Spectrum

Where code meets capital

Q3 2025

Federated Hermes Limited

www.hermes-investment.com
For professional investors only

Introduction

In the changing world of finance, where instinct, emotion, and experience have all-too-often held sway, quantitative investing has emerged as both a disruptor and a guidepost since its conception in the mid to late twentieth century. It's a field of investing where data doesn't just support decision-making – it drives it.

As we stand at the crossroads of a tech revolution powered by artificial intelligence (AI), the role of a quant investor has evolved from being a quiet force behind the scenes to taking centre stage.

In this paper, the MDT team at Federated Hermes trace the origins of quantitative investing, exploring the rise of systematic strategies and the core principles that have stood the test of time – chief among them, the disciplined search for return-driving factors. The team explain why a forward-thinking approach, which considers factors like company age, economic moats, and other structural characteristics, presents advantages over traditional equity strategies, and what sets their approach apart.

We hope this report offers insight and inspiration for investors who recognise the exciting new possibilities data can unlock.

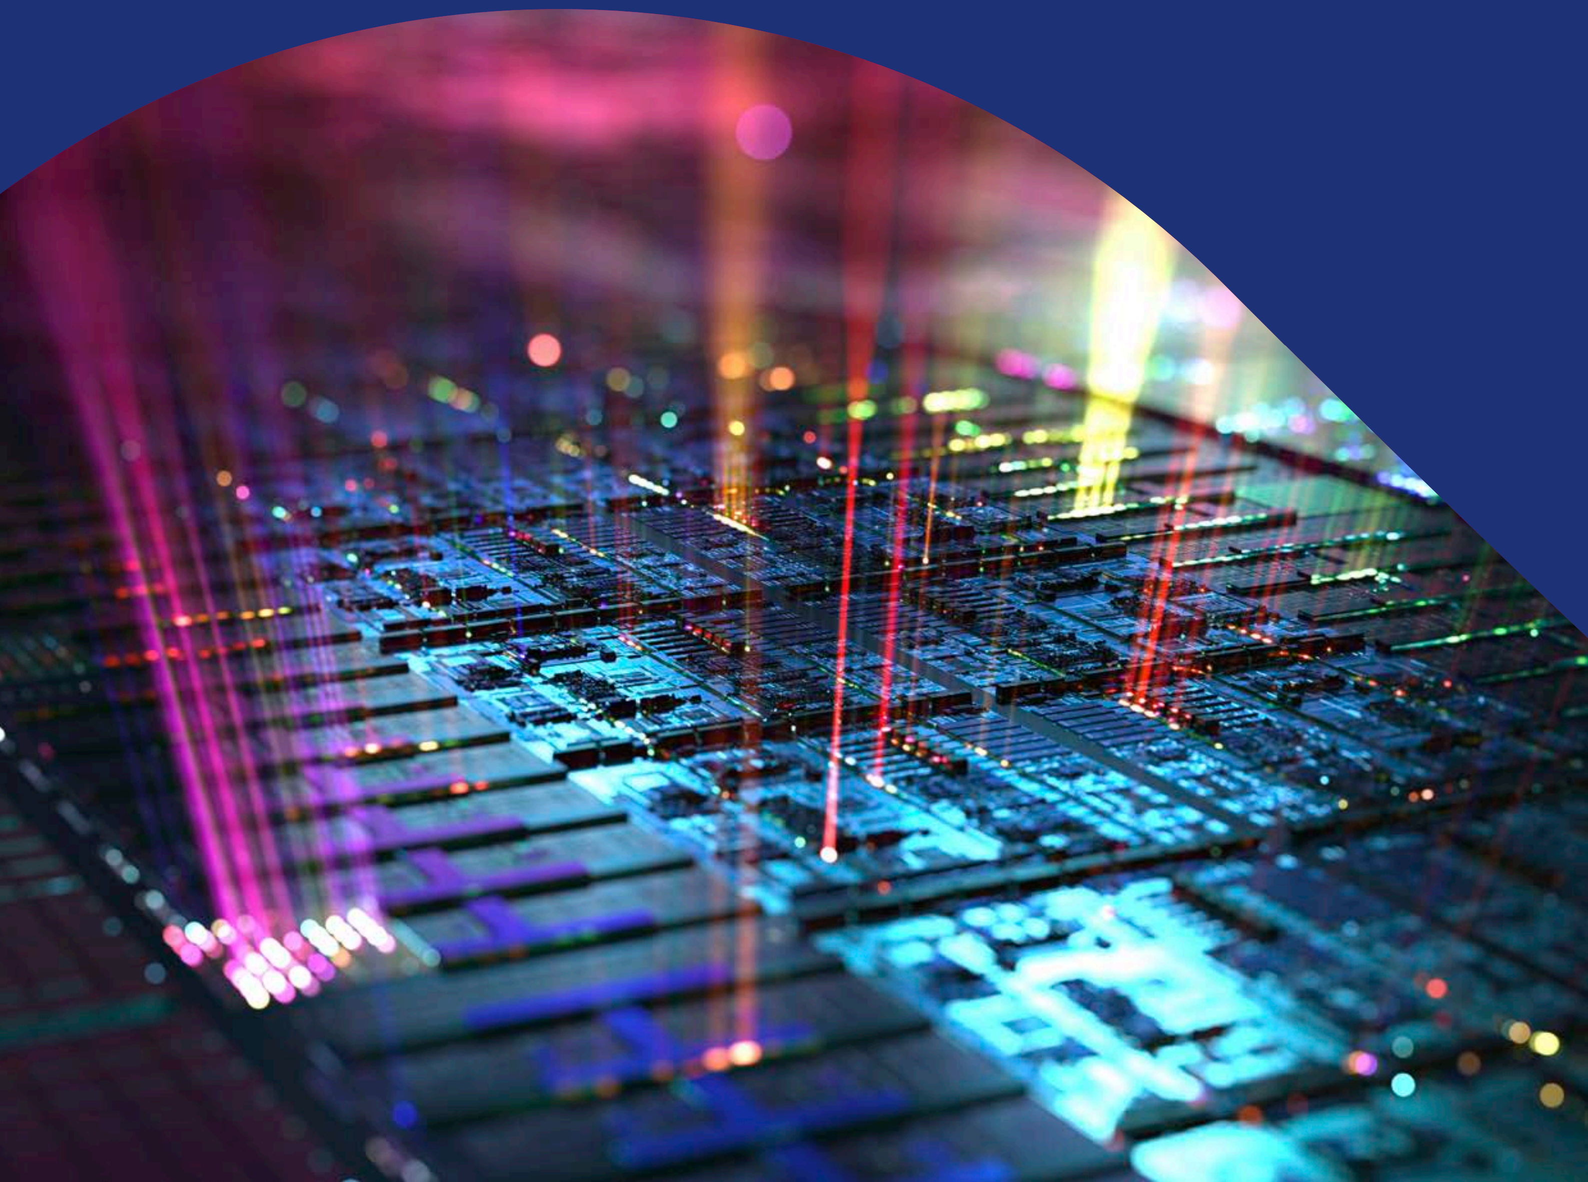

Contents

Part one:
Setting the scene

The history of quant

Quantitative investing (quant) encompasses a broad range of strategies that use data analysis, mathematical modelling and automated transactions.

Over more than a century, it has evolved from a purely theoretical concept to a practical approach to investing in financial markets. Ideas that were once confined to the academic world have been implemented by numerous investment strategies, often with remarkable success. Along the way, there have been some high-profile failures too. This has led to a degree of scepticism and even cynicism towards quant strategies.

But advances in computing power and an extraordinary abundance of data allow today's quant managers to achieve insights that were previously unimaginable. Each new day provides quant processes with a wealth of new data. And as new datasets achieve sufficient maturity to offer genuine predictive power, quant investing is making extraordinary strides in its reach and scope.

Each new day provides quant processes with a wealth of new data.

A selective chronology of quant

- **1900** Louis Bachelier's *Theory of Speculation*
- **1952** Harry Markowitz's *Portfolio Selection*
- **1960s** Development of the capital asset pricing model (CAPM)
- **1966** Victor Niederhoffer's *Market Making and Reversal on the Stock Exchange*
- **1969** Edward O. Thorp launches Convertible Hedge Associates
- **1970s** Introduction of computerized trading to the New York Stock Exchange
- **1973** Fischer Black and Myron Scholes's *The Pricing of Options and Corporate Liabilities*
- **1976** Stephen A Ross's *The Arbitrage Theory of Capital Asset Pricing*
- **1980** Victor Niederhoffer launches the NCZ Commodities fund
- **1982** Founding of Renaissance Technologies
- **1984** Breiman et al's *Classification and Regression Trees* (CART)
- **Mid-80s** Major investment banks set up quant desks
- **1988** Founding of D.E. Shaw
- **1992** Eugene Fama and Kenneth French's three-factor model
- **1998** Collapse of Long-Term Capital Management
- **2007** The 'quant quake'
- **2008** The global financial crisis
- **2000s** Ongoing revolutions in computing power, data storage and algorithms
- **2010s** Increasing use of machine learning
- **2020s** AI begins to come into its own, offering ever great capacity to effectively interpret datasets.

Federated Hermes MDT Advisers

What makes us different?

The Federated Hermes MDT Advisers (MDT) investment approach is armed with an analytical edge, evaluating nearly every stock in the US equity market using the unique combination of intuitive, fundamental and technical factors that we believe are most important to forecasting each stock's future performance.

Our research-intensive approach is differentiated from other quantitative equity strategies, and we believe the modelling techniques and objective, unemotional nature of our investment process presents advantages over traditional fundamental equity strategies.

- **Differentiated alpha engine:** Using sophisticated predictive modelling, including decades of experience applying machine learning techniques to forecast stock returns, we seek to identify the most relevant factor combinations when evaluating the alpha potential of each company. This can provide greater breadth of alpha sources to build portfolios with, creating the potential to outperform in various market environments and ultimately deliver more consistent active performance.
- **Dynamic, risk-managed process:** Stock forecasts and portfolio positions are updated daily, enabling the MDT strategies to adapt and take advantage of timely market opportunities and help to ensure portfolios reflect our strongest, most current alpha ideas. Highly diversified portfolios are built to limit unintended risks and focus active positions in diverse alpha sources. We believe this can lead to stronger portfolio resilience through the market cycle and ultimately improved risk-adjusted returns over time.
- **Strong performance results:** Historical outperformance achieved independent of market direction or any style leadership, has led to durable, all-weather performance outcomes driven by skill-based alpha.

Federated Hermes MDT investment team

MDT manages more than **US$23bn**[1] across several US equity strategies, including nine long-only and one equity market-neutral strategy. Certain strategies are available in other vehicles, including mutual funds, SMAs, ETFs and a CIT.

The MDT Investment Team is purposely structured to seek excellence in three areas: research, portfolio construction/trading, and process analytics, all of which we view as essential to our ability to pursue compelling performance outcomes for our clients.

Unlike other equity teams that may have members covering different markets and models, every member of our team is focused on one market (US equities) and one model that drives the process for all our portfolios. This focus has led to a highly collaborative culture where everyone is pulling in the same direction – leading to efficient management of our investment process.

> **The MDT Investment Team is purposely structured to seek excellence in three areas: research, portfolio construction/trading, and process analytics**

[1] As of August 2025.

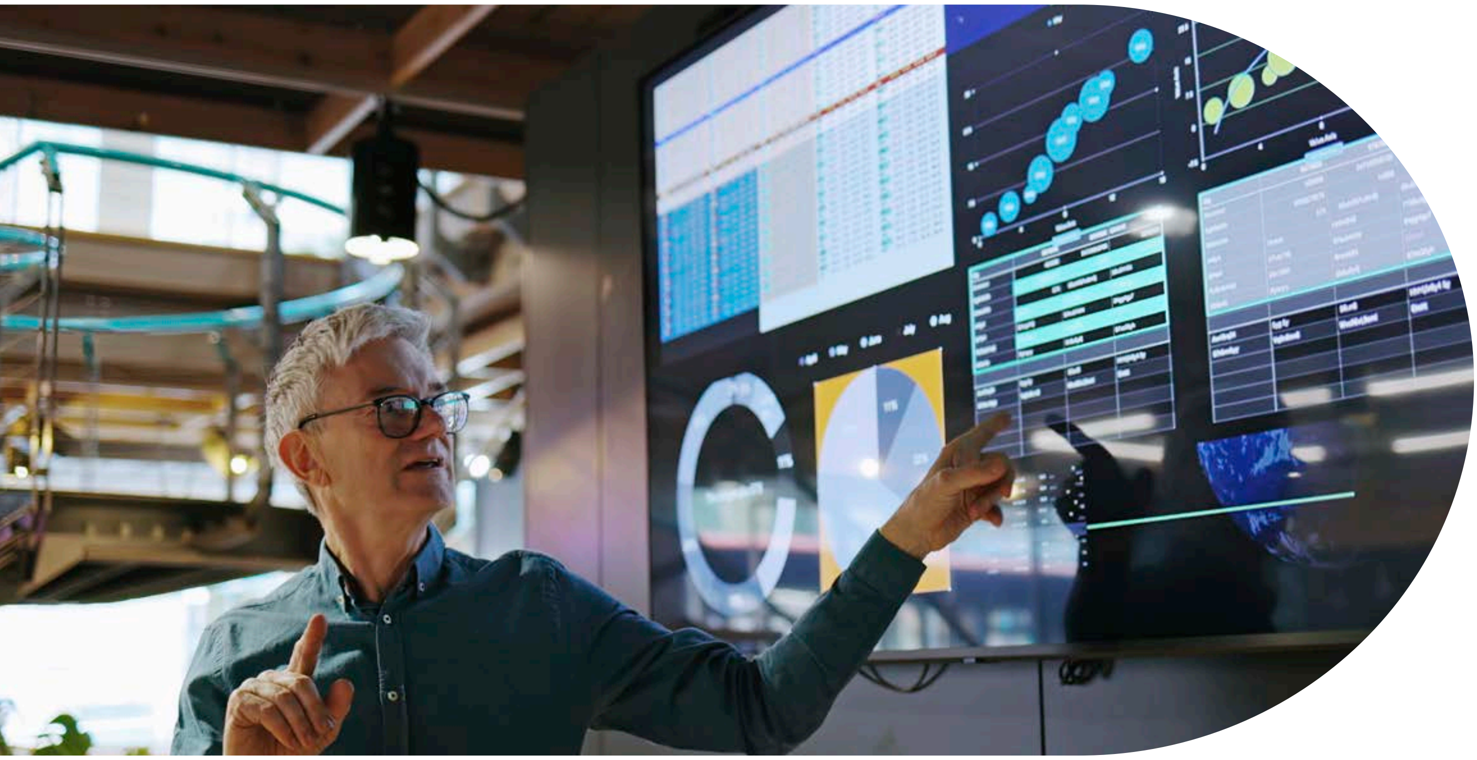

The AI revolution: How machine learning has transformed everything

Key takeaways

- Teaching machines how to 'think' has been a goal since the early days of computing, and recent advances in generative AI (artificial intelligence) have shown how close we may be to achieving it.
- ChatGPT and the interest it spawned is a recent development in the broader field of artificial intelligence, which has had a long and colourful history.
- Generative AI has substantial promise, but the current generation of models poses significant challenges for use in investing.
- As this field evolves, MDT Advisers will seek to identify ideas worth importing from AI research to improve our investment process.

As a leader in the use of machine learning in the investment field, MDT Advisers (MDT) has been a quantitative investing-practitioner since the 1980s and a proponent of technology in stock selection and portfolio construction. The recent-excitement around generative AI applications like ChatGPT has led to many questions about our thoughts on AI in-general and its application in investment management.

Teaching machines how to 'think' has been a tantalising goal since the early days of computing, and recent advances in generative AI have brought us closer than ever to achieving it. By leveraging vast datasets and computing power, researchers have developed models that enable computers to have conversations, draw pictures and perform other human-like tasks that were thought impossible even a few years ago.

Generative AI did not appear out of nowhere. It's the latest step in the broader field of AI, which has had a long and colourful history of progress and setbacks. Some early successes in the lab failed to translate into practical applications. Some promising ideas led to dead ends. Some seemingly dead ends, such as the humble Perceptron,[2] were later resurrected as core components of the neural networks that undergird modern generative AI models.

Many early attempts at AI centred around rules programmed by humans. These approaches ran up against the difficulty of managing the complex tangle of rules that inevitably result from trying to navigate the nuances of the real world.

What has instead become the dominant approach is to have a machine learn by distilling the rules from large datasets containing inputs and desired outcomes. The result of that learning can be encoded as an equation, a forest of decision trees, a neural network or any other flexible model. The better the model is at capturing the true nature of the rules at play, the more useful that model is.

Generative AI is the result of this latter approach taken to the extreme. With more data, more computational power and larger models than ever, we have not yet reached the limits of this approach.

But even as generative AI provides the most compelling proof of the effectiveness of machine learning, it also illustrates some of its pitfalls. As generative AI upends our notions of what a machine can do, it's worthwhile to explore how it could be used to enhance our own field of investment management. At MDT, we have a long history of incorporating machine learning innovations while seeking to avoid the pitfalls.

What is artificial intelligence (AI)?

- AI is a field of computer science that covers a wide range of algorithms and approaches, with many subfields.
- It's focused on the development of machines that mimic functions associated with the human mind and that can perform human tasks such as understanding speech, playing games and driving cars.
- Colloquially, AI is often used as a nebulous term that encompasses big data, machine learning or artificial, neural networks.

The generative AI revolution

Much of the current buzz around AI stems from the release of ChatGPT in 2022. Its remarkable ability to converse and respond to a wide array of topics with seemingly human-like intelligence provoked questions from practicality to ethics. Can it do my homework? Can it take over my job? Does it exhibit consciousness?

Under the hood, the original ChatGPT used a variant of a large language model (LLM) called GPT-3.5, which was trained on a dataset containing hundreds of billions of words collected from the internet. Perhaps surprisingly, given its apparent capabilities, the model is trained only to predict the next word in a sentence as accurately as it can. Some have even called it a glorified autocomplete. Though complex in the details, the training itself amounts to having the model guess the next word, checking how big of a mistake it made, and adjusting some of its billions of model parameters ('neurons') so it makes a smaller mistake in the future.

Under the hood, the original ChatGPT used a variant of a large language model (LLM) called GPT-3.5, which was trained on a dataset containing hundreds of billions of words collected from the internet.

[2] A perceptron is a fundamental building block of artificial neural networks, inspired by the structure and function of biological neurons. It's a single-layer neural network that performs binary classification by learning a linear decision boundary.

Yet this contrast is emblematic of the success of machine learning's data-driven approach. The simple high-level technique of applying massive amounts of data to train a very large model made of simple individual components can lead to impressive results – generally, the bigger the model, the better the result – if there is enough data.

Just as a model can be trained to predict the next word, so too can it be trained to predict the next sound, the next pixel or some other combination, leading to audio, image and multi-modal generative models. These models are also improving, and a new generation of open-source models now rivals even the best proprietary implementations.

One particular strength of generative LLMs is their ability to summarise long blocks of text, which could help investors parse lengthy regulatory filings for critical bits of information. A less obvious but more intriguing possibility is for the model to distil that summary into a rating or other quantitative metric.

For example, suppose we'd like to assign a sentiment score to a news article about a company. Historically, such tasks have used rudimentary machine learning models or counted the number of positive- and negative-sounding words. Now these tasks can be accomplished by telling the LLM to read the relevant article and then asking it to respond with a number for sentiment. Once you know what you want to see, you have the potential to use generative AI to unlock that information from regulatory filings, earnings calls, satellite photographs and other kinds of unstructured data. While this has a lot of promise, there are problems as well.

The 'but'…

One only needs to read a few passages of AI-written content to notice that, in most cases, it's a simulacrum. When asked to write an essay, the program may line the words up, but the text may lack humour, nuance or deeper insight. It can struggle to solve simple math problems. Sometimes the model will even come up with an excuse to avoid doing what you asked. Missing the mark can be harmless and amusing if there's nothing at stake, but investing based on unverified AI input presents a real problem.

These AI systems can sometimes generate misleading or factually incorrect responses with complete conviction. Termed 'hallucinations', these responses continue to afflict large AI models in part because they're trained on vast swathes of human-generated data on the internet that is itself replete with inaccuracies presented as fact. In this regard, the models are remarkably human-like. While some adjustments can be made to the models after they're trained, it's impossible to manually check the entire training data to remove inaccurate information.

So, can AI help investors?

Given all of the above, the question of whether AI-generated signals can be used for investing becomes an empirical one. Can the errors made in hallucination or incorrect interpretation be outweighed by the value of the signals they extract?

Additionally, we need to be aware that the reasoning process that these AI models use to arrive at answers is a black box. You could ask a model how it arrived at an answer, but as the model is incapable of genuine introspection, the answer will still be based on a best guess of the likeliest next word. Thus, if an AI-generated signal indicates a counterintuitive trade, it will be difficult to understand the true chain of reasoning that led to that trade.

Finally, for systematic investors, back testing is an important tool that gives essential information about the effectiveness of a strategy. Generative AI models are trained on data up to the present day, and their enormous sizes mean that they can memorise a great deal of historical information, including market information. Using a signal generated from such a model risks look-ahead bias: contaminating simulated trading in a past time period with information from the future.

Using a signal generated from such a model risks look-ahead bias: contaminating simulated trading in a past time period with information from the future.

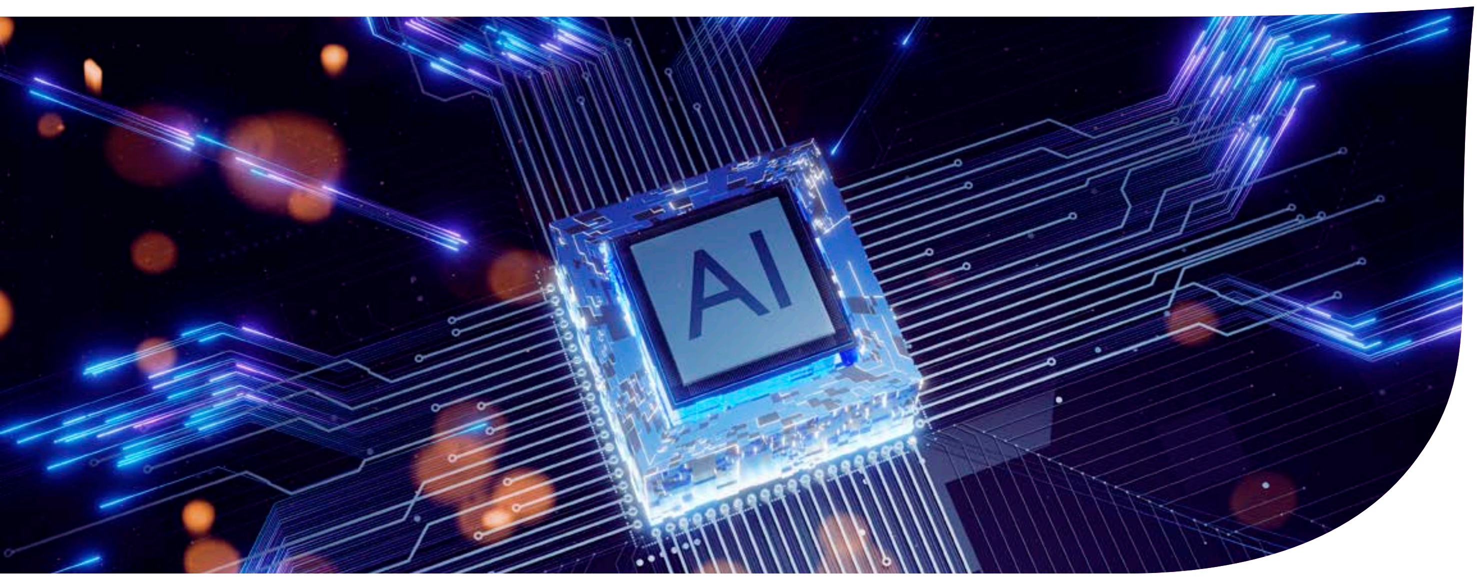

Part two:
Factor Focus

FACTOR FOCUS

Company age – How to measure it and why it matters

Key takeaways

- Academic research indicates that there is a relationship between a company's age and stock returns.
- Calculating company age is more subjective than it would seem.
- We found that the interaction of company age with other factors in our decision trees could yield significant benefits beyond the contributions of company age alone.

The research team at MDT Advisers casts a wide net to find new ideas to test as possible enhancements to our investment process. In 2019, we came across a paper entitled 'Age Matters',[3] primarily authored by a Ph.D. student in statistics at the University of Waterloo. The paper was not published and did not make much of a stir online (3.5 years later, it has been cited on Social Science Research Network (SSRN) only once). However, we were intrigued by the paper's finding that there was a relationship between company age and stock returns – particularly by the nature of that relationship. The report used standard regression tools to uncover the relationship and performed additional analysis to show that the effect was significant only among the younger half of firms. We hoped that by applying our decision tree modelling to this factor, highly differentiated from our other factors and non-linear in its nature, we would be able to make a significant improvement in the accuracy of our alpha forecasting.[4]

How to measure age

At first, we followed the paper's authors in simply using the pricing data from the Center for Research in Securities Pricing (CRSP) to determine age. Doing it this way makes age a function of how long the security has been traded on a major stock exchange. However, as we often find to be the case, the construction of a factor, even one that represents such a seemingly straightforward idea as a company's age, can be improved with craftsmanship. For example, for a company that goes bankrupt, delists, and later returns to the stock market, the CRSP dataset will split this into two separate securities. Is it right to treat the stock of a company that has emerged from bankruptcy as having the same age as the stock of a recent IPO? Similarly, should a company that emerged from a merger with a SPAC (Special Purpose Acquisition Company) be treated as older than a company with an IPO simply because the SPAC vehicle traded on the exchange for months, if not years, before consummating the merger?

Whether a company is "young" or "old" is relatively easy to define for most firms, and the paper's definition certainly suffices. However, for a non-trivial subset of companies, calculating age is a subjective exercise. Our company age factor utilises not only pricing data but also various pieces of information from a firm's financial statements as additional means to capture the essence of how old the company behind a stock listing truly is. We continue to look for ways to refine and enhance the factor as we come across examples in trading our portfolios where the calculation of age conflicts with our intuition.

Why age matters

Our research generally agreed with the paper's findings that company age does have a relationship with future returns, at least within the younger cohort of companies. The paper's authors, interestingly, did not perform the standard asset pricing model tests in presenting their results. Instead, they offered some evidence that at a minimum, the size factor was not responsible for their findings. However, it is not unreasonable to suspect that other known factors could explain some of the "age effect" found in the paper. For example, certain measures of value are correlated with company age. So, while we were pleased to see directionally similar results to the paper, it was not surprising to find that the "age effect" was weaker in a multi-factor framework. What we hadn't anticipated was that the interaction of company age and our other factors in the context of our decision trees would yield significant benefits from an alpha-modelling perspective – much greater than the contributions of company age on a standalone basis. One of the powerful features of using a decision tree for stock picking is that not all companies get scored the same way. The algorithm figures out the most important questions to ask of a particular type of company (and, conversely, it determines what questions are not essential to ask). By integrating company age into our factor lineup, we have given the trees a mechanism to discover that certain factors are more important for younger companies (e.g., price and analyst-based sentiment) and other factors are more important for older companies (generally speaking, value and quality measures).

> **Our research generally agreed with the paper's findings that company age does have a relationship with future returns, at least within the younger cohort of companies.**

[3] Guo, Danqiao & Boyle, Phelim & Weng, Chengguo & Wirjanto, Tony, 2019. "Age matters," MPRA Paper 93653, University Library of Munich, Germany, revised 01 May 2019.

[4] As part of our ongoing research, in 2023 we discovered a previously published work discussing the company age factor which was not cited in the 2019 research paper. We want to note the work done by that author. Zhang, X. Frank, (April 20, 2004), "Information Uncertainty and Stock Returns".

Figure 1: Applying company age in a regression tree

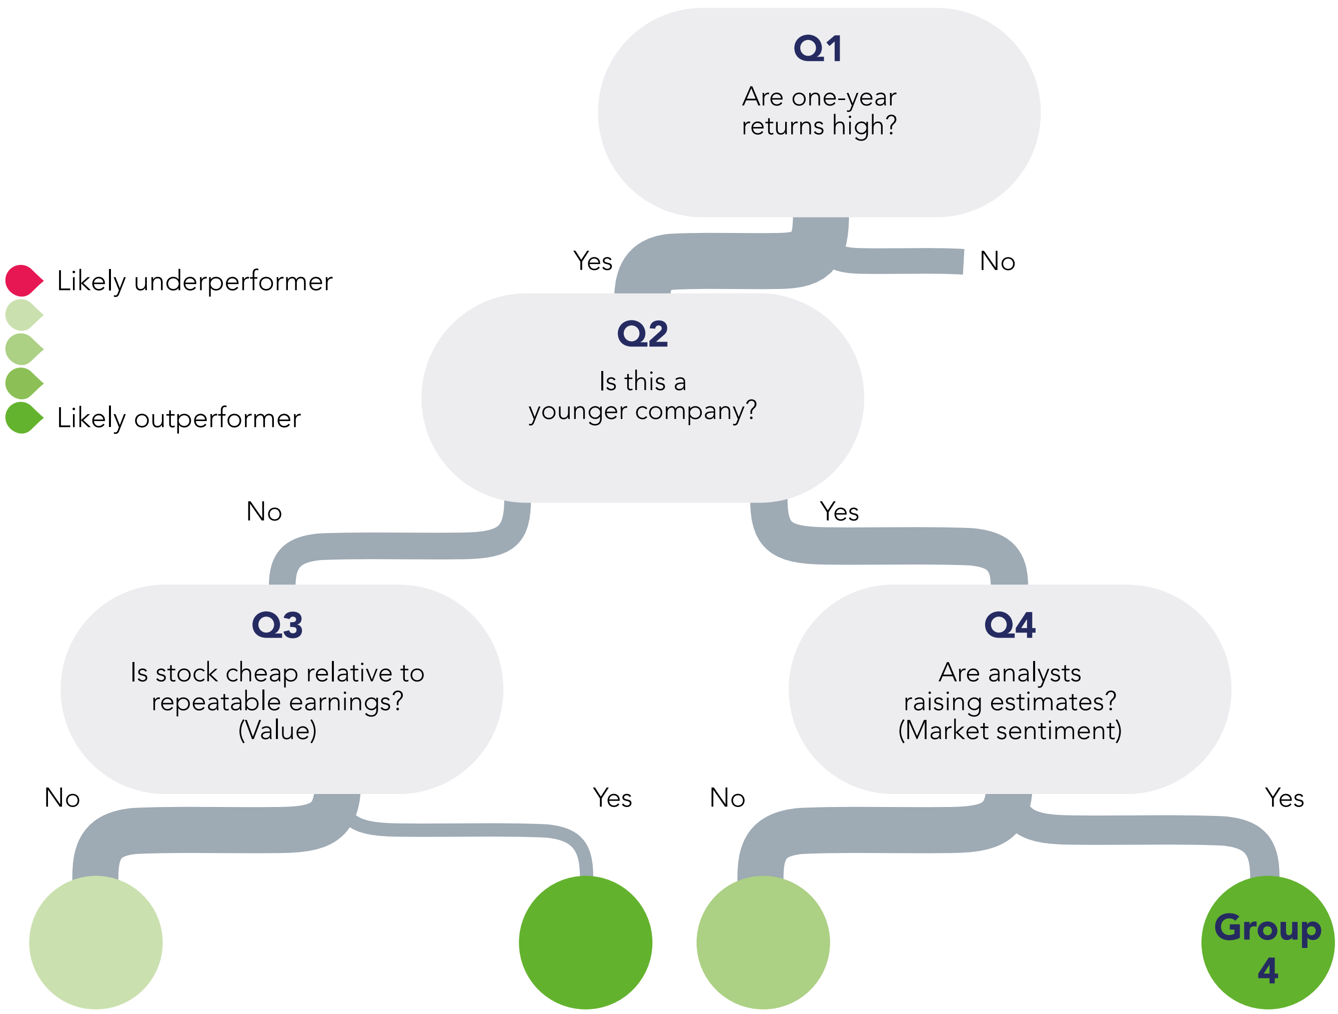

This partial tree shows how company age fits within a more extensive regression tree analysis. After finding a group of companies with high one-year returns (the "Yes" path out of Q1), the best question to ask those companies concerns company age (Q2). After answering the age question, the best question for older companies is about value (Q3), while the best question for younger companies is different, about analyst-based sentiment (Q4). The diagram shows how the company age factor is meaningful enough that the successive questions for young and old companies are very different, as well as showing that answers higher on the tree determine a set of subsequent questions tailored to a company's own characteristics. We have found that inclusion of the company age factor into our investment process significantly improves the forecasting accuracy of our alpha model and the simulated returns of our model backtests.

Conclusion

Our ongoing search for ways to improve stock selection can lead to unusual places. Company age may seem unlikely to have predictive value on its own or in a multi-factor framework, however its interaction with other factors in our regression tree analysis has increased the trees' predictive power. Our decision tree model continues to show that, over time, some factors are more relevant to certain companies than others, and an unusual factor like company age may become more valuable inside a forest of decision trees. We will continue to evaluate new factors and enhance others in order to try to unlock the predictive powers of our model.

FACTOR FOCUS

Bringing economic moats to quantitative investing

Key takeaways

- Channelling the results of successful investors like Warren Buffett through a quantitative process has been problematic.
- Popularised by Buffet, companies with a wide "economic moat" that can help reduce competitive threats is appealing as an enhancement to traditional value investing.
- In 2023, Federated Hermes MDT began incorporating an industry moat factor to help identify companies whose out-of-favour status is likely only temporary.
- We believe this factor can help us identify companies that are more likely to overcome negative sentiment/ momentum and potentially generate strong future returns.

Value investing in the US has gone through a rough stretch in recent years. Over the five years ending September 30, 2024, the Russell 3000® Growth Index[5] outperformed the Russell 3000® Value Index[6] by 8.48% annually. Tech giants with generally rich valuations continue to dominate the list of the largest stocks in the US by market capitalisation and have driven these returns. The largest stock that most investors would consider to be a value stock has only about one-third the market capitalisation of the largest growth stocks. Interestingly, that value stock is none other than Berkshire Hathaway, the conglomerate run by Warren Buffett, perhaps the world's best-known value investor. But even Buffett has long moved on from buying fair companies at a wonderful price to buying wonderful companies at a fair price.

What makes a company wonderful? An investing concept widely associated with Buffett – though he would credit his longtime friend and business partner, Charlie Munger – is to find companies with a wide economic moat, where a company has developed advantages that allow it to defend its profitability against encroaching competition. Brand identity (created through advertising and marketing spending) and patents (created through research and development spending) are two ways companies typically accomplish this. Coca-Cola, which Berkshire Hathaway has had a long-standing stake in, is a notable example of the former. Patent protection for pharmaceutical companies, an example of the latter, enables them to invest in the costly endeavour of developing new drugs so that competitors cannot imitate them.

One reason to suspect that wide moat firms might be mispriced in the market is that standard accounting treatment generally expenses "moat building" in the current period versus capitalising the cost over a period of years. This lowers a firm's income and thus appears at first glance as value "destruction," not value creation. A more enlightened accounting treatment appreciates that such spending can create lasting value for the firm and possibly improve profitability for many years. Just think of all the decades-old advertising pitches and jingles the average consumer has rattling around in the back of their mind!

The challenge: Defining a moat

With this in mind, in recent years many systematic investors have tried to improve standard value factors by adjusting them to incorporate the ongoing value of economic moat-oriented spending. However, a lack of consensus indicates the elusiveness of replicating Buffett.[7,8]

MDT has favoured a different approach. Because we use decision trees in our stock-picking, we can introduce an economic moat as a standalone factor and let the data drive the specific contexts where using economic moats can help improve our investment decision-making. The premise of this process is that certain factors – or more importantly, combinations of factors – will be more relevant to certain company types than other factors. In September 2023, we added an economic moat factor to our model that seeks to capitalise on a company's spending on potential moat-building activities. We aggregate estimated moats across a company's industry to smooth out inequalities in individual company data reporting, and call the resulting factor "industry moat." When added to our stock selection models, we found industry moat to be an additive source of excess returns.

The industry moat factor can help us most in improving stock selection among highly out-of-favour stocks – which are not precisely the same as value stocks, although there tends to be a substantial overlap. Our research has indicated that the stronger return potential among those stocks is associated with companies having a wider economic moat. In other words, companies in narrow-moat industries that offer a commoditised good or service with low value-added features (e.g., banks, airlines, mining) are less likely to rebound from a negative shock than those in wide-moat industries (e.g., software, retailing, pharmaceuticals).

[5] Russell 3000® Growth Index: Measures the performance of those Russell 3000® companies with higher price-to-book ratios and higher forecasted growth values.

[6] Russell 3000® Value Index: Measures the performance of those Russell 3000® Index companies with lower price-to-book ratios and lower forecasted growth values.

[7] Feifei Li, Intangibles: The Missing Ingredient in Book Value (April 29, 2021). Available at SSRN: https://ssrn.com/abstract=3686595 or http://dx.doi.org/10.2139/ssrn.3686595.

[8] Savina Rizova and Namiko Saito, Internally Developed Intangibles and Expected Stock Returns (July 27, 2021). Available at SSRN: https://ssrn.com/abstract=3697452 or http://dx.doi.org/10.2139/ssrn.3697452.

Applying the concept

We believe the industry moat factor works particularly well within our decision tree process. Companies with high market sentiment, those that have beaten earnings estimates quarter after quarter or those with other strong quality characteristics are typically already well appreciated by the market. So, the presence (or lack) of a wide economic moat tends not to have significantly impacted expected returns for those companies.

However, even successful companies can fall out of favour for various reasons, some within their control, but others beyond their control. An established moat can help out-of-favour companies keep competitors at bay while re-establishing operating results. Thus, we have begun to use this factor to help identify those companies that may be more likely to overcome negative sentiment/momentum to return to in-favour status, seeking to avoid companies whose stock prices continue to tailspin. To paraphrase Buffett, we believe it helps us identify wonderful companies at wonderful prices.

> **To paraphrase Buffett, we believe it helps us identify wonderful companies at wonderful prices.**

The quantitative advantage

While Warren Buffett comes across as the furthest thing from a quantitative investor, there is no doubt that there are systematic components to his investing process, even as his concept of value remains challenging to define. With Federated Hermes MDT's decision-tree-based quantitative process, portfolios can gain exposure to the concept of value without a fixed exposure to the systematic value factor, where investors' portfolios will tend to be overweight all cheap stocks and underweight all expensive ones.

More importantly, decision trees can find opportunities within particular combinations of factors that simple factor tilts cannot. Every investor should be looking for the key that unlocks value. With a nod toward Omaha, we think the industry moat factor has helped us unlock value in certain overlooked companies that are temporarily out of favour.

FACTOR FOCUS

The short term is not the enemy of the long term

Key takeaways

- While long-term fundamentals are crucial, information may potentially be gleaned from short-term market trends to address near-term performance.
- We have found that incorporating factors like price trends and market sentiment alongside fundamentals-based metrics may improve investment performance over time.

"In the short run, the market is a voting machine, but in the long run it is a weighing machine." – Benjamin Graham

Investors seeking to determine a security's 'true' long-term value and invest in those priced below that value can employ various methods. Some may build dividend discounting models or apply multiples to earnings forecasts. Others may come from the top down, estimating total addressable markets, market shares and profit margins. One commonality these investors share is that there will inevitably be periods when their methods appear to malfunction.

As a result, one classical way to interpret Graham's words above (and presumably the one intended by Graham, widely considered the father of value investing) is essentially: Have confidence, ye value investors! The short-term beauty contest that is the stock market will often overlook those rugged, less universally loved stocks in your portfolio. Still, over the long run, the market's weighing machine can come around to recognise the value in those stocks' fundamentals.

And there is a reason this advice has stood the test of time – generally, the worst time to abandon a well-reasoned investment process is when the voting machine seems to have gone haywire and is voting for all the "wrong" kinds of stocks.

But what if – rather than looking at the voting machine as just a source of noise in the markets (in effect, as the enemy of the weighing machine) – we read Graham's words as prescriptive: as a reminder not to ignore the short run while waiting for the long run to arrive. Perhaps a portfolio could have better outcomes by harnessing the voting machine when it favoured some of the stocks and the weighing machine when it favoured others. In 2015, several strategies at Federated Hermes MDT (MDT) underperformed their benchmarks. A key factor in the outcome was our decision to underweight certain fast-growing but expensive tech stocks while overweighting companies whose businesses, in retrospect, were on a collision course with those same stocks. The companies being disrupted seemed to offer strong value based on a variety of metrics. As they underperformed during the year, many seemed to represent even better value as their stock prices retreated. The dominance of a few tech companies was such a prominent feature of the stock market in 2015 that the media coined a nickname for those exciting growth stocks – FANG (Facebook, Amazon, Netflix, Google).

> What if we read Graham's words as a reminder not to ignore the short run while waiting for the long run to arrive? Perhaps a portfolio could have better outcomes if the voting machine favoured some of the stocks while the weighing machine favoured others.

It seems clear in hindsight that had we locked into the modelling that drove that negative outcome in 2015, our investors would have been in for a much bumpier ride over the ensuing years, as the FANG companies and others like them have continued to outperform and the acronym evolve. Fortunately, our ongoing research efforts, driven by the desire to determine whether investing in companies with qualities like the FANGs could improve portfolio results over the long run, helped yield a better outcome.

Since 2015, MDT has added a variety of purely price-based, or technical, factors to our stock selection models. The added factors consider price trends over a wide range of intervals, from as short as a few months to as long as half a decade. Some detect momentum-type effects, where companies that have done well over a particular horizon have continued to perform well (and vice-versa). Others detect reversal effects, where the largest losers may become the biggest winners (and vice-versa). What they all have in common, of course, is that they are pure expressions of market sentiment – results tabulated by the voting machine.

Adding price-based factors to our existing lineup of fundamentals-oriented factors in MDT's decision-tree-based stock forecasting model has uncovered interesting relationships. We expected that certain companies with substantial value and/or quality characteristics, combined with good technicals, would, on average, produce appealing outcomes, which turned out to be true.

However, we also encountered favourable outcomes from companies with strong fundamental characteristics and extremely weak recent performance. In the same vein, we found that the outcomes can be favourable over the short- to medium-term for certain companies with very strong technical characteristics, even without validation from most of our value and quality characteristics.

A benefit of managing a portfolio of investments is that diversification may help to improve risk-adjusted outcomes. Adding price-based factors to our investment process has not only helped us find some new and interesting types of stocks in which to invest but, perhaps even more importantly, from a portfolio construction point of view, some of those new opportunities are significantly differentiated from what we were previously able to find. In the many years since 2015, our portfolios have benefitted from stock contributions because the voting machine loved them or because they were extremely out of favour on our price-based metrics. We still find plenty to like about value and quality characteristics – an investment strategy with only voting machine factors can likewise be enhanced by adding the weighing machine.

Like active and passive investment strategies, there is no reason short-term technical and long-term fundamental factors cannot co-exist. There can be complementary benefits. We see no contradiction and no reason to choose exclusively between the short and long run. In our opinion, the best response to Ben Graham's observation on the behaviour of markets is this: "Why not both?"

FACTOR FOCUS

Two ways to be wrong in equity portfolio management (and how to help mitigate them)

Key takeaways

- Equity portfolio management often involves tempering the optimism needed to be successful with humility and risk management.
- Investors should consider adverse outcomes, in addition to the positive ones, as no one can perfectly predict all factors affecting securities prices.
- Effective risk management involves, among other things, diversification, appropriate bet sizing and avoiding unintended bets.
- We believe diversifying risk exposure in a portfolio – trying to leverage multiple, differentiated risk premia – can be a powerful tool for improving risk-adjusted returns.

Managing stock portfolios is generally a business for optimists. Over time, economies grow, stock markets tend to go up and taking risks in both the business world and in markets should be rewarded. Because a stock's price stops going down when it reaches US$0, and upside price potential has no similar limitation, stock investors considering a particular investment may often ask, "What could go right?" before asking, "What could go wrong?" When thinking about a portfolio of stock investments, though, humility is important, hence the need to give the latter question meaningful consideration. Here, we look at portfolio risk management through the lens of two different frameworks of negative outcomes in the stock-picking process and discuss methods of potentially mitigating those outcomes at a portfolio level.

#1 – Being wrong

Investments in all but the safest securities are inherently risky. No investor can perfectly chart the factors affecting securities prices, such as interest rates, energy prices, corporate regulation or geopolitics. At the security level, uncertainty surrounds the prospects for individual companies' product launches, the emergence of future competition or shifts in customer preferences. When evaluating historical data, is an identified pattern something that can be relied upon to repeat in the future, or is it merely a statistical artifact unlikely to lead to future profitable investment decisions?

Uncertainty is what makes markets. If it were easy to predict all the potential impacts on the value of a security, then investors would quickly agree on a security's value and volatility would disappear from the market for that security—a truly efficient market. Fortunately for investors, many portfolio managers are aware that the thesis behind a particular investment may not play out. Many also come well-armed to deal with some of the uncertainty surrounding the securities in their portfolios.

Mitigating risk

Outlined below are some of these tools, familiar to anyone involved in risk-taking activities:

- **Diversification** – The adage against putting all your eggs in one basket. If there is risk that any particular investment may not work out the way an investor hopes, a prudent approach is to have a portfolio of assets with positive expected but independent outcomes.
- **Bet sizing** – The size of a bet, all else equal, should be inverse to the uncertainty surrounding that bet. From a portfolio perspective, managers should take smaller positions in securities where they believe the potential variance of returns is relatively greater.
- **Avoiding unintended bets** – Trying to minimise exposure to uncertainty that is unrelated to one's investing edge. It would be undesirable for an investor with skill at picking stocks from the bottom up to end up with a portfolio where all the stocks are in the oil and gas production business. The future path of energy prices will probably be a more significant determinant of portfolio outcomes than the fortunes of the individual companies in the portfolio.

As portfolio managers, we are keenly aware that not every security we own will play out according to plan. Properly applied risk measures, such as those outlined above, can help reduce the individual impact of a poorly performing security on the overall portfolio.

#2 – Being right, but early

Adding a layer of complexity to the portfolio risk puzzle, portfolio managers must confront investment theses that may ultimately work out but move in the wrong direction in the near term. And as difficult as it may be for investment managers to stick with (or, better yet, to augment) an investment decision that has moved against them, a more significant challenge may lie in convincing current or prospective investors that the correct course is being followed. How can an outsider reliably distinguish whether their manager has experienced a loss of skill or has just gotten unlucky in the short run?

The tools mentioned above are helpful here but can be costly. For example, in the August 2007 meltdown in quant strategies, when many stocks inexplicably dropped 10% or more,[9] many investors reduced their positions due to higher perceived risk surrounding those stocks. When the liquidity event that caused the selloff subsided, and those stocks bounced back a few days later, investors that had reduced risk at unfavourable prices were left less than whole by the event. This is a good example of where the downside of being heavy-handed with the risk mitigation tools identified above can adversely affect returns. Reducing risk when positions are only temporarily out of favour may limit the potential upside from when they come back into favour.

Popular approaches for investment managers to mitigate the risk of investors abandoning out-of-favour strategies at unfavourable times involve compulsion—whether tangible, in the forms of capital lockups and gates, or intangible, in the form of persuasion. "No pain, no premium," goes a pithy saying to help investors stay the course through short-term underperformance. Left unsaid is how investors should know if there really is a light at the end of the tunnel. Not every investment strategy will deliver a satisfactory outcome in the long run. Despite the urgings from every corner, past underperformance will be a signal to some that there may not be a positive outcome ahead. Asset owners may approach the portfolio risk puzzle differently. In hiring multiple managers with different investing styles and time horizons, they know that when a particular manager's investments are out of favour ("right, but early"), other managers in the total portfolio could be in favour, balancing out any drag on overall portfolio outcomes. But even this approach has potential pitfalls.

First, trying to hire multiple skilled managers may have added costs and risks. Managers may not have perfectly consistent styles over time, so the hoped-for diversification benefits across managers may be lower than expected. Also, this approach may lead to sub-optimal capital deployment because the bets of multiple independent managers may inadvertently cancel each other out, leading to a more index-like portfolio at an active management fee level.

We believe there is an alternative that can help solve some of these issues. By using sophisticated optimisation and risk management techniques, a single manager with multiple diversified alpha-seeking engines can potentially benefit from a multimanager type of approach while reducing the frictions of utilising multiple independent sub-portfolios.

Federated Hermes MDT's approach

At Federated Hermes MDT Advisers, we have spent more than 30 years developing and refining our systematic process to picking stocks and building portfolios, with the goal of delivering alpha to our clients with as much consistency as possible.

We realise that avoiding bad outcomes in the investment business is not a job for risk controls alone, although they are a critical piece of the equation. We believe diversifying risk exposure in a portfolio—trying to leverage multiple, differentiated risk premia— can be a powerful tool for improving risk-adjusted returns.

As we discuss in our next article, in 2001, we discovered that employing a decision tree in stock picking can be a powerful means of seeking diverse alpha sources for portfolio construction. Decision tree algorithms search down every branch of the tree for the means to explain the best and worst potential outcomes within that branch. If a tree splits on value, then the algorithm tries to find the characteristics associated with not only the best and worst value stocks but also, separately, for those characteristics associated with the best and worst "not-value" (growth) stocks.

Over the past 20 years, understanding how to use these tools to help us pick stocks and build portfolios has evolved tremendously, but always with the same goal: the creation of highly diversified, resilient portfolios for our clients.

[9] Amir E. Khandani and Andrew W. Lo What Happened To The Quants In August 2007?: Evidence from Factors and Transactions 2008.

Part three:
The MDT Approach

THE MDT APPROACH

A deliberate approach to stock selection

Key takeaways

- Regression trees are a series of yes/no questions based on explanatory (independent) variables that lead to a prediction.
- Regression trees can be used to forecast individual stock performance versus the universe.
- After implementing decision trees in 2001 and monitoring their effectiveness, MDT at Federated Hermes identified key areas for improvement. We determined that multiple decision trees give us a more flexible decision-making framework.
- Bagging and boosting are two key methods we used to rectify the identified shortcomings of a one-tree model.
- In 2020 we implemented an algorithm that utilised both ideas: boosted random "forests".

Overview

The formal academic background for classification and regression trees came out of Stanford University research by Professor Leo Breiman (Jerome Friedman, Charles J. Stone, and R.A. Olshen), who published a book entitled "Classification and Regression Trees" in 1984. The methodology for building regression trees and the software provided through this research to create them became a standard tool in the insurance industry and in the physical sciences. The difference between regression trees and the more commonly known decision trees is simple: regression trees predict a number, while decision trees predict an endpoint state or "classification."

As we will outline below, MDT uses them in an effort to forecast how much a stock will outperform or underperform its universe.

Classification and regression trees (CART) background

Regression trees are a series of yes/no questions about a set of explanatory (dependent) variables, chosen using an algorithm, with the goal of predicting a target (independent) variable.

The algorithm figures out the best questions to ask at each point by rote. It tests every possible question (every explanatory variable, every point at which you can split the dataset into observations with a higher value and lower value) and chooses the one that yields two subgroups with different average values of the target measure and a minimum sum of squared errors versus those new averages. The algorithm uses the same process to choose the next best question for every subgroup.

Using CART in an effort to forecast stock alphas*

In 2000, MDT began a program of research based on the CART technique. We believed that CART might work well for selecting stocks, as we liked the non-linear nature of the analysis – the fact that a regression tree didn't allow a characteristic that wasn't important to a specific company to affect the outcome. We evaluated the technique with back tests and saw that the regression trees improved the results, so we added trees to our live strategies in 2001.

We provide an example tree below to show how the regression trees can be used to forecast alpha. This is a very small tree, but it illustrates the concepts and advantages of using a tree for this purpose. This illustration does not represent any of the regression trees in our strategies.

**Alpha in this document refers to a stock's excess return versus a strategy's universe.*

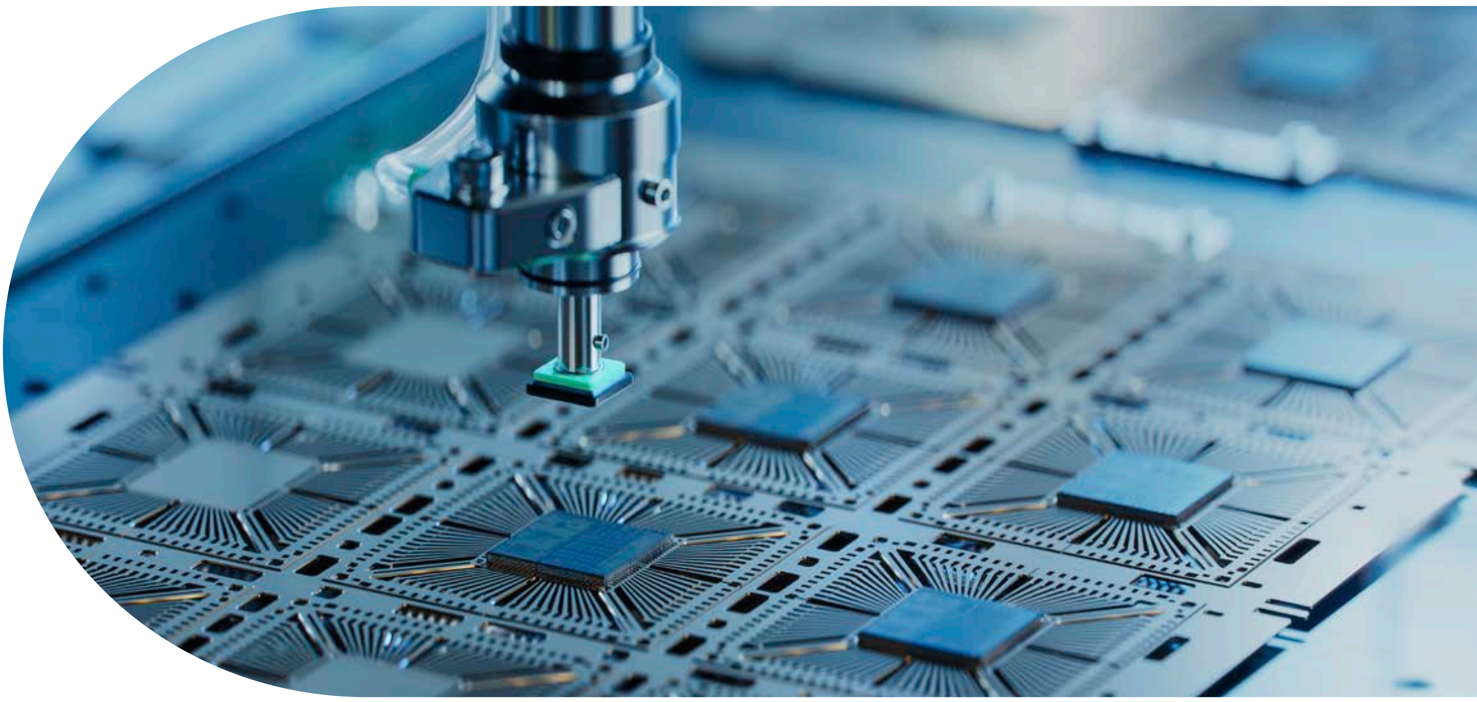

Figure 2: Using CART to select companies

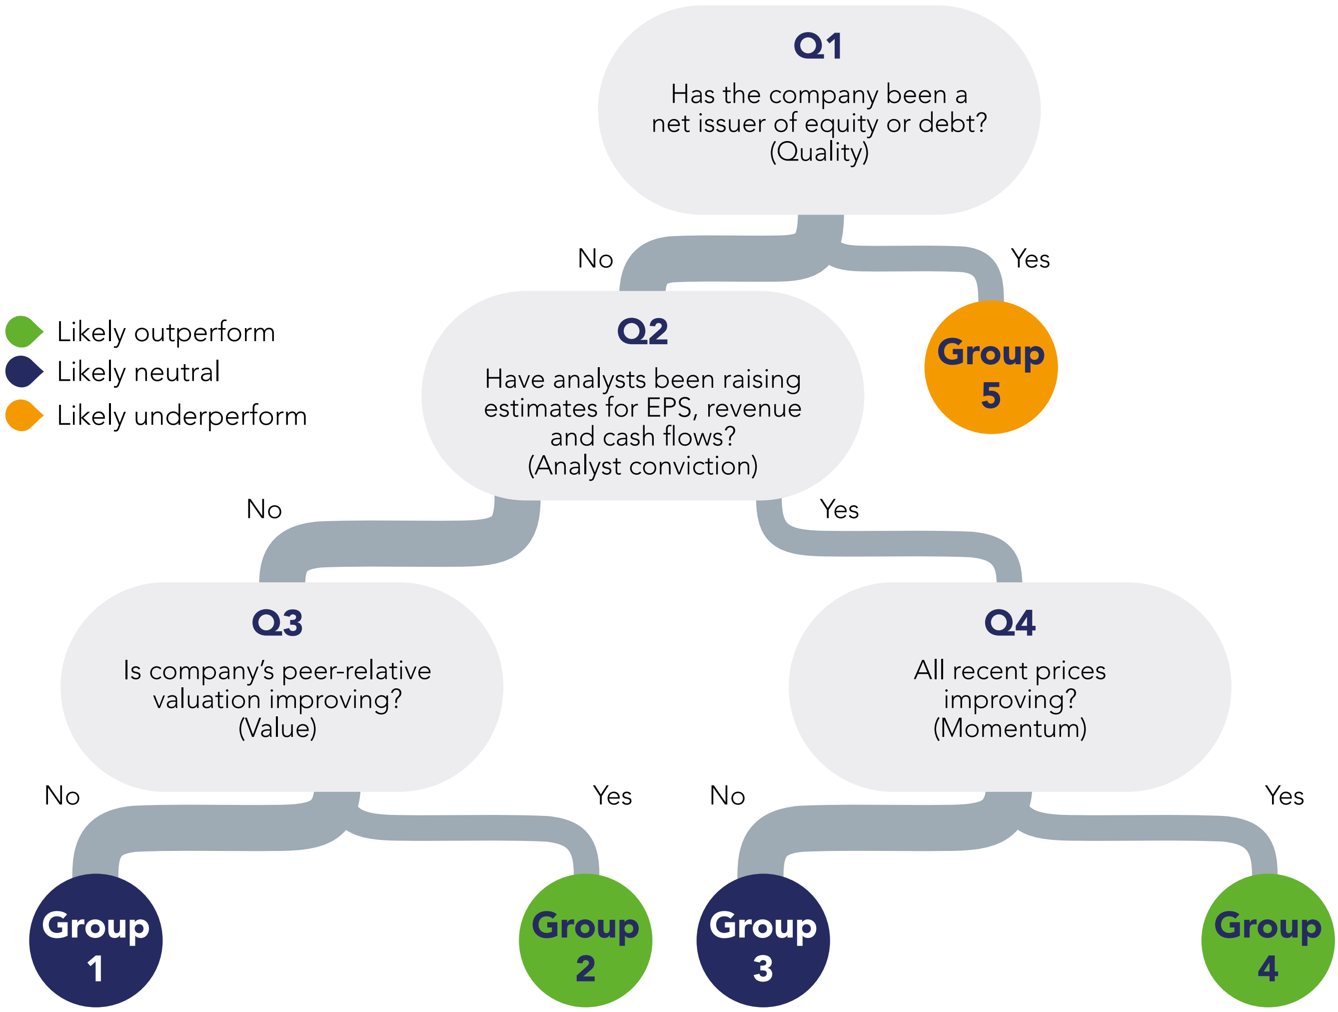

In this tree, the first question is whether or not the company has recently been a net issuer of equity or debt. On average, companies that use a substantial amount of external financing tend to underperform. Note that the question chosen doesn't have to split the data 50/50. In the example tree, it is a minority of companies (Group 5) that use a substantial amount of external financing, but it provides a strong signal that (again, on average) those companies are likely to underperform.

The second question in this sample tree is about sell-side analyst conviction. This question is asked of only the companies that answered "No" to the first question, so we already know that these companies aren't relying on excessive financing. Next, the algorithm determines the best question to ask these specific companies, and the question it finds is "Have analysts been raising estimates for earnings-per-share, revenue and cash flows?". That question we have found is most useful for growth-oriented companies. The question divides the remaining data observations into two smaller groups, with somewhat more companies answering that question with "No" than "Yes."

At this point, we know more about the companies going into questions three and four. The companies going into question three did not have high analyst conviction, and the algorithm finds that the best next question for those companies is about a value-oriented variable. The companies going into question four do have high sell-side analyst conviction, and the best next question is about momentum, another factor that tends to do well for growth-oriented companies.

After each company has answered the relevant questions in this simple example tree, we now have five groups of companies with differentiated alpha estimates.

Advantages of using a regression tree to forecast stock alpha:

Versus a more traditional/linear approach

- Regression trees sift through a vast amount of data to find companies with combinations of characteristics that have foreshadowed price movements relative to the universe over its history.

- Regression trees use explanatory variables non-linearly. A high value of one explanatory variable may be suitable for one company but bad for another with different characteristics.
- Only the questions relevant to each company are asked. If a company is a value company, the algorithm doesn't waste time asking questions that are more relevant to growth companies.
- Highly scored companies won't all have the same characteristics. In the sample tree shown on the previous page, both Groups 2 and 4 are forecast to outperform. Group 2 has more value-oriented names, while Group 4 has more growth-oriented names with substantial price momentum. That makes it easier to build a portfolio with better risk characteristics than if you had high-scored companies from a linear model where the companies with the highest scores had similar characteristics.
- From anecdotal evidence, few investment shops make regression trees a central part of the investment process. We believe that our trades are less crowded than those of other managers.

Versus other machine learning techniques

- Regression trees are transparent. It is easy to understand why companies get high scores by looking at their values of the explanatory variables. That means it is easy for us to understand why or why not a company gets a high score, and it is easy for us to review the model's daily trades and understand why they are being made – a valuable quality control.
- Regression trees are relatively easy to build. There are a modest number of parameters to be specified and what those parameters do is fairly intuitive.
- Regression trees are robust to input data. Outliers aren't a problem as all that matters is which companies are above or below the split point. The data doesn't need to be normalised.

The drawbacks of using a regression tree to forecast stock alpha:

Single tree problems

As outlined above, there are many advantages of using a classification and regression decision tree (CART) to forecast stock returns over traditional quantitative approaches. However, as can often be the case with systematic investing, after observing the behaviour of decision tree models in day-to-day practice and over the course of many years, we noticed a variety of potential areas for improvement in how we utilise technology.

Forecasts are too sensitive to specific factor values

The valuable characteristic of decision trees, where the modelling can be different on either side of each decision point, can be problematic in a dynamic setting. This is because slight changes in factor values from one day to the next can lead to answering questions differently, causing potentially large differences in forecasting, leading to overtrading. We were able to devise ad hoc techniques to mitigate this issue, but they did not fully cure the underlying problem.

Overfitting

As a decision tree is built deeper with each additional level of questioning, the algorithm has smaller and smaller pools of data to work with as it selects the best question to ask. This can lead to overfitting, where a fit model does poorly on new (out-of-sample) data, which is especially true when the data has a high level of noise, as is the case with stock returns. CART comes with tree pruning mechanisms to help users determine the appropriate depth to which a tree should be built, but our experience was that those mechanisms were not particularly helpful. Therefore, a significant amount of hand-tuning was required in our tree selection process, which made it hard to backtest the process properly.

Underfitting

Only a single question is asked of all companies, at the top of the tree. If there are multiple questions that would be useful to ask broadly of all companies – as seemed likely with our application of decision trees – the model will be underfit relative to what would be possible with other techniques. Furthermore, because the trees will necessarily be kept shallow due to overfitting issues, only a limited number of questions will be asked of each company. At most, we used trees of depth 5 or 6, which put fairly strict constraints on how sophisticated the modelling could be. In fact, a single tree can even overfit and underfit at the same time when the number of questions is limited but it chooses the wrong questions to ask (questions that reflect some statistical feature in the historical data but do not generalise to unseen future market environments).

Noisy model updates

Decision tree construction is a greedy algorithm – that is, when the questions are selected by the algorithm, only the single best question gets incorporated into the tree. However, our experience was that there would often be one or more runner-up questions using a different factor that were almost as good. Then, when we would update our model with additional market data, one of the runner-up questions might take the lead and replace the winner question in the tree. This could lead to significant differences in the tree structure and predictions, especially if it happened near the top of the tree. This is an undesirable product of small changes in the question selection process, as it would lead to excessive turnover when we updated our models.

The solution

An ensembles of trees – "bagging" and "boosting" come to the rescue

We suspected that if we were able to build models that employed multiple decision trees, many of the challenges we faced with our existing trees would be lessened, while still hopefully retaining the attractive features of decision tree modelling. Fortunately, the research community had provided a few potentially attractive paths to pursue.

Bagging

In the late 1990s, a handful of researchers began experimenting with introducing randomness into the decision tree construction algorithm in order to improve predictive accuracy in the face of problems like the ones described above. Among this work is the 2001 paper by Leo Breiman (author of the original CART paper) introducing the concept of "random forests." Central to the construction of a random forest of decision trees is bootstrap aggregating, or "bagging," which is the idea of constructing multiple decision trees where each tree is constructed using a small random sample of the overall pool of data. The overall model forecast is the average of the forecast of the ensemble of bagged decision trees. The key idea is that even though each individual tree in the forest is a "weak learner" – that is, it is not as good as a tree that is trained on the full dataset – the collective wisdom of all these individual attempts to model the data leads to a more robust set of forecasts that will perform better on new data.

This approach systematically addresses the issue of overfitting. No two trees are built with the same sample of data, but the samples are all drawn from the same overall pool of data. Thus, each tree contains some questions that are specific to its own particular dataset and other questions that generalise to the larger pool. Averaging across these trees reduces overfitting by "averaging out" the sample-specific questions, which are less likely to generalise to out of sample data, and enhances the prominence of the questions that are more likely to generalise.

With overfitting less of a concern, it was possible to build individual trees with far greater depth, allowing each tree to ask many more questions that can help make more nuanced forecasts. By allowing more questions to be asked, the accuracy of the overall forest of trees is no longer constrained by the number of questions, which addresses the issue of underfitting.

Furthermore, while the primary goal of using multiple trees was to improve the accuracy of our alpha forecasts by addressing underfitting and overfitting, the forest approach had a secondary benefit of also smoothing the forecasts so that a small change in a model input no longer tended to cause a large change in the model forecast. This smoothing comes as a direct consequence of using averaging to reduce overfitting. Because overfit models tend to produce large and overconfident, but idiosyncratic, forecasts, averaging out these predictions across multiple trees allows the more modest but more generalizable forecasts to prevail. This helps reduce not just day-to-day turnover but also model change turnover, as the model is closer to capturing the underlying patterns that predictably drive stock returns over the long term, which tend to change slowly.

MDT Advisers introduced a random forest-based alpha model with 500 trees into our investment process in 2013, after our backtesting research showed there to be significant improvement compared to a single tree-based alpha model. While there were obvious transparency and calculation costs to implementing this enhancement, the backtested results were highly compelling, and we were able to mitigate the transparency issues by developing new tools to allow us to visualise how a change to the value of a particular actor would impact the forest's alpha prediction.

Figure 3: Predictions are unique to each company and factor

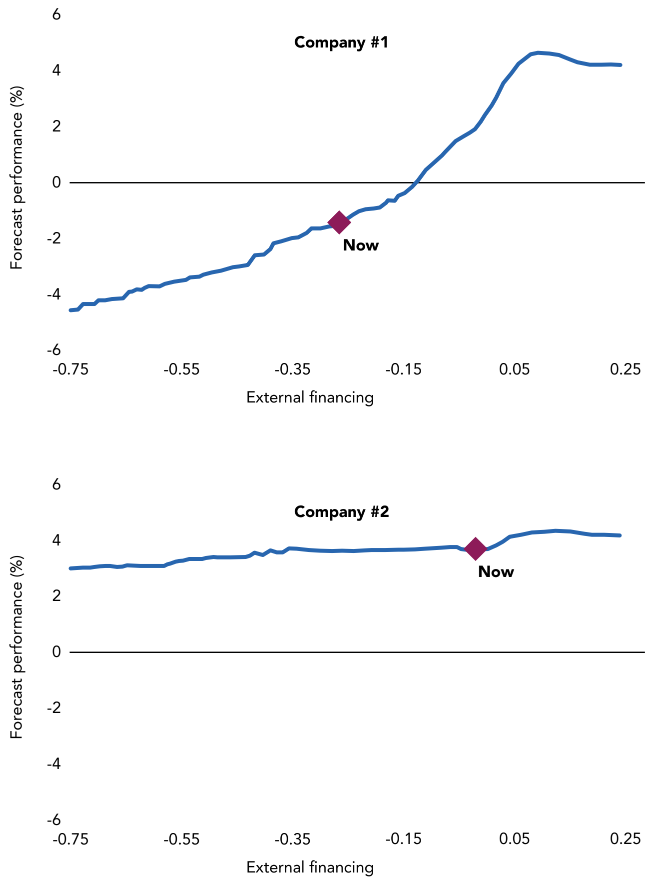

Source: Federated Hermes MDT Advisers. Representative daily factor analysis.

For any company, we can generate graphs for each investment factor. In the charts above, the y-axis shows forecast performance, while the x-axis shows the possible values of a particular factor. The purple diamond shows the company's current values for the factor and the forecast performance. The blue line displays how the company's forecast performance today would change should the factor value increase or decrease from its present value.

We can see here how two companies respond differently to changes in one of our investment factors, given the nonlinearity introduced by the multiple-tree structure. The external financing factor has a major impact on scoring for Company #1, while Company #2 is fairly insensitive to that factor. The apparent smoothness and lack of discontinuities in the blue lines results from the large number of splits on each factor across the trees in the full forest. If we were to produce a graph like this for a single tree, the blue line would be horizontal, perhaps with a sudden jump to a new plateau if the tree had a split on that factor in the company's path through the tree.

Over time, as had been the case with single decision trees, we became more aware of the downsides of the random forest algorithm. Principally, we found that whereas the features of a single decision tree model were overly sharp, the features of the random decision tree forest were too dull, as a byproduct of the averaging of all the individual trees. Model predictions adjusted slowly to changes in our factors – perhaps, too slowly.

Boosting

One of Professor Breiman's co-authors on the original CART paper, Jerome Friedman, had a different take on how to improve CART models using multiple decision trees, introducing the concept of gradient boosting in a set of papers in 2000 and 2001. Unlike random forests, where the trees are all built independently from one another, boosted trees are built in sequence, which reduces the ability of parallel processing to efficiently construct the ensemble of trees. Here, each individual tree is intentionally built smaller than would be optimal in a single decision tree context (to create a "weak learner"), but the subsequent trees are able to learn from the trees that came before it. So, rather than start from scratch, the second tree incorporates some information about the predictions made by the first tree, the third tree incorporates information from the first two trees, and so on. While this technique helps to cure many of the problems faced by single decision trees or random forests, it is exceptionally prone to overfitting issues, so mitigation procedures (regularisation) are highly important in order to achieve robust outcomes. Boosting really took off with the introduction of the XGBoost statistical package, which quickly became the zeitgeist of the machine learning world.

Inspired by the success that others had found using boosted decision trees across a wide variety of machine learning problem spaces, we began a program of research, culminating in the 2017 implementation of an XGBoost-based decision tree forest for our alpha model. We saw significant improvement relative to our random forest-based models, and it gave us a better platform for discovering the subtle interactions between fundamental and technical factors over the model updates to come.

The main improvement over random forests was in the issue of underfitting. Boosted decision trees tended to make more accurate forecasts by using the historical data in a more efficient way. Because each tree is able to build off of the patterns identified by the trees that came before, it does not need to start from scratch and identify those patterns again. It can, instead, find a new combination of questions to ask where the forecast can be improved. In contrast, in a random forest, each tree must start from scratch and many end up identifying similar lines of questioning over and over again.

Figure 4: Forest of decision trees

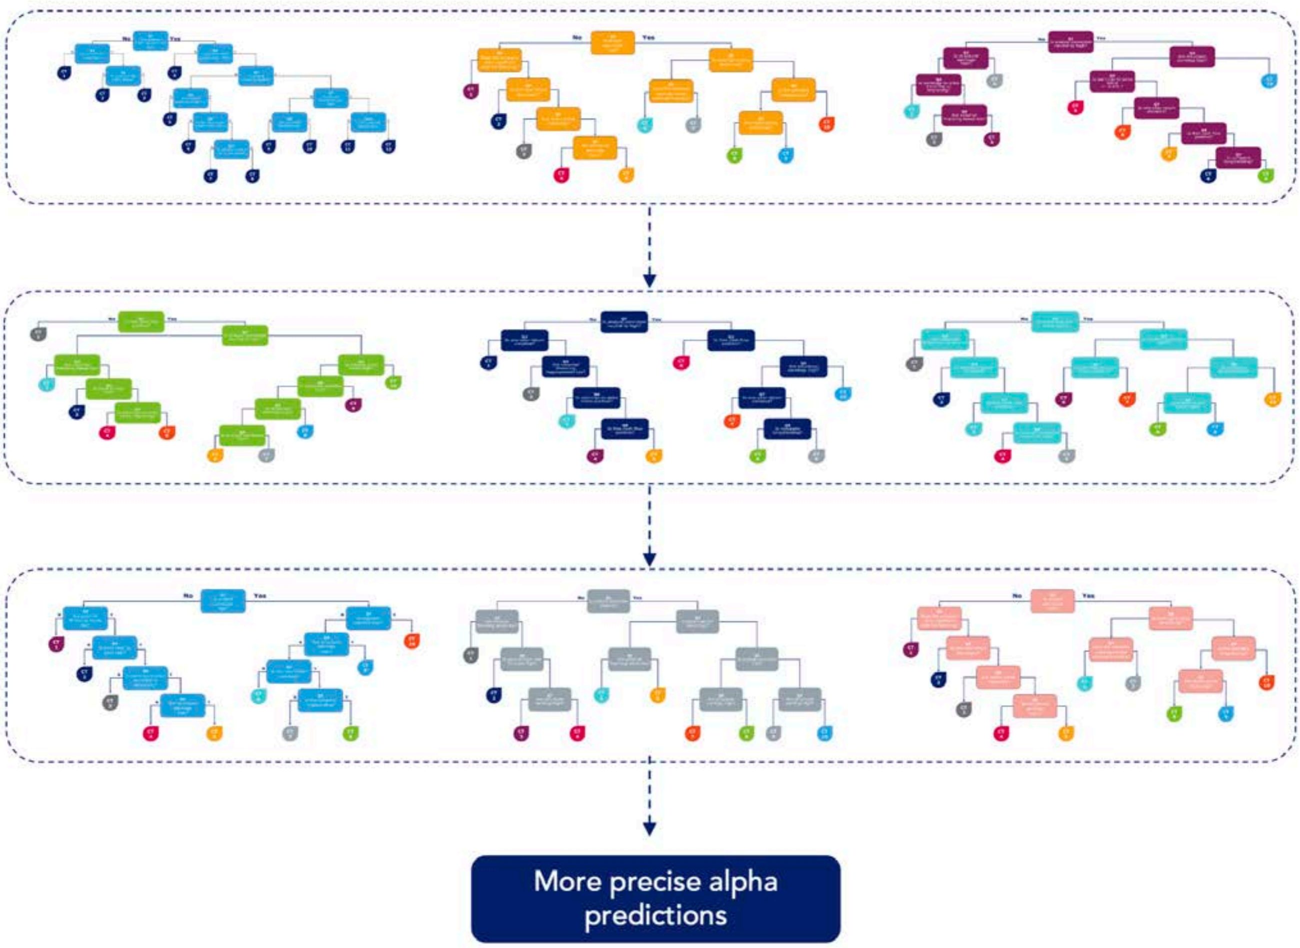

Bringing it all together

While it may have been convenient to simply put the idea of bagging out to pasture with the improved results seen in our model and across the machine learning world by using boosting techniques, our experience of using random forests for a number of years had given us a fair bit of appreciation for the potential value of bagging. We began a relatively novel exploration program of researching the interaction of the boosting and bagging, and in 2020, we implemented a forest-building algorithm that utilised both ideas – boosted random forests.

In contrast to traditional boosting, where a single small decision tree is built as a weak learner in between each round of boosting, our algorithm builds small bagged forests at each level, as illustrated above. Rather than building a forest of 500—2,500 subsampled trees as had been done in our random forest-based models from 2013-2017, we build 25-tree bagged forests in between boosting steps. The increased robustness of the modelling at each step, combined with parallel processing each of the 25-tree bagged forests, means that we are able to construct these bagged-and-boosted forests in roughly the same run time as our single-tree boosting models. Empirically we have found that this combination of techniques produces models with a more robust fit on out-of-sample data than the individual techniques we had previously used.

MDT's work on more sophisticated machine learning techniques has paid significant dividends in terms of the effects of improved alpha modelling on the return profiles of our various investment strategies, but there have been other benefits as well. The increased robustness of our model predictions to the data has allowed us to update our models more frequently – every six months rather than 24 – because the significant "model update noise turnover" that we had previously seen has all but disappeared. Being able to incorporate our new ideas into the model in a timelier manner is a valuable benefit for our clients. Additionally, better forecasting has helped us tighten up the risk controls and the expected tracking error of our strategies. With a single decision tree, we had a much more limited set of alpha opportunities, which led to costly trade-offs in market environments where those opportunities were concentrated, say, in a small number of sectors. With multiple trees, we see a much wider set of companies with alpha potential and are therefore able to tightly control our exposure to sectors (risk) without a significant return cost.

While building investment models using cutting-edge machine learning techniques is not without costs in terms of somewhat reduced transparency, increased computational challenges, and greater difficulties in client communications, we have seen highly satisfactory benefits to our clients' portfolios over the past decade of using multiple decision trees in our investment process. As machine learning algorithms and techniques improve over the coming years, we expect to continue to evolve our process to leverage those advancements wherever we see the potential for improving our strategies' outcomes.

The value of investments and income from them may go down as well as up, and you may not get back the original amount invested.

Past performance is not a reliable indicator of future results.

BD015852 0019262 08/25

Federated Hermes

Limited

Federated Hermes

Federated Hermes is a global leader in active, responsible investing.

Guided by our conviction that responsible investing is the best way to create long-term wealth, we provide specialised capabilities across equity, fixed income and private markets, multi-asset and liquidity management strategies, and world-leading stewardship.

Our goals are to help people invest and retire better, to help clients achieve better risk-adjusted returns and, where possible, to contribute to positive outcomes that benefit the wider world.

Our investment and stewardship capabilities:

- **Active equities:** global and regional
- **Fixed income:** across regions, sectors and the yield curve
- **Liquidity:** solutions driven by five decades of experience
- **Private markets:** private equity, private credit, real estate, infrastructure and natural capital
- **Stewardship:** corporate engagement, proxy voting, policy advocacy

For more information, visit **www.hermes-investment.com** or connect with us on social media: