Invesco

Risk & Reward
Research and investment strategies

#03
3rd issue 2021

Photo by Mahdis Mousavi on Unsplash.

Risk & Reward #03/2021

Marty Flanagan
President and CEO
of Invesco Ltd.

In this edition of Risk & Reward, we present two examples of how our quantitative analysts help us achieve two goals: striving to deliver strong performance to clients while supporting the ongoing transition to sustainable energy.

In this edition's feature article, we look at Natural Language Processing (NLP) in the context of Energy Transition. Three of my colleagues have developed a systematic investment approach based on the analysis of company news reports. Applying a sophisticated methodology, they let computers search first for relevant keywords and then for companies mentioned in connection with them. Further filtering and human input ultimately results in a portfolio of companies that are likely to be key players in the ongoing transition to sustainable energy.

In the words of one author: "Our systematic approach is the reason why we are pioneers in ESG." I think this summarizes quite succinctly our values when it comes to sustainable investing. You can read the full interview in this edition.

Another study deals with low volatility and low carbon: conventional wisdom claims that it is impossible to construct a portfolio that fulfils both goals. However, with the appropriate quantitative techniques, portfolio stability can indeed be achieved without the need to invest in carbon-intensive utilities.

Finally, we discuss the role of fixed income ETFs at the height of the pandemic-related market stress in March 2020. We argue that they provided at least some stability in the midst of the storm. Find out why and what this means for investors in the future.

I hope you enjoy the new edition of Risk & Reward.

Best regards,

Mark L. Flanagan

Marty Flanagan
President and CEO of Invesco Ltd.

Sustainable investing meets Natural Language Processing

A systematic framework for building customized theme portfolios

By Yifei Shea, Margit Steiner and Erhard Radatz

We lay out a systematic investment process for sustainable theme portfolios, presenting an Energy Transition portfolio as a case study. Using Natural Language Processing (NLP) techniques, we first define relevant subthemes and compile a theme-specific dictionary. This allows us to select relevant companies and narrow down the investment universe using Environmental, Social, and Corporate Governance (ESG) data before constructing the portfolio.

Sustainable investing has gained substantial investor attention in recent years, and Covid-19 has accelerated this trend.[1] In Q1 2021, flows into the global sustainable universe[2] reached all-time highs for the fourth quarter in a row. Often, sustainability metrics such as ESG ratings and carbon emission scores are integrated into the general investment process. Alternatively, investors may focus on specific sustainability themes. For instance, an investable portfolio can be constructed for the theme Energy Transition, which refers to the changeover from fossil fuels (mainly coal, natural gas and oil) to renewable energy sources.

Energy consumption is responsible for approximately 73% of global greenhouse gas emissions.[3] As such, Energy Transition is directly related to United Nations Sustainable Development Goal 7 (Affordable and Clean Energy),[4] and effectively managing energy consumption is instrumental for tackling climate change. Considering the amount of energy consumed globally, the International Renewable Energy Agency identifies a need for additional investments of USD 15 trillion into the energy sector and reallocation of a further USD 15 trillion from fossil sectors into infrastructure, renewable energy and energy efficiency.[5]

Identifying theme-relevant companies via NLP
Effective investment in a theme such as Energy Transition requires comprehensive and timely identification of relevant companies and thus calls for processing large quantities of data, including textual data from a variety of sources. To this end, we use techniques from Natural Language Processing (NLP). This subfield of artificial intelligence (as well as computer science and linguistics) enables computers to process and analyze text in a manner similar to human beings.

Our approach is dictionary-based, which speeds transparency and facilitates easy review by analysts and portfolio managers. One way to create the dictionary is to rely on domain expertise and use manual assignment of keywords and key phrases. Our method of choice, however, is to leverage techniques from NLP while employing human oversight. Below we describe two of the key NLP methods we applied: topic modeling and keyword extraction.

Topic modeling for subtheme extraction
To identify companies relevant for a broad thematic portfolio, subthemes are helpful for structuring the dictionary. It makes sense to pursue a data-driven approach for extracting relevant subcategories which are likely to make up the overall theme, especially when they are not explicitly defined, such as in the case of Energy Transition.[6] One such approach is a statistical learning method known as 'topic modeling', and Latent Dirichlet Allocation (LDA) is one of the most popular topic modeling techniques. Below we focus on the intuition behind LDA[7] as well as practical considerations for applying this model to identify key topics within the Energy Transition theme.

We start with a collection of documents relevant to Energy Transition, including academic literature, non-profit/think tank publications and sector reports, bearing in mind that there might not be any clear label or category that applies to each document.[8] Figure 1 shows the high-level essence of LDA: we can think of each document as a mixture of a small set of topics, or latent (i.e., not directly observable) variables, which drive the generation of words.

One practical consideration is that key phrases are better suited for identification of sensible topics than individual words. Therefore, we use bigrams and trigrams as observed data for LDA[9] – figure 1 shows

Technical aspects of Latent Dirichlet Allocation
In statistics terminology, LDA is a probabilistic topic model (Blei, 2012). More precisely, it is a hierarchical Bayesian model (Gelman et al., 2013), such that a Dirichlet distribution is used as a prior to the topic distribution for each document, and once a topic is selected, another Dirichlet distribution is used as a prior to the word distribution for the given topic.[10] Several approximate inference algorithms have been developed to compute the posterior distribution of latent variables given the observed documents, including one method based on variational Expectation Maximization (Blei et al., 2003) and another using Gibbs sampling (Steyvers and Griffiths, 2006). The estimated probability distributions enable allocation of the words of the document to various topics.

Figure 1
Structure of Latent Dirichlet Allocation
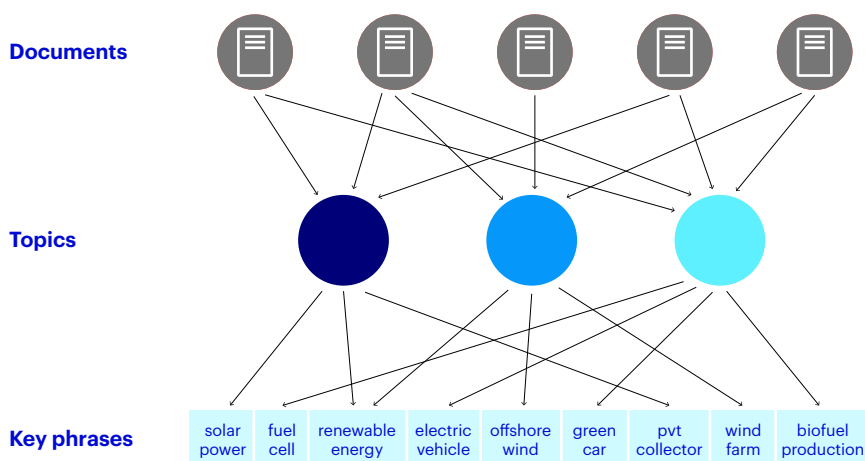
Source: Invesco. For illustrative purposes only.

Figure 2
Coefficient of determination (R^2) for different numbers of topics

Source: Invesco; as of May 31, 2021. We measure the goodness-of-fit of the LDA model for Energy Transition as a function of the number of topics specified a priori, using a generalized form of coefficient of determination (denoted R^2, see Jones, 2019). The model's goodness-of-fit generally improves as the number of topics increases, while incremental improvement tends to be larger below the threshold of 10 topics.

such examples for Energy Transition. For instance, phrases such as 'solar power', 'renewable energy' and 'pvt collector' could be the manifestations of the leftmost topic, which could then be labeled 'Alternative Energy – Solar'.

A limitation of LDA is that the potential number of topics (in our case, subthemes) is unknown and needs to be specified a priori.[11] There are statistical tools to guide this decision, such as measures of topic coherence and the coefficient of determination. Topic coherence measures such as the UCI metric (Newman et al., 2010) and the UMass metric (Mimno et al., 2011) assess the quality of the learned topics by calculating the degree of semantic similarity between high scoring words for each individual topic. The coefficient of determination (R^2) is a well-known goodness-of-fit metric for ordinary least squares regressions. Jones (2019) extended the definition of R^2 to LDA by using its geometric interpretation.[12] As we increase the number of subthemes, the goodness-of-fit of the LDA improves quickly when the number of subthemes

approaches 10, suggesting that, for Energy Transition, an appropriate number may be around 10 (figure 2).

Even though these statistical metrics provide guidance for LDA evaluation, choosing the right number of topics is part art and part science. Note that it may be worthwhile to trade off model goodness-of-fit against the interpretability of topics, which requires human discernment and domain knowledge. In practice, we review the LDA output for Energy Transition at varying numbers of topics. Figure 3 shows the nine subthemes emerging from LDA, which fall broadly in three categories: Alternative Energy, Green Mobility and Energy Utilization.

Some of the most frequent key phrases associated with each of the nine subthemes then become the seed phrases for the theme dictionary. For instance, the Green Mobility subtheme includes phrases such as 'electric vehicle', 'fuel cell', 'green car' and 'biofuel production', to name a few.

Figure 3
Key subthemes for Energy Transition

Alternative Energy	Green Mobility	Energy Utilization
Solar	Green mobility	Management and storage
Wind		Green building infrastructure
Hydrogen		Energy transition
Other renewables		Energy efficiency

Source: Invesco; as of May 31, 2021.

Figure 4
Word co-occurrence graph

	clear	renewable	energy	efficiency	electrification	drivers	deep	decarbonization
clear	1							
renewable		1	1					
energy		1	2	1				
efficiency			1	1				
electrification					1			
drivers						1		
deep							1	1
decarbonization							1	1

Source: Invesco. For illustrative purposes only.

> Our dictionary and subthemes also evolve over time given that they reflect evolving literature and associated news.

Refining theme dictionary through keyword extraction

Note the key phrases suggested by LDA only serve as a starting point and require further refinement. To this end, we introduce RAKE (short for: Rapid Automatic Keyword Extraction), a simple yet efficient NLP method for text summary. The intuition behind RAKE is that keywords typically contain several words, but rarely any standard punctuation or stop words (Rose et al., 2010). Therefore, the essence of RAKE boils down to parsing sentences into sequences of contiguous words by phrase delimiters and stop word positions. Additionally, RAKE calculates keyword scores, such as frequency and ratio of degree to frequency, by constructing a word co-occurrence graph. To illustrate, consider the following excerpt of a recent article from the World Economic Forum:[13] "It's now *clear* that *renewable energy*, *energy efficiency* and *electrification* must be the *drivers* of the *deep decarbonization* we need." The candidate key phrases are in italics while the rest are the stop words.[14] We then construct the word co-occurrence graph of the keywords, where the numbers indicate frequency of co-occurrence in a phrase (figure 4). Note that the only keyword with frequency 2 is 'energy' since it appeared twice in the key phrases.

Using the co-occurrence graph, we can compute the 'degree' of individual words, which is the sum of all frequencies in its row (or column). The degree-to-frequency ratio for each key phrase is then defined as the sum of its member words' scores. For instance, the key phrase 'renewable energy' has a degree-to-frequency ratio of 2/1 + 4/2 = 4. In general, degree-to-frequency favors words that predominantly occur in longer candidate keywords.

In practice, we utilize additional NLP techniques such as annotation (e.g., part-of-speech (POS) tagging) prior to RAKE and employ several methods of text summarization to get a more meaningful list of key phrases. Similar to how we utilize LDA, our approach is to combine the outputs from automatic key phrase extraction methods with human judgement.

Once the dictionary containing the refined key phrases is created, we utilize news data to identify companies involved in the Energy Transition theme.[15] We think that if a company is often mentioned in the theme (measured by news hits of the key phrases in the dictionary), it suggests that the theme is relevant for the company. The same idea also motivates our portfolio weighting methodology to be discussed in the next section.

The curated dictionary allows us to efficiently and dynamically identify global companies that are relevant for Energy Transition. Our dictionary and subthemes also evolve over time given that they reflect evolving literature and associated news.

Practical considerations for constructing a sustainable theme portfolio

In the last section, we described an NLP framework flexible enough to identify companies relevant for any theme.[16] For investing in a sustainable theme such as Energy Transition, this serves as an important first step as it provides a broad universe of companies associated with the theme across sectors and regions. As the NLP framework relies on news flow data, it could also pick up companies in the early stage of Energy Transition. Common examples of this would be an oil company, which might have generated some news flow in articulating a vision for a transition towards newer sources but is still heavily reliant on fossil sources both in terms of revenues and capex. Those companies would not be expected in a socially responsible portfolio.

Next, we discuss how we differentiate and incorporate various ESG scores in a thoughtful manner to narrow down the investment universe for Energy Transition, and how we leverage news information and Energy Transition scores to devise an appropriate portfolio weighting scheme.

The nature of ESG data

Generally, the available ESG data falls into three main categories: 1) activities undertaken by a company; 2) policies and

procedure to ensure compliance with certain standards; and 3) controversies in doing business. The first category is further split into controversial activities and sustainable goods or services. Interestingly, a company may have activities in both areas at the same time: for example, a utility company can own coal-based power plants as well as renewable energy assets. Furthermore, there is a large divergence in ESG data from different vendors,[17] relating to different methodologies as well as materiality mappings.

It is important to understand and combine different types of data to implement the desired ESG profile. In practice, a sustainable portfolio will thus combine a broad set of ESG metrics. It may be screened for revenues in controversial activities or controversies such as violations of the United Nations Global Compact,[18] and it may be tilted towards ESG leaders.

Quantitative approaches naturally lend themselves to ESG integration: given a universe of hundreds or thousands of companies, a quantitative strategy can be formulated to achieve the investment objective while adhering to customized ESG considerations. This applies to factor investing strategies[19] as well as thematic strategies. Specifically, investing in an ESG theme requires a meaningful set of filters to ensure alignment with the sustainability targets – a concept that resonates with the upcoming regulation enforcing the 'Do no significant harm' (DNSH) principle.

One example of an excluded company is an electric vehicle manufacturer. Though one might expect to see the company as a leader in green mobility within an Energy Transition portfolio, and indeed the NLP process alone does pick up the company, it is nonetheless excluded from the strategy due to its less than favorable ESG rating. Both the Governance and the Environmental components of its ESG rating are lagging. The CEO has been involved in multiple investigations by US government agencies in relation to its business practices and external communication (on social media). On top of that, this company fails to have sustainable solutions when it comes to waste management and controling carbon emissions from its Fremont, California production site. While, businesswise, the company might look like a good fit, a strong ESG filter prevents the stock making it into the strategy.

Putting it together: ESG and NLP for portfolio construction

To form the Energy Transition theme portfolio, we first identify all relevant companies across the globe using our theme dictionary and news data. Consistent with our ESG philosophy, we redefine the investable universe by eliminating stocks based on criteria including (but not limited to) business involvement in controversial areas, significant controversies and low aggregate ESG rating.

For the remaining stocks, we apply an innovative proprietary weighting methodology based on news hits, accounting for both the relevance of a company for the Energy Transition theme and the relevance of the theme for a company. Additionally, the weighting of each stock is favorably adjusted if the company ranks high based on the Energy Transition score provided by Vigeo Eiris, which measures a company's ability to transition to a low carbon economy based on current results and strategy. Finally, the portfolio construction process is subject to

> It is important to understand and combine different types of data to implement the desired ESG profile.

> We apply an innovative proprietary weighting methodology, accounting for both the relevance of a company for the Energy Transition theme and the relevance of the theme for a company.

Figure 5
CO_2 intensity in comparison

■ Energy Transition portfolio   ■ MSCI AC World Index
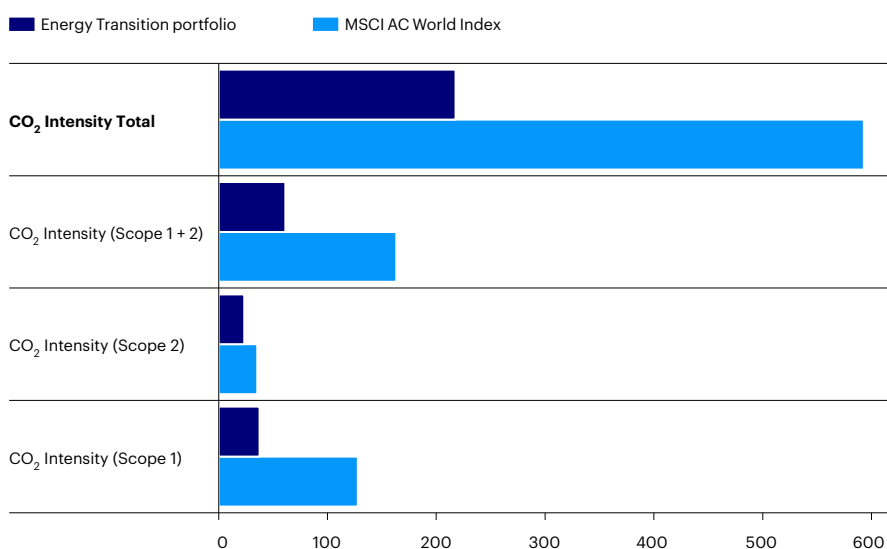
Source: Invesco; as of May 31, 2021. CO_2 intensity as measured by tonnes of CO_2 emissions per USD 1 million revenue. Scope 1 is direct emissions from owned or controlled sources and Scope 2 is indirect emissions from the generation of purchased energy.

Figure 6
Sector weights in comparison
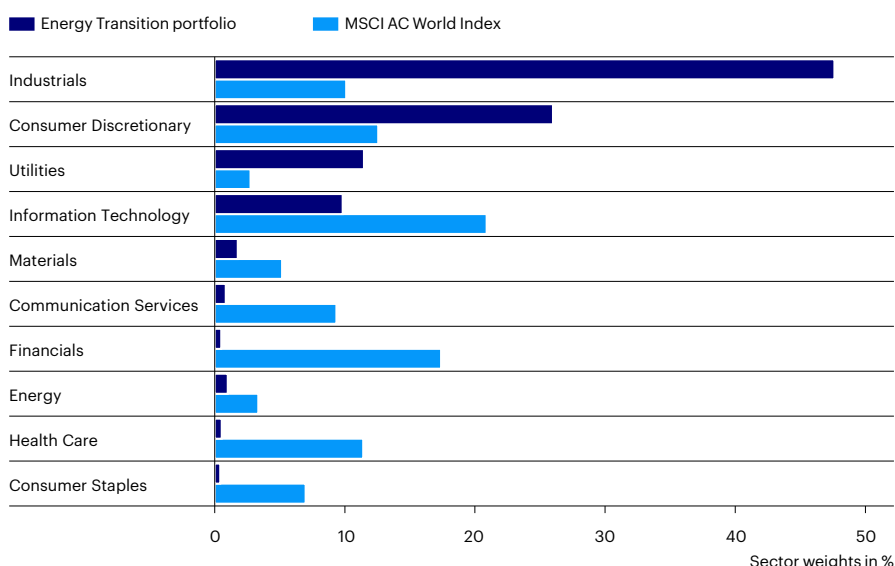
Source: Invesco; as of May 31, 2021.

diversification and transaction cost constraints to ensure investability.

Figure 5 shows the Energy Transition portfolio as of May 31, 2021. It exhibits a massive reduction of carbon exposure compared to the market capitalization weighted MSCI AC World Index. The greenhouse gas intensity is reduced to 218 tonnes CO_2 equivalent per USD 1 million revenue – less than 40% of the benchmark, whose intensity is 593 tonnes per USD 1 million as of May 31, 2021. Similarly, greenhouse gas intensity is also reduced meaningfully if we focus on Scope 1 (direct emissions from owned or controlled sources) or Scope 2 (indirect emissions from the generation of purchased energy) CO_2 intensity. The portfolio also demonstrates superior ESG characteristics, with an overall ESG rating of 6.8 compared to 6.0 for the market index.[20]

The sectorial profile of the portfolio is also interesting (figure 6). There is little exposure to sectors less impacted by the Energy Transition theme: healthcare, financials, communication services and consumer staples might have to change their energy sources, but they won't drive the change. Companies in those sectors are likely to change little even after the broader society has transitioned to sustainable energy usage. Another interesting sector is energy: according to

the GICS definition, this sector represents fossil energy, especially oil and gas. Energy Transition encompasses a divestiture from such unsustainable sources. The portfolio is thus massively underweighting this sector.

On the other hand, enablers of Energy Transition can be found in the utilities sector, which includes renewable energy. The consumer discretionary sector plays a crucial role in the green mobility theme. Since Energy Transition requires massive capital expenditure, the most important sector is industrials.

Summary
ESG-minded investors who wish to participate in one or multiple sustainability themes can benefit from using textual data, NLP tools and ESG information to construct an investable theme-tracking portfolio. In particular, the NLP process facilitates identification of all companies involved and active in the relevant theme. We have used Energy Transition as an example, but the same framework can be applied to other themes such as those tracking the United Nations Sustainable Development Goals. Our process can capture new subthemes as they emerge, and it is fully customizable for the desired investment universe, factor overlay and weighting scheme.

Notes

1 For instance, see "COVID-19 accelerates ESG trends, global investors confirm" by Fiona Reynolds, CEO of Principles for Responsible Investment, available at: https://www.unpri.org/pri-blog/covid-19-accelerates-esg-trends-global-investors-confirm/6372.article.
2 https://www.morningstar.com/lp/global-esg-flows
3 Ge et al. (2020).
4 In 2015, the United Nations set a collection of 17 interlinked Sustainable Development Goals (SDGs), including themes such as Good Health and Wellbeing (SDG 3), Gender Equality (SDG 5) and Affordable and Clean Energy (SDG 7).
5 IRENA (2019).
6 Another example is 'smart agriculture' as a theme.
7 Blei et al. (2003) provide a detailed introduction.
8 Note that pre-processing of the documents, including removing named entities and context-appropriate stop words (i.e., words appearing frequently and considered to have minimum lexicon meaning, such as 'of', 'the'), is important for generating meaningful grouping using LDA. Leung and Gupta (2021) summarize the commonly used NLP pre-processing techniques.
9 In linguistics, a 'bigram' is a pair of consecutive written units, such as letters, syllables or words, and a 'trigram' is a group of three consecutive written units.
10 In Bayesian statistics, a 'prior' represents our guess about the probability before we see any available data (e.g., the documents), while the 'posterior' is the updated probability distribution given the observed data. Dirichlet distribution is a commonly used prior distribution in Bayesian statistics; it is a conjugate prior of the multinomial distribution used in the generative process for topics and words in LDA.
11 Blei (2012) discusses how a non-parametric Bayesian topic model can address this limitation of LDA.
12 A general definition of R^2 is $1 - SS_{res}/SS_{tot}$, where SS_{res} (residual sum of squares) can be viewed as the total squared-Euclidean distance from each observation to its predicted value under a statistical model and SS_{tot} (total sum of squares) can be viewed as the total squared-Euclidean distance from each observation to the mean. In the context of LDA, the observed value (a document) is a vector of integers counting the number of appearances of each word.
13 https://www.weforum.org/agenda/2021/04/why-renewables-are-the-cornerstone-of-the-energy-transition
14 Stop word list from SMART (Salton, 1971).
15 A few iterations of dictionary refinement may be involved as we carefully check for key phrases which may drive false positives in identifying companies for the theme.
16 Elsaesser et al. (2020) describe how a similar NLP framework can be used for investing in innovation themed companies.
17 See Berg et al. (2020) for a detailed discussion.
18 The UN Global Compact is a voluntary initiative based on CEO commitments to implement universal sustainability principles and to take steps to support the UN goals.
19 Please refer to the cited article by our colleagues (Elsaesser and Nerlich, 2020) on this subject.
20 ESG scores are based on MSCI ESG industry-adjusted company ratings. Scores range from 0 to 10, with 10 being the highest score and 0 being the lowest score.

References

Berg, Florian, Julian F. Koelbel, and Roberto Rigobon (2020): Aggregate confusion: the divergence of ESG ratings. Working paper, Massachusetts Institute of Technology.

Blei, David M. (2012): Probabilistic topic models. Communications of the ACM.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan (2003): Latent Dirichlet Allocation. Journal of Machine Learning Research 3 (Jan): pp. 993-1022.

Elsaesser, Georg, and Jennifer Nerlich (2020): Multi-factor strategies and ESG – perfect partners. Invesco.

Elsaesser, Georg, Martin Kolrep, Alexandar Cherkezov, and Michael Rosentritt (2020): NLP in portfolio management: an innovative approach to finding innovative companies. Invesco Risk & Reward, #4/2020, pp. 7-11.

Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin (2013): Bayesian data analysis. Chapman & Hall, London (3rd edition).

Ge, Mengpin, and Johannes Friedrich (2020): 4 Charts Explain Greenhouse Gas Emissions by Countries and Sectors. World Resources Institute.

IRENA (2019): Global energy transformation. International Renewable Energy Agency (IRENA), Abu Dhabi.

Jones, Tommy (2019): A coefficient of determination for probabilistic topic models. arXiv preprint, arXiv:1911.11061.

Leung, Edward, and Tarun Gupta (2021): Machine learning: building factors from unstructured data. Invesco Risk & Reward, #2/2021, pp. 11-16.

Mimno, D., H. M. Wallach, E. Talley, M. Leenders, and A. McCallum (2011): Optimizing semantic coherence in topic models. In Proc. of the Conf. on Empirical Methods in Natural Language Processing, pp. 262-272.

Newman, David, Youn Noh, Edmund Talley, Sarvnaz Karimi, and Timothy Baldwin (2010): Evaluating topic models for digital libraries. In Proceedings of the 10th annual joint conference on digital libraries, JCDL '10, pp. 215-224, New York, NY, USA. ACM.

Rose, Stuart, Dave Engel, Nick Cramer, and Wendy Cowley (2010): Automatic keyword extraction from individual documents. Chapter 1 of Text Mining: Applications and Theory edited by Michael W. Berry and Jacob Kogan, John Wiley & Sons, Ltd.

Salton, Gerard (1971): The SMART Retrieval System—Experiments in Automatic Document Processing. Upper Saddle River, NJ, USA: Prentice-Hall, Inc.

Steyvers, Mark, and Tom Griffiths (2006): Probabilistic topic models. In T. Landauer, D. McNamara, S. Dennis and W. Kintsch (eds) Latent Semantic Analysis: A Road to Meaning. Lawrence Erlbaum.

> ## "Having pioneered ESG investing for three decades, our latest efforts leverage Natural Language Processing techniques to inform our Energy Transition strategy."

Interview with Manuela von Ditfurth, Erhard Radatz and Yifei Shea

> Our systematic approach is the reason why we are pioneers in ESG.

> The main advantage of NLP is its holistic view of all companies engaged in the theme.

Risk & Reward spoke to Erhard Radatz, Manuela von Ditfurth and Yifei Shea about the use of Natural Language Processing (NLP) in sustainable investing, Invesco's ESG approach and the future of clean energy.

Risk & Reward
We are talking about ESG today. What is the benefit of a systematic approach to ESG investing?

Erhard Radatz
Systematic investing and ESG are ideal partners. Systematic strategies seek certain outcomes – be it factor exposures or alignments to specific themes. The individual carriers of these outcomes, a single stock for example, can be easily replaced if an asset fails to adhere to given sustainability criteria.

Manuela von Ditfurth
Our systematic approach is the reason why we are pioneers in ESG. We launched our first ESG portfolio in 1990. Over the years we've developed a range of best-in-class, balanced and low volatility strategies and launched a number of ESG portfolios. In 2019, our first actively managed ESG ETF came to the market. All these strategies are managed following strict ESG criteria and have been awarded renowned labels certifying their ESG standards – and all are based on our integrated ESG process.

Risk & Reward
How does an integrated ESG process work in a quantitative portfolio?

Erhard Radatz
ESG presents itself in various facets. For instance, ESG is certainly an important element when managing risk. This is captured by the 'adverse ESG momentum' and 'ESG exposure control' mechanisms in our investment process. Beyond such quantitative elements, our proprietary

Quality factor captures elements of governance. We practice active ownership, acknowledging our responsibility as an investor. This is an important tool to encourage companies to be more responsible and transparent regarding climate change risks and climate change action.

Manuela von Ditfurth
We engage in regular dialogue with carefully selected companies and have pursued this kind of active engagement for roughly 12 years now. Invesco's proxy voting approach is governed by our proprietary Global Proxy Voting Policy. But as certain idiosyncratic elements can hardly be encoded in a policy – shareholder proposals for example – the portfolio management team reviews the policy output and can override it if necessary. Specifically, this includes voting decisions on climate change reporting, climate change action and other relevant issues. In this way we can enhance the overall ESG performance of the companies we invest in. Given our history and background, it was only natural to develop the Energy Transition strategy, our first ESG-aware theme portfolio built on Natural Language Processing.

Risk & Reward
For someone interested in investing in a sustainable theme, what is the advantage of NLP as opposed to a more traditional approach?

Yifei Shea
The main advantage of NLP is its holistic view of all companies engaged in the theme. A traditional approach would rely on analysts' expertise – which is usually concentrated on one sector or industry. An analyst specializing in the energy sector, for instance, would have missed a significant number of relevant players. Moreover, for portfolio construction we are able to apply a relevance-based weighting methodology using NLP techniques.

> The beauty of the approach is its ability to flexibly respond to new developments in technology or society.

Risk & Reward
Could you provide more insight regarding the relevance-based weighting scheme?

Yifei Shea
Relevance-based means we look at how often Energy Transition is mentioned in news reports about a specific company – and check how often the company is mentioned in reports about Energy Transition. The two scores may well be different. Our relevance weighting is a balance of these two perspectives, adjusted for ESG information such as the company's Energy Transition score from Vigeo Eiris. In a final step, revenue-based exclusions ensure that no company still reliant on revenues from fossil fuel industries is included. This weighting scheme results in a portfolio of companies ranging from 'pure plays' – focusing on a specific subcategory within Energy Transition – to those active in multiple subthemes or even the entire spectrum.

Risk & Reward
How do you weight the subthemes? And which Energy Transition subthemes have the largest weights in the overall portfolio?

Yifei Shea
The weight of each subtheme is determined from the bottom up: as the sum of the weights of all underlying stocks. When a company engages in activities across multiple subthemes, its weight is split between them. As a result, whether we have, say, nine subthemes or three, this has little impact on the constituents and how each stock is weighted in the overall portfolio. The subthemes are nevertheless very helpful for us when it comes to organizing the theme dictionary and dissecting the theme portfolio.

Erhard Radatz
Generally, we expect a balance between the different themes. Currently, the highest weights are in green mobility and wind energy, as well as energy management and storage, whereas hydrogen is one of the themes in a developing phase. It will be interesting to see the themes evolve over time. Certainly, we will see some new themes emerge and others possibly changing in weight. The beauty of the approach is its ability to flexibly respond to new developments in technology or society.

Risk & Reward
We have heard a lot about machine learning. Could you share your thoughts and tell us how it relates to NLP?

Yifei Shea
Technically speaking, our approach for extracting relevant subthemes uses an unsupervised machine learning algorithm. When people talk about machine learning, they often have artificial neural networks in mind. Depending on the specific NLP application, it is sometimes sensible to use a pre-trained model. For example, we use pre-trained models from publicly available NLP libraries to perform Named Entity

Recognition. But pre-trained models will not always be directly useable as context-specific documents may be needed for the training. In this case, we can choose between training a model from scratch or using a pre-trained model and fine tuning it to the specific task. The second approach has recently made considerable progress: one of the milestones is Google's pre-training technique BERT – short for Bidirectional Encoder Representations from Transformers – open sourced in 2018.

In general, machine learning consists of a variety of statistical tools and predictive models. We need to select the appropriate toolkit for the specific task and apply it with care. In a recent paper published in The Journal of Financial Data Science,[1] a few colleagues and I discussed some of the advantages and pitfalls of applying a non-linear rather than linear machine learning models for stock selection. For our NLP framework for sustainable theme investing, we think human supervision is important for identifying the subthemes and evolving the dictionaries, whichever model we use. Once the right tools are in place, we prefer a systematic process to incorporate ESG information and determine company weights.

Risk & Reward
You set strict ESG quality criteria for stock selection. Can you provide some details?

Erhard Radatz
This strategy is for investors who want to do something good while potentially increasing their wealth. Doing something good means first and foremost: avoiding negative impacts. This is similar to the EU taxonomy's 'do no significant harm' criterion. An ESG-aware strategy must have a holistic view of companies. A company engaged in dangerous arctic drilling activities while also investing in an offshore wind farm is unlikely to forcefully drive Energy Transition. It is naïve to assume major producers of greenhouse gases are simply not yet aware of their effects on global warming – but will be, if only there is some shareholder engagement. Companies that still target capex toward fossil energy projects are part of the problem, not the solution.

Risk & Reward
Speaking of engagement, how much active involvement goes on in terms of the Energy Transition strategy?

Manuela von Ditfurth
We engage on topics related to the Invesco Quantitative Strategies priority ESG themes, which include climate change. We select investee companies that, due to their size and stage of development, are likely to be influenceable. The objective is to identify weaknesses in the company's sustainability management and discuss these with the management, aiming to improve its ESG performance in the medium to long term. Discussions can take place via telephone calls, personal meetings and written communication. Engagements are

followed through for years, if necessary. We also leverage Invesco's firm-level engagement and actively support the Climate Action 100+ initiative.

Risk & Reward
Considering the strong performance of renewable energy last year, do you think it is still worth investing?

Erhard Radatz
Clean energy is better for the planet and for humanity than fossil fuels. Clean energy avoids greenhouse gas emissions, delivers cleaner air and brings energy to marginalized communities.

Manuela von Ditfurth
Although fossil fuels still dominate, growth in renewables is accelerating rapidly. Statistics show that wind and solar energy have by far been the fastest-growing energy sources worldwide in the past several years according to the International Energy Agency.[2] Renewables have already or will soon become cheaper than fossil fuels. The costs, for instance, to build solar power or wind plants have come down, and their efficiency and technology has improved significantly. Without a doubt, innovations in the renewables sector will help impel new technologies and support Energy Transition. These trends should continue and lead to attractive investment opportunities in the future.

Erhard Radatz
Even the International Energy Agency acknowledges this need: they are forecasting a four times higher demand for renewable energy, an increase of electric vehicles to 60% of new registrations by 2030 and additional energy efficiency measures. The projected need for investment of USD 5 trillion a year may provide significant investment opportunities.

Risk & Reward
These are some exciting prospects. Thank you for your time.

Notes
1 Leung, E., Lohre, H., Mischlich, D., Shea, Y., and Stroh, M. (2021): The Promises and Pitfalls of Machine Learning for Predicting Stock Returns, Journal of Financial Data Science, 3 (2) 21-50.
2 IEA, Installed power generation capacity by source in the Stated Policies Scenario, 2000-2040, IEA, Paris https://www.iea.org/data-and-statistics/charts/installed-power-generation-capacity-by-source-in-the-stated-policies-scenario-2000-2040.

About the authors

Manuela von Ditfurth
Senior Portfolio Manager, Invesco Quantitative Strategies
Manuela von Ditfurth is responsible for the management of global and European equity portfolios and is an expert in the field of Responsible Investing.

Erhard Radatz
Senior Portfolio Manager, Invesco Quantitative Strategies
Erhard Radatz manages equity, multi-asset and credit portfolios using systematic approaches with a strong focus on ESG considerations.

Yifei Shea, Ph.D., CFA
Senior Quantitative Research Analyst, Invesco Quantitative Strategies
Yifei Shea conducts research into quantitative models that drive the investment decisions for multi-factor equity products. Her current focus includes equity factors, machine learning and natural language processing.

Margit Steiner, Ph.D.
Senior Quantitative Research Analyst, Invesco Quantitative Strategies
Margit Steiner conducts research on quantitative equity investment strategies. Her current work focuses on natural language processing and machine learning.

Low carbon portfolios – why defensive is not always dirty

By Marcus Axthelm and Erhard Radatz

With traditional low volatility portfolios loading up on high carbon emitters like utilities and materials companies, limiting portfolio volatility while pursuing environmental goals seems difficult. To address this issue, we have developed a new and flexible approach that successfully combines low volatility with low carbon exposure.

About 5% of stocks by market capitilization account for 75% of total S&P 500 carbon emissions.

Climate change is among the most prevalent environmental challenges of our times. Scientific consensus that anthropogenic greenhouse gas emissions, especially carbon emissions, are a major contributing factor to global warming,[1] is reflected in the Paris Agreement, which aims to substantially reduce global greenhouse gas emissions. Many analysts deem companies with high emissions riskier as they likely face additional costs in the transition to a low carbon economy.

Carbon intensity, defined as carbon emissions per unit of output, is not equally distributed across the stock universe. Rather, it is concentrated in a minority of high-emitting companies and sectors. One way to illustrate this is a Lorenz curve (figure 1), plotting cumulative carbon emission intensity when sorting companies from high to low carbon intensity. The steepness in the beginning, as well as the flatness in the end, implies a very tilted distribution towards a few major contributors. Indeed, about 5% of stocks by market capitilization account for 75% of total S&P 500 carbon emissions. The distribution looks almost identical for the MSCI World.

Figure 2 shows the carbon intensity of different sectors in the S&P 500, which is heavily tilted towards certain industries. The utilities sector contributes 49% to the total carbon emissions while accounting for only 2.7% of the market cap. When adding in the materials, industrials and energy sectors, 73.8% of the carbon exposure in the index can be attributed to sectors accounting for only 8.2% of market

Figure 1
A handful of companies is responsible nearly all carbon intensity

Cumulated carbon intensity, %

Source: MSCI, S&P, ISS Ethix; as of March 31, 2021, using 2019 emissions data.
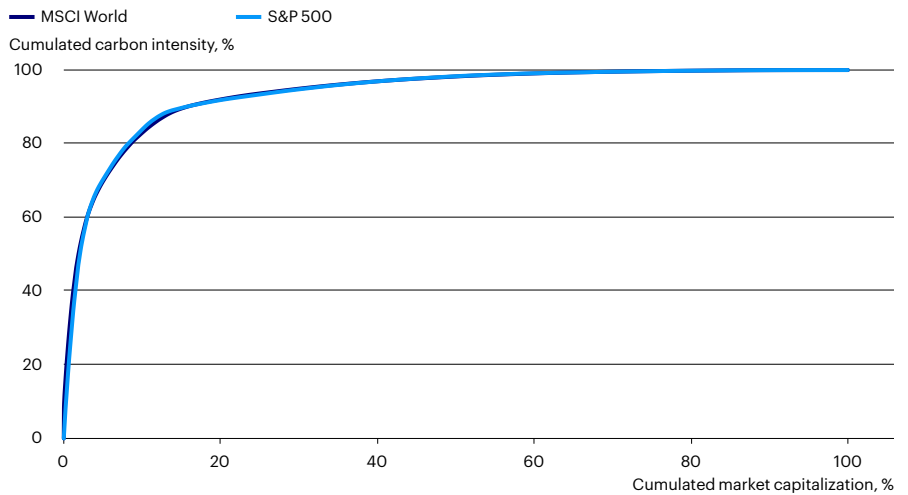
Figure 2
Carbon intensity is heavily tilted to a few sectors
Sector contributions to S&P 500 carbon intensity

Source: S&P, ISS Ethix; as of March 31, 2021, using 2019 emissions data.

"

A sustainable society requires all industry sectors.

capitalization. This indicates that investors looking to reduce their carbon footprint should carefully monitor sector risk or take steps to mitigate it.

Excluding high-emitting industry sectors from portfolios may seem an easy solution. But this would more than likely result in unintended bets and higher volatility. Several other arguments also speak against this approach: first, a sustainable society requires all industry sectors. Modern life depends on the energy (at least currently) and utilities sectors, and they are part of the energy transition. Second, some companies in 'dirty' sectors are arguably driving change and therefore should not be excluded. Last, some weight in high-emitting companies helps with an active ownership approach, urging companies to become sustainable leaders. In some cases, a targeted engagement

approach may lead to more desirable outcomes than outright divestment.

Carbon intensity and investment styles
Style factors such as value, quality or momentum have become an integral part of analyzing portfolio exposures, risks and returns. Whether explicitly targeted through factor strategies or not, all portfolios have factor exposures. And understanding how such (systematic) deviations from the market alter a portfolio's carbon footprint is often the first step in a client's journey towards integrating ESG.

Figure 3 shows the average carbon intensity of major US factor indices compared to the S&P 500. Specifically, investors attempting to reduce portfolio risk, either by simply investing in low volatility stocks or by choosing security weights to minimize the total portfolio volatility (considering stock correlations) face a dilemma: both minimum variance and low volatility indices have much higher carbon intensities than the overall market, as they overweight utilities with low betas and high carbon exposure.

How can we resolve the apparent trade-off between the competing objectives of lower volatility and lower carbon exposure? Given that the minimum variance index has the highest carbon intensity, we will use this as an example to outline a "low-vol low carbon" strategy and then extend the framework by including other factors.

Constructing a low-vol low carbon portfolio
Low volatility investing is popular with investors who want to defensively position their equity portfolios. In times of market distress, the lower beta can limit losses. Also, on a risk-adjusted basis, low volatility has been shown to outperform market benchmarks over long horizons – giving rise to the 'low volatility anomaly'.

There are two approaches to systematic low volatility investing: first, investors may

Figure 3
Low volatility often corresponds to high carbon intensity
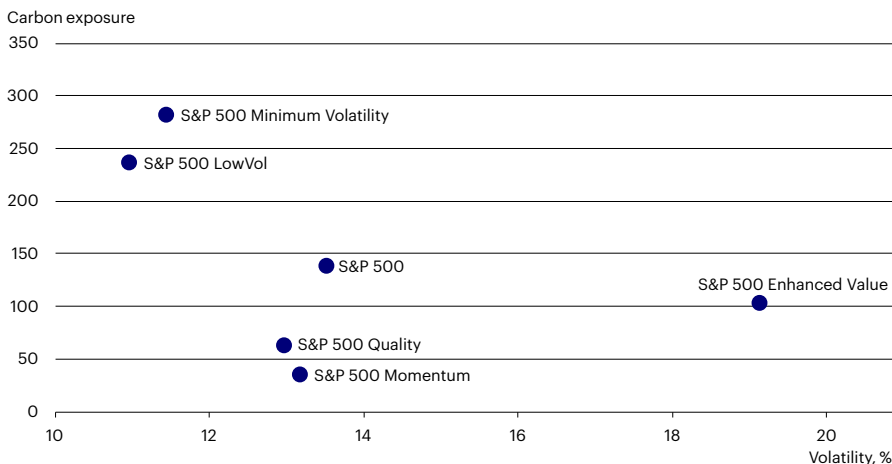
Source: S&P, ISS Ethix; selected factor indices, monthly data, 10-year horizon using 2019 emissions data; as of March 31, 2021.

choose stocks based on their individual return volatility, e.g., overweight stocks with a lower standard deviation of returns. Alternatively, they may wish to minimize overall portfolio volatility, considering return correlations between stocks, which is technically a minimum-variance (MinVar) strategy. Whereas the first approach has the advantage of being straightforward and transparent, the second one takes a more holistic view and allows for explicit control of total portfolio risk. Standard & Poor's uses the S&P 500 Low Volatility Index for the first approach and the S&P 500 Minimum Volatility Index for the second.

Our methodology follows the minimum-variance concept, even though we refer to it as "low-vol low carbon". The goal is to construct a portfolio with a lower carbon exposure and a better risk-return profile than the market. We use an optimization technique because it allows for seamless integration of a variety of constraints. For example, investors may not wish to take on large amounts of unrewarded industry risk resulting from a low carbon bias. In the case of a rules-based, step-by-step investment approach, we would have to fine tune the rules to meet all applicable investment objectives. With optimization, however, limits for regions, sectors and individual securities, as well as turnover controls (such as transaction cost limits) are easier to implement. For a minimum-variance portfolio, we therefore optimize for the lowest total risk given the constraints.

We build three portfolios, all with limited active security, industry and sector weights to mitigate concentration risks. The portfolios are rebalanced monthly.

- **Market 50:** We use the market portfolio as a starting point but cap the carbon footprint at 50% of the market. This makes sense given that a 50% carbon reduction has been identified as one of the key requirements for EU benchmarks to align with the long-term targets of the Paris Climate Agreement.[3] Tracking error relative to the market is minimized, i.e., the portfolio aims to behave like the overall market (beta = 1), but with lower carbon intensity.

- **MinVar:** The second portfolio is the standard minimum-variance portfolio based on the market portfolio universe, minimizing total risk. Since total volatility is minimized, this portfolio will have attractive defensive features, but also a relative high carbon footprint.

- **MinVar 50:** For the MinVar portfolio as well, we add a carbon constraint, again capping the carbon footprint at 50% of the original market portfolio.

Figure 4 and table 1 depict the performance and risk of these three portfolios and compare to the S&P 500 as a proxy for the market. The simulation period is from 2006 to March 2021: starting two years before the global financial crisis, it also includes the market downturn and subsequent recovery during the COVID-19 pandemic.

As intended, the Market 50 portfolio is remarkably similar to the S&P 500, except for the carbon footprint. Similarly, MinVar 50 is very similar to the MinVar portfolio. This indicates that desired risk and return features can be maintained while pursuing a lower carbon footprint.

The key to satisfying both investment objectives, low volatility and low carbon intensity, is diversification. If the equity universe is large enough and imposed constraints do not unduly limit the opportunity set, there are many ways to select and weight stocks that result in a portfolio close to the unconstrained portfolio. An additional carbon constraint can be easily satisfied by choosing stocks with similar risk/return characteristics but different carbon exposure.

> Desired risk and return features can be maintained while pursuing a lower carbon footprint.

Figure 4

Performance of the carbon-controlled portfolios and their conventional counterparts is almost indistinguishable
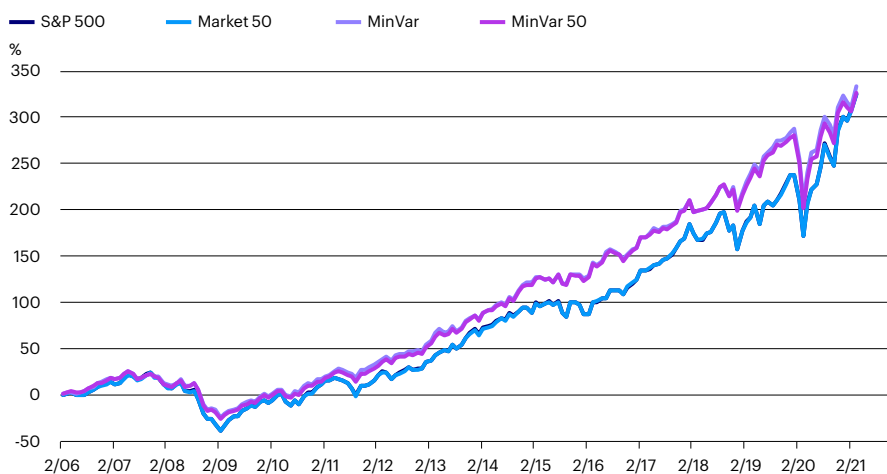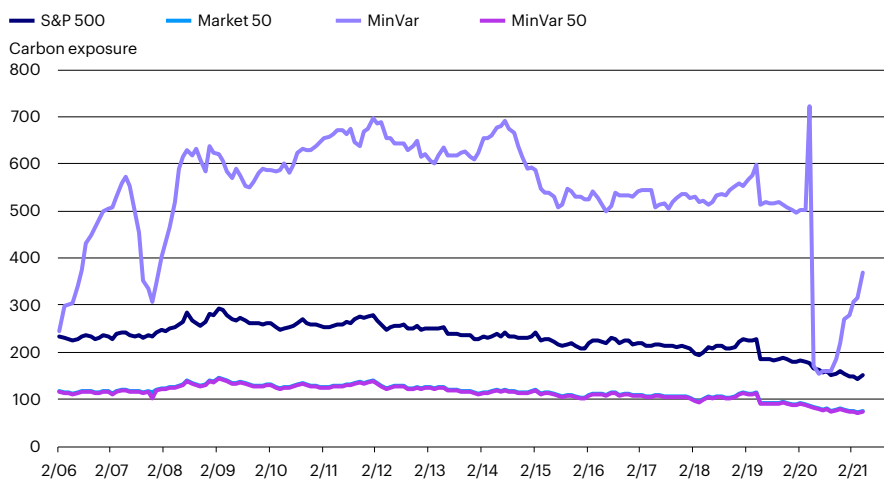

Source: Invesco, as of March 31, 2021.

Table 1
Descriptive statistics

	S&P 500	Market 50	MinVar	MinVar 50
Return p.a. (%)	10.0	10.0	10.2	10.0
Volatility p.a. (%)	15.0	15.0	11.7	12.1
Ann. turnover (%)	5.1	9.5	58.9	61.3
Avg. # of assets	502	549	236	242
Carbon intensity (t CO2/ mUSD)	230.6	115.3	539.1*	115.2
Max. drawdown (%)	-70.9	-70.7	-50.4	-53.2
Sharpe ratio	0.67	0.67	0.87	0.83
Tracking error to S&P 500 (%)		0.16	6.57	5.86

* The carbon intensity of the MinVar portfolio differs from that of the S&P 500 Minimum Volatility index. First, our mean-variance optimization allows for wider sector constraints (10% vs. 5% for the S&P 500 Minimum Volatility Index). Second, our market universe includes up to 25% non-benchmark names in addition to S&P 500 constituents
Source: Invesco, simulation period: January 31, 2006 – March 31, 2021. **Past performance is not a guarantee for future results.**

Figure 5
MinVar is more carbon-intensive than the market – and the carbon-controlled strategies reduce emissions by 50%
Market 50 and MinVar 50, due to their carbon constraint relative to the market, have almost identical carbon intensities over time:

Source: Invesco, as of March 31, 2021.

Adding factor exposures

As we have shown, a low carbon approach can be compatible with low volatility. Investors who want to not only capture the low volatility anomaly, but also target other style factors – such as quality, momentum or value – may want to know whether a low carbon constraint is compatible with such a multi-factor approach, too.

Fraikin, Gerard and Roberts (2020) demonstrate how to integrate low volatility exposure into a core multi-factor approach. We apply a similar methodology to construct core multi-factor portfolios that target quality, momentum and value (QMV) while being defensively positioned, i.e., focusing on low volatility (L).

Our QMVL portfolio comes without a low carbon feature, whereas QMVL 50 restricts carbon intensity to 50% of the overall market. Table 2 and figure 6 show the results of the two new portfolios and the

overall market. From 2006 to 2021, performance of the three portfolios is very similar – and in line with our earlier simulations. The time period includes the well-known factor underperformance of recent years, particularly of the value factor. Longer time horizons and extensive academic research strongly support the notion of positive factor premiums and therefore long-term outperformance compared to the market.

The QMVL portfolios maintain the defensive feature of the original MinVar approach. Most importantly, however, QMVL 50 has almost the same risk and return features as the non-carbon-controlled QMVL portfolio. This supports our claim that a low carbon approach, at the very least, does not hurt performance. Since stock returns are volatile but correlations are often low, targeting well-known and researched factor premia can result in a defensive portfolio with lower volatility and a significantly reduced carbon

Table 2
Descriptive statistics

	Market	QMVL	QMVL 50
Return p.a. (%)	10.0	10.1	10.1
Volatility p.a. (%)	15.0	12.9	13.2
Annual turnover (%)	5.1	91.4	90.7
Avg. # of Assets	502	298	286
Carbon intensity (t CO2/ mUSD)	230.6	389.1	86.4
Max. drawdown (%)	-70.9	-58.9	-59.7
Sharpe ratio	0.67	0.78	0.77
Tracking error to S&P 500 (%)		3.86	3.49

Source: Invesco, simulation period: January 31, 2006 – March 31, 2021. **Past performance is not a guarantee for future results.**

Figure 6
Again, no significant performance impact of lower carbon intensity
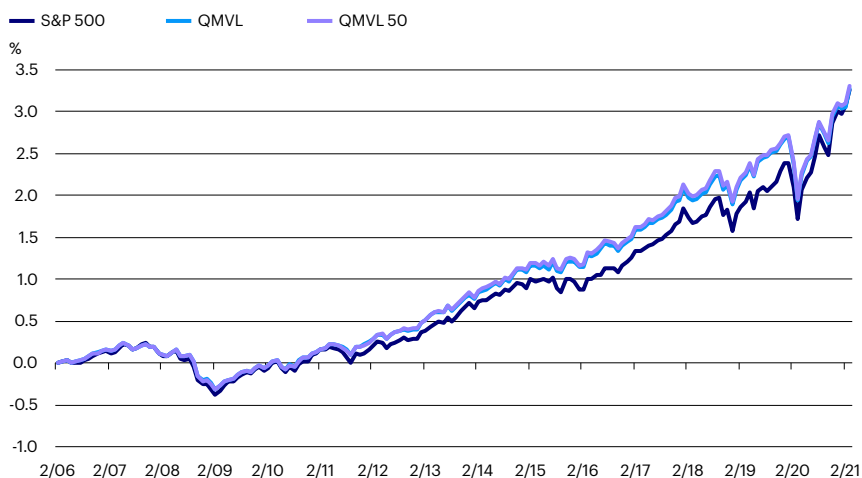
Source: Invesco, as of March 31, 2021.

footprint, yet the potential for long-term outperformance.

Climate scenario modeling
While the inclusion of a carbon reduction goal can be considered a desirable purpose by itself considering the climate crisis, it is also interesting to simulate forward-looking effects, e.g., through climate scenario modeling. Decarbonizing the economy will require policy changes: companies will face financial consequences when their carbon emissions have a price tag attached to them. Von Ditfurth et al. (2021) show how a broad range of ESG exclusions and best-in-class filters can lead to less pronounced drawdowns in the event of such changes.

In this context, we use Vivid Economics/ Planetrics scenarios to simulate the consequences of three possible developments: 'Hot house world' assumes a continuation of current policies, which would lead to falling short of the Paris climate goals; 'Immediate transition' assumes that policies immediately align with the Paris goal of limiting global warming to not more than two degrees compared to pre-industrial levels;

'Delayed transition' assumes only limited climate action today, requiring more drastic measures after 2030.

The effects of these scenarios on individual companies are then modeled through transmission channels, like carbon prices, a changed energy mix or the share of electric vehicles versus internal combustion engines. Generally, companies that have not taken steps towards decarbonization are deemed more vulnerable under the more drastic scenarios and may therefore experience declining valuations.

Even though the simulations use backward-looking data (reported or modeled emissions data), the analysis shows that carbon-controlled portfolios are less vulnerable to policy shocks (figure 7). They benefit from an immediate or delayed energy transition, while the non-controlled portfolios (including the market portfolio) suffer from declining valuations as a result of transition risks.

Conclusion
Integrating ESG considerations into factor portfolios can be a challenge. There are no universal ESG standards. Ratings differ, as

> " Carbon-controlled portfolios are less vulnerable to policy shocks.

Figure 7
Simulated valuation changes – carbon control pays off

■ S&P 500	■ MinVar	■ QMVL
■ Market 50	■ MinVar 50	■ QMVL 50

Simulated valuation changes (%)
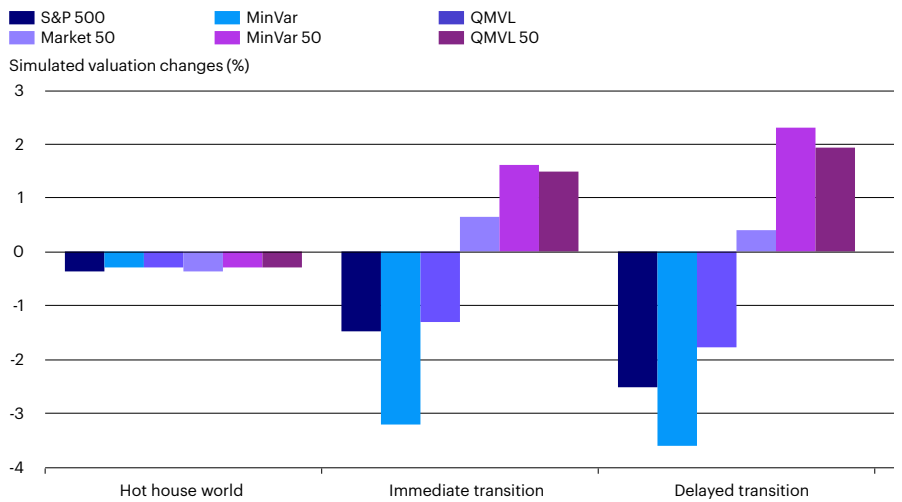
Source: Invesco, Vivid Economics/Planetrics, as of March 31, 2021.

do client objectives, and there are various competing frameworks for the materiality of measurable non-financial metrics. Additionally, ESG metrics correlate with classic style factors, presenting clients with trade-offs between achieving certain factor outcomes and an intended ESG risk reduction.

Nevertheless, we have shown that low volatility and low carbon approaches can be combined without altering the risk/return profile of a portfolio. A 50% carbon reduction compared to the market is easily achievable, even though a relatively unconstrained minimum volatility approach would have a carbon intensity more than 150% above market, or 250% of the market.

Our results also remain stable when integrating other factors such as quality, value and momentum into our low-vol low carbon approach. Such a portfolio is expected to outperform over longer periods while maintaining below-market volatility.

Notes
1   For simplification, the term carbon emissions will be used throughout, but the correct term is greenhouse gas emissions. The six greenhouse gases defined in Annex A of the Kyoto Protocol (1997) are Carbon dioxide (CO_2), Methane (CH_4), Nitrous oxide (N_2O), Hydrofluorocarbons (HFCs), Perfluorocarbons (PFCs) and Sulphur hexafluoride (SF_6). The equivalent carbon emissions are calculated using the global warming potential (GWP) for each gas as accepted by the Intergovernmental Panel on Climate Change (IPCC).
2   EU regulations for Paris-aligned benchmarks also stipulate an annual carbon reduction target and further requirements, which have not been considered here.

References

Von Ditfurth, Manuela, Thorsten Paarmann and Erhard Radatz (2021): Low Volatility and ESG investing combined: Invesco's holistic approach, Invesco Risk & Reward, #2/2021.

Fraikin, Michael, Xavier Gerard and André Roberts (2020): Integrating low volatility style exposure into core equity factor investments, Invesco Risk & Reward, #1/2020.

> " Low volatility and low carbon approaches can be combined without altering the risk/return profile of a portfolio.

About the authors

Marcus Axthelm
Americas Director
Systematic and Factor Investing
Marcus Axthelm helps clients understand
and benefit from the extensive factor
investing capabilities across Invesco
and works with portfolio managers and
research experts on ESG integration.

Erhard Radatz
Senior Portfolio Manager
Invesco Quantitative Strategies
Erhard Radatz manages equity, multi-asset
and credit portfolios using systematic
approaches with a strong focus on ESG
considerations.

Fixed income ETFs: guiding light in market stress

By Justin Danfield, Patrick Galvin and Eric Pollackov

We take a look back at March 2020, when credit markets suffered one of the sharpest downturns ever recorded. COVID-19 had caused a left tail event that may influence the pricing of credit and liquidity risks for decades to come. But even in this extraordinary situation, fixed income ETFs provided at least some stability. Find out why and what this means for their future.

Six of the ten worst performance days ever recorded for the Investment Grade index occurred in March 2020.

In early March 2020, we witnessed the instantaneous earnings impact of social distancing in a consumption-driven economy. Revenue streams tied to human interaction evaporated, causing an abrupt repricing of credit risk as the world desperately sought to "flatten the curve". Credit market volatility dwarfed that of the 2008 global financial crisis, creating horrific liquidity conditions as credit investors rushed for the exits all at once.

During the early phase of the pandemic, much of the focus was on equity market volatility – and rightfully so. The VIX closed over 80 for the first time since 2008 and we witnessed the largest rolling one-month drawdown since October 1931. However, less attention was paid to the equally violent price action in credit markets: during the financial crisis of 2008, the iBoxx Investment Grade Corporates index fell by 15.7% peak to trough within in 278 days. But in March 2020, it declined by 19% in just 14 days (figure 1)!

Six of the ten worst performance days ever recorded for the Investment Grade index occurred in March 2020, including a 5% decline on March 18 – the worst ever single-day decline. Credit volatility was truly unprecedented as investors struggled to assess the impact of social distancing, mortality and credit solvency.

Investors rushed for the exits
Credit risk appetite evaporated in early March 2020 as investors flocked to safe assets such as cash. Money market portfolios saw record inflows while corporate bond funds lost assets amidst the selloff. Mutual funds were hit particularly hard, exacerbating the underlying market stress as portfolio managers were forced to sell bonds to meet mounting redemptions (figure 2). The two weeks from March 16 – 27 saw the largest weekly outflows ever in investment grade, high yield, municipal and global bond funds.

Due to regulatory obligations, the US banking system was better capitalized in 2020 than in 2008. The cash bond market thus continued to function. If investors needed to sell, however, liquidity came at a steep price. Redemptions had flooded dealers with supply, and record credit volatility magnified the liquidity risk premium for dealers to warehouse the inventory. It is worth noting that this all happened while many offices were being shut down, forcing buy-side and sell-side traders to adjust to work-from-home conditions and the associated technological challenges. It is difficult to measure the direct impact of this, but we assume it further fragmented liquidity.

How did fixed income ETFs hold up?
During the stressful period from February to April 2020, fixed income ETFs were not immune to the selling pressure. Nevertheless, outflows totaled just around 2.3% of assets in March – a fraction of mutual fund outflows.

Figure 1
Drawdown comparisons

Source: Bloomberg, as of June 30, 2021.

Figure 2
Outflows from credit funds, February 21 to April 1, 2020

■ Mutual funds ■ ETFs

USD billion
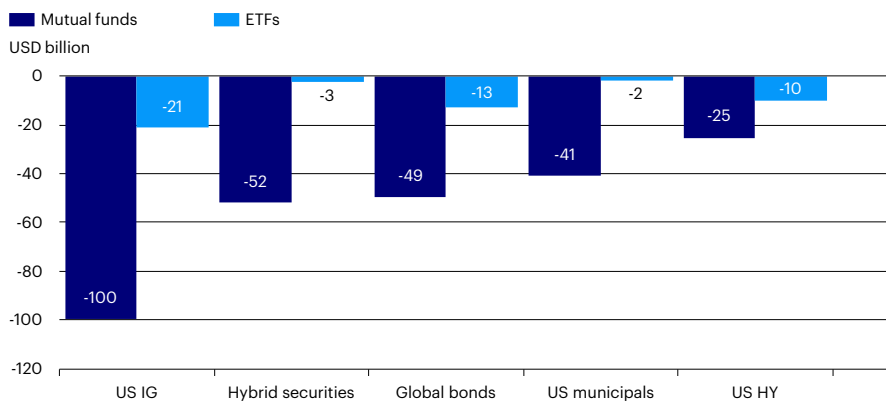
Source: Bloomberg. Data period: February 24, 2020 – April 1, 2020.

Figure 3
ETF flows in the first half of 2020

■ Corporate credit ETF volume ■ Corporate credit ETF net flows

USD billion
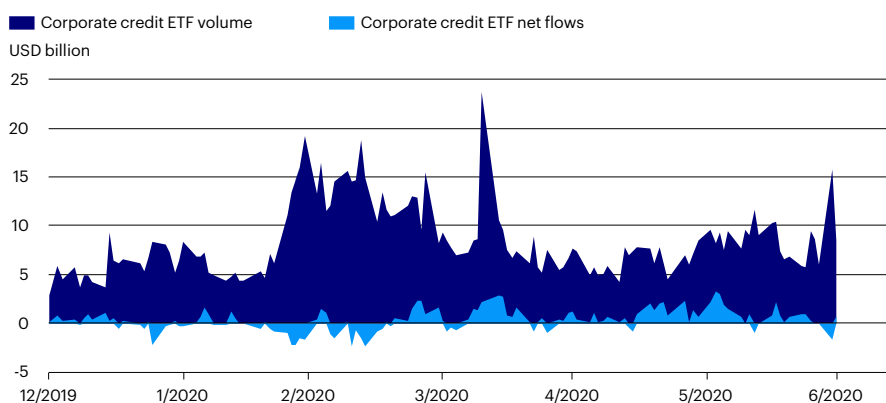
Source: Bloomberg. Data period: December 31, 2019 – June 30, 2020.

Figure 3 shows the total volume and daily net flows of US corporate credit ETFs in the first half of 2020. Trading volumes were exceptionally high, but redemptions were limited.[1] It appears that many investors saw the ETF market as a welcome alternative to the underlying bond market, where liquidity was fragmented. Opportunistic buyers were partially offsetting the selling pressure from liquidity takers. But, given the magnitude of market stress, liquidity conditions and uncertainty at the time, why didn't we observe a flood of credit ETF redemptions? Why did buyers come in and absorb the sales?

In discussions with clients, we often heard that even large institutions with dry powder to deploy found it quicker, easier and more price transparent in spring 2020 to buy credit ETFs on the exchange than to scrape dealer inventory lists to gather up individual bonds. Of course, there were also opportunities to pick up bonds at fire sale prices, but large amounts of capital could more easily be invested in ETFs.

Transaction data from regulatory filings of insurance companies suggests that insurers were one of the client groups tactically buying credit ETFs at this time. They were net buyers of USD 706 million

in credit ETFs in March, on the back of nearly USD 6 billion in credit ETF trading volume.[2] Indeed, the liquidity needs of an insurer are not necessarily tied to market conditions – cash inflows from insurance premiums and cash outflows to process claims for a property & casualty insurer aren't driven by the corporate credit market. This enabled some to put cash to work and capitalize on the market volatility.

Other sources of opportunistic buying during market stress can come from hedging and rebalancing. Examples include long/short credit and global macro hedge funds using credit ETFs for short beta positions, but also for liquidity purposes. When these investors liquidate their short positions to take profits, that can stabilize markets in volatile times, as can rebalancing: pension funds, target date funds and ETF model portfolios are all examples of investors with set rebalancing parameters that can create counter sentiment flow by design.

A diverse investor base is essential for the market to remain liquid. The melting pot of liquidity needs, time horizons and rebalancing parameters magnify this market liquidity – particularly in times of stress. As more and more client types have adopted credit ETFs, the natural

> A diverse investor base is essential for the market to remain liquid.

matching of buyers and sellers has increased, creating an additive layer of liquidity over the underlying market.

What about the discounts to NAV in March 2020?

Some investors have pointed to dislocations between fixed income ETF prices and their net asset values (NAV) as a sign that the ETF structure malfunctioned in March 2020. Figure 4 shows that a broad range of fixed income ETFs indeed traded at a discount to NAV throughout the COVID-19 crisis. These discounts, however, were not an indication of malfunction in the ETF wrapper, but a reminder of the appraisal bias that can cause a lag in NAV for over-the-counter (OTC) markets.

ETF trading in the secondary market is an efficient form of price discovery, as ETFs are traded on exchanges and bound by arbitrage opportunities. If the ETF price is too low, liquidity providers can buy ETF shares at a discount to NAV, redeem those ETF shares for the underlying basket of securities and sell this basket at a higher price. Liquidity providers are in competition for this arbitrage opportunity and are motivated to quickly close any dislocation between the ETF price and the

price at which they can buy or sell the underlying basket of securities. A key consideration here is that OTC-traded holdings in a portfolio do not have official closing prices on exchange, and actionable prices may differ from the evaluated or appraised end-of-day prices in the underlying basket.

In the US, although corporate and municipal bonds trade OTC, these trades are reported each day to the Trade Reporting and Compliance Engine (TRACE) and the Municipal Securities Rulemaking Board (MSRB), respectively – so there is some transaction data to help assess the end-of-day valuation for these bonds. Nevertheless, data remains scarce and price evaluations remain difficult, meaning that any two pricing agents can arrive at entirely different prices for the same bond. In figure 5, we break down the daily TRACE-reported trade data from H1 2020 for a sample of more than 400 investment grade bonds in the Bloomberg Barclays US Intermediate Corporate Bond Index (1-10 yr). In March 2020, nearly 30% of these bonds did not have a TRACE-reported trade, while more than half of the sample did not have a mid-sized or round lot TRACE trade each day. Valuing these

Figure 4
Discount to NAV during the crisis
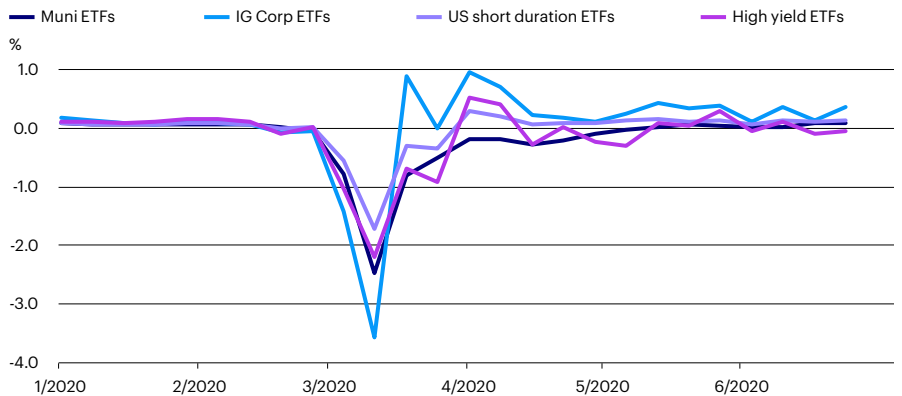
Source: Bloomberg. Data period: January 5, 2020 – June 28, 2020. **Past performance is not a guarantee of future results.**

Figure 5
Investment grade bond trading in H1 2020 according to TRACE
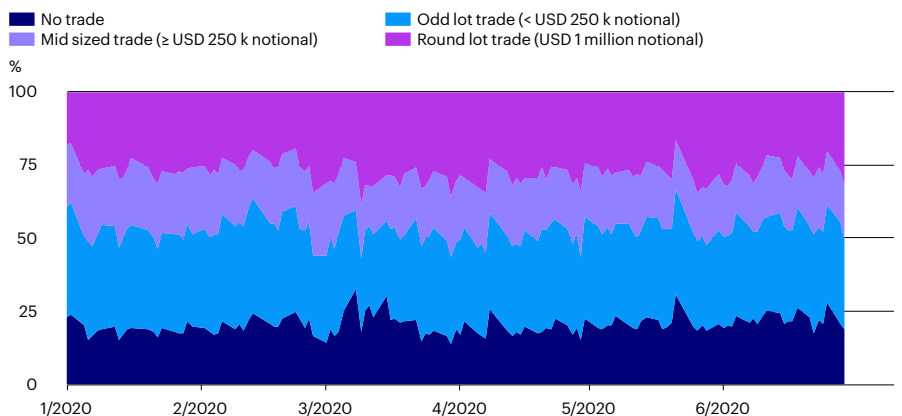
Source: Bloomberg. Data period: January 2, 2020 – June 30, 2020.

bonds with limited pricing signals from live trades can be a difficult task and contributes to a smoothing effect in evaluated prices. This resembles a concept more familiar to investors: the appraisal bias observed in private real estate holdings.

For these bonds, pricing agents can gather indicative prices from dealer quotes as well as traded prices of similar bonds (e.g., sector, credit rating, maturity). These evaluated prices can naturally lag market movements, particularly in periods of abnormal market volatility as dealers may be slow to mark down inventory. In this scenario, which as we can see above could apply to a large subset of bonds, indicative prices may remain well above actionable levels. In March 2020, it was common to see bonds trading several percentage points away from the indicative dealer bid. This is entirely understandable given that we noted some of the steepest daily declines on record and news of initial US COVID-19 cases was unfolding hour by hour. This was a main contributor to the NAV lag, or appraisal bias, that resulted in some fixed income ETFs trading at a perceived discount.

For bonds that do have trade data, the evaluated price can still naturally lag actionable levels in periods of abnormal market volatility. If, for instance, TRACE trade data showed a five-year Company F, US automotive manufacturer corporate bond trading at 10:00 a.m. ET for a price of USD 101.20, yet hawkish Fed minutes released at 1:00 p.m. ET send five-year Treasury yields up five bps and Company G, a US automotive competitor unveils a new pickup truck at 3:00 pm ET forecast to disrupt Company F's pick-up truck sales in the coming year, actionable bids for that five-year Company F bond could sit well below USD 101.20 at the end of the day.

Assessing the end-of-day value for a bond can also be complicated by the details of the reported trade for each TRACE print.

For example, does the price of an odd lot trade reported to TRACE have an associated liquidity premium, and should that price be excluded when assessing the end-of-day value for the bond? What is the right threshold of trade size to constitute an odd lot (<USD 1mm, <USD 500,000 notional)? Should trades that are completed as the result of a firm-specific liquidation event such as a fire sale to meet redemptions be excluded from the valuation? Should these fire sale trades be considered arms-length transactions? These are some of the many issues that complicate the valuation of a bond each day, which tend to cause a lagging or smoothing effect in evaluated prices relative to live bids and offers.

An example of this can be seen in figure 6, with the March 2020 TRACE trade data and evaluated price of a 2024 bond, which we refer to as 'Company A'. In the first week of March, the bond traded every day and the evaluated price hugged the TRACE-reported volume weighted average price (VWAP). In the second and third weeks of March, there were multiple days without TRACE-reported trades and just a few days over the entire month with round lot TRACE prints. Due to the scarcity of trade data and the issues we previously discussed, the evaluated price is smoothed relative to the volatile handful of odd lot prints observed during the depths of the crisis.

ETFs make prices more transparent
Envisioning a hypothetical ETF that holds just this one bond, and stepping into the shoes of a liquidity provider tasked with quoting live bids and offers for this ETF on exchange, it would be reasonable to quote the ETF at a 5 - 10% perceived discount to NAV. The liquidity provider has real money risk when quoting live bids on screen and must thus determine in real time the value at which the ETF shares could be redeemed, delivery of the underlying bond taken and the bond sold in the open market. As a result, the liquidity provider must price-in actionable market levels

> " In March 2020, it was common to see bonds trading several percentage points away from the indicative dealer bid.

Figure 6
Appraisal bias in OTC corporate bond prices
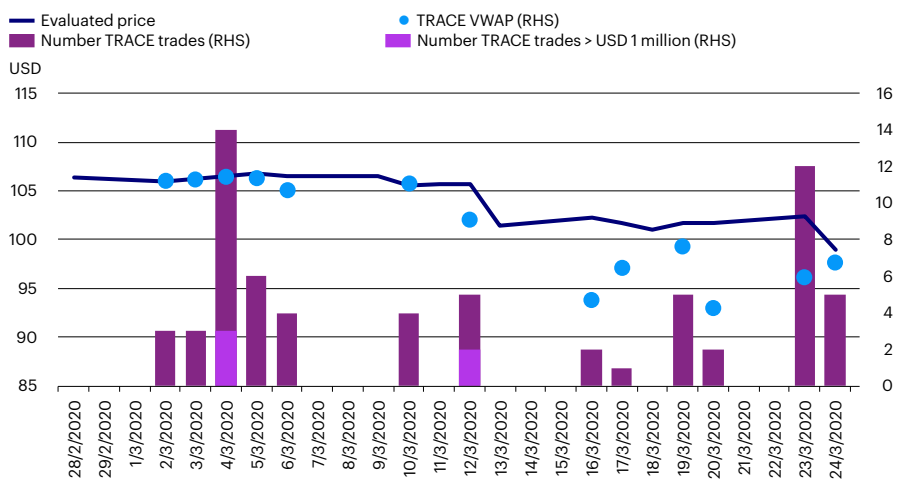
Source: Bloomberg and Invesco. Data period: February 28, 2020 – March 24, 2020.

Figure 7
Fixed income ETF AUM by strategy

- Aggregate
- Bank loans
- Convertible
- Corporate
- Government
- Inflation protected
- Mortgage Backed
- Municipals
- Preferred
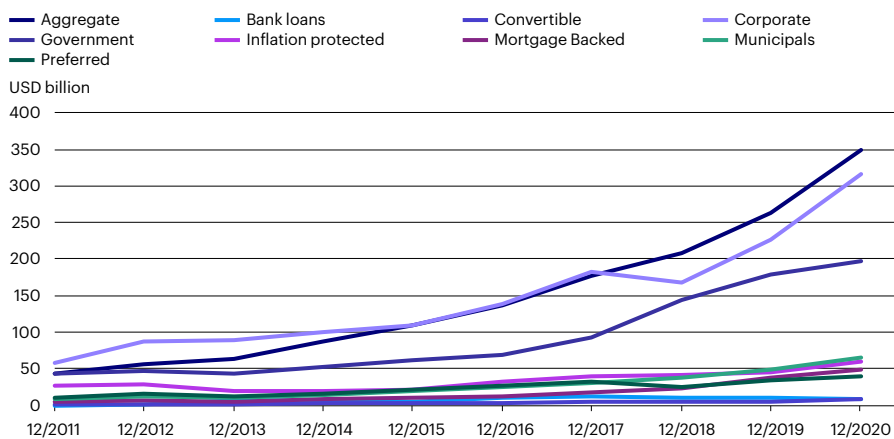
Source: Bloomberg. Data period: December 31, 2011 – December 31, 2020.

Figure 8
Total AUM of actively managed fixed income ETFs
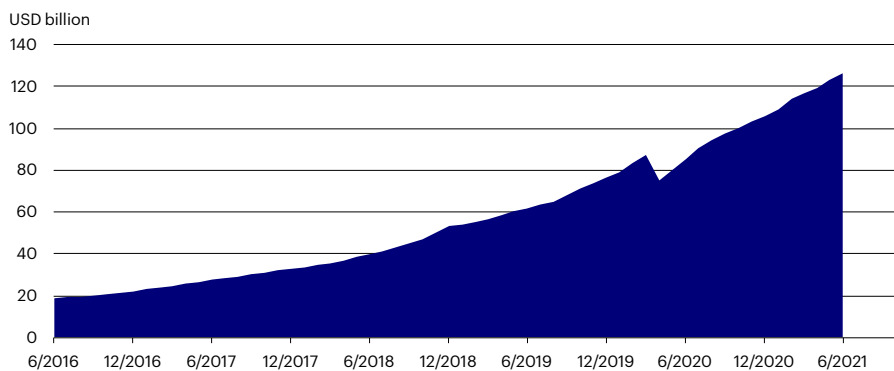
Source: Bloomberg. Data period: June 8, 2016 – June 8, 2021.

---

rather than indicative or evaluated marks for this bond. In this way, fixed income ETFs act as a guiding light for pricing transparency in the underlying market, particularly in periods of fragmented liquidity conditions. This price transparency assists in externalizing transaction costs to each individual ETF investors' on-exchange activity, so that remaining shareholders do not bear the cost of other investors leaving the fund.

A mutual fund holding this same bond would need to sell the bond at market prices below the evaluated marks, the difference being a transaction cost realized inside the portfolio and borne by the remaining shareholders. One might argue that the mutual fund could meet the redemption with cash on hand rather than selling the bond at market prices. However, this would create a cash drag relative to an ETF with the same holdings. As the mutual fund internalizes these costs to remaining shareholders, the ETF structure pushes these transaction costs outside the portfolio in the form of a premium/discount to each individual investor trading their shares of the ETF on exchange. The upshot is that the remaining shareholders do not pay the cost of others leaving the fund. This cost externalization and pricing

transparency, coupled with liquidity, ease of access and breadth of coverage, are key characteristics of the ETF wrapper.

What does the future hold?
Growth and broader adoption will continue as fixed income ETFs continue to prove their utility and resilience. Total assets now surpass USD 1.1 trillion across more than 450 fixed income ETFs.[3] These products provide investors access to a wide variety of fixed income subsectors that were previously difficult to reach. Direct access to market segments such as bank loans historically required significant resources, expertise and institutional infrastructure. ETFs have democratized access to these segments, broadening the range of tradeable asset classes for most investors.

The toolbox of fixed income ETFs continues to expand with a wave of new product innovations coming to market. Investors are beginning to branch out from bulk beta fixed income ETF offerings to embrace concentrated subsector, maturity and credit exposures. One notable area of the fixed income ETF marketplace is the rapid growth of actively managed fixed income ETFs, as shown in figure 7. In 2020 alone, 20 new actively managed fixed income ETFs came to market (figure 8),

and total assets in the segment breached USD 100 billion for the first time. Thus far in 2021, 17 more products have come to market, and actively managed fixed income ETF assets now sit above USD 125 billion.

Conclusion

The structural benefits of the ETF wrapper were on display throughout the COVID-19 crisis. The pandemic brought about some of the sharpest credit market drawdowns ever recorded – much worse than during the global financial crisis of 2008. Fixed income ETFs proved resilient amidst horrific underlying market conditions, trading just as they were designed to in the secondary market – with losses and subsequent gains similar to those of the indices they were designed to follow, but without the sudden illiquidity observed in individual credit and without the massive outflows many credit mutual funds suffered.

The ability of ETF buyers and sellers to match-off on exchange helped alleviate pressure in the underlying market, creating a critical additive layer of liquidity in stress. Fixed income ETFs provided the price transparency and liquidity necessary to weather the storm, serving as a guiding light in this historic period of market volatility. This event may serve as a reminder of the many benefits that fixed income ETFs provide as well as a catalyst for further adoption in the space.

Notes
1   Even though ETFs are traded on the stock exchange, so that for every seller there is also a buyer, the number of ETF units does not remain constant. When there are more sellers than buyers, a market maker (called 'authorized participant') takes units off the market and sells the underlying securities; if there are more buyers than seller, the authorized participant buys the underlying securities and issues new units.
2   Source: NAIC via S&P Global Market Intelligence; as of June 30, 2021.
3   Source: Bloomberg; as of June 30, 2021.

About the authors

Justin Danfield, CFA
Fixed Income ETF Strategist
Justin Danfield researches and develops product-specific strategies and creates thought leadership to position and promote the fixed income ETF lineup.

Patrick Galvin, CFA
Institutional ETF Specialist
ETF & Indexed Strategies Team
Patrick Galvin serves as an ETF specialist providing research, thought leadership and ETF product expertise to institutional clients.

Eric Pollackov
Global Head of ETF Capital Markets
Eric Pollackov develops relationships with sell side desks, implements ETF liquidity strategies and helps clients achieve their goals.

About risk
The value of investments and any income will fluctuate (this may partly be the result of exchange rate fluctuations) and investors may not get back the full amount invested. There are risks involved with investing in ETFs, including possible loss of money. Index-based ETFs are not actively managed. Actively managed ETFs do not necessarily seek to replicate the performance of a specified index. Both index-based and actively managed ETFs are subject to risks similar to stocks, including those related to short selling and margin maintenance.