The Kernow Journal

The Alternative UK Equity Market View

Kernow Contemplations

Research Spotlight: What Is The Right Level Of Concentration?

Company Spotlight: Saga plc

Chart Of The Month

KERNOW
ASSET MANAGEMENT

Series 1, Volume 1 | June 2022

CONTENTS

A note from
THE EDITORS

Welcome to the first edition of the Kernow Journal!

This is our attempt to share some of the insights that we come across in our stewardship of UK listed assets. We find our work fascinating and are conscious that the vast majority of it does not see the light of day beyond our beloved office.

When reading through the various research reports, practitioner white-papers, financial news and editorial articles that invariably form part of our investor psyche, I often think of the following quote.

"Time changes everything, except something within us, which is always surprised by change"
Thomas Hardy

During the day-to-day dialogue at Kernow, our discussions vary from straightforward evaluations to the outright metaphysical. This is one of the many reasons that makes being a professional investor so interesting. Whilst it is impossible (and perhaps inadvisable) to elaborate on all our conversations, we thought that passing on some of these would potentially spark some interesting conversations amongst our clients and contemporaries.

This journal therefore serves two purposes: we hope that it provides a novel insight 'under the covers' as we navigate the economic landscape, and provides some food-for-thought on selected topics that have interested us in recent weeks.

Whilst we anticipate the structure of the journal to be relatively free flowing, in this instalment we highlight some general research themes that have piqued our interest. We also introduce The Kernow Index; this forms a core component of our research process and so we anticipate revisiting the properties (and evolution) of this framework in future volumes of this journal, along with some examples of how we use it to inform our investment process.

Much of our research is somewhat specialist and therefore perhaps of only specialist appeal, however, our aim is to summarise a selected project that we have developed recently that is perhaps of more general interest. In this instalment we have focused on the topic of portfolio concentration, addressing both the measurement and management of concentration using some basic examples, and then showing how these theoretical ideas can be applied. We find that achieving a 'sweet spot' of concentration, measuring

concentration appropriately, and allowing portfolios to organically evolve into more concentrated positions can each be beneficial. We also take a quick look into a single company, highlighting some of the more esoteric lines of enquiry we have recently taken when conducting our evaluations.

Finally, to cite an old adage, a picture is worth a thousand words. On occasions during the course of our working day, we come across particular charts or visualisations that we find particularly compelling.

We close this edition by paying homage to these works-of-art by showcasing and briefly discussing our chart of the month. We hope that you find the varied contents of this journal as interesting and illuminating as we do.

Happy reading!

Kernow Asset Management

Kernow
CONTEMPLATIONS

Over recent weeks it has been almost impossible to avoid discussions on the current inflationary environment, rises in interest rates and looming (if not already present) cost-of-living elevations. The conflict in Ukraine and political rhetoric seems to be pivoting away from discussing the merits of globalisation, towards more focus on operational security – whether this be in moving towards energy or food independence.

Clearly the ramifications are somewhat counter to the desire to create an economic environment geared towards growth. We feel that this has some relatively profound consequences on medium and long term horizons, although admittedly, it's unclear how those may play out.

Quite a few of our team conversations have been loosely based on this theme; recently our CEO, Edward Hugo, coined the term 'economic entropy' in an attempt to describe the way in which the globalisation of supply chains appears to be the natural order and

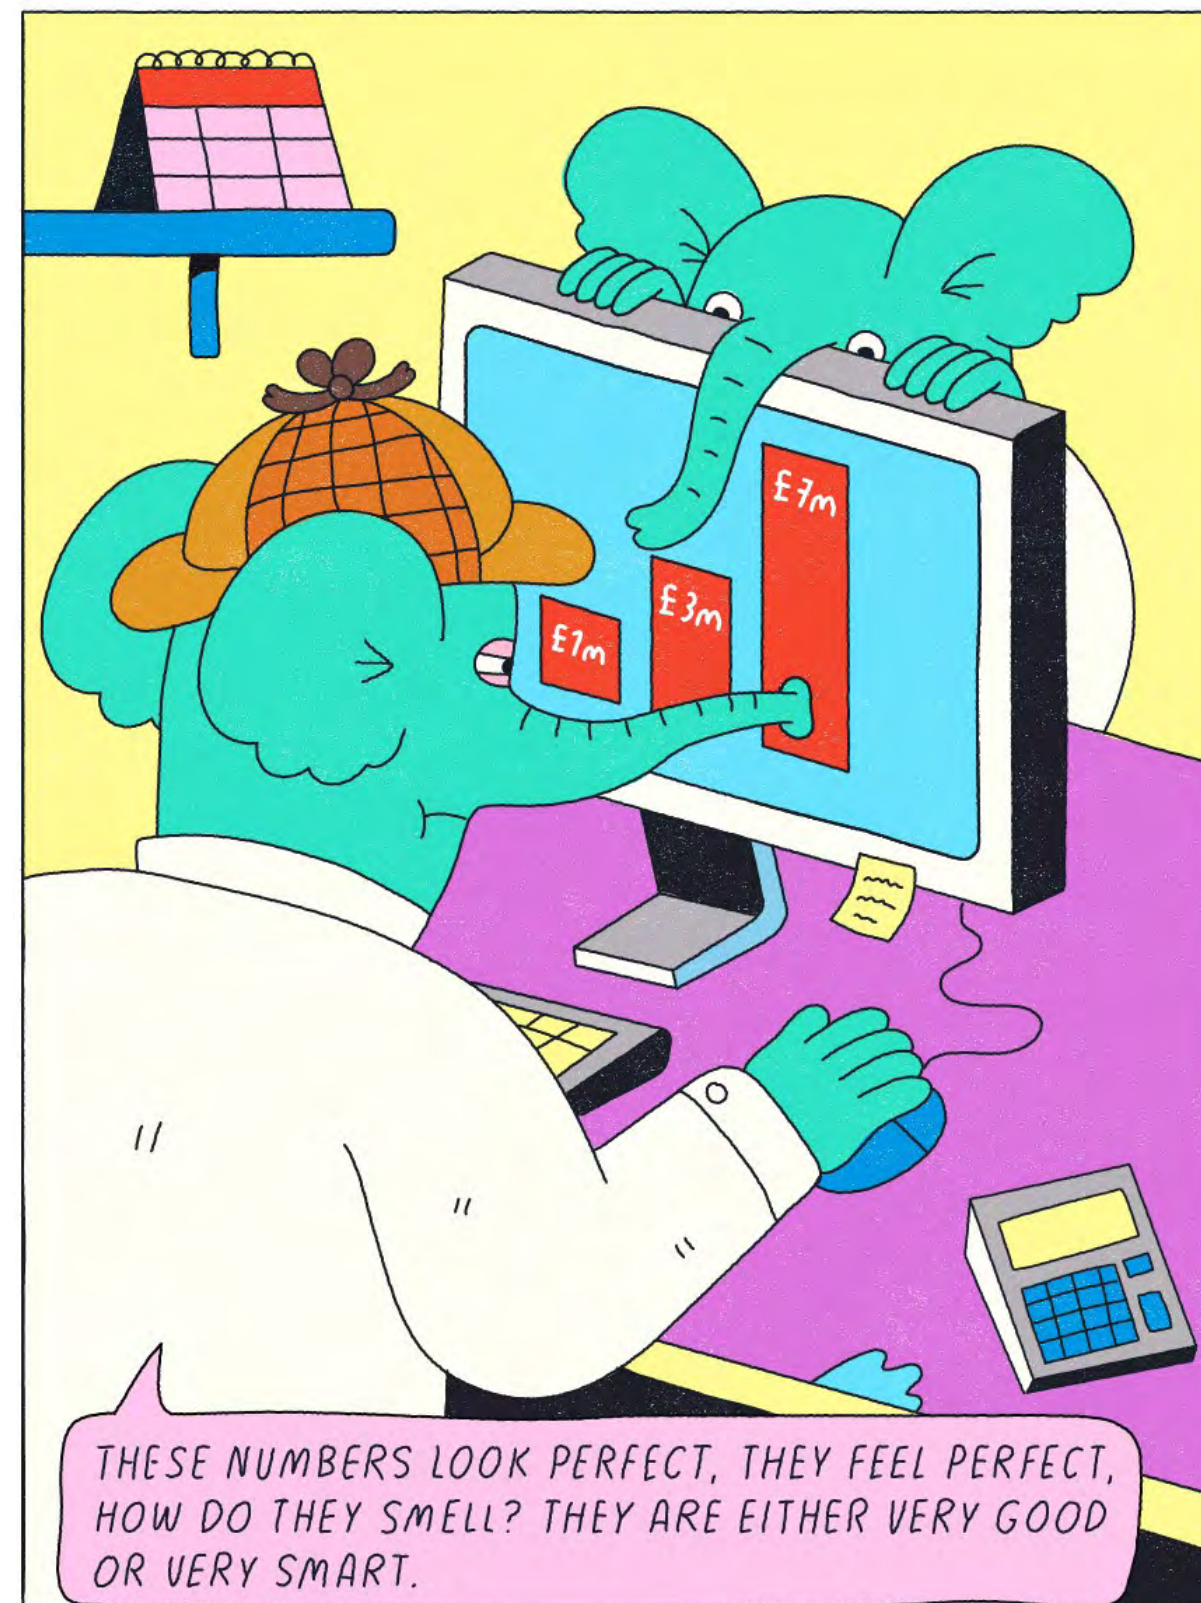any attempt to reverse this through neo-mercantilism will be extremely difficult.

The concept of economic entropy, it turns out, is discussed in a handful of academic articles spanning back to the mid-1970's, with only scant attention from the research community. However, recent developments in the fields of econophysics have yielded several highly practical outcomes, some of which we use within our risk analysis.

It is fascinating to see how the cross-pollination of concepts, ideas and techniques can be used

effectively in different contexts. In general, ideas that draw from multiple disciplines have always appealed to us; we would generally consider them to have a higher likelihood of success, or at least, a good chance of finding something truly novel.

Tangentially, we have also recently been discussing the detection of anomalies, or manipulated numbers, particularly within accounting figures. This led us to the simple but often overlooked mathematical metric - Benford's law.
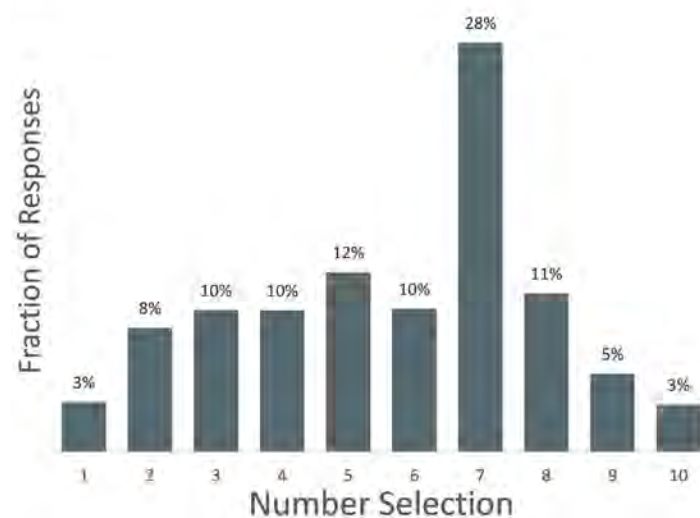
This law outlines the expected distribution of

THESE NUMBERS LOOK PERFECT, THEY FEEL PERFECT, HOW DO THEY SMELL? THEY ARE EITHER VERY GOOD OR VERY SMART.

Kernow
CONTEMPLATIONS

Fig 1: The distribution of responses from a survey asking 1,000 individuals to select a random number between 1-10.

numbers to their 'expected' distribution, we can relatively straightforwardly scan for potentially manipulated data.

When people are asked to make up a list of random numbers, they often include a disproportionately large number of sevens (see fig.1). This arises from an emotional bias - the number feels more random than it actually is.
This basic comparison has been used in several applications to fraud detection, including within criminal trials, election data, and macroeconomic data. Whilst there are limitations to this approach, it has provided us with a simple way to quickly screen numerous datasets. This

approach is one of the many tests that can be used for detecting accounting fraud.

Finally, in recent weeks we have discussed a number of topics relating to the Regulatory News Service (RNS) disclosures by UK listed companies. When analysing textual data, often a lot of attention is placed on computing various sentiment measures, this makes sense, as defining a positive or negative tone has been shown (at least, historically) to be indicative of future company performance.

However, it is relatively straightforward to reverse-engineer disclosures such that they tick various sentiment boxes. A slightly

different angle, that we have found interesting, is to instead calculate measures of readability of the text. Anecdotally, we have found that RNS announcements that are presenting underwhelming results tend to dance around a little, and end up being more convoluted than they otherwise would be.

Perhaps it is a natural human instinct to sugar-coat negative news, and this manifests as long-winded, obfuscated text. Fortunately, there are a number of interesting, and relatively simple metrics that have been developed over the decades that attempt to quantify text readability. Perhaps the first of these, and still one of the more

widely-used, is the Gunning-Fog index.

When we apply this metric to the various RNS disclosures, adjusting our metrics for different types of announcements or different businesses. It has been fascinating to see that in general, our 'sugar coating' hypothesis appears to be accurate.

Whilst promising, this has caused us to further question corporate disclosures. Do the companies now know the machines are watching them? Now that we have computers checking that humans are not robots, it is unclear who is watching who!

significant digits in a set of numbers. Specifically, it shows that the first set of numbers extracted from a diverse set of records does not have uniform distribution: their incidence diminishes (exponentially)

digits become larger. Thus, the occurrence of '1' happens approximately 30% of the time, whilst '9' occurs only 5% of the time. By comparing the distribution of significant digits of sets of

The Kernow
INDEX

We get lot of questions on how we track valuations and the inefficiency spread within UK equities - so we thought we would start by sharing a summary visual.

Fig. 2: The Kernow Index showing the UK equity market over and undervaluations for each company.
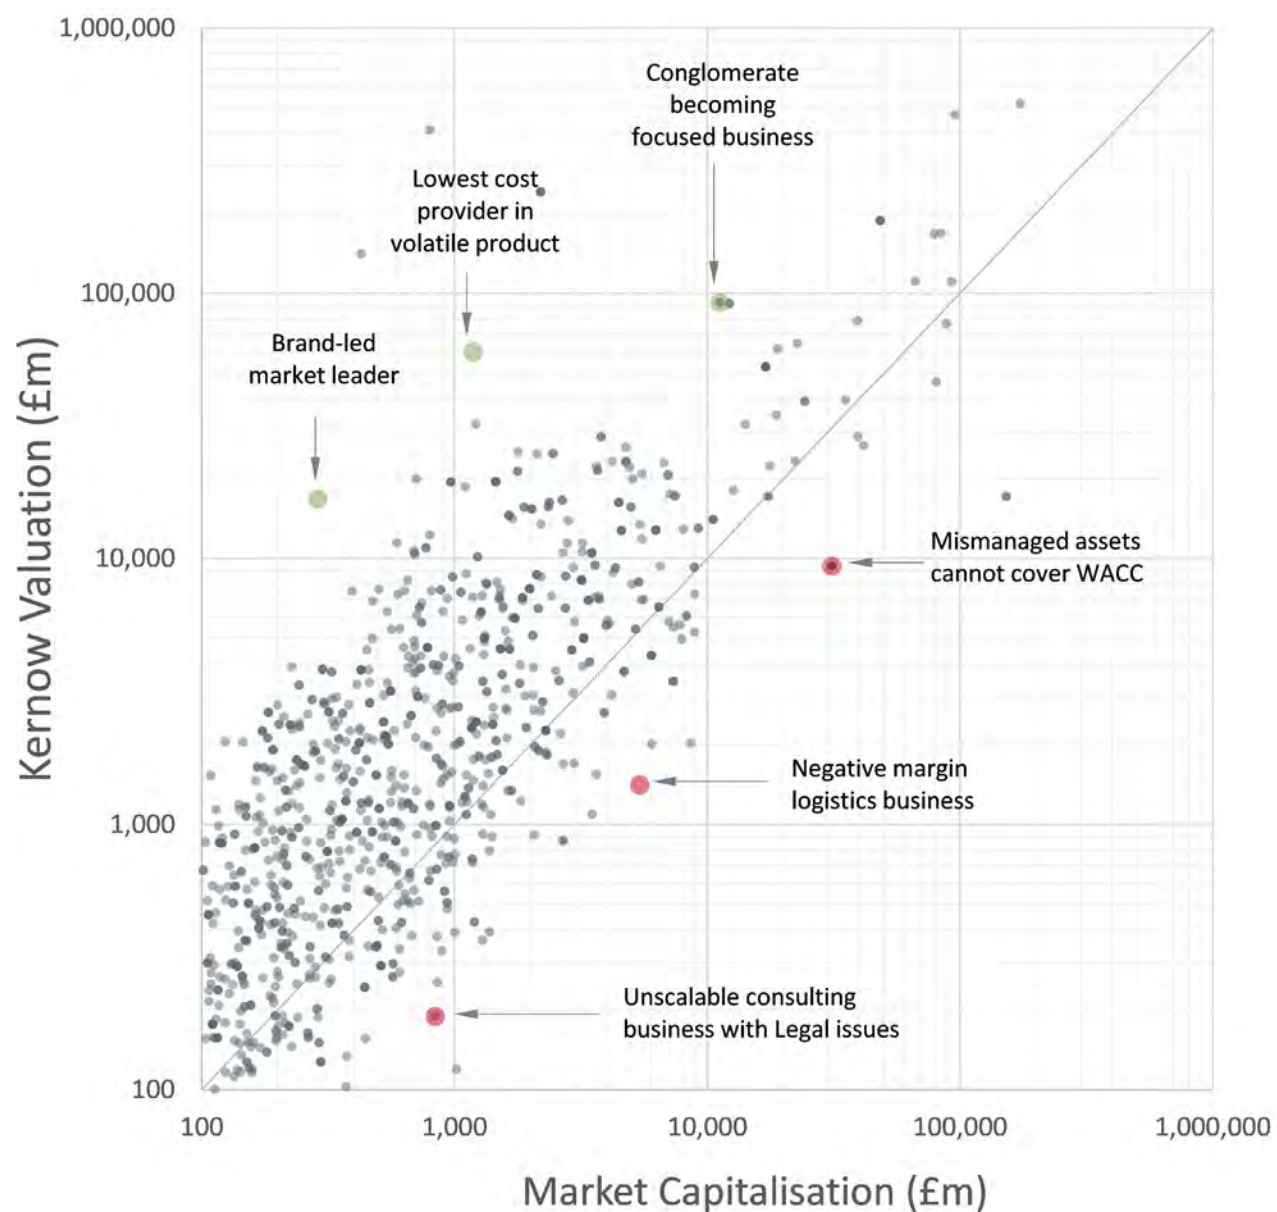
This chart (fig. 2) shows the current levels that make up the Kernow index. Each point on the scatter corresponds to a company, and the placement indicates whether companies are under or over-valued. The companies are valued using a proprietary fundamental bottom-up process.

The points highlighted are some companies currently being analysed. By observing the clustering of points and their dynamics, we can assess where opportunities are likely to arise within our investment universe.

At Kernow we are unapologetically contrarian in our approach to investing: we find companies that are intrinsically mispriced, potentially due to investor fear or ignorance, which therefore present a high potential upside for us. There are a number of nuances to our style of investing and we continually try to be self-questioning in our assessments of intrinsic company values, asking; 'My valuation model actually indicating that this company is significantly mis-priced - is my model

Fig. 3: The distribution of implied UK equity premium, partitioned by TRBC Sectors.
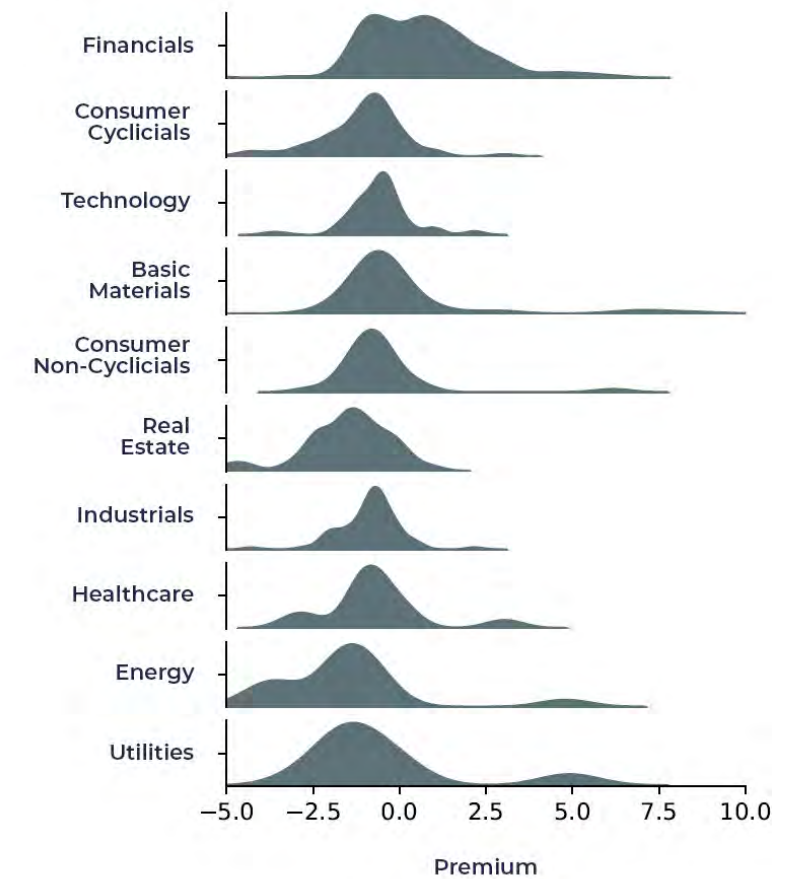
missing something, or is this an opportunity?'. A lot of time and effort goes into our research process, which has been progressively honed over many years, and across numerous investment decisions.

We certainly haven't been correct on all of our decisions, but we have learnt from each of them,

and our investment process and insights are enhanced each time.

For all companies within our eligible investment universe (which we define as being any company traded in the UK with a market capitalisation greater than £100m), we leverage the Kernow Analytics platform to model the possible

The Kernow
INDEX

ranges of intrinsic values of a company. To this aim, we employ a multi-faceted (football field) approach, incorporating a number of valuation models to calculate a range and expected value of each company. We combined this with our deep company analysis in order to determine what we feel is an appropriate valuation for each company. We have to be relatively disciplined to ensure that we are not unduly distracted by the current market valuation of the company, or recent moves in price.

The result of our labour is what we call the Kernow Index - an evolving model which computes our view on the intrinsic value of each company. By combining this with the market capitalisation of companies, we can easily see where our contrarian opportunities arise. The index thus provides us with a wealth of information regarding the potential premium associated with particular investments.

As an example, we may use this to quantify the potential upside in our current portfolio, and compare this to historical levels. It also allows us to highlight sections of our investment universe that are likely to yield more upside than others. In fig.3, we show the current upside potential of FTSE All-Share constituents, partitioned by Sector. Using this visualisation we notice a number of interesting attributes. In particular, we find multi-modality in several of these distributions, which indicates that there are likely to be sub-populations of companies within the sectors, we also see systemic differences in sectors, but a natural tendency of the Kernow Index to have negative premia.

Fig. 3 shows one of the many uses of the index, we also look at the way in which this evolves through time, and observe how our portfolio appears on a scatter plot of the Kernow index vs market capitalisation.

We look forward to introducing several more use cases in future instalments of this journal!

"The result of our labour is what we call the Kernow Index - an evolving model which computes our view on the intrinsic value of each company."

What is the Right Level of Concentration?

In recent weeks we have spent some time looking into what we feel is an optimal level of concentration within our portfolio.

At face value, this seems like a relatively straightforward subject to tackle, but when digging deeper a number of nuances arise. Firstly, when discussing concentration we first need to have a clear metric that allows us to quantify it. Our basic assumption (which we will later show isn't quite correct) is that higher concentration leads to higher risks. We have summarised some interesting results below, from the original report (see Kernow Asset Management, 2022).

This issue is of critical importance for anyone striving to produce a well-balanced portfolio, where risks are appropriately balanced, and return potential is maximised. Our analysis began by reviewing the academic and practitioner literature, and we identified a long and rich set of articles which investigate various subjects relating to concentration, conviction and risk within investment portfolios.

Balancing Diversification and Conviction

Generally speaking, this independent research suggests that managers who hold higher-conviction (i.e. more concentrated) portfolios have tended to do better in terms of absolute returns, and higher concentrations have not led to increases in drawdown or portfolio volatility. These findings generally resonate with our prior viewpoints, and aligns with our fundamental approach to investing. Moreover, from a practitioner perspective, the world's largest sovereign wealth fund, Norges Bank Investment Management (NBIM) published a detailed article about their experiences dealing with external managers (Slyngstad, Hilde, and Management, 2020), this article covers many aspects of manager selection, but on the topic of concentration it states;

'Managers running more concentrated portfolios, whether measured as sector concentration, active share or weight of largest over-weights, fared better than more diversified managers.'

The authors also state the following perspectives;

'Our managers focus on thorough fundamental company research, such as visiting factories and scrutinising reports. This often leads to a concentrated portfolio of companies. This concentration of investments reduces the governance risk, as most aspects of the companies in the portfolio are analysed.'

At the core of our investment process, we perform deep analysis of companies to ensure that we have the best possible conviction when we decide to allocate some of our capital to a particular security. We have found that a portfolio with approximately 20 holdings on the long side and up to 20 on the short side, with limited leverage and a modestly net long exposure produces a portfolio with desirable characteristics. Whilst the specifics of this vary with the macroeconomic environment, our portfolio generally looks like this. From a purely theoretical perspective, we can show that the combination of uncorrelated assets produces a strong diversification benefit when holding very few securities, but this benefit diminishes relatively rapidly with larger numbers of holdings. On the other hand, adding more securities tends to have a dilutive effect on conviction - if we have 20 high conviction holdings, what impact will the 21st have? The result of this tends to be a 'sweet

Fig. 4: A toy model showing the balance between the benefits of diversification and dilution of conviction under three correlation regimes.
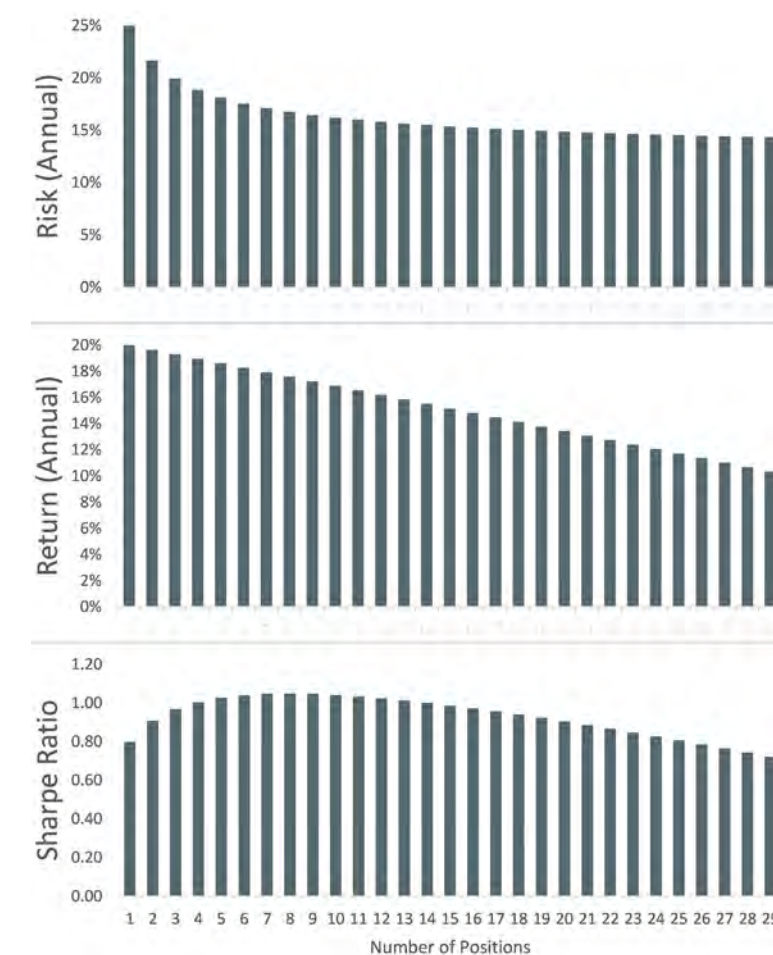
Research
SPOTLIGHT

spot' of approximately 10-30 holdings, with some dependence on the overall market conditions.

By creating a theoretical model using some basic financial theory, in fig. 4 we show different patterns of risk and return for portfolios with varying numbers of holdings, and different levels of assumed correlation - this basic model highlights the interplay between conviction and diversification. Within our analysis we set about to first come up with a definition of concentration - it turns out that there is no standard definition, but there are a number of commonly used

metrics; variations on the Herfindahl Index, the Gini Index and the Coefficient of Variation all yield somewhat different results.
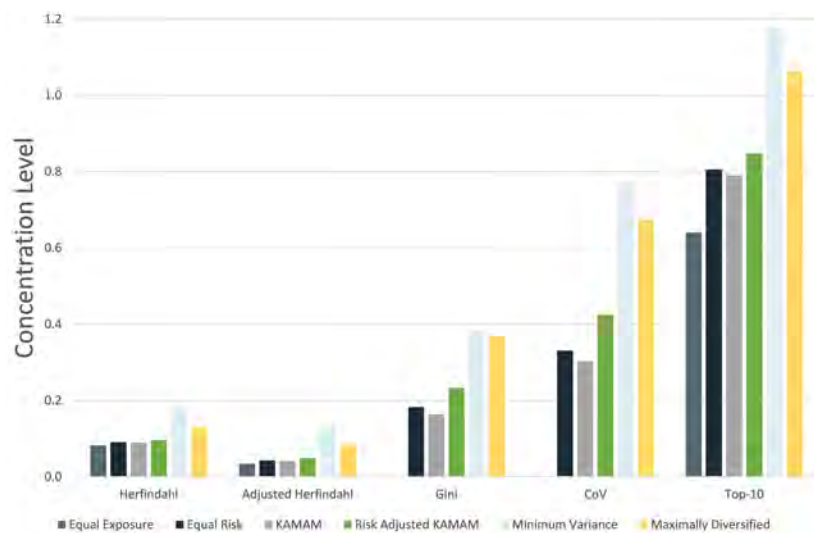
Measuring Concentration

At a high-level, concentration measures attempt to describe the departure of weights in an index (or portfolio) from a 1/N equally-weighted scheme. By this definition, equally weighted portfolios have the lowest concentration (for a given number of holdings). However, these concentration measures are

relatively generic and do not account for a number of key considerations when constructing a well-balanced portfolio - these are;

· What about long-short positions, where portfolio weights can be negative and exposures can be offsetting?

· What about the underlying riskiness of individual securities?

· What about the economic relationships between securities?

· How do we account for leverage, where the gross and net exposure can vary (sometimes significantly) from unity?

To this aim, we have developed our own concentration measure that accounts for each of these, we call this the Kernow Concentration Metric (KCM). Our general conclusion though, is that we feel that there is no singular measure that adequately quantifies portfolio concentration in a way which is a good proxy for the spread of risks within a portfolio.

Fig. 5: The resulting concentration scores using 5 commonly used concentration measures, for 6 alternative allocation methods.

"Our general conclusion... is that we feel that there is no singular measure that adequately quantifies portfolio concentration in a way which is a good proxy for the spread of risks within a portfolio."

Research
SPOTLIGHT

In order for us to highlight the issue when using concentration as a measure of investment risk we construct a number of example portfolios using different allocation methodologies and perform a scenario analysis of each. We enumerate 6 different allocation frameworks; a 1/N (equal weighted) scheme, the current target portfolio weights (KAMAM), and equal risk scheme, a risk-adjusted target weight, a minimal-variance and a maximally- diversified variant. We then compute the average concentration of each of these schemes using the common concentration metrics as outlined above.

Given the relative levels of riskiness, and relationships between securities, we find that the portfolios designed to minimise volatility

have considerably higher concentrations than the other variants, irrespective of the underlying concentration metric used. In Fig. 5 we show this, where the minimum-variance portfolio is considered to be the most concentrated portfolio, followed by the maximally-diversified portfolio.

Whilst we caveat that the calculation of these weighting schemes is done across the full-sample period, and therefore any scenario analysis has a degree of hindsight bias, in Fig. 6 we show the headline performance and risk statistics for each portfolio. We calculate our statistics for a historical scenario running from 2015-present. We find that the minimum variance and minimum volatility construction have the lowest risk (despite

having the highest concentration).

Finally, we notice that the drawdowns are significantly reduced, and payoff ratios improved, however, we note that due to hindsight-bias in the construction and evaluation of the dynamic portfolios. We conclude from this analysis that conventional concentration metrics are not a good proxy for the spread of risk within portfolios.

The Impact Of Long and Short Exposures

A critical component of hedge-fund positioning is in the offsetting risk characteristics of long and short positions. Within the CG Kernow Equity Navigator Fund (KNF) portfolio we strive to produce

Fig.7 The high-level correlation and risk statistics for the long, short and long-short portfolios for comparison. With correlation and beta relative to the FTSE All Share benchmark.

	Correlation	Volatility (%)	Beta
Long	0.72	14.66	0.67
Short	-0.39	14.89	-0.37
All	0.65	10.82	0.45
FTSE All Share	1.00	15.76	1.00

balanced exposures with long and short positions allowing us to capitalise on market directionality whilst constraining overall portfolio risks. Clearly, a concentration measure on the long-only holdings therefore misses the overall spread of risk within this portfolio configuration. To highlight this, in Fig. 7 we show the risk characteristics of the long and short side of our current KNF holdings through time, and of the full (long-short) portfolio. We find that the correlation of the long-only portfolio is 0.72 when compared to the benchmark, with a comparable (but slightly lower) volatility of 14.66% giving a beta of 0.67. However, this is offset by the short portfolio, which has a correlation of -0.39 and comparable (but slightly lower) volatility of 14.89%, yielding a beta of - 0.37. In combination, the long-short portfolio has

significantly lower risk (10.82%) than any of the components, or the index, which is a direct consequence of the offsetting exposures on the long and short side. This lower volatility and reduced

correlation therefore produces a modest beta of 0.45.

Moreover, to demonstrate another way in which we view risk in the KNF portfolio, in Fig. 8 we show a network representation of the correlation structure within our current holdings.
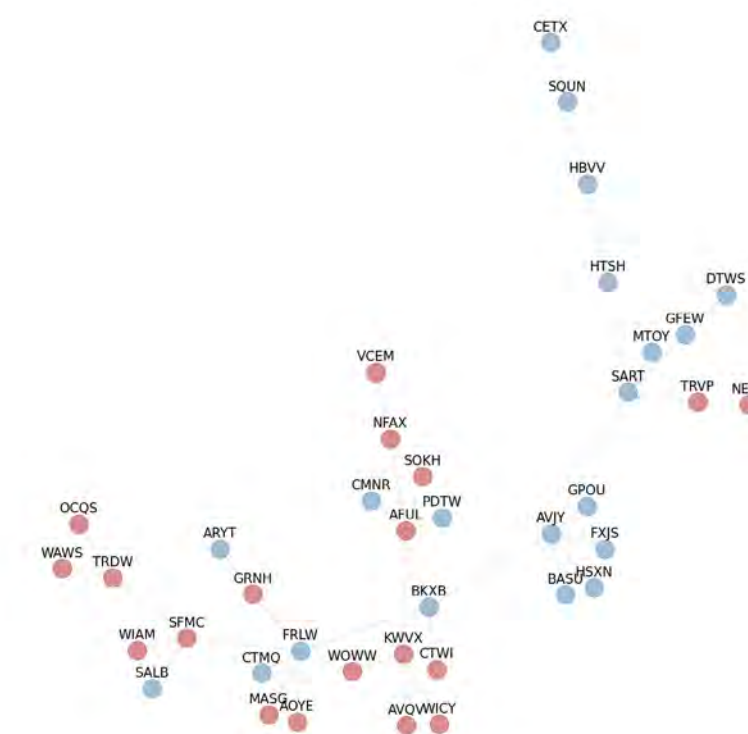
The Impact of Leverage on Concentration?

In the examples shown in Fig.5, we considered

Fig. 8: A force-directed network representation of the minimum spanning-tree transformation of the adjacency matrix formed

Fig 6: The summary (gross) performance and risk statistics for the historical scenario analysis of target portfolios with different allocation schemes (and therefore different concentration levels).

	Equal Exposure	Equal Risk	Kernow Model (KAMAM)	Risk Adjusted KAMAM	Minimum Variance	Maximally Diversified
Return(%)	18	18	17	18	20	18
Risk (%)	20	19	20	20	17	17
Sharpe	0.90	0.96	0.84	0.89	1.16	1.05
Max Drawdown (%)	67	60	66	60	44	61
Hit Rate (%)	54	54	54	54	54	54
Payoff Ratio	1.01	1.00	1.00	1.00	1.02	1.02

"We propose a new measurement method - the Kernow Concentration Metric (KCM), which is designed to account for mutual correlations, security riskiness, the impact of leverage and long-short portfolios."

the basic example of an equally- weighted, unit-leverage, long-only portfolio. Increased leverage further tightens the severity of the constraints; for example, if our portfolio above has modest leverage of 120%, and 20 equal positions, the top-N metric equates to 60% before any price moves of the securities. Under this leveraged scenario, if the winning securities on-average made 20%, the top-N exposure would become 72%, therefore we see that this amplification effect causes the concentration measures to often look exaggerated.

Using this example, we argue that it is cleaner to decouple concentration measures from leverage. We therefore recommend using the total risk capital (defined as the gross portfolio exposure in order to scale concentration metrics consistently. Thus the concentration score would be invariant to the level of leverage.

When To Take Profits And Cut Losses?

A key component of portfolio management is in determining the appropriate entry and exits

of certain positions. The disposition effect typically means that some managers tend to hold onto their losing positions, and also trim their winning positions for longer than optimal. However, it has been shown in numerous studies that diligent management of positions; reducing or eliminating exposure to losing positions and holding the winning positions significantly enhances returns (see Shulka, Rushdi, and Katiyar, 2020 and references therein for a short review).

To illustrate this, consider that a portfolio manager has a fixed hit-rate, equating to a stock picking ability of 50% which is an extremely fine edge. Statistically, if no systematic process takes place to reduce exposure to losing positions and maintain exposure to winning positions, the payoff ratio tends towards unity, whilst diligent portfolio management can push this significantly further upwards. Considering that $Rp = HR \times PR$ where Rp is the portfolio return and is the result of the product of the hit-rate HR and payoff ratio PR, at a fixed hit-rate the success of a portfolio of investments is proportional to the payoff ratio, and thus

the portfolio position sizing.
An Analysis Of Rebalancing

To illustrate how concentration constraints can become artificially binding we show how trimming exposures to winning positions, and topping-up exposures to losing positions has a detrimental impact on portfolio performance.

Using the KAMAM weights as our target, we run a historial scenario analysis which rebalances a portfolio to these weights on different frequencies (daily, monthly and quarterly). By rebalancing daily, we ensure that the portfolio always maintains the target weights, however this does mean that winning positions are trimmed and losing positions are topped-up, monthly an quarterly rebalancing allows the relative weights of the portfolio to drift toward higher weightings for the winning positions and lower weightings to the losing ones until the portfolio is rebalanced. Clearly, the lower the rebalance frequency, the more pronounced these effects are.

Running our scenario analysis historically since

Research
SPOTLIGHT

Fig. 9: A historical scenario analysis, where we use current target weights on a portfolio and allow those weights to adjust with the security movements, weights being rebalanced (brought back to target) on different time horizons.
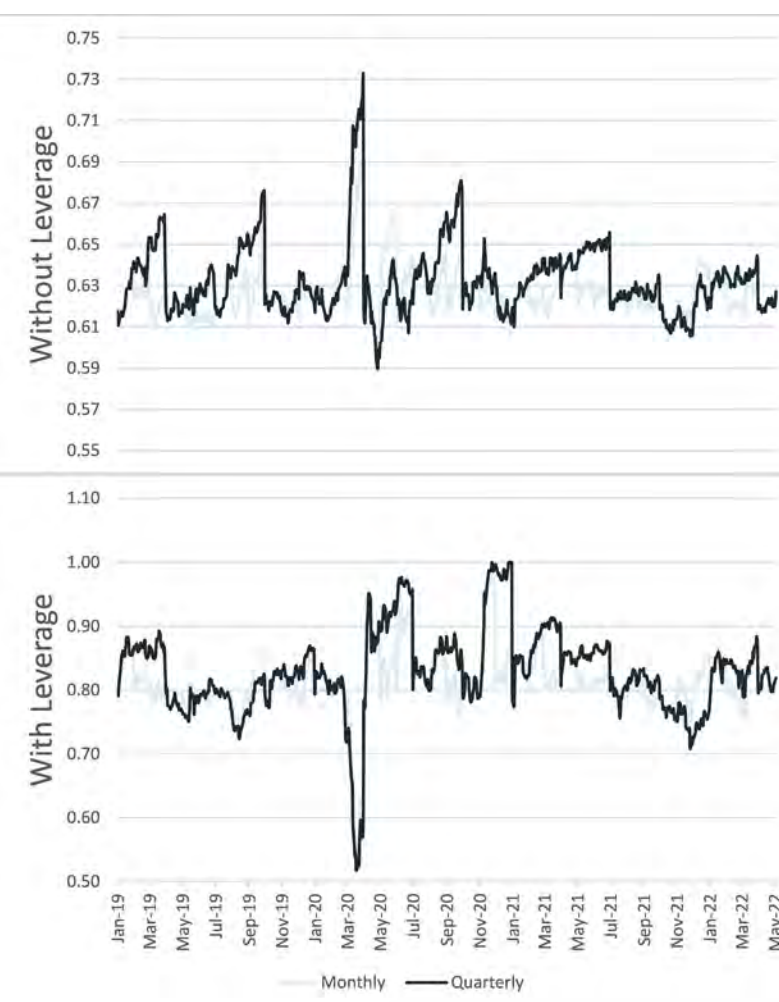
1-Jan-2019, in Fig. 9 we show the time-series of portfolio concentration (as measured by top-10 allocation) with different rebalance frequencies. We notice the monthly rebalancing diverges from this, generally towards higher concentrations, with a time-averaged concentration of 81%, and quarterly rebalancing showing a time- averaged concentration of 83%, with peaks at 98% and 100% respectively.

Whilst this is a deliberately basic illustration of how concentration comes about through the relative performance of the underlying holdings, we also note that the overall portfolio performance is impacted positively by the lowering of the rebalancing frequency.

We summarise the performance and risk characteristics of the different rebalancing regimes in Fig. 10, finding that with fewer rebalances (and therefore more concentrated holdings, and weights tilted towards winning positions & away from losing positions); the overall returns improve, whilst the risk remains relatively stable resulting in significant increases in Sharpe ratio. Moreover, the drawdown of the portfolios reduces significantly; we see that the hit-rate remains stable whilst the payoff ratio increases significantly – generating the improved portfolio returns.

Conclusions

Using the current target holdings of the KNF as our basis we explored a variety of concentration measures and allocation schemes, showing that conventional concentration measures do not sufficiently capture the nuances of individual and offsetting risk characteristics of securities within a portfolio.

We conclude that whilst these measures are helpful in understanding certain aspects of a portfolio, they are generally not good proxies for quantifying the prudent spread of risk. Given our critique of conventional concentration measures, we propose a new measurement method - the Kernow Concentration Metric (KCM), which is designed to account for mutual correlations, security riskiness, the impact of leverage and long-short portfolios.

By producing portfolios with different objectives (such as minimal variance, maximal diversification, equal risk etc.) we find each shows markedly different risk and return profiles under a historical scenario analysis, the concentration measures of these portfolios does not correspond to levels of portfolio risk – providing further evidence for the inadequacy of using conventional concentration measures to quantify risk within portfolios.

Moreover, when managing the relative sizing of positions as a result of the performance of the securities, we find that holding the winners and cutting the losers yields portfolios with better performance on both an absolute and risk-adjusted basis, concentration naturally rises in portfolios which are allowed to hold winners and cut losers - the portfolio risk remains relatively invariant to the resulting concentration changes.

We therefore conclude that concentration constraints are likely to be detrimental to performance when they force the trimming of winning positions and topping up of losing ones.

Our findings are in broad agreement with Yeung et al., 2018; in this paper the authors constructed highly concentrated portfolios by extracting the 'high conviction' holdings from reported portfolio holdings (in a similar way to Anton, Cohen, and Polk, 2021), and found that creating increasingly concentrated portfolios generated increasingly high performance with only modest increases in volatility; thus the generated Sharpe ratios increased despite the decreasing breadth. Importantly, when considering the spread of risks, the authors find that higher concentration does not directly equate to higher risk.

Fig. 10: The summary performance and risk statistics for the historical scenario analysis of target portfolios with different rebalance frequencies

	Daily	Monthly	Quarterly
Return(%)	17	39	41
Risk (%)	20	22	22
Sharpe	0.84	1.79	1.85
Max Drawdown (%)	65	42	38
Hit Rate (%)	54	54	54
Payoff Ratio	1.00	1.14	1.17

Company
SPOTLIGHT

Saga plc

The Kernow team recently visited the headquarters of Saga Plc (SAGA LN).

Saga provides products and services primarily for customers over the age of 50, with its diverse product range spanning from insurance and finance products to cruise holidays and a magazine. We consider the company to be somewhat of a national treasure - providing people of a more senior age (golden oldies), with a raft of services. This demographic also happens to be the wealthiest and fastest growing segment of UK society.

We met with its Chairman, Chief Executive Officer, Chief Financial Officer, and one of its Non-executive Directors, as well as appreciating their striking office in Folkestone.

During our visit, we were able to consolidate our independent analysis, and gain insight into the company culture and personality of Saga's leadership team. You can do all the desk work you want, but nothing comes close to getting to know companies up close and personal!

During our discussions, it was particularly reassuring to hear management's passion for Saga's clients is firmly back. Their rekindled focus gave us some confidence that they are again moving along the right trajectory.

The company lost its way in recent years following a period of private equity ownership and its eventual IPO in 2014. The returning pre-PE Chairman (and son of the group's founder in 1951) and the new team appear to have a steady hand on the tiller.

Having survived the Covid storm, which put the company's very existence into doubt, management is now looking to deliver on its new incentivised vision to 2026 with a minimum share price target of £6. The new marketing push and branding story has gathered traction and rich data improvements will allow it to better cross-sell and enable new products and services to begin to flourish in a light CAPEX manner.

As we see it, there are five levers for share price value creation, which we enumerate below:
- **Transfer of value from debt-to-equity holders as debt is paid down over the next five years before next due date of 2027**
- **A margin pivot**
- **Fostering compoundable growth**
- **Developing new service lines**
- **A perception change in the market from ascribing a conglomerate discount to a re-rating based on its future position as a brand-led market darling.**

This trade therefore ticks all our boxes and feels to us a bit like buying Berkeley Group in 2010 or Games Workshop in 2016. As always though, the success will come down to execution and people. To be clear, this is a short snapshot about the company which does not discuss the risks and other nuances: it is not a recommendation to buy Saga.

Chart of
THE MONTH

At Kernow, we appreciate great data visualisation; the ability to present potentially complex information in a concise, visually stimulating way requires a great deal of skill and creativity. In recent years various publications have taken the humble infographic to new heights. Visually displaying data is nothing new however, in this first edition of the journal we thought we would introduce the chart of the month by giving a brief history of data

Prior to the 17th century, the visual representation of data existed mainly in maps, which potentially showed semi-abstracted forms. However, it is believed that the first true visualisation of statistical data came about in 1644 when Michael Van Langren (a Flemish astronomer) produced a one- dimensional graph of longitudinal differences between Toledo and Rome.

Relatively rapidly following this, and largely attributed to William Playfair many of the currently-recognisable graphs we use today were

developed, including the line, bar, and pie chart.

Histograms, scatter plots, contour plots and others were also invented during this period.

Perhaps one of the most famous examples of early data visualisation is the chart we show in Fig. 10. This chart is widely considered one of the most influential charts of all time, we will attempt to provide some explanation of why.

In 1812, Napoleon marched to Moscow in order to

conquer the city. It was a disaster: having started with around 470,000 soldiers, he returned with just 10,000. The chart tells the story of that campaign. The map details the out-and-back journey of Napoleon's troops. The width of the line represents the total number of soldiers and the colour represents the direction (yellow for towards Moscow, black for the return trip). Below the central visualisation is also a simple temperature line graph illustrating the rapidly dropping winter cold. The chart was created
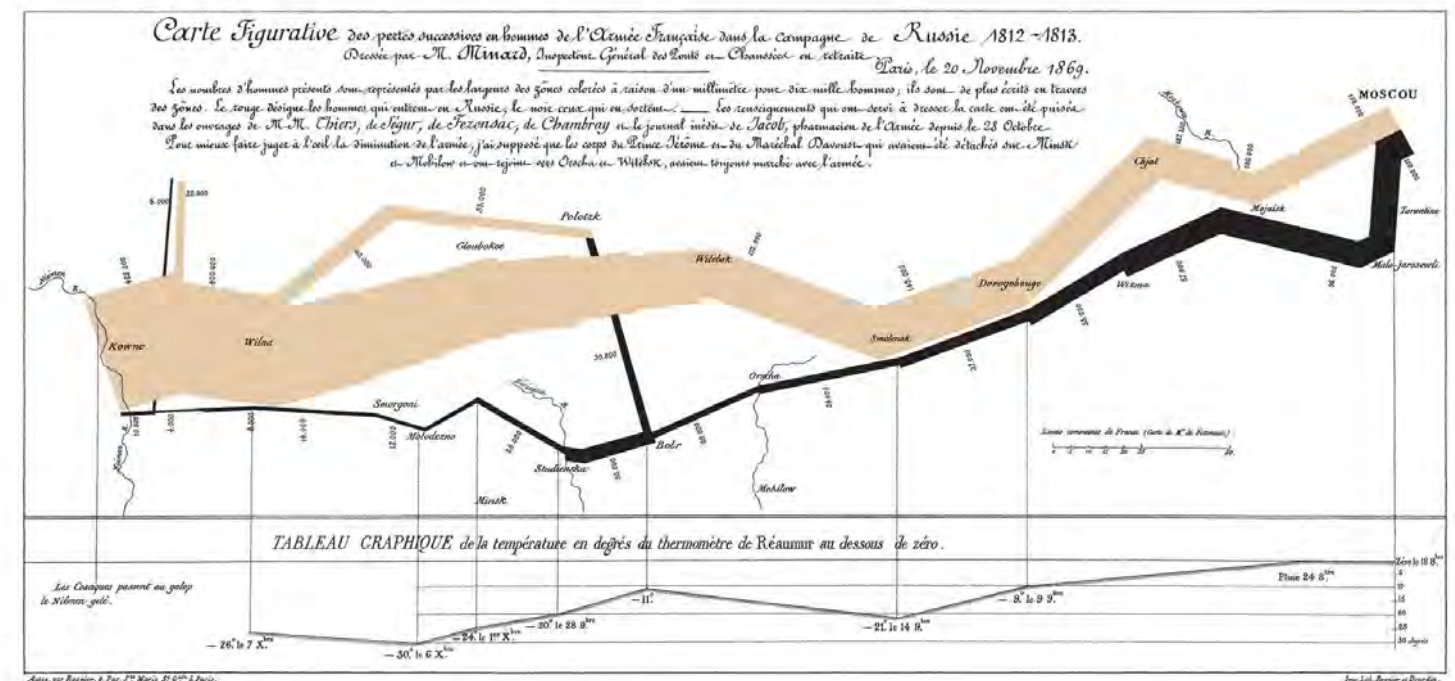
by Charles Minard, who is considered to be the godfather of information graphics.

The chart is effective in detailing the staggering picture of the journey's devastation and it's elegance comes from it's clear displaying of a multitude of dimensions; loss of life at different times and locations, temperature, geography and historical context.

Most charts even today only deal with two or three dimensions. There is a lot

of critical commentary about this chart. A lot of it is reasonable criticism, but this remains a hugely influential and successful chart, one which excels in telling a story with rich detail at each data point and encourages curiosity. Whilst we aim to look at visualisations representing data relevant to our investment landscape, we thought it fitting to introduce our chart-of-the month section with a nod to history.

ENDNOTES

1. Du, Ning and David Budescu (Dec. 2017). "How (Over) Confident are Financial Analysts?" In: Journal of Behavioral Finance 19 (3). URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3499357.

2. Yeung, Danny et al. (Jan. 2018). "Diversification vs Concentration ... and the Winner is?" In: UTS Working Paper Series 18 (1). URL: https://www.uts.edu.au/sites/default/files/wp18.pdf.

3. Shulka, Anurag, Neala Jamal Rushdi, and Rakesh Chandra Katiyar (Apr. 2020). "Impact of Behavioral Biases on Investment Decisions: A Systematic Review". In: International Journal of Management 11 (4).

4. Slyngstad, Yngve, Erik Hilde, and Norges Bank Investment Management (Apr. 2020). "Investing with External Managers". In: Internal Whitepaper 1 (1).

5. Anton, Miguel, Randolph B. Cohen, and Christopher Polk (June 2021). "Best Ideas". In: HBS Working Paper 1 (1). URL: https://www.hbs.edu/faculty/Pages/item.aspx?num=58365.

6. Kernow Asset Management (May 2022). "Concentrated Portfolios: Conviction vs. Diversification". In: Internal Whitepaper.

DISCLAIMERS

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.

KERNOW
ASSET MANAGEMENT

www.kernowam.com

+44 (0)132 676 1826

LONDON OFFICE
33 St James's Square
London
SW1Y 4JS

CORNWALL OFFICE
Admiralty House
2 Bank Place
Falmouth, TR11 4AT