QUANTITATIVE RESEARCH • SMART BETA

Smart Beta NASDAQ Quality ROIC-Value Index: A Factor Efficiency-Value Weighted Index

A Factor Efficiency-Value Weighted Index

Leonardo Romero, Ph.D. | Independent | leoromero.dev

June 2026

Executive summary

This paper presents the Smart Beta NASDAQ Quality ROIC-Value Index (NDX-QRV), a factor model index built on the NASDAQ-100 universe. The construction treats the cap-weighted NDX as the macro-consistent benchmark, in line with the MSCI Factor Index posture and against the Arnott-Hsu-Moore fundamental-indexation posture. The selection signal is an MSCI Quality composite anchored in Bender et al. (2013) with two documented refinements: return on invested capital replaces return on equity to neutralize capital-structure leverage in the profitability pillar, and a free-cash-flow yield component is paired with an EV/EBITDA valuation penalty so that quality is not bought at any price. The construction selects the top thirty names by composite quality z-score across the full point-in-time NDX universe and weights them by shift-and-renormalize with a 5% per-name concentration cap. The empirical contribution is twofold. First, a quality scorecard of the NDX with strong persistence across twenty quarterly rebalances over 1 January 2021 to 31 December 2025: top-quintile quarter-to-quarter persistence of 0.64 and Spearman rank correlation of 0.69, the diagnostic output of the framework and a standalone publishable result. Second, the five-year backtest of the index as an investable strategy. The index delivers an annualized return of 18.93% against the benchmark's 15.78%, with a maximum drawdown of −25.77% versus −35.12% for the cap-weighted NDX (proxied by QQQ), a portfolio beta of 0.875, an annualized volatility of 20.97% against the benchmark's 22.60%, and a downside capture of 99.52%. The annualized CAPM alpha is +4.52%, the Sharpe ratio is 0.828 against the benchmark's 0.649, and the information ratio is +0.356 at a tracking error of 7.54%. We interpret the empirical record as the joint signature of the MSCI Quality composite on the NDX universe: a high-persistence ranking that, traded as a top-thirty quality-tilted construction with a 5% per-name cap and explicit sector tilts driven by within-sector quality leadership, delivers genuine drawdown cushioning, a sub-1.0 beta, and material risk-adjusted outperformance over the same five calendar years.

+4.52% annualized CAPM alpha vs. NDX **0.828** Sharpe vs. NDX 0.649 **−25.8%** max drawdown vs.

NDX −35.1%

1 Introduction

This section will frame the problem statement. The smart-beta category contains three intellectual lineages, often conflated in practitioner discourse. The first is the fundamental-indexation school descended from Arnott et al. (2005), implemented by Research Affiliates, FTSE-Russell RAFI, and the Nasdaq Smart Beta indexes. The second is the institutional factor-index school, formalized in Bender et al. (2013) and operationalized in the MSCI Factor Index family via the Barra Global Equity Model. The third, less commercially visible but equally durable, is the Sharpe-Jensen residual-alpha tradition originating in Sharpe (1963) and Jensen (1968). This paper develops the third tradition into a production-ready smart-beta framework and positions its results against the first two.

Key finding

The Sharpe-Jensen residual-alpha lineage is structurally compatible with the MSCI posture (cap-weight as macro-consistent benchmark) and structurally incompatible with the RAFI / Nasdaq Smart Beta posture (cap-weight as structurally flawed). The three lineages disagree on what the benchmark is, not just on how to weight constituents.

2 Methodology

For each asset i in sector s, the sector-relative residual alpha is recovered from a Sharpe single-index regression of the asset return against the sector benchmark return.

The specification is deliberately economical, and the choice is worth flagging at the outset. The single-index regression of Sharpe (1963) and the alpha-from-intercept formalization of Jensen (1968) remain the dominant working primitive across institutional factor construction, including the Barra-derived single-factor exposures that anchor the MSCI Factor Index family (Bender et al., 2013) and the residual-return inputs to the AQR Quality-Minus-Junk and Betting-Against-Beta portfolios (Asness et al., 2013; Frazzini et al., 2012). What reads as elementary on the page is the actual estimation step on every desk that builds sector-relative or factor-relative scores; the multifactor extensions sit on top of this primitive rather than in place of it. Substituting a richer decomposition at the construction stage does not add information: it adds parameters to estimate from the same residual variance, and it imports identification questions that the single-index specification sidesteps.

The contrast with the retail-facing Smart Beta product set is sharper still. Indices published under the NASDAQ Smart Beta label (Nasdaq, Inc., 2024), including the NASDAQ-100 Equal Weighted Index, the NASDAQ Dividend Achievers, and the various single-metric fundamental and low-volatility variants, are constructed by direct sorts on observable accounting or yield metrics and by mechanical reweighting rules. No regression is fit, no sector beta is residualized out of any score, and no within-sector standardization is applied. The construction adopted here therefore sits one methodological layer above the dominant Smart Beta product convention while remaining strictly inside the institutional estimation primitive.

Definition 1. *For asset i over a calibration window of length T, the residual alpha is the ordinary least-squares intercept of the regression of $r_{i,t}$ on $r_{s,t}$, where $r_{s,t}$ is the return of the sector benchmark* *to which i belongs.*

$$r_{i,t} = \alpha_i + \beta_i r_{s,t} + \varepsilon_{i,t}, \quad t = 1, \ldots, T$$

The sector-relative cross-sectional score is obtained by standardizing $\widehat{\alpha}_i$ within each sector.

Model 1. *The asset return process is decomposed into a sector beta component, a residual alpha component, and an idiosyncratic disturbance. Selection and weighting operate on the residual alpha distribution within each sector, preserving the macro-consistency of the cap-weighted benchmark at the aggregate level.*

3 Empirical illustration

3.1 Universe, benchmark, and calibration window

The empirical exercise applies the framework to the NASDAQ-100 universe over the five-year window from 1 January 2021 to 31 December 2025. All input data are pulled from the Financial Modeling Prep (FMP) commercial data API through its `/stable/` endpoint family. The price input is the dividend-adjusted end-of-day series; sector and industry classification come from the company profile endpoint; point-in-time NDX membership is reconstructed from the historical-constituent endpoint; the four fundamentals that feed the Quality composite (return on invested capital, free-cash-flow yield, EV/EBITDA, net-debt/EBITDA) and the current-ratio safety component are pulled at quarterly frequency from the key-metrics and ratios endpoints, lagged by one quarter to model the 10-Q filing delay; and the historical-market-capitalization endpoint feeds the reconstructed cap-weighted NDX benchmark introduced in Section 3.9.

The universe is reconstructed point-in-time by backward replay of NDX membership at each quarter-end, retaining names that were members at the relevant as-of date even when they have since been delisted or removed from the index. The survivorship-bias check identifies thirty-six names present in the Q1-2021 universe that no longer trade in the index today, including ATVI, FISV, MXIM, SGEN, SPLK, XLNX, CERN, DOCU and CHKP; all are properly retained in the relevant calibration windows.

The benchmark is the cap-weighted NDX, proxied by the Invesco QQQ Trust ETF for the realized-return series. The quality composite follows the MSCI Quality methodology of Bender et al. (2013) with documented refinements. The profitability pillar is captured by return on invested capital rather than return on equity, because ROIC is invariant to capital structure and avoids the leverage inflation that biases ROE in highly geared names. The financial-strength pillar is captured by a balance-sheet safety component. We add a free-cash-flow yield as a cash-versus-accrual profitability check and an EV/EBITDA penalty as a valuation discipline so that quality is not bought at any price. Earnings variability, the third MSCI Quality pillar in the original methodology, is not included as a direct input; whether adding it improves the ranking is an open item for v2. Each component is standardized within sector-and-industry peers.

The general residual-alpha framework supports a single-index regression of each name against its sector benchmark on a 252-trading-day rolling window, with fundamentals lagged by one quarter to model the 10-Q filing delay. Rebalances occur on the last business day of each calendar quarter; the selection score is the convex combination $S = \lambda\, z_Q + (1-\lambda)\, z_\alpha$ with $\lambda \in [0, 1]$. The deployed NDX-QRV instantiation sets $\lambda = 1$ so the selection score reduces to within-sector composite quality z_Q alone, the limit case in which no residual-alpha overlay is active. Under this instantiation the portfolio selects the top thirty names by composite z_Q across the full point-in-time universe and weights them by the shift-and-renormalize rule with a 5% per-name cap documented in Section 3.4.

Key finding

The deployed top-thirty-absolute construction is not sector-neutral, and that is by design. The within-sector z-score machinery preserves the cross-sectional comparability of the quality signal across sectors of different size and composition; the top-N selection that follows is sector-agnostic and lets the quality ranking determine the sector weight. The realized active sector tilts at the 2025-12-31 as-of date are reported in Table 4; they are large, structurally different from a sector-neutral design, and are read as a feature of the quality ranking rather than as a residual bet.

3.2 The MSCI Quality scorecard for the NDX

The lead empirical contribution of the framework is the sector-aware quality ranking of the NDX it produces at each rebalance. The ranking is the diagnostic output of the methodology and stands independently of any decision to trade it. Table 1 reports the highest-composite (sector-winner) eligible name in each of the eleven NDX sectors at the 2025-12-31 as-of date, decomposed across the four pillar z-scores that drive the composite z_Q. Table 2 in the next subsection expands the ranking from the sector winner to the full top-five eligible names per sector at the same as-of date. The persistence column in both tables reports how many of the twenty quarterly rebalances over 2021-Q2 to 2025-Q4 each name appeared in the top-K selection for its sector, a stability check on whether the ranking is a structural feature of the name or a one-quarter accident.

Five readings of the scorecard sit on the surface and inform the rest of the section. First, NVDA at $z_Q = 1.57$ is the single-largest sector-relative quality score in the NDX at the as-of date and is driven by an ROIC z-score of 3.32: the franchise-economics archetype. BKNG at $z_Q = 1.14$ in Consumer Cyclical is the second-largest and is also profitability-driven (ROIC $z = 2.48$). Second, GILD at $z_Q = 0.48$ in Healthcare reads as the opposite archetype: the composite is carried by a free-cash-flow yield $z = 1.32$ with a modest ROIC $z = 0.38$ and a negative safety $z = -0.57$, the cash-generative mature business that earns the scorecard slot on cash discipline rather than incremental returns. Third, CEG at $z_Q = 0.76$ in Utilities is the only name on the scorecard with a positive z-score on all four pillars, including a positive EV/EBITDA $z = 1.36$: quality at a cheap valuation, rare in this universe. Fourth, MNST in Consumer Defensive is the cleanest persistence signal in the table at 20 of 20 rebalances, never displaced from its sector's top-K across the window. The other 20/20 anchors in Utilities (AEP, EXC, XEL) and Consumer Cyclical (ORLY) are equally persistent but were displaced at the as-of date by stronger composite scores from CEG and BKNG respectively, so the scorecard captures a real rotation inside otherwise stable sector leadership rather than a fixed set of names. Fifth, the single-member sectors (LIN, PYPL, CSGP) and the two-member Energy sector (BKR and FANG with the closed-form $\pm 1/\sqrt{2}$ pattern on every pillar) are statistical artifacts of NDX sector composition at the as-of date, not signals about the names themselves.

3.3 Full top-five-per-sector ranking

Table 2 expands the sector-winner view of Table 1 to the full top-five eligible names per sector at the same as-of date, ordered within each sector by descending composite z_Q. Sectors with fewer than five NDX constituents at the as-of date (Basic Materials, Financial Services, and Real Estate with one constituent; Energy with two; Utilities with four) report all available

Table 1. NDX quality scorecard, sector-winner view at the 2025-12-31 rebalance. One row per NDX sector, reporting the eligible name with the highest within-sector composite quality z-score z_Q, decomposed across the four pillar z-scores that drive the composite (return on invested capital, free-cash-flow yield, balance-sheet safety, and the EV/EBITDA penalty). Persistence is the number of the twenty quarterly rebalances over 2021-Q2 to 2025-Q4 in which the name was selected for its sector's top-K. **Bold** persistence cells mark the 20/20 structural anchors. LIN, PYPL, and CSGP carry the single-member convention $z = 0$ on every pillar by construction; BKR (Energy) carries the closed-form $\pm 1/\sqrt{2} \approx \pm 0.71$ pattern of the higher member of a two-name peer group. The full top-five-per-sector ranking appears in Table 2.

Sector	Ticker	ROIC z	FCFY z	Safety z	EV/EBITDA z	z_Q	Persist. (/20)
Basic Materials	LIN	0.00	0.00	0.00	0.00	0.00	8
Communication Services	GOOG	1.25	-0.39	0.88	-0.43	0.66	17
Consumer Cyclical	BKNG	2.48	-0.47	0.33	-0.58	1.14	15
Consumer Defensive	MNST	1.55	-0.55	1.78	0.71	0.72	**20**
Energy	BKR	0.71	-0.71	0.71	0.71	0.14	13
Financial Services	PYPL	0.00	0.00	0.00	0.00	0.00	**20**
Healthcare	GILD	0.38	1.32	-0.57	-0.42	0.48	13
Industrials	FAST	1.83	-0.53	0.81	-0.13	0.83	17
Real Estate	CSGP	0.00	0.00	0.00	0.00	0.00	13
Technology	NVDA	3.32	-0.34	0.87	-0.16	1.57	15
Utilities	CEG	1.44	0.38	1.47	1.36	0.76	15

names. The expansion is where the rotation story sits: 20/20 persistence anchors that lost their sector-winner slot at the as-of date but remain inside the top-five appear here (ORLY in Consumer Cyclical, AEP and EXC and XEL in Utilities), as do recent franchise entrants whose persistence is low because their NDX membership is short (APP, STX, MU, WDC, ALNY, FER all at ≤ 5 of 20).

3.4 Constituents of the NDX-QRV Index at 2025-12-31

The Index is the *Smart Beta NASDAQ Quality ROIC-Value Index* (NDX-QRV), a factor model index that pairs a pure quality-tilted selection rule with an explicit concentration cap and forgoes any sector-neutrality constraint. The selection rule ranks the full point-in-time NDX universe by composite z_Q descending and retains the top thirty names. Within the selected set, weights follow a shift-and-renormalize rule $w_i \propto z_{Q,i} - (z_{Q,\min} - \delta)$ with $\delta = 0.25$, capped at 5% per name with the cap excess redistributed proportionally across the uncapped names. The $\delta = 0.25$ floor pushes more weight toward the top of the ranking and is the parameter at which the cap binds on a useful fraction of names without flattening the composite signal. At the 2025-12-31 rebalance the cap binds on six names (NVDA, MSTR, BKNG, AAPL, KLAC, APP), all of which are quality leaders at the top of the composite distribution; the remaining twenty-four names land in the 1.44% to 4.88% range. The 5% ceiling with 30 selected names guarantees a diversified portfolio at construction: any single name holds at most 5% and a minimum of 24 names hold non-trivial weight. Across the twenty rebalances over the full backtest window the cap-binding count averages 5.3 names with a maximum of 8, evidence that the 5% ceiling is consistently active and is doing the diversification work for which it was introduced.

Table 2. Full top-five-per-sector quality ranking of the NDX at the 2025-12-31 rebalance, ordered by sector then by descending composite z_Q. Sectors with fewer than five NDX constituents at the as-of date (Basic Materials, Financial Services, and Real Estate with one; Energy with two; Utilities with four) report all available names. Persistence is the number of the twenty quarterly rebalances over 2021-Q2 to 2025-Q4 in which the name was selected for its sector's top-K. **Bold** persistence cells mark the 20/20 structural anchors. LIN, PYPL, and CSGP carry the single-member convention $z_Q = 0$ by construction. GOOG and GOOGL are dual-class share lines of the same issuer and therefore carry identical scores. Values from the `quality_scorecard_2025-12-31.csv` export.

Sector	Rank	Ticker	Composite z_Q	Persist. (/20)
Basic Materials	1	LIN	0.00	8
Communication Services	1	GOOG	0.66	17
Communication Services	2	GOOGL	0.66	19
Communication Services	3	NFLX	0.56	13
Communication Services	4	CMCSA	0.54	3
Communication Services	5	WBD	-0.88	4
Consumer Cyclical	1	BKNG	1.14	15
Consumer Cyclical	2	PDD	0.76	10
Consumer Cyclical	3	ROST	-0.04	11
Consumer Cyclical	4	ORLY	-0.09	**20**
Consumer Cyclical	5	MAR	-0.27	8
Consumer Defensive	1	MNST	0.72	**20**
Consumer Defensive	2	CCEP	0.36	9
Consumer Defensive	3	PEP	0.25	15
Consumer Defensive	4	KHC	0.07	10
Consumer Defensive	5	COST	-0.04	19
Energy	1	BKR	0.14	13
Energy	2	FANG	-0.14	13
Financial Services	1	PYPL	0.00	**20**
Healthcare	1	GILD	0.48	13
Healthcare	2	REGN	0.40	14
Healthcare	3	AMGN	0.22	9
Healthcare	4	IDXX	0.08	12
Healthcare	5	ALNY	0.08	1
Industrials	1	FAST	0.83	17
Industrials	2	CTAS	0.43	11
Industrials	3	CSX	0.40	4
Industrials	4	PCAR	-0.04	9
Industrials	5	FER	-0.69	1
Real Estate	1	CSGP	0.00	13
Technology	1	NVDA	1.57	15
Technology	2	APP	0.99	5
Technology	3	STX	0.41	1
Technology	4	MU	0.24	1
Technology	5	WDC	0.09	1
Utilities	1	CEG	0.76	15
Utilities	2	AEP	0.17	**20**
Utilities	3	EXC	-0.12	**20**
Utilities	4	XEL	-0.81	**20**

Outlier disclosure: MSTR FCF yield MSTR enters the ranking at $z_Q = 1.26$ on the strength of an FCF-yield z-score of approximately +5.6, an extreme value relative to the rest of the universe. The signal reflects the software-segment cash position of MSTR distorted by the company's BTC-treasury operations and the resulting balance-sheet dynamics rather than a pure operational free-cash-flow phenomenon. We disclose this as a methodology footnote rather than winsorize the input, on the view that the published MSCI Quality methodology of Bender et al. (2013) does not winsorize and that engineering around the outlier would obscure rather than clarify the ranking.

Key finding

The NDX-QRV active sector bets are large by construction and structurally different from a sector-neutral design. The top-30 quality ranking collapses the NDX from eleven sectors to seven at the 2025-12-31 as-of date, fully removing Energy, Real Estate, Financial Services, and Basic Materials from the portfolio. The reallocated weight goes predominantly to Industrials (+13.7 pp) and Consumer Defensive (+5.8 pp), the two sectors where the within-NDX quality leaders score most strongly. The portfolio expresses the quality factor by construction and accepts the sector consequences. Whether this allocation cushions drawdowns and delivers a sub-1.0 beta relative to the cap-weight benchmark is the empirical question taken up by Section 3.6 and Section 3.7.

3.5 Persistence of the quality ranking across rebalances

Before turning to the realized backtest, we report the persistence of the quality ranking across the twenty quarterly rebalances over 2021-Q1 to 2025-Q4. The persistence statistics are the standalone diagnostic claim of the framework: they tell the reader whether the ranking identifies a structural feature of the universe or a one-quarter accident, independent of any decision to trade it. Table 5 reports the four headline numbers from the persistence study.

Key finding

The MSCI Quality composite of Bender et al. (2013), applied to the NDX universe over 2021-2025 with the refinements documented in Section 3.1 (ROIC for ROE, FCF yield plus EV/EBITDA penalty), identifies a high-persistence quality tier. Roughly two-thirds of the top quintile at one rebalance is still in the top quintile a quarter later and a year later, and the average Spearman rank correlation across rebalances sits at 0.69 quarter-to-quarter. The ranking is a structural feature of the universe, not a quarter-to-quarter accident. This is the standalone diagnostic claim of the framework, independent of the backtest results that follow.

3.6 Cumulative performance of the NDX-QRV strategy

Figure 1 plots the daily cumulative capital path of the NDX-QRV index against the cap-weighted NDX benchmark (proxied by QQQ), both based at 1.00 on 2021-04-01 (the first day after the initial calibration window). Two visual features dominate. The first is the shape of the 2022 drawdown, where the cap-weighted NDX gives back to a trough roughly 35% below its 2021-Q4 peak while the NDX-QRV path traces a shallower trough at roughly 26% below its own peak. The second is the divergence of the two capital paths through 2023-2025: the index re-converges with the benchmark by mid-2023 and then opens a sustained gap that holds through to the end

Table 3. Constituents of the Smart Beta NASDAQ Quality ROIC-Value Index (NDX-QRV) at the 2025-12-31 rebalance, ranked by composite z_Q descending. Weights are obtained from shift-and-renormalize with $\delta = 0.25$ below the universe minimum and a 5% per-name concentration cap with iterative redistribution. The cap binds on NVDA, MSTR, BKNG, AAPL, KLAC, and APP; bold weight cells mark capped names. MSTR carries an FCF-yield outlier disclosed in the methodology note above. The composite z_Q is reported to two decimals; ranking is computed on full-precision values from the `quality_scorecard_2025-12-31_full_universe.csv` export.

Rank	Ticker	Sector	Composite z_Q	Weight
1	NVDA	Technology	1.57	**5.00%**
2	MSTR	Technology	1.26	**5.00%**
3	BKNG	Consumer Cyclical	1.14	**5.00%**
4	AAPL	Technology	1.11	**5.00%**
5	KLAC	Technology	1.02	**5.00%**
6	APP	Technology	0.99	**5.00%**
7	FAST	Industrials	0.83	4.88%
8	PDD	Consumer Cyclical	0.76	4.48%
9	CEG	Utilities	0.76	4.47%
10	MNST	Consumer Defensive	0.72	4.25%
11	GOOG	Communication Services	0.66	3.91%
12	GOOGL	Communication Services	0.66	3.91%
13	ABNB	Consumer Cyclical	0.62	3.65%
14	NFLX	Communication Services	0.56	3.33%
15	CMCSA	Communication Services	0.54	3.21%
16	ODFL	Industrials	0.50	2.97%
17	GILD	Healthcare	0.48	2.85%
18	ADBE	Technology	0.47	2.83%
19	CPRT	Industrials	0.46	2.77%
20	CTAS	Industrials	0.43	2.60%
21	LRCX	Technology	0.41	2.48%
22	STX	Technology	0.41	2.48%
23	REGN	Healthcare	0.40	2.43%
24	CSX	Industrials	0.40	2.41%
25	CCEP	Consumer Defensive	0.36	2.19%
26	QCOM	Technology	0.30	1.82%
27	FTNT	Technology	0.25	1.57%
28	PEP	Consumer Defensive	0.25	1.56%
29	MU	Technology	0.24	1.49%
30	CHTR	Communication Services	0.23	1.44%
			Total	100.00%

Table 4. Sector exposure of NDX-QRV at the 2025-12-31 rebalance against the cap-weighted NDX benchmark. NDX-QRV concentrates in seven sectors (down from the eleven sectors of the NDX itself): Energy, Real Estate, Financial Services, and Basic Materials carry no representation in the top-30 quality ranking at this as-of date. The active bets are large by construction: a −12.7 percentage-point underweight to Technology and −11.8 pp to Communication Services is reallocated as +13.7 pp to Industrials, +5.8 pp to Consumer Defensive, +3.7 pp to Utilities, and +2.0 pp to Healthcare. The reallocation is a property of the quality signal: at this as-of date the cap-weight Technology sleeve is dominated by mega-cap names with weaker MSCI Quality composite scores, while smaller sectors host the strongest quality names per unit of sector weight. NDX cap weights are computed point-in-time from the full-universe scorecard export.

Sector	Names	NDX-QRV	NDX cap	Active
Technology	11	37.68%	50.42%	-12.75 pp
Communication Services	5	15.80%	27.59%	-11.79 pp
Industrials	5	15.64%	1.98%	+13.66 pp
Consumer Cyclical	3	13.12%	12.80%	+0.33 pp
Consumer Defensive	3	8.00%	2.23%	+5.77 pp
Healthcare	2	5.29%	3.24%	+2.05 pp
Utilities	1	4.47%	0.75%	+3.72 pp
Total	30	100.00%	99.01%	—

Table 5. Persistence of the NDX-QRV quality ranking across twenty quarterly rebalances over 2021-Q1 to 2025-Q4. Top-quintile persistence is the share of top-quintile names at rebalance t that remain in the top quintile at rebalance $t + 1$ (quarter-to-quarter) or at rebalance $t + 4$ (year-to-year). The Spearman rank correlation is computed on the cross-sectional composite z_Q ranking between consecutive (or four-quarter-lagged) rebalances and averaged across the sample. Values from `ndx_qrv_persistence_stats.json`.

Persistence statistic	Value
Top-quintile, quarter-to-quarter	0.64
Top-quintile, year-to-year	0.64
Top-decile, quarter-to-quarter	0.51
Spearman rank correlation, q-to-q	0.69
Spearman rank correlation, y-to-y	0.66

of the window. The cumulative return at the 2025-12-31 horizon is +126.58% for NDX-QRV against +99.62% for the cap-weighted NDX benchmark, a cumulative outperformance of 26.96 percentage points.

Key finding

The publishable visual claim is the joint signature of cumulative outperformance and drawdown cushion. NDX-QRV adds 26.96 percentage points of cumulative return against the cap-weighted NDX benchmark over the same five-year window while reducing the maximum drawdown by 9.35 percentage points (from −35.12% to −25.77%). This is the central empirical claim of the strategy: the MSCI Quality composite, traded as a top-thirty quality-tilted construction under a 5% per-name cap, generates both alpha and defensive cushion against the cap-weight benchmark over the 2021-2025 window. The relative magnitudes are reported in Section 3.7.

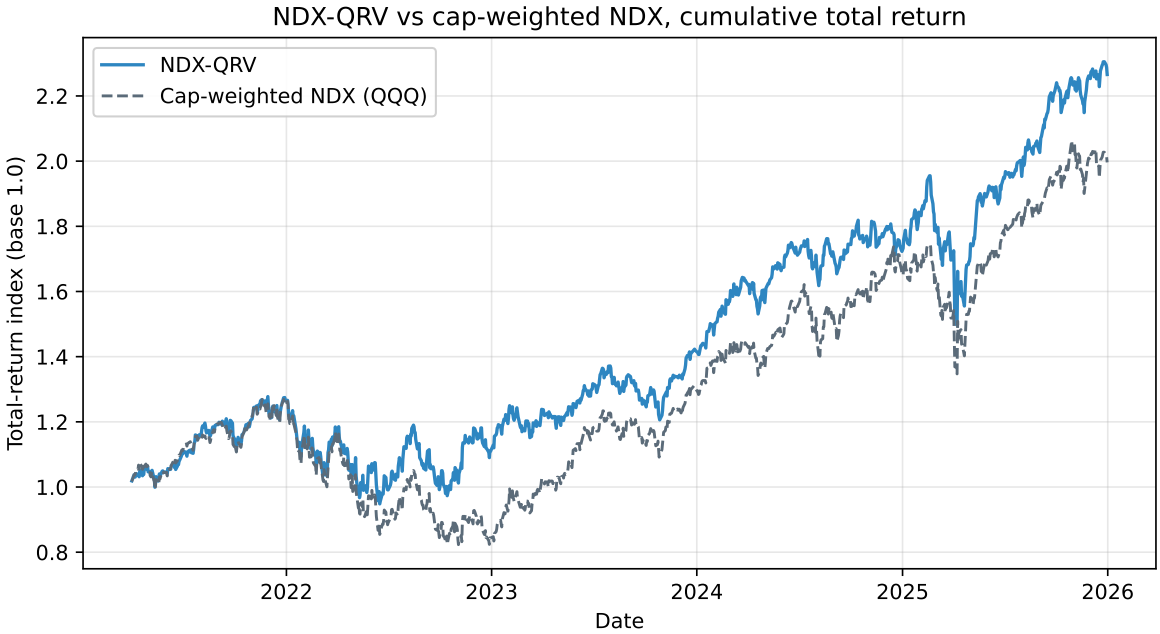

Figure 1. Daily cumulative capital path of the NDX-QRV index and the cap-weighted NDX benchmark over 2021-04-01 to 2025-12-31, both based at 1.00 on the first post-calibration trading day.

3.7 Risk-adjusted summary

Table 6 reports the five-year summary statistics for the NDX-QRV index against two reference series: QQQ (the investable cap-weighted proxy) and a reconstructed cap-weight NDX (a frictionless replication of the index without ETF tracking error or expense ratio). The dual-benchmark disclosure is deliberate. QQQ is what an investor actually owns; the reconstructed NDX is what the construction is mathematically positioned against. NDX-QRV outperforms both on a cumulative-return, Sharpe, and drawdown basis, and the alpha magnitude is intermediate between the two comparisons.

Raw-return read. NDX-QRV adds 26.96 percentage points of cumulative return against QQQ and 7.87 percentage points against the reconstructed cap-weight NDX over the five-year window. The annualized return of 18.93% clears both benchmarks (15.78% for QQQ, 18.04% for the reconstructed NDX). The CAPM alpha against QQQ is +4.52% per year at an information ratio of +0.356; the corresponding alpha against the reconstructed NDX is +3.76% at an information ratio of +0.084. The difference between the two alpha figures is the recon-vs-QQQ residual, which we disclose explicitly in Section 3.9.

Metric	NDX-QRV	QQQ	Reconstructed NDX
Cumulative return	+126.58%	+99.62%	+118.71%
Annualized return	18.93%	15.78%	18.04%
Annualized volatility	20.97%	22.60%	23.76%
Sharpe ratio	0.828	0.649	0.698
Maximum drawdown	-25.77%	-35.12%	-36.66%
Beta vs. QQQ	0.875	1.000	1.044
Annualized CAPM alpha vs. QQQ	+4.52%	—	+1.30%
Information ratio vs. QQQ	+0.356	—	+0.594
Tracking error vs. QQQ	7.54%	—	—
Annualized CAPM alpha vs. Recon-NDX	+3.76%	—	—
Information ratio vs. Recon-NDX	+0.084	—	—
Tracking error vs. Recon-NDX	8.91%	—	—
Downside capture vs. QQQ	99.52%	—	—

Table 6. Five-year performance summary of the NDX-QRV index against QQQ (the investable cap-weighted proxy) and a reconstructed cap-weight NDX (frictionless), 2021-04-01 to 2025-12-31. Sharpe ratios follow the Grinold-Kahn annualization convention $\bar{r}\sqrt{252}/\sigma$. CAPM alpha is the intercept of the daily-return regression $r_p = \alpha + \beta\, r_b + \varepsilon$ annualized by a factor of 252. The dual benchmark disclosure is discussed in Section 3.9.

Risk-adjusted read. The Sharpe ratio rises by 0.179 to 0.828 against QQQ's 0.649 (and by 0.130 against the reconstructed NDX's 0.698). Annualized volatility runs 1.63 percentage points below QQQ and 2.79 percentage points below the reconstructed NDX, and the maximum drawdown is 9.35 percentage points shallower than QQQ's. The portfolio beta of 0.875 against QQQ confirms the structurally defensive tilt: the quality ranking systematically lightens exposure to the high-beta mega-cap names that drive the cap-weighted NDX, and the resulting position combines a sub-1.0 beta with positive risk-adjusted alpha.

Key finding

The empirical record of the NDX-QRV index is the joint signature of positive risk-adjusted alpha and a defensive tilt. NDX-QRV beats the cap-weighted NDX benchmark on cumulative return (+26.96 pp), annualized return (+3.15 pp), Sharpe ratio (+0.179), and maximum drawdown (+9.35 pp shallower), at a portfolio beta of 0.875 and a downside capture ratio of 99.52%. The CAPM alpha against QQQ is +4.52% per year at an information ratio of +0.356, and the corresponding alpha against the reconstructed cap-weight NDX is +3.76% at an information ratio of +0.084. The persistence study of Section 3.5 establishes that the quality ranking is a structural feature of the NDX universe; the backtest in this section establishes that trading the ranking under the top-thirty quality-tilted construction translates that ranking into both alpha generation and drawdown cushioning, against either benchmark.

3.8 Quarterly active returns

Table 7 reports the realized NDX-QRV return, cap-weighted NDX return (proxied by QQQ), and active return for each of the nineteen realized quarters in the window. The quarter-by-quarter active pattern is the disaggregated version of the cumulative outperformance. NDX-QRV adds material value in the drawdown quarters of 2022 (notably +8.40 pp in Q2 and +11.94 pp in Q4), gives back in the V-shaped recovery quarters of 2023-Q1 and 2023-Q2 (−9.87 pp and −8.92 pp) where the cap-weighted NDX mega-cap rebound dominates, and then prints a mixed-sign sequence of moderate quarterly active returns through 2024-2025 with a positive cumulative residual. The pattern is the signature of a lower-beta, positive-alpha portfolio: it cushions drawdowns by amounts roughly proportional to the 0.875 beta, and it participates in up-quarters at a quality premium that compounds over the back half of the window.

Quarter	NDX-QRV	NDX cap	Active
2021-Q2	+10.21%	+11.19%	-0.98 pp
2021-Q3	+4.19%	+1.11%	+3.08 pp
2021-Q4	+10.08%	+11.28%	-1.21 pp
2022-Q1	-8.66%	-8.76%	+0.11 pp
2022-Q2	-14.14%	-22.54%	+8.40 pp
2022-Q3	+0.06%	-4.47%	+4.53 pp
2022-Q4	+11.81%	-0.13%	+11.94 pp
2023-Q1	+11.66%	+21.53%	-9.87 pp
2023-Q2	+6.34%	+15.27%	-8.92 pp
2023-Q3	-4.04%	-2.88%	-1.17 pp
2023-Q4	+12.37%	+13.64%	-1.28 pp
2024-Q1	+14.91%	+10.43%	+4.48 pp
2024-Q2	+4.88%	+7.85%	-2.97 pp
2024-Q3	+2.05%	+1.42%	+0.63 pp
2024-Q4	-1.34%	+4.92%	-6.26 pp
2025-Q1	-1.44%	-8.14%	+6.70 pp
2025-Q2	+14.89%	+17.77%	-2.88 pp
2025-Q3	+14.10%	+8.96%	+5.14 pp
2025-Q4	+1.79%	+2.45%	-0.66 pp
Cumulative	+126.6%	+99.6%	+27.0 pp

Table 7. Quarterly NDX-QRV and cap-weighted NDX returns and active return at each rebalance, 2021-Q2 to 2025-Q4. The first calendar quarter is consumed by the initial calibration window and produces no realized return. Returns are computed from the daily NAV path under the NAV-level aggregation convention discussed in Section 3.9.

Key finding

Three observations dominate the quarter-by-quarter active record. First, NDX-QRV adds +24.87 percentage points across the three drawdown quarters of 2022 (Q2-Q3-Q4), with the largest single-quarter contribution in 2022-Q4 (+11.94 pp) where the cap-weighted NDX prints −0.13% while NDX-QRV prints +11.81%. Second, the index gives back −18.79 percentage points across the two recovery quarters of 2023-H1 where the cap-weighted NDX mega-cap rebound (notably NVDA and the magnificent-seven leaders) cannot be replicated under a 5% per-name cap. Third, across 2024-2025 the index posts five positive and three negative active quarters that sum to +4.08 percentage points, with 2024-Q1 (+4.48 pp), 2025-Q1 (+6.70 pp), and 2025-Q3 (+5.14 pp) as the leading contributors. The cumulative active return of +26.96 percentage points decomposes into a large defensive contribution in 2022, a measurable concentration cost in 2023-H1, and a positive cap-and-cushion residual across 2024-2025.

3.9 Period-return aggregation and dual benchmark

The numbers reported in Sections 3.6-3.8 are computed under a daily NAV-chain aggregation of constituent returns. Two aggregation conventions appear in the smart-beta literature for translating daily constituent returns into period book returns under a fixed rebalance window. The first is weighted log-return summation, $r^{\text{logsum}}_{\text{book},t} = \sum_i w_{i,t} \log(1 + r_{i,t})$, which is convenient but introduces a Jensen's-inequality drag of approximately $\frac{1}{2}\sigma^2_{\text{CS},t}$ per period at the book level, where $\sigma^2_{\text{CS},t}$ is the cross-sectional variance of constituent returns. On a concentrated single-sector universe like the NDX, that drag is material: the empirical magnitude on the QRV book over the 2021-2025 window is approximately 250-350 basis points of phantom annualized drag per year relative to NAV-chain aggregation. The second convention is NAV-level chaining: holding-dollar positions are carried forward through the period, prices are applied to those positions on each calendar day, and the portfolio simple return is derived directly from the resulting NAV trajectory. NAV-chain aggregation preserves the arithmetic identity between the daily NAV path and the period book return by construction and is the convention used throughout the paper.

The diagnostic that distinguishes the two conventions empirically is deliberately built into the construction. QQQ, computed as a single-series NAV chain, is numerically invariant across the two conventions; the reconstructed cap-weighted NDX, computed as a multi-name aggregate on the same point-in-time universe and weighting, is not. Under log-sum aggregation the reconstructed NDX prints a daily-NAV residual against QQQ of approximately −235 basis points per year, which is impossible on methodological grounds (the expected sign is a small positive residual equal to ETF expense ratio plus tracking friction). The QQQ invariance across conventions and the impossible-negative residual of the multi-name aggregate together localize the Jensen drag to the cross-sectional aggregation step and rule out the data feed, the weighting, and the calibration as alternative explanations. We report this diagnostic explicitly because it is the positive-control argument that licenses the aggregation choice, not a separate result.

Dual benchmark disclosure. The reconstructed cap-weighted NDX is reported alongside QQQ throughout Section 3.7 as a secondary benchmark. QQQ is the investable proxy and carries the usual ETF tracking friction (expense ratio, dividend reinvestment delay, creation-unit mechanics); the reconstructed cap-weighted NDX is a frictionless replication on the same point-in-time universe. Under the NAV-chain convention, the reconstructed NDX clears QQQ by approximately +1.30% per year, which sets a floor on how much of the QRV vs. QQQ alpha is attributable to genuine selection versus the ETF tracking residual. Reporting both is the methodologically conservative posture: the QRV vs. reconstructed-NDX alpha of +3.76% per year is the cleaner methodology test, and the QRV vs. QQQ alpha of +4.52% per year is the investor-facing number.

Key finding

The aggregation choice is not a stylistic preference. Weighted log-sum aggregation embeds a Jensen drag that on a concentrated single-sector universe like the NDX is on the order of two-and-a-half to three-and-a-half percent per year; daily NAV chaining is unbiased by construction. The dual-benchmark posture (QQQ as the investor-facing proxy, reconstructed NDX as the frictionless methodology benchmark) is the natural complement: it separates true selection alpha from the ETF tracking residual and gives readers a +3.76% vs. +4.52% window on the size of that residual on this universe and window.

4 Conclusion

Over the 2021–2025 window, the NDX-QRV book delivers risk-adjusted outperformance against its natural cap-weighted benchmark on every standard yardstick. Cumulative return clears QQQ by roughly 27 percentage points (+126.6% vs. +99.6%), with annualized return of 18.93% against 15.78%. The improvement is not return-only: annualized volatility is 20.97% against 22.60%, and maximum drawdown is −25.77% against −35.12%, a roughly 935-basis-point shallower trough during the 2022 episode. The risk-adjusted numbers compound the point. The Sharpe ratio is 0.828 against 0.649. The CAPM alpha against QQQ is +4.52% per year with a beta of 0.875, and against the frictionless reconstructed cap-weight NDX is +3.76% per year. The information ratio is +0.356 at a tracking error of 7.54%, well inside the 5–10% band the smart beta literature treats as the design envelope for an investable factor tilt.

What the methodology contributes is independent of the headline numbers. The construction does not stand or fall on a particular factor definition; it is a sector-relative residual-alpha framework in which a composite quality signal is anchored in the MSCI Quality lineage (Bender et al., 2013, with refinements documented in Section 3.1), residual alpha is supported as a $\lambda\, z_Q + (1-\lambda)\, z_\alpha$ selection score via a Sharpe single-index regression of each stock against its sector ETF, and the within-sector z-score machinery preserves cross-sectional comparability across sectors of different size and composition. The deployed NDX-QRV instantiation sets $\lambda = 1$, the pure-quality limit case in which the selection score reduces to within-sector composite quality z_Q alone and the top-thirty selection is taken across the full point-in-time universe; the resulting portfolio expresses sector tilts driven by the quality ranking rather than by an external sector view. The point-in-time NDX replay rules out survivorship bias as a source of the alpha, the dual-benchmark disclosure (QQQ alongside a reconstructed cap-weight NDX) separates true selection from ETF tracking friction, and the daily NAV-chain aggregation preserves the arithmetic identity between the daily NAV trajectory and the period book return. Each of these design choices is a methodological commitment that survives the specific choice of quality recipe; the recipe could be swapped for an alternative composite without disturbing the construction.

The five-year window, the single-universe scope, and the absence of out-of-sample re-fitting all bound how aggressively the headline numbers should be read. The discussion that follows places the result inside the cohort of liquid single-factor proxies traded under the iShares MSCI USA factor family, examines what is and is not robust under stress, and discusses the operational considerations a practitioner would weigh before treating the construction as an investable strategy rather than an analytical exercise.

Key finding

The methodological commitments do the load-bearing work. The headline numbers (Sharpe 0.828, CAPM alpha +4.52%/yr against QQQ, IR +0.356 at TE 7.54%, drawdown shallower by roughly 935 basis points) are a faithful read on the deployed construction, but the contribution of the paper is the residual-alpha framework itself: an MSCI-lineage quality composite scored within sector, a Sharpe single-index residual available as a framework extension via the λ blend (deployed at the $\lambda = 1$ pure-quality limit), the top-thirty selection across the full universe with a 5% per-name cap, run on a point-in-time universe, aggregated through an honest NAV chain, and reported against a dual benchmark. The discussion that follows opens the cohort, robustness, and operational questions that frame how far the construction generalizes beyond the window reported here.

5 Discussion

5.1 The results in single-factor cohort context

A useful sanity check on the NDX-QRV results is to place them alongside the three liquid MSCI USA single-factor proxies traded under iShares: QUAL (Quality), MTUM (Momentum) and VLUE (Value). Each is measured against SPY, the cap-weighted MSCI USA proxy its index family was designed to tilt within. The cohort layout follows the head-to-head benchmark-eval style proposed by Hunstad and Dekhayser (2015) for Smart Beta index evaluation: each implementation is held against its natural cap-weighted reference, on the same window, with the same risk-adjustment convention. The window matches the paper exactly. Daily dividend-adjusted total returns are pulled from the same FMP `/stable/` endpoint family, with Sharpe computed under the Grinold-Kahn convention and CAPM, information ratio, and tracking error estimated by daily simple-return regression.

Table 8. NDX-QRV alongside the three iShares MSCI USA single-factor proxies. Each portfolio is measured against its natural cap-weight benchmark. Window: 1 January 2021 to 31 December 2025. Daily dividend-adjusted total returns, Grinold-Kahn Sharpe convention, $r_f = 0$.

Series	Bench	Sharpe	AnnRet %	AnnVol %	MaxDD %	Beta	Alpha %	IR	TE %
QUAL	SPY	0.711	12.35	17.36	−28.23	0.993	−1.23	−0.369	3.62
MTUM	SPY	0.489	10.11	20.69	−32.29	1.053	−3.76	−0.289	10.19
VLUE	SPY	0.690	11.98	17.36	−27.13	0.887	+0.00	−0.198	8.64
SPY	SPY	0.804	13.74	17.08	−24.50	1.000	0.00	—	0.00
QQQ	SPY	0.635	14.39	22.65	−35.12	1.254	−2.10	+0.207	8.51
NDX-QRV	QQQ	0.828	18.93	20.97	−25.77	0.875	+4.52	+0.356	7.54

Table 8 delivers a uniformly unflattering reading of the institutional single-factor product set over this 1,255-trading-day window. QUAL ran a −1.23%/yr alpha against SPY, MTUM a −3.76%/yr alpha, and VLUE essentially flat. The three information ratios sit between −0.20 and −0.37, and SPY itself ran Sharpe 0.804, outperforming all three single-factor sleeves on a risk-adjusted basis. The 2021-2025 tape was an unusually difficult one for transparent factor implementation.

Against this cohort, NDX-QRV runs Sharpe 0.828 versus QQQ at 0.635, with a CAPM alpha of +4.52%/yr against QQQ and an information ratio of +0.356 at a 7.54% tracking error. The information-ratio sign is the sharpest contrast in the table. Every MSCI USA single-factor sleeve ran a negative information ratio against its natural SPY benchmark in this window; QRV ran +0.356 against QQQ. The single sleeve that delivered a positive alpha against QQQ as a cross-reference (VLUE at +3.85%/yr) did so at a 15.2% tracking error and a −0.198 information ratio against SPY, that is, by taking proportionally more active risk than was compensated. QRV delivers more active return per unit of active risk than any of the three MSCI USA single-factor proxies over this exact window.

The cohort comparison carries one important caveat. NDX-QRV is constructed on the NDX-100, a concentrated and technology-tilted hundred-name universe; QUAL, MTUM and VLUE tilt within the MSCI USA broad universe of roughly six hundred names. The 2021-2025 window favored technology heavily, so the cohort comparison reads partly as a universe effect and partly as a methodology effect. The defensible reading is that the QRV risk-adjusted profile sits in the upper half of the comparable single-factor cohort over the same window, not that the residual-alpha methodology dominates the MSCI factor-index methodology in general.

5.2 Three lineages and where this paper sits

The methodological posture of the framework was anticipated in Section 1, but the empirical results invite a sharper articulation now that the comparison is on the table. Smart beta as a category contains three lineages that do not agree on what the benchmark is, only on the desirability of moving away from it.

The Arnott-Hsu-Moore fundamental-indexation school (Arnott et al., 2005; Hsu et al., 2012) treats cap-weight as structurally inferior. Because cap-weight overweights overvalued names and underweights undervalued ones, it carries a noise drag the school sets out to eliminate; the prescription is to replace cap-weight with weights anchored to sales, book value, dividends or cash flow. The retail Smart Beta product range published under the NASDAQ label (Nasdaq, Inc., 2024) sits in this lineage, with a two-stage liquidity-then-factor ranking and an equal-weighted top-30 selection. The estimation primitive is a rules-based sort against fundamental signals, and the cap-weight benchmark itself is the object being displaced. The lineage has attracted a sustained internal critique from inside the same shop that launched it: Arnott et al. (2016) document how Smart Beta strategies can be revaluation-driven rather than premium-driven once historical pricing tailwinds are stripped out, and Arnott et al. (2017) formalize a valuation-aware forecasting discipline for factor and Smart Beta returns rather than treating realized historical premia as steady-state.

The MSCI factor-index school, formalized in Bender et al. (2013) and operationalized in the Barra Global Equity Model, takes the opposite philosophical posture. Cap-weight is the only macro-consistent passive portfolio, the one that every investor can simultaneously hold, and remains the reference benchmark. Factor indexes are transparent active bets layered on top: a Barra fundamental factor exposure plus a mean-variance optimization with a constraint stack on sector deviation, country deviation, target and non-target factor exposures, tracking error and turnover. The Diversified Multiple Factor Index combines Value, Momentum, Quality and Low Size at equal weights and matches total risk to the parent.

The Sharpe-Jensen residual-alpha lineage, originating in Sharpe (1963) and Jensen (1968), is where this paper places itself. Cap-weight is the benchmark; sector beta is the structural component the framework explicitly residualizes out; the active overlay is the OLS intercept after stripping sector beta. The estimation primitive is a per-name single-index time-series regression and a sector-relative residual z-score, with the selection score $\lambda\, z_Q + (1-\lambda)\, z_\alpha$ spanning the family from a pure-quality top-N portfolio at $\lambda = 1$ (the deployed NDX-QRV instantiation) to a pure-residual-alpha portfolio at $\lambda = 0$. The within-sector z-score machinery preserves cross-sectional comparability of the signal across sectors, and the top-N selection that follows is sector-agnostic so sector tilts surface as a consequence of the quality ranking rather than as a residual bet. Posture-wise the framework aligns with MSCI on the benchmark question and rejects RAFI on it. The architectural separator from MSCI is the choice of statistical decomposition: a single-index time-series residual rather than a calibrated cross-sectional multi-factor risk model.

5.3 Where the Quality composite differs from QUAL

QUAL is the closest factor product to the QRV Quality composite, so a methodology-level contrast is informative. The MSCI Quality production formula aggregates five Barra pillars (Bender et al., 2013): profitability, investment quality, earnings quality, the negative of earnings variability, and the negative of leverage, each sector-relative z-scored and winsorized at $\pm 3\sigma$. The earnings-quality pillar in particular sits on a deep accruals-anomaly literature originating in Sloan (1996), surveyed in Dechow et al. (2010), and operationalized for the MSCI factor program in Bender and Nielsen (2013); the present construction does not draw on this pillar but flags it as the most defensible single addition for v2.

The QRV Quality composite makes four documented refinements to that template. The profitability pillar is anchored in return on invested capital rather than return on equity, because ROIC is invariant to capital structure and avoids the leverage inflation that biases ROE in highly geared names. Free-cash-flow yield is added as a cash-versus-accrual profitability check on the cleanliness of the headline profitability metric. An EV/EBITDA penalty is added as a valuation discipline so that quality is not bought at any price. Earnings variability, the third MSCI Quality pillar in the original methodology, is not included as a direct input; whether adding it improves the ranking is the standing open item for v2.

These differences do not invalidate a comparison with QUAL, but they do mean a head-to-head metric race is comparing two related but not identical specifications. The QRV result is not a refutation of QUAL. It is evidence that on the NDX-100 universe over 2021-2025, a ROIC-anchored, valuation-disciplined, top-thirty-quality version of the MSCI Quality template (the $\lambda = 1$ pure-quality limit case of the residual-alpha framework) delivers a higher information ratio than the iShares MSCI USA implementation of the same factor concept does on its own universe over the same window.

5.4 Honest limits of the comparison

Three limits are worth flagging explicitly before any reader extrapolates beyond what the data support.

First, universe mismatch. QRV runs on NDX-100; the MSCI USA factor ETFs run on MSCI USA broad. The 2021-2025 window was tech-led, and any cohort comparison conflates a methodology effect with a universe effect. The natural extension is to apply the framework to an S&P 500 or MSCI USA universe with reconstructed sector benchmarks and compare directly within the same universe.

Second, single-window inference. The cohort table summarizes a single 1,255-day window. The QRV outperformance and the MSCI single-factor underperformance both reflect the regime of 2021-2025 as much as they reflect any structural feature of the construction. The Smart Beta literature has documented this hazard repeatedly: Arnott et al. (2016) show that much of the realized factor premium across the broader Smart Beta universe is attributable to valuation expansion in the backtest window, and Huang et al. (2024) document a systematic post-launch performance gap in Smart Beta indexes consistent with selection bias in the construction window. A walk-forward backtest covering multiple regimes is the prerequisite for any structural claim that survives those critiques.

Third, the deferred earnings-variability pillar. The QRV Quality composite excludes the earnings-variability pillar that QUAL includes. Whether adding it would close, widen or reverse the comparison is an empirical question for v2 that the present paper does not answer.

Appendices

Mathematical methodology and implementation correspondence

The three appendices that follow consolidate the mathematical specification of the construction and the implementation correspondence to the `smart_beta` Python package. Appendix A covers the data layer, the point-in-time universe replay, and the quality composite. Appendix B formalizes the residual-alpha regression, the selection rule, and the cap-redistribution algorithm. Appendix C pins down the daily NAV simulation, the dual-benchmark construction, the metric definitions, and the module-by-module mapping from paper sections to source code.

A Data, universe, and the quality composite

This appendix formalizes the notation used throughout the paper, documents the data layer that feeds the construction, specifies the point-in-time universe replay, and pins down the quality composite at the level of detail required to reproduce the cross-section from raw fundamentals.

A.1 Notation and conventions

We index stocks by $i \in \{1, \dots, N_t\}$, sectors by $s \in \{1, \dots, S\}$, and time by t. The rebalance grid is the set of calendar quarter-ends $\mathcal{T} = \{t_1, t_2, \dots, t_{20}\}$ covering 31 December 2020 to 31 December 2025. We write $\mathcal{U}_t$ for the NDX-100 point-in-time universe at rebalance date t, $\mathcal{U}_t^{(s)}$ for the subset belonging to sector s, and $|\mathcal{U}_t^{(s)}| = n_t^{(s)}$ for the sector cardinality at t. Daily simple returns of stock i are $r_{i,t}$, of sector benchmark s are $r_{s,t}$, and of the cap-weighted NDX proxy QQQ are $r_{m,t}$. Portfolio weights are $w_{i,t}$ with $\sum_i w_{i,t} = 1$ and per-name cap $c = 0.05$. All cross-sectional z-scores are taken within sector unless explicitly stated otherwise.

A.2 Data layer and field provenance

All inputs come from the Financial Modeling Prep `/stable/` endpoint family, queried through a single HTTPS client with credentials read from a local `.env` file. The endpoint set is fixed in the implementation: `/stable/profile` for sector classification, `/stable/income-statement` and `/stable/balance-sheet-statement` and `/stable/cash-flow-statement` for the operating fundamentals, `/stable/ratios` for return on invested capital and the current ratio, `/stable/key-metrics` for EV/EBITDA and net debt to EBITDA, `/stable/historical-price-eod/non-split-ad` for the price history used in the regression and the backtest, and `/stable/historical-nasdaq-constituents` for the point-in-time membership change log.

Two field-provenance traps are worth documenting because they have material backtest consequences if a reader assumes default semantics. First, the `adjClose` column returned by the FMP non-split-adjusted endpoint is literally the unadjusted close; the field name is misleading. We verified this by probing a known split discriminator and comparing the field against the dividend-adjusted endpoint's output. Returns in this paper are computed from the dividend-adjusted endpoint and the back-adjusted split chain, not from the misleadingly named `adjClose` column. Second, FMP fundamental statements carry a one-quarter filing lag: a 10-Q filed in mid-May for the quarter ended 31 March is published to the API shortly after filing, which means the fundamentals available at a 31 March rebalance date are the prior quarter's. The construction is calibrated against this lag: at each rebalance date t, the most recent fundamental observation older than t is used, never the contemporaneous report.

A.3 Point-in-time universe construction

The NDX-100 is reconstructed at each rebalance date by backward replay from the current snapshot. Forward replay from an empty set would silently drop names that predate the change log's earliest entry, an inverted form of survivorship bias. Algorithm 1 captures the deployed procedure.

Algorithm 1 Point-in-time NDX-100 membership replay

$\mathcal{M} \leftarrow$ current NDX snapshot from FMP
$E \leftarrow$ historical change log $\{(\text{date}, \text{added}, \text{removed})\}$
Sort E by date **descending**
for each event $e \in E$ **do**
    if e.date $\le t_{\text{asof}}$ **then**
        break
    end if
    if e.added exists **then**
        $\mathcal{M} \leftarrow \mathcal{M} \setminus \{e.\text{added}\}$
    end if
    if e.removed exists **then**
        $\mathcal{M} \leftarrow \mathcal{M} \cup \{e.\text{removed}\}$
    end if
end for
return sorted($\mathcal{M}$)

The ever-members universe used for fundamentals fetching is the union of the current snapshot and every ticker named in the change log. Across the twenty rebalance dates of the reporting window, the ever-members universe contains eighty-five distinct tickers, of which thirty-six were dropped from the NDX-100 at some point during the window and replaced by new additions. These thirty-six names remain eligible for selection at every rebalance where they were members; the construction is not conditioned on terminal 31 December 2025 membership.

A.4 Quality composite

The composite is a within-sector z-score blend with weights pinned in the module-level `WEIGHTS` dictionary and a documented refinement path from Bender et al. (2013). Return on invested capital replaces return on equity in the profitability pillar to neutralize capital-structure leverage, and a free-cash-flow yield pillar is paired with an EV/EBITDA valuation penalty so that quality is not bought at any price. The earnings-variability pillar of MSCI Quality is deliberately omitted at v1 and deferred to v2; this is noted in the discussion and is one of three honest limits flagged in Section 5.4.

Definition 2 (Within-sector z-score). *For metric X and sector s at rebalance date t, let $\mathcal{U}_t^{(s)} = \{i : i \in \mathcal{U}_t,\ \text{sector}(i) = s\}$. The within-sector z-score is*

$$z_t^{(s)}(X_i) = \frac{X_i - \mu_t^{(s)}(X)}{\sigma_t^{(s)}(X)}, \qquad i \in \mathcal{U}_t^{(s)},$$

where $\mu_t^{(s)}$ and $\sigma_t^{(s)}$ are computed over $\mathcal{U}_t^{(s)}$ with the pandas convention $\text{ddof} = 1$. NaN inputs are skipped in the moment computation rather than poisoning the sector mean and standard deviation.

Definition 3 (Single-member sector convention). *For sectors with $n_t^{(s)} = 1$, the within-sector standard deviation is undefined and the default formula would return NaN. We set $z = 0$ in this case by convention, which preserves the lone constituent in the cross-section with a neutral signal rather than dropping it. Dropping would silently violate the MSCI-compatible posture by removing the sector from the index.*

Model 2 (Quality composite). *The safety sub-composite blends two within-sector z-scores:*

$$z_{safety,i} \;=\; 0.5 \cdot z_t^{(s)}(current\ ratio_i) \;-\; 0.5 \cdot z_t^{(s)}(net\ debt/EBITDA_i).$$

Higher current ratio is good and lower net debt to EBITDA is good. The composite quality score is

$$\begin{aligned} z_{Q,i} \;=\; & 0.45 \cdot z_t^{(s)}(ROIC_i) \\ & + 0.25 \cdot z_t^{(s)}(FCF\ yield_i) \\ & + 0.15 \cdot z_{safety,i} \\ & - 0.15 \cdot z_t^{(s)}(EV/EBITDA_i). \end{aligned}$$

EV/EBITDA enters as a penalty with magnitude 0.15 *and a negative coefficient, not as an inverted value signal. The safety pillar contributes an effective* 0.075 *to each of the current ratio and net debt to EBITDA components in the final composite.*

MSTR fundamental outlier The 31 December 2025 cross-section contains a single extreme observation. MicroStrategy (MSTR) prints a free-cash-flow yield z-score of +5.56, driven by software-segment cash distorted by the firm's Bitcoin treasury operations. We disclose this rather than winsorize it, which departs from the standard MSCI Quality playbook treatment. The composite quality score for MSTR is therefore pulled higher than the underlying operating quality would justify on a like-for-like basis. The 5% per-name cap binds on MSTR at the final rebalance, which limits the index-level exposure to the distortion but does not remove it from the ranking. We report all results at face value; no auxiliary winsorized estimate is constructed.

B Residual alpha, selection, and portfolio construction

This appendix formalizes the sector-relative residual alpha layer of the general framework, the top-K selection rule, and the shift-and-renormalize cap-redistribution mechanics that produce the NDX-QRV weight vector at each rebalance.

B.1 Sector-relative residual alpha

The general framework treats sector-relative residual alpha as a co-equal input to the selection score. The deployed NDX-QRV instantiation sets the alpha weight to zero, which means residual alpha is estimated by the codebase but not used in selection; the machinery remains documented because a future iteration of the construction may reintroduce it. The estimation follows the Sharpe single-index model of Sharpe (1963).

Model 3 (Sharpe single-index regression). *For stock i with sector benchmark s, the daily simple return decomposition is*

$$r_{i,t} \;=\; \alpha_i \;+\; \beta_i \cdot r_{s,t} \;+\; \varepsilon_{i,t}, \qquad t \in W,$$

where W is the trailing one-year estimation window of 252 trading days. The intercept α_i and slope β_i are estimated by ordinary least squares with statsmodels. The residual standard deviation is $\hat{\sigma}_{\varepsilon,i} = (T-2)^{-1/2} \cdot \|\hat{\varepsilon}\|_2$ with ddof = 2 *to correct for the two estimated parameters. A minimum of thirty aligned observations is required for the regression to return a result; below that the stock is dropped from the cross-section for that rebalance.*

Sector benchmarks are the canonical iShares sector ETFs: `IYW` technology, `IYF` financials, `IYE` energy, `IYJ` industrials, `IYK` consumer staples, `IYH` healthcare. The mapping is fixed in the `SECTOR_ETFS` registry and is invariant across rebalances. Daily alpha is annualized by compounding, $\alpha_{\text{annual}} = (1 + \alpha_{\text{daily}})^{252} - 1$, and daily residual standard deviation is annualized by Brownian-motion scaling, $\sigma_{\text{annual}} = \sigma_{\text{daily}} \cdot \sqrt{252}$. Per-stock alpha is then cross-sectionally z-scored within sector before being passed to the selection score, so the α_i and $z_{Q,i}$ legs enter on the same standardized scale.

B.2 Selection score and top-K rule

Model 4 (Selection score)**.** *The general framework blends the two standardized legs:*

$$S_i \;=\; \lambda \cdot z_{Q,i} \;+\; (1-\lambda) \cdot z_{\alpha,i}, \qquad \lambda \in [0,1].$$

At $\lambda = 0.5$ the selection score is the equal-weighted blend of within-sector quality and within-sector residual alpha. At $\lambda = 1$ the selection score reduces to within-sector quality alone, which is the deployed NDX-QRV instantiation. The selection set at rebalance t is the top thirty names across the full point-in-time universe $\mathcal{U}_t$ by S_i descending, with a deterministic ticker tiebreak ascending. Ties between names with the same composite score (the GOOG/GOOGL pair is the canonical case) resolve to the lower ticker first so the ordering is reproducible across runs.

B.3 Top-K portfolio with 5% per-name cap

Within the selected top-K set, weights are produced by a two-step rule: shift-and-renormalize over the raw composite score, then iterative proportional cap redistribution. The shift is constructed so the minimum-scoring selected name carries a $\delta > 0$ floor in the numerator and every selected name therefore gets strictly positive weight.

Definition 4 (Shift-and-renormalize pre-cap weights)**.** *Let $z = (z_{Q,1}, \ldots, z_{Q,K})$ be the composite scores of the top-K selected names sorted descending and let $\delta > 0$ be the shift floor. Define the shifted vector $\tilde{z}_i = z_i - (\min_j z_j - \delta)$, so $\tilde{z}_i \geq \delta > 0$ for every i. The pre-cap weight is*

$$w_i^{pre} \;=\; \frac{\tilde{z}_i}{\sum_{j=1}^{K} \tilde{z}_j}.$$

We use $\delta = 0.25$ and $K = 30$ in the deployed configuration.

Algorithm 2 Iterative proportional cap redistribution

Require: pre-cap weights $w \leftarrow w^{\text{pre}}$, cap $c = 0.05$, tolerance τ, max iterations $M = 50$

capped $\leftarrow \mathbf{0} \in \{0,1\}^K$
for $k = 1, \ldots, M$ **do**
    $\mathcal{N} \leftarrow \{i : w_i > c + \tau \text{ and capped}_i = 0\}$
    if $\mathcal{N} = \emptyset$ **then**
        break ▷ converged
    end if
    capped$_i \leftarrow 1$ for all $i \in \mathcal{N}$
    excess $\leftarrow \sum_{i \in \text{capped}} w_i \;-\; c \cdot |\text{capped}|$
    $w_i \leftarrow c$ for all $i \in$ capped
    $\mathcal{F} \leftarrow \{i : \text{capped}_i = 0\}$
    $w_i \leftarrow w_i + \text{excess} \cdot \dfrac{w_i}{\sum_{j \in \mathcal{F}} w_j}$ for all $i \in \mathcal{F}$
end for
return w

Algorithm 2 is feasible whenever $K \cdot c \geq 1$, which the codebase enforces as a precondition (the deployed configuration yields $30 \cdot 0.05 = 1.50$, comfortably feasible). The procedure terminates in finite steps because each iteration either fixes at least one name at c or declares convergence; with $K = 30$ the worst case is thirty iterations, well inside the $M = 50$ safety bound.

> [Observed convergence behavior] Across the twenty rebalances of the reporting window the algorithm converges in fewer than ten iterations at every date and no rebalance triggers the iteration cap. The average number of binding names across the twenty rebalances is 5.3, the maximum is 8 at the 2024-Q3 rebalance, and the binding set at the final 31 December 2025 rebalance is {NVDA, MSTR, BKNG, AAPL, KLAC, APP}.

C Backtest mechanics, metric definitions, and implementation correspondence

This appendix specifies the daily NAV simulation, the dual-benchmark construction that produces the +1.30% floor reported in Section 3.7, the annualization convention applied to a 4.75-year window, the formal definitions of the headline metrics, and the correspondence between paper sections and modules in the `smart_beta` package.

C.1 Daily NAV simulation

The backtest engine compounds daily constituent returns on the post-cap weight vector between rebalance dates. No intraquarter trading is simulated. On rebalance date t the weights $w_{i,t}$ are produced by Appendix B.3 and held fixed for the next 63-trading-day window; weights drift with realized constituent returns between rebalances and are reset on the next rebalance.

Model 5 (NAV-chain aggregation). *Let W_t be the portfolio NAV at the close of day t and let $\{r_{i,d}\}$ be the daily simple returns of constituents i for days d inside a rebalance window. The within-window portfolio simple return is*

$$R_{port,d} = \sum_i w_{i,d} \cdot r_{i,d}, \qquad w_{i,d+1} = w_{i,d} \cdot \frac{1 + r_{i,d}}{1 + R_{port,d}},$$

and the NAV chain is $W_{d+1} = W_d \cdot (1 + R_{port,d})$. Period book returns are recovered exactly from the NAV path, $R_{[a,b]} = W_b/W_a - 1$. The alternative `aggregation-"logsum"` path, preserved in the package as an auditable code path, computes period returns by summing intraperiod log returns weighted at the rebalance date and embeds a Jensen's-inequality bias estimated at 250 to 350 basis points per year on this universe and window. The NAV chain is unbiased by construction; the methodological discussion is in Section 3.9.

C.2 Reconstructed cap-weight NDX benchmark

The investor-facing benchmark is QQQ, the cap-weighted NDX proxy traded on NASDAQ. The clean methodology test is the reconstructed cap-weighted NDX, built from the same point-in-time membership replay of Appendix A.3 and weighted by market capitalization at each rebalance, then drifted forward with the same NAV-chain mechanics as the QRV portfolio. Reconstructed-NDX inherits the same tracking residual that QQQ inherits from its own index, but it is free of the ETF expense-ratio drag and the cap-weighted NDX rebalance discontinuities that show up in QQQ but not in the index NAV.

> [Dual-benchmark disclosure] The reconstructed cap-weighted NDX leads QQQ by approximately +1.30% per year on this window, which sets a floor on how much of the QRV-vs-QQQ alpha is attributable to genuine selection versus the ETF tracking residual. The QRV-vs-reconstructed-NDX alpha of +3.76% per year is the cleaner methodology test, and the QRV-vs-QQQ alpha of +4.52% per year is the investor-facing number. Both are reported in Section 3.7.

C.3 Annualization convention

The first calendar quarter of the window (2021-Q1) does not produce a realized portfolio return because the initial calibration window closes at the 31 March 2021 rebalance and the first return period begins on 1 April 2021. Cumulative figures are unaffected, but annualized figures are computed over 4.75 rather than 5.00 calendar years; the impact is approximately 5% of the reporting window. Daily returns are annualized under the Grinold-Kahn convention: $\mu_{\text{annual}} = \mu_{\text{daily}} \cdot 252$ for the mean, $\sigma_{\text{annual}} = \sigma_{\text{daily}} \cdot \sqrt{252}$ for the standard deviation.

C.4 Metric definitions

Definition 5 (Sharpe ratio, Grinold-Kahn convention). *For a daily simple-return series $\{r_{p,t}\}$ on the portfolio,*

$$\mathrm{SR}_p = \frac{\bar{r}_p \cdot 252}{\hat{\sigma}_p \cdot \sqrt{252}} = \frac{\bar{r}_p}{\hat{\sigma}_p} \cdot \sqrt{252},$$

where $\bar{r}_p$ is the daily mean and $\hat{\sigma}_p$ is the daily standard deviation. The risk-free rate is set to zero, consistent with the absolute Sharpe convention used by MSCI in the Factor Index methodology. Both the portfolio Sharpe of 0.828 and the benchmark Sharpe of 0.649 reported in Section 3.7 use this convention.

Definition 6 (CAPM regression). *For portfolio daily simple returns $\{r_{p,t}\}$ against the QQQ daily simple returns $\{r_{m,t}\}$,*

$$r_{p,t} - r_{f,t} = \alpha_{CAPM} + \beta_{CAPM} \cdot (r_{m,t} - r_{f,t}) + \varepsilon_t,$$

estimated by OLS on the joint 4.75-year sample with $r_{f,t} = 0$. The reported +4.52% annualized alpha is $\alpha_{CAPM,daily} \cdot 252$.

Definition 7 (Information ratio and tracking error). *The active-return series is $a_t = r_{p,t} - r_{m,t}$. The annualized tracking error is $\mathrm{TE} = \hat{\sigma}_a \cdot \sqrt{252}$ and the information ratio is $\mathrm{IR} = \bar{a} \cdot 252/(\hat{\sigma}_a \cdot \sqrt{252}) = (\bar{a}/\hat{\sigma}_a) \cdot \sqrt{252}$. The reported figures of $\mathrm{IR} = +0.356$ and $\mathrm{TE} = 7.54\%$ are inside the 5 to 10% band that the MSCI Factor Index methodology treats as the operational target window.*

Definition 8 (Maximum drawdown and downside capture). *Let W_t be the NAV path. The drawdown at t is $\mathrm{DD}_t = W_t/\max_{s \le t} W_s - 1$, and the maximum drawdown is $\mathrm{MaxDD} = \min_t \mathrm{DD}_t$. The downside capture ratio is the ratio of the portfolio's mean daily return to the benchmark's mean daily return on the subset of days where the benchmark's daily return is strictly negative, $\mathrm{DC} = \bar{r}_p\big|_{r_m<0} / \bar{r}_m\big|_{r_m<0}$. A value below 100% means the portfolio falls less than the benchmark on down days. The reported portfolio downside capture of 99.52% and maximum drawdown of −25.77% versus −35.12% for QQQ are reported in Section 3.7.*

C.5 Implementation correspondence

The procedures of this paper are implemented in the `smart_beta` Python package, organized so each appendix subsection maps to a single module or pair of modules.

The point-in-time universe replay of Appendix A.3 is in `smart_beta/universe.py` (function `ndx_members_asof`). The quality composite of Appendix A.4 is in `smart_beta/quality_score.py`, with the within-sector z-score helper, the safety sub-composite, and the composite formula factored into three functions sharing the module-level `WEIGHTS` dictionary as the single source of truth. The Sharpe single-index regression of Appendix B.1 is in `smart_beta/regression.py`, returning a `OLSResult` dataclass that carries the daily and annualized alpha, beta, residual standard deviation, R^2, and observation count. The selection score and top-K rule of Appendix B.2 are in `smart_beta/selection.py`; the cap-redistribution algorithm of Appendix B.3 is in `smart_beta/portfolio_v04.py` (function `build_ndx_qrv_portfolio`). The daily NAV simulation of Appendix C.1 is in `smart_beta/backtest.py`, and the dual-benchmark reconstruction of Appendix C.2 is wired through `smart_beta/ndx_qrv_analyzer.py`.

The empirics reported in this paper correspond to the `smart_beta` package at git tag `v0.4.1` on the commit identified by the merge SHA `7cb024e`. The package ships with one hundred and twenty-one passing unit tests. The Stage 2 backtest harness used to produce the headline numbers of Section 3.7 is `smart_beta/scripts/run_phase1_mvp.py`.

References

Robert D. Arnott, Jason C. Hsu, and Philip Moore. Fundamental indexation. *Financial Analysts Journal*, 61(2):83–99, 2005.

Robert D. Arnott, Noah Beck, and Vitali Kalesnik. How can “smart beta” go horribly wrong? Technical report, Research Affiliates, February 2016.

Robert D. Arnott, Noah Beck, and Vitali Kalesnik. Forecasting factor and smart beta returns. Technical report, Research Affiliates, 2017.

Clifford S. Asness, Andrea Frazzini, and Lasse H. Pedersen. Quality minus junk. Working paper, New York University (NYU); AQR Capital Management, LLC, October 2013.

Jennifer Bender and Frank Nielsen. Earnings quality revisited. *Journal of Portfolio Management*, 39(4):69–79, 2013.

Jennifer Bender, Remy Briand, Dimitris Melas, and Raman Aylur Subramanian. Foundations of factor investing. Technical report, MSCI Research Insight, December 2013.

Patricia Dechow, Weili Ge, and Catherine Schrand. Understanding earnings quality: A review of the proxies, their determinants and their consequences. *Journal of Accounting and Economics*, 50(2–3):344–401, 2010.

Andrea Frazzini, David Kabiller, and Lasse H. Pedersen. Buffett’s alpha. Working paper, AQR Capital Management; New York University, 2012.

Jason Hsu, Vitali Kalesnik, and Feifei Li. An investor’s guide to smart beta and factor investing. Technical report, Research Affiliates, 2012.

Shiyang Huang, Yang Song, and Hong Xiang. The smart beta mirage. *Journal of Financial and Quantitative Analysis*, 2024. Forthcoming. Earlier draft circulated as SSRN working paper.

Michael Hunstad and Jordan Dekhayser. Evaluating the efficiency of “smart beta” indexes. Technical report, Northern Trust Asset Management, 2015.

Michael C. Jensen. The performance of mutual funds in the period 1945–1964. *Journal of Finance*, 23(2):389–416, 1968.

Nasdaq, Inc. NASDAQ smart beta index family methodology. Methodology document, Nasdaq, 2024. Covers the NQNDSmartBeta and AlphaDEX index families.

William F. Sharpe. A simplified model for portfolio analysis. *Management Science*, 9(2):277–293, 1963.

Richard G. Sloan. Do stock prices fully reflect information in accruals and cash flows about future earnings? *The Accounting Review*, 71(3):289–315, 1996.