The Kernow Journal

The Alternative UK Equity Market View

Series 1, Volume 4 | January 2024

KERNOW
ASSET MANAGEMENT

CONTENTS

PUBLISHED BY

KERNOW
ASSET MANAGEMENT

www.kernowam.com
contact@kernowam.com

A note from THE EDITOR

Welcome to the fourth edition of the Kernow Journal

In this issue, we take an in-depth look at data.

Recent years have seen a proliferation of data, to such an extent that the past two years alone have almost produced more data than the entirety of human history that preceded them. With the rise of AI, the rate of data collection will continue to accelerate and, in turn, so too will the replication of any errors in that data. So how can we best collect, screen and interrogate this vast new expanse of data?

In Kernow Contemplations on page 8, we explore the complexities of assessing data and consider the comparative merits and drawbacks of quantitative versus qualitative information. We also discuss the importance of qualitative information in evaluating the authenticity of corporate behaviour, and the various means by which such qualitative data can be collected.

On page 15 we continue our exploration of the Kernow Valuation Framework. We demonstrate the accuracy of the valuations produced by our framework, using the 'factor' approach. This allows us to explore many of the properties of our valuation framework in a way that naturally controls for various biases and allows for comparability across multiple periods, investable universes and valuation metrics. This approach allows us to demonstrate the efficacy of our valuation framework.

On page 24, we discuss how best to interrogate data. We present several fascinating observations that we uncovered in collaboration with the University of Exeter Mathematics Department in a study of the UK Regulatory News Service dataset. We highlight some of the approaches we took to process the data and the insights we gleaned.

By this point, we've talked quite a bit about how to analyse and interrogate different types of data. But there are also the bigger questions of how data should be obtained and why data sources are important. In recent decades, the business of data vending has grown exponentially, with vendors promising to process and standardise the enormous quantities of data supplied by primary sources. Although this offers significant utility, the data sold by data-vendors can become debased or corrupted in many ways.

On page 45, we use several real-world examples to show that it is important to construct models using data directly from primary sources, with data-vendors used only for validation or comparison purposes. There is a significant opportunity for investors who are able to identify where mistranslations of data occur – a considerable 'inconvenience premium'.

Lastly, on page 52 we continue our exploration of data visualisations by delving into geospatial analysis. As an illustrative example, we examine the evolving impact of human activity in specific geographic regions. The elegant combination of appealing aesthetics with highly topical information is tough to beat – and who doesn't love a map?!

We sincerely hope that, in these pages, you find valuable insights into our perspectives as UK equity investors.

Ed Hugo

Edward Hugo
CEO

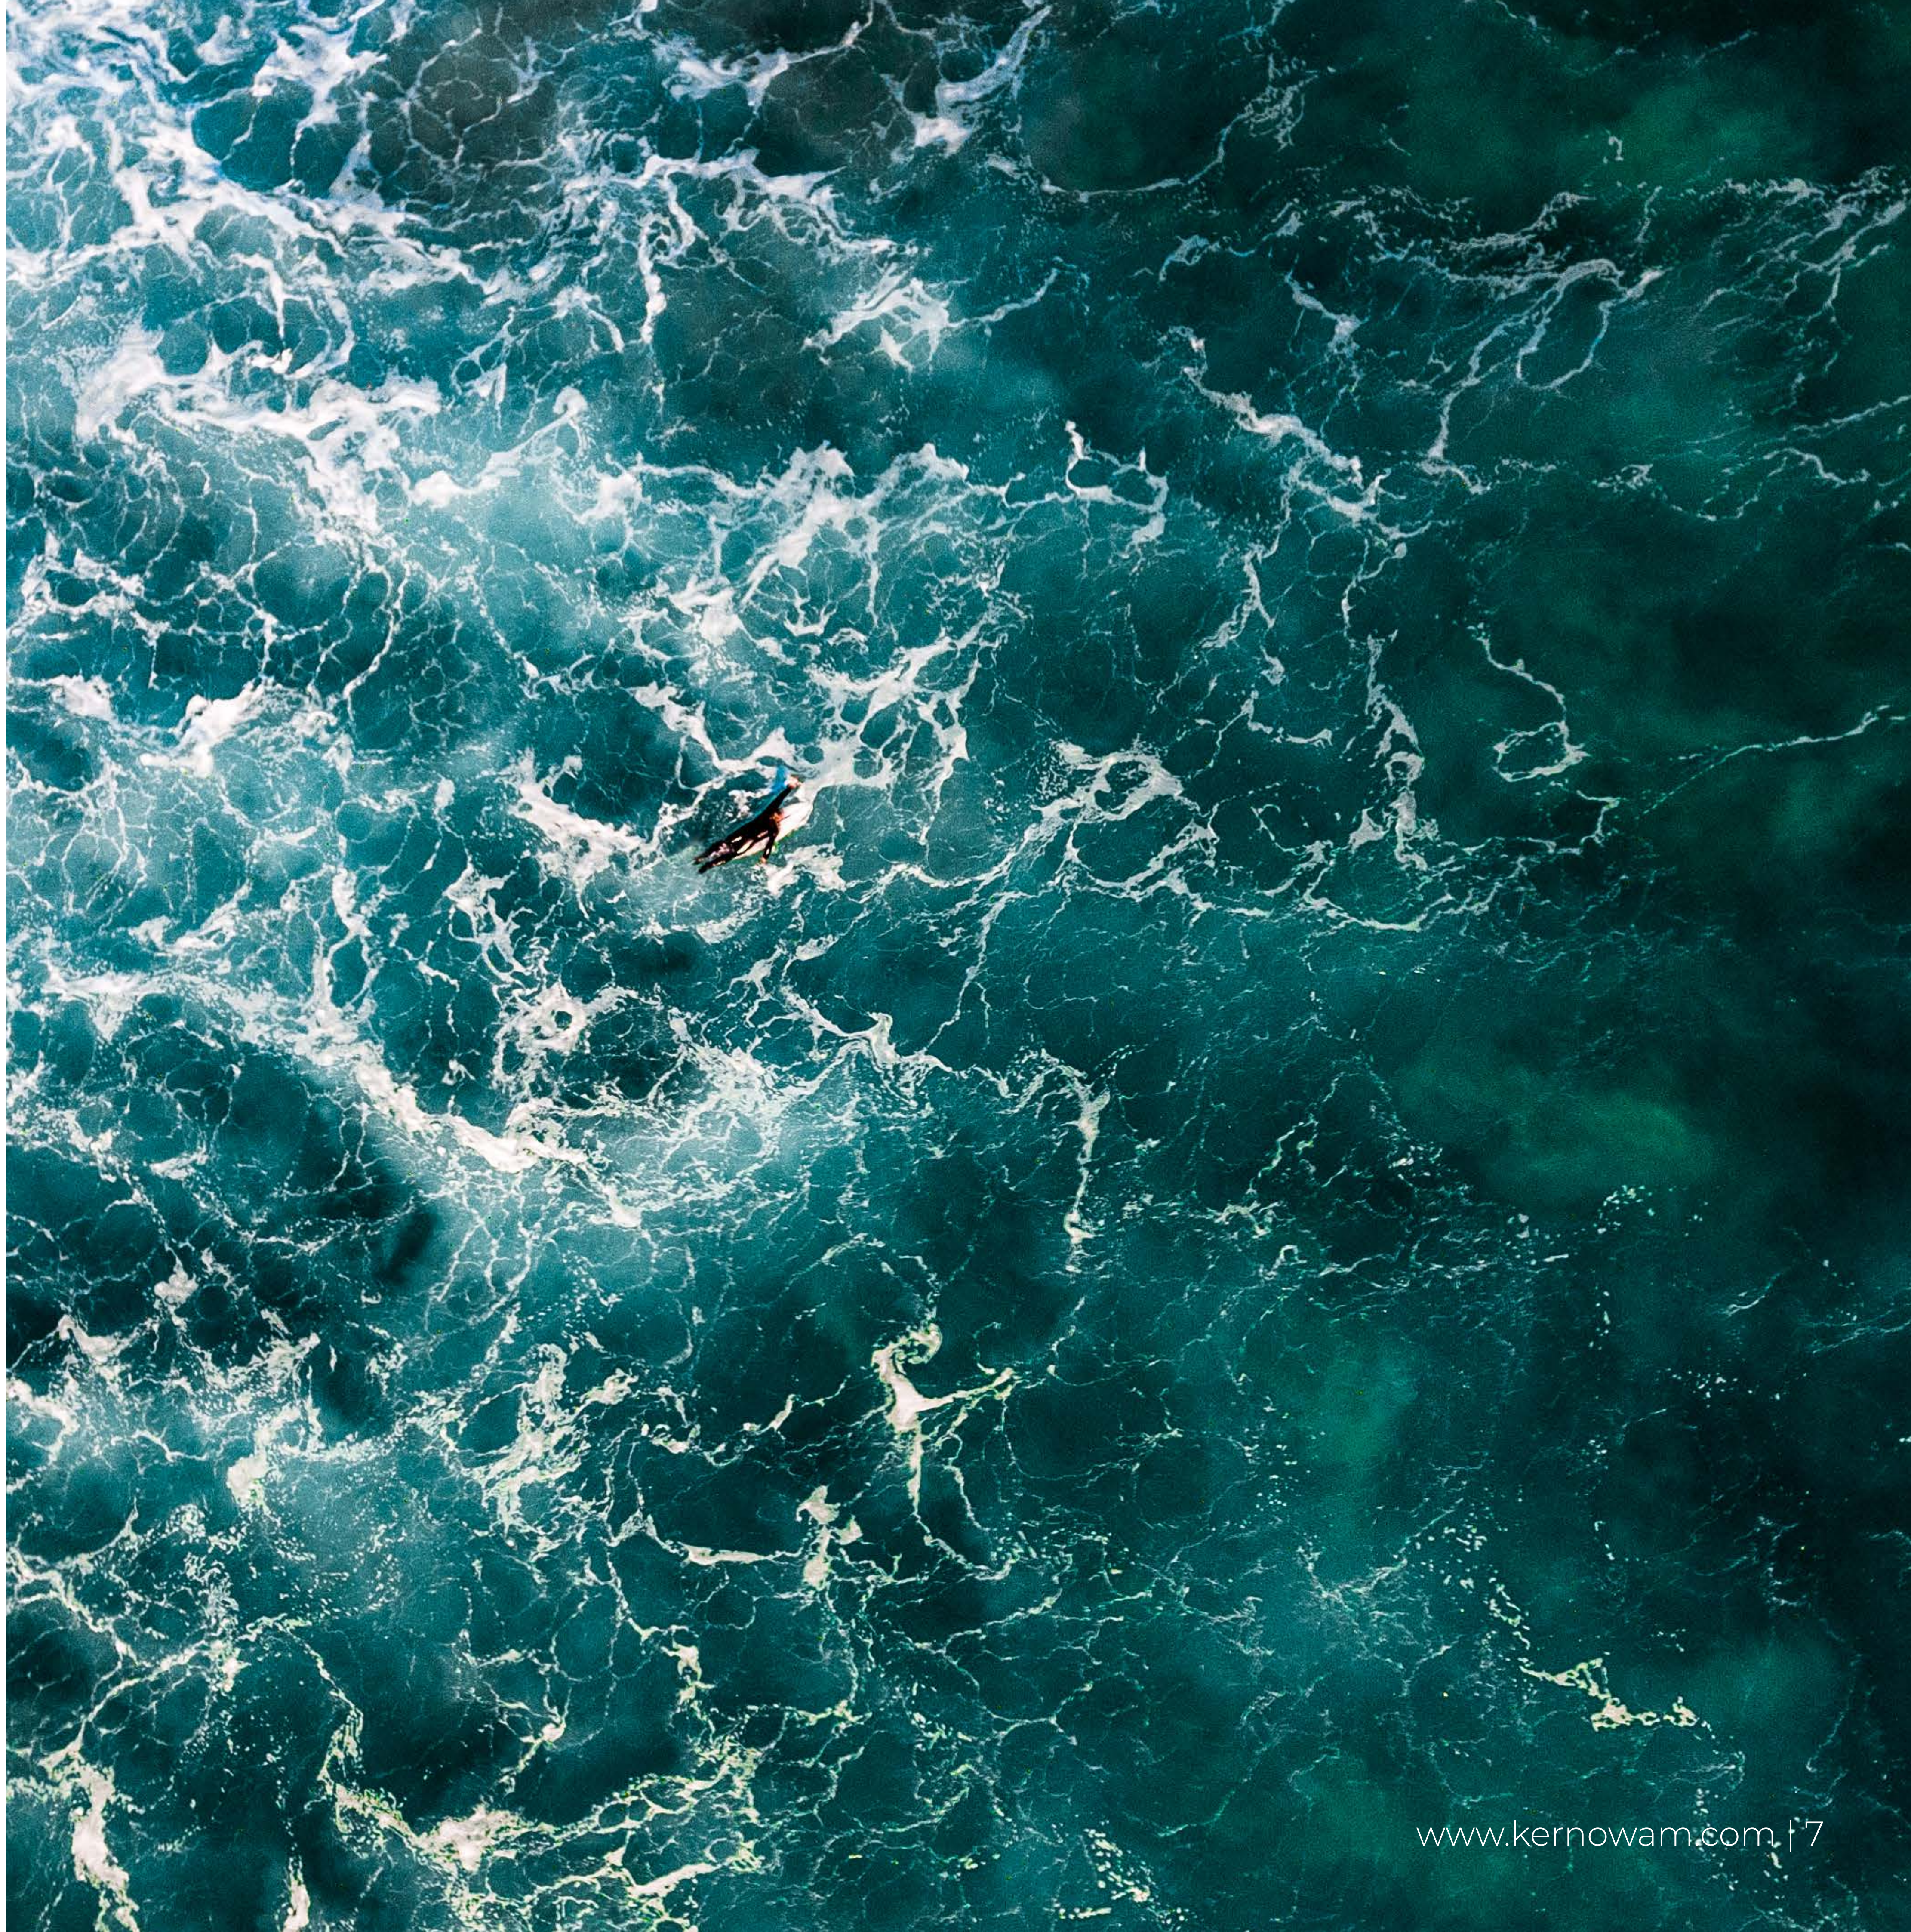

Kernow
CONTEMPLATIONS

Weighing the Pros and Cons of Qualitative and Quantitative Information

Humans naturally favour quantitative information, which enables us to tally, measure and compare with relative ease. In finance, we naturally consider market trends and patterns, company fundamentals, analyst expectations, technical market data patterns and other quantitative information. We also ponder the context of this information in terms of 'how' and 'why' a company has a particular characteristic.

Increasingly, we are finding that quantitative information processing is becoming commoditised.

This is because qualitative information is considerably more difficult to process and compare. However, there is a significant wealth of valuable edge available in qualitative information. The way that human minds work is that stories shape our brains and influence our actions. Investor psychology is important with regards to the way it interacts with a company's perceived prospects. Therefore, it is the combination of both quantitative and quantitative approaches which we find yields the most insightful perspectives.

For example, consider a company facing a sudden decline in customer satisfaction, as indicated by a significant drop in customer survey scores, i.e. quantitative data. This might seem like a straightforward problem and a quantitative approach could suggest solutions like reducing response times or increasing support staff.

However, a qualitative analysis could reveal the real issue. Upon collecting qualitative data from in-depth interviews and customer feedback, it becomes apparent that customers are frustrated because they feel that their concerns are not being adequately addressed.

Customers express dissatisfaction with the tone and unempathetic language used in staff responses and the failure to satisfactorily resolve customer issues. In this example, qualitative analysis would unveil the underlying problem, which requires a fundamental shift in customer service culture, training and communication style.

This anecdote underscores the value of qualitative analysis in uncovering the nuances and context behind numerical data. While quantitative data can point to a problem or opportunity, qualitative data often provides the rich detail necessary for a more comprehensive and effective solution.

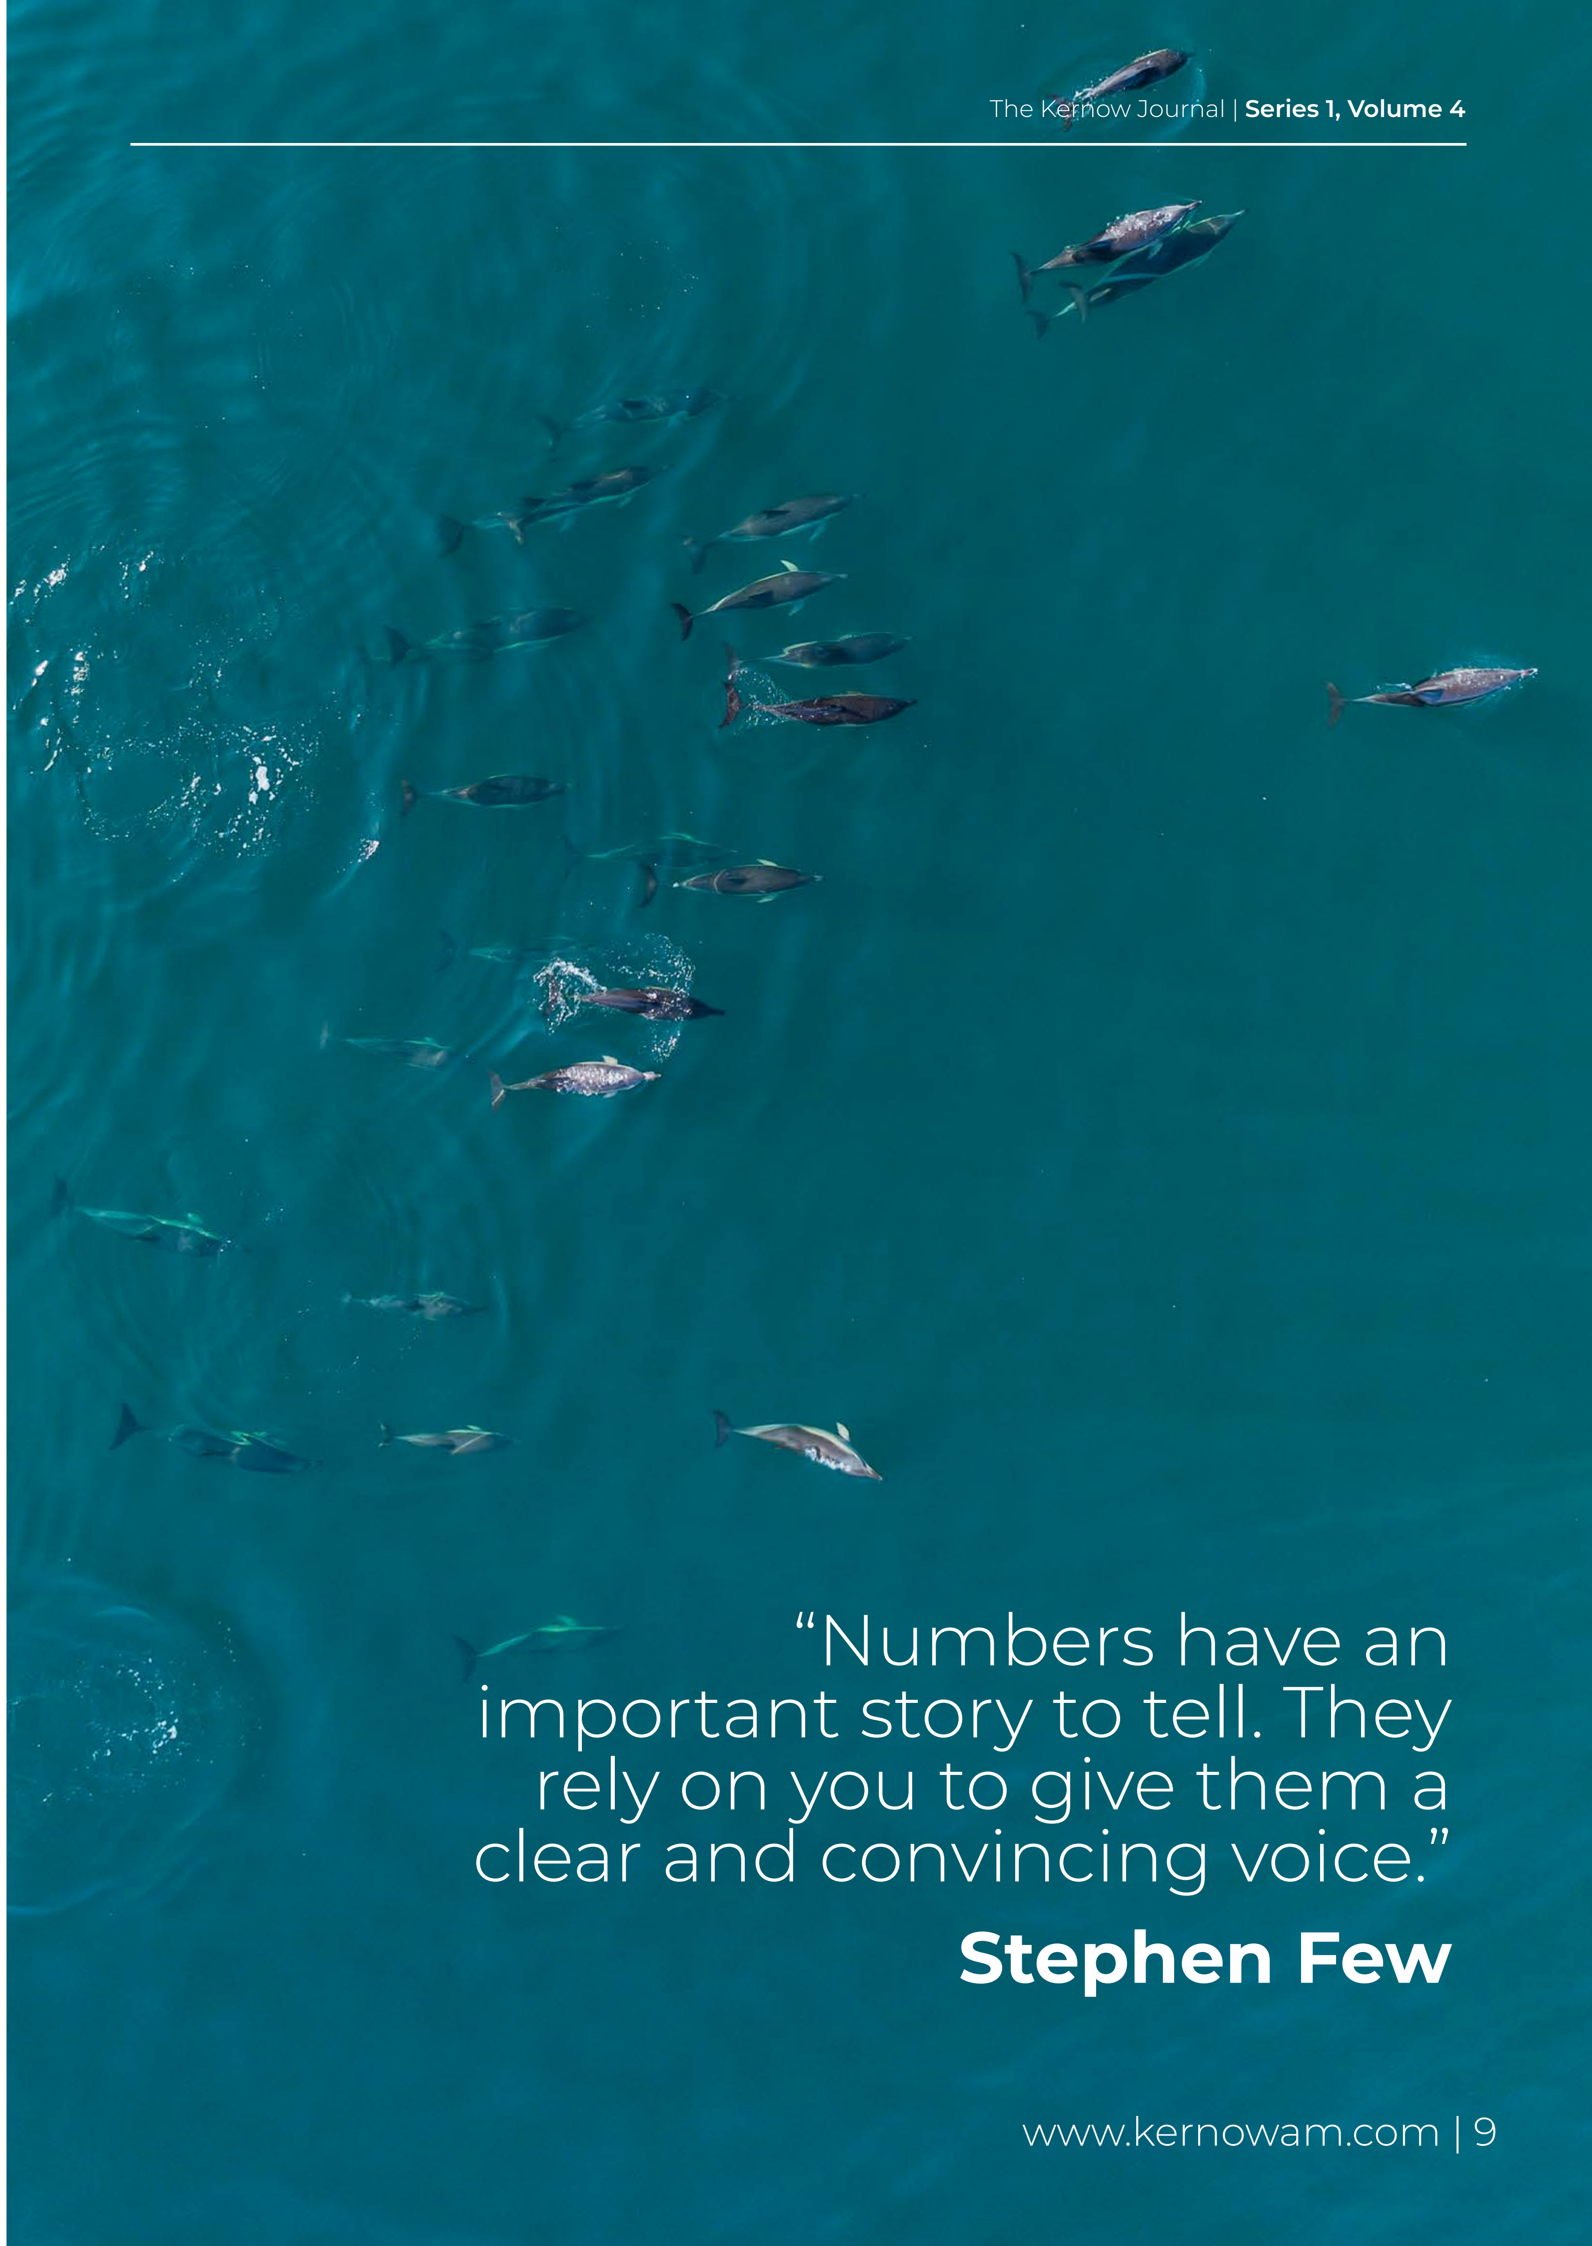

"Numbers have an important story to tell. They rely on you to give them a clear and convincing voice."

Stephen Few

Collecting Qualitative Information to Assess Authenticity in Corporate Behaviour

Yes, it's harder to collect qualitative data, but that doesn't mean you should avoid it. A pertinent example of a critical qualitative attribute is the idea of authenticity, for individuals and organisations. Measuring authenticity is complex, as the metrics are often subjective, multi-faceted and deeply personal. Typically, we seek out companies whose business objectives are authentic and, as a result, their business leaders carry a high degree of authenticity.

Authenticity is critical for businesses for many reasons, given that maintaining a healthy business almost always involves building trust-based relationships, leading to strong customer and supplier relationships.

In an increasingly crowded market, we also find that authenticity can positively impact the development of a differentiated business with a strong brand identity.

Whilst there is significant cross-over with corporate culture, authenticity also tends to trickle down from the senior management to the company's employees – and they tend to feel happier when working for companies with a sincere purpose.

So, how do we go about measuring authenticity? The short answer is we don't!

We think boiling down such a complex attribute into a single number misses the point. Moreover, it is not a binary trait. Instead, we focus on identifying behavioural patterns and qualities that indicate authentic or inauthentic practices.

As Simon Sinek aptly puts it, "authenticity is the key to building loyalty. Be yourself, be genuine, and be true to what you believe in."

For example, we speak to management and interview boards (as discussed in Kernow Journal edition 2). We also conduct referencing, including third party interviews with company employees, customers and suppliers. We have an internal catalogue of topics for further examination. For instance, a company website tells us about a company's culture and how it wishes to present itself. Is it genuine or superficial? How has the company vision and biography changed over time, when viewed using the Wayback Machine? (https://archive.org/). Does the company use stock images instead of photos of their own personnel? This is part of Kernow's intangible value-add, beyond just number crunching and code.

Our firm belief is that authenticity isn't just a value but the cornerstone of trust and lasting relationships. It forms a critical component of our evaluation process for the companies in whom we invest. Conversely, where we find patterns that indicate inauthenticity, we can often uncover more profound challenges within the company under analysis.

Kernow
CONTEMPLATIONS

The Consequences of Research Misconduct

How fragile is scientific discovery? There have been several instances of misleading or fraudulent scientific research throughout history. While these cases are relatively rare compared to the vast body of legitimate scientific work, they serve as cautionary tales and underscore the importance of rigorous peer review and ethical conduct in research.

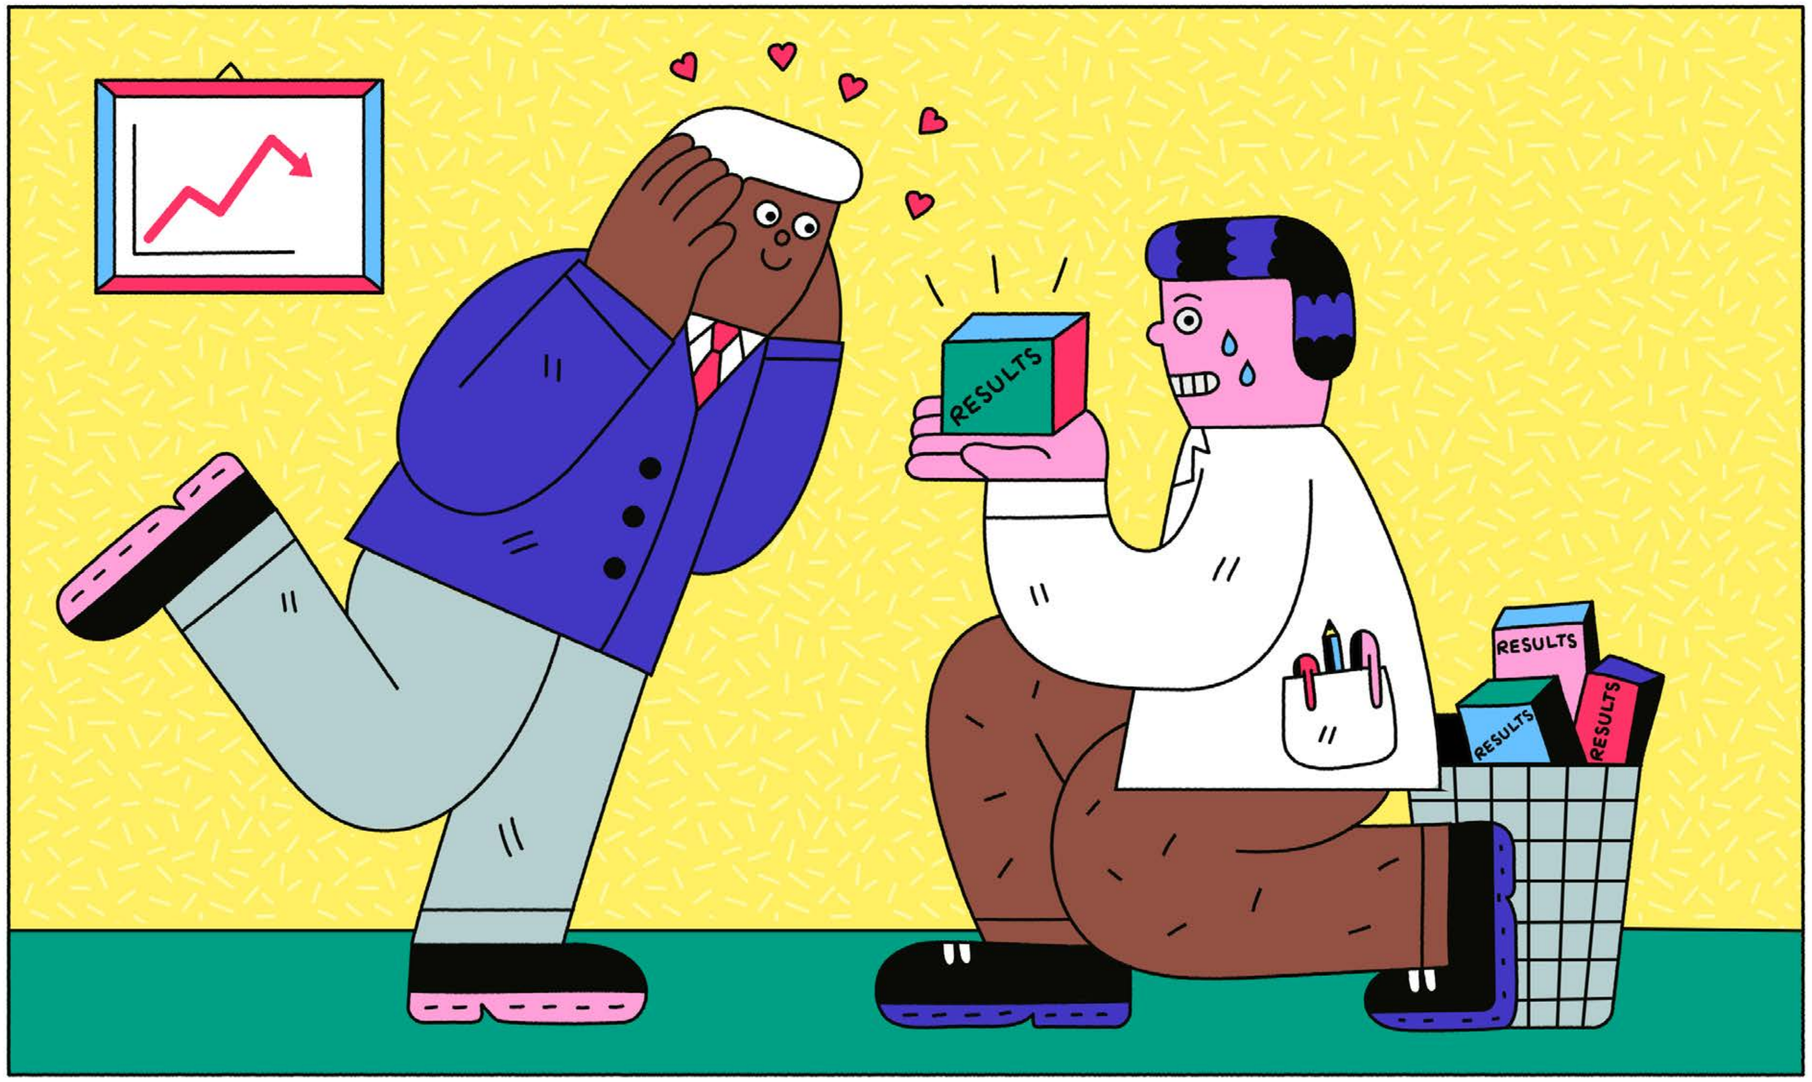

"ACADEMIC ILLUSIONS: WHERE DATA BENDS TO THE WILL OF AMBITIONS, AND FACTS PLAY HIDE-AND-SEEK WITH TRUTH."

Recently, we witnessed a rapid wave of hype as a research team based in Korea published a pre-print article suggesting they had developed a room-temperature superconductor.

The implications of such a discovery are highly profound: revolutionising energy transport and storage - with enormous implications for the electrification efforts of the net-zero strategy - as well as improving efficiency in many electronic goods. A breakthrough would also significantly impact applications in medical imaging, quantum computing, and levitation.

As the news unfolded, hype grew, but due to the highly specialised nature of the discovery, only a relatively small number of teams could attempt to verify the results. As the days passed, anticipation increased but ebbed away as it became clear that the alleged superconducting effect could not be reproduced. This example highlighted the power of the scientific process in independently validating the results, but questions remained as to why the authors published the pre-print in the first place.

A related issue plagues financial research. The so-called 'replication crisis' relates to many studies that have been published that have supposedly uncovered a statistically significant economic relationship that cannot be replicated independently. Moreover, given the nature of financial research, there are several accusations of data-snooping effects – whereby researchers trawl through many permutations and manipulations of data until they find a 'significant' relationship, which they publish. This is almost certainly due to publication bias. Clearly, these two effects are highly detrimental to the canon of financial literature, and for practitioners, it is often difficult to determine which research is legitimate and which is bogus. We use experience and draw our independent conclusions from internally derived research to mitigate these issues.

In closing, we've delved into the importance of qualitative insights in our otherwise increasingly data-driven world. We've contemplated the value of authenticity in shaping relationships and we've scrutinised the challenges of ensuring the integrity of scientific and financial research.

As we conclude this edition's contemplations, let us keep these questions in our minds: How can we balance the quantitative and qualitative to reveal deeper insights? How can we efficiently navigate the complexities of our relationships and research, striving for truth and integrity? These questions drive us onward in our pursuit of insights and authenticity and assist us on the unwavering journey towards making better-informed investment decisions.

"In science, there are no shortcuts to truth. The integrity of research and the meticulousness of the scientific method are our best tools for understanding the world." Carl Sagan

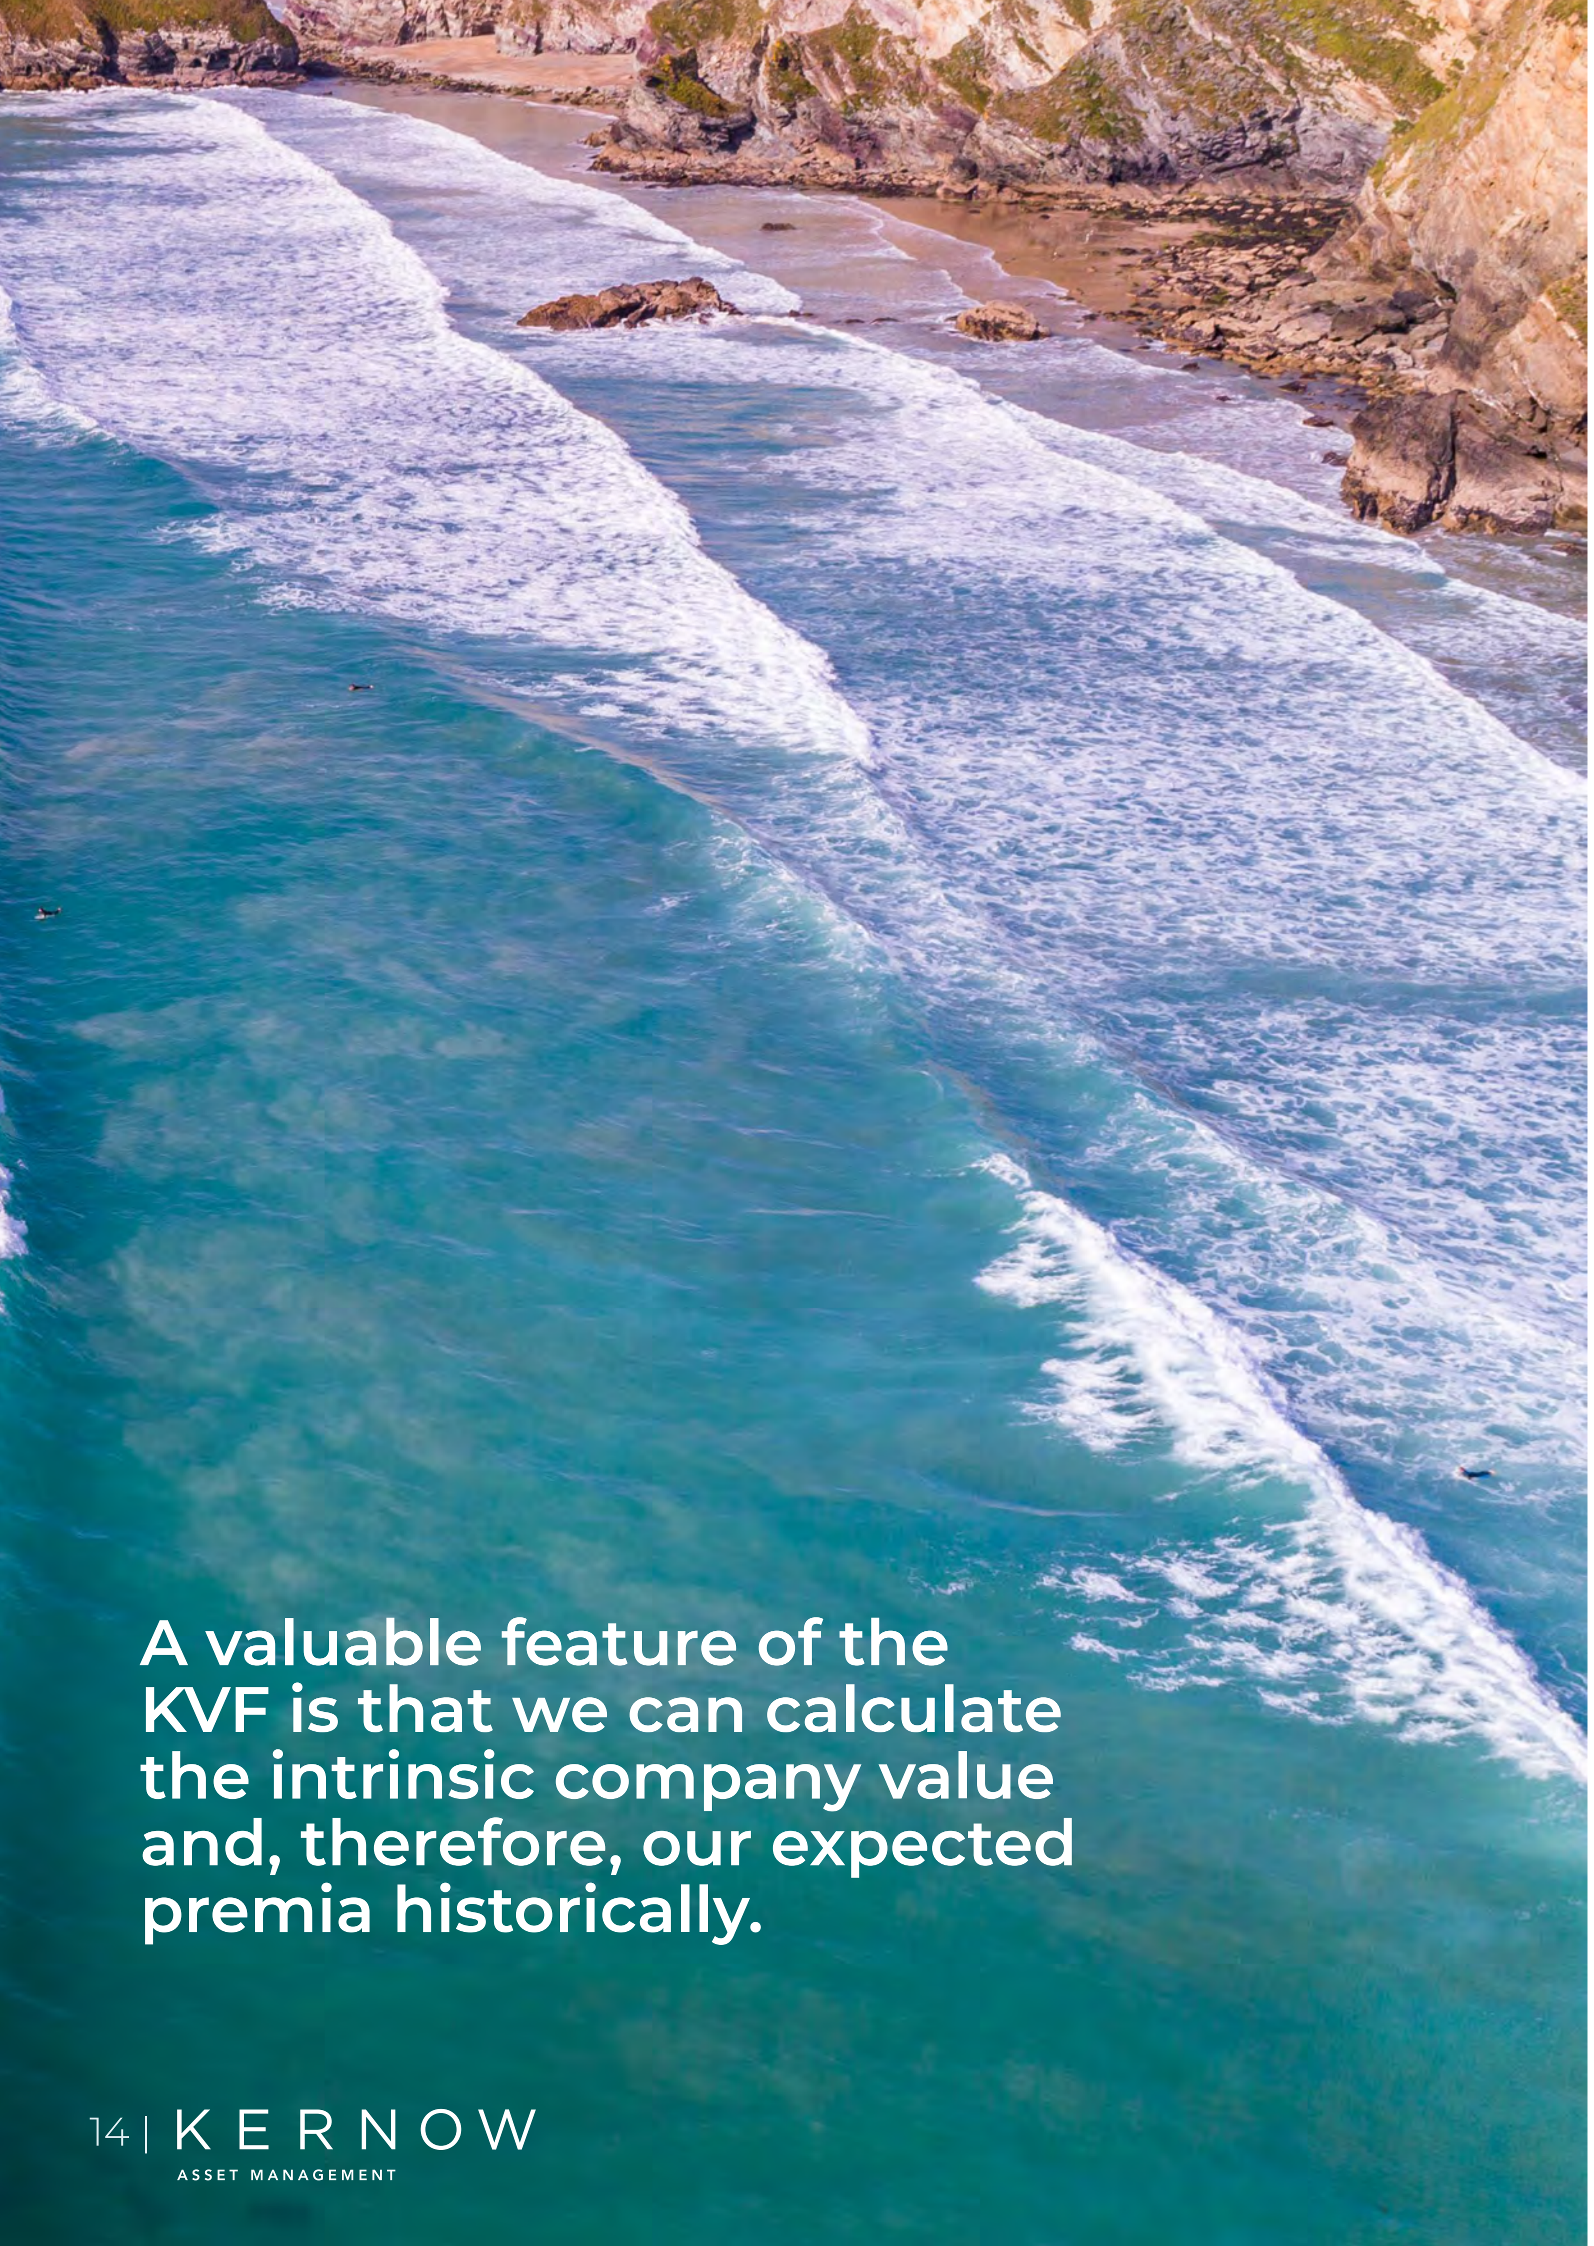

A valuable feature of the KVF is that we can calculate the intrinsic company value and, therefore, our expected premia historically.

The Kernow VALUATION FRAMEWORK

As explained in the previous editions of this journal, the Kernow Valuation Framework (KVF) computes our view on the intrinsic value of each company listed on the UK equity market. By combining this with the market capitalisation of companies, we can easily see where our contrarian opportunities arise.

The KVF provides us with a mental map of the evolving valuation characteristics of our portfolio universe and highlights daily mispricing opportunities. It is our secret weapon for sifting through, comparing and keeping a close watch on companies' intrinsic valuations and related premiums within our investable universe. The KVF can be applied to a range of securities including retrospectively, allowing us to track the ebbs and flows of valuations and pockets of premia as they evolve.

Not all statistical models are created equal. How can we assess whether the Kernow valuation model is a good one? One way to test that is through factor analysis. In this article, we outline a straightforward transformation of the valuation premia provided by our model to produce a 'factor score'. Even though the financial world has been buzzing about factor-based models, the practice of factor analysis is not widely understood. In this article, we will walk through the factor-based analysis of the Kernow valuation model, to assess whether or not it is better at valuing companies than the stock market over time. If the Kernow value factor is positive, then our valuation model is a profitable signal - but if it is negative, then the model is random in terms of profitability.

What is a Factor?

Factor analysis is a technique for identifying patterns and relationships among many variables. It is commonly used in finance and economics to analyse complex datasets and identify underlying (often latent) factors driving market trends or investment returns. In factor analysis, a large dataset is reduced to a small set of variables, called factors, that capture the most important information or variation in the original dataset.

The most common factor in equity investing is the relevant market – securities tend to have a relationship to their headline market. Under the 'Capital Asset Pricing' model, this relationship is conventionally depicted by the regression coefficient, beta, calculated between the security and its parent market. In this way, a factor approach considers that the 'market factor' can explain a proportion of the cross-sectional variation in security returns. For the UK market, this equates to approximately 25% of variability.

Factors can be statistical (using techniques such as principal component analysis), market-based, sector-based or style-based. In this article, we focus on style factors, which typically refer to a characteristic or attribute that is believed to be associated with a higher investment return or a lower investment risk on a forward-looking basis.

Many different style factors have been identified and studied over the years, including value, growth, sentiment, quality and volatility, among many others. These factors are based on different underlying characteristics, metrics and themes.

Factor analysis can be used for a variety of purposes in finance and investing, including uncovering alpha, portfolio construction, risk management and performance attribution. For example, factors can be used to identify the key drivers of investment returns or to construct optimised portfolios.

Whilst the topic can become highly nuanced and sophisticated, three general steps are typically involved in the construction of an investment factor, as follows.

1 Identify the characteristic

The first step in constructing an investment factor is to identify the characteristic or attribute that is believed to be associated with higher investment returns or lower investment risk. This may involve analysing historical data, academic research, or other sources of information to determine which factors have been shown to be statistically significant in explaining returns. The valuation premium output by the Kernow model is the characteristic we identify for this factor analysis. The model tracks mean reversion: in general, the market value of companies should revert to their intrinsic valuation over time. Both over- and under-valuations are likely to show mean-reversion characteristics.

2 Define the factor construction parameters

Once the metric is derived, the next step is to define the set of eligible securities for comparison (i.e. 'the universe') and define any groupings on which to stratify the data – such as industry, geography and capitalisation. For this example, we tend to stratify by industry, as the valuation premia are only comparable across companies operating in similar areas.

3 Calculate factor scores

Factor scores are then calculated for each security or asset in the dataset. This involves forming a cross-sectional comparison within stratification groups to standardise the metric scores. Several methods, including the 'z-score' or percentile-rank-based methods, can be applied here. For this example, we use a ranking and normalisation method to create normally distributed scores in a way which is robust to outlier data.

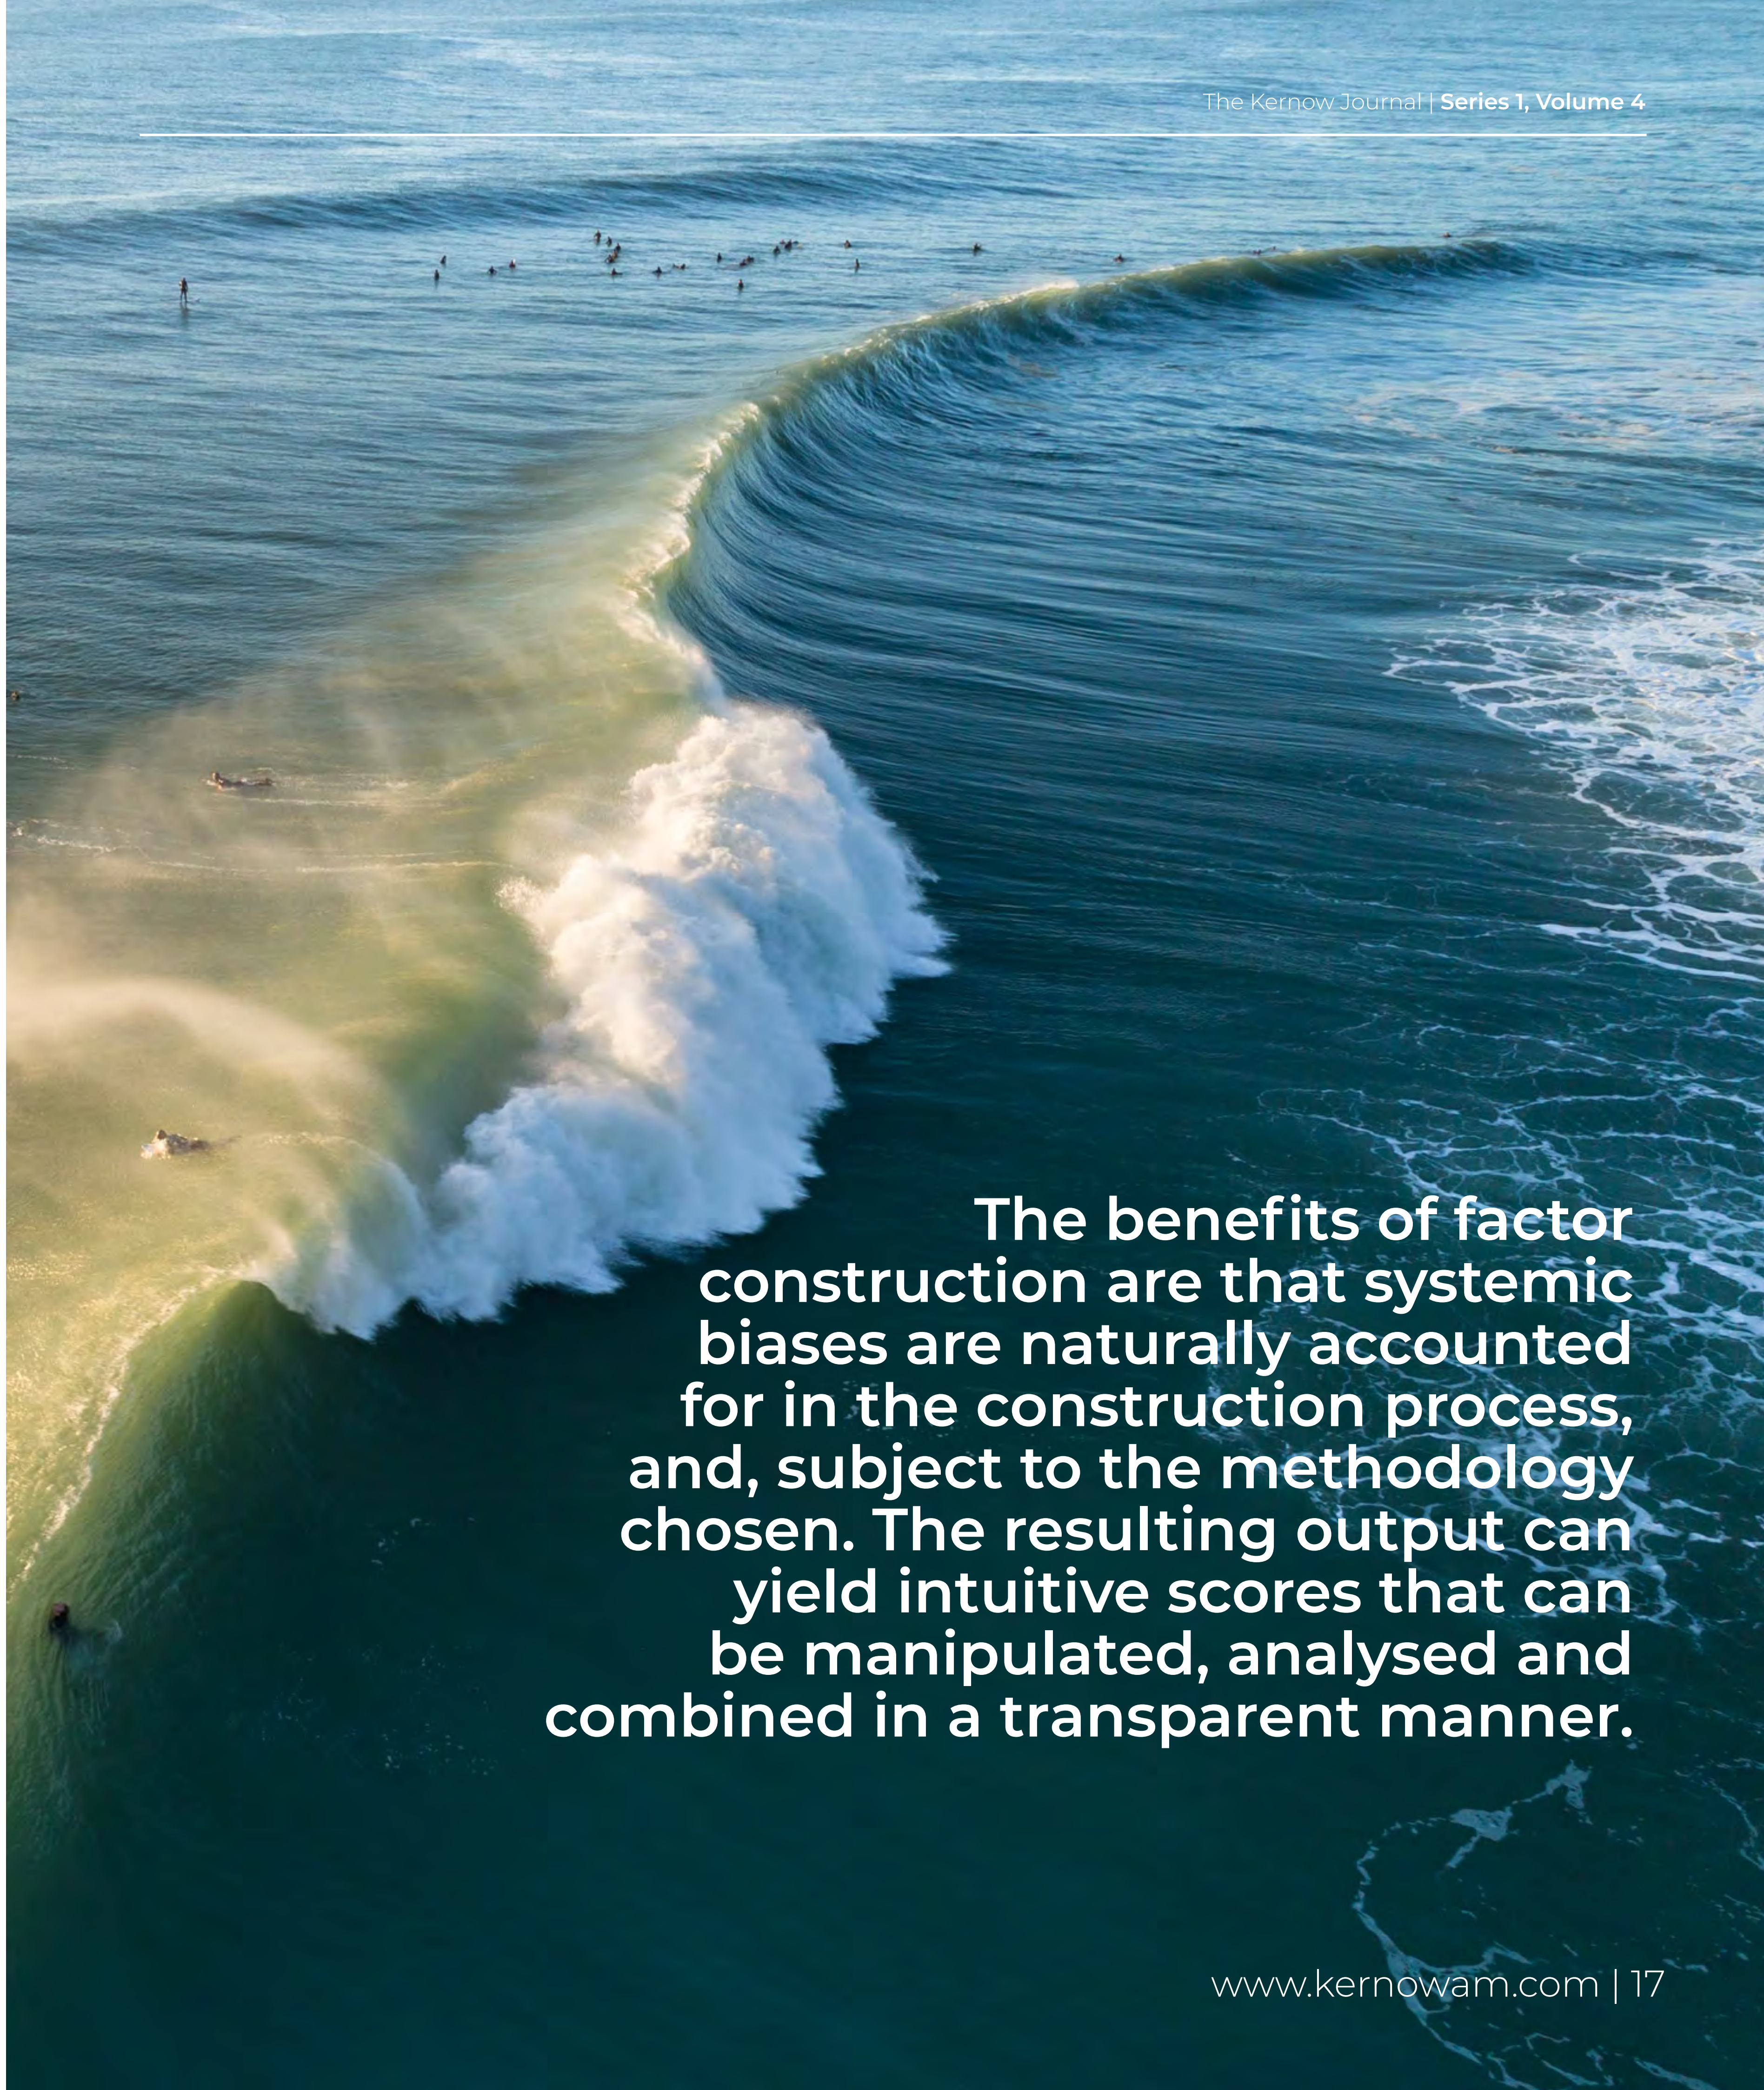

The benefits of factor construction are that systemic biases are naturally accounted for in the construction process, and, subject to the methodology chosen. The resulting output can yield intuitive scores that can be manipulated, analysed and combined in a transparent manner.

The Kernow VALUATION FRAMEWORK

What are the characteristics of the Kernow value factor?

By performing a statistical standardisation of the Kernow valuation model premia, which is calculated across all eligible securities on each date, we can see how the 'Kernow value factor' evolves, including its size and reliability. Overall, we find that the Kernow value factor is positive, meaning that the KVF is effective in explaining the relative fortunes of companies. In other words, the Kernow valuation model is more likely than not to be accurate in its assessment of true company valuations.

Examining the results of this factor analysis is helpful in many other ways. For example, we can gain insights into the historical under- or over-valuation of specific sectors and their evolution over time. We can also assess the relative value of the Kernow valuation framework by comparing the Kernow value factor to traditional style factors.

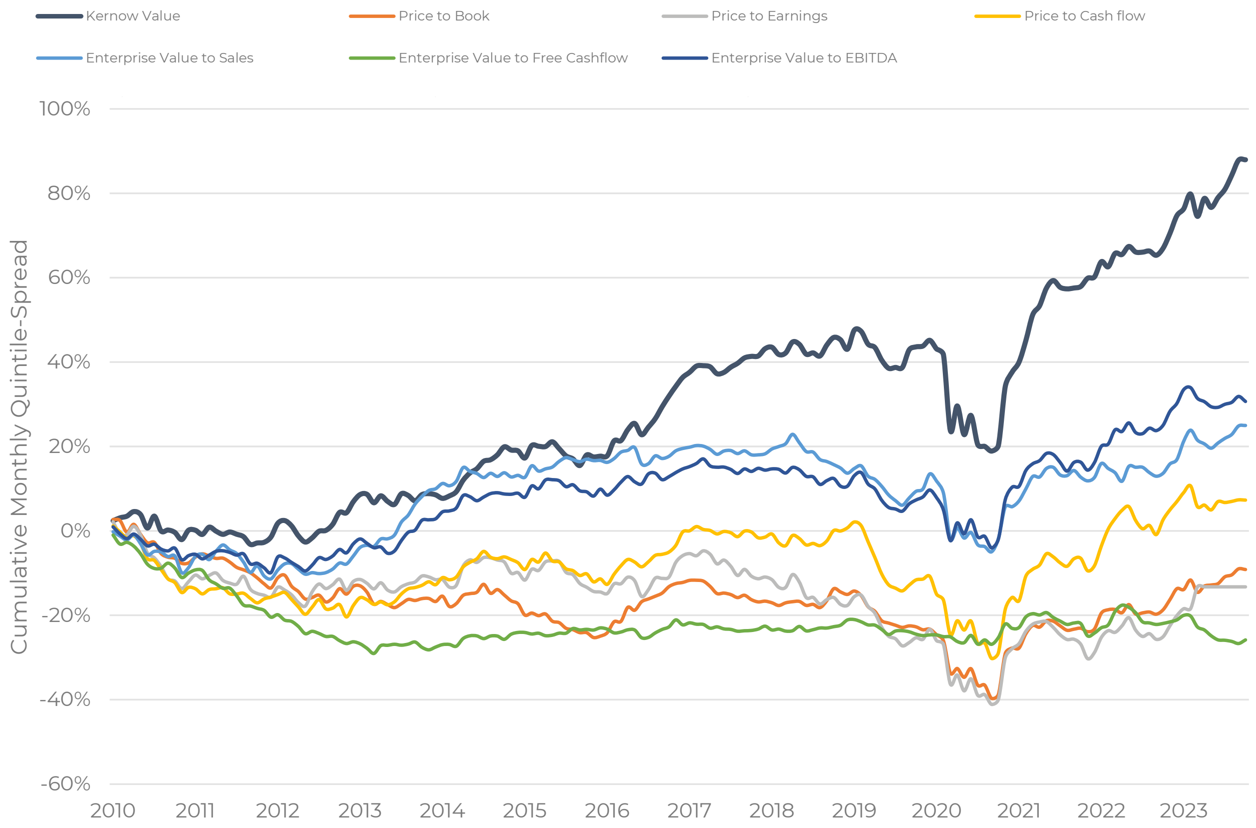

Figure 1: This figure shows the cumulative performance of each style factor, using the quintile spread measure of performance, which takes the average return of the top 20% of securities, as specified by the factor scores, and subtracts the average return of the bottom 20% of securities. In essence, this equates to tracking a cash-neutral portfolio with equally weighted holdings of the top 20% of securities on the long side and equally weighted shorts of the bottom 20% of securities. Thus, the net exposure is 200% and the pseudo-portfolio contains approximately 100 longs and 100 shorts.

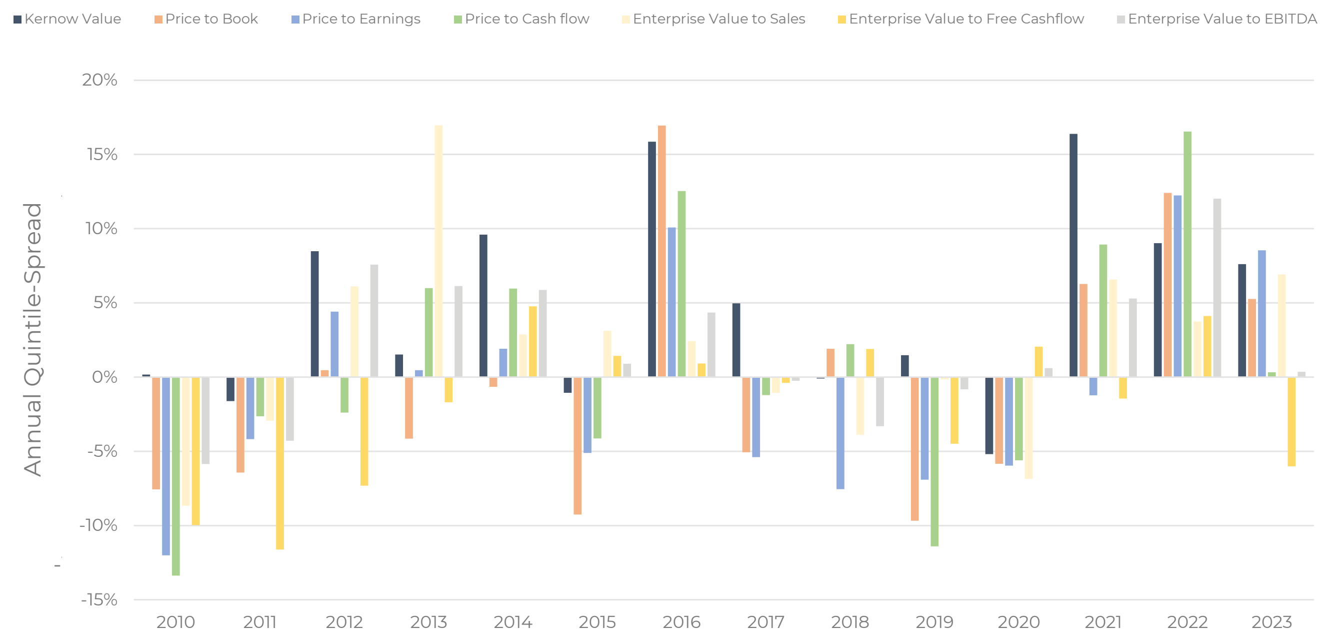

Figure 2: The annual quintile spread, calculated for the Kernow value factor as compared to the 'price-value' metrics (comprising price-to-book, price-to-earnings, and price-to-cashflow), and 'enterprise-value' metrics (comprising EV-to-sales, EV-to-free cashflow, and EV-to-EBITDA).

In Figure 1, we show the cumulative performance of the Kernow value factor compared to traditional style factors computed using standard valuation ratios as metrics and identical factor construction methodologies. The Kernow value factor performs significantly better than the other factors.

There has been a lot of discussion in the finance world about 'the lost decade' for value investors in the 2010-2020 period. Nevertheless, the Kernow valuation model still yields positive results. The period since COVID-19 (2021-present) has been notably strong for value investors, and the Kernow value factor has also excelled. We attribute a degree of this out-performance to the unique Kernow company valuation method.

In Figure 2, we compare the Kernow value factor on an annual basis against metrics based on price-valuation ratios or enterprise-value ratios. These are calculated using quintile spreads for the same universe of eligible companies using identical factor construction methods. The Kernow value factor maintains a limited drawdown with limited loss-making years, and positive performance in eight years and significant out-performance in three

We further illustrate the performance comparison in the table below. We compare the summary statistics on a long-term basis and a shorter-term (post-COVID-19) basis. Interestingly, whilst the conventional valuation metrics suffered a 'lost decade' from 2010-2023, the Kernow value factor returned a modest but encouraging 0.74 Sharpe ratio. Performance since 2021 has corresponded to a resurgence in the efficacy of value investing. As such, we have seen the Kernow value factor

	Annual Return	Annual Risk	Sharpe Ratio	Maximum Drawdown	Sortino Ratio
2010-2023					
Kernow Value Factor	**4.40%**	**6.00%**	**0.74**	**21.00%**	**1.05**
Price to Book	-1.40%	6.50%	-0.21	45.70%	-0.33
Price to Earnings	-1.50%	7.80%	-0.19	45.60%	-0.29
Price to Cash flow	0.30%	7.30%	0.04	37.60%	0.06
Enterprise Value to Sales	1.60%	5.00%	0.32	24.20%	0.46
Enterprise Value to Free Cashflow	-2.00%	4.60%	-0.43	31.20%	-0.67
Enterprise Value to EBITDA	2.30%	5.20%	0.43	20.30%	0.7
2021-2023					
Kernow Value Factor	**11.30%**	**5.30%**	**2.13**	**4.70%**	**3.6**
Price to Book	7.90%	6.60%	1.19	5.80%	1.91
Price to Earnings	5.10%	7.50%	0.68	15.90%	1.12
Price to Cash flow	10.60%	7.60%	1.4	7.40%	2.24
Enterprise Value to Sales	5.80%	5.20%	1.12	6.00%	1.78
Enterprise Value to Free Cashflow	1.40%	5.70%	0.25	7.80%	0.42
Enterprise Value to EBITDA	8.60%	5.00%	1.72	5.10%	2.94

Figure 3: The headline performance summary statistics across a long-term and more recent period, comparing the Kernow value factor to more commonly derived value factors. Data computed using daily returns.

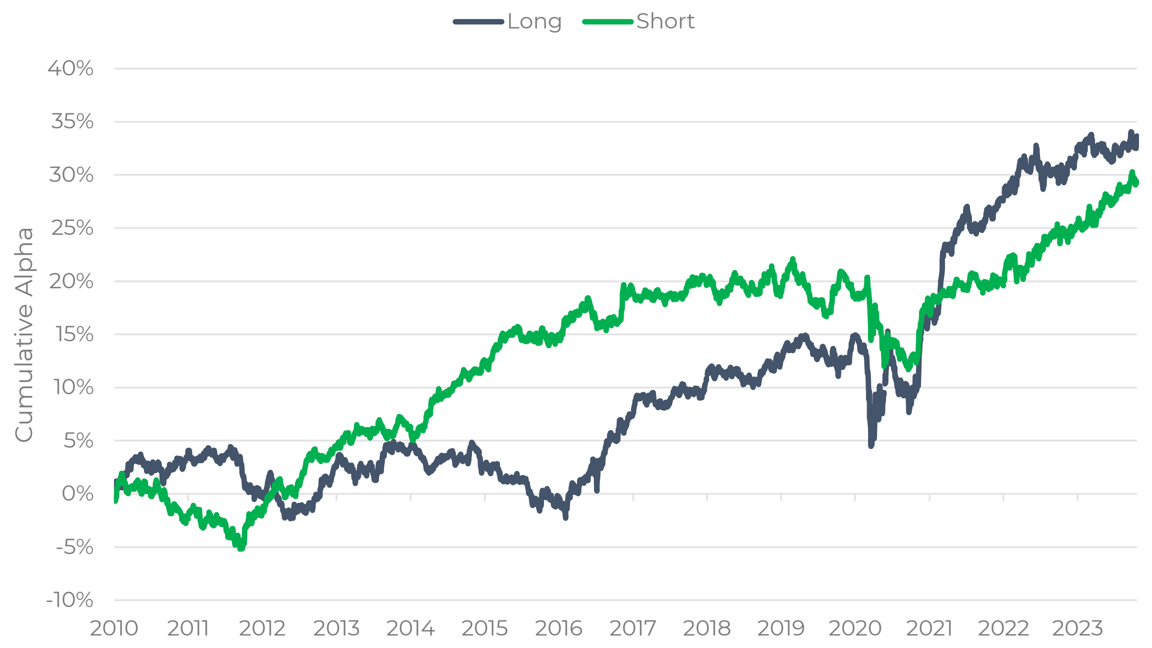

Figure 4: The long and short cumulative alpha (market-relative performance) of the Kernow value factor. Calculated with daily returns data and non-compounded. The short side is inverted for clarity.

consistently produce outsized returns, yielding an impressive Sharpe ratio above 2.

A critical aspect of value investing that receives less attention is understanding the underlying source of out-performance. Figure 4 shows the average returns of the top-third and bottom-third of companies in the KVF relative to the overall market. In other words, this figure shows the out-performance of undervalued companies to the market, as well as the under-performance of overvalued companies to the market. In summary, both the long and short sides of the Kernow valuation model consistently produce a similar amount of alpha.

To avoid cherry-picking, we also analyse the reliability of the valuation model for all companies in the UK market. Specifically, we focus on how securities perform in different tranches. The figure below shows the average annualised return over the analysis period for score quintiles.

Figure 5 shows a significant differential between the over-valued and under-valued companies, which indicates a broad positive expectancy of performance from the Kernow valuation model. Interestingly, there is a slight inflection in the most extreme under-valued companies. We interpret this as most likely to stem from companies ‘trapped’ in their under-valued states.

At Kernow we consider it essential to consider the potential for value traps when evaluating the investment prospects of a company that appears to be undervalued. It is important to note that this analysis primarily emphasises the pure effectiveness of the KVF.

In real-world applications, other crucial considerations include the company’s

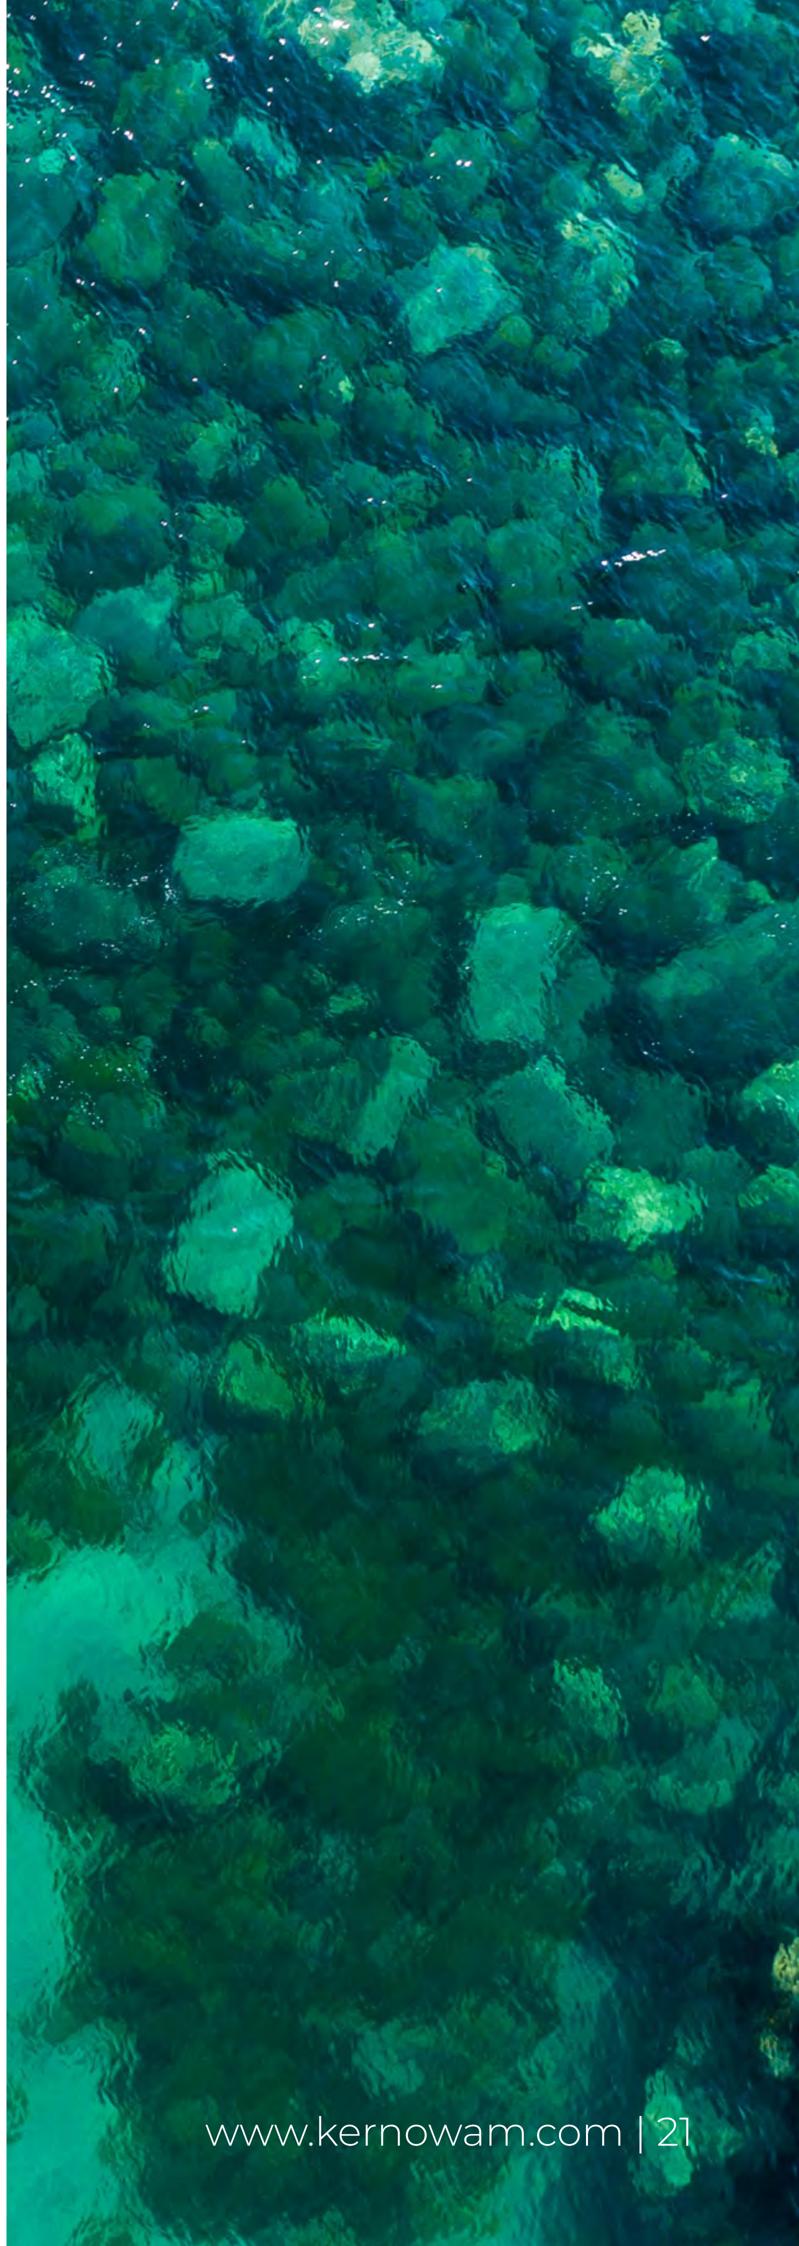

quality, strategic direction, prevailing investor sentiment, competitive landscape among peers, and how a specific security interacts with other holdings in a portfolio. These considerations significantly influence how this approach is put into practice. In our Navigator product, we take a holistic view when evaluating the suitability of a particular security within our portfolio and employ a catalyst-based approach to reduce the chance of falling into value traps.

Finally, to evaluate the efficacy of the Kernow value factor across different sectors, we consider each sector as a separate sub-portfolio with market-neutral net exposure. Figure 6 shows the average annual returns for each sub-portfolio. The heatmap shows a broad dispersion in fortunes across sectors and years. Overall, every single sector shows positive returns when using the Kernow valuation model over the cumulative period. There are additional risk management benefits from the dispersion of sector returns that could be accessed in portfolio construction.

Conclusion

The Kernow Valuation Framework has significant value, as demonstrated by the positive nature of the Kernow value factor and outperformance when compared to other factors.

This is a distinct competitive advantage for Kernow: under-valued companies, as measured by our proprietary valuation model, tend to outperform. Kernow's investment process incorporates information on catalysts that may accelerate the movement of companies towards their intrinsic valuation. Examining the historical efficacy of the Kernow Valuation Framework validates our approach beyond a simple screening tool.

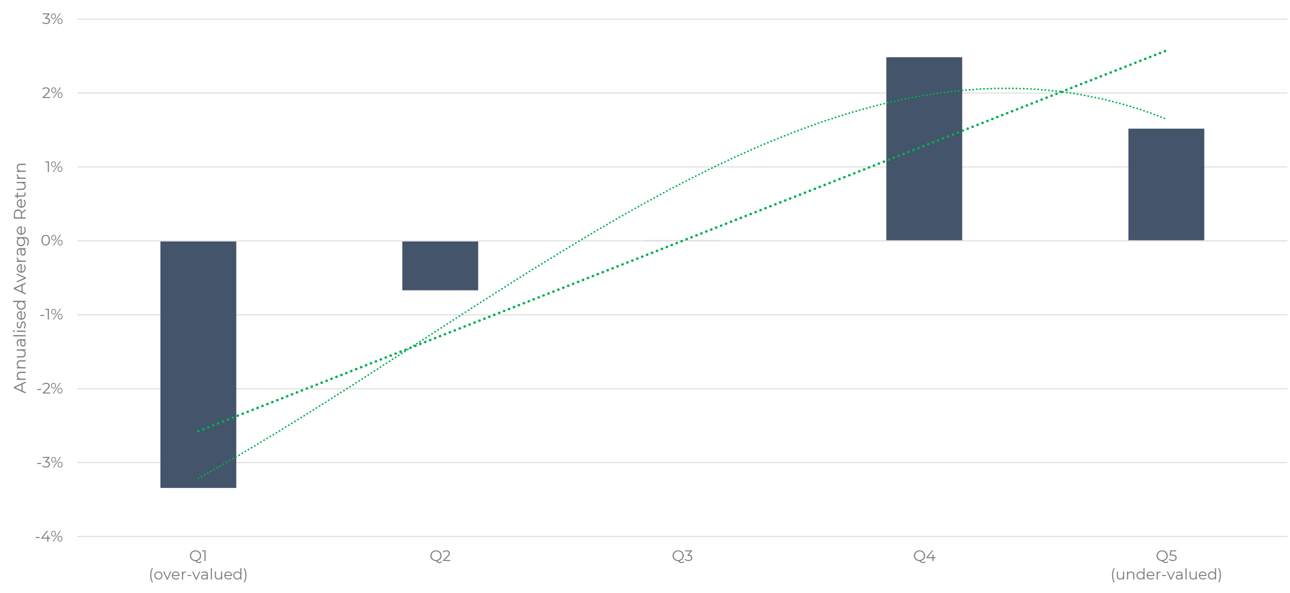

Figure 5: The quintile returns, showing the typical (annualised) performance of the most under and over-valued companies in the investable universe. The general shape of the cross-section is encouraging, but the inflection at Q5 indicates a potential value-trap effect in the most under-valued companies.

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	All Periods
Basic Materials	27%	-4%	-27%	-38%	-11%	-32%	111%	18%	10%	-18%	14%	20%	51%	10%	9%
Consumer Cyclicals	-11%	0%	9%	-4%	9%	2%	-3%	4%	-2%	7%	-21%	29%	7%	9%	2%
Consumer Non-Cyclicals	-3%	-10%	-6%	-5%	-7%	2%	22%	-12%	7%	-9%	-4%	33%	29%	35%	5%
Energy	-10%	33%	0%	31%	-3%	-29%	-26%	2%	-16%	-4%	-30%	25%	-7%	41%	1%
Financials	6%	-7%	13%	1%	0%	4%	12%	6%	-1%	7%	3%	10%	-7%	5%	4%
Healthcare	18%	-3%	-7%	18%	22%	58%	84%	17%	19%	16%	-7%	9%	109%	54%	29%
Industrials	-5%	9%	13%	10%	27%	3%	16%	1%	-6%	-5%	-25%	6%	21%	4%	5%
Real Estate	-3%	13%	3%	10%	10%	-8%	19%	13%	2%	-5%	14%	9%	10%	-9%	6%
Technology	-19%	-32%	29%	20%	30%	-6%	15%	-1%	8%	-9%	1%	30%	37%	0%	7%
Utilities	-17%	21%	21%	-20%	94%	4%	-9%	6%	-8%	9%	-26%	24%	14%	2%	8%

Figure 6: The annual performance of sub-factors of the Kernow value factor from 2010-2023. Highlighting the array of performance fortunes of these subsets over time and the diversification benefits of a broad universe.

Research
HIGHLIGHT

In partnership with University of Exeter

Regulatory News Service: A Critical Component to UK Investing

Introduction

As investors in publicly traded companies, we depend on various information channels to craft a thoroughly informed evaluation of a specific company for our investment goals.

Data can be collected from diverse sources, including news articles, corporate disclosures, regulatory filings, presentations, transcriptions of investor calls, social media and various other channels. The volume of financial data is continuously rising with increasing online information availability. While most of this information is presented in conventional tabular and quantitative formats, valuable insights can also be derived from other sources such as text, images, recordings, location data, or more complex and less readily processed formats.

Within the extensive array of datasets accessible to investors, we can categorise the primary data into several overarching segments. These include market-related information, fundamental company data, analyst forecasts, industry-specific key performance indicators, in-depth product information, as well as stock ownership, peer analysis and data concerning notable stakeholders.

The UK Regulatory News Service (RNS) dataset is the main means by which UK companies communicate with their investors. It contains a vast amount of information, which is available to everyone. But how can we use our time efficiently to manage that flow of data, spot trends and areas of interest for further study? The timing and sector characteristics of this dataset provides interesting insights.

The process of analysing and extracting meaningful insights from text data presents a multitude of challenges. Financial texts often exhibit unique characteristics in terms of their structure, differing from generic text sources. Text-based datasets frequently encompass a mix of quantitative and qualitative information. Additionally, documents may include structural elements such as graphics, embedded tabular data, distinctive sections, chapters and document-specific content.

Alongside the acceleration in the availability of text-based datasets, the field of natural language processing

(NLP) has developed rapidly. There are a growing number of applications of NLP, including the calculation of sentiment, linguistic complexity, entity extraction, language translation and document summarisation.

Whilst certain NLP techniques are deceptively basic, in recent years, more sophisticated large language models (LLMs) based on state-of-the-art machine learning approaches have gained significant notoriety in their ability to parse text-based information and retrieve insights by learning from vast corpora.

For this analysis, we worked with the University of Exeter Mathematics Department to explore several applications of NLP methods across a historical archive of the RNS dataset. We have summarised some of the pertinent findings from our initial research.

The Regulatory News Service: A Primer

For UK-listed companies, the RNS is a regulatory and financial news dissemination service operated by the London Stock Exchange.

Its primary purpose is to facilitate the timely and accurate distribution of regulatory announcements and company news to the financial markets and the investing public.

Publicly listed companies in the UK must make specific regulatory announcements, such as financial results, director appointments, major acquisitions or disposals and other events that may materially impact their share price or operations. These announcements are usually made to keep shareholders and the broader market informed about the company's activities and financial performance.

The RNS platform serves as a central repository for these regulatory announcements, ensuring that all market participants have simultaneous access to crucial information. It is crucial in promoting transparency, fairness and efficiency in the UK financial markets.

As specialists in analysing UK-listed companies, the RNS dataset presents us with a distinctive information channel. Across most geographies, company disclosures are made during scheduled regulatory filings, reported ad-hoc on corporate websites or news articles. The RNS mechanism provides a timely, accurate, transparent and standardised disclosure channel.

It is generally true to say that investors focused on UK-listed companies acquire a significant proportion of their real-time information from the RNS service. Typically, most RNS articles are released at 7am GMT, an hour before the UK market opens, providing investors with a window to digest information before the day's trading commences.

A note of caution, however. Although regulatory obligations require certain material to be published on the RNS, these disclosures are curated by the company submitting them. As such, a degree of creative wording may be used to distract, obfuscate or sugar-coat information. There can also be lively discussions between company management and advisors, debating what material should be included in an RNS announcement, particularly when facing negative news.

With our decades of experience, it is possible to read between the lines of RNS submissions to uncover hidden patterns, or a lack of disclosure. Our aim in this analysis is to provide quantitative rigour to reveal subtle patterns and evolutions of these across the bulk RNS dataset.

Data Preliminaries - Filtering

Our dataset contains 1,822 companies in the FTSE All Share index with a market capitalisation over £100m.

The RNS dataset comprises over 326,000 RNS news articles from 2002 through 2022, corresponding to 5,307 unique dates.

Given the diverse disclosure of RNS information, we consider the dataset to be 'semi-structured'. As several of the RNS articles in this dataset are semi-automated, text analysis thereof is unlikely to yield meaningful insight. These submissions generally relate to changes in share ownership, price monitoring extensions or other standard regulatory reporting obligations. We filter the dataset to omit these RNS articles in our analysis, yielding 115,460 relevant RNS submissions (the Dataset).

Although significant detail is available within the text of RNS articles, they have no formal classification: only the text body and subject line indicate their broad taxonomy. We have therefore built our own classification scheme to categorise RNS articles into six categories. These are as follows.

1 Results

Relating to interim and full year results, financial statements and dividend declarations.

2 Trading Updates

Relating to news relevant to the operations of the company and a financial outlook statement.

3 M&A Activity

Relating to corporate activity, including share offers, acquisitions, mergers, tenders and other corporate actions.

4 Management Changes

Relating to changes in senior management and board members.

5 Capital Markets Activity

Relating to financing, raising capital, loan details and credit facilities.

6 Contracts & Partnerships

Relating to financing, raising capital, loan details and credit facilities.

Purchase of Own Shares	TR-1	Form 8.1	AIM Rule 17	Appendix 5b
Transaction in Own	Price Monitoring	Form 8.2	Listing Rule 15.6	Scrip
Exercise of Options	Grant of Share Awards	Form 2.9	Block Listing	Listing Rule

Figure 7: Title-fragments corresponding to the automated RNS events that we exclude from our text analysis.

To implement our categorisations, we apply an algorithm across RNS subject texts using a set of key phrases associated with each category. We aim to maximise the matching of each article with one or more categories without incurring significant false-positive matches. Any RNS article whose subject does not have a sufficient match is associated with a specific 'miscellaneous' category.

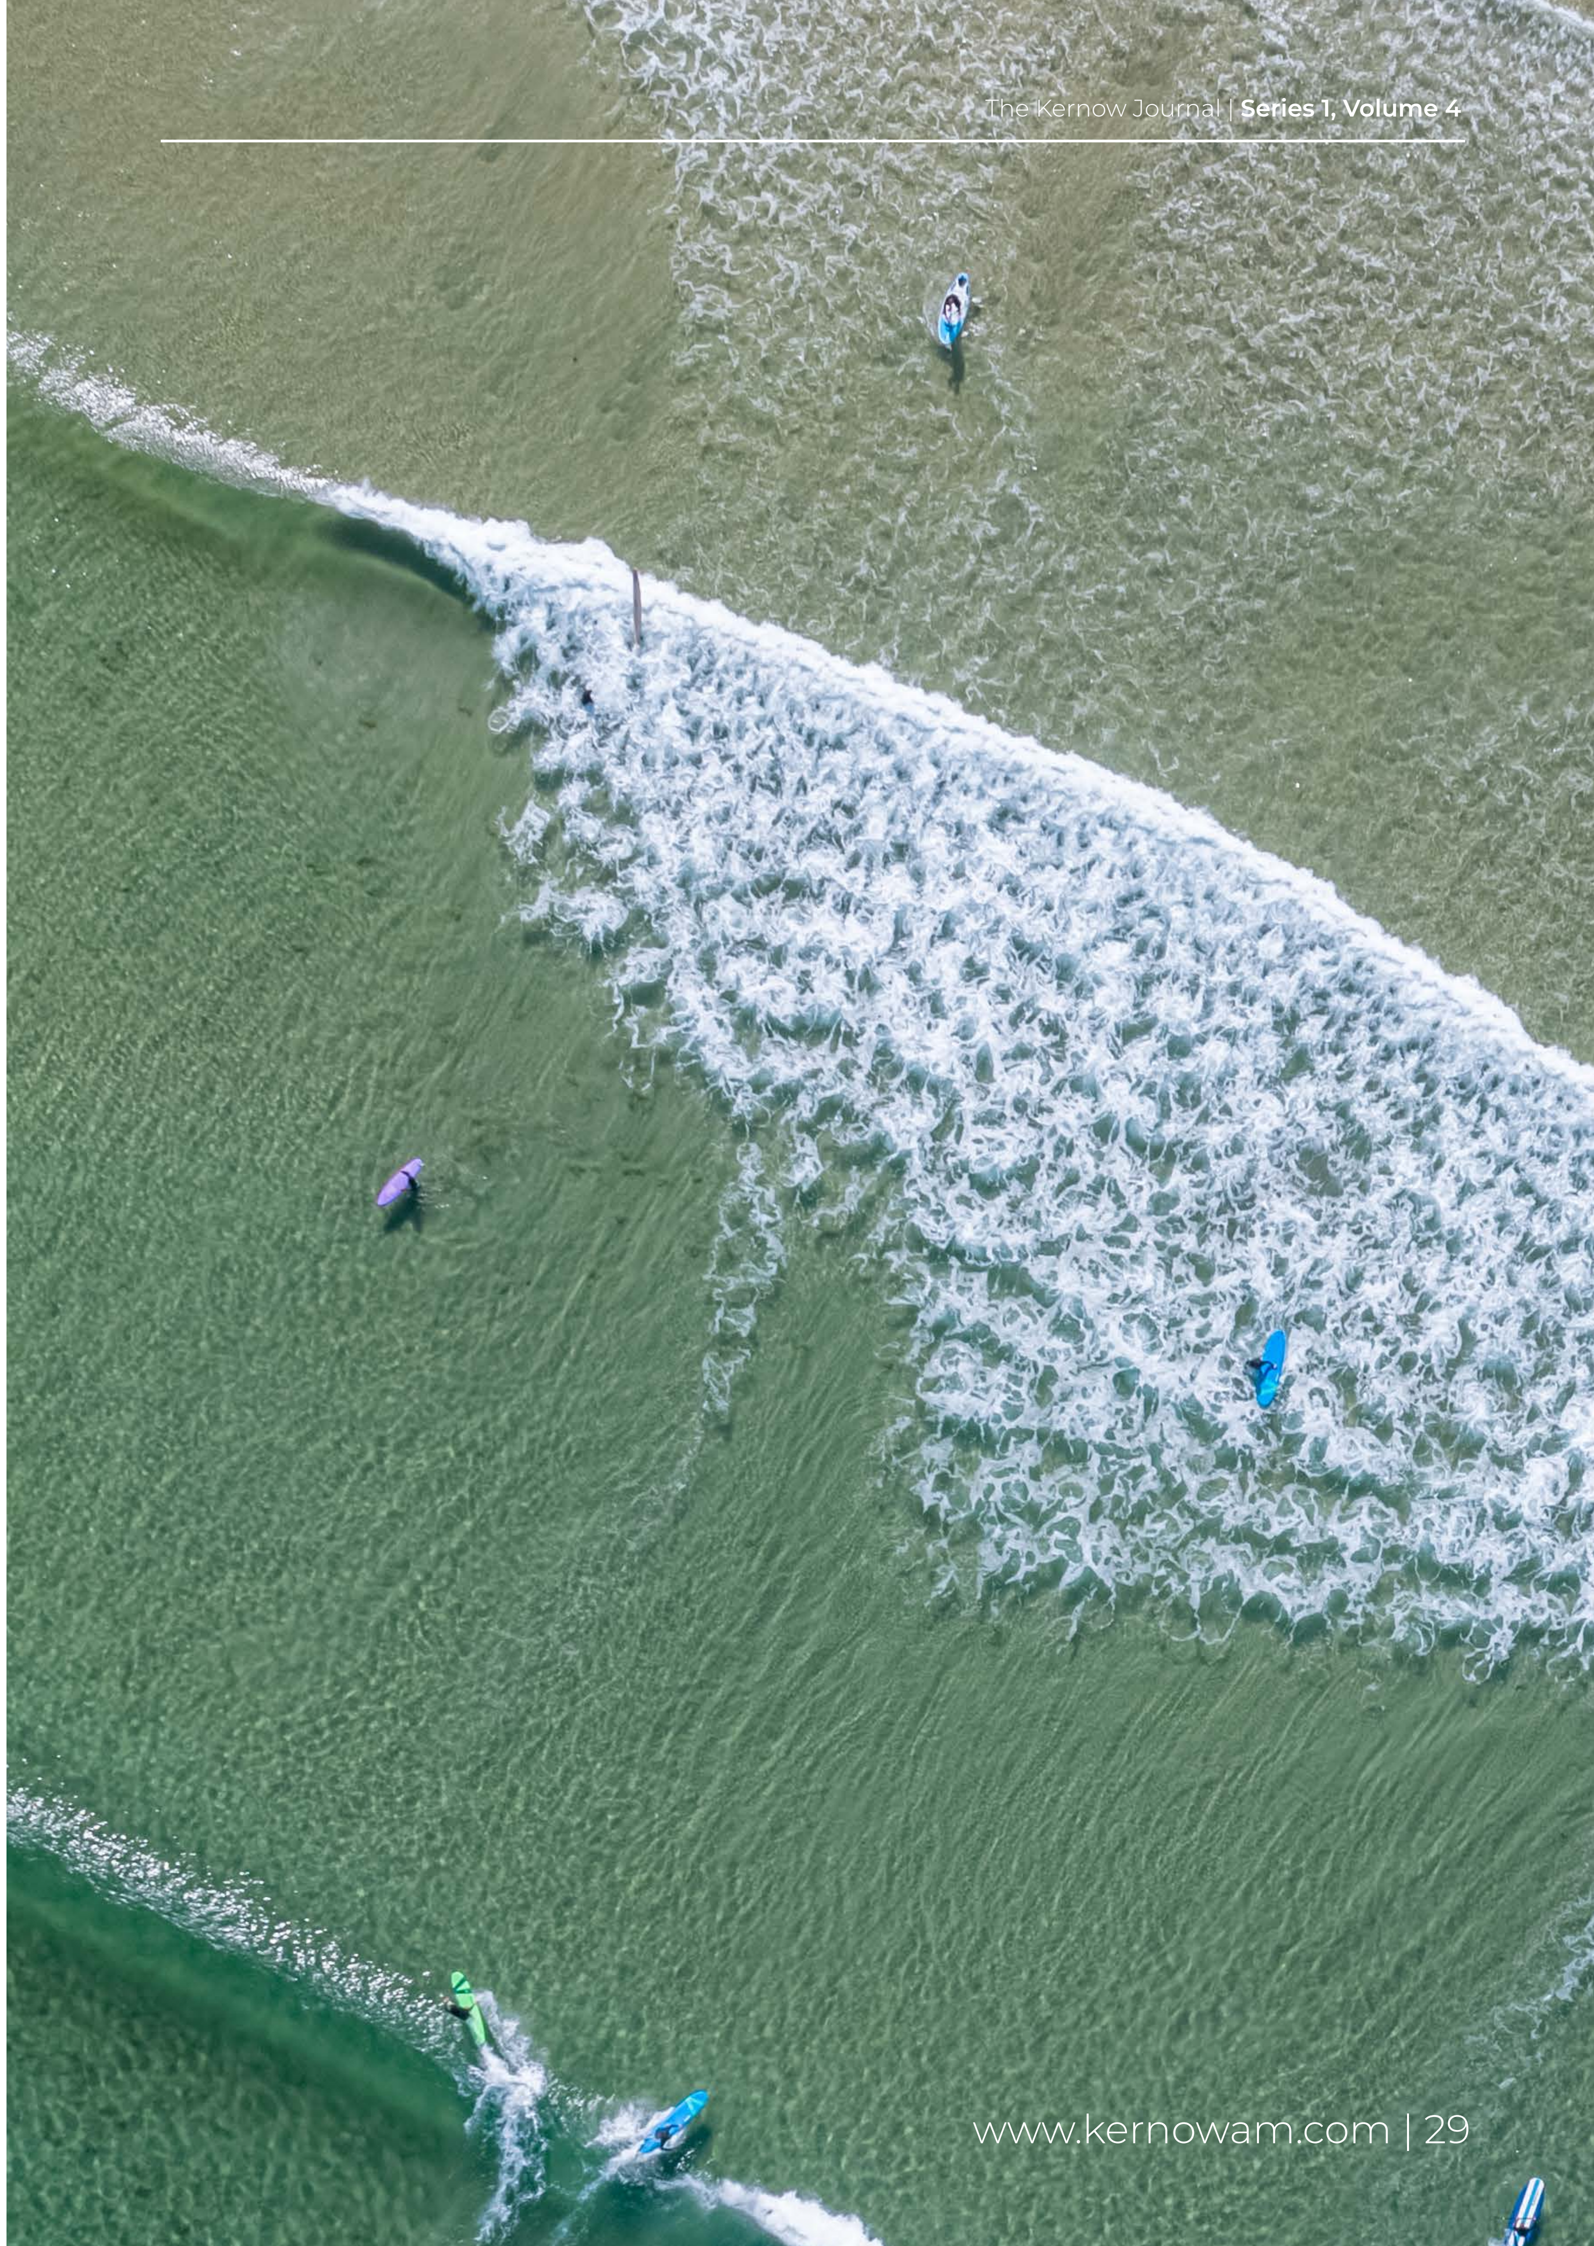

Research **HIGHLIGHT**

In partnership with University of Exeter

Dataset Characteristics

Our analysis has highlighted certain characteristics of the RNS dataset, which are helpful for the management of the flow of data in real-time.

Not only does this assist the efficient use of investor time, but it also enables one to draw insights from company behaviour that falls outside of the norm, such as issuing RNS statements at unusual times or with unusual frequency.

First, we assess the timing of RNS releases. In Figure 8, we show that there is a seasonal pattern to RNS submissions,. which tends to peak during earnings season. It also shows a spike in RNS submissions during the onset of the Covid-19 pandemic and a small upward trend across the period which reflects the increase in universe coverage.

To expand on our seasonality analysis, we also assess seasonality on three scales. Figure 9 displays the modest seasonality of RNS submissions in the Dataset regarding the time of the day, day of the week and month. The time-of-day seasonality is particularly interesting; Figure 9 shows a spike representing a 'bulk daily update' at 7am. It is interesting to speculate whether RNS submissions out of market-hours are likely to have a significant chance of delivering negative news than conventionally-scheduled news events.

Next, we consider the typical frequency of RNS submissions for companies within different economic sectors. Figure 10 shows that healthcare companies tend to have a higher number of RNS submissions, perhaps equating to more frequent disclosures of medical trials and other industry-specific news. On the other hand, we find that consumer cyclical companies tend to submit less frequently, indicating their more 'business as usual' approach to company operations.

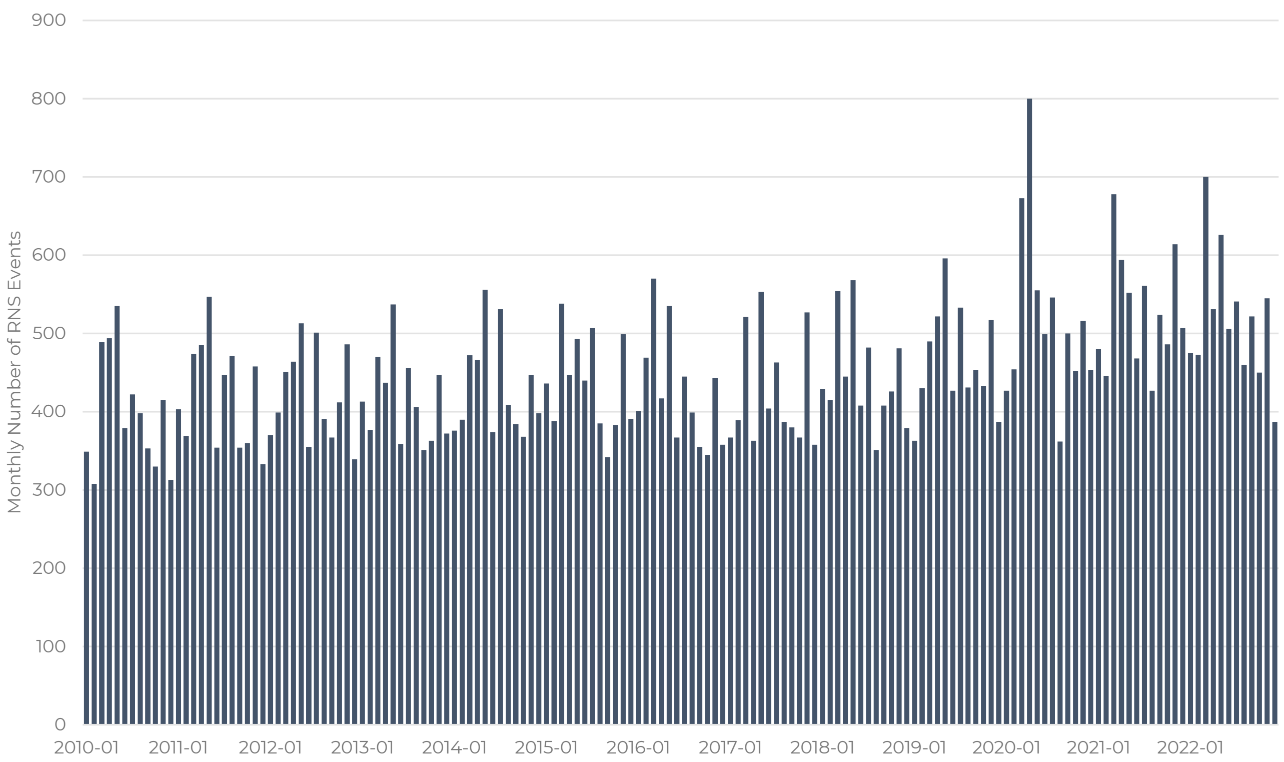

Figure 8: The monthly number of RNS articles in the Dataset, 2010-2022

Event Studies

Is the mere presence of an RNS announcement likely to lead to a period of out-performance for a particular company?

In Figure 11 we investigate this question by showing the average out-performance of companies in the Dataset in the week after submitting an RNS, grouped by RNS category and company sector. We see several interesting characteristics. For example, RNS announcements relating to the Contracts & Partnerships and Capital Market Activity categories' yield modestly positive post-announcement drift, whilst those relating to the Management Change category yield negative drift. This is somewhat unsurprising, as we may naturally expect news of contracts to be positively viewed by investors and news about management changes to be considered in a more cautious light, at least in the short-term.

We also find differences across sectors, with RNS announcements for utilities companies tending to be followed by positive drift, irrespective of their RNS category, which may indicate

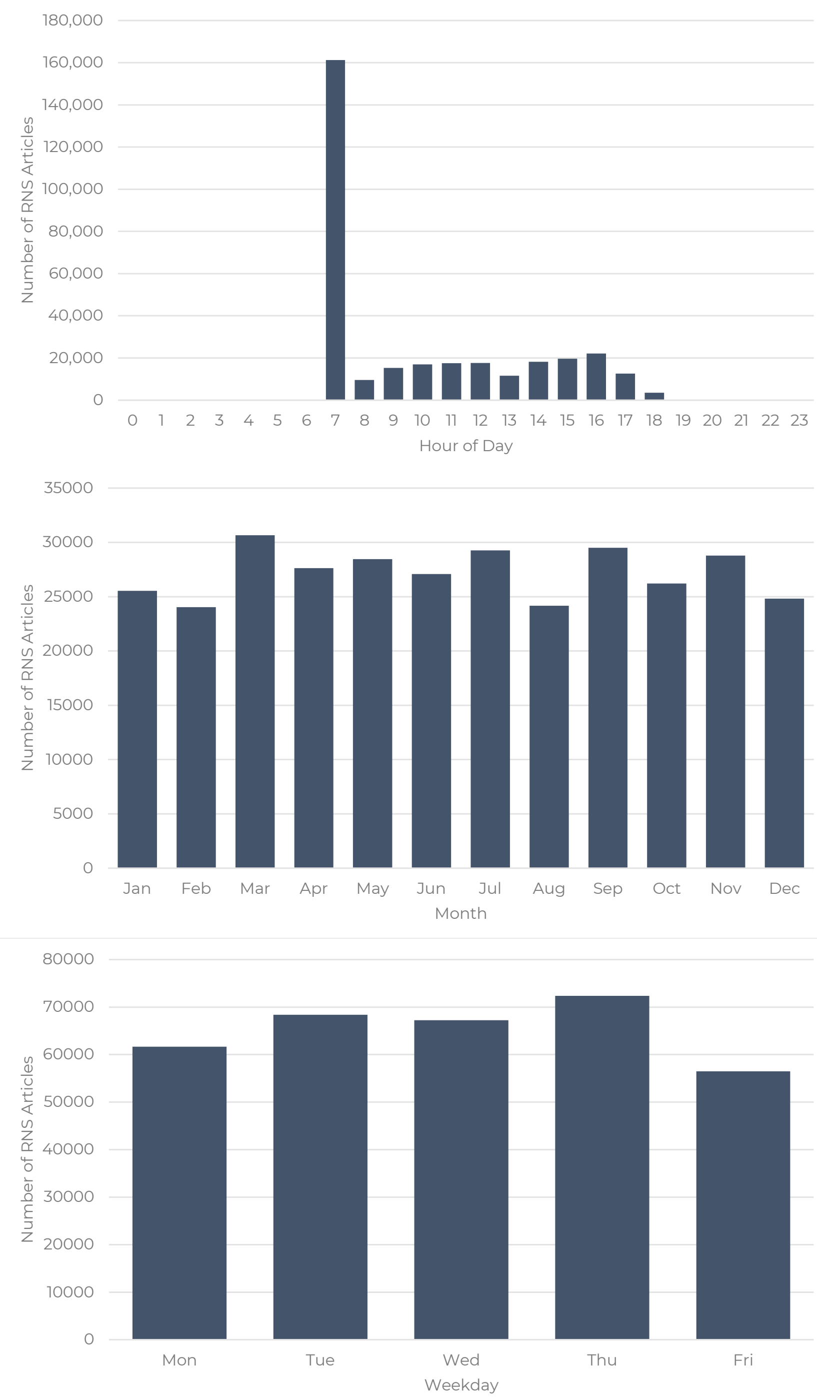

Figure 9: The RNS submissions are highly seasonal. Typically, they have been submitted during the 7am pre-market update, with more updates in mid-week times, as opposed to Mondays and Fridays, and during earnings seasons across the calendar.

Research
HIGHLIGHT

In partnership with University of Exeter

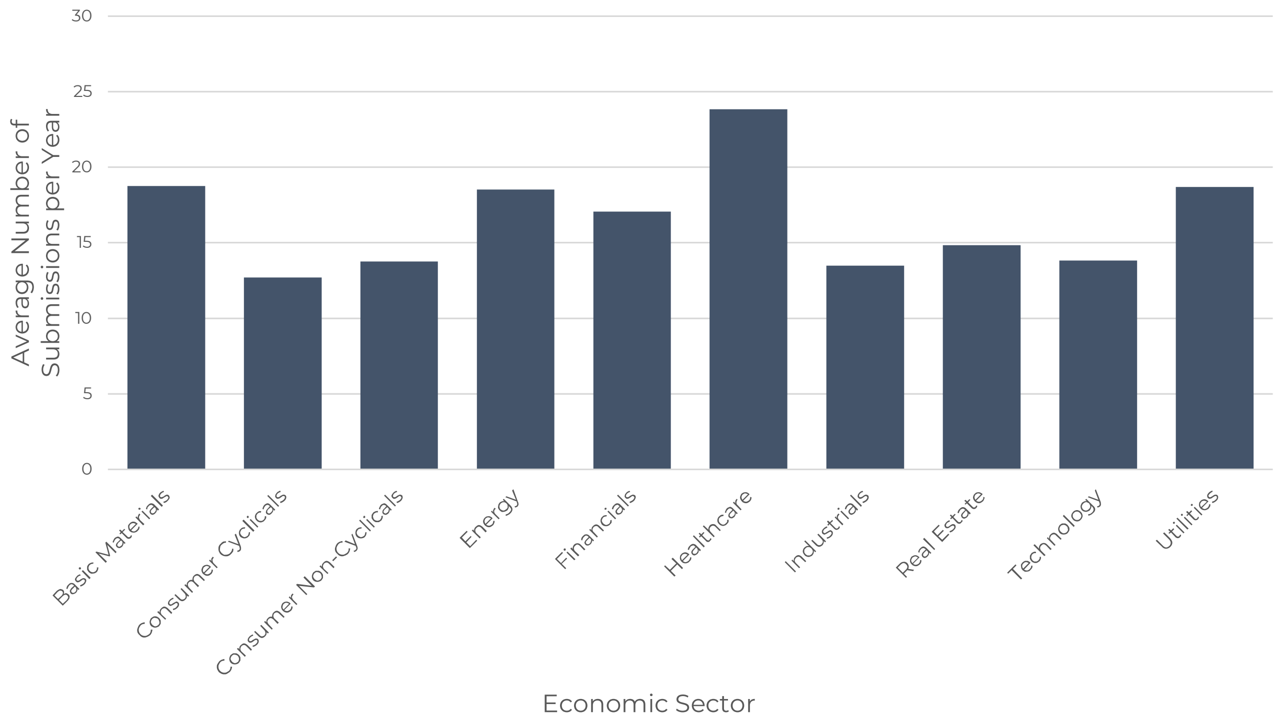

Figure 10: The typical number of RNS submissions per year in the Dataset, for securities partitioned by economic sector

relief rallies around news announcements. We see broad dispersions in effect across RNS categories for the Energy and Healthcare sectors, indicating that these RNS articles may indeed contain market-moving information.

Whilst this basic review of the price-action around RNS events helps build intuition, there is a vast amount of helpful information within the RNS article text that can provide further insights. The next step is to consider how the content of the RNS articles and other metadata around their submission can provide hints about the underlying nature of the disclosure and how it is likely to be perceived by market participants.

RNS Sentiment Analysis

The underlying sentiment contained within an RNS article may be useful in determining whether the information being disclosed by a company is broadly favourable or unfavourable to its operations.

There are a vast number of NLP techniques that have been developed to extract the level of sentiment from documents. For this initial study, we use a bag-of-words sentiment measure which tallies the frequency of word matches to the Loughran & McDonald (2011) keyword dictionary of positive and negative words. Despite several known limitations of this approach, the dimensionless sentiment ratio derived from this has been shown in other contexts to be able to broadly quantify the level of sentiment. Figure 12 shows the typical sentiment calculated for different RNS categories and sectors in the Dataset. Immediately, we notice the effects of self-disclosure of RNS announcements - specifically, we find that articles relating to management changes have positive sentiment.

This is not surprising, as these RNS submissions generally celebrate a departing individual or evangelise an incoming individual, irrespective of the underlying rationale for the personnel changes. We also see broadly positive sentiment in RNS articles in the Contracts & Partnerships category. This self-disclosure nature of RNS articles naturally results in a bias towards positive sentiment within these articles. Interestingly, we also find that Energy companies tend to be distinct in their sentiment characteristics.

Whilst the aggregate level of sentiment reveals interesting details about the nature of the information within them, considering the impact in the cross-section of different sentiment, for given RNS categories is something that we can use potentially as a leading indicator in our investment process.

Correspondingly, Figure 13 shows the tertile spread, a simple performance measure, for RNS articles, partitioned by their category and shown across three time-horizons. Interestingly, we see that sentiment is informative in the cross-section for most RNS categories, except for Contracts & Partnerships and Management Change. Perversely, we find that the highest sentiment RNS for Management changes lead to the worst subsequent price performance. Despite these quirks (which we may attribute to self-disclosure effects), we find that sentiment is broadly connected with positive performance on a relative basis, although the results are somewhat muted.

	Basic Materials	Consumer Cyclicals	Consumer Non-Cyclicals	Energy	Financials	Industrials	Real Estate	Technology	Utilities	Healthcare	
Capital Market Activity	0.19%	0.40%	0.70%	0.68%	0.16%	0.97%	0.25%	0.11%	1.58%	1.02%	0.61%
Contracts & Partnerships	0.27%	0.79%	1.06%	1.36%	0.26%	2.19%	0.54%	0.48%	0.92%	-0.30%	0.76%
M&A	0.96%	0.42%	0.81%	-0.15%	0.03%	0.14%	1.02%	0.18%	1.19%	0.75%	0.53%
Management Changes	0.24%	-0.32%	-0.10%	-0.75%	-0.07%	-0.13%	-0.17%	0.18%	0.34%	-0.25%	-0.10%
Misc.	-0.15%	0.29%	-0.30%	-0.08%	-0.61%	0.49%	0.24%	0.00%	0.20%	-0.14%	-0.01%
Results	0.33%	0.42%	0.39%	0.42%	-0.31%	-0.16%	0.46%	-0.01%	1.08%	0.34%	0.30%
Trading Update	-0.20%	0.12%	-0.16%	-0.97%	-0.16%	0.07%	0.05%	0.21%	0.96%	-0.10%	-0.02%
	0.24%	0.30%	0.34%	0.07%	-0.10%	0.51%	0.34%	0.16%	0.90%	0.19%	

Figure 11: The typical idiosyncratic returns on a 1-week horizon post-RNS submission. Partitioned by the RNS category and the economic sector of the company.

	Basic Materials	Consumer Cyclicals	Consumer Non-Cyclicals	Energy	Financials	Industrials	Real Estate	Technology	Utilities	Healthcare	
Capital Market Activity	0.06	0.03	0.03	-0.21	-0.04	0.08	-0.09	-0.16	-0.04	-0.31	-0.07
Contracts & Partnerships	0.41	0.43	0.43	-0.16	0.40	-0.09	0.43	0.26	0.26	0.09	0.24
M&A	-0.06	0.00	0.16	-0.14	-0.05	0.04	-0.01	0.05	0.01	0.06	0.00
Management Changes	0.58	0.60	0.57	0.40	0.44	0.55	0.46	0.54	0.61	0.39	0.52
Misc.	0.06	0.09	0.15	-0.19	0.02	0.13	0.10	0.13	-0.19	-0.01	0.03
Results	0.08	0.19	0.08	-0.15	-0.09	0.10	-0.08	0.07	-0.12	0.05	0.01
Trading Update	0.07	0.17	0.25	-0.18	0.03	0.09	0.08	0.11	0.08	0.07	0.08
	0.17	0.22	0.24	-0.09	0.10	0.13	0.13	0.14	0.09	0.05	

Figure 12: The average sentiment for RNS articles, calculated on a scale -1 to +1, partitioned by economic sector and RNS Category

Figure 13: The average tertile-spread, calculated as the difference in median return for securities in the top vs bottom third of the investable universe, ranked on RNS sentiment. Results shown partitioning by the RNS News Category, shown across three time-horizons.

RNS Readability Analysis

Another interesting characteristic of text data that has been extensively studied is readability.

The readability of a passage of text simply refers to the ease in which the informational content of the text can be extracted. In the context of financial disclosures, readability is crucial as the purpose of RNS articles is to aid transparency by disclosing relevant information. Nevertheless, the self-disclosed nature of RNS submissions may result in companies obfuscating negative information.

There is a rich history of studying the readability of text data, spanning over a century. Several straightforward formulae have been proposed that broadly measure the complexity of a passage of text. These include the Flesche Reading Ease Index, the Gunning-Fogg Index, the Smog Index and the Automated Readability Index.

Whilst these measures differ in their detail, they have a generally high correspondence. We focus on the Gunning-Fogg index (GFI) methodology, which uses the average sentence length and the percentage of complex words within a passage of text as a proxy for its complexity. For example, a typical children's story would have a GFI of approximately 4-7, newspapers would generally have a GFI of 8-12 and highly technical articles would have GFI of 20-25.

Interestingly, the readability indices across all relevant RNS articles each have correlations over 0.8 and we generally find a slight negative relationship between sentiment,

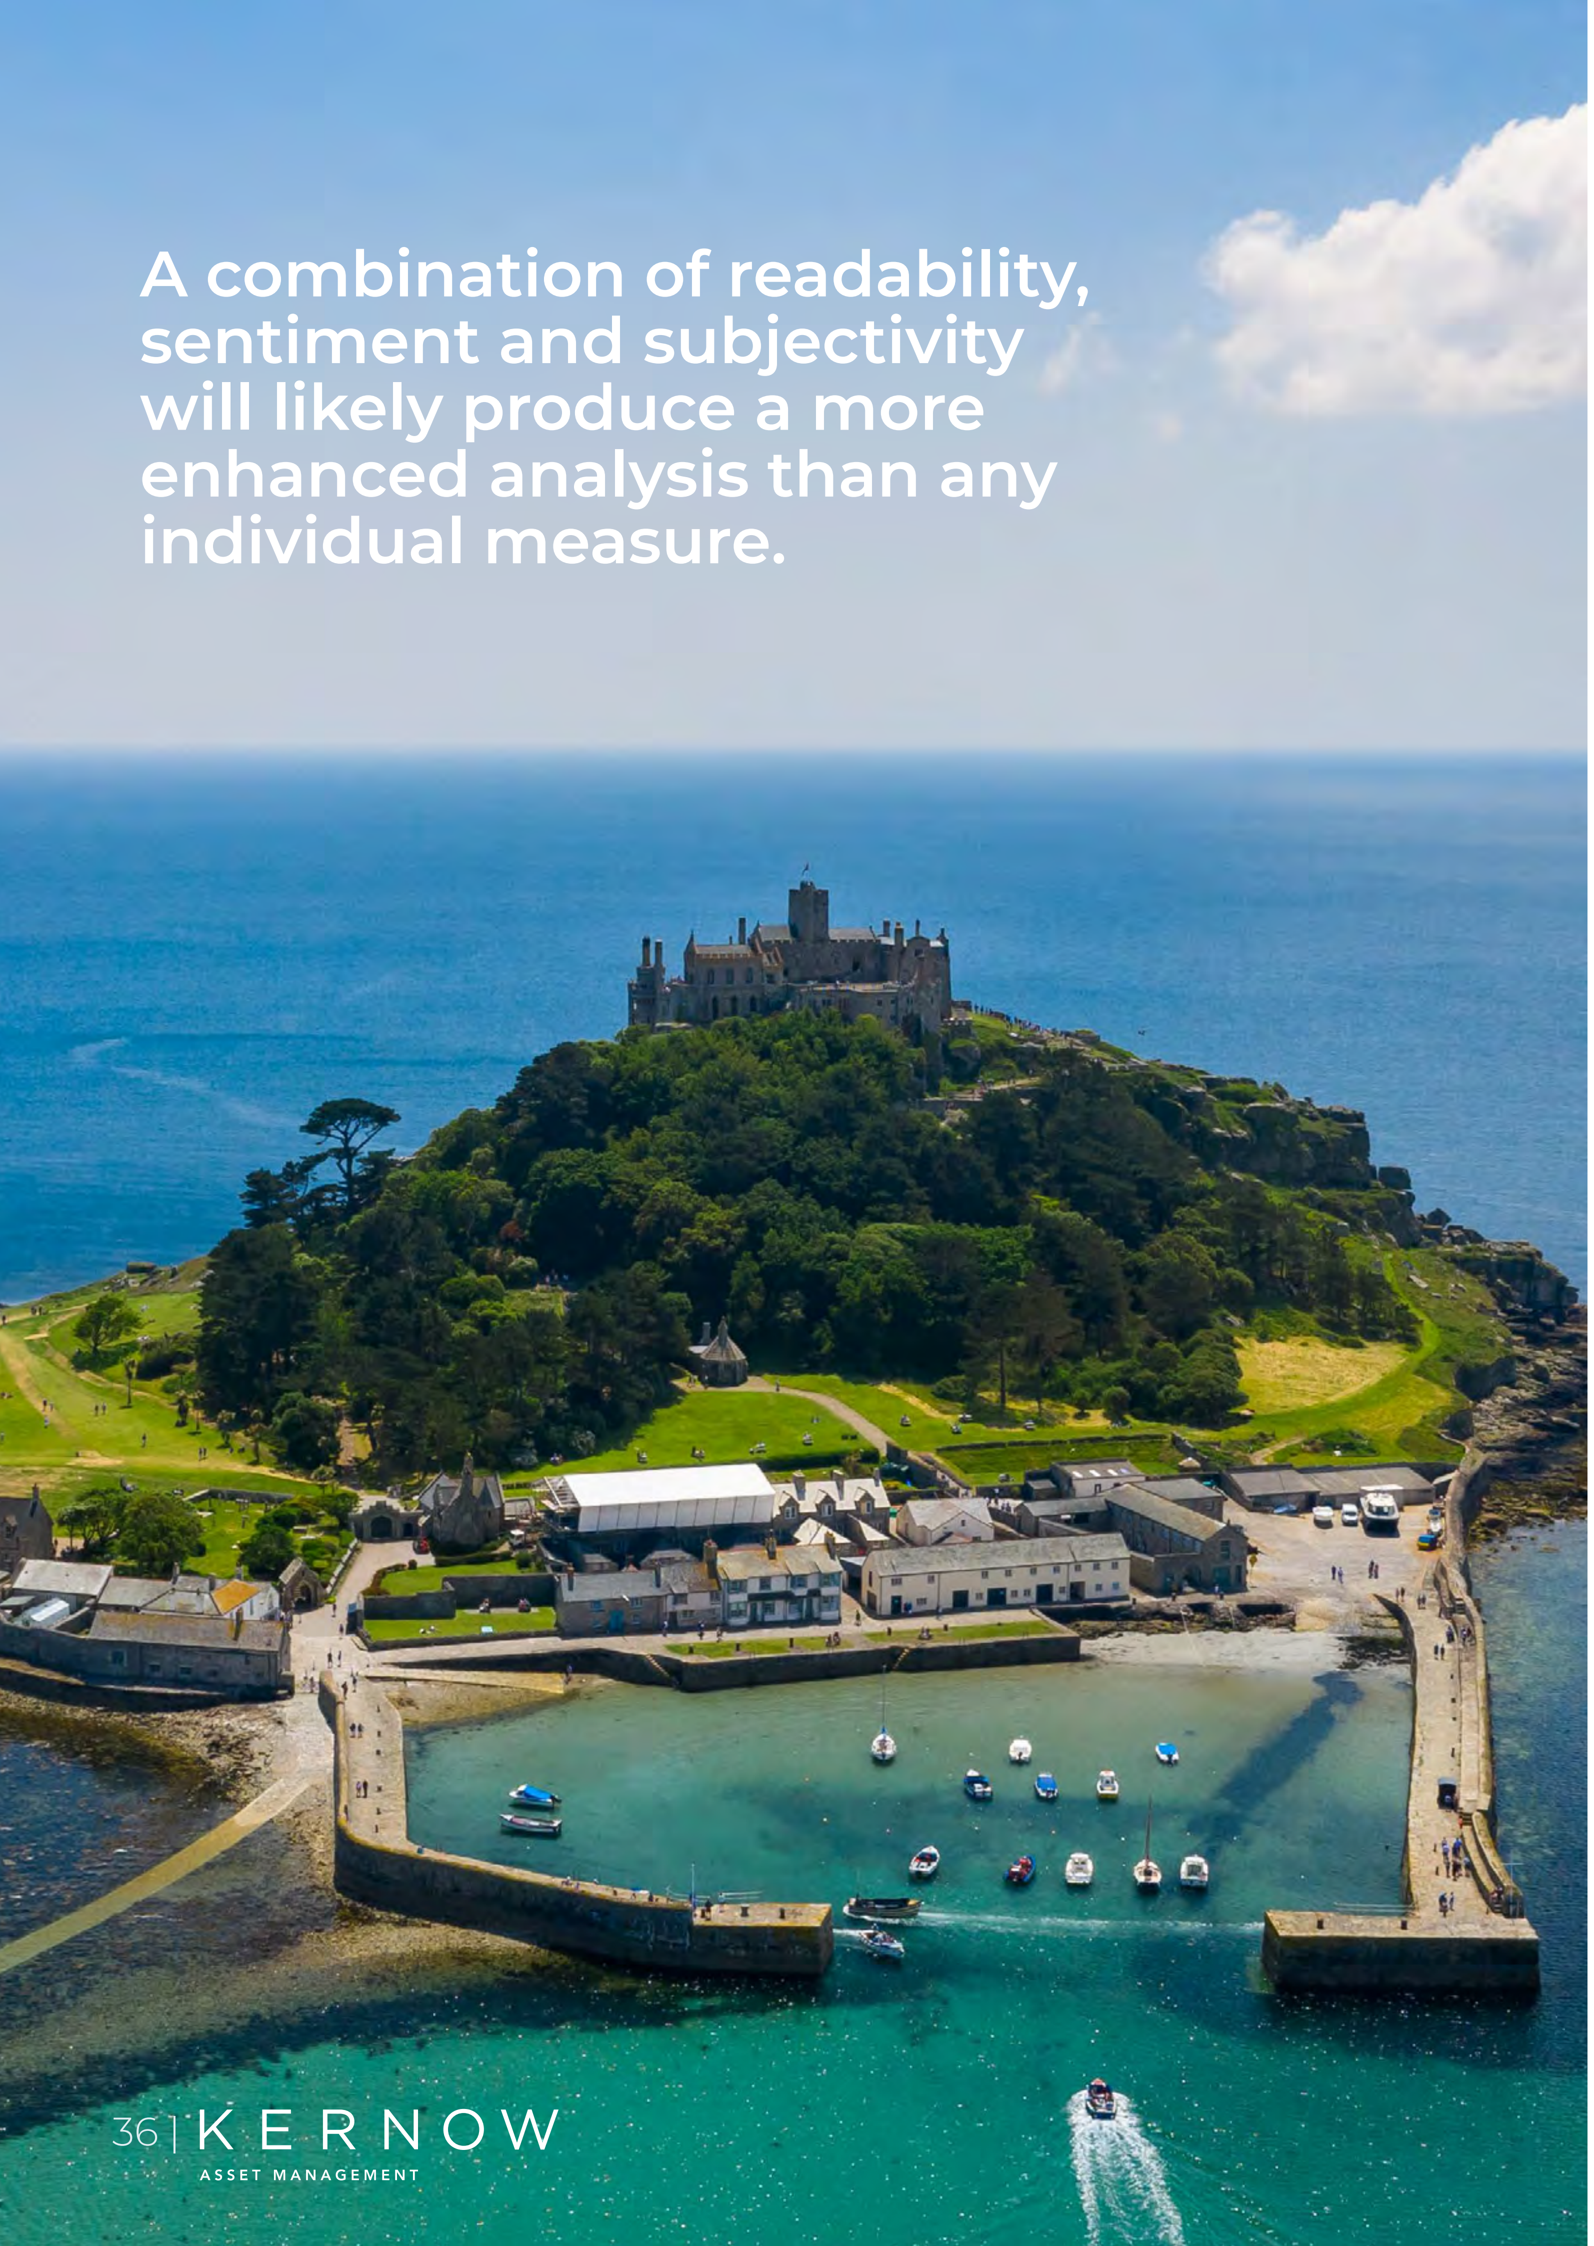

A combination of readability, sentiment and subjectivity will likely produce a more enhanced analysis than any individual measure.

subjectivity and readability. Put simply, articles with more positive sentiment also tend to have higher levels of subjectivity and are easier to read. This finding supports our obfuscation hypothesis and indicates that a combination of readability, sentiment and subjectivity will likely produce a more enhanced analysis than any individual measure.

	Readability	Sentiment
Sentiment	-0.06	
Subjectivity	-0.02	0.13

Figure 14: The correlation of sentiment, subjectivity, and readability across all RNS articles.

In analogy to the sentiment analysis in Figure 13, in Figure 15 we show the tertile spread performance across the RNS articles, partitioned by their category, ranked by their complexity. Interestingly, we see no significant relationship between subsequent company performance and the complexity of the language. We can see that for M&A RNS articles, the more complex the language, the more likely the company is to under-perform, but this effect is reversed in other categories, showing that this effect is fragile and unlikely to be statistically significant.

In conclusion, we find a relationship between company performance (on a relative basis) and the underlying sentiment of the article, we also find modest but notable relationships between sentiment, subjectivity, and readability in RNS articles. A combination of these features, along with the incorporation of metadata, such as RNS category and other company specifics, provide useful additional insight to an investment process.

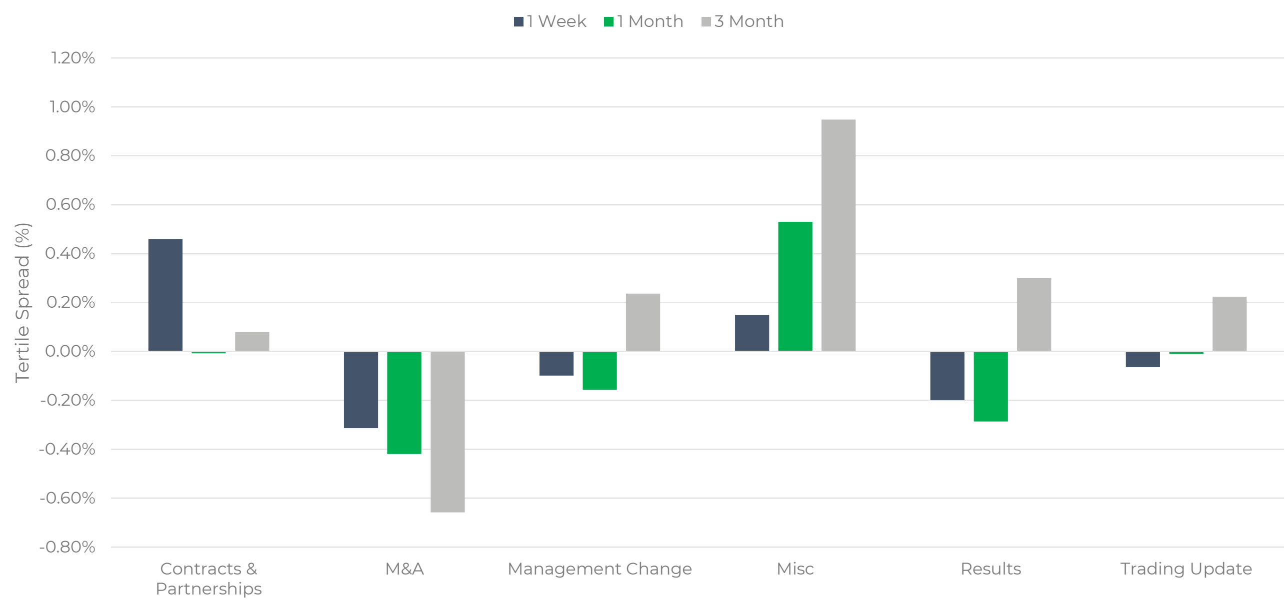

Figure 15: The average tertile-spread, calculated as the difference in median return for securities in the top vs bottom third of the investable universe, ranked on the RNS readability. Results shown partitioning by the RNS News Category, shown across three time-horizons.

Research HIGHLIGHT

In partnership with University of Exeter

Uncovering Hidden Connections

A relatively novel use of NLP techniques is to quantity the similarity between texts. Being able to calculate similarity between texts has numerous uses, allowing us to uncover connections between the companies submitting the texts, or to look for clusters of similar disclosures. We may also consider that period submissions by a particular company may naturally be relatively similar, so spotting outliers allows us to highlight novel information.

Although it is possible to measure word frequency and the co-mentions of particular words across text, these basic methods can create spurious results as they do not account for context, synonyms or other linguistic features. By contrast, new language models can use a method called embeddings to encode the semantic structure of a language. Document embeddings are numerical representations of text documents that capture their semantic content and context. They are used to enable efficient processing and analysis of text data for tasks such as document classification, information retrieval and similarity analysis.

The historical development of document embeddings has seen a progression from simpler methods to more sophisticated techniques for representing and understanding text documents. Early approaches to document representation primarily relied on basic techniques such as bag-of-words models, in which documents were represented as collections of individual words without considering order or meaning.

The field saw a significant advancement with the emergence of word embeddings, which assigned numerical vectors to words, capturing their semantic relationships. Introducing models like Word2Vec and GloVe enabled researchers to extend these ideas to documents. Doc2Vec, a variation of Word2Vec, allowed for creating document embeddings that considered both word and document-level information. More recently, the advent of large language models like GPT-3 and BERT revolutionised the field by providing pre-trained models capable of generating highly contextual and semantically rich document embeddings.

These models have been fine-tuned for various NLP tasks, offering state-of-the-art solutions including for document classification, summarisation and sentiment analysis.

For this analysis, we used the Doc2Vec embeddings model, which we have trained across the cleaned text data within all our RNS articles. In this way, our version of this model is calibrated to the type of documents that we are processing, rather than retrofitting a more generalist language model. This approach enables us to use relatively straightforward quantitative techniques to assess document similarities. For example, we use the cosine-similarity measure to generate a score between -1 and +1 that quantifies the level of similarity between documents. Calculating this similarity across all documents allows us to uncover specific patterns across companies, sectors and RNS categories over time.

Figure 16 shows the distributions of similarity scores for RNS articles partitioned by their sector. We also break these down by comparing companies within the same sector or across a different sector. Although there is a broad dispersion in similarity, indicated by the wide tails of the boxplots, we see that RNS announcements generally have a higher similarity when compared with stocks of the same sector. We also notice that certain sectors are more nuanced. For example, Healthcare RNS releases tend to be more self-similar than other sectors, forming a cluster of their own linguistic structure, whilst Real Estate RNS announcements tend to be distinct from the typical language used by companies in other sectors.

Although correlation between stocks is a critical metric for portfolio design, it has several known challenges. These include the assumption of linear relationships, the ignorance of latent factors, instability in the metrics and measurement errors. As such, it is important to also consider alternative ways of characterising relatedness between companies.

Using the RNS text similarity metric, we find that companies who are similar in their RNS disclosures tend to have higher correlations in their returns, but text similarity provides a stable complement to correlation measures. Figure 17 shows a network representation of the degree

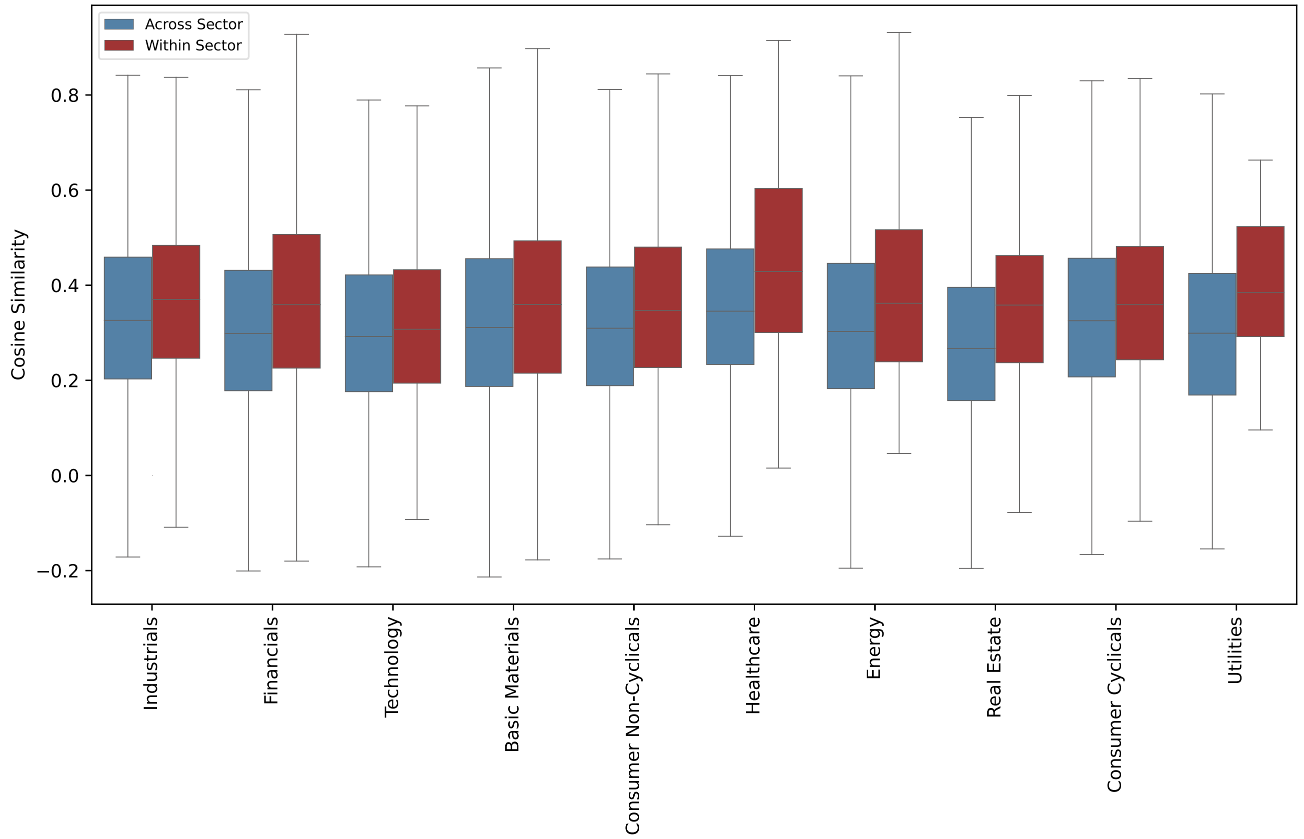

Figure 16: The boxplot showing the distribution of similarity between RNS submissions of stocks across different sectors, as compared to submissions from stocks within the same sector (red) or in different sectors (blue).

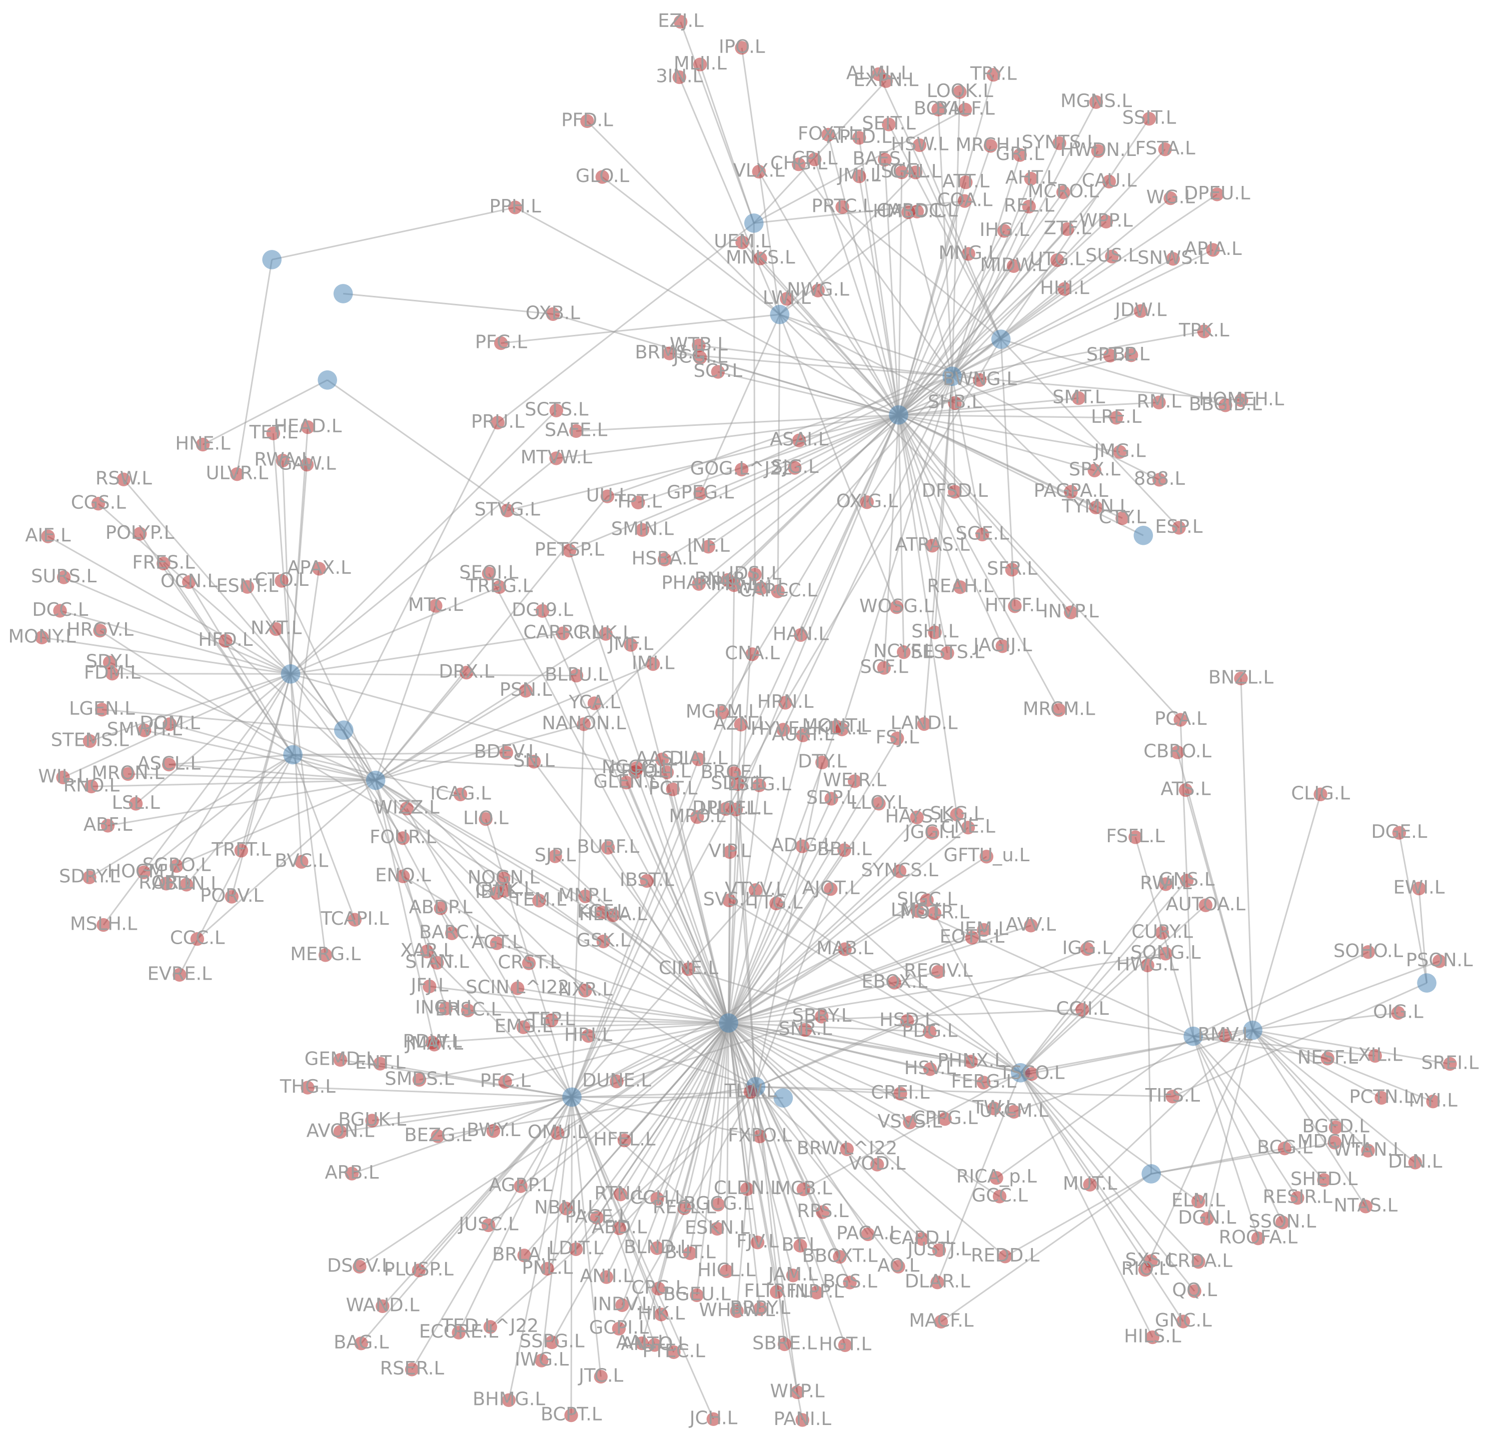

Figure 17: The minimum-spanning-tree representation highlighting similarities between Kernow holdings and related companies via their RNS text. Blue nodes correspond to Kernow holdings, and red nodes correspond to eligible securities which have a degree of similarity to them. Edges represent the distance (similarity) between companies, given their RNS text.

of similarity in a broader sense. The visualisation shows portfolio holdings in blue with similar stocks in red. The length of the edge linking stocks relates to the level of similarity, with higher text-similarity represented as closer distances between them.

This representation highlights the overall topology of the connectedness between securities. The positioning algorithm naturally results in securities with a higher overall similarity in the more central regions, whilst the more distinct securities on the periphery. We also see clusters created where certain themes emerge within RNS articles.

Conclusion

The RNS dataset provides UK investors with a unique information channel that is guaranteed to be of relevance to company operations.

There are several nuances to this dataset and specialist handling is required. Within the brief analysis presented in this article, we have shown the impact of RNS articles on a company stock prices and that, in general, measures of sentiment can be an indicator of future company returns. We also have shown that sentiment, subjectivity and readability are related – with good news typically being delivered in a more succinct manner than bad news.

Finally, by employing a deep-learning algorithm to create vector embeddings from RNS articles, we have computed text similarity scores and used these to search for novel news, related companies, and other more collective themes. Although this article has only scratched the surface, we hope it gives you a taste of the power of text-based analysis of RNS data. There is ample opportunity for further exploration, which continue to develop in support of Kernow's fundamental-based investment process.

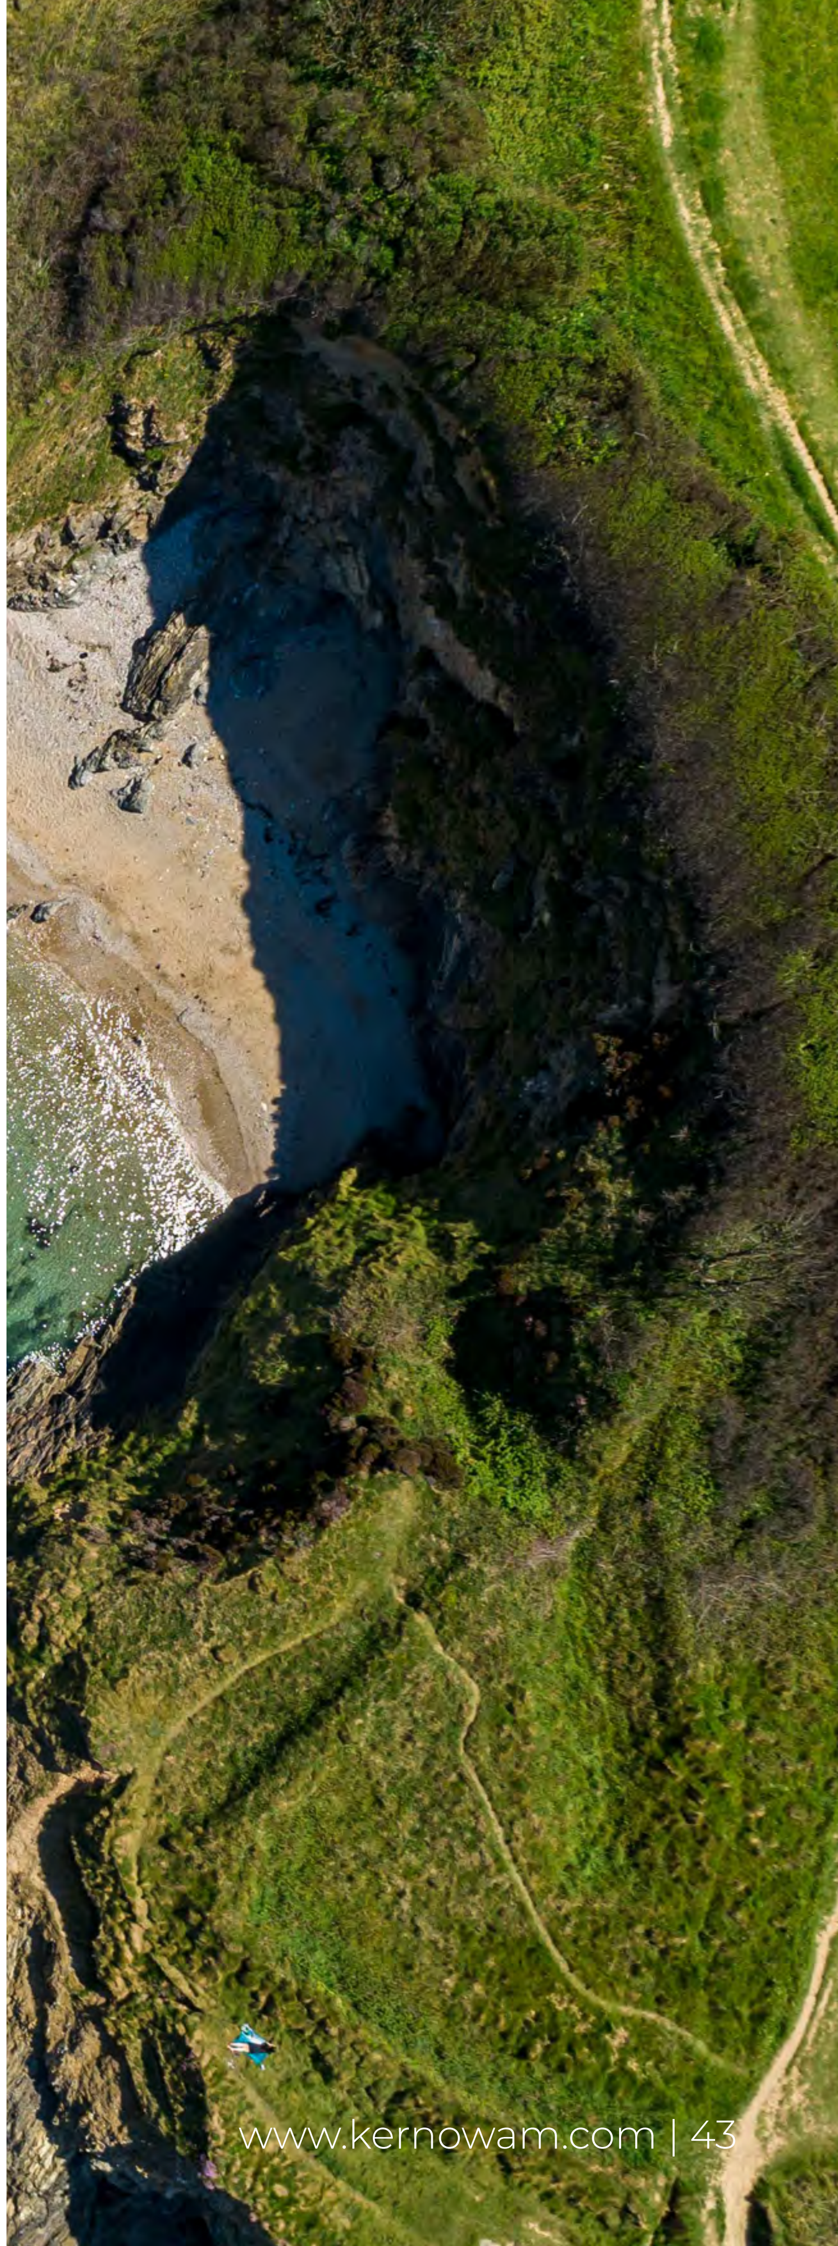

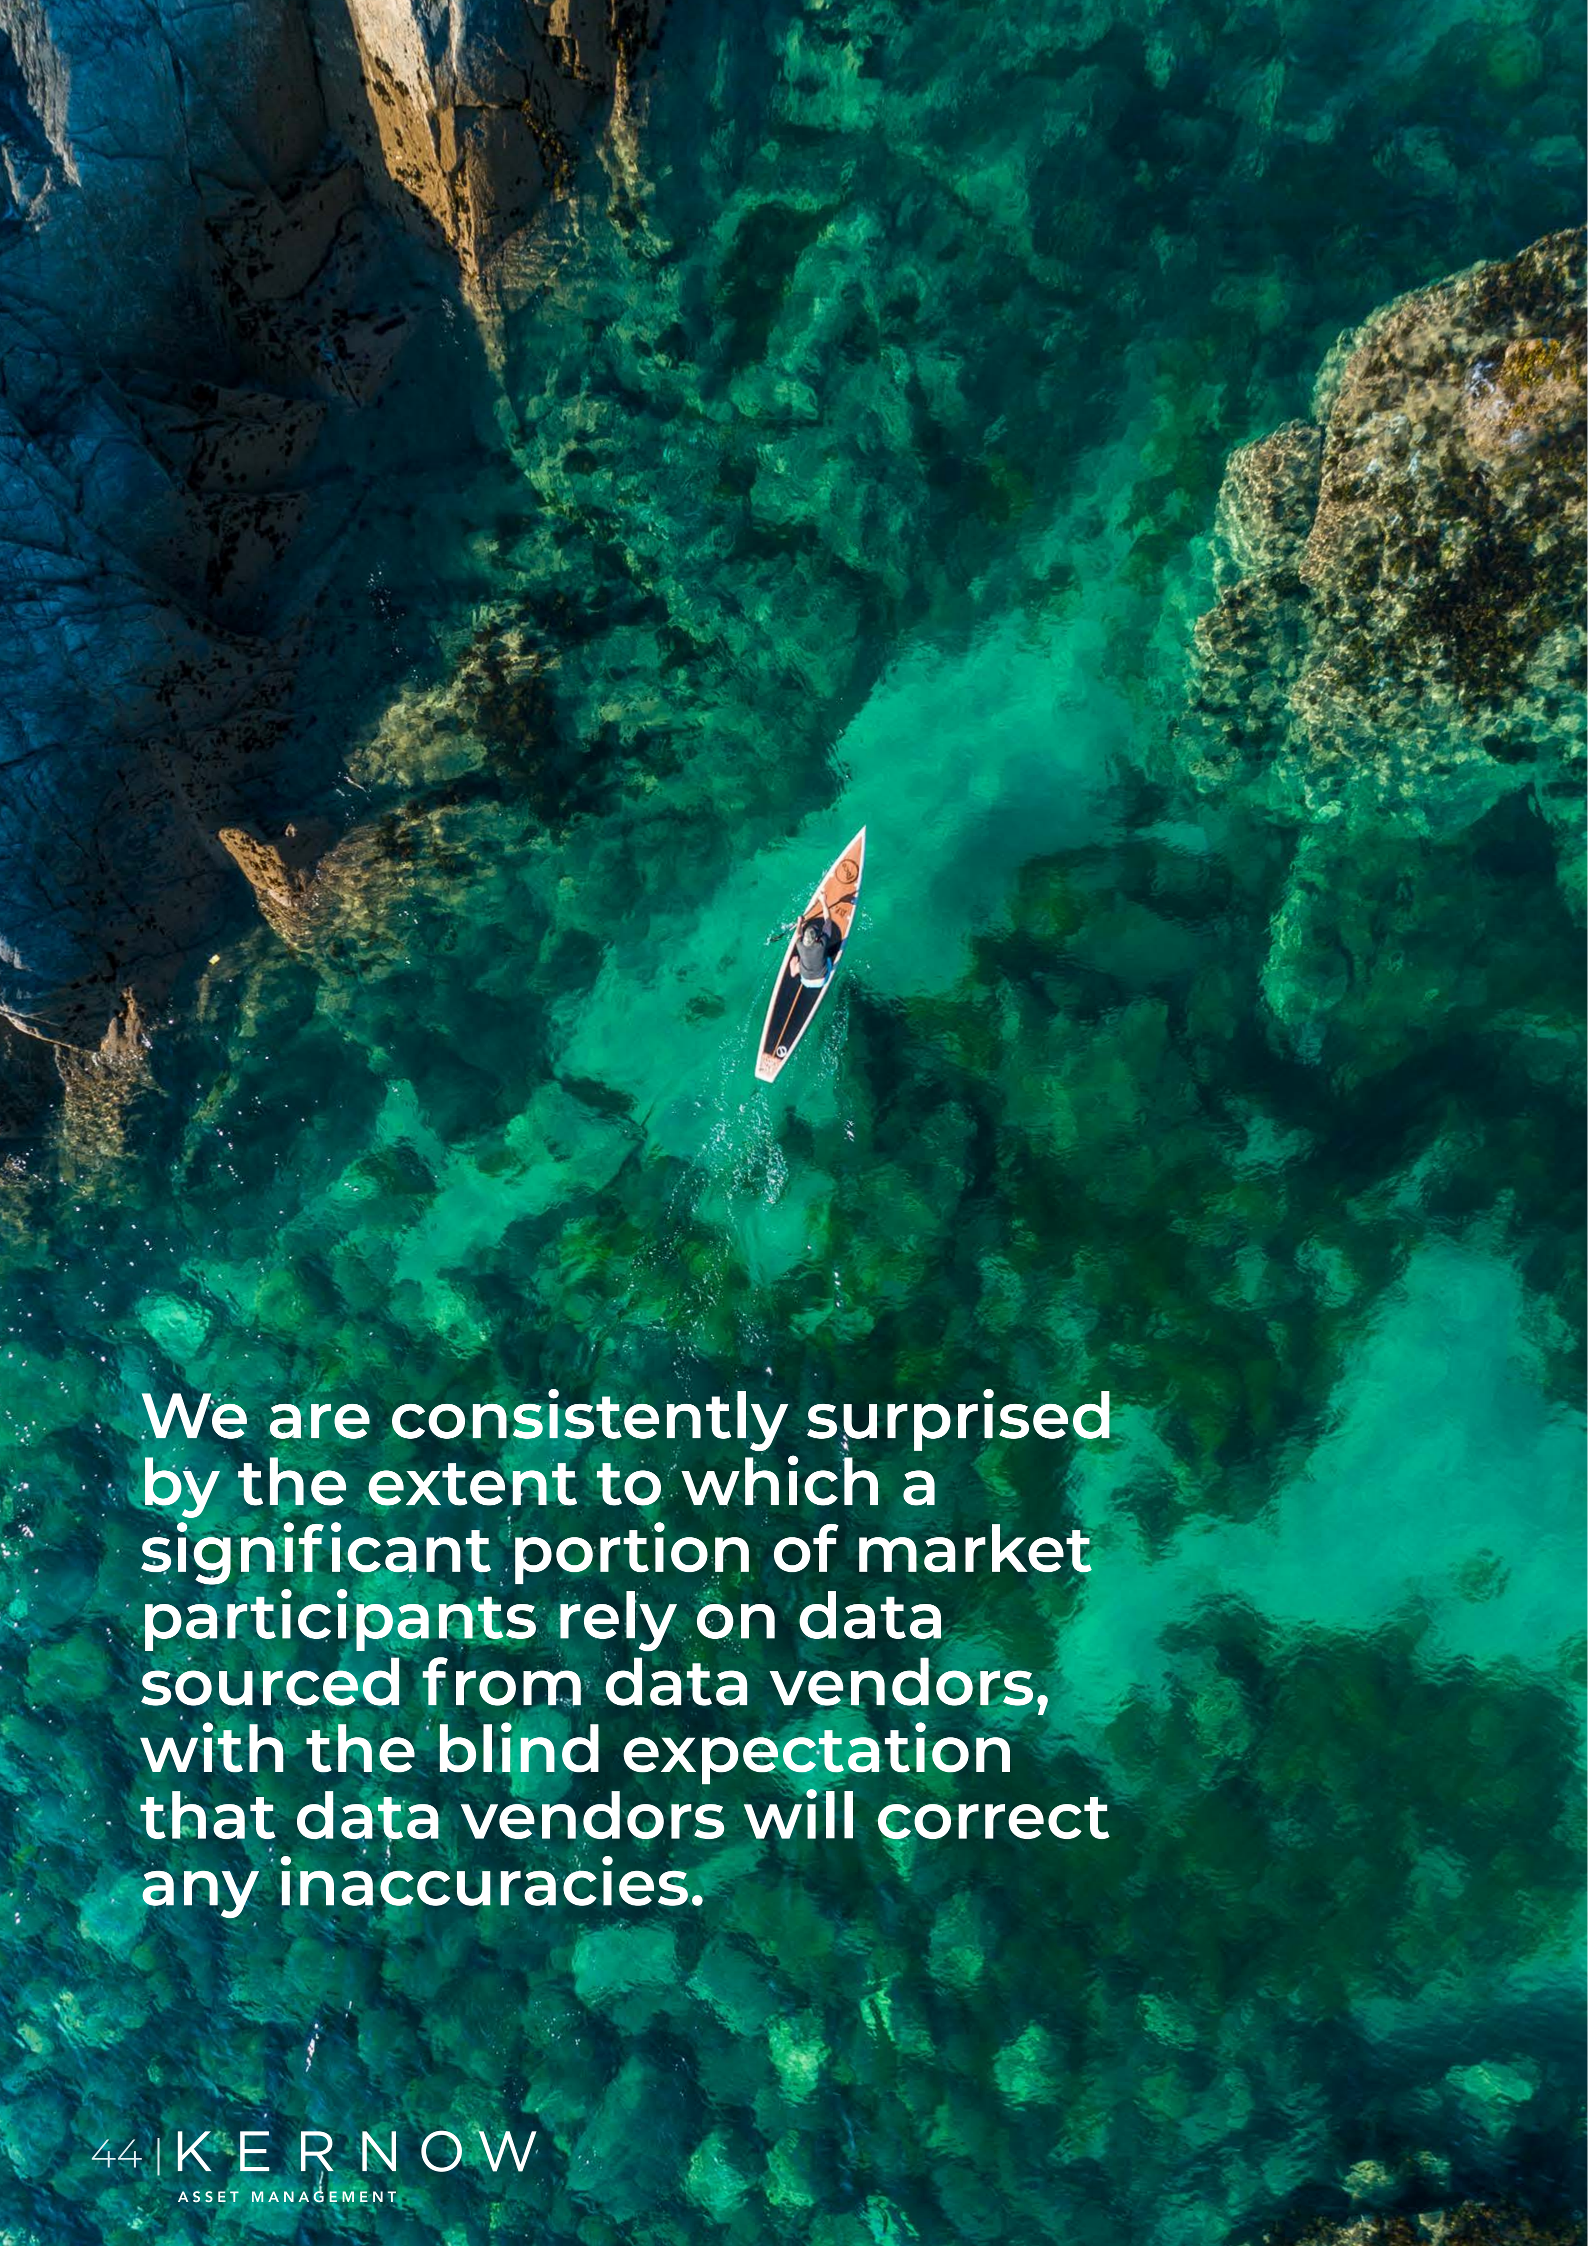

The Devil is in the Details

At Kernow, we place strong emphasis on the importance of sourcing data directly from primary sources, as opposed to relying on data provided by third-party vendors.

Data vendors have played a pivotal role in facilitating broader access to ever-expanding volumes of data over the past two decades for market participants. This has enabled investors to obtain a standardised, increasingly extensive perspective across a progressively vast array of metrics and a growing universe of securities. In part, data vendors have enabled many of the systematic approaches to investing.

However, there are numerous data integrity issues within such data-vendor datasets. These issues may come about either by propagating errors from primary sources or incorrect data parsing to standardise information intrinsically resistant to conformity. We are consistently surprised by the extent to which a significant portion of market participants rely on data sourced from data vendors, with the blind expectation that data vendors will correct any inaccuracies.

By contrast, we use primary data sources to construct financial models, delving into company annual reports and utilising data from the UK Companies House. We find this approach is fundamental to gaining in-depth insights into the inner workings of a particular company. We firmly believe that there are no quick shortcuts to this endeavour, and this meticulous process is imperative for thoroughly comprehending a company's operations.
Then of course there are oversights in primary data as well. Often, we tend to give companies the benefit of the doubt, attributing most errors to innocent mistakes, such as copy-paste errors, rounding inconsistencies, or other minor oversights. However, it is essential to recognise a connection between inaccuracies in financial accounts and the potential presence of fraudulent activities or negligence of duty within certain companies.

By way of illustration, we have highlighted five common errors, which can be identified - and appropriately addressed - by collecting and interrogating

Fundamental ANALYSIS

1 Impairments

It's crucial to emphasise that accounting errors are not limited solely to small companies with smaller in-house finance teams. For instance, consider Antofagasta, a company with a market capitalisation of £14 billion, which, as of October 2023, contained an erroneous sign in its cash flow reconciliation regarding an impairment provision producing an error of £161.6 million (i.e. -£80.8m should be +£80.8m).

A) Reconciliation of profit before tax to cash flow from continuing operations

	20221 ($m)	**2020 ($m)**
Profit before tax	**3,477.1**	1,413.1
Depreciation	**1,078.7**	1,048.7
Net loss on disposals	**9.2**	6.3
Net finance (income)/expense	**-16.0**	103.4
Net share of results from associates and joint ventures (exc. exceptional items)	**-59.7**	-5.1
Provision for impairment	**117.6**	-80.8
Decrease/(increase) in inventories	**10.9**	-13.6
Increase in debtors	**-206.8**	-259.9
Increase in creditors	**55.7**	31.0
(Decrease)/increase in provisions	**-19.0**	26.4
Cash flow generated from continuing operations	**4,507.7**	2,431.1

Figure 20: Note from Antofagasta's 2021 Annual Report. an error of £161.6 million - still available on its website!

2 Rounding Errors

There is a legal requirement to publish data in company results to a particular decimal precision. However, it is common to find rounding errors published in company results, which arise from truncating, rather than rounding the underlying information. This can cause error propagation at higher aggregations.

For example, the screenshot to the right from a recent annual report by Victoria Plc which contains such a rounding error. In the table, we demonstrate that adding the software cost items as-stated actually yields a total of £3.8m, which is different from the £3.6 erroneously published by Victoria Plc.

IT Software (£m)
3.1
0.9
0.0
-0.20
~~3.60~~

£3.8m

Figure 21: Victoria Plc 2023 annual report, highlighting the potential for cascading rounding errors in aggregated data.

3 Material historical restatements (corrections)

This is where a past set of accounts had errors that are corrected the following year in the next set of accounts. In the following example we show that CMC markets has a 'correction of error' line that comes at a cost of over £1m. If we spot them we point them out to companies, and sometimes they are corrected, and sometimes we are ignored. We also notice in many occasions, where we would not be able to decipher the 'non-balancing' errors in primary data sources and these restatements of particular components show red-flag activity.

9. Property, plant and equipment

£'000	Leasehold Improvements	Furniture, fixtures and equipment	Computer hardware	Right-of-use assets	Construction in progress	Total
At 31 March 2021 (As previously reported)						
Cost	19,273	9,656	36,249	19,146	-	84,324
Accumulated ammortisation	-14,393	-8,795	-27,235	-7,796	-	-58,219
Correction of Error	-	-	-	-1,134	-	-1,134
Carrying amount at 1st April 2021 (Restated)	4,880	861	9,014	10,216	-	-24,971

Figure 22: CMC Markets with the addition of a "Correction of Error" line to its PPE

4 Data Entry Errors

In other instances, data is pasted in from other sources, presumably internal spreadsheets from company finance departments. In the below example from Amadeo Air Four's, there is a missing "1" in the bottom line. Adding the figures gives a total of £113,384,108, which is visually similar to the number published albeit missing the leading number. Unfortunately, this makes the headline number approximately 99% mismatched!

Deprecation of the current year based on previous year residual values	111,930,032
Amortisation of acquisition costs on aircraft	756,519
Adjustment due to change in useful life *	-1,246,006
Adjustment due to change of residual value	1,943,563
Net depreciation charge on all aircraft for the year	~~13,384,108~~

£113,384,108

Figure 23: Amadeo Air Four Plus, 2023 annual financial statement, showing a somewhat dramatic data entry issue.

5 Incorrect categorisation

Construction business Galliford Try has an issue with mis-categorisation at several data providers. Although its housebuilding business was sold in January 2020, many vendors still categorise it as a housebuilder; and hence many people also believe it still is.

The successful execution of its business plan in the infrastructure sector is, therefore, being ignored by many and has not yet been fully rewarded by the market. The company is also listed in various places as having debt, when in fact, it is debt-free, and the "debt" is merely leasing liabilities. This means it misses out on many screens by investors looking for debt-free companies to invest in.

However, the majority of third-party data vendors have incorrectly labelled this as conventional debt and as such many investment quality metrics for investors are likely to misrepresent the company's leverage.

NOT conventional debt

Lease Liabilities

	2023 (£m)	2022 (£m)
Current	**14.9**	9.9
Non-current	**24.2**	14.9
Total Lease Liabilities	**39.1**	24.8

The consolidated income statement shows the following amounts relating to leases for continuing operations.

Figure 22: The Galliford Try income statement, highlighting their lease liabilities, mislabelled as conventional debt by third-party data vendors.

Conclusion

Gaining a competitive advantage as an investor can stem from various sources. Adhering to the fundamentals with meticulous attention to detail allows us to foster a stronger sense of confidence in our investment decisions. We view the use of primary source data as an essential element of our investment approach. Ignoring them could be a perilous oversight.

To illustrate this, in 2020, we observed well over 100 restatements among UK-listed companies, representing roughly 7% of our investable universe. Having a way to spot and track this drift from what a company says with numbers, that then change when no one is looking a year later, is clearly important in building a trading edge. These examples underscore the importance of accessing primary data and scrutinising data points incorporated into a model.

Visual
ELEGANCE

Human activity indelibly impacts the natural world we live in. As a society, we are showing more appreciation of this and increasing our understanding of how our activities can lead to alarming consequences.

Impacts have intensified as the global population rapidly advances and civilisation grows in complexity.

There is a broad range of activities that significantly impact our environment, including intensive mechanised farming and rapid urbanisation, including construction of buildings, utilities and transport networks.

It is only relatively recently (on civilisation timescales) that society has realised how vital our impact is on the global environment.

The below graphic shows the extent of humanity's impact on the world from 1993-2009,using information collated in the paper 'Global Terrestrial Human Footprint Maps for 1993 and 2009'. Within this study, the authors compiled a set of eight variables that specify the degree of human pressure: built environments, population density, night-time lights, croplands, pasture, roads, railways, and navigable waterways. These effects were then normalised and weighted in a way which balanced the impact of each. These were plotted in the map below.
This visualisation is particularly elegant, allowing the viewer to apply their prior knowledge of global geography to the data represented. Geospatial data is often overlooked as it requires relatively specialist handling, but clearly, a wealth of information can be obtained. For example, overlaying this visualisation with population statistics (density, affluence, etc) can uncover unprecedented insights.

Back to the visualisation, we see that the overall environmental footprint occurs around dense urban centres. Though there were still corners of the Earth with little to no human impact in 2009, changes in demographics, politics, and consumption could have an outsized effect on humanity's current and future footprint.

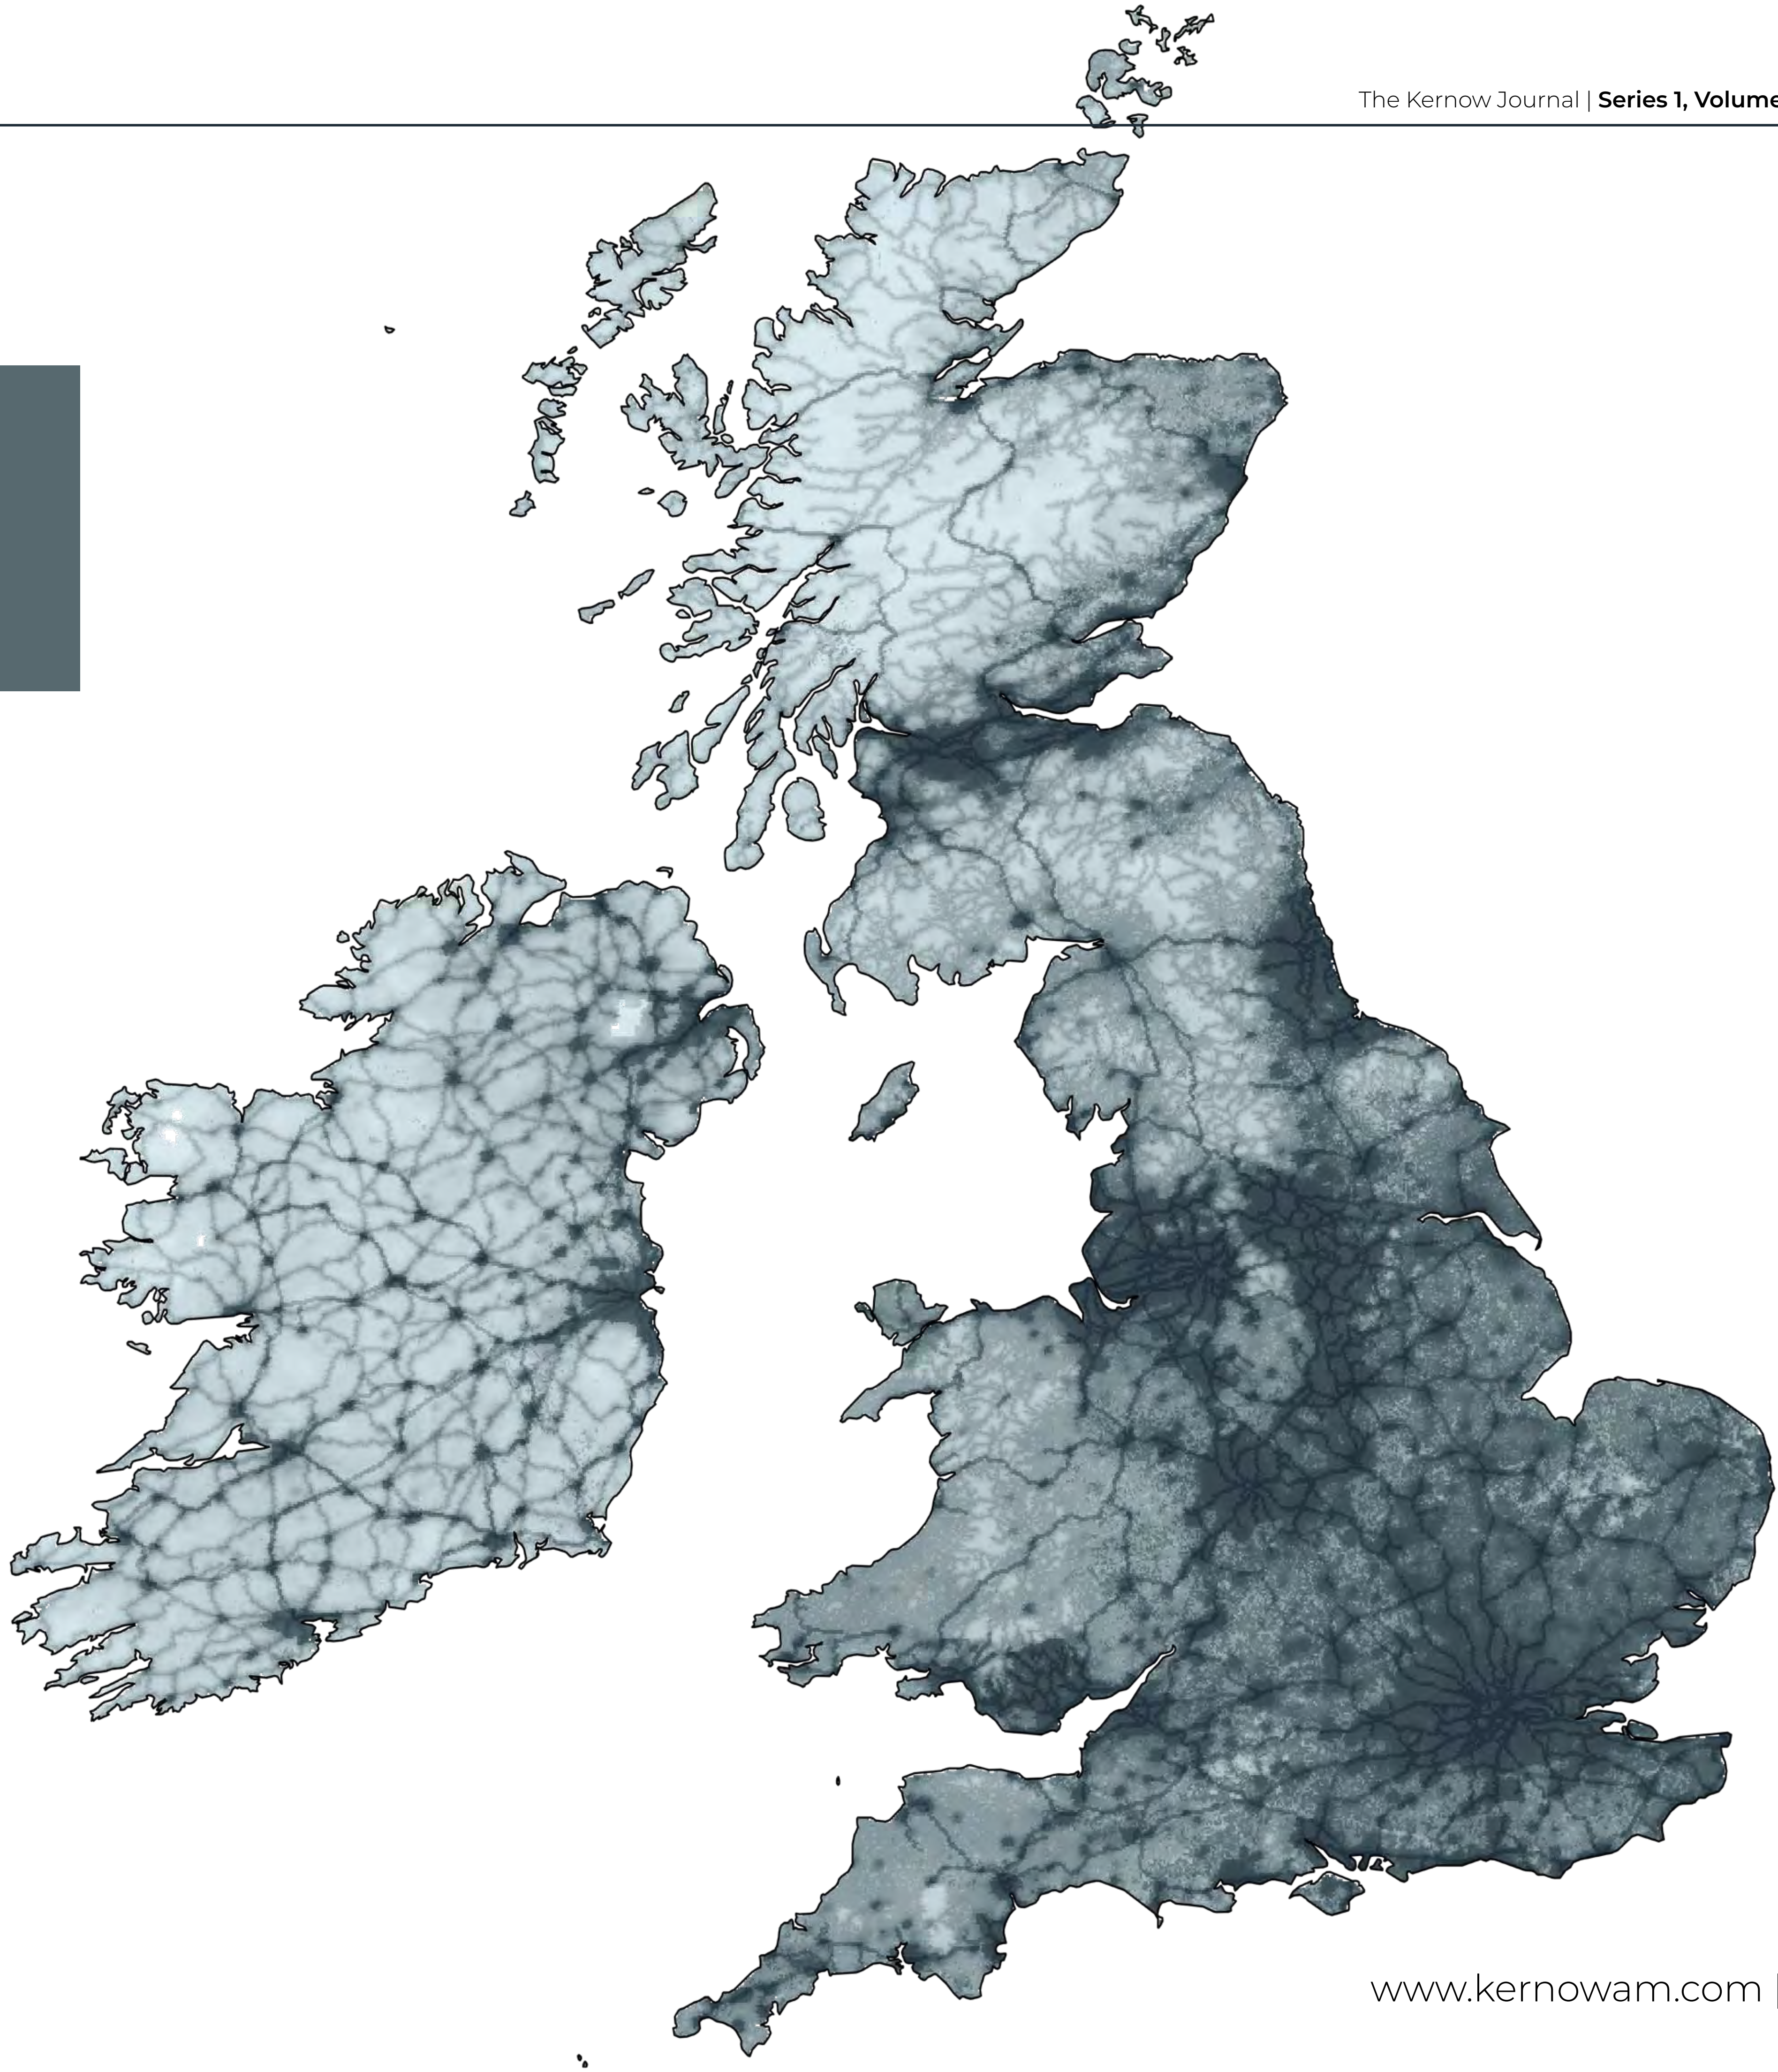

PythonMaps by Adam Symington

Meet
OUR TEAM

Alyx Wood CHIEF INVESTMENT OFFICER
alyx@kernowam.com

Alyx is responsible for trading and investment research.

He has been investing since the age of fifteen and considers himself fortunate to be able to describe this fascinating activity as his job. Kernow is the realisation of a lifelong dream to work alongside talented people in applying their aggregated investment knowledge in an uncompromised form, putting investors at the forefront of decision making.

Previously Fund Manager at Downing LLP, Vice President at Deutsche Bank AG and Management Consultant at KPMG LLP. Qualified Chartered Accountant (ICAS) and holds the Investment Management Certificate and a BA in Accounting and Economics.

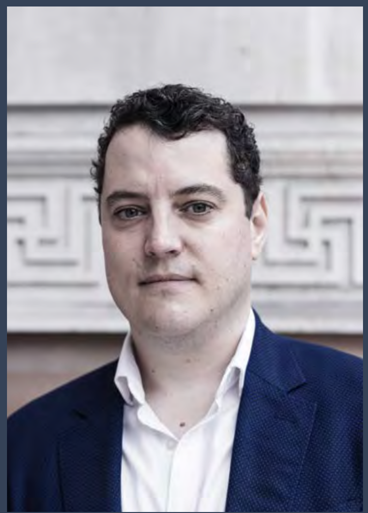

Edward Hugo CHIEF EXECUTIVE OFFICER
edward@kernowam.com

Edward is responsible for managing the resources of the company and executing Kernow's strategy. He began his career analysing technology and cleantech start-up businesses, before starting his own company in the consulting sector. Covering the alternative energy, food and agriculture sectors as an equity analyst provided him with an excellent insight into the global financial markets. He was attracted to the opportunity of returning to his entrepreneurial roots, backing his own ideas and working with a great team.

Previously Head of Equity Research at boutique investment bank VSA Capital. Holds an MSc in Microsystems & Nanotechnology and a BSc in Mathematics & Artificial Intelligence.

Dr Michael Cook HEAD OF DATA AND ANALYTICS
michael@kernowam.com

Having spent over a decade in investment teams spanning both systematic and fundamental investing, Mike has a wealth of experience in quantitative analysis, data science and investment operations. Before joining Kernow, Michael worked within the Advantage team at Lazard Asset Management, focusing on equity research across all major global regions. Prior to this, he was a quantitative analyst at Man AHL, focusing principally on incorporating alternative data within systematic investment processes. Michael began his career in finance at Lloyds Banking Group, independently valuing exotic cross-asset derivatives.

He holds a PhD in Astrophysics from SISSA (Italy) and an MSci in Physics with Astronomy (first class honours) from the University of Nottingham.

CREDITS

We would like to extend our gratitude and congratulations to the University of Exeter students for successfully completing their Data Science Masters course. Our collaboration on RNS analysis and exploration of carbon emissions demographics among UK listed companies has provided numerous valuable insights. Special appreciation goes to Professor Stuart Townley, Professor Markus Mueller, and the Applied Mathematics department for their warm hospitality and engaging discussions. We are also extremely grateful for the beautiful images captured by Aerial Cornwall (https://aerialcornwall. com/), expert graphic design by Grace Kennard, and skilful branding by Holly Carberry (https://www. hollycarberry.com/). Finally, we express our gratitude to Adam Symington for his valuable contributions to our Visual Elegance segment. His expertise in geospatial data analysis and creation of captivating visualisations have established new standards in his field. We look forward to the possibility of future collaborations with enthusiasm.

DISCLAIMERS

KERNOW
ASSET MANAGEMENT

www.kernowam.com

contact@kernowam.com

CORNWALL OFFICE
1st Floor Great Stable,
Trelowarren,
Helston,
Cornwall, TR12 6AF