

Letter to *Partners*

Q 1 · 2 0 2 6

April 7, 2026

To the Partners of the Umatiya Investment Group Funds,

For the quarter ending March 31, 2026, the Fund returned 4.21% net of expenses, compared to - 4.33% the S&P 500. The result should not be misconstrued or extrapolated to any measure. As I mentioned in our prior letters, we do not intend to emphasize short-term performance; however, given the early stage of the Partnership and in the interest of transparency, we believe it is appropriate to include brief updates, for the time being. Our investment approach remains oriented toward long-term compounding, and quarterly results are not indicative of underlying performance.

On Capital Stewardship & the Philosophy of Long-Term Compounding

In 1626, the Dutch purchased Manhattan from the Lenape people for about \$24 worth of trade goods. It is one of history's most cited "bad deals." But run that \$24 forward at a 7% annual return, a rate any disciplined long-term investor might reasonably expect, and by today it compounds to \$24 trillion. The lesson is not that the Lenape got a fair price. The lesson is that time is the most underappreciated variable in investing, and most people never give it the chance to work.

The historical record is clear on this point. A few weeks ago, I read Larry Fink's annual letter to investors. He made an interesting point that needs reiterating: families who stayed invested through depression, war, and financial crisis didn't win because they predicted the recoveries, they won because they didn't sell during them. In other words, holding for the long term is a bet on the economy within which you reside and ties you to the secular growth of your countries. When savings are invested over long periods, capital is put to work building businesses, infrastructure, and jobs. When this happens domestically, your financial future becomes tied to your country's progress: you help fund its growth, and it, in turn, supports yours. This mutli-decade mindset matters more than most investors realize, and it sits at the center of how we think about capital stewardship.

This discipline extends to how we view cash. I have always wanted to have optionality built into the fund which is why I have intended that the fund hold 10% in cash most of the times (except scenarios where we absolutely cannot afford to stay on the sidelines). Our cash position through the end of this quarter stands at 20% of the fund. Now before this baffles you and I start receiving phone calls it is important you understand the following. The reasoning is not because we anticipate a downturn, and not because we lack ideas. The businesses we would most want to own for the next decade are simply not available at

prices we find reasonable. Cash, in this context, is not a position. It is the result of not willing to compromise. I am reminded this constantly, but most recently via Buffett's latest interview. Buffett was asked whether the recent market pullback had made things look cheaper. His answer was characteristically blunt: a 5 or 6% decline doesn't move the needle when nothing was attractively priced to begin with. He was not predicting a crash nor was he making a macro call. He was simply saying that the opportunities aren't there yet, and that liquidity is what allows you to act decisively when they are. That is precisely our thinking. We are not waiting for a signal. We are waiting for a price. Lastly, as the markets broadly remain expensive, please remember that patience is not the enemy of returns, it is often the source of them.

The Macro Backdrop: A Confluence of Pressures

I want to be clear about the purpose of this section before proceeding. The following is not a prediction but more so an observation. We do not attempt to forecast geopolitical events or time economic cycles, and we have no edge in doing so. What we do care about are the forces that affect the fundamentals of the businesses we own and give us a better understanding regarding them. It is for that reason alone that we examine the macro environment.

With that said, the world has changed significantly since our last letter. The escalation of the conflict involving Iran has introduced a supply shock of historic proportions, one whose consequences extend well beyond oil prices and will take time to fully appear. We are not here to alarm. We are here to think carefully about what comes next.

Macro events matter to us only to the extent that they shape economic reality on the ground. Evaluating the businesses we own requires understanding who their customer is, what pressures that customer is facing, and what the broader environment looks like for the businesses we do not yet own but would like to.

Right now, the consumer is being squeezed from an angle worth acknowledging. The ongoing energy disruption has taken over 3% of global LNG supply offline for the next three to five years following damage to Ras Laffan, Qatar's largest LNG facility. According to the IEA this is the largest supply shock in oil market history. In a market that typically operates with a 2–3% spare capacity buffer, this is a material issue. Any future disruption, however minor, now lands on a system with no excess LNG, which means the resilience of global energy supply that markets have grown so used to may be diminished for the foreseeable future.

Elevated energy costs are not absorbed equally. Lower-income households spend nearly twice the share of their budget on gas and electricity as higher-income ones. When that bill rises, something else gets cut. It is almost always discretionary spending. The second-order effects compound this further. Natural gas feeds fertilizer production, which feeds food prices, which flows through logistics and downstream through the cost of nearly everything. The consumer in this environment is not just paying more at the pump, they are gradually paying more for most things, while their actual purchasing power quietly erodes. If spending grows at 2% but prices are rising faster, the consumer is effectively consuming less even as they spend more. That gap does not show up in any single month, but it accumulates and eventually shows up in the revenues of businesses that depend on that consumer feeling financially comfortable enough to spend beyond their basic needs.

To stretch this out further, for illustrative purposes, are the third-degree effects from this fallout. Qatar produces roughly one-third of the world's helium as a byproduct of LNG processing. With Ras Laffan offline, that supply is effectively gone for the foreseeable future. Helium is not a commodity most investors think about. But it is irreplaceable in semiconductor manufacturing and medical imaging where there is no meaningful substitute, and stockpiles are finite. The bottleneck this creates has nothing to do with oil prices and everything to do with the domino effect of concentrated supply chains. It is a useful reminder that the consequences of a single disruption rarely stay within the boundaries of the industry where they start. They travel indirectly, and often in directions that markets are slowest to price.

This is exactly why we have always been selective about consumer exposure. Less than 15% of the fund sits in businesses dependent on discretionary consumer spending, not because we are making a macro call, but because these businesses are cyclical by nature and require a higher burden of proof. When we do own them, we want a structural reason why they hold up when the consumer is stretched. Our most notable holding in this category is CZR. Regional, drive-to casino destinations have historically proven resilient in exactly this kind of environment as they capture the person who skips an expensive vacation but still seeks an experience close to home. That resilience is a feature of the business model itself and not a bet on the environment it operates in. It is the kind of distinction we care about deeply when researching any consumer-facing business. This is relevant to any business whose customer is the everyday American consumer.

Beyond what we own, we are paying close attention to what we do not own yet. Stocks broadly remain priced for a degree of optimism that the current environment does not obviously support. Geopolitical uncertainty, higher input costs, and a consumer under pressure are not conditions that usually justify

premium valuations, yet here we are. We are not attempting to predict when or how that resolves. But we are aware of where we stand, and that awareness is reflected in our positioning.

AI: Real, Remarkable, and Richly Priced

Howard Marks said it better than most: AI is very real, its potential is more likely underestimated than overstated, and none of that makes AI investments a bargain. I want to start there because it captures exactly the tension every serious investor should be sitting with right now. Most aren't.

Most AI commentary lives in one of two places. The evangelists argue compute curves and trillion-dollar TAMs and conclude the bullish case is self-evident. The skeptics state the dot-com bubble and conclude the whole thing is priced for perfection. Neither camp is doing the harder work, which is thinking about the technology carefully enough to figure out where it actually creates value, where it destroys it, and what any of that is worth at today's prices.

The Wrong Question

Nearly every investor, sell-side analyst, and technology journalist covering AI is organized around the same question: who wins? Which model, which platform, which chip captures the most value as this scales? It is a valid question. It is also the most competed question in markets right now, which means whatever answer you arrive at is most likely already in the price.

We spend very little time on it. The question we find more useful is: when AI compresses the cost of something that used to be expensive to produce, where does the value go?

This is not rhetorical. It has a historical answer, and the answer is not appealing to the businesses commanding the highest multiples today.

The Productivity Surplus Nobody Is Considering

When a productivity technology enters a competitive market, it creates a surplus. The question is always who captures it. The company deploying the technology? The workers using it? The customers buying from it?

History is one-sided here. In competitive markets, productivity gains flow to the customer. The first mover gets a temporary margin advantage. Competition forces that advantage into lower prices. The

margin reverts. Customers win. This happened in manufacturing, software, and logistics. It happens almost every time a productivity technology moves to broad commercialization.

AI is moving fast toward broadly accessible. The models are converging. The APIs are cheap and getting cheaper. Switching costs between providers are low. In that kind of environment, AI stops being a competitive advantage and starts being a cost of doing business, like electricity or cloud storage. The businesses priced as though their AI lead is a durable moat are, in many cases, being priced on an assumption the technology's own trajectory is undermining.

What the Market Is Missing

The more productive analysis, in our view, is not identifying who benefits from AI. That list is long and already priced. It is finding where the market has applied the right thesis to the wrong businesses, and where it has applied the wrong thesis to the right ones.

Start with enterprise software. The SaaS model is built on a per-seat pricing assumption. More employees at a customer means more seats, more seats mean more revenue. It is a business model that is structurally long human headcount. AI is structurally short it. When a financial services firm can perform the analytical work of twenty employees with five, it does not just cut headcount. It cuts seats across every platform those employees were using. That demand destruction does not show up in one quarter. But it is directional and durable, and it is not factored in most models.

Here is where the analysis needs precision, because the market is making a category error.

SaaS is not one business. It is two different businesses sharing a multiple.

The first is software as a tool. Applications that make a task faster, priced per seat, sitting lightly on top of a workflow with no meaningful data accumulation over time. This category is exposed in exactly the way the “SaaS apocalypse” narrative suggests. When AI makes the underlying task cheap enough to automate, the seat evaporates. The SMB end of most SaaS verticals lives here, and the disruption is not coming. This is already happening.

The second is software as infrastructure. These platforms sit between a company's internal systems and its end customers. They manage compliance, security protocols, and integration layers that a large enterprise cannot rebuild in-house without absorbing years of operational risk. More importantly, they accumulate proprietary data with every single interaction: call transcripts, resolution outcomes, customer

sentiment across thousands of engagements over years. That data trains AI to work better over time. It compounds. And it belongs entirely to the platform.

An enterprise that switches does not just change software vendors. It abandons the institutional memory it spent years building. No competitor can hand that back. No AI model trained on generic data can replicate it. The switching cost is not contractual. It is informational, and that is a different kind of moat.

The market is discounting both categories on the same logic because from the outside they look identical: SaaS, per-seat, headcount-dependent. That compression is where the mispricing lives. The first category deserves the pressure it is receiving. The second may be receiving it by association. That distinction is doing real analytical work, and almost nobody is pondering this.

We are not avoiding enterprise software categorically. We are asking, for every business we examine, which of these two things it actually is and whether the current price reflects the right answer.

There is a second category where pressure is building, and it is less visible to markets because the businesses feeling it most directly are private. The consequences, though, are public.

The business model of a law firm, a management consultancy, or a large accounting practice is not just selling expertise. It is a leverage model. Junior professionals do high-volume cognitive work billed at rates that subsidize senior partner time and fund the partnership's margin. AI does not make the junior associate faster. It attacks the economic rationale for having as many of them. The headcount does not collapse from the top; senior judgment, accountability, and client relationships are genuinely hard to replicate. It collapses from the bottom, which is exactly where the leverage lives.

These are largely private businesses, so markets cannot price the disruption directly. But the publicly traded companies that depend on them as customers, the workflow tools, the data platforms, the software built around that labor model, carry the exposure too. When this hierarchy compresses, the tools built to serve every layer of it compress with it. This is something worth paying attention to over time.

What AI Cannot Touch

There is a version of this analysis that slides into dystopia, and I want to avoid it, not because it would be alarmist, but because it would be imprecise.

AI makes it cheaper to produce analysis. What it does not make cheaper is judgment: deciding when the facts are incomplete, taking a position others disagree with, and staying with it when there is real downside. The payoffs to that kind of judgment have always shown up where things are unclear, the consequences matter, and the results take a long time to be known. Better AI models do not change any of that.

As AI makes analysis abundant and cheap, the scarcity of the judgment sitting on top of that analysis actually increases. For a fund like ours, this matters practically. As AI democratizes screening, modeling, and pattern recognition, the informational edge that came from pure analytical work decreases across the industry. The edge migrates toward the things that cannot be replicated: sitting with uncertainty, holding a position against consensus, distinguishing between a business that looks broken and one that is not. That is a more favorable environment for the kind of investing we do, not a less favorable one. We say that not to be self-serving but because we think it is true, and because it shapes how we think about where durable value in the AI economy ultimately concentrates.

Why We Are Not Buying

The question of buying AI largely comes down to price.

These are not cheap businesses. They are extraordinary businesses priced as if extraordinary is now permanent, as if the early leads that made them remarkable are durable features of a fixed competition rather than current positions in a race that has barely started. Buying at today's prices means buying every optimistic assumption built in the valuation at once, with no margin of error on any of them.

What we do not know is still considerable: which architectural approaches dominate at scale, how competitive dynamics shift once model performance plateaus, what regulatory intervention looks like as labor displacement becomes an issue, and whether current capital expenditures by hyperscalers will produce the returns investors are expecting. Current prices reflect the narrow end of that range. The version where everything works. That is not the base case we assign the highest probability to.

We are doing the work now so that when the price reflects the uncertainty that actually exists, we can act with conviction rather than scrambling to catch up. That is the only honest description of where we stand.

Lessons From FOUR

Long-term conviction is only as sound as the work beneath it. This past quarter, I spent much time doing a deep analysis on Shift Four Payments (FOUR). The kind of work that forces you to pressure-test what you think you know about a business. I can say hands down this was the toughest internal conversation I have had with myself regarding any business we have researched.

Shift4 Payments is a vertically integrated payments technology company that processes transactions for restaurants, hospitality, and entertainment venues, bundling point-of-sale software with payment processing under one platform.

We initiated a position in FOUR on the premise that the business was undervalued relative to its potential, specifically, management's guided path to a \$1 billion adjusted free cash flow run rate. After a week of deep research and several dedicated hours stress-testing the thesis, I closed the position. This was a mistake in my initial analysis, and I want to walk you through my reasoning transparently.

The core of the thesis depended on a question I ultimately could not answer with conviction: what percentage of acquired merchants would actually migrate to Shift4's end-to-end platform, and what conversion rate underpins management's \$1 billion FCF target? Without that answer, the entire model is dependent on an assumption I cannot verify from the outside.

This led me to a more uncomfortable realization about the business model itself. Shift4 grows primarily through acquisition rather than organic sales and marketing. That means acquisition spend is not a discretionary line item but more so their customer acquisition cost. Yet accounting conventions and management's own framing treat it as an investing activity, which meaningfully flatters reported earnings and free cash flow. When you adjust for this, the true economics of the business look materially different.

Layering second and third-order thinking on top of this revealed further fragility. The acquisition machine only functions if acquired restaurants stay open long enough to recoup the purchase price, if conversion rates hold, and if debt markets remain accommodating. Shift4 operates in discretionary consumer verticals — restaurants and entertainment — with a leveraged balance sheet. That is a combination that becomes vulnerable in an economic downturn, precisely when revenue softens and interest expense hits the hardest. More acquisitions to sustain growth only compound the risk, not resolve it.

In short: the business is running on a high-speed treadmill financed by debt, and the asymmetric upside I require to take on that complexity and risk was not present.

What I take from this is a reinforcement of a principle I hold firmly: process over outcome. I did not have a view I could defend with conviction on the key variables that drive the investment. Owning that uncertainty and acting accordingly is not a failure; it is the discipline this job demands. As Buffett has long said, you do not have to swing at every pitch. I know I could be very wrong and FOUR could end up being a great business. But regardless of the outcome I would be content knowing that FOUR fell outside my circle of competence, and I would rather acknowledge that honestly than construct a narrative to justify holding.

Closing

Although the Fund has gotten off to a decent start, early results should be interpreted with a grain of salt. We measure our progress over years, not quarters, and encourage our partners to do the same. In the short term, markets often swing between extremes for reasons that have little to do with underlying business value. Over the long run, however, the quality of the businesses we own and the prices we pay for them matter far more.

We feel the same way about our current holdings. Some have appreciated, some have not, and others, we believe, still have a long runway ahead. The work we are doing today is towards a simple goal: to be prepared when great businesses can be purchased at attractive prices. In the end, investing is not much more complicated than that.

Thank you for your continued trust and partnership.

Sincerely,

Samit Umatiya

Disclosures

This document has been prepared by Umatiya Investment Group LLC (“UIG”) and is provided for informational purposes only to existing partners and qualified prospective investors. The information contained herein is confidential and may not be reproduced or distributed without the prior written consent of UIG.

Performance information presented herein is unaudited and reflects returns net of operating expenses, unless otherwise indicated. Such returns do not reflect annual audit expenses or incentive fees, which are calculated and applied at year-end. Accordingly, actual realized returns may differ, potentially materially. Past performance is not indicative of future results.

The S&P 500 is a market capitalization-weighted index consisting of 500 leading publicly traded companies in the United States and is widely regarded as a proxy for the overall U.S. equity market. The index reflects the performance of large-cap equities across multiple sectors. It is unmanaged, does not incur fees or expenses, and is not available for direct investment.

Certain statements contained herein constitute forward-looking statements, including, but not limited to, statements regarding market conditions, portfolio positioning, and investment opportunities. These statements are based on UIG’s current views and assumptions and are subject to change without notice. Actual results may differ materially.

The Fund employs a concentrated investment strategy and may, at times, maintain meaningful cash positions, which may increase volatility and result in performance that deviates from broader market indices. Investments are subject to market, economic, and issuer-specific risks, including the possible loss of principal.

Any market commentary, including references to specific securities, industries, or macroeconomic conditions, is provided for illustrative purposes only and does not constitute investment advice or a recommendation to buy or sell any security, nor should it be relied upon as a forecast of future events.

Certain information contained herein has been obtained from third-party sources believed to be reliable; however, UIG does not represent or warrant the accuracy or completeness of such information and accepts no liability for any errors or omissions.

This document does not constitute an offer to sell or a solicitation of an offer to buy any security. Any such offer will be made only pursuant to the Fund’s definitive offering documents and in accordance with applicable securities laws.

Nothing contained herein constitutes investment, legal, or tax advice. Recipients are solely responsible for evaluating the merits and risks associated with any investment and should consult their own advisors prior to making any investment decision.