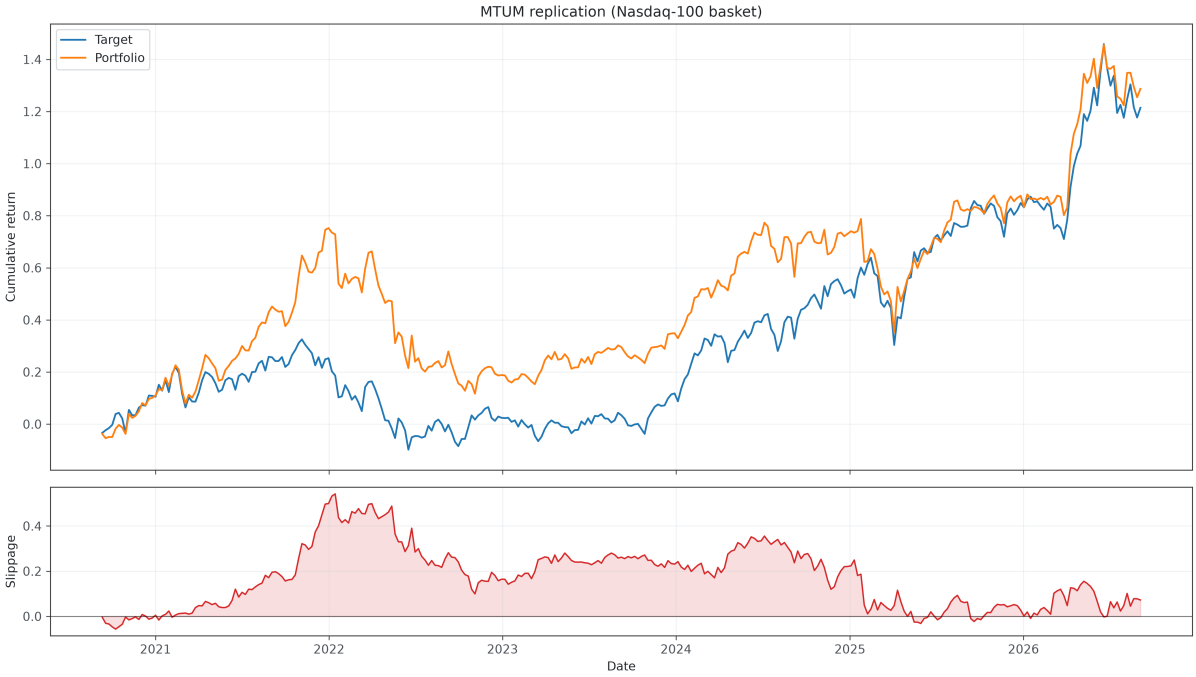

Figure 2: Momentum factor replication — target MTUM, Nasdaq-100 basket, weekly rebalance. Correlation 0.765, tracking error 14.2%, R^2 58.5%, hit rate 80.5%, slippage 7.25% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

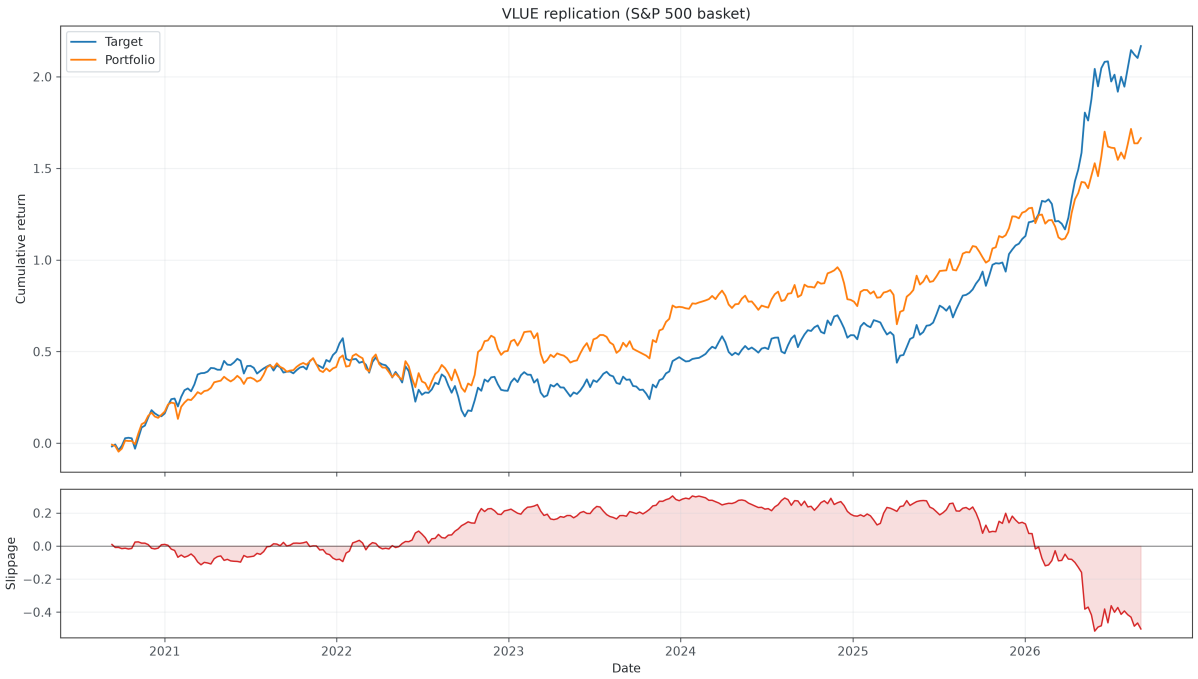

Figure 3: Value factor replication — target VLUE, S&P 500 basket, weekly rebalance. Correlation 0.820, tracking error 10.4%, R^2 67.3%, hit rate 79.9%, slippage -50.20% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

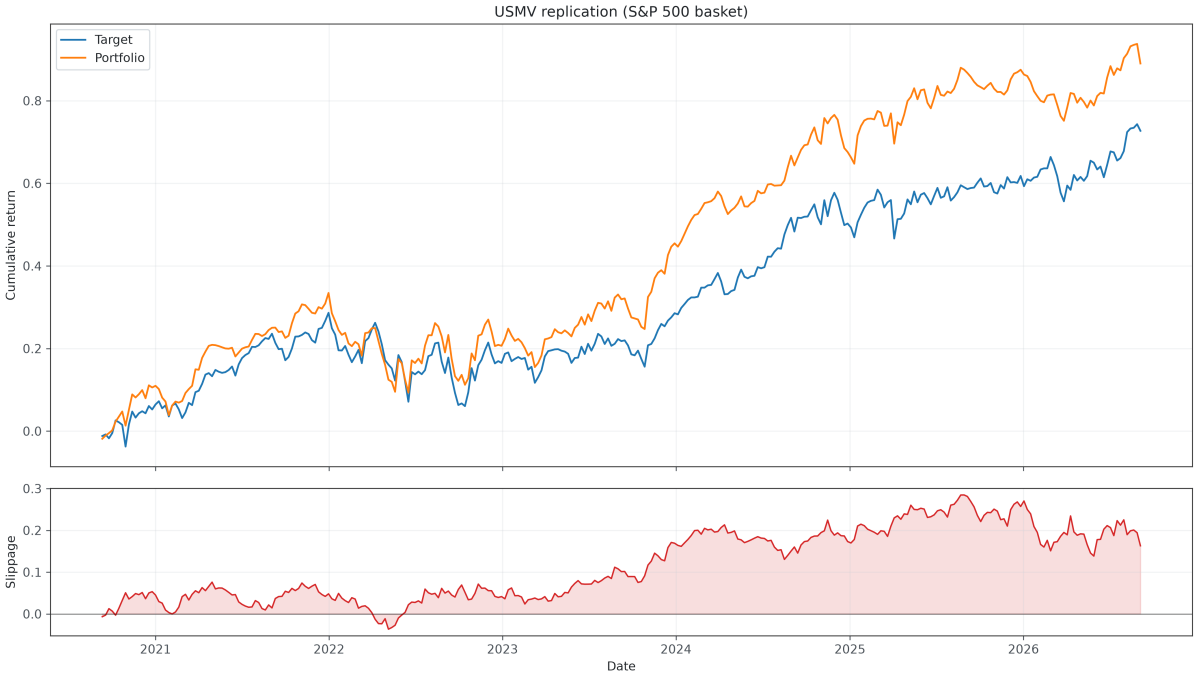

Figure 4: Minimum-volatility replication — target USMV, S&P 500 basket, weekly rebalance. Correlation 0.848, tracking error 6.7%, R^2 72.0%, hit rate 81.5%, slippage 16.31% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

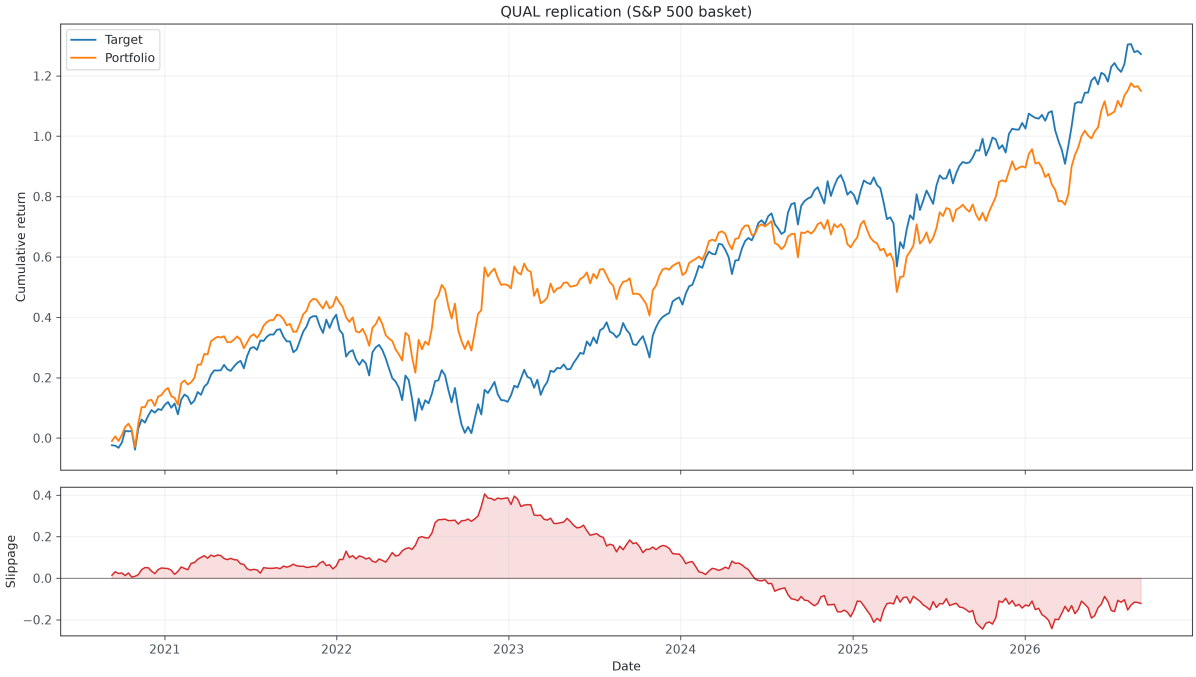

Figure 5: Quality factor replication — target QUAL, S&P 500 basket, weekly rebalance. Correlation 0.843, tracking error 9.0%, R^2 71.0%, hit rate 79.6%, slippage -12.19% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

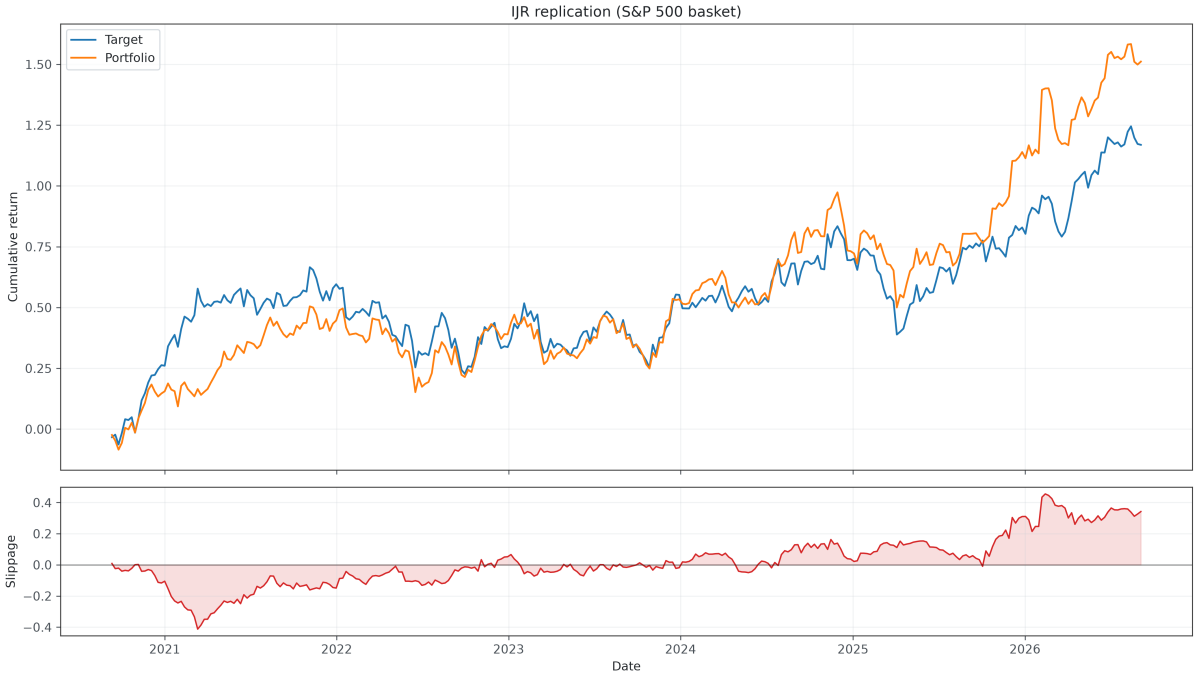

Figure 6: Small-cap replication with US listed ETFs — target IJR, S&P 500 basket, weekly rebalance. Correlation 0.795, tracking error 12.9%, R^2 63.2%, hit rate 78.9%, slippage 34.23% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

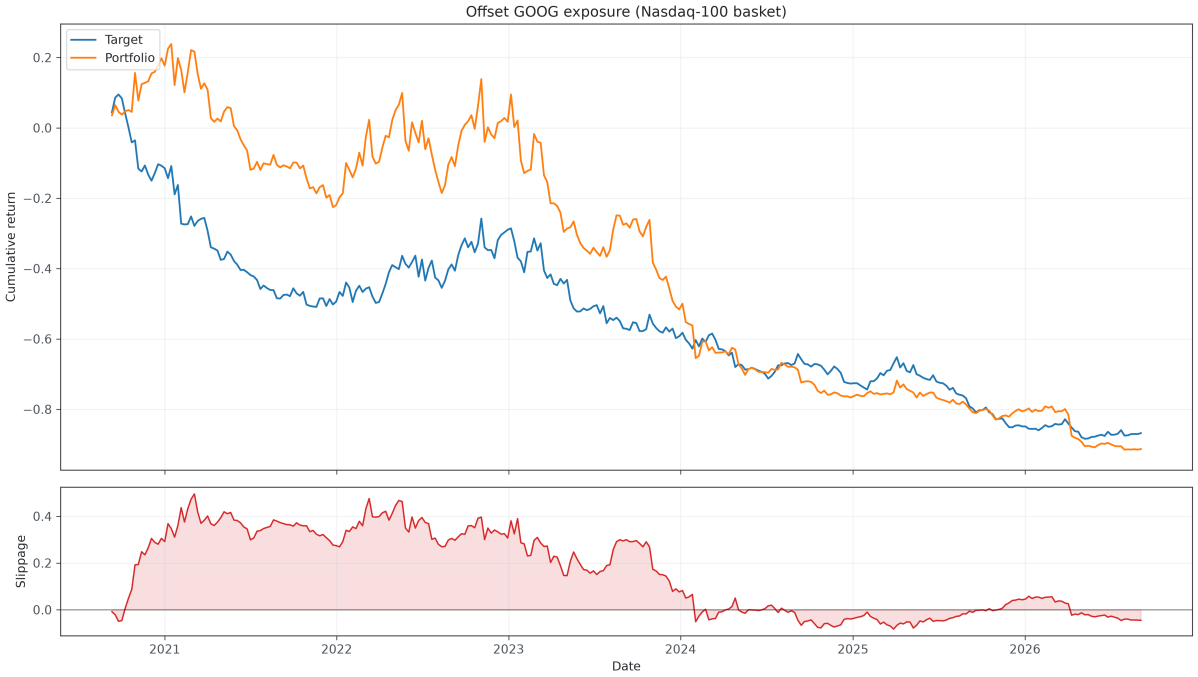

Figure 7: Alphabet offset — target GOOG, Nasdaq-100 basket, weekly rebalance. Correlation 0.502, tracking error 32.9%, R^2 25.2%, hit rate 72.8%, slippage -4.52% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

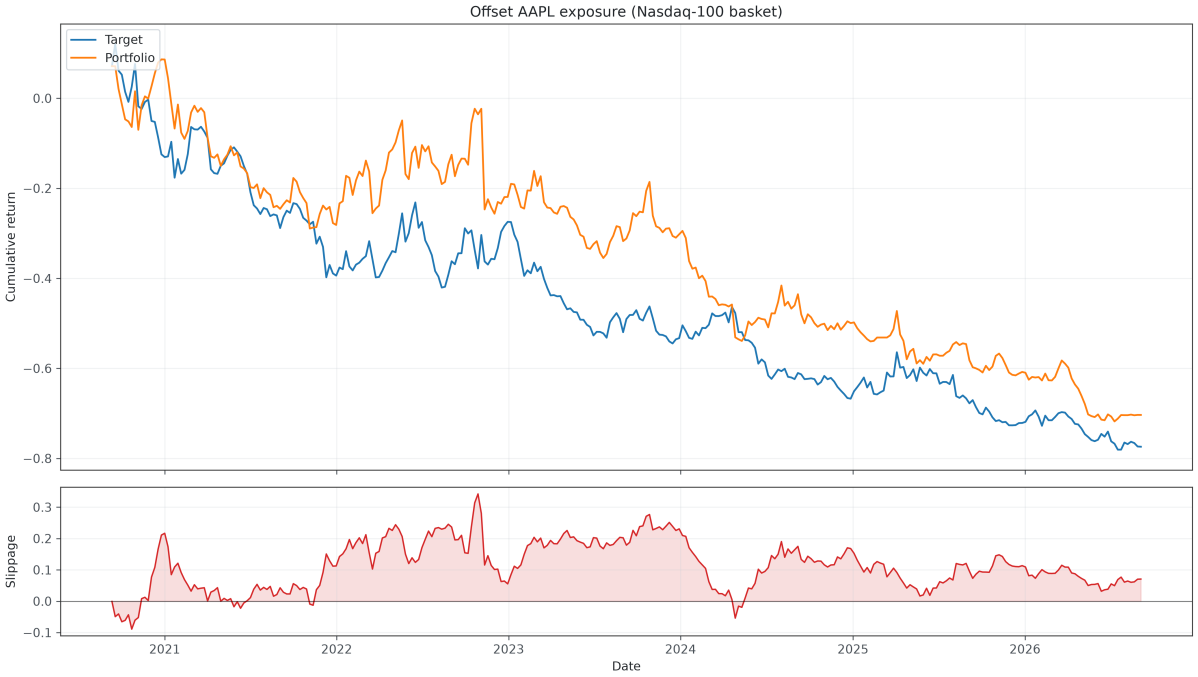

Figure 8: Apple offset — target AAPL, Nasdaq-100 basket, weekly rebalance. Correlation 0.463, tracking error 28.0%, R^2 21.4%, hit rate 70.6%, slippage 7.06% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

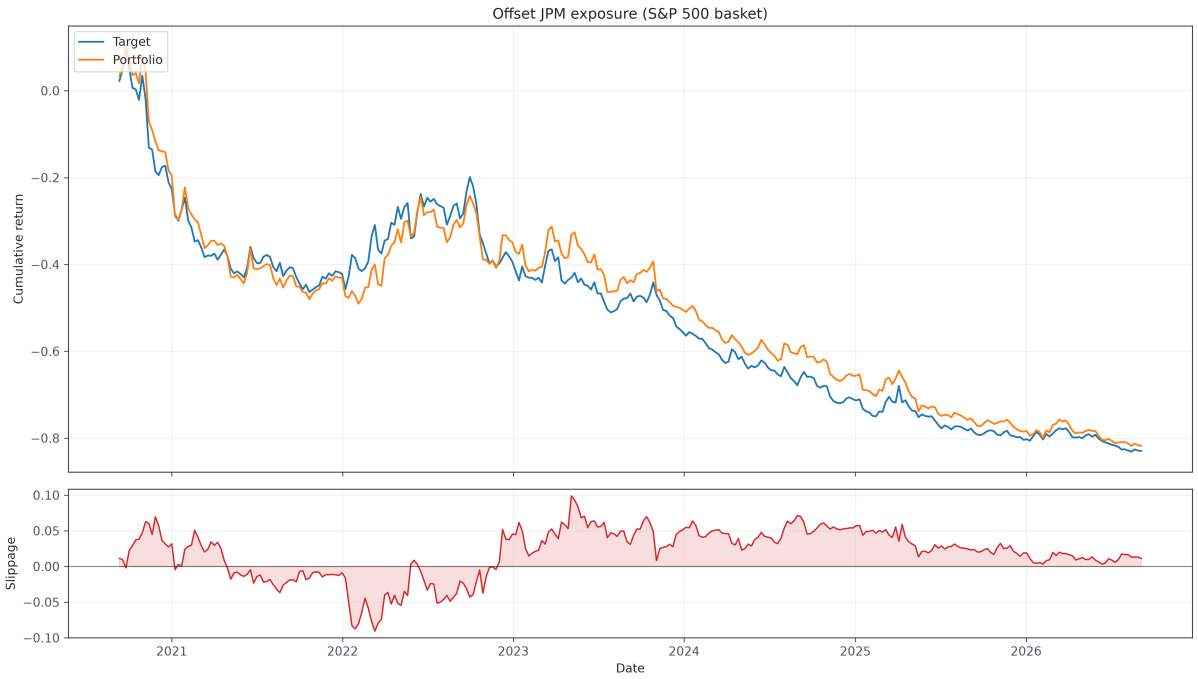

Figure 9: JPMorgan offset — target JPM, S&P 500 basket, weekly rebalance. Correlation 0.827, tracking error 15.0%, R^2 68.4%, hit rate 88.2%, slippage 1.13% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

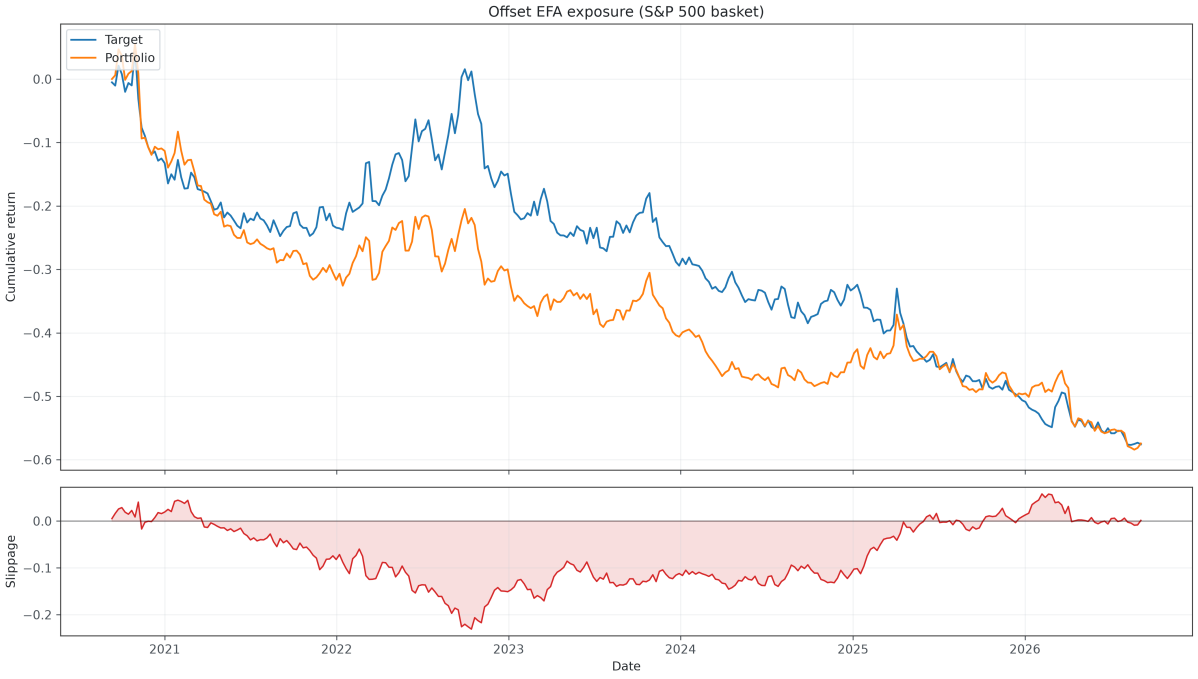

Figure 10: International equity ETF offset — target EFA, S&P 500 basket, weekly rebalance. Correlation 0.758, tracking error 11.3%, R^2 57.4%, hit rate 75.4%, slippage 0.14% over 2020-09-11–2026-09-04 (313 weekly obs.). Illustration only.

- **correlation** — Pearson correlation between target and portfolio returns over the sample.
- **tracking_error** — annualized standard deviation of $(r_t^P - r_t^T)$.
- **information_ratio** — mean active return per unit of tracking error (annualized).
- **hit_rate** — fraction of periods where target and portfolio share the same sign.
- **slippage** — cumulative return difference between portfolio and target.
- **beta**, **alpha_annual**, **r_squared** — linear association statistics on aligned returns.
- **avg_assets**, **turnover**, **gross_exposure** — portfolio structure summaries.

These metrics support side-by-side comparison against in-house models on identical return inputs—the fastest path to evaluate fit before committing integration effort.

8 Typical integration path

1. Align target and basket returns to a common daily calendar (your data warehouse or research store).
2. Call `POST /v1/validate` to check payload shape and date coverage.
3. Run `POST /v1/replication` with a preset universe (`sp500`, `nasdaq100`, `ucits_etfs`, etc.) or a custom basket (Professional tier).
4. Persist weights and metrics into your existing backtest or risk stack.
5. Schedule batch reruns when targets, universes, or rebalance policies change.

Authentication uses Bearer tokens. Plan tiers differ by call volume and basket flexibility; see the pricing page on the Replicurve site.

9 Scope, limitations, and what we do not disclose

- **No market data licensing.** You provide returns; Replicurve estimates weights.
- **No execution or advisory services.** Outputs are descriptive quantities for research systems.
- **Past tracking does not guarantee future tracking.** Structural breaks, corporate actions handled differently in your inputs, and liquidity frictions can widen error.
- **Proprietary estimation.** Candidate ranking rules, adaptive filter internals, stability checks, and numerical thresholds are intentionally omitted from this document and from the public API schema. They may evolve in server releases without client migration.

10 Summary

Dynamic replication and offset are recurring research maintenance tasks: targets change, regimes shift, and static regressions do not scale. Replicurve centralizes the estimation loop behind a REST API—aligned returns in, time-varying weights and tracking metrics out—so engineering teams can focus on data and downstream systems rather than reimplementing adaptive replication and offset math.

For interactive charts, OpenAPI reference, and API access: `https://replicurve.com`.

Important disclaimers

Regulatory status. Replicurve provides a portfolio replication API. Replicurve is *not* a registered investment adviser, broker-dealer, bank, or tax adviser, and does not solicit or recommend the purchase or sale of any security or financial instrument. Nothing in this document constitutes investment, legal, tax, or accounting advice.

Purpose of this document. This whitepaper describes what the Replicurve API computes and returns—time-varying weights, aligned return series, and tracking metrics—for integration into *your* research, risk, or backtest systems. It explains methodology at a high level only. Outputs are descriptive quantities; how they are used downstream is entirely the customer's responsibility.

Illustrations and hypothetical results. All tickers, universes, and pairings shown (including tables and charts) are **illustrations only**. They are not recommendations to replicate, offset, buy, sell, or hold any instrument. Statistics and charts are produced from historical return data over the periods stated in each caption. **Past tracking quality does not predict future tracking quality.** Correlation, tracking error, slippage, and related metrics can change materially with market conditions, input data errors, corporate actions, liquidity, and parameter choices.

Model and estimation risk. Replicurve outputs are *estimates* from a statistical engine, not guaranteed replications of live fund NAV, ETF intraday prices, or executable fills. The engine does not model transaction costs, market impact, borrow availability, short constraints, regulatory limits, or tax treatment unless you supply that logic separately. Proprietary estimation details may change between server releases without notice.

Data and third-party content. You supply return series; Replicurve does not license market data on your behalf. Accuracy of outputs depends on the quality, alignment, and timeliness of inputs you provide. Replicurve makes no representation regarding withholding tax, UCITS eligibility, listing rules, or cross-border investability—those questions require your own counsel and compliance review.

No warranty. This document and any cited metrics are provided "as is" without warranty of any kind, express or implied, including fitness for a particular purpose. To the fullest extent permitted by law, Replicurve disclaims liability for decisions made or actions taken based on this document or on API outputs.

Terms. Use of the Replicurve API is subject to the applicable service terms at `https://replicurve.com/terms`. In the event of conflict, those terms govern.

Document version: September 2026. For questions: `mailto:api@replicurve.com`.