Dear Partners:

As long-time clients know, I am not prone to hyperbole, but I believe this letter is the most important we have ever written, and I hope it conveys our genuine excitement. AI has increased our productivity by a factor of ~5 – 10x (no, that's not a typo). This has completely transformed our investment process, already delivering tangible results – yet I believe we are in the very early innings of fully harnessing and utilizing the power of AI. I've attached two examples of recent AI-enabled research – a 50-page memo on Grocery Outlet (GO) and a 10-page update on Fitlife (FTLF) – to demonstrate the depth and quality of work that can be accomplished in a fraction of a time.

AI is a generational shift that enables previously inconceivable scalability of brainpower, insight, and judgment. It automates tedious manual work, freeing human time and energy for far more valuable activities. In my view, this is comparable to the step-change from poring through physical books or newspapers seeking information… to just Googling it. Particularly for small firms like ours that don't have an army of internal analysts and access to expensive paid services or research, AI levels the playing field, allowing us to cover vastly more ground.

We've always viewed our edge as insight and judgment rather than in the gathering of raw information. AI, used correctly, is the equivalent of having a limitless team of interns or junior analysts, allowing us to spend vastly more (and, eventually, perhaps nearly all) of our time on value-added insight and judgment, rather than manual tedium.

Regardless of workflow or investment philosophy, I believe learning to fully leverage AI is one of the most valuable and highest-return activities a public equity investor can undertake right now. Adoption among value investors has been surprisingly slow, perhaps due to the rapid development of the technology (which as I discussed in my prior letter, seems to have evolved from "fun toy" to "powerful tool" within merely quarters).

Indeed, the most effective ways to use AI are often non-obvious, and treating it as "intelligence" in the conventional sense tends to lead to frustration – again, as I discussed last time, AI does not really "think" or "reason" in the way that you or I do, even when it appears to be doing so. Nonetheless, those who restructure their processes to be AI-first will, at least for the next several years, have a substantial competitive advantage over those who do not.

I remain convinced that human judgment and insight are essential, and I am skeptical that AI can replace them. What it can do – and do very, very well – is complement them, drawing on its statistical, big-data approach in ways that enrich the decisions we make. AI is not making our investment choices or substituting for our judgment, nor do we expect it to do so. It is automating busywork and data collection, so I can spend more of my time on the work of a senior analyst or portfolio manager: interpreting the facts and deciding what to do with them.

This shift has already proven its worth. During our proxy contest at AstroNova, AI tools helped us produce the extensive ISS deck and other materials that would have been vastly more costly to prepare otherwise (i.e., we would have needed to spend an incremental six figures, and/or developed far lower quality work-product). We earned endorsements from both ISS and Glass Lewis, and the incumbent board requested the CEO's resignation. I believe, though it is of course impossible to verify, that we may well have run the first AI-powered proxy contest in history.

Although we are still early in implementing this technology, and further advances (such as GPT-5) are arriving in real time, we are mostly through rebuilding all of our core tools we have relied on for a decade to be AI-first – from our research documents to our watchlist. Some of this is still under construction, which is why we ask for your patience for another month or two before we return to full tear-sheets and portfolio commentary.

Our goal is to explain how we're using AI, and how our investment approach is evolving because of it.

Section 1: How We're Using AI On A Daily Basis

There are at least three key vectors along which AI is helping our process, directly or indirectly; of course, they are not mutually exclusive. The full extent of all of our use cases is beyond the scope of this letter, but hopefully this provides a flavor.

There are some areas where I will be a little vague because I believe it is in our, and clients' best interests, to not reveal some secret sauce. I will, of course, have more to say on these topics in the future.

1. Custom Software Development & Automation Workflows

Purpose: eliminate manual friction and reclaim research time for higher-order thinking.

Background: I could not code my way out of a cardboard box if my life depended on it. If you offered me ten million dollars to create a program that says "hello world," I couldn't do it.

My lack of tech savvy notwithstanding, I have been able to develop and successfully deploy a significant amount of custom software – both locally and in the cloud – to automate routine tasks. Some of this involves direct use of AI / LLMs via APIs; the rest of it does not – but could not have been cost-effectively created without the help of AI.

Examples, small and large:

Small: file naming/categorization: AI bulk-created folders for essentially all U.S. public companies in our investable universe and built automation to standardize date formats and route files correctly by ticker and company name. This prevents lost notes and saves the hours that would be wasted hunting for them.

Large: investment letter ingestion/scoring: I've built a tool that reads investment letters of other fund managers, ignore all macro / philosophical discussion, extract only single-ticker investment ideas, summarize them, and scores them against a rubric based on our historical priorities. That rubric was – drumroll – drafted by AI after reading years of our letters, then subsequently refined by me.

This is something that we could not previously have done, because there are simply far too many investment letters out there, and the signal-to-noise ratio – even for talented investors – is generally low. We'd have to read hundreds of pages before encountering an idea that potentially aligns with our process. Now, if another smart investor discusses a fat pitch in their letter that's right in our sweet spot, we won't miss it.

Tangible impacts: Time spent on certain tasks cut by 95 – 100%. New workflows enabled that could never have been justified due to low signal-to-noise ratio.

2. Mass Information Identification, Ingestion & Analysis

Purpose: rapidly surface relevant facts and KPIs – including custom, qualitative KPIs or metrics disclosed infrequently or haphazardly – from a flood of data, allowing us to focus human attention where it matters most.

Background: Although LLMs have a well-documented tendency to hallucinate, not all AI tools are LLMs; some tools, such as NotebookLM, may miss context or misinterpret things, but essentially never hallucinate or make things up. We have identified which workflows to route to more creative tools like LLMs that can identify and synthesize information that may not be in our data-set – and which workflows require high-fidelity tools such as NotebookLM that stick to the truth, and nothing but the truth.

Real-life examples from July/August:

A. Grocery Outlet – getting up to speed on a complex new business in ~3 days.

Grocery Outlet is a business we'd never studied in depth before last month. On the surface, it appears to be just another grocer, but it has two unique traits that bear deep analysis: its opportunistic purchasing model (buying excess goods from CPGs, like Ross or TJ Maxx do in apparel) and its quasi-franchise independent operator structure, which creates a substantially different economic model than corporate-owned stores. Add in management and strategic changes that required diligence – including studying new CEO Jason Potter's background at The Fresh Market – a botched ERP rollout, and inflation, and it's a complex story that would normally have taken us nearly two weeks to diligence in depth.

Prior to AI, I often had to read documents at least two or three times – the first time or two to grasp the basics, then circling back to dig in on key issues. To develop a full, comprehensive understanding of the business, it would likely have taken up to two weeks assembling the full picture. Instead, in just three days, we had a completed 50-page investment memo with detailed analysis of each key issue. We were fortunate to establish a position before the company's recent 40%+ pop after a strong earnings report and extremely impressive conference call from new CEO Jason Potter. Without AI, we would've filed it under "interesting, but no time" and missed the move entirely.

As an even more specific example, one item that came up in our research was that Grocery Outlet was being sued over reference pricing claims (i.e. "we sell this for $1.99, others sell it for $3.99." Instead of spending hours searching for press coverage, AI quickly identified comparable lawsuits, summarized the claims, settlement patterns, and outcomes, and linked every point back to source material. Within fifteen minutes, we knew the suit was unlikely to be material. Conducting such research manually could easily have taken two hours or more.

B. Fitlife acquisition of Irwin out of bankruptcy – thousands of pages distilled in hours.
Last week, Fitlife Brands announced it was acquiring Irwin Naturals out of bankruptcy, doubling the size of its business. AI allowed us to produce a comprehensive update merely ~5 hours after the transformative acquisition. First, it located the bankruptcy filings in minutes, along with various news articles and financial disclosures of Irwin that otherwise would have taken us an hour or two to locate.

Second, after we had pulled a few key documents, NotebookLM analyzed these documents at our behest to tell us which of the other 700+ filings likely contained important information. It referenced a few dozen, which we then downloaded to update its data-set (imagine how much time this saved vs. clicking download on 700 documents!). NotebookLM built a chronological narrative from thousands of pages of legal text, letting us focus on the filings that truly mattered. A manual approach would have taken days, yet still missed details buried in obscure exhibits.

C. Updating old research – instant reorientation.

Over the past decade, I have produced extensive analysis and research on likely hundreds of public companies, but many of these reports have not been updated in years. If we last analyzed a business two or three years ago – let alone 5 or 10 – it can often be a daunting task to even get back in our original headspace, requiring re-reading of all of our original work, as well as hours of trying to get up to speed on how various segments have performed, what may have changed in terms of acquisitions/dispositions, the competitive landscape, and so on. Now, AI instantly summarizes the original analysis, compares it with subsequent disclosures, flags where we were right or wrong, and lists the key new issues. That allows us to decide in minutes whether a name deserves a fresh deep dive, rather than losing a day just to get back to square one.

Tangible Impacts:

- Days / weeks of research shrunk to hours / days, enabling us to reinvest research hours into breadth (covering more companies) or depth (delving into topics that previously might have been beyond the scope of analysis).

3. Writing & Documentation

Purpose: preserve our thinking for future reference, automate tedious transcription, and speed up both internal and external communication without losing valuable depth or nuance.

Background: writing takes a lot of time, but is a critical and central part of our process. Early on we learned that if we did not write down our thoughts at the end of a research project, we would forget most of what we had learned a few months later. The problem is that writing takes time; some of it is enjoyable while other parts are pure drudgery. I enjoy writing my personal thoughts and interpretations, but do not particularly enjoy transcribing factual details like the number of locations a company has or the exact purchase price of an acquisition. These facts matter, but they are tedious to document, especially in bulk.

Automating the tedious parts.
For companies we know well, such as Franklin Covey or LGI Homes, I had largely abandoned writing formal quarterly updates, relying on quick notes or sometimes skipping updates altogether. That saved time, but it risked missing subtle developments or losing the thread of how my thinking evolved over multiple quarters. AI changes this by handling the repetitive parts of quarterly reporting. I am restructuring our research format so that KPIs and expectations are tracked in a way that allows AI to auto-generate the factual quarterly updates. That frees me to spend my time focusing on understanding key shifts and qualitative nuances that AI cannot see, and writing down only these portions of the analysis. (AI now writes ~75% of our research, at least as a first draft; I simply go through and edit, add, or subtract as needed until it matches what I have in my head – but this is still vastly faster than doing it by hand.)

Making summaries less painful.
Even with long-form writing, the most tedious part is often distilling it into a concise summary. For research documents, I now use AI to create a tight executive summary and extract KPIs and other critical factors from lengthy, detailed, nuanced writeups, simply editing the summary as needed. This keeps my voice intact while saving hours of current and future time.

Creating investor letters.
My letters are famously long; I wrote this one from scratch and it was about 20 pages originally, then AI helped substantially improve the structure and organization, and cut the length in half. As the apocryphal quote goes, "*I didn't have time to write a short letter, so I wrote a long one.*" True no more!

This approach extends to our portfolio commentary and statistics. Historically, producing these was time-consuming, especially the seemingly simple elements like standardized financial tables. The same goes for rewriting detailed internal investment memos for an audience that needs more of the highlights. Once our upgrades to our documentation process are complete (likely within the next few months), this should turn our writing of portfolio commentaries – which has often consumed a weekend or longer –into a process that takes just a few hours.

Tangible impacts:

- More complete and consistent documentation of our thinking over time.
- Significant time savings on quarterly updates and client materials.
- More mental energy for the parts of documentation that require my personal touch.

Section 2: Trade-Offs: Are We Missing Something?

A devil's-advocate point of view is that we must intrinsically be missing something – that even if we use high-fidelity tools that don't hallucinate, surely something is lost when you take the human factor out of the equation. (Incidentally, if anyone tries to rely on ChatGPT to do investment analysis, all I have to say is good luck: it frequently gets very basic facts wrong, such as the stock price, let alone key details. That doesn't mean that it doesn't have incredible use cases, but on its own, it's a recipe for disaster.)

It's reasonable to ask: how can we be confident in our analysis if we haven't done it all ourselves?

My answer stems from my comfort with the idea of trade-offs. I am not inherently a perfectionist; I am a *satisficer*. Some people criticize Roombas because they don't get the corners well. I rejoice: for a relatively low cost, I now never have to vacuum 90% of the house again. I just have to Swiffer and sweep the corners. The Roomba is not a perfect tool, and in fact it is often a very stupid one that requires my intervention to function properly. But it is still night and day better than manual vacuuming.

Implicit in the concern that we might "miss something" by utilizing AI is the idea that human analysis is perfect. It is not. Any analyst may miss something during research – they may mistranscribe a number, or simply not notice an important development. Even for the most competent among us, we are sometimes biased, distracted, or undercaffeinated.

In fact, I'm not ashamed to admit that my eyes glaze over when I read legalese. Few or no people read every securities filing of a portfolio company, line by line from start to finish – it would simply be impossible. We quickly scan/skim looking for the meat of it, but if we're tired or distracted, we might miss something; indeed, companies often take advantage of this by burying necessary required disclosures in the footnotes or in filings that many people don't think to look at. Technology certainly has its flaws, but getting tired, distracted, or not having enough time in the day is certainly not one of them.

Similarly, if we spend hours researching any given company ourselves – with our own eyes – we are by default *missing* the opportunity to research *every other company* that we could have spent that time on. Errors of omission are just as costly as errors of commission. Indeed, we have *missed* many successful investments that we *could have made* due to a lack of time or energy – we were simply a week or a month away from getting around to that research project, or we started it but never finished it, due to time constraints.

Expounding on our earlier discussion of Grocery Outlet and Fitlife, researching a new company in the depth that I believed was required usually took me three days to two weeks, depending on the company and the industry; updates on existing research (depending on the situation and the required depth) could take a day to numerous days or even a week. Mathematically, then, it was impossible for us to do a deep dive on more than 50 ideas per year. We of course had heuristics that attempted to prioritize the best ideas, but these were imperfect.

I believe it is inevitable that there will be some flaws in our new approach that we will learn about over time. We will likely identify areas in which we can improve the interplay between human and technology. However, our existing approach was hardly foolproof, as evidenced by our poor results in recent years – a significant part of which is due to my own errors in judgment.

I firmly believe that broadening the top of the funnel will allow us to more rapidly identify the fat pitches, the obvious ideas that don't have the complexity of some of the ideas we have ended up struggling with in recent years. Following more of these businesses in real time, whether on our watchlist or as smaller portfolio positions, will allow us to scale up if they get to a truly compelling price.

Section 3: Evolutions in our Investment Approach

3.1: Core Investment Philosophy

Let's start with our core premise, which will never change: we are fundamentally value investors looking to build a portfolio of good to great businesses at good to great prices, primarily in the small and micro-cap space. We aim to thoroughly understand the key factors that drive business performance over multi-year periods, and purchase an ownership position in those businesses when market prices do not reflect our estimate of their intrinsic value.

We believe that intrinsic value is the discounted present value of a company's future cash flows, although we recognize that there are certain sectors or situations in which many/most other investors are – over extremely long periods of time, rather than just recent months or years – willing to use less conservative discount rates or more aggressive growth assumptions than we are, which we sometimes factor into our analysis in a modest way (i.e., being willing to accept slightly lower returns or pay a slightly higher valuation for businesses with extremely attractive fundamental characteristics, even beyond what we believe those characteristics are worth).

We believe that it is better to run a concentrated portfolio of a relatively smaller number of attractive companies, where every position is compelling as well as needle-moving, rather than running a highly diversified portfolio of 50+, let alone 100+, companies.

Finally, we are active managers rather than buy-and-hold investors. To minimize risk and maximize future return, we consistently evaluate the attractiveness of existing portfolio positions against other opportunities, both within and outside of the portfolio, and make swaps/replacement when needed.

Since COVID, we have discussed numerous evolutions in our thinking around this core premise, including:

1. An increased level of diversification

2. An increased focus on business quality
3. An increased focus on management quality
4. A recognition, as referenced above, that it is often worth paying somewhat more for items 2 and 3 (although we will never be "at-any-price" investors, and in fact we just entirely sold out of Bel Fuse, despite thinking it is a good business managed by an EXCELLENT manager, Farouq Tuweiq, because the stock is simply too expensive in our view.)

3.2: Our Watchlist

Our watchlist has always been the backbone of our investment process, allowing us to track potential opportunities until the moment is right to act.

AI has transformed how we manage and scale this watchlist. Although we had over 100 names on our watchlist, many had not been updated in years. Our original intention was that we should have done enough work for any name on our watchlist to be able to get up to speed quickly and make a purchase decision, but as years have gone on, this was not met. At any given time, there were roughly more like 50 companies we could get to a "buy" decision on within a day or less.

With AI-driven productivity gains, I've decided that our new goal is to be able to get to that "buy" decision within half a day for any company on the watchlist. We've trimmed the watchlist down to 31 names as of today – essentially our current portfolio and a small number of other names we know very well. Through a combination of new companies (such as Grocery Outlet) and those we've followed for many years (such as Korn Ferry), I intend to rapidly rebuild the watchlist.

My goal is to reach ~150 names by this time next year, and ~250 the following year, but with the key proviso that all names on the watchlist will have *relatively recent and actionable research.*

AI-driven processes make this possible by producing updates every 3–24 months depending on our affinity for a company and its proximity to fair value. Portfolio companies and imminently-actionable opportunities will receive more frequent, detailed refreshes, whereas companies trading well above fair value may get a quick, mostly AI-generated update just to keep them in view.

If you visualize our investment process through the lens of a typical "sales funnel" – where the top of the funnel is our watchlist and the bottom is our portfolio – this higher-frequency, broader-coverage approach expands our top of funnel dramatically.

We believe this will have numerous benefits. First, we will have much broader and deeper research, which will expose us to more investment opportunities.

Second, the watchlist also enhances our ability to act with timeliness. Our old watchlist only flagged companies for review when their prices declined, but there are often many interesting opportunities where a company's stock stays flat – or even rises – but value creation has outpaced the stock price, such that the stock is even cheaper despite a flat or higher price. (Fitlife is a great example of this, where the stock is flat to prices it traded at earlier this year, even though the acquisition boosts our fair value estimate by ~30%.)

By maintaining current, contextualized views on a much larger set of companies, we are better positioned to strike when catalysts emerge or a business turns the corner. Our recent investment in Grocery Outlet is a good example.

Third, I believe the watchlist process will feed a virtuous cycle of judgment improvement. Think about investment research like a tennis serve. If we produce a beautiful piece of research – but don't check back in to see where the ball landed after we hit it – then we're missing valuable opportunities to learn. Every update is an opportunity to check prior forecasts against actual results. That learning compounds over time, improving decision quality not just for watchlist names but across the portfolio.

Fourth and finally, the watchlist amplifies the advantages of being a generalist. Patterns we see in one domain often transfer to another. By following a wider range of companies in real-time, we build a much broader bottom-up view of how developments like AI, tariffs, and other key trends are actually impacting companies on the ground. Those cross-currents often surface opportunities in places that look unrelated at first glance.

This rebuilt and AI-enhanced watchlist will not only broaden the range of opportunities we can track but also change the shape of our portfolio. With more ideas vetted and actionable at any given time, we can make better relative comparisons between existing holdings and new candidates, and we can spread capital across a slightly larger number of compelling opportunities without diluting conviction. That leads directly to the next evolution in our approach: a measured increase in diversification.

3.3: Meaningfully More Diversification

As referenced, I strongly believe in concentration. However, I also believe that our beliefs should be based on facts – and if the facts change, so should our beliefs.

To me, the philosophical premise of concentration is essentially *"one person is only capable of having so many good ideas."* Such a premise is actually not inconsistent, in certain aspects, with having a portfolio of hundreds or even thousands of securities. As a toy example, imagine that you have one good idea, which is a quantitative algorithm that looks for certain patterns or market inefficiencies and takes advantage of them. You're still staying within your lane, even though you may own hundreds or thousands of stocks whose businesses you know nothing about.

That's not us; we're fundamental business analysts. At the same time, something I've come to believe very strongly in the past few quarters is that increasing the size of the top of the funnel will substantially improve the likelihood of us achieving superior investment returns. Allow me to walk you through a little thought experiment.

Let's say, for ease of math, that prior to AI, I was able to look at – in some unquantified minimum level or threshold of depth – at 80 stocks per year. (Obviously I look at far more stocks than this on an annual basis, but I'm talking about some level of research that goes beyond simply reading a little, and instead seriously contemplating the business model and valuation. 80 is as good a number as any, equating to about 1.5 per week, which equates to perhaps one new name and one company update or follow-up on an existing watchlist name per week.)

Let's also say that I owned a portfolio of 12 such stocks, named A, B… L. While reasons for owning securities vary and analysis of which is a better risk-reward is inherently subjective, for the sake of argument, let's (to keep the thought experiment clean) assume that stocks can be ranked objectively. Therefore, out of the 80 stocks I looked at, we own the 12 best, or the top ~15%.

But there are thousands of publicly-traded stocks that at least nominally fit our potential investment criteria. What I choose to work on is often somewhat random. As referenced before, there are many ideas that I know are attractive, but I don't have time to look at them. And then there were the ideas I never saw. For example, going back to my letter-reading tool, let's say I have friends named John and Jim who both write interesting investment letters. Maybe I read John's letter, but didn't have time to read Jim's, so I never saw his interesting investment ideas, even though he had ideas that were just as great as John's.

Let's say that AI allows us to look at 240 stocks per year in the same level of depth, rather than 80, with about half of these actually ending up on our watchlist. Let us further assume that there are only *half* as many really good ideas in the new stocks analyzed versus the old, just for argument's sake.

If this is the case, then assuming a normal distribution of sorts, there should be at least 12 stocks (M, N… X) in the new batch of 160 that are equally or more attractive than our existing 12 positions, and certainly the less-attractive of our existing positions. If that is the case, then there's really no justification for us only owning 12 stocks A through L, as opposed to some larger subset of A through X.

In practice, concentration helped us substantially pre-COVID, and seems to have hurt us since, sometimes even in the same names (such as Franklin Covey). I believe that our historical concentration has been a product of the

phenomenon I described – not so much only having so many good ideas, but rather *not having enough time to review all the ideas that we had.*

I've previously discussed what I believe are the numerous benefits of a happy medium, i.e. a concentrated but not *too-concentrated* portfolio. However, with the benefits of AI allowing us to cover far more ground, my plan is to migrate to a portfolio that looks more like 20 stocks, plus or minus a handful. (I still believe it would be hard for us to have truly differentiated insights on 25 - 30+ stocks, and I also still believe that we will frequently find names that are sufficiently more attractive than others such that they deserve a high single digit allocation. Finally, although this is a behavioral bias, I also believe I would just find it hard to personally care about that many positions, or spend that much time, energy, and effort for all of our positions to be 2-3% allocations which ultimately might not move the needle if we get it right. While this is clearly a personality flaw on my part, it is at least one that I am aware of. These three factors, in concert, explain why I currently have no plans to diversify further.)

Most positions should be in a range of 4 – 6%, which makes mathematical sense (100 divided by 20 is 5, so a 20-stock investment portfolio, if fully invested, must have an average position size of 5%). We believe this will allow us to allocate more capital to those ideas that really deserve it, rather than every idea becoming a larger position somewhat out of necessity.

Again, for clarity: we are still perfectly willing to take larger positions when we believe the risk-reward is disproportionate – indeed, we were recently buying more Fitlife and Franklin Covey even though they were already our largest positions in the 8 – 10% range. But we will do so more judiciously, reserving high single digit allocations for truly special opportunities (and we'll have a better ability to judge what is truly special because we'll be kissing a lot more frogs before deciding which ones to marry).

20 is not a magic number and perhaps it will be as high as the low 20s, or as low as 15, if we find ourselves in an environment with a paucity of opportunities. Simultaneously, given what we believe is an increasing ability to track a much broader range of names – and catch and take advantage of any temporary dislocations driven by liquidity or other factors – we're also somewhat less opposed to temporarily holding some cash than we once were – to be clear, I'm not talking about routinely having meaningful cash, or probably ever having a double-digit cash allocation, or attempting to time the market in any way. However, given what we believe is a rapidly increasing ability to identify new investments – with a particular emphasis (that we will discuss in a section below, as well as future letters) on investments that are *timely* – we are more comfortable having some transitory cash than we have been for the past 7-ish years.

3.4: Modestly Increased Preference for Liquidity

Not entirely related to AI, given the performance of the past few years, we've decided that we have a slightly increased preference for liquidity. While we love micro-caps, there is nothing inherently wrong with investing in small or even mid-caps. We've even invested in and done well in the occasional large cap, although that is neither our preference nor something that will ever be a large part of our philosophy.

There is an AI angle here given that liquidity provides the benefit of being able to take advantage of attractive opportunities when they arise – but the motivation here is more simply factor diversification, as really small and micro-caps have seemed to have a really tough go of it lately, while other similar (but slightly larger) businesses have performed better in the market based on similar results. At the end of the day, our focus is on price-to-value and not arbitrary market cap for the sake of market cap, so while we strongly believe that small and micro caps are vastly more inefficient, we plan to spend slightly more time on small caps as well as micro caps – Grocery Outlet being an excellent example.

3.5: More Qualitative Assessment of Fair Value

This is not an explicitly AI-driven change. Most of our clients are familiar with how my watchlist and portfolio spreadsheets are structured with various price-to-fair value metrics, specifically my "3-year CAGR" estimate (i.e., the investment returns we would achieve if the investment traded at our estimate of fair value in 3 years).

All models are wrong; some models are useful. I have previously discussed how in certain cases, I've been led astray by focusing too much on this metric (or more broadly, our own fair value estimates.) That is – to be very clear – **not to say that valuation is not important.** I believe that **valuation is critically important.** However, there are other factors at play.

Bob Whitman at Franklin Covey – a distinguished and highly successful real estate investor – once told me that not all CAGRs are equal. While this was a long time ago, it was quite thought provoking.

One can, of course, try to incorporate quality explicitly into fair value estimates – which we, as most investors, obviously do to some degree. At the same time, leaning too heavily into this can lead to silly outcomes. At the extreme, are we really willing to argue that any reasonably mature business has a fair value of 30x its free cash flow? What about 35? What about 40? Obviously not.

Conversely, businesses with extremely stable or resilient demand obviously have a higher chance of meeting their forecasts than those that do not; even if a grocery store and a high-end steakhouse, let alone a high-end hotel in a far-flung tourist destination, somehow had identical financial figures, the first has much more certainty of outcome than the second or the third, no matter what may happen in the world.

It is of course not so obvious in the real world, but I've learned over time to respect qualitative factors – whether business quality or management quality – that cannot be explicitly captured in a valuation model (or, as discussed, can be captured but only by putting in silly premiums that lead to a whole other set of problems, i.e. runaway justification of owning assets "at any price," so long as they satisfy some arbitrary set of quality or management criteria.)

There are bright lines, i.e. you are unlikely to ever find us owning mature assets at 25-30x free cash flow unless something very spectacular is happening. And of course, that's hardly the threshold; we are usually uncomfortable continuing to hold businesses at multiples meaningfully below that. For example, we just entirely sold out of Bel Fuse, despite thinking it is a good business managed by an EXCELLENT manager, Farouq Tuweiq, because the stock is simply too expensive in our view.

Even if we pencil in pretty good 2026 numbers – and value the stock based on these – it appears to trade at ~13.5x EV/EBITDA and 20x EV/NOPAT on a forward basis. These seem like quite full multiples to us, and we know of numerous businesses (both ones we own and do not own) that either trade at a similar multiple with much better business quality, or offer similar business quality at a much lower multiple.

While BELFB's multiple is not out of line with industry peers, the business does still have some economic and cyclical exposure, not to mention meaningful exposure to tariffs and a supply base in China. These same risk factors caused the stock to trade down to ~$60 in the wake of Liberation Day… merely ~4 months later, the same business trades at ~$130. Although Bel could certainly justify its current valuation and make us regret selling it, we think that most of the low-hanging fruit has been picked, and it will be incrementally more difficult for Bel to drive the kinds of substantial self-help improvements that it has achieved in recent years; conversely, a weakening macro or more tariff volatility could send the stock back down. Although we will continue to track and admire the company from afar, it no longer fits our value-investing philosophy.

Conversely, we held on to most of our Grocery Outlet (GO) stock after the 40%+ post-earnings pop, only making a modest trim, because as we continued to conduct more due diligence, including a half-dozen calls with former key leaders, we concluded that it is a special and competitively differentiated business that has suffered from a combination of excessive growth expectations (which led to a lack of focus on getting the basics right), as well as a botched ERP implementation (which caused ripple effects throughout the business).

With new CEO Jason Potter having an excellent history of turning around grocery peer The Fresh Market (TFM) – and a refreshed Board including former Sprouts CFO Chip Molloy – the setup here reminds us a lot of Jack Sinclair's early days at Sprouts. Given the habitual nature of grocery purchase, it takes time for operating improvements to flow through to comp growth, but there seems to be a multi-year runway for driving numerous

blocking-and-tackling self-help improvements, ranging from private label penetration to better merchandising to better technology tools for stores to manage their forecasting / inventory ordering. We sold Sprouts far too soon, and don't intend to make the same mistake here.

Finally, this is a business that benefits from many of the key challenges in the world – it passes along inflation to consumers, and as an extreme-value grocer, it sees not only stable but rapidly growing demand during economic downturns (it comped something like 27% cumulatively in 2008/2009 during the GFC). Therefore, although the stock is not screamingly cheap at ~17.5x our forward estimate of NOPAT, we believe that its very stable demand profile – and a long runway of improvements – make it worth continuing to hold a meaningful position.

Again, I want to be careful to thread the needle here; we would likely not hold it if it traded up to (for example) 22 or 25x our forward earnings estimates, as we would feel at that point that too much future improvement was being priced in – but just because we're only willing to assign a fair value of ~20x EV/NOPAT doesn't mean that we think we should be selling aggressively at ~17.5x to fund purchases of vastly inferior businesses, even if those businesses trade at a meaningfully greater discount to our fair value estimate, given what we view as the very high probability that this business will achieve our expectations over the next several years.

3.6. FIRM RULE: DO NOT BET AGAINST AI

As a value investor, it's always tempting to be a contrarian, but I've learned over time not to bet against disruption. I think this is very true in AI.

Recently, I've seen a number of businesses that I – and others – previously believed to be very high quality – that have sold off dramatically due to the perception of AI disruption risk. (In particular, there are a lot of content or research-based businesses in the crosshairs.) While several of these are quite tempting, I've decided to implement a rule. Rules exist to keep you from doing dumb things even when you want to do them. And my rule is, *don't bet against AI*. Situations like Chegg show that the disruption risk is just too great.

Of course, this is a subjective measure. In the extreme, frothing-at-the-mouth tech-bro AI-maximalist take, artificial general intelligence (AGI) is right around the corner and every job (other than worshiping our AI overlords) will be eliminated. Perhaps some of our portfolio companies might be more affected by AI than we currently anticipate. Nonetheless, if I see an obvious risk, I will respect it, and I will err on the side of perceiving more risk rather than less. There are plenty of businesses out there for whom AI is neutral or even a tailwind, so it doesn't make sense to me to make the anti-AI bet.

If you see me breaking my rule, please yell at me!

Conclusion

I've never been more excited about investing. I love learning about companies and how they operate. I hate sifting through boilerplate legalese looking for the one line of real content, or doing manual data entry to stitch together multiple years of disclosure of a key metric beyond the income or cash flow statement. AI is excellent at reducing this drudgery, meaning that I get to spend my time on the *fun* stuff – contemplating businesses and making judgments about their futures.

Although I believe our investment results would have improved materially even without the tailwind of AI, due to improvements in our process we had already made, I believe AI is a massive accelerator to those efforts, and should materially improve our returns while hopefully reducing our risk. This is, of course, not a guarantee of investment performance, but does represent my sincere belief.

Westward on,

Samir