



ANNUAL UPDATE

2025

2025 Results

During 2025, the Saga Portfolio (“the Portfolio”) increased 44.2% net of fees compared to the overall increase for the S&P 500 Index, including dividends, of 17.9%.

Since inception on January 1, 2017, the cumulative return for the Saga Portfolio is 417.9% net of fees compared to the S&P 500 Index of 255.7%. The annualized return since inception for the Saga Portfolio is 20.1% net of fees compared to the S&P 500’s 15.1%. Please check your individual statement as specific account returns may vary depending on the timing of any contributions throughout the period.

Performance			
	Saga Portfolio (gross)	Saga Portfolio (net) ⁽¹⁾	S&P 500 Index ⁽²⁾
2017	16.0%	14.3%	21.8%
2018	2.1%	0.6%	-4.4%
2019	65.6%	63.2%	31.5%
2020	123.8%	120.5%	18.4%
2021	-9.6%	-10.9%	28.7%
2022	-84.7%	-84.9%	-18.1%
2023	209.2%	204.6%	26.3%
H1'24	40.4%	39.4%	15.3%
H2'24	53.4%	52.2%	8.4%
2024	115.4%	112.2%	25.0%
Cumulative	651.3%	551.4%	272.2%
Annualized	23.6%	21.8%	14.8%

(1) Saga Portfolio serves as a model for client accounts. Net returns assume 1.5% management fee, or 0.375% applied at beginning of each quarter.

(2) S&P 500 performance includes dividends. Please see disclaimer at the end of this letter regarding comparison to indices.

Why Share Prices Tell Us Very Little

People often describe the stock market as a casino, with prices bouncing around for no obvious reason. It can feel that way if one focuses on what a stock will do next and consumes financial media designed to make every day sound like a crisis that demands immediate action. But at its core, the stock market is a system for price discovery and capital allocation. It allows ownership in businesses to change hands and helps fund companies that are solving real problems. When it works well, it supports innovation, productivity, and rising living standards.

Every day, the market sets prices for thousands of companies. Those prices reflect millions of competing explanations and expectations about business value, compressed into a single number. But that number is not a precise measurement of what a business is worth. A share price is better understood as the market’s current conjecture about value, expressed in dollars. Like any conjecture, it is incomplete, fallible, and often wrong.

Stock prices change as people update their explanations and act on them. Sometimes those updates reflect real changes in long-term earning power. Often, they are driven by shifting narratives even when the underlying business has changed little. When uncertainty is high, small changes in interpretation can move prices significantly. That is why much of what passes as analysis does little to improve our understanding of how

businesses actually create value. For that reason, I treat short-term performance as a description of market behavior, not a reliable scorecard for investment quality. Saga is structured to think and act over multi-year periods, without the constraints that push most investors toward short-term decision making. Much of the market is focused on predicting how prices will move over the next few weeks or months, rather than what returns might be if one owned the business for decades.

Many investor letters explain which holdings “contributed to” or “detracted from” results and then retrofit explanations for those moves. That can sound analytical, but it often implies a level of short-term predictability that does not exist. The narratives are usually constructed after the fact, making outcomes feel more knowable than they were at the time decisions were made.

If we cannot reliably predict stock prices, how should we decide which businesses to own?

Explanations, Not Predictions

The answer is to base decisions on explanations of businesses, not on short-term predictions about stock prices.

If stock ownership is taken seriously as partial ownership of a real business, the right question becomes: how would this company be evaluated if it had to be owned for the next ten or twenty years with no ability to sell? This constraint shifts attention away from forecasting near-term earnings or share price movements and toward understanding what the business can become over time.

Thinking in decades is a humbling constraint. Very few businesses endure, and far fewer thrive. Many advantages that look durable in the short run turn out to be temporary. Knowledge is hard to create, but once discovered it tends to spread.

That is why excess returns are rare. Competitive markets act as an error-correcting system. When a company earns unusually high profits on capital, it often signals a better way of solving a problem or a constraint others have not yet overcome. Eventually, those profits attract attention, capital, and imitation. Competitors learn, customers demand more, and alternatives emerge. Returns compress unless the company keeps improving faster than competing solutions.

Sustained excess returns usually require something more durable: capabilities others cannot replicate, coordination problems that remain hard even after a winner sets the standard, or constraints that are difficult to copy, such as trust, distribution, regulation, or deep integration into customer workflows. Even these protections are not permanent in a world where invention and competition persist. The most durable advantage is often not a static moat, but the ability to keep creating and applying new knowledge to better serve customers.

Given those realities, very few companies meet the standard required for long-term ownership. The key is having a good explanation for how a business creates value over time. What matters is not precision, but whether the explanation is hard to vary.

Easy-to-vary explanations are vague, non causal, and hard to test. They are often only loosely connected to how a business creates value, which allows them to be adjusted to fit almost any outcome and to explain winners only after they have already won. High return on invested capital is not an explanation. It is the result that needs to be explained. More commonly, easy-to-vary explanations lean on outside conditions like rates will fall or the macro environment will improve, or rely on labels like network effects or brand, without specifying the underlying

mechanism or what evidence would refute the thesis.

A hard-to-vary explanation is the opposite. It is a causal account of the mechanisms that generate value, constrained by how the world actually works. Changing key elements should break the logic of the thesis. It should answer a few essential questions: What problem does the company solve? Why do customers choose it? Why is that difficult to replicate? How can the business continue to improve? What evidence would refute the thesis?

The explanation comes first. Forecasts can follow, but they remain fallible because the future can contain genuinely new and unpredictable knowledge.

Once there is a working explanation, valuation becomes practical. How large could the business become? What must be true for growth to continue? Does the advantage strengthen or weaken with scale? What could break the model? In practice, risks like reliance on a single gatekeeper, regulatory exposure, leverage, or loss of trust often matter more than small changes in discount rates.

The next step is to compare that explanation to the one implied by the current share price. The price reflects the market's prevailing conjecture. A spreadsheet cannot predict the future, but it can make assumptions explicit. If the market's implied explanation matches the explanation in hand, there is little reason to expect excess stock returns. If it does not, and the explanation in hand is better, that is where opportunity exists.

Returning to the thought experiment of owning a stock for twenty years helps filter out noise. What share prices will do along the way is unknowable, so portfolio management becomes the discipline of enduring volatility without being forced into poor decisions. That is why Saga avoids margin, options strategies, and short selling. The goal is not certainty. It is resilience: building a portfolio grounded in sound explanations, with position sizes that allow for error without catastrophic loss.

Why Beating the Market is Hard

With that framework in mind, the next question is why outperformance remains so rare in practice.

The [S&P SPIVA After-Tax Scorecard](#) illustrates how difficult this is. Over the ten years through 2024, 95.5% of actively managed "All Domestic" U.S. equity mutual funds underperformed their benchmarks after taxes. Over twenty years, 97.7% underperformed.

The idea that there is a simple way to beat the market is self-limiting. Any method that is both simple and repeatable would be copied, competed away, and ultimately reflected in prices. That is why the most obvious bargains that existed decades ago, such as companies trading at large discounts to conservatively appraised assets or special situations with a clear catalyst, appear far less frequently today. It is also why pattern investing and rules of thumb tend to disappoint. Past correlations can look convincing, but they rarely explain what is actually driving value or why the market is misjudging it.

The hard reality is that long-term stock returns are lopsided. Over any decade, most stocks deliver ordinary results and many disappoint, while a small minority drives a disproportionate share of wealth creation. That skew is exactly what you would expect in a world where progress depends on new knowledge. Creating something genuinely better is hard, and competition erodes anything that is easy to copy. In mature categories, differentiation typically fades and excess profits drift toward average. Many public companies that are considered "high quality"

are simply executing well understood playbooks in mature markets. Even a great business can be a mediocre investment if the share price already assumes a great outcome.

The rare long-term winners tend to emerge when a company finds a genuinely better way to solve an important problem and then scales that advantage to many more customers. When the reason for success is deep and causal rather than a temporary tailwind, advantages can compound. Learning curves improve the product, scale lowers unit costs, reputation and distribution expand its reach, and switching costs or network effects increase customer stickiness. These forces are not linear. Once a company gets ahead, the same mechanism that created the lead can make it easier to widen. Early on, that trajectory is hard to value because the evidence is noisy and the end state is uncertain. That is why meaningful mispricings can persist for long periods.

This is the central challenge in active investing. It is not enough to find good companies. To outperform, you need two things. First, a resilient mechanism for why the business can keep solving problems better than the alternatives as it learns and adapts. Second, a share price that implies a meaningfully weaker future than the one the mechanism suggests. Even when you have both, the path is rarely smooth. Markets continually revise their guesses, and share prices can swing wildly for reasons that may have little to do with long-term fundamentals.

Institutional portfolio management makes beating the market that much harder. Most professionals are judged on short horizons, so strategies are selected to look safe quarter to quarter and stay within the comfort zones of clients and committees, not to be right over a decade. That pressure pushes portfolios toward wide diversification and consensus holdings. Add benchmark constraints, tracking error limits, position size rules, and investor flows that tend to arrive at the worst times, and it becomes extremely difficult to own the rare winners at meaningful size and hold through enough volatility to capture outperformance.

A system optimized for short-term optics tends to produce closet-index portfolios. It is no surprise that long-term underperformance after fees is common in professional investment management and that passive index funds have attracted so much capital.

Saga Case Studies: Carvana and GoodRx

A philosophy only matters if it improves decision making in the face of conflicting information, especially when narratives shift and share prices fall. The key distinction is between volatility that creates opportunity and new information that breaks a thesis.

Carvana

The used car market is enormous, but its structure has historically limited the scale of any single player. Dealerships are built around local inventory, local reconditioning, local financing, and local delivery, so selection is constrained and unit costs stay high.

Carvana is easy to misunderstand because it is not simply selling cars on a website. It is an integrated system designed to break those constraints by centralizing inventory, industrializing inspection and reconditioning in purpose-built facilities, building a dense logistics network, and embedding financing into the same flow.

The mechanism is cumulative learning and operating leverage. Each part of the system improves the others, which is why the advantage compounds rather than scales linearly. Higher throughput lowers delivery and reconditioning cost per unit. Broader selection improves conversion. Faster turns reduce capital intensity. More

transactions improve pricing, logistics, and credit decisions. Over time, this creates a system that is hard to copy because it requires years of operational learning, national infrastructure, funding relationships, and enough volume for the machine to improve through repetition and cover high upfront fixed costs.

When we invested in 2019, the market's expectations embedded in the stock price were far weaker than what the mechanism suggested. Carvana was still reporting losses, so the easy story was a cash-burning retailer that would eventually hit a wall. The question I kept returning to was simple: is the mechanism beginning to show up where it should? I looked for improving unit economics in older cohorts, a strengthening customer proposition as selection and delivery reliability improved, evidence of operating leverage as utilization rose, and a balance sheet that could fund the company until it reached scale and profitability. Over time, the evidence increasingly pointed to an integrated system whose economics should improve with density. It also pointed to the difficulty competitors would have replicating the full stack cost structure and customer experience once an operator had already moved down the learning curve.

Then 2022 arrived. Used car affordability deteriorated, interest rates rose sharply, volumes fell, and execution issues hit at the wrong time. The share price collapsed. The key question was not whether prior expectations were wrong, but whether the causal explanation had been refuted.

The thesis would break if demand structurally shrank, if unit economics failed to improve as utilization recovered, or if competitors matched the experience and economics without comparable infrastructure. I did not see evidence of that. The mechanism looked intact, and the company had enough liquidity to resize costs and fund the business through the downturn. We held and added opportunistically at depressed prices. As management improved efficiency and restored profitability while maintaining customer service, fundamentals improved and shares recovered from their lows.

GoodRx

GoodRx is a useful counterexample. I originally bought it because the value proposition looked durable as a pharmaceutical marketplace. It aggregated discount card pricing across multiple pharmacy benefit managers (PBMs), helping cash pay consumers access negotiated prices. It also appeared to have a defensible advantage through scale and a patented interface that enabled contracting with multiple PBMs in one place. The flywheel was intuitive. Better prices attract more consumers, which pushes PBMs to compete harder for that demand.

One of the key risks was the pharmacy point of sale. Pharmacies were already squeezed on insured prescriptions and historically relied on higher margin cash pay transactions. Discount cards pushed those prices down and shifted bargaining power away from pharmacies. In the first half of 2022, I learned a thesis breaking fact. Kroger stopped accepting GoodRx coupons at the counter.

Management previously claimed that pharmacies were effectively required to honor these discounts through PBM contracts. In reality, they could opt out and impose a price floor. This was not a temporary headwind. It changed the industry structure by making GoodRx's advantage contingent on pharmacy participation in a way I had underestimated. I sold our shares.

To date, I consider GoodRx one of the biggest mistakes of commission in the portfolio, but the mistake was the purchase, not the sale. Selling at a loss was the right decision because the thesis was broken and the original explanation no longer held. It was not about the stock price.


The contrast between holding Carvana through extreme volatility and selling GoodRx is the point of the framework. A falling stock price is not a reason to sell. A broken causal explanation is.

Reflections on Nine Years Managing the Saga Portfolio

Since inception at the beginning of 2017, the Saga Portfolio has compounded at 20% per year net of fees, 5% ahead of the S&P 500. To put those results in context, I screened Morningstar's database of U.S.-focused mutual funds. Among roughly 1,100 funds with a nine-year track record, the Saga Portfolio would rank near the top over the same period.

The chart below shows the distribution of annualized returns for those U.S.-focused mutual funds. The key point is how skewed the outcomes are. A small minority outperform the S&P 500, and those that do typically outperform by only a few percentage points. Most funds underperform.

One important point is survivorship, which likely understates how widespread underperformance really is. The distribution reflects only funds that still exist and have long track records. Many others were merged or liquidated over time, biasing the results upward.



Source: Saga Partners, Morningstar

It also helps to understand how the portfolio has evolved. In 2016, I left my job to manage a portfolio full-time with a simple goal: invest my own capital in the way that made the most sense to me, write openly about the process, and partner with investors who shared that philosophy. In 2017, Michael Nowacki and I launched the Saga SMID Cap Portfolio as co-portfolio managers, focused on small and mid-cap stocks where we expected to find the most overlooked opportunities.

At the beginning of 2019, Michael decided to focus on providing wealth management services, and I became the sole manager of the Saga Portfolio. Since then, I have managed the portfolio more opportunistically, trimming or exiting positions that became less compelling and concentrating more heavily in the opportunities I think were best, regardless of market cap. That shift increased concentration and reduced turnover. It also increased short-term volatility, but I expected it would improve the portfolio's long-term compounding potential.

Since 2019, the Saga Portfolio has compounded at 24% per year, net of fees, ranking 4th out of approximately 1,500 U.S.-focused mutual funds in Morningstar's database with a seven-year track record.

Portfolio	7 Yr CAGR (2019-2025)
1 Baron Partners Fund	30.8%
2 Baron Focused Growth Fund	26.0%
3 Alger Focus Equity Fund Class I	24.3%
4 Saga Portfolio	24.0%
5 Baron Opportunity Fund	23.2%
6 Virtus Silvant Focused Growth Fund	22.3%
7 Fidelity Advisor Growth Opportunities Fund	22.2%
8 Federated Hermes MDT Large Cap Growth Fund	22.1%
9 Alger Capital Appreciation Portfolio	22.0%
10 Voya Russell Large Cap Growth Index Portfolio	21.9%

Source: Saga Partners, Morningstar

Unlike firms that launch multiple strategies and later close underperformers, Saga has run a single portfolio with a single approach. The record includes the full distribution of outcomes, including mistakes.

It also includes the 2022 drawdown, for which I take full responsibility. That period was a reminder that markets can move to extremes when prices are marked daily and narratives shift faster than fundamentals. Some of our companies faced real challenges, but share prices can overshoot in both directions. The worst outcome during periods of heightened skepticism is being forced to sell and lock in losses that later prove temporary. The key is having the flexibility to hold through pessimism and, when appropriate, act opportunistically when prices fall well below any reasonable estimate of long-term value.

If the Saga Portfolio has an edge that has contributed to its long-term results, it is not better forecasting or superior resources. It is structural. Many professional investors operate under constraints that encourage decisions that are easier to explain, protect careers, and require continuous activity. Those constraints make meaningful outperformance difficult.

The Saga Portfolio is managed with the sole intention of compounding capital over long horizons. A small number of exceptional businesses drive most long-term returns, but their share prices can be volatile. My decisions are guided by my understanding of those businesses and a willingness to look wrong for extended periods when the underlying explanation remains intact.

In practice, that means I spend most of my time reading, observing, and thinking. Action is infrequent. I only act when I think it will improve the portfolio's long-term returns, not to appear active or prudent.

Portfolio Review

Activity in 2025 was consistent with the portfolio's underlying philosophy. Changes were not driven by near-term sentiment or a need to do something. I invest only in businesses where I have a clear understanding of how they create value over time and where that understanding differs meaningfully from the market's current pricing.

As a result, the portfolio consists of a small number of businesses whose share prices can be volatile, especially when sentiment shifts quickly, even as the underlying businesses continue to compound value. The key question is not the magnitude of that volatility, but whether the underlying explanations have changed and how share prices compare to them.

During 2025, I did not initiate any new positions. The most notable change was increasing the position in The Trade Desk as its share price declined and the market's implied expectations became increasingly pessimistic relative to the company's underlying competitive position and long-term opportunity. I include an overview in the appendix, along with a more detailed write-up published separately for those interested in a deeper analysis of the digital advertising ecosystem ([link](#)).

I also increased the position in Wise, where the underlying business continues to strengthen and the long-term opportunity remains substantial, while the market's view remains well below my estimate of its long-term value.

The capital to fund these changes came primarily from Carvana. Carvana remains a large position, and I continue to think it is undervalued. However, as the share price appreciated toward the end of the year, the gap between price and long-term value narrowed, while the opportunity in The Trade Desk and Wise became more attractive, leading me to reallocate capital accordingly.

More detailed updates on each holding are included in the appendix.

Conclusion

If you follow markets closely, you will never run out of reasons to worry. There is always a new alarming headline and a new prediction that "this time is different." Even into the first quarter of 2026, attention has shifted to new concerns, with share prices reflecting those concerns. Some risks are real and can affect the value of specific businesses. More broadly, pessimism tends to dominate because the future is inherently uncertain and there are an infinite number of things that can go wrong.

Yet none of that prevents progress. It comes from businesses solving real problems for their customers. Markets do not always reflect that progress in real time. Sentiment shifts constantly, often driven more by perception and narrative than by changes in underlying fundamentals. These shifts are not abstract. They shape how people feel, how they act, and ultimately how capital is allocated.

As investors, we operate within that system. Investing on behalf of others adds another layer of complexity. Every decision is visible. Outcomes are judged continuously. The prospect of loss carries a different weight when it affects more than just your own capital. After nine years managing the Saga Portfolio, I have felt that pressure, especially during 2022. Periods like that test not just the portfolio, but the framework behind it.

What has kept me focused on Saga's core philosophy is a simple principle. I manage the portfolio with the same rigor I would apply to my own capital, because it is. In practice, that means owning fewer businesses, accepting

more volatility, and at times holding positions that may appear uncomfortable or controversial—not because I seek that out, but because the best opportunities rarely feel obvious or easy in the moment.

Even with that principle, managing other people’s capital has, at times, constrained my willingness to concentrate further when opportunities appear especially attractive. When the possibility of being wrong affects others, it inevitably shapes decision-making. In hindsight, outcomes often appear clear. In real time, they rarely are.

The experience of 2022 made that reality more concrete. Long-term investing is not just about identifying good businesses, but owning businesses that can endure difficult periods and emerge stronger. Periods of declining stock prices are when the underlying explanation is most tested and criticized. Rarely does a falling share price come with reassuring news. It often feels uncomfortable, and at times it should.

Yet those periods are also when expectations reset and opportunities are often most attractive. For long-term owners, lower prices are not a problem if the thesis remains intact, but an opportunity to allocate capital more effectively. That is why I would rather see our companies make strong fundamental progress while their share prices remain depressed, even if it makes the Saga Portfolio’s results appear weaker in the interim.

My responsibility is not to avoid volatility. It is to ensure the portfolio can endure it and, when appropriate, benefit from it.

When I assess the portfolio today, I am highly encouraged by its prospects. The businesses we own generated record operating profits in 2025 and are expected to do so again in 2026. They continue to strengthen their positions, while valuation multiples across much of the portfolio remain near historical lows. I do not know what share prices will do over the next year, but the gap between price and value appears compelling over the long term.

I remain fully invested alongside you and am grateful for your trust as we continue to compound value over time. Please feel free to reach out if you have any questions or would like to connect.

Sincerely,
Joe Frankenfield

Appendix


Carvana

The used car market has historically been fragmented and locally constrained, with high costs, limited selection, and a cumbersome customer experience. Each dealership operates its own inventory, reconditioning, and logistics, which limits scale and keeps unit economics inefficient.

Carvana approaches this differently by building an integrated system that centralizes inventory, industrializes inspection and reconditioning, and coordinates logistics and financing into a single platform. This expands selection, reduces friction, and lowers costs while improving the customer experience.

As the system scales, each additional unit improves efficiency. Higher throughput lowers cost per unit, increases inventory density, improves conversion, and accelerates inventory turns. This creates operating leverage and embeds learning into the system, making it increasingly difficult to replicate.

In 2025, Carvana grew retail units by 43% while demonstrating strong profitability, generating EBITDA of \$2.2 billion and an EBITDA margin of 11%, even as industrywide used car volumes were roughly flat.



Source: Saga Partners, Company filings

Carvana has shown it can grow while expanding profitability, in part because it spent the last several years building infrastructure for volumes well above its current run rate. As throughput rises, the fixed cost base is better utilized and margins expand. 2025 provided further evidence that the system is working as intended. The company also continued to reduce debt and strengthen the balance sheet, ending the year with net debt to EBITDA of 1.2x and committed liquidity of \$4.4 billion.

The key question is how much of the roughly 40 million unit U.S. used car market can migrate toward a model like Carvana's. E-commerce penetration varies significantly by category. Total online retail penetration in the U.S. is nearly 20%, but more complex and logistically intensive categories generally see lower adoption, and used cars fall into that group.

However, the traditional offline used car buying experience is far more cumbersome than in most retail categories, suggesting the relative improvement from a model like Carvana's may be greater than in categories that have already shifted online. While used cars are more difficult to move online than categories like books or electronics, Carvana has demonstrated that it can buy, sell, recondition, and transport vehicles at scale with attractive

economics, resulting in a meaningfully better experience across convenience, selection, and transparency. As a result, a larger share of used car sales could migrate toward this model.

Given the difficulty of replicating its infrastructure and the absence of a competitor with comparable scale, brand, and capabilities, Carvana is well positioned to capture a significant share of that shift. Building a competing system would require years of infrastructure investment, operational learning, and sufficient volume to support the economics.

Despite the share price recovery following Carvana's operational progress, I continue to view the shares as attractive relative to their future earnings power, which does not require aggressive assumptions. Carvana already has the capacity to sell 1.5 million retail units annually and expects to increase that to roughly 3 million with modest incremental capital investment within its existing real estate footprint. At 3 million units and approximately \$4,000 of EBITDA per unit, that would imply roughly \$12 billion of EBITDA.

Demand should exceed 3 million units over time. While growth beyond that level will require additional real estate investment, the business should be able to fund that expansion as it scales.

If the system continues to develop as expected, Carvana should improve both its cost structure and customer experience, reinforcing its position in the market. The current share price does not appear to reflect the company's ability to approach 3 million units in the intermediate term, let alone its potential beyond that.

Roku

As television shifts from linear to streaming, the most valuable position may not be owning content, but controlling how viewers discover and consume it. In a fragmented ecosystem with dozens of apps and an expanding range of choices, the interface that organizes attention becomes increasingly important.

Roku sits at that layer. As the operating system for connected TV, Roku helps viewers navigate content across apps while helping content providers reach audiences. That position gives Roku influence over discovery, engagement, and monetization across the broader streaming ecosystem.


Its advantage comes from sitting between viewers and content suppliers. Through search, recommendations, home screen placement, and navigation, Roku can shape what gets watched without needing to own the content itself. As more viewing flows through the platform, Roku gathers more data, improves the user experience, and increases its ability to monetize attention over time. In that sense, Roku benefits less from picking winners in streaming than from being the platform through which viewing decisions are made.

That position did not emerge by accident. Roku's focus on simplicity, low-cost hardware, and broad content access helped drive early adoption. Its distribution strategy, particularly through embedded TV operating systems, then allowed it to scale rapidly across households. Unlike competitors with conflicting priorities across hardware, content, and advertising, Roku remained focused on the operating system layer, which helped it establish a durable position in the U.S. market.

Today, Roku reaches roughly half of U.S. broadband households and accounts for nearly half of all TV streaming hours in the U.S. That scale matters because it makes Roku increasingly important to both sides of the ecosystem:

viewers want easy access to content, and content providers want access to viewers.

In 2025, Roku made progress on the two things that matter most for the thesis: engagement and earnings power. Streaming hours increased 15% to 146 billion, while Platform gross profit grew 15% to \$2.1 billion. Roku also reached positive net income and record trailing twelve-month free cash flow, making the earnings power of the model more visible.



Source: Saga Partners, Company filings

The value of controlling the connected TV operating system remains underappreciated. One example is The Roku Channel, Roku's owned-and-operated channel that aggregates ad-supported and premium content. In December 2025, it accounted for more than 6% of total U.S. streaming time.

More important than the number itself is what it suggests about Roku's position. The Roku Channel became one of the most engaged apps on the platform, ranking second in the U.S. by engagement beginning in early 2025, ahead of Netflix. That occurred despite significantly lower content spend, suggesting that control over distribution and discovery can be as important as content ownership itself.

The opportunity becomes even clearer in the context of television advertising. U.S. TV advertising remains a roughly \$90 billion market, with about \$30 billion already shifted to connected TV and roughly \$60 billion still tied to linear television. Yet streaming now accounts for nearly half of viewing. As that mismatch closes, advertising dollars should continue moving toward connected TV. Streaming also enables more precise targeting and measurement at the household level, allowing connected TV to take share not only from linear television, but also from portions of digital video advertising.

Roku is well positioned to benefit from that shift. It accounts for a large share of streaming hours and typically receives 15–30% of advertising inventory on its platform. Combined with its role in subscription distribution, that points to earnings power materially above what the current share price appears to reflect.

The key question is not whether Roku becomes something fundamentally different. It is whether Roku continues strengthening a position that becomes more valuable as viewing and advertising migrate toward streaming. If it does, the company's long-term earnings power is likely higher than the market currently assumes.

The Trade Desk

Advertising has become a decisioning problem. Instead of choosing among a small number of media channels, advertisers now allocate budgets across millions of opportunities each second across websites, apps, streaming platforms, and devices. In that environment, the valuable position is not simply owning media, but helping advertisers decide where spending should go.

The Trade Desk was built to fill that role. As an independent demand-side platform, it helps advertisers evaluate and allocate budgets across the open internet. Unlike Google’s demand side platform DV360 or Amazon’s DSP, it does not own media or represent sellers. Its role is to act on behalf of the buyer.

That distinction matters because most digital advertising still flows through large closed platforms such as Google, Meta, and Amazon. Those platforms have grown not necessarily because they offer better opportunities, but because they offer simpler systems. By controlling inventory, data, and measurement, they reduce friction and make campaigns easier to execute. The open internet, by contrast, offers broader reach and more independence, but is harder to navigate.

The Trade Desk’s role is to make that market more usable and efficient. Through initiatives such as OpenPath and broader supply-chain optimization, it is working to reduce unnecessary intermediaries, improve data consistency, and make pricing more transparent. If that happens, more of each advertising dollar should reach working media, improving the economics of the open internet over time.

Recent results raised questions about whether this is still playing out. In 2025, gross spend increased 11% and revenue increased 15%, below historical levels, while closed platforms grew faster. But those results likely overstate how much has actually changed. Over a longer period, however, the business has continued to scale meaningfully, with gross spend and revenue compounding substantially over the past decade.



Source: Saga Partners, Company filings

The Trade Desk’s customer base is concentrated in large global advertisers, particularly in categories such as consumer packaged goods and automotive, which faced unusual pressure. Results were also affected by political advertising comparisons and internal execution changes, including the rollout of Kokai. Taken together, the slowdown appears more consistent with cyclical and execution-related factors than with a structural break in the thesis.

The deeper question is whether independent decisioning still matters. Platforms such as Google’s DV360 and

Amazon's DSP both own inventory and influence how budgets are allocated, creating a clear economic incentive to direct spending toward their own properties. The Trade Desk does not have that conflict. Just as important, outside the walled gardens there are very few scaled independent alternatives. The remaining demand-side platforms generally lack the data, infrastructure, and cross-channel capabilities needed to compete effectively for large global advertisers. In practice, the set of realistic options is narrower than it first appears.

The opportunity remains significant. Global digital advertising is approaching \$800 billion, with roughly \$200–250 billion already allocated across the open internet, compared to The Trade Desk's approximately \$13 billion of gross spend in 2025. Even if the open internet does not improve much from here, it is already a large market that needs independent, scaled decisioning, and The Trade Desk is well positioned to continue taking share within it.

The upside may be larger still. The open internet remains under-monetized relative to time spent, with consumers spending meaningfully more time there than inside the major closed platforms. If the supply chain becomes more efficient and friction declines, more advertising dollars should follow. In that scenario, The Trade Desk would benefit not only from share gains within the existing market, but also from the expansion of the market itself.

At the current share price, the market appears to be assuming a more fragile business and a more limited opportunity than I think is likely. The key point is that the thesis does not require a particularly optimistic outcome. There is already a large market that needs independent, scaled decisioning, and The Trade Desk is uniquely well positioned to serve it. If the open internet becomes more efficient, the opportunity expands further. That is why the risk-reward appears better than the market currently assumes.

Given the amount of recent skepticism and negative headlines surrounding The Trade Desk, I also wrote a separate, more detailed piece that explores these issues more fully ([link](#)).

Trupanion

Pet insurance is simple in concept, but often poorly aligned in practice. Many policies rely on reimbursement, include coverage limits, and reprice over time, particularly as a pet becomes more expensive to insure. That creates uncertainty for pet owners and incentives for insurers to optimize for short-term profitability rather than long-term value.


Trupanion was built differently. Rather than treating insurance as a series of one-year contracts, it prices policies based on the expected lifetime cost of care. Coverage does not reset by condition, and claims are paid directly at the point of care through integration with veterinary clinics. When a pet owner visits the vet, Trupanion pays its portion of the bill immediately, reducing both financial friction and uncertainty at the moment it matters most.

This structure changes the economics of the model. Instead of optimizing individual policy years, success depends on accurately pricing for the lifetime cost of care and retaining pets over long periods. As the insured base grows, Trupanion gathers more data on claims and outcomes, improving underwriting and strengthening lifetime value. The system can improve with scale, not just grow.

In 2025, the business was still working through the effects of elevated veterinary cost inflation in 2022 and 2023. Higher costs required price increases to realign premiums with underlying claims costs, while new pet acquisition

was deliberately slowed until unit economics stabilized.

That tradeoff weighed on near-term growth, but improved profitability. New pet additions increased only 5% in 2025, while adjusted operating income increased 33% to \$152 million. Subscription segment adjusted operating margin reached a record 16.5% in the fourth quarter.



Source: Saga Partners, Company filings

That dynamic has contributed to investor skepticism and the decline in the share price. Slower pet additions can look like weakening demand, even as the underlying economics improve. But Trupanion now appears to be moving through that transition. As margins and lifetime value per pet have improved, the company has begun to reinvest more aggressively in acquisition.

The opportunity remains significant. Pet insurance penetration in the United States is still low at roughly 4%, compared to closer to 25% in more mature markets such as the United Kingdom. While market structure differs across countries, the gap suggests the U.S. market is still in the early stages of adoption.

At current prices, the market appears to be extrapolating slower pet growth and a smaller long-term opportunity. That view underweights the improvement in unit economics and the resumption of investment in growth. If the system continues to develop as expected, Trupanion should be able to grow both pet count and earnings power over time. The current share price does not appear to reflect the long-term earnings power of the business.

Wise

Cross-border payments remain one of the largest and most inefficient areas of the global financial system. Even today, money typically moves through a network of correspondent banks, passing through multiple intermediaries before reaching its destination. Each step adds cost, delay, and opacity, with fees often embedded in exchange rates rather than clearly disclosed. Despite advances elsewhere, sending money internationally can still take days and cost several percentage points of the transaction value.

Wise was built to change that. Instead of routing payments through intermediaries, Wise connects directly into local payment systems and matches flows internally. When a customer sends money across borders, funds are paid in locally on one side and paid out locally on the other, with Wise managing settlement across its own network. The result is a system with fewer steps, faster delivery, and meaningfully lower cost.

The advantage is not tied to any single feature. It comes from how the system is constructed. Over time, Wise has built a set of interlocking capabilities, including regulatory licenses, direct integrations into payment systems, and automated operations, that together create a system optimized for speed, transparency, and cost. Any one piece may be easy to copy, but the full system is much harder to replicate.

At its core, cross-border payments are a commodity, where advantage accrues to the lowest-cost provider. Wise is structured around that reality. As volume grows, more flows can be matched internally, reducing reliance on external settlement and lowering marginal cost. Rather than expanding margins, Wise passes those efficiency gains back to customers through lower prices. Lower prices attract more volume, which further improves efficiency, creating a self-reinforcing cycle.


Over time, that cycle compounds. As infrastructure improves, costs continue to decline, allowing Wise to keep lowering prices while maintaining its position as the lowest-cost provider. This creates increasing pressure on competitors, particularly those operating on legacy systems or relying on higher fees. Matching Wise on price ultimately requires matching its cost structure, which in turn requires replicating years of infrastructure, regulatory work, and operational complexity across multiple markets.

As the business grows, it is also evolving. What started as a transfer product is increasingly becoming a broader international account. Customers are using Wise to hold, manage, and spend money across currencies, with balances now exceeding £25 billion. That deepens engagement and makes Wise more embedded in how customers actually operate financially.

At the same time, Wise is extending its infrastructure to other institutions through Wise Platform. Rather than competing with banks across every dimension, it is increasingly enabling them by allowing banks and fintechs to plug into its network to offer faster and lower-cost cross-border services. As Wise's cost advantage grows, the opportunity cost of not using its infrastructure rises, shifting the dynamic from competition toward integration. Over time, this positions Wise as a potential routing layer for global cross-border payments.

The opportunity is large. Global cross-border payment flows exceed £30 trillion annually, with roughly £170–190 billion generated through fees and foreign exchange spreads. Wise today processes only a small fraction of that, with low single-digit share in consumer transfers and very limited penetration in business payments.

Recent results are consistent with how the system is designed to work. Cross-border volumes are expected to exceed £180 billion in fiscal year 2026, growing roughly 25% year over year, with underlying gross profit expected to increase 26% to £1.3 billion. Wise generates operating margins of approximately 15% and strong free cash flow. Customer adoption remains strong, and balances continue to grow, driven largely by word of mouth and increased usage of the Wise account.



Source: Company filings, Saga Partners

Note: Wise's fiscal year ends March 31. 2026E reflects consensus estimates, including Q4 FY2026, to present full-year figures consistently.

Despite this, the market remains focused on declining take rates, rising competition, and the contribution of interest income. While understandable, these concerns largely reflect how the model is designed to function rather than evidence of deterioration. Lower pricing is a direct result of efficiency gains being passed back to customers, which drives volume and reinforces the system. Competition is real, but matching Wise requires replicating its cost structure across infrastructure, regulation, and operations, which is neither quick nor easy. While interest income has contributed to recent results, the underlying economics are driven by volume growth and improving efficiency rather than pricing power.

One question that comes up is whether new technologies, such as stablecoins, could change the landscape. While these systems may improve certain parts of the process, they primarily affect how value moves between currencies, not the local integrations required to fund and deliver payments. Much of Wise's advantage sits in those connections, along with its pricing model and operational infrastructure. If new rails become more efficient, they are more likely to be incorporated into Wise's system than to replace it.

The key question is how much of the global cross-border payments market can shift toward a system like Wise over time. A large portion still runs through infrastructure that is slower, more expensive, and less transparent. If Wise continues to expand its network and maintain its cost advantage, more of that volume should migrate to its platform.

At the current share price, the market appears to be discounting a more competitive and lower-return future. That view underweights a system whose economics improve as it grows and whose importance increases as more participants rely on it. If that continues, Wise has the potential to become a more central part of how money moves globally.

DISCLOSURES & DISCLAIMERS

This document is provided for informational purposes only and should not be relied upon as the basis for any investment decision. Past performance is not indicative of future results, and there can be no assurance that any investment objective will be achieved or that any investor will avoid losses.

Any investment decision should be based on an investor's own due diligence and review of all applicable offering, advisory, custodial, and account documentation, as well as consultation with the investor's own legal, tax, regulatory, accounting, and investment advisers.

The Saga Portfolio is a concentrated, long-only public equity strategy managed by Saga Partners, LLC. The strategy involves substantial risk, including the risk of loss. Because the portfolio is concentrated in a limited number of securities, its performance may be more volatile than that of a diversified portfolio. Portfolio holdings and allocations are subject to change without notice.

The information contained herein reflects the views of Saga Partners, LLC as of the date indicated and is subject to change without notice. Certain statements contained in this document may be forward-looking in nature. Forward-looking statements are based on current expectations, assumptions, and judgments and are subject to a wide range of risks and uncertainties. Actual results may differ materially from those expressed or implied in such statements.

Unless otherwise indicated, performance information presented herein is unaudited. Gross returns are presented before the deduction of management fees. Net returns are shown after applying a 1.5% annual advisory fee, reflected as a 0.375% fee applied at the beginning of each calendar quarter. Net returns shown are intended to provide a general illustration of the effect of advisory fees and may differ from the actual performance experienced by any individual account.

The Saga Portfolio is presented as a model portfolio. References to portfolio holdings, allocations, and performance are for illustrative purposes and are based on the Saga Portfolio model and/or a representative account implementing the strategy. Actual client accounts may differ materially in size, timing of investment, cash flows, custodial arrangements, fees, tax circumstances, and portfolio implementation. As a result, the performance of any individual client account may differ materially from the performance shown herein.

Performance figures assume the reinvestment of dividends and other distributions where applicable. Performance results may reflect periods of unusual market volatility or favorable or unfavorable market conditions and should not be viewed as indicative of the results that may be achieved in the future.

Any references to specific securities are provided solely to illustrate the application of the strategy and do not constitute a recommendation to buy, sell, or hold any security. There is no assurance that any securities discussed will remain in the portfolio or that securities purchased in the future will be profitable.

Comparisons to the S&P 500 Index are provided for general informational purposes only. The S&P 500 Index is an unmanaged index of widely held common stocks and is not available for direct investment. Index returns do not reflect management fees, transaction costs, or other expenses. The holdings, characteristics, and risk profile of the Saga Portfolio differ materially from those of the S&P 500 Index.

This material does not constitute an offer to sell or a solicitation of an offer to buy any security or investment advisory service in any jurisdiction where such offer, solicitation, or sale would be unlawful. Any offer or advisory relationship may be made only pursuant to definitive legal documentation and in accordance with applicable law.