Sustainable Investing From a Practitioner's Viewpoint: What's in Your ESG Portfolio?

Jeffrey Bohn*, Lisa R. Goldberg† and Simge Ulucam‡

November 8, 2021§

Abstract

Many investors have shifted their asset allocations to account for Environmental, Social, and Governance (ESG) issues. While we welcome this shift from an ethical perspective, the financial and non-financial benefits of ESG investing as well as best practices for portfolio construction are subjects of heated debate. We look at aspects of the debate through a series of practical examples. First, we illustrate the trade-off between risk control and unwanted exposures in energy and "vice" stock exclusions, which have exhibited inconsistent performance at a ten-year horizon. Next, we show how recent underperformance of a gender lens portfolio has been confounded by technology stocks. Finally, we explore how ESG score disparities lead to important differences in portfolios constructed with these scores. In aggregate, our examples point to the inherent complexity of ESG investing, which will benefit from better data, transparency, customization, and an acknowledgement that doing good does not necessarily lead to doing well. An important theme throughout this paper is that everything should be made as simple as possible, but no simpler.

*Consortium for Data Analytics in Risk, University of California, Berkeley and One Concern, email: jeffrey.r.bohn@gmail.com.

†Consortium for Data Analytics in Risk and Department of Economics, University of California, Berkeley, and Aperio Group (part of BlackRock), email: lrg@berkeley.edu.

‡Aperio Group (part of BlackRock), email: Simge.Ulucam@blackrock.com.

§We are grateful to Lloyd Han, Mark Bateman, the participants of the *Journal of Investment Management* Spring 2021 conference and the June 2021 AperioX seminar.

1 Introduction

Trends in investing typically follow a predictable cycle that begins with a new innovation, which finds its way into the investment policies of early adopters. If the innovation gains traction in the broader investment community, a wider cross-section of investors and institutions incorporate the innovation into their strategies. The innovation then becomes overly promoted, which eventually leads to some degree of disillusionment. At this point, researchers, analysts, vendors, and portfolio managers take a step back and reflect on the sources of concern. In some cases, the innovation disappears, while in other cases, it finds a sustainable path toward widespread adoption.

Environmental, social, and governance (ESG) investing is a good example of a recent investment innovation that appears to be generating some disillusionment, even as it is embraced by the financial services industry. Advocates make strong claims regarding the financial performance of ESG portfolios and their ability to affect social change. Naysayers point to inconsistent ESG metrics, particularly noisy data, and the entanglement of ESG factors with standard risk factors, as reasons to question much of the reported ESG portfolio performance. Moreover, lack of standardized ESG metrics creates opportunities for greenwashing by companies less interested in pursuing legitimate ESG objectives, than in building brand by meeting minimum criteria for a metric that may be easy to game. Finally, the vast, internally inconsistent array of ESG issues makes it difficult or perhaps impossible to design ESG investment strategies with broad appeal. These considerations make ESG investing inherently complex.

In this article, we take the complexity of ESG investing as a given, and focus on trade-offs that arise when constructing ESG portfolios with techniques used routinely by financial practitioners. We consider both market-cap-weighted and optimized portfolios representing different ESG themes. As we walk the reader through our examples, we identify factors that may affect outcomes, but may not be apparent at first glance. Each extension to a given analysis improves our ability to identify potential anomalies that might arise from ESG investing. We illustrate, for example, that the return and risk profiles of strategies that exclude unwanted securities from an index depend materially on how admissible securities are weighted. Specifically, the process of cap-weighting admissible securities leads to a sometimes substantial large-cap bias relative to benchmark. In contrast, optimizing a portfolio of admissible securities to reduce tracking error leads to a sometimes substantial correlation bias, overweighting securities that are correlated with exclusions, and potentially just as objectionable. We show that unintended sector biases can drive the performance of ESG portfolios. In addition, we highlight portfolio-level differences that arise from relying on different ESG scores.

Our results demonstrate that even simple rules, such as excluding a subset of securities based on straightforward criteria, can lead to difficult trade-offs, unintended factor bets, time-sample dependent performance that clouds evaluation, and outcomes that need not align with the investor's original motivation.

Our results also show the challenges associated with using ESG scores that may reflect more sophisticated methodologies, but still face complexities in implementation, evaluation, and interpretation. These results highlight the need for transparency and best practices in ESG investing.

Portfolios that include ESG considerations may or may not generate better risk-adjusted returns than purely financial-return-focused portfolios. Naturally, a subset of ESG portfolios will outperform, and unsurprisingly, outperformance depends on time, geography and countless other factors. Often, performance of an ESG portfolio can be explained entirely with standard risk factor exposures. In any case, transparent and consistent ESG portfolio construction is one way to do good in the world, whether or not it leads to doing well.

Before we describe our ESG portfolio construction approaches and our efforts to explore outcomes based on these different approaches, we highlight some of the related findings in the published literature related to ESG portfolios. This short non-exhaustive survey of related literature provides background and context related to the complexities illustrated in our analysis.

2 Related Literature

Over the past decade, the number of ESG-related publications has grown dramatically. Several papers highlight issues similar to the ones we address in our paper. We summarize a few of the more salient findings related to the results we describe in our paper. An overview of ESG investing by the OECD discusses empirical findings, practices, and motivations, and it includes an extensive list of references, especially to the practitioner literature; see Boffo & Patalano (2020).

The idea that some forms of ESG may be good for a corporation date back at least to Friedman (1970), and there is a lively discussion of which ESG attributes are financially material to which companies rages on as providers of ESG ratings compete for market share. While at least some investors have social or ethical reasons for ESG investing, Amel-Zadeh & Serafeim (2018) report that 82% of investment professionals they surveyed believe that incorporating ESG information into an investment evaluation process is financially material to investment performance. Respondents also indicated they typically use ESG information to assess a firm's risk. An emerging literature suggests that funds may use ESG investing to repair damage to reputation; see, for example Khanna & Warburton. (2021).

The extent to which ESG considerations affect a company's risk profile is widely debated. Early work by Chia et al. (2009) identified the challenges associated with extracting a "green" factor related to renewable energy. These authors find evidence that such a factor does exist after controlling for other risk factors. A decade later, Giese et al. (2019) argue that highly-scored ESG companies have lower cost of capital and lower idiosyncratic risk, in addition to higher profitability and lower exposure to extreme risk. More recently, using methods from machine learning, Goldberg & Mouti (2021) find that ESG factors do not contribute to the predictive power of models that forecast drawdown with

standard financial indicators.

Madhavan et al. (2021) explore how ESG components correlate with typical risk factors (e.g., market, size, value, quality, momentum, etc.) Their assessment demonstrates the importance of disentangling ESG and standard risk factors. These authors find for a large sample of active equity funds that "...ESG exposure was rewarded—especially for funds with high environmental scores associated with large quality and momentum factor loadings. But the link between high ESG ratings and high returns is only through the ESG components that are correlated with factor components. Other ESG components unrelated to factors carry insignificant excess return premiums that are economically small." (p. 85) These results highlight the point that an ESG portfolio may be masquerading as a portfolio with specific factor exposures.

One aspect of the ESG debate related to the interplay of expected and realized returns can be quite subtle, but still very important as highlighted in the recently published work of Pastor et al. (2019). The argument's narrative begins with the idea that a high-ESG firm should see its cost of capital fall as more investors buy the firm due to its benefits (e.g., reduced long-term risk, better quality, etc.) As cost of capital falls, the high-ESG firm's expected return should also fall. That is, "good" firms should enjoy the benefit of having more investors willing to provide capital. Thus, in equilibrium, high-ESG firms should have lower expected returns.

The same authors point out in Pastor et al. (2021), however, that realized returns for a green factor for the past several years (2012 to 2020) have been high and reflect outperformance. What explains this counter-intuitive result? It likely arises from the "surprise" in how "green" investments have suddenly become more fashionable. This narrative arc and accompanying results highlight the complexity we are exploring. Thus, high-ESG firms may suddenly see outperformance as a consequence of surprises that may result from a "good" firm that was initially ignored or mistakenly deemed "bad" and is suddenly deemed high-ESG thereby attracting new capital. Alternatively, the notion of ESG investing may become appealing enough that suddenly new capital chases high-ESG firms. These papers illustrate how sorting out time-sample, geographic, measurement, and market dependencies contribute to ESG-investing complexities. It is possible, of course, that trends in investor sentiment play a role in observed ESG outcomes; see for example Serafeim (2020) and Mahmoud & Meyer (2020).

3 Portfolio Construction

As discussed in Geddes et al. (2015), Cooper et al. (2016), Bender et al. (2018), Henriksson et al. (2019), Branch et al. (2019) and many other publications, the method for constructing an ESG portfolio plays an important role in determining its factor exposures and risk and return profile. We expand on this theme in the examples studied below.

Many ESG strategies exclude unwanted stocks from a diversified index.[1] Popular exclusions include oil stocks, tobacco and other "vice" stocks, and stocks with poor records on social issues such as diversity or labor practices. Motivations for avoiding securities include both ethical and financial considerations. The exclusion strategies considered in this article are based on the MSCI ACWI Index, which serves as both a universe and a benchmark. Given a set of unwanted stocks in the index, we construct a Simple Exclusion by cap-weighting the remaining stocks, and an Optimized Exclusion by weighting the remaining securities to minimize forecast tracking error to the benchmark. Simple and Optimized Exclusions based on the same securities can have very different return and risk profiles.

When scores are available for securities in a benchmark, we can construct ESG portfolios by tilting away from the benchmark. For example, an investor may ask for a minimum tracking error portfolio whose carbon footprint is a fraction of the benchmark's. Portfolios of this type are sensitive, of course, to the way the scores are calculated.

Tracking error forecasts used to construct Optimized Exclusions and tilts are generated by the Barra Global Equity Model (GEMLTL), and we use Barra's Portfolio Manager to construct and rebalance optimized portfolios on a monthly basis.

4 Energy Exclusions

An investor who objects to fossil fuel companies on ethical or financial grounds may wish to exclude the GICS Energy sector, which includes oil producers such as Exxon as well as oil drilling companies such as Schlumberger and Halliburton. We look at the return and risk profiles of Energy Exclusions over the period from August 1995 to December 2020.

Rolling ten-year active returns of Simple and Optimized Energy Exclusions are shown in Figure 1. Consistent with its construction principles, the Optimized Exclusion hugs the benchmark more closely than the Simple Exclusion.

[1] An overview of ESG exclusion strategies is Dimson et al. (2020*b*). Atz et al. (2021) argue in a meta-study that ESG integration performs better than exclusion.

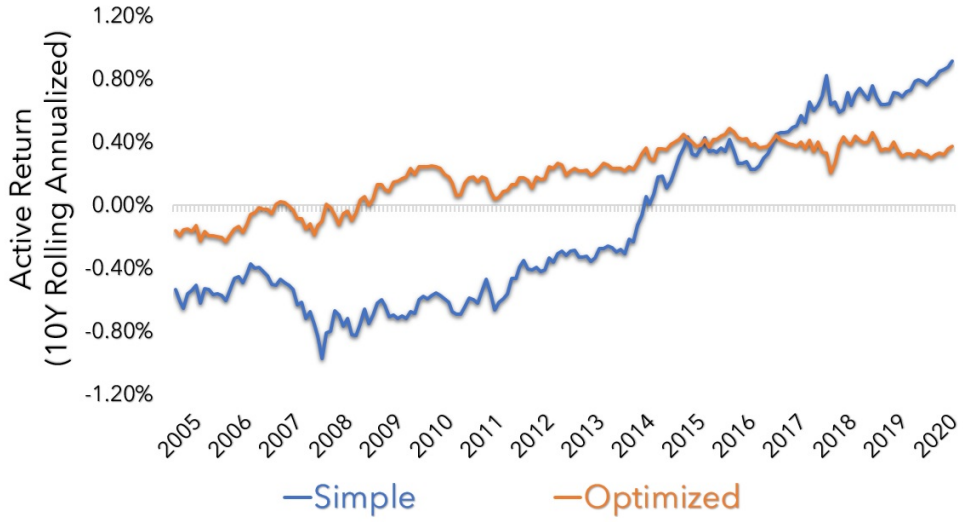

Figure 1: Rolling ten-year annualized active returns to Simple and Optimized Energy exclusion strategies benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

As shown in Figure 2, the rolling ten-year annualized active returns to the Energy Exclusion strategy have mirrored the Energy sector weights. Thus, the weight of the Energy sector in the index at any time explains, to a great extent, the active return to a Simple Energy exclusion strategy over the preceding ten years.[2]

[2] An OLS regression of ten-year returns to the Simple Energy exclusion strategy onto the end-of-period sector weight over our study period yielded an R-squared value of 0.81. The analogous figure for the Optimized Energy exclusion strategy was 0.49.

Figure 2: Weight of the GICS Energy sector in the MSCI ACWI Index and ten-year rolling annualized active returns of a Simple Energy exclusion strategy benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

Turning to risk, Figure 3 shows the evolution of tracking error for the Simple and Optimized Energy Exclusions. Both rolling three-year realized tracking error (panel a) and forecast tracking error (panel b) were uniformly higher for the Simple Exclusion strategy than for its Optimized counterpart.[3] The differences were greatest during the 2008–09 financial crisis, when tracking error was elevated.

[3]The three-year history used to estimate realized tracking error explains why it is more muted than forecast tracking, which is tuned to a shorter horizon.

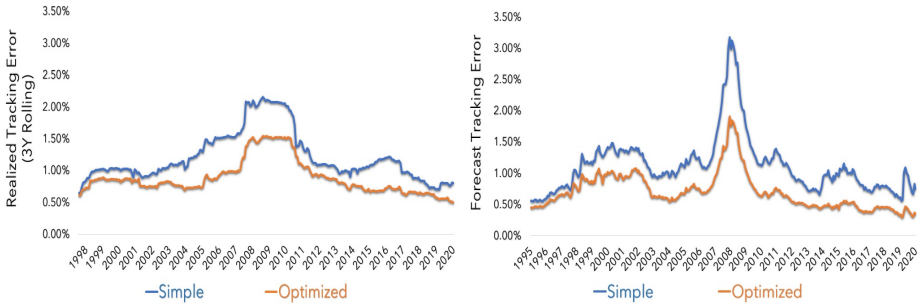

(a) Rolling 3-year realized tracking error (b) Forecast tracking error

Figure 3: Tracking error of Simple and Optimized Energy exclusion strategies benchmarked to the MSCI ACWI Index. Realized tracking error estimates are based on three years of monthly data. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

A feature that distinguishes an Optimized Energy Exclusion from its Simple counterpart is unintended sector bets. By design, both strategies fully exclude the GICS Energy sector, leading to an active weight of -7.87% in each. In the Optimized Exclusion, tracking error is reduced by placing excess weight on sectors that are correlated with Energy, such as Materials and Utilities. This leads to substantial average overweights for these sectors as shown in Figure 4.

In the other direction, every Simple Exclusions carries a large-cap bias. To see why, suppose that the excluded securities account for a fraction γ of the index weight. Then the active weight w_a of an admissible security in a Simple Exclusion is given by

$$w_a = \frac{\gamma}{1-\gamma} w, \tag{1}$$

where w is the weight of the security in the index. In other words, securities with larger market capitalizations have larger active weights in a Simple Exclusion. The same holds for sub-industries, industries and sectors.

Figure 4 shows modest average overweights for the relatively small Materials and Utilities sectors in the Simple Exclusion, and more substantial average overweights in the relatively large Information Technology and Financials sectors. Relative to the Simple Exclusion, active weights of large admissible sectors are diminished in the Optimized Exclusion, leading to a lower tracking error. The performance statistics in Figure 5 indicate that the Optimized Exclusion had a higher Information Ratio than the the Simple Exclusion.

Active Weights	Energy	Materials	Utilities	Financials	Information Technology
Simple	-7.87%	0.56%	0.35%	1.81%	1.10%
Optimized	-7.87%	2.79%	2.54%	0.35%	-0.05%

Figure 4: Average active weights of GICS Energy, Materials and Utilities Sectors in Simple and Optimized Energy exclusions benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

Energy Exclusions	Index	Simple	Optimized
Return Measures			
Total Return (ann.)	7.72%	7.78%	7.90%
Active Return (ann.)		0.06%	0.19%
Risk Measures			
Volatility (ann.)	15.55%	15.56%	15.58%
Beta		1.00	1.00
Tracking Error (avg. forecast)		1.12%	0.70%
Risk-Adjusted Return			
Information Ratio		0.05	0.27

Figure 5: Performance of Simple and Optimized Energy exclusion strategies benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

5 Vice Exclusions

Current ESG investing practices have roots in the avoidance of companies that promote or profit from the exploitation of human vices, such as alcohol, tobacco, gaming, adult entertainment and firearms.[4] Vice exclusions have historically been standard practice for faith-based investors.[5] More recently, many secular investors have followed suit, and new categories of vice, such as animal cruelty, oil, national defense, private prisons, processed foods and sugar, have emerged.

As ESG investing has gained traction, some vice companies, such as Philip Morris International and British Petroleum, are attempting to re-brand themselves as ESG exemplars, and this raises the question of what should and should

[4]The seminal study on the impact of tobacco exclusions is Hong & Kacperczyk (2009).

[5]See, for example, the guidelines for socially responsible investing for the United States Conference of Catholic Bishops `https://www.usccb.org/about/financial-reporting/socially-responsible-investment-guidelines` and the screening used by the Evangelical Lutheran Church in America `https://www.elca.org/Resources/Corporate-Responsibility#Policy`).

not be excluded from a portfolio on ethical grounds.[6] Additional complications stem from the exposures of vice stocks to standard risk factors.[7] Still, given the popularity of Vice Exclusions, it is useful to look at their performance.

In what follows, we look at return and risk profiles of Simple and Optimized Vice Exclusions between August 1995 and December 2020, the same date range we used to study Energy Exclusions. We omit the GICS Chemicals and Tobacco industries as well as the GICS Brewers, Distillers and Vintners, Casinos and Gaming, Packaged Foods and Meats and Soft Drinks sub-industries from the MSCI ACWI Index. These industries and sub-industries belong to Consumer Staples, Consumer Discretionary and Materials sectors. A diagram showing the structure of the Consumer Staples sector and average active weights of its industries and sub-industries over the study period is in Figure 6.

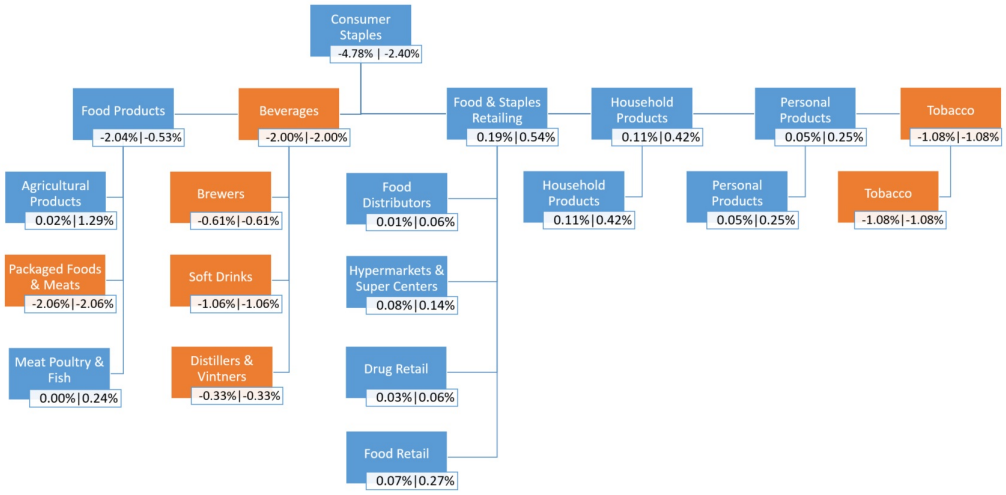

Figure 6: Industries and sub-industries of the Consumer Staples sector along with their average weights in the Simple Exclusion (left) and Optimized Exclusion (right). Excluded industries and sub-industries are marked in orange. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

Figures 7, 8 and 9 are the Vice Exclusion analogs of Figures 1, 2 and 3 for Energy Exclusions. Some qualitative conclusions in the two examples are similar. Optimized Exclusions hug their benchmarks more closely than Simple Exclusions, whose time series of ten-year active returns mirror the time series of exclusion weights. Tracking error is regime dependent, and it is materially higher for Simple Exclusions than for their Optimized counterparts.

[6] See, for example, Philip Morris International (2019) and Dow Jones Sustainability Index (2019).

[7] More information about the exposures of vice stocks to standard risk factors can be found in Blitz & Fabozzi (2017), Branch et al. (2019) and Blitz & Swinkels (2021).

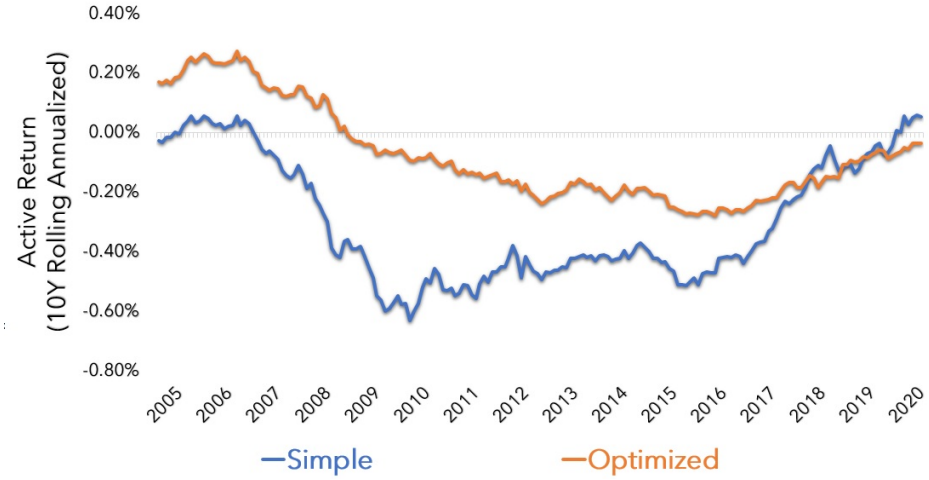

Figure 7: Rolling ten-year annualized active returns to Simple and Optimized Vice exclusion strategies benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

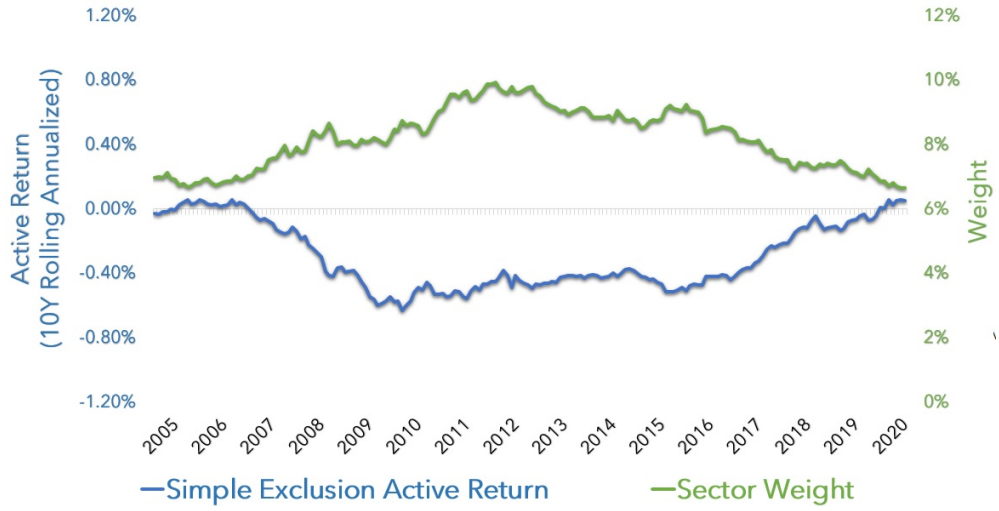

Figure 8: Weight of Vice stocks in the MSCI ACWI Index and ten-year rolling annualized active returns to a Simple Vice exclusion strategy benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

(a) Rolling 3-year realized tracking error (b) Forecast tracking error

Figure 9: Tracking error of Simple and Optimized Vice Exclusion strategies relative to the MSCI ACWI Index. Realized tracking error estimates are based on three years of monthly data. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

An important distinction between the Energy and Vice Exclusions can be seen by comparing Figures 4 and 10. In the former case, we excluded an entire sector, which led to relatively large overweights in correlated sectors for the Optimized Exclusion, but not the Simple Exclusion. In the latter case, exclusions were industries and sub-industries of the Consumer Staples, Consumer Discretionary and Materials sectors. To lower tracking error, admissible industries and sub-industries such as Agricultural Products, which is part of Consumer Staples, were overweighted in the Optimized Exclusion, as indicated in Figure 11. These overweights explain why the magnitude of the Consumer Staples underweight is greater in the Simple Exclusion than in the Optimized Exclusion.

Active Weights	Consumer Staples	Materials	Consumer Discretionary	Financials	Information Technology
Simple	-4.78%	-2.21%	0.53%	1.80%	1.46%
Optimized	-2.40%	-2.01%	0.58%	0.27%	0.31%

Figure 10: Average active sector weights of the Vice Exclusion strategies in Simple and Optimized exclusions benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

As for the Energy Exclusions, the large-cap bias for the relatively large Information Technology and Financials sectors is pronounced, as shown in Figure 10. A more nuanced aspect of the large-cap bias is shown in Figure 11 for Diversified Banks, the sub-industry with the largest average active weight between August 1995 and December 2020. In the Optimized Exclusion, active weights of large, non-excluded sectors, industries and sub-industries were diminished, relative to the Simple Exclusion, to minimize tracking error.

Consider our exclusion of the Casino & Gaming sub-industry of the Consumer Discretionary sector. The average weight of the sub-industry was 0.25% between August 1995 and December 2020. Still, the Consumer Discretionary sector had a positive average active weight of 0.53%, as shown in Figure 10, due to the large-cap bias, which averaged 0.78%.

Active Weights	Agricultural Products	Diversified Banks
Simple	0.02%	0.81%
Optimized	1.29%	0.13%

Figure 11: Average active sub-industry weights of the Vice Exclusion strategies in Simple and Optimized exclusions benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

The performance statistics in Figure 12 indicate that the Simple Exclusion underperformed on a risk-adjusted basis, while the Optimized Exclusion had the same return as the benchmark.

Vice Exclusions	Index	Simple	Optimized
Return Measures			
Total Return (ann.)	7.72%	7.59%	7.72%
Active Return (ann.)		-0.12%	0.00%
Risk Measures			
Volatility (ann.)	15.54%	15.87%	15.51%
Beta		1.02	1.00
Tracking Error (avg. forecast)		0.58%	0.33%
Risk-Adjusted Return			
Information Ratio		(0.21)	(0.00)

Figure 12: Performance of Simple and Optimized Vice Exclusion strategies benchmarked to the MSCI ACWI Index. August 1995–December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

6 Diversity and Inclusion

Emerging evidence suggests that US corporate boards are in the process of becoming more heterogeneous. To focus on a single dimension of this broad subject, Gorley et al. (2020) reports that representation of women on US corporate boards increased from an average of 11% in 2014 to to 19% in 2019.

Even so, we have not reached parity. This is shown in Figure 13, which displays the distribution of percentages of women on boards of companies in the MSCI ACWI Index as of December 2020.

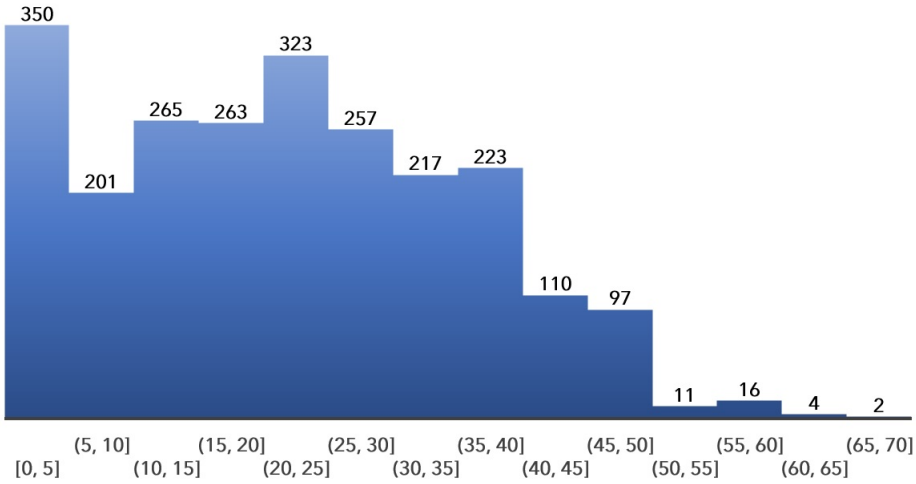

Figure 13: Distribution of percentages of seats on corporate boards in the MSCI ACWI Index occupied by women. December 2020. Source: Aperio Group and MSCI ESG Manager.

Taking the perspective of an investor who prefers companies with gender diverse boards of directors, we looked over calendar year 2020 at the performance of Simple and Optimized strategies that excluded any company whose board was less than 30% female from the MSCI ACWI Index. As of December 31, 2020, there were 1542 such companies, accounting for 52% of the index by count and 47% by market capitalization.[8]

Figure 14 shows the 2020 active returns to the Simple and Optimized Exclusion strategies (orange bars) along with the largest contributors to performance. The underperformance of these strategies, which was worse for the Simple strategy than the Optimized strategy, was driven by the exclusion of a few stocks, notably Apple and Tesla, which had excellent returns in 2020.

[8]There is look ahead bias in our gender exclusion strategies, as we used data on gender composition of boards as of the end of December 31, 2020 to run strategies through calendar year 2020. Given this data limitation, we chose not to run the strategy prior to 1 January 2020.

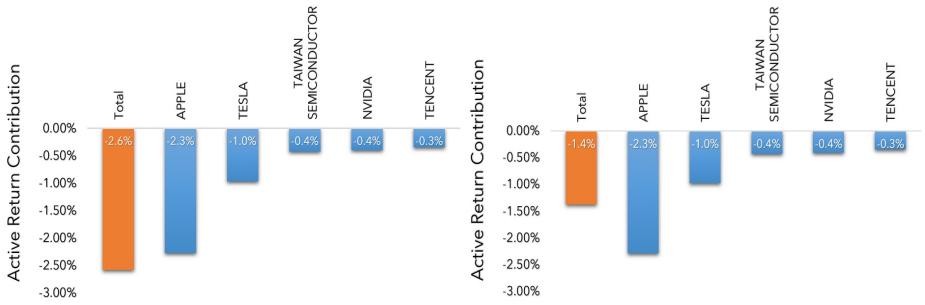

(a) Simple Exclusion (b) Optimized Exclusion

Figure 14: Active return of Gender Exclusion Strategies (orange bars) benchmarked to the MSCI ACWI Index and large security contributions to active return (blue bars). The strategies exclude companies whose boards have less than 30% women. December 31, 2019–December 31, 2020. Source: Aperio Group, MSCI ESG Manager and MSCI Barra Portfolio Manager.

Figure 15 provides a sector-based view of our Gender Exclusion strategies. The outperformance of the Information Technology sector in combination with its underweight in the Simple Exclusion strategy explains, to a great extent, the strategy's underperformance. Relative to the Simple Exclusion, active weights were muted in the Optimized Exclusion, as was the underperformance.

Active Weights	Information Technology	Industrials	Consumer Staples
Simple	-3.59%	-1.64%	2.23%
Optimized	-1.79%	0.13%	0.74%
Active Return	29.28%	-5.08%	-8.00%

Figure 15: Average active sector weights of the Gender Exclusion strategies as well as sector active returns. December 31, 2019–December 31, 2020. Source: Aperio Group, MSCI ESG Manager and MSCI Barra Portfolio Manager.

From these examples, we conclude that Gender Exclusion strategies may be driven by unintended bets, and that these bets may be muted by optimization. The performance statistics in Figure 16 summarize the risk-adjusted underperformance of both the Simple and Optimized Gender Exclusions. Any interpretation of these statistics, however, should take account of the short period over which they were collected.

Gender Exclusions	Index	Simple	Optimized
Return Measures			
Total Return	16.83%	14.25%	15.45%
Active Return		-2.58%	-1.38%
Risk Measures			
Volatility	26.02%	25.70%	25.24%
Beta		0.99	0.97
Tracking Error (avg. forecast)		2.02%	1.29%
Risk-Adjusted Return			
Information Ratio		(1.28)	(1.07)

Figure 16: Performance of Simple and Optimized Gender Exclusion strategies benchmarked to the MSCI ACWI Index. December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

7 ESG Scores

While ESG investing has roots in the exclusion of unwanted securities, the growing availability of more nuanced ESG data facilitates portfolio construction techniques that go beyond exclusions. ESG scores, which can be obtained from many different providers, are at the heart of many sustainability-focused portfolios. Different implementations of a single set of scores, however, can lead to different portfolios. We illustrate this by applying three standard construction methodologies to scores of firms in the MSCI ACWI Index. Our scores are obtained by averaging MSCI's numerical scores for the E, S and G pillars,[9] and standardizing the result to have mean 0 and standard deviation 1.

Our first two strategies are Simple and Optimized Exclusions, where we omit all but the top 30% of companies in our ESG ranking from the MSCI ACWI Index. Our third strategy is an Optimized Tilt, which minimizes tracking error while matching the average ESG score in the Optimized Exclusion. The Optimized Tilt includes all the securities in the MSCI ACWI Index.[10]

Figure 17a shows the cumulative active returns of the three strategies, while Figure 17b shows forecast tracking errors over calendar year 2020. Consistent with its construction principles, the Optimized Tilt hugged the benchmark more closely than the other two strategies.

[9] For more information about the pillar scores, please refer to MSCI (2020).

[10] In practical situations, portfolio managers may constrain the number of positions in a tilt.

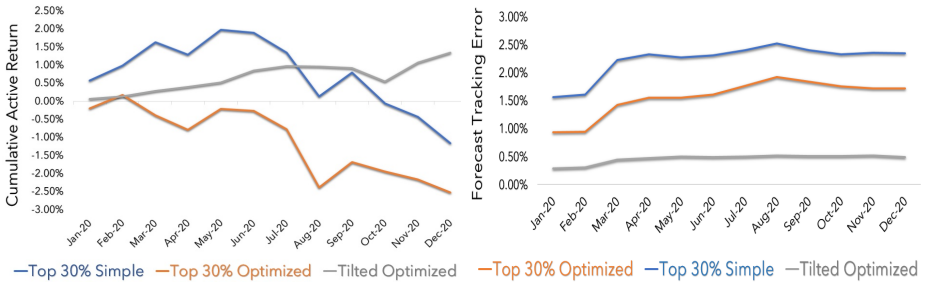

(a) Returns (b) Forecast tracking error

Figure 17: Active return and forecast tracking errors of ESG exclusions and tilt benchmarked to the MSCI ACWI Index. December 31, 2019–December 31, 2020. Source: Aperio Group, MSCI ESG Manager and MSCI Barra Portfolio Manager.

The performance statistics in Figure 18 suggest that the Optimized ESG Tilt materially outperformed the ESG Exclusions on a risk-adjusted basis. This conclusion should be tempered, however, by the short period over which the data were collected.

ESG Score	Index	Simple	Optimized	Tilted Optimized
Return Measures				
Total Return	16.83%	15.67%	14.31%	18.17%
Active Return		-1.17%	-2.52%	1.34%
Risk Measures				
Volatility	26.02%	24.55%	25.77%	26.24%
Beta		0.94	0.99	1.01
Tracking Error (avg. forecast)		2.22%	1.56%	0.46%
Risk-Adjusted Return				
Information Ratio		(0.53)	(1.62)	2.88

Figure 18: Performance of ESG exclusions and tilt strategies benchmarked to the MSCI ACWI Index. December 2020. Source: Aperio Group and MSCI Barra Portfolio Manager.

The three ESG strategies considered above were constructed by applying different portfolio construction methodologies to a single score. However, the sometimes substantial disagreement across ESG scores from different providers, documented in Berg et al. (2020) and Dimson et al. (2020*a*), adds another dimension of ambiguity to ESG investing. We explored this by looking at the distance between portfolios tilted on different ESG scores. First, we constructed a family of alternative scores from the MSCI E, S and G pillars by varying pillar weights between 0 and 100% in increments of 10%. Then we constructed tilted portfolios using these scores while minimizing tracking error and achieving an

ESG score equal to 1.3 times the score of the benchmark, the MSCI ACWI Index. Finally, we measured the forecast tracking error between each tilt and the tilt based on the equally-weighted ESG score. The histogram in Figure 19 shows that the tracking error between most alternative tilts and the equally-weighted tilt ranged between 0.50% and 1.10%.

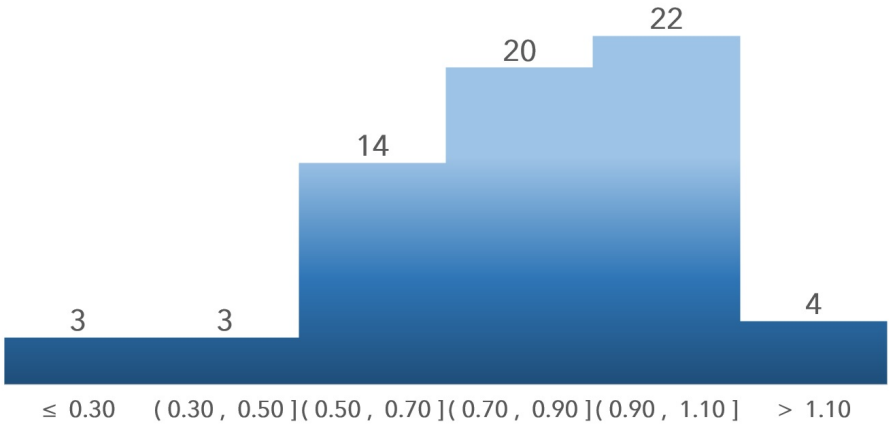

Figure 19: Histogram of tracking errors between alternative ESG tilts and an equally weighted tilt, all calibrated to 1.3 times the benchmark, the MSCI ACWI Index. December 2020. Source: Aperio Group and MSCI ESG Manager

8 Discussion

Since it aspires to achieve both financial and social goals with market-based tools, ESG investing is inherently complex. Without consistent metrics, portfolio construction with ESG objectives becomes even more complex as the investor may not achieve the objective he/she desires.

The experiments reported in this paper highlight three important points that should be more prominent in discussions of ESG-based portfolio construction.

- Even if the ESG objective is clear (e.g., avoid carbon-emitting firms or seek out exposure to companies that treat women well), practical strategy implementation is complicated by inconsistent ESG analytics and challenges arising from biases from exclusion strategies.

- Even when a portfolio strategy can be consistently implemented, risk-adjusted returns do not necessarily outperform more conventional strategies. Empirical evidence fails to support outperformance of ESG strategies.[11]

[11] See, for example, Boffo & Patalano (2020) for a recent discussion of the performance of ESG indices from different providers. Bolton & Kacperczyk (2021) find that stocks of firms with higher total carbon dioxide emissions (and changes in emissions) earn higher re-

- ESG portfolios may, in fact, reflect exposure primarily to other risk factors. That is, after controlling for standard risk factors, ESG risk may not be significant. This result does not mean ESG objectives should not be incorporated into portfolio construction, but rather investors should be clear about what risks they are taking.

These findings highlight the need for transparency and best practices in ESG investing.

Before the widespread development of ESG metrics, exclusion-based strategies reflected an approach to implement constraints that may not directly relate to maximizing risk-adjusted returns. Branch et al. (2019) describe a more sophisticated approach to incorporate ESG objectives into portfolio construction in their "A Guide to ESG Portfolio Construction". Our paper illustrates aspects of how transparent risk disentanglement can improve understanding of the actual risks in an ESG portfolio. Robust portfolio management requires understanding of actual risk exposures regardless of high-level portfolio objectives.

As demonstrated with the different strategies described in these analyses, one objective (e.g, exclude a cohort of energy or vice companies per a particular sector mapping) potentially leads to unintended bets such as going overweight in materials. While optimized strategies with a given ESG objective such as excluding an undesirable asset sub-portfolio may mitigate some of the knife-edge behavior arising from exclusion strategies, unintended bets may still appear as a particular strategy with a given investment opportunity set will eventually hit tracking-error lower bounds. Thus, another dimension of ESG-related investing complexity arises from the difficulty in addressing all (or even a large subset of ESG concerns) simultaneously. Moreover, the resulting strategy may ostensibly focus on one theme e.g., avoid vice stocks, only to become a different mix of seriously active bets.

In these contexts, optimization approaches constitute one way to minimize some of the challenges our results illustrate. Moreover, these examples demonstrate transparency in a way that can contribute to better sorting out when results arise from time-specific and/or sector-specific characteristics such as Big Tech (i.e., (F)acebook, (A)mazon, (A)pple, (M)icrosoft, (G)oogle– FAAMGs) driving returns. Once a key driver shifts (e.g, changes in thresholds, ESG metric construction, investment opportunity set, and geography) related to a particular portfolio's construction and performance evaluation, the outcomes can be highly variable. Our analysis shows how one can add more transparency to better understand the implications of ESG objectives, which could include unintended risk profiles and risk-adjusted underperformance.

Using ESG ratings (in this case from MSCI) is not as straightforward as many have initially assumed. As others have pointed out (see Berg et al. (2020)), ESG ratings are not consistent across vendors. We did not address this dimension of ESG-related complexity in this paper. Rather, we focused on the fact that using

turns, controlling for size, book-to-market, and other return predictors. Brandon et al. (2021) find a positive association between returns and disagreement in environmental ratings across providers.

one analytical approach can still generate quite different portfolio outcomes depending on how a strategy is implemented. Thus, a simple application of ESG ratings to build ESG portfolios is more complicated than it appears.

An important theme that threads through this analysis reflects the regular misalignment of objectives and actual risk exposure. That is, a particular ESG portfolio often reflects unintended risk bets and may not even align with the original ESG objective. Disentangling the risk exposures and understanding the underlying drivers of ESG portfolio outcomes contributes to more effective approaches to addressing ESG portfolio construction complexity.

9 Conclusion

Achieving ESG objectives in portfolio construction is a noble goal. These results illustrate why achieving this goal is not as easy as early adopters assumed. As more data become available and ESG analytics become more standardized and more clearly defined, we will have an easier time addressing the complexity of ESG investing. In the meantime, transparency in profiling the risk of a given ESG portfolio is essential to understanding which portfolios are actually outperforming on a risk-adjusted basis. Many ESG investors will find they are taking unintended risks and will often underperform. These outcomes should not deter proper ESG investing; but rather, focus the discussions on why particular objectives are targeted. Doing well (in terms of risk-adjusted performance) by doing good (in terms of targeting consistently defined ESG objectives) will happen sometimes, but not all the time. Setting aside financial performance, we conclude that we should do good whether or not we do well, and have transparency to ensure that we understand an ESG portfolio's risk profile.

References

Amel-Zadeh, A. & Serafeim, G. (2018), 'Why and how investors use ESG information: Evidence from a global survey', *Financial Analysts Journal* **74**(3), 87–103.

Atz, U., Liu, Z., Bruno, C. C. & Holt, T. V. (2021), Does sustainability generate better financial performance? review, meta-analysis, and propositions. working paper.

Bender, J., Bridges, T., He, C., Lester, A. & Sun, X. (2018), 'A blueprint for integrating ESG into equity portfolios', *Journal of Investment Management* **16**(1).

Berg, F., Kölbel, J. F. & Rigobon, R. (2020), Aggregate confusion: The divergence of ESG ratings. MIT Sloan Working Paper.

Blitz, D. & Fabozzi, F. J. (2017), 'Sin stocks revisited: Resolving the sin stock anomaly', *The Journal of Portfolio Management* **44**(1), 105–111.

Blitz, D. & Swinkels, L. (2021), Does excluding sin stocks cost performance?

Boffo, R. & Patalano, R. (2020), Esg investing: Practices, progress and challenges. OECD Paris.
URL: *https://www.oecd.org/finance/ESG-Investing-Practices-Progress-Challenges.pdf*

Bolton, P. & Kacperczyk, M. (2021), 'Do investors care about carbon risk?', *Journal of Financial Economics* .

Branch, M., Goldberg, L. R. & Hand, P. (2019), 'A guide to ESG portfolio construction', *The Journal of Portfolio Management* **45**(4), 61–66.

Brandon, R. G., Krueger, P. & Schmidt, P. S. (2021), ESG rating disagreement and stock returns. ECGI Working Paper Series in Finance, No. 651/2020.

Chia, C.-P., Goldberg, L. R., Owyong, D. T., Shepard, P. & Stoyoanov, T. (2009), 'Is there a green factor?', *The Journal of Portfolio Management* **35**(3), 34–40.

Cooper, L., Evnine, J., Finkelman, J., Huntington, K. & Lynch, D. (2016), 'Social finance and the postmodern portfolio: Theory and practice', *The Journal of Wealth Management* **18**(4), 9–21.

Dimson, E., Marsh, P. & Staunton, M. (2020*a*), 'Divergent ESG ratings', *The Journal of Portfolio Management* **47**(1), 75–87.

Dimson, E., Marsh, P. & Staunton, M. (2020*b*), 'Exclusionary screening', *The Journal of Impact and ESG Investing* **1**(1), 66–75.

Dow Jones Sustainability Index (2019), 'Industry Leader Report — British American Tobacco PLC'.
URL: *https://www.spglobal.com/esg/csa/static/docs/DJSI_IndustryLeader_TOB.pdf*

Friedman, M. (1970), 'The social responsibility of business is to increase its profits'.

Geddes, P., Goldberg, L. R., Tymoczko, R. & Branch, M. (2015), Building a carbon-free equity portfolio. Aperio Group Working Paper.

Giese, G., Lee, L.-E., Melas, D., Nagy, Z. & Nishikawa, L. (2019), 'Foundations of esg investing: How esg affects equity valuation, risk, and performance', *The Journal of Portfolio Management* **45**(5).

Goldberg, L. R. & Mouti, S. (2021), Sustainable investing and the cross-section of drawdown. CDAR working paper.

Gorley, T. A., Gupta, V. K., Matsa, D. A., Mortal, S. C. & Yang, L. (2020), The big three and board gender diversity: The effectiveness of shareholder voice. ECGI Finance Working Paper No 714/2020.

Henriksson, R., Livnat, J., Pfeifer, P. & Stumpp, M. (2019), 'Integrating ESG in portfolio construction', *The Journal of Portfolio Management* **45**(4), 67–81.

Hong, H. & Kacperczyk, M. (2009), 'The price of sin:the effects of social norms on markets', *Journal of Financial Economics* **93**, 15–36.

Khanna, V. & Warburton., A. J. (2021), Insuring against wrongdoing? socially responsible investing by mutual funds.

Madhavan, A., Sobczyk, A. & Ang, A. (2021), 'Toward ESG alpha: Analyzing esg exposures through a factor lens', *Financial Analysts Journal* **77**(1), 69–88.

Mahmoud, O. & Meyer, J. (2020), The anatomy of sustainability.

MSCI (2020), MSCI ESG ratings methodology.

Pastor, L., Stambaugh, R. F. & Taylor, L. A. (2019), Sustainable investing in equilibrium. NBER Working Paper 26549.

Pastor, L., Stambaugh, R. F. & Taylor, L. A. (2021), Dissecting green returns. Becker Friedman Institute Working Paper NO. 2021-70.

Philip Morris International (2019), Progress toward a world without cigarettes. Integrated Report.

Serafeim, G. (2020), 'Public sentiment and the price of corporate sustainability', *Financial Analysts Journal* **76**(2).
URL: *https://doi.org/10.1080/0015198X.2020.1723390*