

One Theme, Four Portfolios

From monopoly-themed ETFs to the stocks behind their returns

Hesham Mashhour | Monopoly MOAT (mplymoat.com)

Correspondence: hello@hesham.us | ORCID: 0009-0008-6985-4289

Working paper 0.3 | 8 September 2026 | Returns through 4 September 2026

Abstract

Monopoly power, economic moats, durable quality: four exchange-traded funds promise some version of the same idea. Do they deliver the same portfolio? We compare MPLY, MOAT, TOLL and the now-closed CZAR, first through their documented strategies and observed returns, then through a simulation of the stocks they recently held. The common live history is short, so we run a separate counterfactual from January 2016 to September 2026, asking how those holdings would have behaved under buy-and-hold and under monthly rebalancing. The results are less about the theme than about a few individual winners. Under buy-and-hold, Nvidia alone accounts for 71.5% of MPLY's total gain and 61.4% of MOAT's. One stock. Rebalancing monthly cuts those direct shares to 21.5% and 11.2%, and replacing each basket's largest contributor with cash still leaves annualized returns of 24.40%, 15.66%, 23.04% and 13.32%, respectively. The rankings themselves move with the starting date. None of this establishes a causal monopoly premium, and a replay of recent holdings cannot reconstruct what managers actually did. What it does show is more useful: one theme, four very different portfolios, with returns decided by the stocks behind the label.

AI assistance: The author directed the research design and interpretation and rewrote the manuscript. OpenAI Codex assisted with data retrieval, code, numerical calculations, consistency checks and earlier drafts. The author reviewed the work and is responsible for its contents.

1. One label, several investment decisions

A fund can own businesses with strong competitive advantages and still behave very differently from another fund sold on the same broad theme. It might favor the largest technology platforms. It might buy established businesses only when they look undervalued. It might hunt for dominant positions across industrial and infrastructure markets. Each of those choices is defensible. Each produces a different portfolio: different stocks, different weights, different rules about how long a winner is allowed to run.

So we work through a sequence of increasingly specific questions. We start with four U.S.-listed ETFs and the rules they say they follow. We then compare what investors in those funds actually earned over the period the funds shared. That history is too short to tell us how their recent portfolios would have behaved over a decade, so we construct a separate holdings-based backtest. Its gains raise a sharper question: how much of the result came from one extraordinary stock? Monthly rebalancing and contributor-exclusion tests then show what changes when that stock cannot dominate unchecked.

The distinction between these questions matters, and we will keep returning to it. An ETF's observed return is an investor experience. A backtest of its recent holdings is a hypothetical experiment selected with hindsight. Both are informative. They cannot be stitched into one track record. The four funds were chosen for this comparison; they are not an exhaustive sample of monopoly-themed investing.

2. Four interpretations of competitive advantage

All four funds say they want durable business positions. What they do with that idea is another matter. Table 1 summarizes the documented rules, rather than inferring a strategy from its ticker. U.S. listing is the common feature; it does not mean that every underlying company or stock listing is American.^{1,2,3,4}

Table 1. Strategy differences

ETF	Selection emphasis	Portfolio construction
MPLY	Active selection for monopolistic attributes and financial strength	U.S.-listed stocks and ADRs; predominantly capitalization weighted with discretion
MOAT	Morningstar wide-moat assessments plus attractive estimated fair value	Rules-based U.S. equity index; valuation matters as well as business quality
TOLL	Durable quality in concentrated markets, with fundamental valuation	Active domestic/foreign portfolio; private exposure in the recent snapshot
CZAR	Sales, profitability, stability, efficiency and balance-sheet screens	Developed markets; sector weights with equal component weights within sectors

MPLY is the active bet on monopoly itself. Its process considers attributes such as network effects, intellectual property, pricing power and barriers to entry, and its predominantly capitalization-weighted construction, subject to manager discretion, can hand very large companies substantial influence over the result. Price-controlled regulated utilities are excluded.¹

MOAT adds an explicit valuation discipline. A Morningstar wide-moat assessment is not sufficient by itself: the index also seeks attractive prices relative to estimated fair value. That single extra rule matters. A portfolio of attractively priced moat businesses can look very different from a portfolio dominated by the largest perceived monopolies.²

TOLL seeks durable competitive positions in concentrated markets using active fundamental analysis. It was renamed from Tema Monopolies and Oligopolies ETF to Tema Durable Quality ETF in June 2025, a change worth recording but not over-reading: a new name is not, by itself, proof of new investment behavior. Its later private Kalshi exposure adds a further distinction, since not all of its assets have observable exchange-traded price histories.^{3,5}

CZAR tracked the Solactive Natural Monopoly Index, screening for sales, profitability, stability, efficiency and balance-sheet strength. Its sector allocation followed the eligible universe, with equal weighting within sectors. It was therefore neither MPLY's capitalization-oriented selection nor MOAT's wide-moat valuation process.⁴

These differences are reasons the funds might perform differently; they do not establish which rule caused any performance gap. Stock selection, weight, valuation and geography all change together across the four funds. We can describe their portfolios and test specified constructions. What this comparison cannot do is isolate a single causal mechanism.

3. What happened while the ETFs traded together?

MPLY, the youngest fund, sets the beginning of the common observed history. We compare closing investments from 16 May 2025 through 2 July 2026, the last U.S. trading session before CZAR's 6 July liquidation announcement. That is about thirteen and a half months of actual ETF history. Not a decade. Cash distributions are reinvested at the ex-date close, and fund operating expenses are already embedded in these returns.^{6,7}

Table 2. Observed ETF market-price returns over the common window

ETF / reference	Total return	Weekly-based volatility	Weekly max drawdown
MPLY	25.96%	15.97%	-12.55%
MOAT	17.83%	14.73%	-12.08%
TOLL	20.23%	14.20%	-11.26%
CZAR	5.02%	11.46%	-8.33%
SPY	27.10%	12.16%	-8.39%
ACWI	28.79%	12.40%	-8.86%

*Source: cached market prices and cash events. Volatility and drawdown use shared weekly observations; weekly sampling can miss larger intraweek declines. These are market-price reconstructions, not validated NAV-return comparisons.*⁷

MPLY leads the thematic group over these dates, followed by TOLL, MOAT and CZAR. And here is the part a marketing page would skip: all four finish below both SPY and ACWI. Whatever else this comparison shows, it does not support the simple claim that monopoly-themed ETFs outperformed the broader market during their shared live history.

The ordering is also sensitive to where you stop the clock. Extending the comparison through CZAR's final exchange close on 28 July puts MOAT ahead of TOLL, while MPLY remains first among the four and CZAR last. That extension includes the announced wind-down, when CZAR was no longer necessarily pursuing its normal objective, so its final exchange price is not substituted for an unverified liquidation payment.⁶

CZAR stays in the study precisely because closure is an investor-relevant outcome. It should not disappear from a comparison of surviving funds, and closure should not be treated as a total loss either. Five daily CZAR closes are missing, although the entry, exit and distribution-date closes are observed. Its sparse trading and pending NAV validation also qualify its apparently low quoted-price volatility.

So the short window establishes one thing: the four products delivered different experiences. It cannot tell us how their recent stock portfolios would have behaved through earlier markets. For that, we need a different experiment, and a different interpretation of what the result means.

4. Extending the question, not the ETFs' track records

Here is the setup. We freeze the full recent holdings and weights of each fund, then allocate an initial \$10,000 to those weights at the January 4, 2016 close. The endpoint is September 4, 2026, a 10.67-year interval. MPLY and TOLL use September 4 holdings files; MOAT uses September 3. Because CZAR is closed, we use its complete June 30 SEC portfolio, the latest full pre-wind-down snapshot recovered, not a claim to have its literal final invested holdings.^{8,9,10,11}

This answers a narrow question: **How did these selected assets perform when held under the specified rules?** It does not answer what the managers would have selected in 2016. The simulation deliberately conditions on later information, including the survival and success of today's holdings. Everything that follows should be read with that in mind.

Table 3. How the starting allocations are represented

Basket	Priced at start	Cash until later history	Private as cash	Other cash
MPLY	93.37%	6.34%	0.00%	0.29%
MOAT	88.27%	11.51%	0.00%	0.22%
TOLL	86.51%	5.16%	7.55%	0.78%
CZAR	90.92%	8.12%	0.00%	0.96%

All figures are percentages of starting capital. Disclosed and residual cash are retained. The categories reconcile before display rounding; unavailable weights are not removed and the remaining stocks are not rescaled.

Across the four snapshots, 236 unique listed securities have usable histories, and 19 of them begin after the initial date. Their reserved allocations sit in zero-return USD cash until the first usable close, then buy at that close. We do not assume access to an IPO offer price, backfill future prices, or equate the beginning of a vendor history automatically with the legal IPO date.

TOLL's private Kalshi position, 7.55% of its starting allocation, stays in zero-return cash throughout. This is a placeholder, not a measured Kalshi return or a proven lower bound. TOLL's result is consequently a public-holdings-plus-cash scenario rather than a complete valuation of every actual holding.¹⁰

Listed-security returns use provider-adjusted prices, with historical currency conversion into USD where necessary. Ordinary market holidays use preceding local closes, subject to stale-price limits. Dividends are treated as reinvested within the security. The baseline excludes ETF wrapper fees, trading costs and taxes. These assumptions are common across the baskets, but they do not make the underlying assets equally liquid or equally easy to trade.

5. The first backtest: buy the basket and let it run

The simplest construction buys each security at the frozen starting weight and lets its weight drift. A stock that compounds much faster than its peers becomes an increasingly large part of the portfolio, and nothing trims it back. Table 4 shows what that rule produces before trading costs and hypothetical wrapper fees.

Table 4. Buy-and-hold current-holdings backtest, 2016-2026

Basket / reference	CAGR	\$10,000 becomes	Max drawdown
MPLY	42.96%	\$452,401	-50.24%
MOAT	27.36%	\$131,888	-36.75%
TOLL	27.92%	\$138,226	-35.96%
CZAR	16.48%	\$50,900	-36.59%
SPY	15.25%	\$45,447	-33.72%
ACWI	12.76%	\$36,003	-33.53%

CAGR is the compounded annual growth rate, not an arithmetic average. SPY and ACWI provide same-date fund benchmarks with their expenses embedded. Holdings-basket drawdowns use daily valuations, unlike the weekly live-ETF estimates in Table 2.




Figure 1. Growth of \$10,000 under buy-and-hold

MPLY's result is striking: \$10,000 becomes \$452,401, a 42.96% annualized return. TOLL and MOAT finish much closer to each other, while CZAR is nearer the benchmarks. But the wealth curves alone leave the most important question unanswered. Is this a broad result across the basket... or is one exceptional holding carrying much of the outcome?

The risk figures sharpen that question. MPLY's greatest observed decline from a prior peak is -50.24%. The highest terminal wealth did not come with the smoothest path. To understand both the gain and the exposure, we need to look inside the portfolios.

6. Looking inside the gains

For a buy-and-hold basket, each security's contribution is its starting weight multiplied by its cumulative total return. Add those contributions and you have the portfolio's total gain; divide the largest security contribution by that gain and you have its share. This is an accounting decomposition, not a measure of excess return relative to a benchmark.

Table 5. The largest buy-and-hold gain contributor in each basket

Basket	Largest contributor	Starting weight	Ending weight	Share of gain
MPLY	NVDA	10.86%	70.1%	71.5%
MOAT	NVDA	2.57%	56.9%	61.4%
TOLL	LRCX	5.37%	17.7%	18.7%
CZAR	AVGO	4.42%	28.7%	34.6%

Gain share refers to total dollar profit above the original investment, not the percentage of ending wealth. Ending weight is shown separately. The largest contributor is identified after observing the full-period results.

Nvidia accounts for 71.5% of MPLY's buy-and-hold gain. Its initial 10.86% allocation grows to 70.1% of ending wealth. In practical terms, the experiment starts as a 100-stock basket and finishes with roughly seven dollars in ten attached to one security. Whatever diversification the original snapshot had, it has not been maintained.

The same mechanism drives MOAT. Nvidia begins at just 2.57%, provides 61.4% of the gain, and ends at 56.9% of portfolio value. A small starting position can eventually dominate a long holding period. So MOAT's high buy-and-hold return cannot be read as evidence that its many holdings contributed equally.

TOLL and CZAR are less dependent on their single largest contributors under this measure. Lam Research provides 18.7% of TOLL's gain; Broadcom provides 34.6% of CZAR's. Those exposures still matter, but the distribution of gains is different from the Nvidia-heavy MPLY and MOAT outcomes.

None of this makes the arithmetic wrong. Letting winners grow is a legitimate portfolio rule. It changes the question, however: the backtest no longer maintains the portfolio weights that initially defined each basket. To understand the selected holdings while containing that drift, the next step is to restore those weights regularly.

7. What changes when we rebalance monthly?

At each U.S. month-end close, we reset every allocation to its original frozen target. Winners are trimmed; other holdings and cash reserves are replenished. This is an illustrative construction rule, not a claim about the funds' actual trading. It was added to the main presentation after inspecting concentration, and that post-analysis choice is disclosed rather than presented as preregistered.

Table 6. Returns and dominant-asset gain shares after monthly rebalancing

Basket	Buy-and-hold CAGR	Monthly CAGR	Gain share: buy/hold	Gain share: monthly
MPLY	42.96%	32.42%	71.5%	21.5%
MOAT	27.36%	17.46%	61.4%	11.2%
TOLL	27.92%	25.64%	18.7%	10.5%
CZAR	16.48%	15.13%	34.6%	13.0%

The dominant security is the one identified in Table 5; it is also the largest direct contributor in the monthly case for each basket. Direct monthly gain shares sum each security's daily dollar profit or loss while held.




Figure 2. How rebalancing changes the dominant asset's direct share of gains

Nvidia's direct share of MPLY's gain falls from 71.5% to 21.5%. Its ending weight falls from 70.1% to 11.3%, near the 10.86% target. Within-month drift is still possible. MPLY's CAGR falls to 32.42%, while its maximum drawdown moderates from -50.24% to -34.66%.

The gain-share decline needs careful reading. Suppose Nvidia generates a profit, that profit is moved into another stock, and the second stock subsequently rises. The later profit is credited to the second stock. Direct attribution follows where gains are earned; it does not trace every dollar back to its original source. A smaller direct share therefore does not mean that removing Nvidia would reduce final wealth by only that percentage.

Monthly rebalancing substantially addresses the concentration problem. It does not remove the contribution of the original winner, or the hindsight in selecting it. To examine that dependence directly, we need to run the portfolio without that security at all.

8. What CAGR remains without the largest contributor?

We now replace each basket's largest full-period contributor with zero-interest cash from the beginning. Every other target weight is unchanged, the remaining stocks are not increased to fill the gap, and the excluded security is never purchased. We then rerun the monthly-rebalanced portfolio over exactly the same dates.

Table 7. Monthly-rebalanced portfolios with the largest contributor replaced by cash

Basket	Removed	Cash replacement	Baseline CAGR	Remaining CAGR	\$10,000 becomes
MPLY	NVDA	10.86%	32.42%	24.40%	\$102,644
MOAT	NVDA	2.57%	17.46%	15.66%	\$47,203
TOLL	LRCX	5.37%	25.64%	23.04%	\$91,334
CZAR	AVGO	4.42%	15.13%	13.32%	\$37,965

The cash replacement equals the excluded security's original target weight. "Remaining CAGR" is recalculated from the new portfolio's wealth path. It is not obtained by subtracting a gain share from CAGR, and it is not a causal estimate of the security's influence on the wider market.

MPLY's annualized return falls from 32.42% to 24.40% without Nvidia. The new ending value is \$102,644 rather than \$199,834. That is a 48.6% reduction in terminal wealth, even though Nvidia's direct gain share under monthly rebalancing was 21.5%. The gap between those two numbers is compounding and reinvestment at work. Nvidia remains economically substantial.

At the same time, the remaining MPLY basket does not collapse into a negligible-return result. It continues to contain other successful companies selected with hindsight. That supports a narrow conclusion: the specified basket's historical outcome was not wholly a single-stock phenomenon. It does not establish that those other winners could have been identified in advance.

MOAT's remaining CAGR is 15.66% without Nvidia, much closer to SPY's 15.25% over the same dates. TOLL retains 23.04% without Lam Research, while CZAR retains 13.32% without Broadcom. Removing an equal number of securities is not removing an equal amount of capital: the four original weights differ, as Table 7 makes explicit.

These tests also separate accounting contribution from robustness. Attribution tells us where the recorded profits were earned. Exclusion tells us what happens under a different, fully specified portfolio. Neither creates an out-of-sample result: the dominant assets were selected for removal only after their returns were known.

9. Does the conclusion survive different dates and costs?

The large differences could still be an accident of the 2016 starting point. So we start fresh portfolios at four additional dates, always using the same frozen weights and September 4, 2026 endpoint. These are new initial allocations, not slices of the wealth already accumulated by the 2016 portfolio.

Table 8. Monthly-rebalanced CAGR by starting date

Start	MPLY	MOAT	TOLL	CZAR	SPY	ACWI
2016-01-04	32.42%	17.46%	25.64%	15.13%	15.25%	12.76%
2018-01-02	30.58%	15.84%	25.51%	13.54%	14.62%	11.65%
2020-01-02	32.55%	13.84%	26.41%	13.81%	15.40%	12.98%
2022-01-03	23.99%	6.93%	24.67%	11.16%	12.23%	11.23%
2024-01-02	36.44%	10.42%	36.95%	15.69%	21.39%	21.16%

MPLY leads the thematic baskets from 2016, 2018 and 2020. TOLL leads from 2022 and 2024. MOAT and CZAR also change order. There is no date-invariant winner in these results, although MPLY and TOLL remain the two highest-returning thematic baskets across these starts.

We also charge 0, 10 or 25 basis points, meaning hundredths of a percent, on every dollar of listed securities bought or sold, including the initial purchase and purchases when later histories begin. Cash transfers are free, orders at a shared listing or rebalance close are netted, and costs are paid from the portfolio rather than injected from outside.

Table 9. Full-window monthly CAGR under illustrative trading costs

Cost per dollar traded	MPLY	MOAT	TOLL	CZAR
0 bps	32.42%	17.46%	25.64%	15.13%
10 bps	32.32%	17.38%	25.55%	15.06%
25 bps	32.18%	17.25%	25.41%	14.96%

The full-window ranking does not change under these cost assumptions. They are sensitivities, not estimates of actual security-level spreads or market impact. Taxes, wrapper fees, terminal liquidation and distribution-reinvestment execution costs are excluded, and no additional synthetic trading cost is charged to the benchmarks.

The five date windows overlap and share the same hindsight-selected assets; they are not five independent confirmations. Likewise, changing a cost parameter is not new evidence that the holdings could have been selected prospectively. These checks make the descriptive comparison more transparent. They do not remove its central selection limitation.

10. What this journey establishes

We began with four funds sharing one broad investment theme. Their documented strategies differed in valuation discipline, security selection, geography and weighting. Their short common live history then showed different investor outcomes, without thematic outperformance of both benchmarks. The longer holdings experiment revealed much stronger hypothetical returns, but also how heavily a buy-and-hold result can depend on an individual winner.

Rebalancing changed both the magnitude and the interpretation of those returns. It stopped Nvidia from becoming approximately 70% of MPLY's ending portfolio, and it spread accumulated profits across the frozen basket. Excluding the largest contributor then showed that substantial historical returns remained in MPLY and TOLL, and smaller returns in MOAT and CZAR. Changing the starting date altered the ranking rather than crowning a universal winner.

So the conclusion is about portfolio implementation. A shared monopoly theme does not specify a single return profile. The selected securities, their starting weights, whether winners are allowed to grow, and the period examined all matter. What the study does not establish is a causal monopoly premium, persistent manager skill or a forecast of future performance. One theme, it turns out, can be four very different portfolios.

Limits that remain material

The simulations select later holdings with hindsight and retain only those portfolios' chosen assets. Their high historical returns must not be confused with returns achievable from information available at the start. Snapshot dates differ, CZAR uses a pre-closure filing, and TOLL's private Kalshi exposure remains unmeasured. Zero-return cash is an explicit completion assumption, not proof that the missing asset would have performed worse.

Provider-adjusted histories are also not fully independent economic reconstructions. Targeted checks corroborate six ordinary split dates and ratios and two dividend amounts; ten close-plus-cash reconstructions agree with adjusted terminal wealth within 0.153%. Same-provider agreement does not validate every corporate action. GE illustrates the issue: an adjustment associated with the GE Vernova spin-off is not the same as separately tracking and retaining the distributed shares.^{12,13}

For actual ETF returns, five CZAR closes remain missing, its cash distribution still relies on secondary-source evidence, and complete NAV histories and the final liquidation payment are unverified. The closed fund is retained without inventing missing prices or proceeds. These limitations qualify publication claims, not merely the precision of a table.

The contribution is a compact, reproducible comparative study, and its strongest result is the sequence itself: similar thematic labels conceal different portfolios; impressive basket returns can be concentrated in a single stock; rebalancing changes that concentration; and removing the largest contributor leaves a measurable, but still hindsight-conditioned, remainder. Detailed assumptions, all scenarios and source receipts are retained with the supporting research files.¹⁴

Funding and interests

This research was self-funded by the author and received no external funding. The author reports no investments in MPLY, MOAT, TOLL, CZAR or the individual stocks discussed, and no financial or professional relationships with the ETF issuers or managers. The author's affiliation is Monopoly MOAT (mplymoat.com).

References

- [1] Strategy Shares. (2025). Monopoly ETF prospectus. Investment selection, weighting and exclusions. [SEC prospectus](#)
- [2] VanEck. (2026 snapshot). Morningstar Wide Moat ETF: fund overview and investment approach. [Issuer overview](#)
- [3] Tema ETF Trust. (2026). Tema Durable Quality ETF amended summary prospectus, January 30. [SEC prospectus](#)
- [4] Themes ETF Trust. (2023). Themes Natural Monopoly ETF summary prospectus. [SEC prospectus](#)
- [5] Tema. (2026 snapshot). TOLL fund page and FAQ, including name history and private Kalshi exposure. [Issuer fund page](#)
- [6] Themes ETF Trust. (2026). Supplement dated July 6 announcing CZAR and other fund closures. Expected timing is not evidence of final cash proceeds. [SEC closure supplement](#)
- [7] Market-price and cash-event sources. Yahoo Finance historical chart feeds; ChartExchange CZAR closes; Stock Analysis/S&P Global CZAR distributions. Dated responses retained in the research cache. [Yahoo historical prices](#); [CZAR closing prices](#); [CZAR cash distributions](#)
- [8] Strategy Shares. (2026). Full MPLY holdings CSV, effective September 4; portfolio values use September 3 closing prices. [Issuer holdings CSV](#)
- [9] VanEck. (2026). Full MOAT holdings dataset, as of September 3. [Issuer holdings dataset](#)
- [10] Tema. (2026). Full TOLL holdings CSV, as of September 4, including Kalshi SPV and cash. [Dated issuer holdings CSV](#)
- [11] Themes ETF Trust. (2026). Form N-PORT, Themes Natural Monopoly ETF, reporting date June 30. Accession 0000894189-26-022852; series S000081024. [SEC portfolio filing](#)
- [12] Issuer and SEC corporate-action records used for the targeted audit. Six ordinary split date/ratio checks and two dividend-amount spot checks; not a full independent audit of all histories. [Nvidia 2021 split completion](#); [Nvidia 2024 split and dividend](#); [Broadcom 2024 split and dividend](#); [Lam Research split and dividend units](#); [Fortinet 2022 SEC filing](#); [KLA 2026 split](#)
- [13] GE Aerospace. (2024). Board approves the GE Vernova spin-off, February 29. Documents the distribution of one GE Vernova share per four GE shares. [Issuer announcement](#)
- [14] Mashhour, H. (2026). Reproducibility materials for One Theme, Four Portfolios. Supplied research bundle: frozen holdings, daily paths, direct dollar-gain attribution, eight exclusion paths, 60 monthly start/cost scenarios, source receipts and technical reports. Offline calculation: `make etf-study`. PDF rendering: `make etf-paper`. The return inputs are frozen; mutable issuer pages are not substitutes for the dated cache.