ESG Investment Insights

Achieving multiple goals through optimization

November 2020

Bloomberg Professional Services

Quantitative Strategies

Bloomberg

ESG Investment Insights

Constructing ESG-tilted equity portfolios

Michael K. Donat, CFA
+1 212 617 5509
mdonat2@bloomberg.net

Steve Hou, PhD
+1 212 617 1357
fhou9@bloomberg.net

We would like to thank the Bloomberg Portfolio & Index Research team for their valued feedback.

By incorporating ESG into their portfolios, investors are increasingly faced with managing multiple objectives in the portfolio construction process. For example, an investor may be looking to gain exposure to securities that are aligned to societal goals while still closely tracking a benchmark. Or perhaps they are interested in constructing portfolios that are both attractive from a risk premia or alpha generating ability and aligned with ESG factors. Some questions naturally arise:

- How can they form reasonable expectations of the performance of their desired exposure?
- How might this investor incorporate their preferences into security selection?
- How can a portfolio achieve multiple goals spanning ESG and other strategic asset allocation objectives?

Using Bloomberg's ESG board governance scores as an example, we first review the traditional rules-based portfolio construction techniques and highlight issues that may arise in interpretation and implementation.

We then show how portfolio optimization and risk models can be used to form an ESG-tilted portfolio that meets multiple objectives and constraints, all while controlling undesirable incidental exposures.

We conclude with an optimization example of an investor who wants to gain exposure to board gender diversity while maintaining exposures to other board governance issues at the benchmark level and controlling active risk factor exposures. We show how the investor can efficiently trade off ESG issue exposure and portfolio active risk.

Finally we point out that previous papers in this series have covered similar ground for ESG-tilted fixed income portfolios. In particular, in *Analyzing the performance of ESG-based corporate bond portfolios*, optimization was used to maximize ESG scores while maintaining various risk characteristics of the benchmark. The work here is complementary in that similar conclusions are drawn for equities, where risk is measured more empirically.

Figure 1: Excess return for different portfolio construction methods (June 2016 – June 2020)

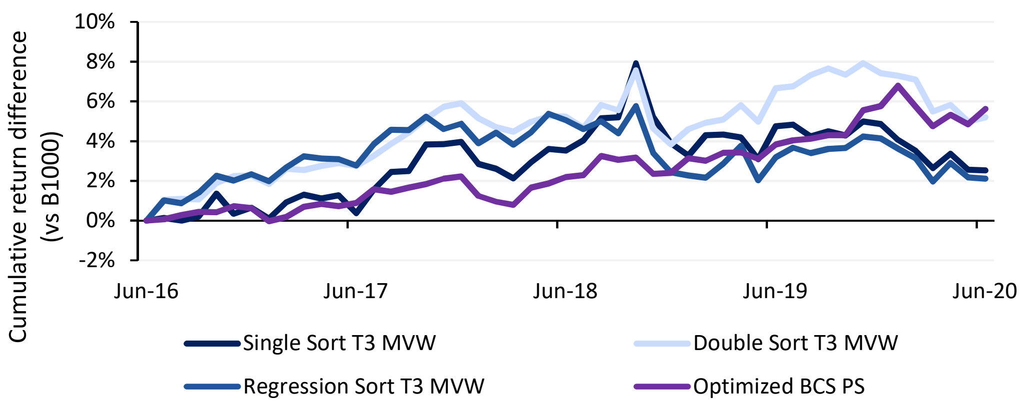

Source: Bloomberg

Introduction

ESG equity investment is of growing importance among asset owners and managers who oversee trillions of dollars. However, there is little consensus on how ESG issues should be systematically incorporated into equity portfolios. One reason is differing investment objectives. While some investors want to orient their investments and engage ESG issues for their own sake, others see ESG as a new class of risk premia. However, the short history of ESG data collection makes the risk premia assessment very difficult and highly uncertain. As a result, most ESG investors employ objectives and constraints beyond strictly seeking ESG exposure.

Historically, index providers constructed portfolios with simple and transparent heuristics such as weight tilting functions or quantile bucketing. In this paper we argue that numerical optimization may be a more natural method for constructing portfolios for ESG investors. An ESG-aware portfolio differs from traditional systematic equity portfolios in its need to satisfy multiple objectives and constraints: in addition to increasing exposure to a given ESG issue or score, the investor may also want to 1) track a benchmark, 2) control for unintended industry and risk factor exposures, 3) control exposures to other ESG issues and 4) manage risk level constraints, just to name a few. Formulating the problem as a constrained portfolio optimization allows us to manage and meet these demands systematically.

We first explore the issues with using the common rules-based portfolio construction methods and show attempts to overcome them. As an example, we construct a portfolio with high Bloomberg Board Composition Scores ("BCS") within the Bloomberg US 1000 Index ("B1000") universe. We show that the rules-based methods leave significant incidental industry and style risk factor biases. We further show that reasonable methods of mitigating these biases, such as double sorting or preprocessing ESG signals with a risk model, are difficult to implement and does not resolve these issues satisfactorily.

We then proceed to present a framework to construct an ESG portfolio using constrained optimization. Using the same board score setup, we use the Bloomberg PORT optimizer and MAC2 equity risk model to isolate the ESG exposures that are not a consequence of exposure to known equity risk factors. Moreover, depending on investors' preferences, exposures to industry and style factors can be managed and systematically incorporated in an optimized ESG portfolio if some exposure is desired.

In the last section of the paper, we make use of the data richness of BCS to design a portfolio focused on a single ESG sub-issue, Gender Diversity. We show how an investor can analyze possible portfolios that exhibit high gender diversity on the board of directors while maintaining benchmark-level exposures to other board governance sub-issues and equity factors.

Data

The primary ESG score used in this paper is the BCS. It is a composite measure of board composition quality within the Bloomberg ESG Governance pillar[1]. We choose this particular score because of its completeness and wide coverage. The BCS is computed as the generalized mean of four constituent *issues*: diversity, refreshment, director roles, and independence. Each of these issues is in turn a composite of several *sub-issues.* A

[1] For more information on Bloomberg ESG scoring please refer to BESG <Go>

hierarchy of the BCS data can be found in the Appendix.

The BCS are backfilled to 2015 and cover over 3000 public companies in approximately 60 countries. In Figure 2, we show the number of companies covered in the top 8 countries each year or about 75% of all companies covered. The scores are based on data published with lags, which can vary due to regional regulations. However, the US and European scores are predominantly available within six months of fiscal year end.

Figure 2: BCS global coverage by year

	2015	2016	2017	2018
All Regions	**3,376**	**3,535**	**3,650**	**3,665**
United States	1,289	1,325	1,358	1,366
Japan	314	342	347	348
Hong Kong	243	259	266	272
China	193	230	239	242
Canada	180	183	187	184
United Kingdom	157	157	146	137
South Korea	102	101	109	111
France	90	94	98	99
Other	808	844	900	906

Source: Bloomberg

In this publication, we focus on companies in the B1000, hence consider the US scores only. In order to have a credible backtest, we assume that the BCS are observable six months after the given calendar year. For example, we assume that the BCS for the year 2015 are available at the end of June 2016. As a result, from henceforth, we will refer to the scores as being available between 2016 and 2019. Additionally, we standardize the BCS each year to have a mean of 0 and standard deviation (SD) of 1 instead of using the scores reported on a 0-10 scale. (Pre-transformation means and SDs are shown in the Appendix).

In Figure 3, we show the snapshots of BCS coverage for B1000 constituents. Close to 90% of the B1000 members have a BCS, which represents about 95% of the total market capitalization of the index.

Figure 3: BCS coverage of the B1000

Year	B1000 Members	BCS Covered	Constituent Coverage	Market Cap Coverage
2016	989	854	86%	92%
2017	989	891	90%	94%
2018	993	927	93%	96%
2019	997	894	90%	96%

Coverage taken during June of each year
Source: Bloomberg

BCS coverage does not significantly vary across business sectors. In Figure 4, we plot the BCS market value coverage of B1000 members by BICS sectors. We can see that while the B1000 has uneven sectoral exposures, there is no noticeable bias in score availability across the sectors.

Figure 4: BCS US coverage of the B1000 by BICS sector 2019

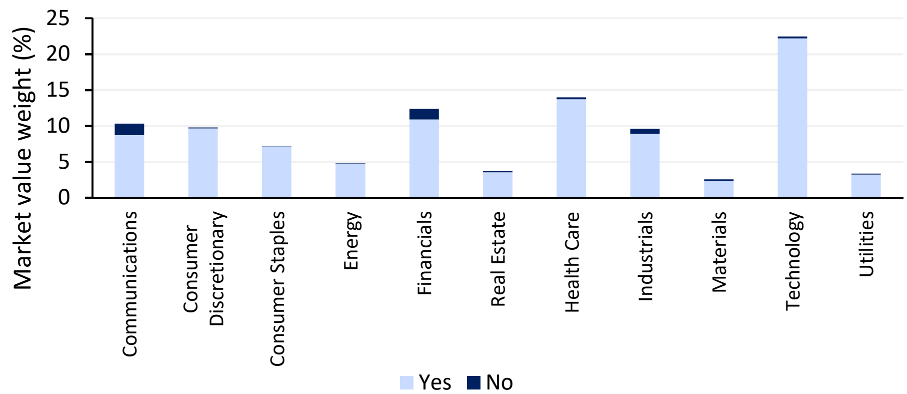

Source: Bloomberg

Finally, since we will investigate and compare the incidental factor exposures of different ESG portfolios, we also examine the correlations between the BCS and some well-known equity style factors. In Figure 5, we show the cross-sectional correlations of the BCS, the four underlying issues, and the loadings of ten equity style factors in the Bloomberg MAC2 risk model. These correlations are computed using 2019 data, but the patterns are very similar and stable over the years. Unsurprisingly, BCS have high and positive correlations with its inputs, the issue scores. Further, the four underlying issue scores have very low pairwise correlations, which suggests that these issues are relatively "orthogonal" to each other and they may each include additional information about the companies. Switching over to the style factors, we note that the pairwise correlations between the scores and the style factor exposures are modest at best. This suggests that the scores might contain information not already explained by the factors.

It is worth noting that there are some modest yet statistically significant correlations between the issues and the style factors. For example, there is a -14% correlation between growth and BCS, suggesting that on average growth companies have a lower BCS. And if we zoom in further, we see that a big contributor to this negative correlation is the -21% correlation between board independence and growth firms, where founder-CEOs might play a relatively more dominant role. We will observe later that these seemingly modest correlations can nevertheless end up having significant impact on the BCS-tilted portfolios.

Figure 5: BCS and issue level correlation* to equity style factors (2019)

	Board Composition	Diversity	Refreshment	Director Roles	Independence	Momentum	Value	Dividend Yield	Size	Trade Activity	Growth	Leverage	Profit	Volatility	Earnings Variability
Board Composition	**100%**	**58%**	**53%**	**40%**	**61%**	3%	2%	0%	**7%**	**11%**	**-14%**	**9%**	-3%	**-7%**	-6%
Diversity		**100%**	**13%**	0%	**25%**	6%	1%	4%	**15%**	**9%**	**-9%**	5%	5%	-4%	**-8%**
Refreshment			**100%**	**-11%**	**17%**	-5%	0%	-1%	0%	**10%**	1%	**12%**	**-12%**	3%	**12%**
Director Roles				**100%**	1%	**8%**	3%	-1%	**-7%**	0%	1%	-3%	2%	-2%	-4%
Independence					**100%**	-3%	0%	-6%	1%	**10%**	**-21%**	**9%**	-1%	**-11%**	**-7%**

*: The boldface correlation numbers are statistically significant at p value of 1% or less.
Source: Bloomberg

A typical rules-based framework

The traditional approach of evaluating a signal or risk factor is forming quantile-based portfolios. Specifically, securities are ranked by their signal values and grouped into quantiles. Portfolios expressing the signal or risk factor can be quickly constructed by selecting the appropriate quantiles. A long-only portfolio can be formed by taking the quantiles with the highest scores. For a market-neutral exposure to the signal, a long-short portfolio can be formed by buying the top quantiles and selling short bottom quantiles, thereby eliminating exposure to the broad market.

While we do not explore weight tilting methodologies here, quantile-based methodologies are qualitatively similar. Both have the same appealing features of *primary* rule simplicity and transparency. However, to implement portfolios based on either quantile construction or weight tilting, *secondary* rules are needed. These secondary rules, meant to address concerns such as security concentrations, turnover, and tracking error, often lead to ad-hoc heuristics and cumbersome (if not opaque) documentation.

Single sorting deciles

Here we place the securities of the B1000 into deciles on an annual basis (our scores are updated annually) based on their BCS. Decile 1 contains the lowest BCS ranked securities and decile 10 the highest. Securities that do not have a BCS are excluded from a decile. There are different ways of handling securities that do not have a score (such as assigning a default value, e.g., the average), but for simplicity we have chosen an exclusion rule. Given the uniformly high security coverage across sectors (Figure 4) the exclusion rule does not impact results significantly. Once securities are assigned a decile, they are weighted equally. Results are shown in Figure 6.

Figure 6: Performance of BCS deciles (June 2016 – June 2020)

Decile: Equal weight	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Annualized return (%)	13.3	11.8	11.7	12.7	13.4	13.5	13.9	10.6	11.4	15.1
Annualized volatility (%)	17.8	18.1	18.8	18.9	18.3	18.1	18.1	17.3	18.9	18.2
Return / volatility	**0.75**	**0.65**	**0.62**	**0.67**	**0.73**	**0.75**	**0.77**	**0.61**	**0.60**	**0.83**
BCS exposure	-1.61	-0.99	-0.61	-0.31	-0.05	0.20	0.45	0.73	1.04	1.56

Source: Bloomberg

From these results we can see that equally weighting securities within deciles based on their BCS does not show a clear pattern of returns, volatility, or risk-adjusted returns. At this initial stage, one would want to see a smooth or monotonic shape of returns which would suggest a degree of performance robustness from using BCS as a security selector. But at this stage, how realistic is it to see smooth or monotonic performance relationships? And even if there were smooth and monotonic relationships on display, how confident can we be that it is the result of the BCS, but not an artifact of weighting scheme or unintended styles or industry exposure? In the absence of a proper performance attribution framework it is difficult to tell. We will return to this important topic at the end of the section.

Single sort portfolios

With the decile portfolios constructed, one can easily examine performance of a long/short portfolio, typically long the top 3 deciles and short the bottom 3 deciles (T3-B3) as a first attempt to neutralize unintended exposures, or compare the performance of the top deciles versus the benchmark as a long only implementation. In Figure 7 we present results of the T3-B3 portfolio using equally weighted securities, T3 deciles using equally weighted ("EW") securities, and the T3 deciles using securities weighted by their market value ("MVW"). These three portfolios highlight that there are often significant changes in performance as one moves from more theoretical portfolio spaces such as T3-B3 EW, to liquid and tradeable spaces like T3 MVW.

Figure 7: Portfolios constructed from deciles (June 2016 – June 2020)

	B1000	T3-B3 EW*	T3 EW	T3 MVW
Annualized return (%)	12.5	-0.1	12.3	13.0
Annualized volatility (%)	15.3	3.0	17.7	15.4
Return / volatility	**0.82**	**-0.03**	**0.69**	**0.84**
BCS exposure	0.28	2.31	1.13	1.15

*This long short portfolio aims to strip away the market and cash rates of return
Source: Bloomberg

In Figure 7 the T3-B3 EW portfolio has an excess (above cash) of -.1% over this time period. Comparatively, the T3 EW and T3 MVW portfolios achieved excess returns (to the B1000) of -.2% and .5%, respectively. Additionally, due to the security weighting scheme and the absence of risk controls, significant levels of tracking error volatility to the B1000 may be introduced with the T3 EW and T3 MVW realizing 4.5% and 2.3%, respectively, over this time frame. Even though the T3 MVW portfolio delivered excess returns to the B1000, on the whole, these results are not great. But we still haven't *really* answered the question: what is driving the returns?

Performance attribution and single sorting issues

We alluded to the fact that forming portfolios of securities on a singular basis, in this case BCS, may lead to unintended exposures to known equity styles, as well as favoring certain sectors over others. By using the Bloomberg PORT Factor-Based Attribution model, we can identify the contribution to returns ("CTR") resulting from industry exposures (one level below sectors), equity style exposures, and security selection as shown in Figure 8. What we would like to see is that as the level of BCS increases, so too does the Selection CTR. To be clear, Selection CTR represents the returns unexplained by modelled risk factors, which our BCS is not one of.

Figure 8: Performance attribution of portfolios constructed from deciles (June 2016 – June 2020)

Decile	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Country CTR (%)	48.8	47.1	47.2	47.8	48.5	48.3	49.1	46.6	46.2	50.4
Industry CTR (%)	2.7	5.3	2.1	0.9	6.5	2.8	7.7	-0.8	-1.0	7.1
Style CTR (%)	-4.9	-5.9	-8.3	-0.2	-5.3	-9.6	-5.6	-12.0	-10.9	-12.3
Time CTR (%)	7.3	7.1	7.1	7.3	7.4	7.4	7.4	7.0	7.1	7.6
Selection CTR (%)	10.5	2.5	7.6	5.1	8.2	17.0	9.2	8.7	12.5	22.6
Cumulative CTR (%)	64.4	56.1	55.6	60.9	65.2	65.9	67.9	49.5	53.8	75.3

Source: Bloomberg

In Figures 9, 10, and 11, we have plotted Industry, Style, and Selection CTRs vs the BCS loading of each decile-based portfolio. As we increase the exposure to BCS, industry exposures (via their CTR) have no relationship to performance, style exposures contribute negatively, and selection is positive.

Figure 9: Industry

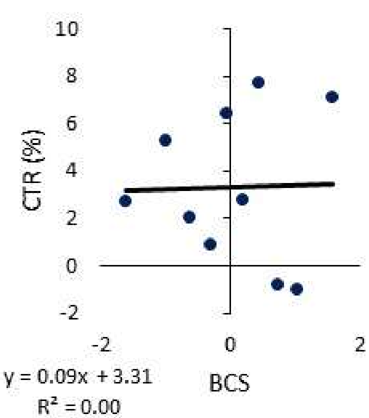

Source: Bloomberg

Figure 10: Style

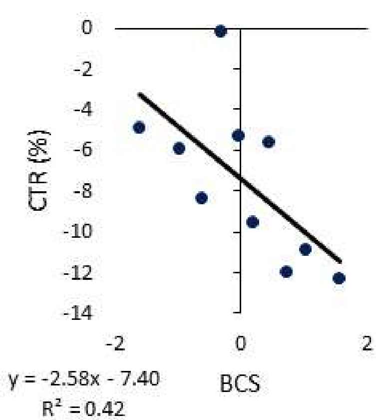

Source: Bloomberg

Figure 11: Selection

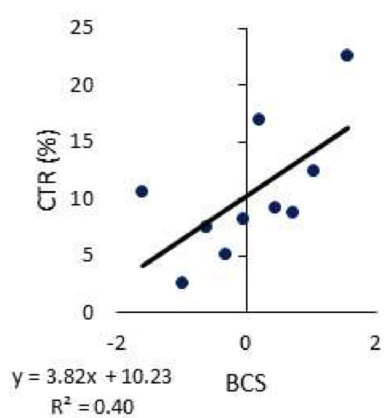

Source: Bloomberg

These three charts highlight the effects from the unintended exposures when we form portfolios from BCS in a traditional rules-based manner.

Double sorting deciles with sectors

In the previous section we have shown that a set of quantile-based portfolios created from a single-rule heuristic may have many unintended exposures to equity styles and industries. One way to try and control for these unintended exposures is to "double sort".

Double sorting refers to sorting and ranking along two dimensions. As an example, we will show results from a common double sorting approach which involves separating the investment universe securities into their sectors, and from within each sector to place the securities into quantiles by their score (just as we did in the prior section). The resulting sector-based quantiles can then be combined such that each sector has an equal weight for a given level of score, or a market value weight in line with the sectors of the benchmark (in our case the B1000).

Double sort results

The results of double sorting along sectors are shown in Figure 12. To keep the number of portfolio construction permutations tractable we present results of the following two portfolios:

1. A long/short T3/B3 double sorted portfolio (DS T3-B3 EW) of equally weighted securities within sectors and equally weighted sectors to assess BCS as an alpha generator.
2. A long only T3 double sorted portfolio (DS T3 MVW) of market value weighting securities along with market value weighting of sectors to assess loading on BCS in the context of managing against a benchmark.

Figure 12: Performance and attribution of double sort portfolios and the B1000 (June 2016 – June 2020)

	B1000	DS T3-B3 EW	DS T3 MVW	Difference to B1000 DS T3 MVW
Attribution				
Country CTR (%)	48.4	-0.2	49.2	0.8
Industry CTR (%)	-0.2	-1.3	-0.1	0.0
Style CTR (%)	-0.6	-2.9	-2.8	-2.2
Time CTR (%)	7.3	0.0	7.4	0.1
Selection CTR (%)	5.4	6.4	11.9	6.5
Cumulative CTR (%)	60.3	2.0	65.5	5.2
Performance				
Annualized return (%)	12.5	0.5	13.4	
Annualized volatility (%)	15.3	2.8	15.4	
TEV (%)			1.9	
Information ratio			0.47	
BCS exposure	0.28	2.25	1.10	0.82

Source: Bloomberg

The results indicate that DS T3-B3 EW portfolio was able to generate a return of 2% with 6.4% coming from selection. The return from unintended factor bets totaled -4.2% (-1.3% from industry exposure and -2.9% from style exposures) is nearly two-thirds the size of the selection component. Additionally, while the DS T3 MVW portfolio was able to mitigate returns attributable to the industry factors, style returns were nearly one-third the size of the selection effect. These two results suggest that while loading on BCS may have improved returns over the four year sample period, double sorting can only go so far in controlling unintended bets.

Double sort issues

While double sorting may in some cases help control exposures to the specific characteristic acting as the first filter (we used sector as a proxy for industry exposures), it is not an easy task to control for more than one characteristic at a time. For example, taking the B1000 and double sorting on sectors means there are only a handful of securities in any one sector-decile bucket. Therefore, one quickly runs out of securities as more sorts are added. Additionally, beyond the ability to market value weight securities within a bucket, it may be difficult to produce portfolios with acceptable levels of turnover and tracking error because these are not explicitly controlled for.

Regression-based sorting

Given the shortcomings of simple sorting techniques in reducing incidental factor exposures, in this section we explain how regression-based sorting may offer an incremental improvement.

The portfolio construction by regression-based sorting consists of two steps. First, we create an "orthogonalized BCS" by "regressing out" the style and industry factor loadings. Specifically, we regress the standardized BCS on Bloomberg's MAC2 risk model factors. The regression residuals are now our scores, and we will refer to them as Board Composite Score Residuals, or "BCSR". By construction the BCSR are orthogonal to the equity style and industry factors. Intuitively, using the residuals from the regression is equivalent to removing the components of the BCS that are already explained by existing risk factors. For example, if companies with high BCS tend to be large growth companies in the tech industry, the "size" and "growth" style factors and the "technology" industry factors would capture this. In the second step, we proceed as before by sorting companies by their BCSR and taking the top three deciles to be in the long only portfolio.

It turns out the BCSR are quite highly correlated with the original BCS as can be seen in Figure 13. This is not a surprising result given the low cross-sectional correlations of BCS and known equity style factors (Figure 5).

Figure 13: BCS residuals vs BCS

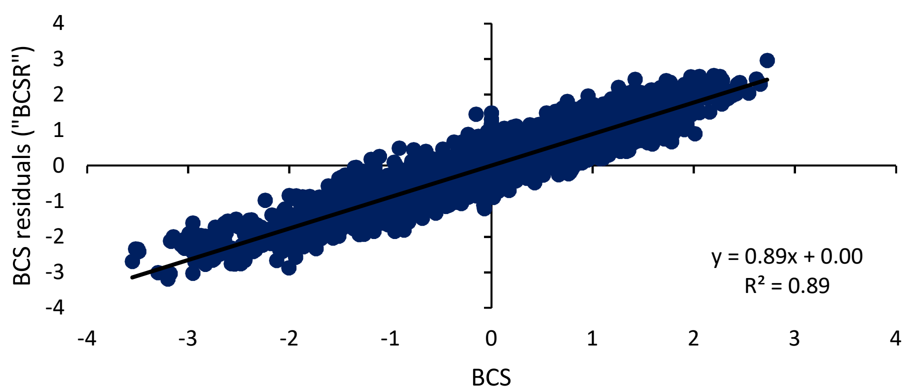

Note: BCS values are standardized
Source: Bloomberg

Orthogonal results

Using the BCSR and the single-sort portfolio construction methodology discussed above, we create three portfolios to assess the impact of regression-based sorting (Figure 14):

1. A long/short T3/B3 portfolio (BCSR T3-B3 EW) of equally weighted securities
2. A long only T3 portfolio (BCSR T3 EW) of equal weighted securities
3. A long only T3 portfolio (BCSR T3 MVW) of market value weighted securities

The BCSR T3 EW outperformed the B1000 11% during our 4 year sample period. This is significantly more than the unorthogonalized BCS portfolios both equal-weighted and value-weighted. This difference can be attributed to two sources: incidental factor exposures and tracking errors. Additionally, we observe that industries and styles have non-zero CTRs for the BCSR T3 EW and MVW portfolios compared to the B1000. This suggests that a regression-based approach may not work as effectively as desired and adds complexity.

Figure 14: Performance and attribution of BCSR portfolios and the B1000 (June 2016 – June 2020)

					Difference to B1000	
	B1000	BCSR LS EW	BCSR T3 EW	BCSR T3 MVW	BCSR T3 EW	BCSR T3 MVW
Attribution						
Country CTR (%)	48.4	-0.2	49.5	48.6	1.2	0.2
Industry CTR (%)	-0.2	-1.6	3.0	-4.8	3.1	-4.7
Style CTR (%)	-0.6	1.1	-7.2	-0.2	-6.5	0.4
Time CTR (%)	7.3	0.0	7.5	7.3	0.2	0.0
Selection CTR (%)	5.4	11.8	18.5	11.5	13.1	6.1
Cumulative CTR (%)	60.3	11.1	71.3	62.4	11.1	2.1
Performance						
Annualized return (%)	12.5	2.7	14.4	12.9		
Annualized volatility (%)	15.3	2.6	17.8	15.7		
TEV (%)			4.5	1.9		
Information ratio			0.42	0.19		
BCS exposure	0.28	2.13	1.04	1.09	0.76	0.81
BCS residual exposure	0.20	2.17	1.04	1.04	0.84	0.85

Source: Bloomberg

Regression sort issues

The comparison of outperformances highlights another problem with using a quantile-based approach: there's no easy way to control tracking error or risk. Any calibration of ex-ante tracking error using a finite sample can be incredibly imprecise. This is commonly the case with ESG data which tend to have relatively short sample periods. On the other hand, using a portfolio optimizer and equipped with a risk model, as we will demonstrate below, we can quite easily control for tracking error as well as factor exposures.

Optimization

Sorting and bucketing of securities are simple tools for evaluating a signal, but it needs to be in the context of a robust performance attribution framework which often shows return contribution from unintended sources. We will now demonstrate how portfolio optimization offers a direct and disciplined approach for assessing signal strength and controlling for factor exposures.

Assessing signal characteristics

To assess the performance and characteristics of a new signal relative to a benchmark, we can run two different optimization backtest analyses. First, we can create a portfolio that maximizes the new signal while taking no active exposures to style or industry factors relative to a benchmark. Second, to understand which factors the signal loads on, we can run an optimization that again maximizes exposure to our signal, but this time allows for active exposures to style and industry factors. The two optimization setups are summarized in Figure 15.

Figure 15: Optimization backtest framework

	Pure signal ("PS")	Relaxed signal ("RS")
Goal	**Maximize signal**	**Maximize signal**
Constraints		
Active Total Risk	1%	2.5%
Active Factor Risk	0%	None
Security Weights	$0\% \leq w_i \leq 100\%$, $100\% = \sum_{i=1}^{n} w_i$	$0\% \leq w_i \leq 100\%$, $100\% = \sum_{i=1}^{n} w_i$

Source: Bloomberg

The pure signal backtest

Using constrained portfolio optimization, one can simultaneously maximize exposure to the signal and precisely control the level and contributors of portfolio risk. In the context of a benchmark, this means honing in on the components of active risk. Specifically, we can control the level of active risk being targeted (1%), and how much of that active risk is taken by known factors (0%).

These two explicit risk-based constraints allow us to pick securities that maximize our intended signal exposure to a benchmark while only taking risks that are idiosyncratic in nature. In short, the excess returns in our backtest will be the result of security selection and not due to systematic factor exposures.

The unconstrained signal backtest

In addition to the pure signal backtest, we run a complementary backtest designed to uncover what known factors our new signal might simultaneously load on. Here the optimization goal is to maximize exposure to the desired signal but this time we relax our active total risk constraint to 2.5% and remove the constraint on active factor risk. This allows our backtest to select the securities that maximizes our signal exposure through both idiosyncratic risk *and* known risk factors. Portfolio attribution can then be used to see what factor exposures have been taken.

Performance results

In Figure 16 we present the performance and return attributions of the two backtests and the B1000 as the benchmark. As expected, the BCS PS portfolio has a much higher BCS exposure, an increment of .9 or almost a full standard deviation, than the B1000. At the same time, over the course of the four years, BCS PS outperformed the B1000 by 5.6% cumulatively while net exposures to known style and industry risk factors are kept close to zero. Most of the excess returns are contributed by the selection effect, or in other words, the tilting towards high BCS firms have properties unexplained by systematic risk factors. In the BCS RS backtest, the BCS exposure improvement (1.74) and the outperformance (14.7%) are even higher. Unsurprisingly the portfolio picked up significant active risk factor exposures: 5.6% industry and -7.0% style exposures relative to the B1000. At the same time, the contribution from the selection effect is also much higher.

Figure 16: Optimization backtest attribution and performance (June 2016 – June 2020)

				Difference to B1000	
	B1000	BCS PS	BCS RS	BCS PS	BCS RS
Attribution					
Country CTR (%)	48.4	49.2	50.7	0.8	2.3
Industry CTR (%)	-0.2	-0.1	5.5	0.0	5.6
Style CTR (%)	-0.6	-1.2	-7.6	-0.6	-7.0
Time CTR (%)	7.3	7.4	7.6	0.1	0.3
Selection CTR (%)	5.4	10.7	18.9	5.3	13.5
Cumulative CTR (%)	60.3	65.9	75.0	5.6	14.7
Performance					
Annualized return (%)	12.5	13.5	15.0		
Annualized volatility (%)	15.3	15.4	16.2		
Realized TEV (%)		1.1	3.1		
Information ratio		0.87	0.80		
BCS exposure	.28	1.19	1.74	.91	1.46

Source: Bloomberg

The BCS PS shows in Figure 17 a steady outperformance of about 5% over the course of 4 years resulting in an information ratio of .87. On the other hand, the BCS RS backtest shows a more volatile return profile (Figure 18); initially underperforming the benchmark before outperforming by 15%. Despite the higher returns of the BCS RS portfolio, the larger TEV budget led to a slightly lower IR than the BCS PS portfolio.

Figure 17: Cumulative return difference of BCS PS vs B1000 (June 2016 – June 2020)

Source: Bloomberg

Figure 18: Cumulative return difference of BCS RS vs B1000 (June 2016 – June 2020)

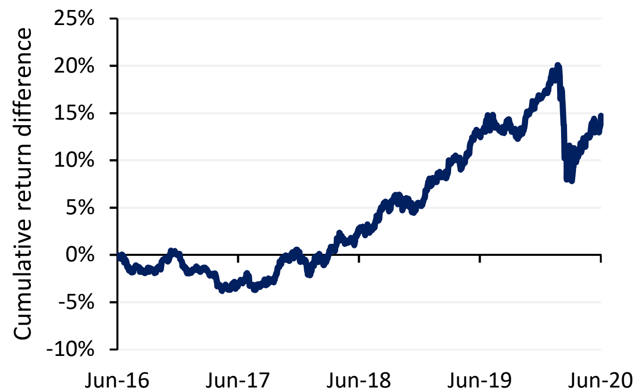

Source: Bloomberg

To assess the unintended equity style exposures that the BCS naturally loads on, we can plot the active exposures of BCS RS to the equity style factors in our attribution model (Figure 19). Here we observe that BCS RS loads on Profit, Value, and Trade Activity style factors, while loading against Growth, Volatility, and Size[2] during our sample period. It should perhaps come as no surprise that firms associated with quality and value tend to have better board governance.

Figure 19: Equity style loadings of BCS RS backtest

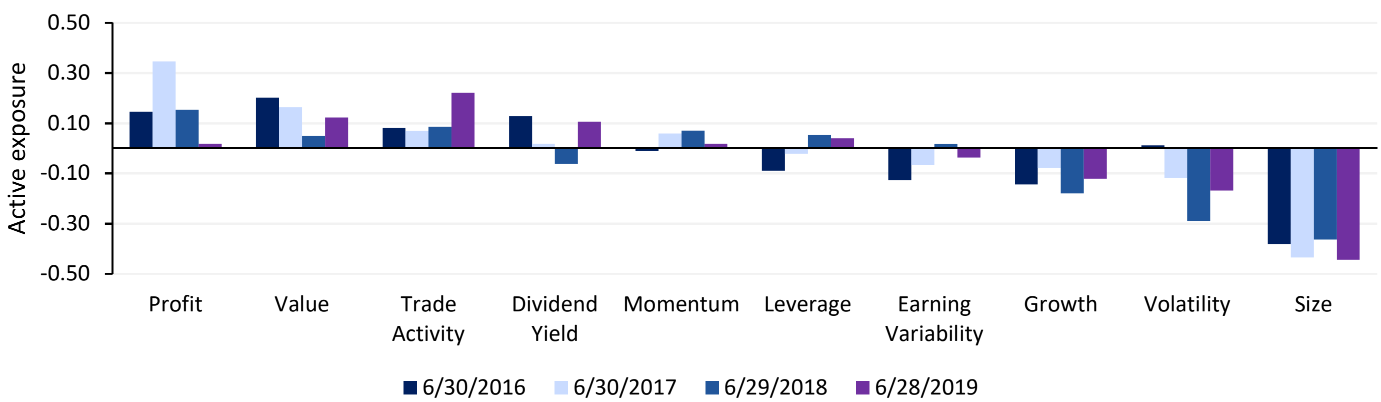

Source: Bloomberg

Sorting vs optimization

In the above sections, we have made a case that optimization produces portfolios that make a more efficient and explicit trade-off between BCS, factor exposure, and risk. Figure 20 displays the performances of the various long-only portfolios considered. While the decile-based portfolios require approximately 2% tracking error to achieve a

[2] Earlier in the paper we showed a weak positive correlation between the BC S and the size factor and here the BCS-maximized portfolio loads negatively on size. This is not a contradiction. The distribution of B1000 firm market caps is such that the vast majority have negative size exposures (MAC2 risk model definition). The correlation between BCS and size is very modest, but the level of size is on average (arithmetic) negative.

BCS exposure above 1, the optimized BCS PS portfolio achieves a 1.2 BCS exposure with a 1.1% realized tracking error and the highest information ratio of .9. Score loading per unit of risk is a point that we will illustrate even more vividly in the subsequent section through a gender diversity example.

Figure 20: Market value weighted sorting portfolios and optimization-based portfolio performance (June 2016 – June 2020)

	B1000	SS T3 MVW	DS T3 MVW	BCSR T3 MVW	BCS PS	BCS RS
Attribution						
Country CTR (%)	48.4	48.7	49.2	48.6	49.2	50.7
Industry CTR (%)	-0.2	-4.9	-0.1	-4.8	-0.1	5.5
Style CTR (%)	-0.6	-3.4	-2.8	-0.2	-1.2	-7.6
Time CTR (%)	7.3	7.3	7.4	7.3	7.4	7.6
Selection CTR (%)	5.4	15.2	11.9	11.5	10.7	18.9
Cumulative CTR (%)	60.3	62.8	65.5	62.4	65.9	75.0
Performance						
Annualized return (%)	12.5	13.0	13.4	12.9	13.5	15.0
Annualized volatility (%)	15.3	15.4	15.4	15.7	15.4	16.2
TEV (%)		2.3	1.9	1.9	1.1	3.1
Information ratio		0.2	0.5	0.2	0.9	0.8
BCS exposure	0.3	1.2	1.1	1.1	1.2	1.7

Source: Bloomberg

Finally, it bears mentioning that decile portfolios and the optimized portfolios are fairly similar to each other. In Figure 21, we show the monthly correlations of active returns of the various portfolios we consider. We can see that all the decile-based portfolios are highly correlated: around 80% and that the optimized portfolio, BCS PS, is moderately correlated with the decile-based portfolios. This suggests they are indeed capturing similar return dynamics. The BCS RS portfolio is least correlated with the sorting portfolios which is expected because the BCS RS portfolio was designed as a test to uncover simultaneous exposures to known factors.

Figure 21: Monthly Active Return Correlations (June 2016 – June 2020)

	SS T3 MVW	DS T3 MVW	BCSR T3 MVW	BCS PS	BCS RS
SS T3 MVW	**100%**				
DS T3 MVW	75%	**100%**			
BCSR T3 MVW	81%	78%	**100%**		
BCS PS	42%	55%	46%	**100%**	
BCS RS	25%	31%	18%	56%	**100%**

Source: Bloomberg

In Appendix 2 we provide a brief summary table highlighting considerations that the different portfolio construction techniques discussed thus far may be well suited to address.

Gender diversity example

Portfolio optimization need not stop at signal performance analysis. To truly highlight the power of optimization and its ability to elegantly express multiple portfolio constraints and objectives, we present an example with the objective of maximizing exposure to a *sub-issue* of the BCS, gender diversity, relative to the B1000. But as we have seen thus far, simply maximizing exposure to one dimension of a portfolio may lead to unintended—or even adverse—exposures along other dimensions.

With portfolio optimization, we can directly control for these unintended exposures through risk-based constraints on known factors and through constraints on other portfolio characteristics that we specify. Therefore, we can refine our portfolio objective to maximize exposure to gender diversity relative to the B1000 *while* maintaining equivalent levels of exposure to other components of the BCS *and* taking no active risk to industry or style exposures (active factor risk).

To keep this example relevant to investors, we run a set of four optimization backtests that increase annualized ex-ante active total risk from .5% to 2% and include a 15% quarterly turnover constraint.

The portfolio optimization setup is summarized in Figure 22:

Figure 22: Example of gender diversity optimization

Objective: Maximize BCS Gender Diversity							
Active Total Risk	**Active Factor Risk**	**Security Weights**	**BCS Director Roles**	**BCS Refreshment**	**BCS Independence**	**BCS Age Diversity**	**Turnover / Quarter**
Four backtest values	0%	$0\% \le w_i \le 100\%$, $100\% = \sum_{i=1}^{n} w_i$	Equal to benchmark	Equal to benchmark	Equal to benchmark	Equal to benchmark	15%

Source: Bloomberg

Results

Results from the four optimization backtest studies are shown in Figure 23. Here we observe that over this brief sample period, loading on gender diversity generates the same or slightly lower historical returns relative to the B1000. However there are meaningful gains to gender diversity score loading, although this benefit grows more slowly (Figure 24) for ex-ante TEVs larger than the initial .5% and 1% level (realized TEV .7% and 1.3%). In real terms, this amounted to an increase of approximately 6.7% and 9.3% respectively of female representation on the board of directors.

Figure 23: Backtest performance (June 2016 – June 2020)

		Ex-Ante Active Total Risk (%)			
	B1000	0.5	1.0	1.5	2.0
Annualized return (%)	12.5	12.5	12.3	11.6	11.1
Annualized volatility (%)	15.3	15.3	15.2	15.3	15.6
Realized TEV (%)		0.7	1.3	1.7	2.3
Information ratio		-0.04	-0.16	-0.50	-0.64
Selection effect (%)	5.4	5.4	4.9	2.4	-0.3
Gender Diversity loading	0.39	1.09	1.44	1.58	1.58
Raw Gender Score	4.24	5.61	6.18	6.44	6.57
Percentage women	25.6	32.3	34.9	36.0	36.5

Source: Bloomberg

Taken together, this case study gives two important implications for investors in the ESG space:

1) Investors may be able to gain significant levels of ESG related utility while taking minimal risk to a benchmark. Taking this further, investors may be able to replace strictly beta investments within their strategic asset allocation with "beta + ESG" considerations.

2) Investors who believe a given ESG factor may lead to alpha generating capabilities can carefully throttle their exposure through optimization without taking unintended bets.

Figure 24: Gender Diversity Score efficient frontier

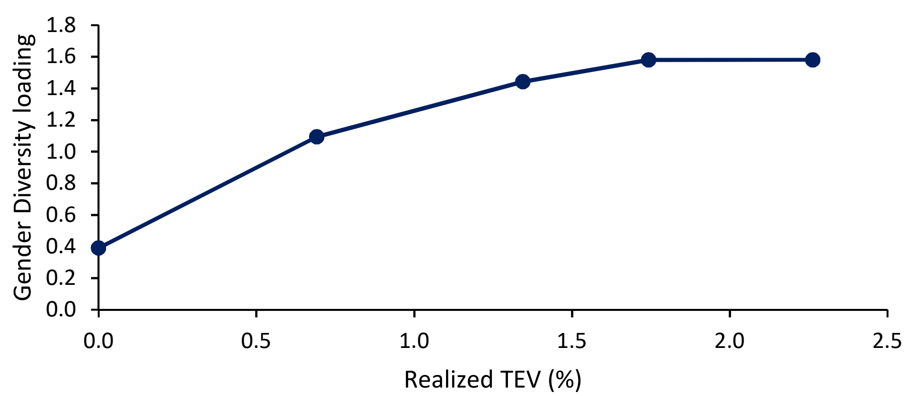

Source: Bloomberg

Conclusion

In this paper, we argue that portfolio optimization equipped with a robust risk model is a powerful and effective way to construct ESG equity portfolios. We show how traditional rules-based simple portfolio construction methods may fail to control for incidental risk exposures and tracking error, potentially leaving various investor objectives unmet. On the other hand, we demonstrate that portfolio optimization facilitates the clear specification of objectives and constraints investors may have. Furthermore, we argue that it provides a more risk-efficient way to gain exposure to ESG issues. Even if portfolio optimization is not ultimately used in an implementation, we believe it is extremely useful for providing comparisons to other portfolio construction methods and can give confidence to developing investment theses. Especially in the case of ESG integration, we believe that the benefits of incorporating risk and optimization are numerous, and come without much loss of transparency.

Appendix

Appendix 1: BCS hierarchy to the issue level

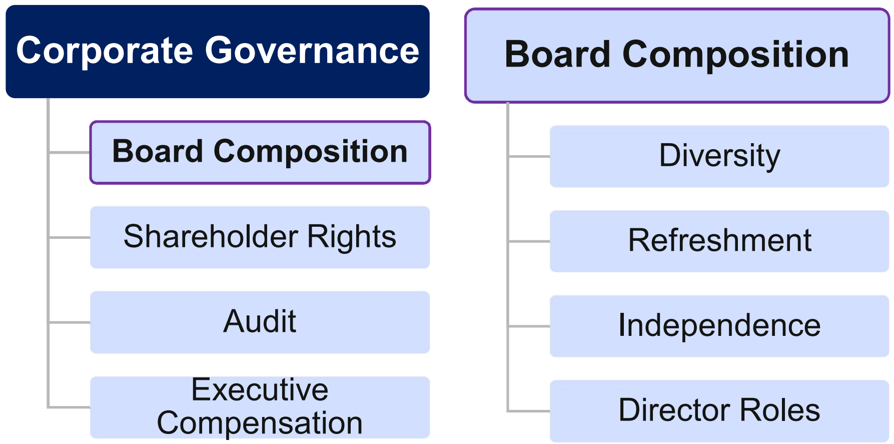

Source: Bloomberg

Appendix 2: Portfolio construction techniques

Control for…	Single Sort	Double Sort	Regression Sort	Optimization
Level of score	X	X	X	X
Unintended exposures		X	X	X
Risk model			X	X
Tracking error				X
Turnover				X

Source: Bloomberg

Appendix 3: BCS standardization adjustments

	Board Composition		Diversity		Gender Diversity		Refreshment		Director Roles		Independence	
Year	Mean	Stdev	Mean	Stdev	Mean	Stdev	Mean	Stdev	Mean	Stdev	Mean	Stdev
2015	5.84	1.25	3.62	1.54	2.54	1.96	6.12	3.16	7.66	2.15	7.64	1.91
2016	5.91	1.27	3.8	1.58	2.78	2.03	6.16	3.13	7.59	2.15	7.72	1.9
2017	6.01	1.28	4.01	1.61	3.06	2.08	6.19	3.09	7.5	2.2	7.83	1.87
2018	6.16	1.31	4.34	1.65	3.51	2.17	6.29	3.03	7.36	2.21	7.98	1.85

Source: Bloomberg

References

1. Cahan E., Ji L., (2016), "US Equity Fundamental Factor Model," Bloomberg LP.
2. Atamtürk A., O'Brien M., (2020), "Portfolio Optimizer (Version 3.25.1)," Bloomberg LP.
3. Gan Y., (2015), "Factor Based Performance Attribution (Version 1.3)," Bloomberg LP.
4. Mendez-Vives D., (2018), "Analyzing the performance of ESG-based corporate bond portfolios," Bloomberg LP.

Important Disclosures and Disclaimer:

Any systematic investment strategies described herein may involve a high degree of risk, including without limitation market risk and other risks inherent in investing in securities, commodities, currencies, derivatives and other financial instruments. The value of and income from investments linked to such strategies may decline in value and loss of the original amount invested can occur. All levels, prices and spreads are historical and do not represent current market levels, prices or spreads, some or all of which may have changed since the publication of this document.

Bloomberg does not represent that the index data, quantitative models, analytic tools and other information ("Content") referenced in this publication (including information obtained from third party sources) is accurate, complete or error free, and it should not be relied upon as such, nor does Bloomberg guarantee the timeliness, reliability, performance, continued availability, or currency of any Content. The Content is provided for informational purposes only and is made available "as is." Because of the possibility of human and mechanical errors as well as other factors, Bloomberg accepts no responsibility or liability for any errors or omissions in the Content (including but not limited to the calculation or performance of any index and/or the output of any quantitative model or analytic tool). Any data on past performance, modelling or back-testing contained in the Content is no indication as to future performance. No representation is made as to the reasonableness of the assumptions made within or the accuracy or completeness of any modelling or back-testing.

Bloomberg shall not be liable for any damages, including without limitation, any special, punitive, indirect, incidental or consequential damages, or any lost profits, arising from the use of or reliance on any Content, even if advised of the possibility of such damages.

Indices are unmanaged and cannot be invested in directly. The development or creation of any product that uses, is based on, or is developed in connection with any index (each a "Product") is prohibited without the prior written consent of Bloomberg. Bloomberg does not sponsor, endorse, sell or promote such Products and makes no representation regarding the advisability of investing in any such Product. Index returns represent past performance and are not indicative of any specific investment. The Content (including any of the output derived from any analytic tools or models) is not intended to predict actual results, which may differ substantially from those reflected.

Information and publications provided by Bloomberg shall not constitute, nor be construed as, investment advice or investment recommendations (i.e., recommendations as to whether or not to "buy", "sell", "hold", or to enter or not to enter into any other transaction involving any specific interest) or a recommendation as to an investment or other strategy. No aspect of the Bloomberg publications is based on the consideration of a customer's individual circumstances. Information provided in the publications should not be considered as information sufficient upon which to base an investment decision. You should determine on your own whether you agree with the conclusions made in the publications.

Take the next step.

For additional information, press the <HELP> key twice on the Bloomberg Terminal®.

Beijing
+86 10 6649 7500

Dubai
+971 4 364 1000

Frankfurt
+49 69 9204 1210

Hong Kong
+852 2977 6000

London
+44 20 7330 7500

Mumbai
+91 22 6120 3600

New York
+1 212 318 2000

San Francisco
+1 415 912 2960

São Paulo
+55 11 2395 9000

Singapore
+65 6212 1000

Sydney
+61 2 9777 8600

Tokyo
+81 3 4565 8900

bloomberg.com/professional