



Machine Teaching:

What I Learned From My Optimizer

December 2023

DE Shaw & Co

www.deshaw.com/library

At the D. E. Shaw group, we've deployed systematic and discretionary investment approaches across global asset classes for more than three decades. We've benefitted from housing multiple disciplines under one roof, in large part because of how they interact with and inform one another.

One key example? We've developed sophisticated approaches to constructing and risk-managing portfolios in our systematic disciplines. This has helped us build optimizers that we utilize across our discretionary investment strategies as well, introducing additional rigor into human-driven processes. We've learned a great deal from working with those tools, and they've sharpened our intuition as investors. We reflect below on those lessons, starting with some basics.


What is Optimization?

It's helpful first to define terms. Different practitioners may have their own definitions but, in essence, an optimizer is a computational tool used to *optimize*, or find the best solution to a constrained problem. This generally takes the form of maximizing a specified utility function by simultaneously considering positive and negative utility components.

For most of human history, our brains have been the most advanced optimizers available to us, and we continue to use them in our daily lives. For instance, with a limited number of hours in your Saturday (constraint), how do you balance the errands you need to run, the chores you need to complete around the house, and the recreation you actually *want* to do? How should you prioritize and order those activities? Your implicit optimization process (reflected in *Figure 1*) might lead you to determine that three hours of errands is optimal: Your utility increases rapidly as you run your most important errands, less rapidly once you've crossed the first few tasks off your list, and then starts to decrease as the opportunity cost of the activities you've neglected starts to predominate (if you spend all your waking hours out running errands, you have no time to go to your neighbor's barbecue or do your laundry).

Because optimization problems can be complex, human brains often rely on heuristics, which can introduce cognitive biases such as overreliance on past experience, binary thinking, excessive optimism, and so on. The invention of calculus in the 17th century, and the discovery of a function's gradient, represented a step forward in our ability to mathematically analyze such

Figure 1: Utility as a Function of Hours Spent Running Errands



problems. As computing power has surged since the mid-20th century, it's become possible and commonplace to harness the power of machines for optimization. Simple optimizers can be built in a spreadsheet, while others may take the form of complex pieces of custom software.

Optimizers play an important role in numerous fields, including:

- designing truck routes to maximize delivery efficiency relative to fuel consumed,
- scheduling sources of electricity generation to maximize reliability relative to costs and carbon emissions, and
- improving the accuracy of machine learning models.

In finance, optimizers are used by a range of market participants to help solve different types of problems, informing asset allocation decisions made by institutional investors and risk allocation decisions within banks and insurance companies. They are also used to inform portfolio construction decisions made by investment managers like us, in both systematic and discretionary contexts.

For purposes of this discussion, we're focused on discretionary investing, in which an optimizer is a tool that can help humans process information and make better decisions. For us, there is no single, uniform "Optimizer" that we deploy across portfolios and markets; rather, we apply a variety of optimization tools tailored to individual investment strategies and use cases. It's worth keeping in mind that the examples in this piece are similarly diverse, representing illustrative (and hopefully informative) lessons rather than an all-encompassing playbook for optimizer use. As with so much in investing, usage depends on context—there's no right way to work with an optimizer.

In these discretionary contexts, we've found that beyond helping humans make decisions about the portfolio at hand, optimizers through their output can also *teach* humans—including how to better design and interact with an optimizer itself, and, importantly, how to counter behavioral traps and cognitive biases when making investment decisions.

Optimize, Understand, Repeat: Optimization is a Conversation

At the heart of building an optimizer, and using and learning from it, is iterative interaction.¹

While the science of optimization is heavily grounded in math, the art of optimization—improving design, inputs, and hopefully outputs—comes from confronting it with human intuition, and *vice versa*. It's this iterative process that provides learning opportunities for the human practitioner.

Imagine a portfolio manager has a new trade idea and an intuition about what size it should hold in their portfolio. They add the trade's forecast and other parameters to their optimizer. If the optimizer's suggested trade size diverges meaningfully from the manager's initial intuition, what does that mean, and what should they do next? The answer is likely to lie somewhere between "do exactly what the optimizer says" and "ignore the optimizer."

In this situation, assuming the optimizer's design is sound, a starting framework might be to assume that the gap between intuition and optimizer can be explained by some combination of (a) inaccurate inputs and (b) biased or incomplete intuition. In other words, the manager might ask: Why does the optimizer's desired trade size differ so meaningfully from what I expected? Did I miss something fundamental when estimating my parameters? If not, what is the optimizer's output telling me about my own assumptions, judgment, and possible biases?

With these questions in mind, the manager may want to iterate with the optimizer to identify variables to which the suggested position size is most sensitive, and then challenge those inputs. How confident are they in their estimates, and if their error bars are wide, how can they communicate that to the optimizer? Are there variables they're implicitly considering that aren't reflected in the optimization? Is the optimizer aware of a relevant variable that they've overlooked?

¹ Certain optimization algorithms themselves rely on iterative computation to arrive at their output, but that's not our focus here.

They might also find it helpful to work backwards, finding sets of inputs that yield a suggested position size from the optimizer that matches their intuition. Are those inputs even plausible? Iterating with an optimizer in this way can challenge the manager's assumptions, and hopefully sharpen their judgment.

Throughout this piece, the observations we share are the result of years of iterative interaction between human and machine—machine teaching.

Optimizers Are Power Tools: Use With Caution

One type of lesson we've learned from working with optimizers relates to the behavior of optimizers themselves. For example, they sometimes produce outputs that are unexpected or counterintuitive. In these cases, there can be an immediate lesson about an optimizer's design. But often there is also an embedded lesson that informs an investor's own intuition.

Flat Curves, Sharp Corners: Optimizers are Sensitive

As many practitioners may have observed, optimizers are prone to putting things in corners.

An optimizer's output suggests the single optimal thing to do—in the case of an investment portfolio, the mix of positions and sizes that maximize the *ex ante* utility the investor has designed their optimizer to target. That output can be surprisingly sensitive to even small differences in inputs, where the optimizer seems to jump erratically from one suggested allocation to a materially different one.

This is because utility curves are typically, well, curved, and the optimal point on a utility curve often falls at a relatively flat portion of its surface. As a result, near the curve's peak value, a substantially smaller or larger allocation to a given position might change the portfolio's net utility only marginally. This leaves the optimizer scope to pursue unexpected corner solutions, making large jumps in sizing based on small input changes.

To illustrate, consider how even a small change in an asset's forecast Sharpe ratio can lead to a big shift in "optimal" position size. As depicted in *Table 1*, imagine a simple portfolio comprising two highly correlated (~0.95) equity indices with identical forecast Sharpe ratios. In this case, the optimizer will suggest a 50/50 allocation to the two.

Now, suppose the manager's estimate for Index B's volatility increases, bringing its expected Sharpe ratio down slightly. As a result of this seemingly insubstantial change, and the fact that the optimizer seeks to capture even marginal utility improvement, the optimizer will now recommend a 75/25 allocation between the indices.

Table 1

	Initial Forecast		Updated Forecast	
	Forecast Sharpe	Suggested Allocation Size	Forecast Sharpe	Suggested Allocation Size
Index A	0.200	50%	0.200	75%
Index B	0.200	50%	0.195	25%

The extent to which this new output should be concerning depends on the manager's confidence in their inputs. The optimizer is in effect asking the manager whether they believe that the small difference in forecast Sharpe ratio represents something substantive or simply reflects noise in the estimates. If it's noise, then the output may be driven by overprecision in the manager's forecasts and assumptions. Underappreciating uncertainty can lead to real differences in a portfolio's composition. We're reminded to (a) challenge, and iterate carefully on, our inputs and/or (b) incorporate uncertainty around them.

If it's not noise, there may be another lesson. While the difference between a 50/50 and a 75/25 allocation might seem dramatic to the manager, from the point of view of the overall portfolio, it's not. Instead, what matters to long-term outcomes is the *utility of the portfolio as a whole*, not necessarily the size of any one trade along the way (on which we might otherwise tend to fixate). This ability to consider a portfolio's overall utility is one of an optimizer's greatest strengths, and one to which we'll return throughout this piece.

Be Careful What You Wish For: Optimizers Obey Constraints—Even Suboptimal Ones

An optimizer is generally obedient to the letter of its code, impervious to the human temptation to ignore or subjectively interpret certain considerations. But like a genie who grants wishes too literally, the consequences of such rigid obedience might not always match the spirit of what's intended.

Portfolio objectives and constraints illustrate this phenomenon. An investor often wishes to impose certain *ex ante* guardrails on their portfolio, and an optimizer can be programmed to do just that—at least *ex ante*.

Imagine that a manager of an active, benchmark-relative equity portfolio wishes to maintain a consistent active risk target for the portfolio (e.g., a tracking error between 4.9% and 5.1%) at all times, while also holding diversified exposures across sectors and industry groups. Much of the time, perhaps, those objectives can coexist without issue. But if the market enters a period of persistently low idiosyncratic volatility, the manager's optimizer may, prompted by its constraints, seek out active risk in unexpected ways.

For example, the optimizer might overweight the most volatile stocks in each sector and industry group, maintaining overall tracking error while preserving targeted diversification. But this could entail hidden costs: the portfolio might now be positioned as a large, *undiversified* bet on volatility as a risk factor, and it might be underexposed to forecasts the manager actually wishes to express. Providing a degree of flexibility around the optimizer's constraints may help mitigate these kinds of unintended consequences. Doing so also acknowledges that even the most careful manager will not be able to fully identify and control for all risk factors—or all possible states of the world.

To take another example, after the Global Financial Crisis, a portfolio manager understandably might have wanted to measure their portfolio's exposure to such an event. If they went a step further and attempted to constrain their portfolio to some *ex ante* loss limit in a GFC-like scenario, it would risk overfitting their positions to how individual assets performed in a highly specific set of historical circumstances. A more robust approach may be to penalize or otherwise constrain the portfolio's expected losses across a *range* of conceivable crises, and to reflect uncertainty as to how a given asset might react in such scenarios. As market participants were most recently reminded in early 2020, the next crisis rarely looks like the one before.

It's tempting as investors to try to constrain portfolios with surgical precision. An obedient optimizer reminds us that it's not so simple in the real world, where every objective comes with tradeoffs. If a constraint is too ambitious or inflexible—or the conditions unfavorable—it can lead to unexpected costs. We're encouraged to keep top of mind the law of unintended consequences when thinking about how to design robust targets and constraints for a given portfolio.

Optimizers Are Power-Ful Tools: Draw On Their Strengths

A question we're occasionally asked is why we use optimizers in discretionary investing at all.

One reason we find it valuable is that, through their objectivity and computational power, optimizers help human traders think through the multilayered implications of their forecasts and positions.

Thinking Local

Optimizers demand inputs. They require traders to quantify their beliefs and assumptions and to analyze each component of their forecast individually.

For example, if in June 2019 a trader of an equity market-neutral portfolio had glanced at EURO STOXX® dividend futures on their Bloomberg terminal, they'd have seen six-month contracts trading at just over €120, whereas contracts eight years out were trading below €95. That trader might have perceived an attractive opportunity in those longer-dated dividend futures: was it reasonable to expect that dividends would be ~20% lower in eight years? What's more, at first glance the trader may reasonably have presumed little fundamental relationship between *long-term* dividend futures and *today's* stock prices. After all, if stocks are down 1% today, does that mean dividends are likely to be down 1% several years from now?

Given their market-neutral mandate, if the trader were using an optimizer, that optimizer would have required as a standard input any trade's beta to equities. In calculating this input for dividend futures, the trader may have been surprised to find that, in practice and for relatively long horizons, dividend futures often have an equity beta of approximately one. Upon further reflection, this makes sense—an investor would ordinarily be able to “purchase” future dividends only by taking on the risk of equity ownership in the interim, and the co-movement in the dividend futures market exists as a result of the same equity risk premium. The discount the trader originally observed in the dividend futures market mirrored a similar discount in the index future itself, driven by expected dividends exceeding interest rates at that point in time.

Equipped with that information, the trader might *not* have been surprised to then learn that their optimizer would not have wanted any of the long-dated dividend futures trade. Considering its equity market-neutral objective, the optimizer would have perceived little net utility in the forecast—it would have wanted to hedge out the equity beta until there was effectively no trade left. In other words, the output would have helped the trader realize that, after accounting for correlations, there was little alpha in the forecast.

By demanding a standard and complete set of inputs, the optimizer caused the trader to confront a reasonable but incorrect assumption (that two instruments shouldn't be fundamentally correlated) and reminded them that correlations can be both hard to intuit and highly impactful to positioning.

Thinking Global

Humans are often disposed to think discretely and sequentially. Similarly, investors tend to think about the returns and risk of their positions individually and day by day.

This focus on individual positions can be useful in certain respects, including in attributing risk and monitoring performance. However, the value of an individual trade may best be understood not in isolation but in terms of its contribution to overall portfolio utility. That in turn depends not only on the characteristics of the trade itself, but also on its relationship to all others already in the portfolio, including asset-level correlations, contribution to tail risk, and even the origin or underlying view of each forecast.

Given the computational complexity involved, it would be challenging—if not impossible—for a human to fully account for all of these relationships on their own. Use of an optimizer can help an investor overcome that complexity, and in doing so help them more generally orient their thinking to the portfolio as a whole.

Trading Fast and Slow: Optimizers Are Helpful with Trading Speed

We find optimizers typically have the upper hand in handling transaction costs and governing trading speed. We believe this is true due to biases to which even an experienced trader may be subject, and the computational complexity associated with transaction costs (as complexity, in any problem, can make biases more likely to appear).

Imagine a trader forecasts that units of Asset A, which currently trade for \$100, will appreciate by 1% over a 30-day horizon. Further assume that the trader and their optimizer are aligned in suggesting that the ultimate size of Asset A in the portfolio should be two million units (approximately \$200 million of market value).

Now assume that the trader has provided their optimizer with reasonable inputs for transaction costs (for purposes of this illustration, those costs comprise only market impact/slippage). Aware of these costs, the optimizer suggests splitting the trade into a series of purchases over 20 days, allowing time for the market impact of each to partially fade before buying the next.

The trader, on the other hand, is inclined to complete their trading more quickly. They may be using an inappropriate trading heuristic, or perhaps psychological effects like salience bias and asymmetric loss aversion mean that the trader anchors to a prior experience when they traded too slowly and missed part of a profitable forecast. Whatever the case, in this example the trader might choose to speed up their trading relative to what the optimizer suggests.

In our experience, the trader's optimizer is likely the better guide. By more objectively balancing the horizon of the trade and its forecast realization, and the expected appearance and decay of market impact, it's likely to slow the rate of trading relative to the trader's inclination. The effects can be material, resulting in some cases in saved slippage costs of a magnitude comparable to that of the forecast profit. What's more, the costs incurred from trading can generally be expected to be realized with greater certainty than the results of the forecast itself.

The lesson here is to provide as much information about expected forecast realization and transaction costs as possible and let the optimizer handle trading speed, although in cases like the above we've found that lesson often translates more simply to "trade slower." Equipped with relatively accurate inputs, an optimizer is less prone to sources of bias to which the human trader is understandably susceptible.

We'd note an important caveat. In the illustration above, we assumed that the trader provided reasonably accurate inputs to the optimizer regarding both expected transaction costs and realization path (which need not be linear). That assumption is easily challenged in practice. As we've seen so far in this piece, and will see again, the output of an optimizer is highly sensitive to its inputs.

Trading Large and Small: Optimizers Are Helpful with Trade Sizing


In the previous example, we alluded to the concept of trade horizon for a single trade to illustrate an optimizer's role in guiding trading speed. We now extend our discussion of trade horizon, using a scenario involving two trades of different expected horizons, illustrating how an optimizer may determine optimal position sizing and challenge a trader's intuition in the process.

Imagine a trader has two forecasts, S and F, on highly liquid (*i.e.*, no- or minimal-slippage) instruments, and is evaluating how to express each in their portfolio. The trader believes that the two forecasts have the same holding period Sharpe ratio (*i.e.*, the expected alpha for each trade is proportional to its cumulative risk), but the trades differ in that Forecast S (the slow forecast) has an expected horizon of 100 days, while Forecast F (the fast forecast) is expected to be realized over 10 days.

In our experience, many traders tend to focus on peak position size, and here the trader might expect that the two trades—which, after all, offer the same ratio of alpha to risk—should be allocated approximately equal risk in their portfolio.

The trader might be surprised by what their optimizer suggests: to size Trade F much larger than Trade S. Given the parameters, the optimizer will aim to allocate the same *cumulative, lifetime* risk to the two trades (because they have the same expected lifetime risk-adjusted returns). The shorter horizon of Forecast F simply means that a much larger position is necessary to achieve the same lifetime risk as Forecast S.

Figure 2: Simplified Example



In our experience, even a trader who intuitively understands this phenomenon is often surprised by how much larger an optimizer wants to size the shorter-term forecast—in this case, approximately 3.2x larger than the longer-term forecast. Another way to consider the difference in desired position size is to note that the *annualized* Sharpe ratio of Forecast F is roughly 3.2x larger than that of Forecast S. This comparison of annualized figures makes clearer the optimizer’s appraisal that Forecast F merits a larger risk allocation during the much shorter period over which it is active.

The overall lesson is straightforward, if challenging to internalize: all else equal, and absent slippage constraints, a shorter-horizon trade is something a portfolio should want more of, often much more of, in terms of peak size, than a longer-horizon trade.

There’s also a more nuanced lesson, introduced earlier in the context of “thinking globally.” Many traders are disposed to think about—and care a lot about—the daily risk of their individual positions. That may be a function of certain human tendencies (e.g., risk and loss aversion) and of institutional constraints (e.g., risk targets or limits). That narrow focus, however, can obscure the fact that it’s the aggregate utility of a portfolio *over time* that matters, which in turn is a function of the lifetime risk and return of the individual trades it comprises.

Define the Relationship: Correlations Matter (a Lot) When Sizing Positions

In the previous illustration, the optimizer's output surprised the trader by liking a forecast much more than the trader intuitively would have, as a function of its trade horizon. Another way in which an optimizer can surprise us is in revealing how correlations impact a position's relative attractiveness in a portfolio. The underpinnings of these two lessons are similar: blind spots often arise due to humans fixating on an individual position's risk and return properties—thinking locally—rather than those of the assembly of positions that make up a whole portfolio.

Imagine a market-neutral portfolio manager has a forecast that, over the next year, interest rates will rise (that bonds will fall) more than the fixed income market is pricing in. They may be tempted to short bonds to express their forecast.

Because bonds have been considered a diversifier for equities for the past few decades, the manager is likely aware that the two asset classes have generally exhibited a moderately negative (approximately -0.25) correlation.²

However, the manager might not think about this relationship at all in the context of the perceived mispricing in the bond market, with some justification. After all, their forecast is on a binary outcome—either rates rise more than expected over the forecast horizon or they don't—and, for the sake of this illustration, they don't have a view on the direction of equities.

A well-designed optimizer, on the other hand, *will* take this relationship into account, especially given the manager's market-neutral mandate. It will recognize that to simultaneously express the forecast on bonds and maintain beta neutrality, it would need to implement a (short) hedge on the S&P 500®. It further will account for the fact that hedging via short equities is a costly proposition over time, given the equity risk premium, and thus might determine not to take the short position in bonds at all.

An optimizer is able to counter the manager's possible bias of implicitly rounding seemingly small numbers to zero, thereby ignoring correlation's effect on the trade's utility. It can also offer a reminder that a position's correlation to equities is fundamentally different than its correlation to other assets, given equities' risk premium and importance to the risk profile of any portfolio. (By contrast, to take a less likely example, lean hogs futures don't have much relationship to the economic cycle—and thus don't command a similar risk premium—so bonds' correlation to those futures may matter only if the manager also has an allocation to lean hogs.)

Traders may grasp directionally the ways in which correlation properties affect a trade's relative attractiveness in a portfolio. On the whole, however, we've found that they tend to underestimate how profound that effect can be, and they tend to be inconsistent at judging those effects when correlations vary over time. After all, it is not intuitive that the threshold for a good trade should change in ways that aren't revealed by looking at the instrument being traded. A well-designed optimizer can help counter those shortcomings and reinforce the importance of correlations, the utility of diversifying assets, and the disutility of assets that generate unwanted or overlapping exposures.

More generally, as we've seen elsewhere, an optimizer can help an investor widen their lens from the narrow properties of a single position to those of an entire portfolio, where an individual trade's utility can only properly be understood in relation to all others.

A Head for Tails: Optimizers Can Help Account for Crash Risk

As just discussed, correlations matter, as do other risk-related inputs to an optimizer. Putting that principle into practice assumes, however, that quantities like correlation and volatility can be reasonably predicted. An optimizer can help account for the fact that conditions can, and often do, diverge from what might otherwise be considered normal (specifically, what would be considered normal under a traditional mean-variance optimization framework).

² This value represents the mean of the 1-year rolling correlations between the S&P 500® and 10-Year U.S. Treasury Note between January 2000 and June 2023, rounded to the nearest 0.05. Applicable data used with permission of Bloomberg.

Given the uncertainty of dealing with such vagaries, a trader might be inclined to let a single aggregate input (such as realized volatility over some reasonably long timeframe) stand in for the range of plausible future states. This tendency may be especially pronounced when attempting to account for asset-level risk properties in a crash scenario, which by its nature can take many different forms at low individual probabilities. If instead an optimizer is designed to incorporate both normal and crash-conditional inputs, the results may reveal just how much conditions matter in constructing a portfolio, sometimes in surprising ways.

Crash Testing Credit

Imagine a portfolio manager is looking to allocate risk between investment grade and high yield credit. Using simple mean-variance optimization, along with return and risk inputs derived using a lookback period of five years, the optimizer may seek a portfolio balanced equally between the two assets:

Table 2

	Forecast Return	Estimated Volatility	Forecast Sharpe Ratio	Correlation	Suggested Allocation Size
Investment Grade	0.3%	1.5%	0.2	0.75	50%
High Yield	1.2%	6.0%	0.2		50%

Much of the time, the manager's inputs may be reasonably accurate. However, this simple mean-variance approach assumes an approximately normal distribution of asset returns, which, as the manager is aware, can be a poor assumption.

Aware of this shortcoming, the manager decides to incorporate a simple crash element, with a low presumed probability of occurring (5%) along with a penalty for exposure to such a crash, into their optimization function:

Table 3

	Forecast Return	Estimated Volatility	Forecast Sharpe Ratio	Correlation	Expected Loss in Crash	Forecast Return per Unit of Crash Risk	Suggested Allocation Size
Investment Grade	0.3%	1.5%	0.2	0.75	-5%	6.0%	40%
High Yield	1.2%	6.0%	0.2		-18%	6.7%	60%

This somewhat more advanced optimizer now seeks 60% exposure to high yield credit and only 40% to investment grade. Accounting for a conditional crash state, even at low likelihood, has a real effect on the desired portfolio.

The *direction* of that effect—allocating more to high yield—might seem surprising, since high yield credit is expected to lose more, in absolute terms, in a crash. Here, the optimizer is providing a lesson in the concept of “return per unit of crash risk.” Even though the two assets in this illustration have the same expected Sharpe ratio (a measure of return per unit of volatility, or normal risk), the manager's forecast of high yield's superior return per unit of crash risk makes it more additive to overall portfolio utility when accounting for a crash-conditional state.

In the example above, the manager incorporated a single crash element into the optimization process. The impact on portfolio positioning may be even more notable if they were to add a range of possible scenarios, with different asset behaviors and uncertainty accounted for. An optimizer, given its computational capacity, can simultaneously consider those multiple plausible states of the world, something difficult for the human brain to accomplish, and can maximize expected portfolio utility across them. It goes beyond the human tendency to focus on a single estimate for any given variable and a single, average state of the world. It also cautions not to under-index to specific or abnormal conditions.

Common Investor Risk is Uncommonly Important

A specific instance when extreme conditions matter stems from common investor risk.

Common investor risk can arise from overlap in positions, forecasts, or investment approaches across market participants. It may manifest as a spike in correlations between specific assets or trades when one or more such managers reduces risk.

The nature of common investor risk is such that (a) it can be hard to discern or model with great accuracy, especially given that it evolves over time, (b) it may exist in some modest amount across many positions all at once, and (c) it presents an acute, if usually infrequent, source of risk. As a result, it's a type of risk that requires a sophisticated optimizer to incorporate.

Imagine a trader has two forecasts, one an idiosyncratic forecast on silver futures, and one a forecast on lean hogs. Assume that the two forecasts are alike in all ways except that the lean hogs trade currently finds strong consensus among speculative investors. The trader's base inclination might be to size the lean hogs trade larger in the portfolio, given the comfort they find knowing that other investors share their forecast.

This intuition isn't necessarily wrong, but it is incomplete. Markets are generally efficient, and there can be useful information in what other participants believe. At the same time, that consensus may indicate the presence of common investor risk.

Even aware of this tradeoff, the trader might be surprised to find that their optimizer suggests taking a relatively small, or even zero, position in the lean hogs trade. This is a more extreme outcome than a trader will typically intuit, because they are generally inclined to think locally and focus on the perceived modest amount of such risk introduced by a single trade.

What an optimizer can consider more readily is the entirety of the portfolio in question. Across that portfolio, there are likely many individual trades, some of which have their own exposure to common investor risk. An optimizer can effectively manage an overall budget for common investor risk. Beyond a certain level, the bar for incurring one marginal unit of such risk can be high.

Because of its ability to factor in low-probability but high-magnitude outcomes at the aggregate portfolio level, an optimizer can help investors appreciate how serious common investor risk can be and how to avoid assessing it too locally.

Concluding Thoughts

The Original Optimizer: Humans Still Play an Important Role

Throughout this discussion, we've largely focused on how an optimizer's comparative advantages can help a discretionary investor achieve better outcomes, and the ways in which an optimizer's outputs can help improve intuition.

We value those takeaways, but we also believe that human practitioners still play a critical role in identifying when the optimizer may be unable to account for an unprecedented development or given state of the world. Sometimes human intervention is justified, and that intervention can take many forms, such as adding a risk factor, changing limits or other parameters, updating an approach to calculating critical inputs, or even putting aside the optimizer altogether.

We believe situations like these should be approached with intellectual humility. Otherwise, the temptation to tinker might override many of an optimizer's benefits we've described above. But we also believe that the more experience humans have working with their optimizers, improving inputs and testing outputs, the better they're positioned to judge when such intervention might be warranted.

Conclusion

Over the course of our wide-ranging investment history, we've found the optimizers we've built and interacted with to be enlightening partners.

Along the way, we've learned about how to better design and inform optimizers, accounting for factors like their sensitivity to inputs and proclivity for corner solutions. We've also learned a great deal about our own nature as investors (and as humans), including the effects of our overconfidence and overprecision in making forecasts, our tendency to under-imagine and under-appreciate risk and correlation, and our limits in intuiting multiple, complex computations at once. Those learnings have come from extensive iteration with our optimizers, and we fully expect we'll continue to learn from them.

It may seem trite in this period of rampant proliferation of artificial intelligence, but we find that a collaborative relationship between human and machine can be powerful in harnessing the comparative advantages of each. We've seen how our traders and portfolio managers internalize lessons from such collaboration, becoming more precise in their communication of assumptions, more aware of their intellectual biases, and more attuned to the tradeoffs inherent to any decision.

As investors, we find these lessons especially powerful given the limits we face in what can actually be known, or reasonably predicted, in financial markets. Even the best human practitioners, working in concert with the most advanced optimizers, won't be able to solve every problem in investing—but we'll happily take that option over going it alone.

>>> Enjoyed this piece? [Share your thoughts!](#)


THIS DOCUMENT IS PROVIDED TO YOU FOR INFORMATIONAL PURPOSES ONLY AND DOES NOT CONSTITUTE INVESTMENT ADVICE OR AN OFFER TO SELL (OR THE SOLICITATION OF AN OFFER TO BUY) ANY SECURITY, INVESTMENT PRODUCT, OR SERVICE.

THE VIEWS EXPRESSED IN THIS DOCUMENT ARE SOLELY THOSE OF THE D. E. SHAW GROUP AS OF THE DATE OF THIS DOCUMENT, ARE SUBJECT TO CHANGE WITHOUT NOTICE, AND MAY NOT REFLECT THE CRITERIA EMPLOYED BY ANY PERSON OR ENTITY IN THE D. E. SHAW GROUP TO EVALUATE INVESTMENTS OR INVESTMENT STRATEGIES. SIMILARLY, THE INFORMATION CONTAINED IN THIS DOCUMENT IS PRESENTED SOLELY WITH RESPECT TO THE DATE OF THIS DOCUMENT (UNLESS OTHERWISE INDICATED) AND MAY BE CHANGED OR UPDATED AT ANY TIME WITHOUT NOTICE TO ANY OF THE RECIPIENTS OF THIS DOCUMENT. THE INFORMATION CONTAINED IN THIS DOCUMENT HAS BEEN DEVELOPED BY THE D. E. SHAW GROUP AND/OR OBTAINED FROM SOURCES BELIEVED TO BE RELIABLE; HOWEVER, THE D. E. SHAW GROUP DOES NOT GUARANTEE THE ACCURACY, ADEQUACY, OR COMPLETENESS OF SUCH INFORMATION. FURTHER, THIS DOCUMENT CONTAINS PROJECTIONS AND OTHER FORWARD-LOOKING STATEMENTS REGARDING FUTURE EVENTS, TARGETS, OR EXPECTATIONS. SUCH STATEMENTS ARE BASED IN PART ON CURRENT MARKET CONDITIONS, WHICH WILL FLUCTUATE AND MAY BE SUPERSEDED BY SUBSEQUENT MARKET EVENTS OR OTHER FACTORS. HISTORICAL MARKET TRENDS ARE NOT RELIABLE INDICATORS OF FUTURE MARKET BEHAVIOR OR THE FUTURE PERFORMANCE OF ANY PARTICULAR INVESTMENT AND SHOULD NOT BE RELIED UPON AS SUCH.

MORE GENERALLY, NO ASSURANCES CAN BE GIVEN THAT ANY AIMS, ASSUMPTIONS, EXPECTATIONS, AND/OR OBJECTIVES DESCRIBED IN THIS DOCUMENT WILL BE REALIZED. NONE OF THE ENTITIES IN THE D. E. SHAW GROUP; NOR ANY OF THEIR RESPECTIVE AFFILIATES; NOR ANY SHAREHOLDERS, PARTNERS, MEMBERS, MANAGERS, DIRECTORS, PRINCIPALS, PERSONNEL, TRUSTEES, OR AGENTS OF ANY OF THE FOREGOING SHALL BE LIABLE FOR ANY ERRORS (AS A RESULT OF NEGLIGENCE OR OTHERWISE, TO THE FULLEST EXTENT PERMITTED BY LAW IN THE ABSENCE OF FRAUD) IN THE PRODUCTION OR CONTENTS OF THIS DOCUMENT, OR FOR THE CONSEQUENCES OF RELYING ON SUCH CONTENTS.

NEITHER THIS DOCUMENT NOR ANY PART OF THIS DOCUMENT MAY BE REPRODUCED OR DISTRIBUTED WITHOUT THE PRIOR WRITTEN AUTHORIZATION OF THE D. E. SHAW GROUP.

COPYRIGHT © 2023 D. E. SHAW & CO., L.P. ALL RIGHTS RESERVED.



D E Shaw & Co

www.deshaw.com