
CLIMATE-AWARE RISK BUDGETING

Brian Jacobsen^{a,b,*}, Eddie Cheng^{a,c} and Wai Lee^{a,d}

Climate change is a risk investors are thinking about, but how can it be practically incorporated into an asset allocation framework? This paper presents two different approaches. One is a traditional approach where the covariance matrix and excess return vector is adjusted to account for climate change. More detail is given for a second approach, a risk-budgeting approach. In this approach, investors adjust their risk budgets based on climate change information.



Asset allocators need to decide on an asset allocation strategy and then how to express that strategy through allocating to other managers, exchange-traded funds, or derivatives. From the asset allocation perspective (the top-down view), climate risk can be included at the asset class and country levels. From the implementation (bottom-up) perspective, climate-focused building blocks can be used.

We illustrate two different approaches to taking this top-down/bottom-up approach to integrating climate-awareness into portfolio construction. The first approach is a “traditional” approach where an investor adjusts the covariance matrix and excess return vector assumptions to account

for climate change. That was the focus of Jacobsen *et al.* (2021), so we will not spend much time on that approach.

A newer and less conventional approach, a risk-budgeting approach, is the main focus of this paper. Hopefully the presentation gives all the theoretical details of risk-budgeting so the reader can be familiar with the technique, but we also want to emphasize some of the practical dimensions associated with choosing and adjusting risk budgets in light of climate change information. Regardless of whether an investor uses a traditional or risk-budget approach, integrating the top-down and bottom-up perspectives can help pursue investors’ financial and sustainability goals.

1 Reframing Climate-Aware Asset Allocation

Climate change is a risk everyone should see coming. While there may be disagreements around the

^aAllspring Global Investments, 100 Heritage Reserve, Menomonee Falls, WI 53051, USA.

^bE-mail: Brian.Jacobsen@allspring-global.com

^cE-mail: Eddie.C.Cheng@allspring-global.com

^dE-mail: Wai.Lee@allspring-global.com

*Corresponding author

nature, timing, and magnitude of its effects, more and more investors are agreeing that climate risk is very real. As a result, it is a risk that needs to be managed. It can be managed from a top-down approach via industry and country selection. It can also be managed from a bottom-up approach via security selection. How to integrate these two approaches is a key function of a portfolio manager.

Asset allocation boils down to trying to optimize for an outcome, subject to different types of constraints. The conventional approach is to start with a risk model (a giant covariance matrix), a vector of excess returns, and then maximize the Sharpe ratio to select the weights to the various assets to get an optimal portfolio. Then, to tune a portfolio to a client's risk-preferences, a manager allocates between a risk-free asset and the optimal portfolio.

Climate change can alter the inputs, changing the elements of the covariance matrix and the vector of excess returns. Climate change may even alter a client's risk-preferences. Jacobsen *et al.* (2021) demonstrate how these changes can be done by boosting the risks of some assets, altering correlations, and modifying excess return assumptions. The condensed version of previous work on this topic is that climate change can bring about more volatility that may not be well compensated. As a result, assets with less climate risk exposure may have marginally lower volatility or marginally higher excess returns resulting in slightly higher Sharpe ratios compared to assets with more climate risk exposure. Regardless of whether the investor believes that climate risks may drive an asset's risks or returns up or down, with N assets, this requires $2N + \frac{N^2 - N}{2}$ modifications, which can be quite a large number.

An alternative approach is to take a risk-budgeting approach. Instead of modifying the covariance matrix—which can be daunting—a portfolio

manager instead can change the risk budget to the N assets and then make an assumption about how to modify the overall portfolio's Sharpe ratio. When modifying the Sharpe ratio, the investor can choose to adjust the excess return, or the volatility. This approach means making $N + 2$ adjustments. It is this approach that we favor and demonstrate, though the two approaches should be tried and be complementary in their results.

2 Risk Budgeting

Because it gets tedious to type *ex ante*, it should be understood that everything that follows—unless explicitly indicated otherwise—is on an *ex ante* basis and represents conditional information.

In Jacobsen and Lee (2020), we walked through the math of risk budgeting and showed a general set of conditions where a risk-parity allocation (i.e., the risk budgets are equal across assets) is also Sharpe optimal. Basically, its optimality requires a specific set of relationships among volatilities and correlations. Using the same notation, we assume that an investor needs to allocate among N risky assets and one risk-free asset. The covariance matrix of the N asset returns is an $N \times N$ matrix. The risk-free rate is a scalar, r_f , and the vector of the N assets' returns is an $N \times 1$ vector, μ . The market aggregated level of risk-aversion is a scalar, λ , which serves as a normalizing constant so that the resulting capital weights sum to one.

The Sharpe (1964) optimal weights (an $N \times 1$ vector w^*) are then given by the following when e is an $N \times 1$ vector of ones:

$$w^* = \frac{1}{\lambda} \Sigma^{-1} (\mu - er_f) \quad (1)$$

One approach to incorporating climate change into this framework, as in Jacobsen *et al.* (2021), is to adjust the elements of Σ and μ . A risk-budgeting approach is different. We start with the variance-covariance matrix, Σ , but instead

of needing a vector of returns we need an $N \times 1$ vector of risk targets, τ , which sums to one.

Risk budgeting then means deriving the capital weights that are consistent with the covariance matrix and the risk targets. The solution is found numerically by solving the following non-linear system of equations where W^* is w^* diagonalized and $w^{*'} is the transpose of w^* :$

$$\tau = \frac{W^* \Sigma w^*}{w^{*'} \Sigma w^*} \quad (2)$$

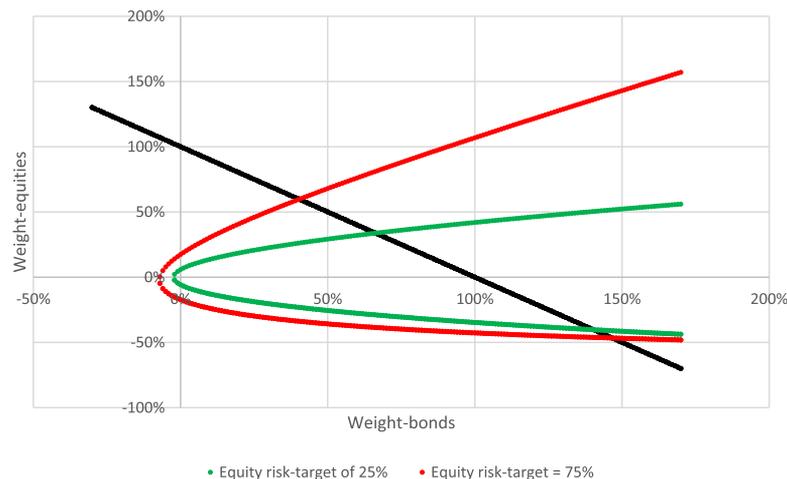
There are many challenges with this approach. The first is that there may not be a unique solution. Let's demonstrate with a two-asset case. If equities have a volatility of 18%, bonds have a volatility of 6%, and their correlation is 0.32, Exhibit 1 shows a straight line that gives the weights summing to one condition and the parabolas represent the weights consistent with two different risk targets. The possible solutions are where the "sum to one" line and parabolas intersect. One solution has both bonds and equities weights being positive, but there is another solution where the bonds weight is positive and the equities weight is negative. In general, with N assets, there are $N - 1$ free variables given the sum-to-one constraint in the system of N quadratic equations. There can

be many solutions. To choose one, the portfolio manager usually needs to impose a belief about the allowable weights, like they must be positive. This belief should be based on evidence.

What makes the more mathematically, and computationally, challenging risk-budgeting approach "better" than the conventional approach that involves forecasting returns? The typical argument is that there are many more stylized facts about risk than the returns. Risks tend to be bounded in a tighter band than returns. Also, realized volatility of equities is most often higher than that of bonds, so the likelihood of being wrong, at least in terms of the ordering of risk, is lower. Plus, the conventional approach requires forecasting risk anyways, so why compound the opportunities for error by also forecasting returns?

A reasonable counterargument is that there is nothing that guarantees the risk targets give capital weights that are optimal. Picking risk targets is like forecasting that the risk targets will give the Sharpe optimal weights. Instead of forecasting returns, the portfolio manager is forecasting risk targets. That too can be fraught with opportunities for errors, though there may be good *ex ante* bands to place on the risk targets, the bands on returns might be wider.

Exhibit 1 Non-uniqueness of capital weights consistent with given risk targets.



Under the risk-budgeting approach, integrating climate risks into the asset allocation process can now begin without the need to directly adjust forecasts of future returns. We can just modify the risk targets in light of information about climate change's likely effect on assets.

Because we are concerned with constructing a portfolio and necessarily in forecasting the returns, we turn next to how to determine these risk-budget modifications.

3 Distilling Down Climate Data to a Single Score

When it comes to managing climate risks, it is common to look for resilience. Which countries and which companies are better situated to manage and adapt to the risks? This question can be broken up into the top-down, asset allocation perspective and then into the bottom-up, security selection perspective. From a top-down perspective the question of resilience is leveled at the country and asset class. We will illustrate how to construct a country climate risk resilience score to inform adjustments to each country's risk budget.

How the risk adjustments are made depend on the investor's belief about what climate risk means for mispricing in the market. Increasing risk allocation to more climate resilient countries reflects a view that, everything else equal, we expect these countries to have relatively higher Sharpe ratios and/or be more diversifying than otherwise. Decreasing risk allocations reflects a belief that an investor thinks higher climate risk equates to lower Sharpe ratios or less diversification opportunities.

There are many governmental, non-governmental, and private sector indicators of climate risk. Though they may differ, they can have a high degree of concordance, or at least be complementary. Many follow guidance from the

Intergovernmental Panel on Climate Change in identifying vulnerabilities and risks from climate change. At the risk of oversimplifying these indicators, they identify the extent to which countries have or will likely suffer under extreme weather scenarios, have populations or industries at risk from the effects of climate change (e.g., coastal flooding, extreme weather affecting agriculture, and productivity), will incur costs in transitioning to a low-carbon future, and have governments or societies with both the willingness and ability to adapt to climate change or to adopt policies or technologies to stave off the effects of climate change. We will highlight some of the common variables considered in constructing climate risk indexes as indicative of the various variables and approaches used to make these assessments.

Geography and economic development can vary greatly across countries and have a material influence on a country's climate risk exposure. Luis and Alvarez (2021) find that geography is a key determinant of the welfare gains or losses from climate change. Mejia *et al.* (2018) summarize the main indicators of climate risks when they show that it is typically hot-poor countries that experience the most adverse effects from climate change. Some cold-rich countries stand to be more resilient, and even potentially benefit, from climate change.

Classifying countries according to the two dimensions of wealth and average temperature oversimplifies the problem. There other variables like the quality of institutions, access to potable water, the responsiveness of political institutions, and many other variables that can play into not only a country's current vulnerabilities, but also its future vulnerabilities. The University of Notre Dame's Global Adaptation Initiative (ND-GAIN at <https://gain.nd.edu/>) provides a very thorough, free, and open source set of data to help identify climate risks based on

exposure, sensitivity, and adaptive capacity. Germanwatch produces a Global Climate Risk Index to analyze “to what extent countries and regions have been affected by impacts of weather-related loss events” (see <https://germanwatch.org/en/cri>). Germanwatch includes a country’s experience with extreme weather events—the frequency of events and magnitudes of events in terms of death tolls and economic damage—to help identify a country’s vulnerability to climate change. Climate change is widely expected to bring about more frequent and more severe weather events.

Based on the preceding, and after considering how highly correlated many of these variables are, Exhibit 2 shows our selected variables and their source:

There are many other variables that could be included. How to weight these variables into a single index is a key problem. We use principal component analysis (PCA) to find the linear combination of variables that explains the greatest amount of common variation of the variables. PCA is often used because variation is information. To better interpret the results, it is advisable to adjust each variable such that they all “point” in the same direction. This means that each variable is recoded if necessary such that a lower value corresponds to low resilience and a higher value corresponds to high resilience. For example, gross domestic product (GDP) is already pointing in the “right” direction since wealthier countries

have more resources available to address climate change. Temperature, however, is pointing in the “wrong” direction as colder countries are commonly viewed as being more resilient to climate change than are hotter countries.

After the variables have been recoded, we rank countries from a cross-sectional perspective based on their trailing three-year median values. This can help us avoid some extreme numbers in some particular years and control for noise in the data. It also puts all the variables on the same scale, which is important for conducting PCA. Without standardizing the scales of the variables, the units of measurement (e.g., measured in dollars or cents) can affect the results.

We now have a matrix of countries and observations on the variables. For each year, we calculate the loadings for the first principal component after calculating a z -score for each variable. Those weights are then used to combine the variables for each country into a climate risk resilience score.

This process of creating a new score each year based on the available data allows us to add variables over time. For example, the climate policy variable did not become available until 2017. Our country climate risk resilience score can be updated as new metrics become available and as some are no longer published.

Let us give an example using two variables, which lends itself to nice visual representations. We

Exhibit 2 Sample set of variables to gauge a country’s climate risk.

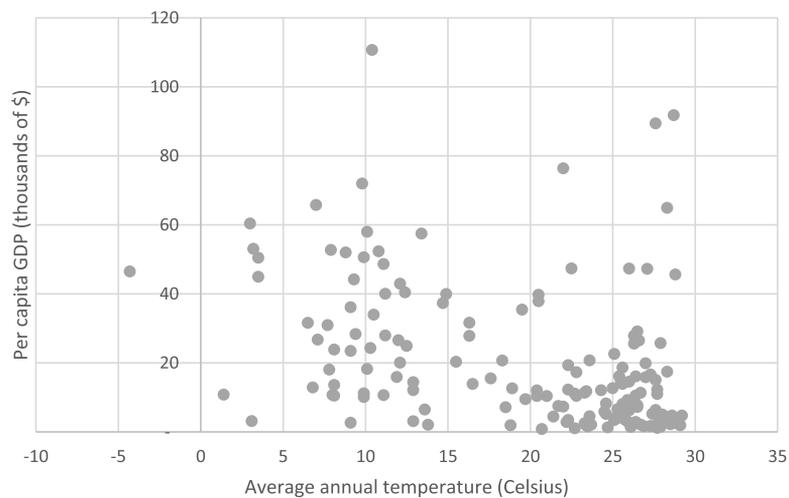
| Variable | Source |
|--|--|
| Gross domestic product (purchasing power parity adjusted) per capita | World Development Indicators from the World Bank |
| Government effectiveness | World Governance Indicators from the World Bank |
| Voice and accountability | World Governance Indicators from the World Bank |
| Average annual temperature | University of East Anglia’s Climatic Research Unit |
| Water | ND-GAIN |
| Exposure to climate change | ND-GAIN |
| Sensitivity to climate change | ND-GAIN |
| Climate policy | Climate Change Performance Index |

choose to illustrate with GDP and average annual temperature. The raw data for 2017 is plotted in Exhibit 3 where the value is the median from years 2015, 2016, and 2017 to help control for noise and outliers. This plot does illustrate the observation that poorer countries tend to be hotter countries.

The Appendix describes how we get the first principal component of this data after ranking and standardizing the data. The first principal component gives the direction of maximum variation in the data and is shown in Exhibit 4.

Now we need to translate these into a score. The case of two positively correlated standardized variables is easy, as PCA becomes unnecessary; we could have just averaged the standardized scores. If they were negatively correlated, we would have just had to negate one of them and then average them. This is unique to the two variable case, though. After averaging their standardized GDP and temperature scores, we simply rank the countries by this new score from PCA. Exhibit 5 shows the relationships between the original component ranks (the temperature and

Exhibit 3 2017 Data on 156 countries for per capita gross domestic product and average annual temperature.



Source: Trailing three-year median value for each country. World Bank for the per capita gross domestic product (\$) data on 156 countries and University of East Anglia for average annual temperature for those same countries.

Exhibit 4 Standardized and ranked data with first principal component.

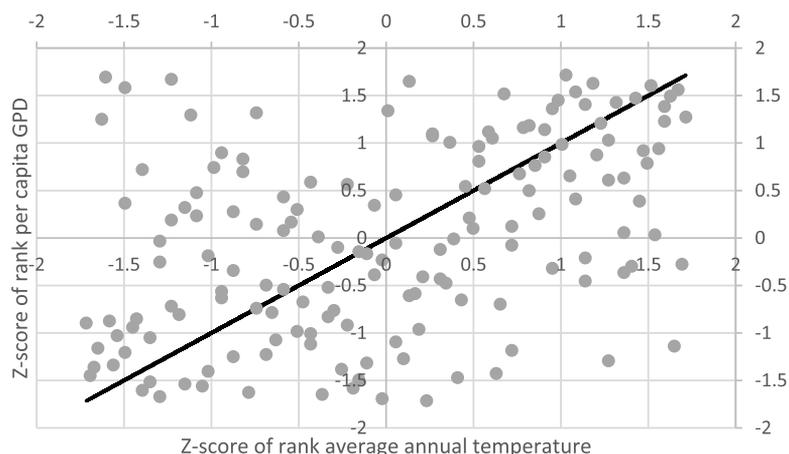
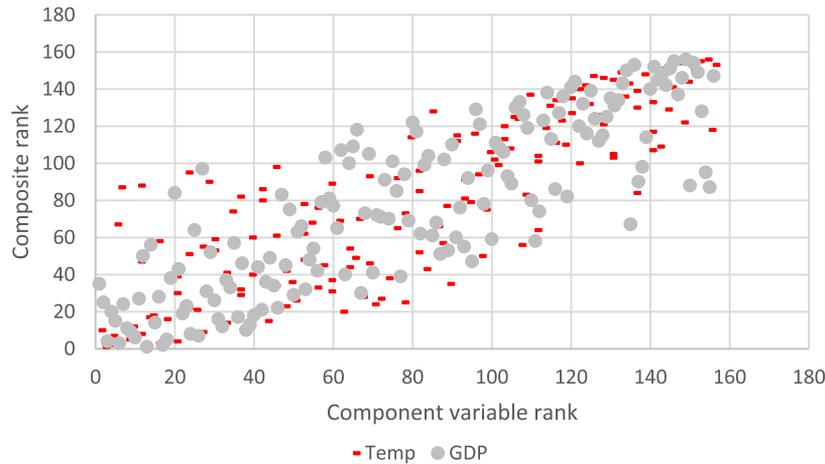
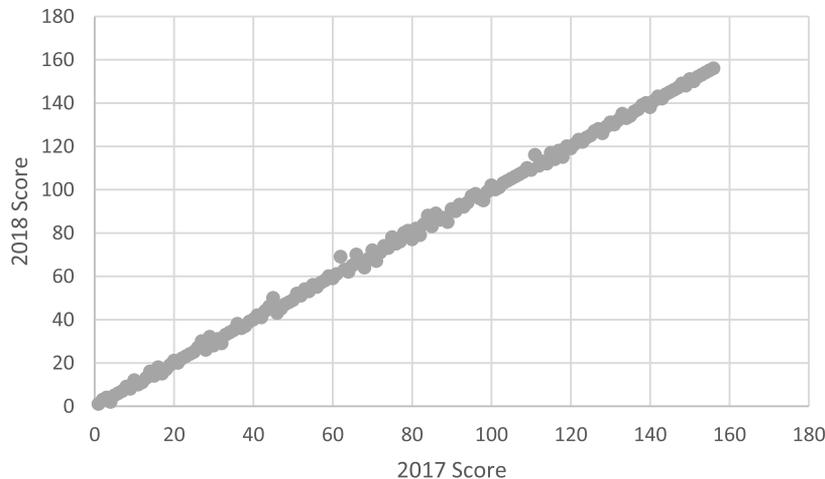


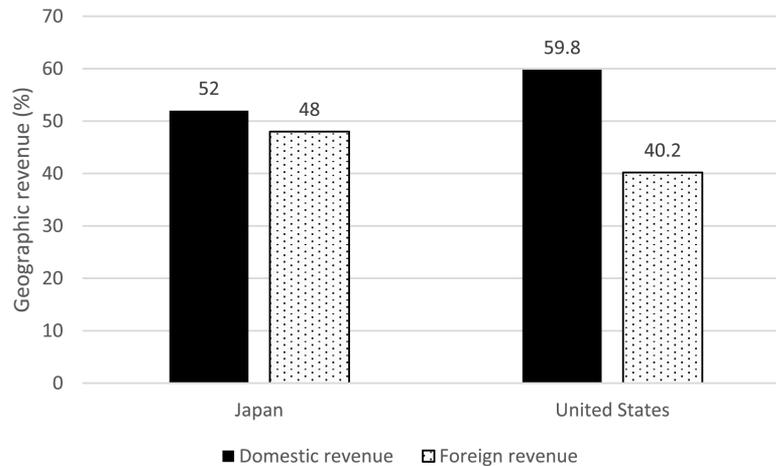
Exhibit 5 Composite score from PCA.**Exhibit 6** Comparing the score from one year to the next.

GDP ranks) and the composite rank from PCA.

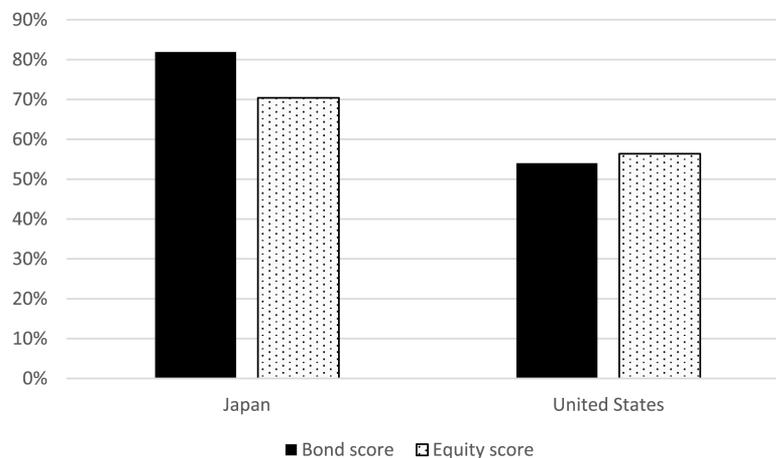
By construction, these composite ranks tend to be fairly stable over time, but they do change as new variables become available. In our two-variable illustration in Exhibit 6, there are some changes in country scores from one year to the next, but it is typically a slow migration. This is because average temperature and level of economic development are usually slowly evolving variables. When adding a climate policy variable, that can result in large changes in rankings.

4 Case Study: Connecting Country Climate Risk to Country Equity and Bond Risk

Countries are not securities. When linking climate risk to investible assets, we must assess how closely connected countries are to investible indexes (e.g., government bond indexes and equity indexes). For government bonds, there is a direct and obvious connection between the country and the investment, so we map the country risk directly to the government bond risk. For equities, the link is different. A company may be domiciled in a country, but have most of its effective economic exposure elsewhere.

Exhibit 7 Geographic revenue exposure for the MSCI Japan and MSCI USA indexes.

Source: FactSet as of 2020.

Exhibit 8 Exposure-based climate resilience scores converted to percentiles instead of ranks.

To get the climate risk of a country's main equity index, we would ideally use data on the operations of all of the equity index constituents. However, due to data availability, we instead use the geographic revenue exposure from FactSet for each index to weight the exposures to each country's climate risk. A country's equity index's climate risk resilience is then a weighted average of the equity index's revenue exposure to each country. This methodology can also be applied to regional equity indexes if the investor is interested in allocating to broad regions instead of specific countries.

Exhibit 7 shows the geographic revenue exposure for the MSCI Japan and the MSCI USA indexes. Exhibit 8 then shows the difference in climate risk resilience for each country's government bonds and its MSCI country index. The score is converted from a percentile to a rank as the country universe is smaller in this illustration with the calculation of the geographic revenue exposure. It must be noted that these scores are only based on the two variables: per capita GDP and average annual temperature. An actual climate resilience score should include many other variables. Japan is more prone to coastal flooding and tsunamis

than the U.S., so this illustration may overstate Japan's resilience. The illustration does show how the bond score of a country often does differ from the equity score due to the effective economic exposure of each type of asset.

Because the explanatory variables are only available with a lag, we need to lag our climate risk indicator by three years to use for current year allocation decisions. Within a risk-budgeting framework, the scores can be used to lower the risk budget to those assets with lower climate risk resilience scores and to increase the risk budget to those with higher climate risk resilience scores. This allows investors to have climate risk resilience reflected in their allocations. If an investor believes that climate risk resilience is negatively associated with prospective returns or positively related to future volatility, then the investor may choose to increase their risk budgets to countries with lower resilience. We believe that the extant evidence favors increasing risk weights with climate risk resilience, but investors do need to regularly review the evidence and ask whether prices have adjusted such that risk weights and climate risk resilience should be negatively related. There is another reason beyond risks and returns to condition risk weights on climate risk resilience scores. Investors may have a strong preference or opinion on the topic, which should be respected.

Based on our composite climate risk resilience score for each country, we use a linear multiplier to over-weight or under-weight risk allocations. A linear rule is one of many possible rules. The specific calculations should reflect the portfolio manager's confidence in the climate risk resilience indicator and other considerations like allowable tracking error to a specified benchmark. One possible adjustment is to calculate each country's asset climate risk resilience score's deviation from the average score and then multiply that deviation by a chosen multiplier. That becomes an

additive (or subtractive) adjustment to the original risk target. Care should be taken to choose the multiplier such that the risk targets remain within allowable bounds. For example, if the original risk target was 1% and that is the minimum allowable risk target, that country's asset exposure may be exempted from adjustment.

5 Bottom-Up: Build Core Portfolio—Connecting Bottom-Up Implementation

While a top-down asset allocation approach incorporating climate risk helps position a portfolio to be resilient to the risk of climate change, informed decisions made by fundamental portfolio managers and analysts are important inputs into the portfolio construction process. A climate aware strategic allocation—whether from a conventional approach or a risk-based approach—can be augmented by carefully selected climate-focused or environmental, social, and governance (ESG)-tilted building blocks to meet investors' specific financial and climate/sustainability goals.

Integrating climate and ESG factors into a strategic allocation is not as simple as just plugging new building blocks into a strategic allocation. The risk–return characteristics of climate- or sustainability-driven building blocks are likely to be quantitatively and qualitatively different from traditional building blocks that do not reflect such considerations. Madhavan *et al.* (2020) find that high environmental scores overlap with quality metrics, strong momentum, and tend to be larger firms. Bruno *et al.* (2021) also find overlap between ESG ratings and quality measures. Interestingly, Pastor *et al.* (2021) find that environmental scores tend to be pro-momentum and anti-value metrics.

We have shown in Jacobsen *et al.* (2019) that ESG-tilted building blocks tend to have less

residual volatility than traditional building blocks after controlling for size, value, and momentum. ESG-tilted building blocks have different factor exposures, so it is important to not simply swap in ESG-tilted building blocks for more conventional building blocks without considering these important differences.

These differences lead to different inputs into the portfolio optimization problem, resulting in different allocations. How are these bottom-up decisions reconciled with the top-down decision? There are three ways to consider. One is a holdings-based approach and two are based on risk models.

Because the strategic allocation gives resulting country and asset class capital weights, managers or securities can be chosen to give the same resulting weights. This can be extended to considering the factor exposures of the building blocks. The problem with this holdings-based approach is that if a portfolio manager is only concerned about the country and asset class exposure, there are many ways to get the same exposure, but with very different risks. For example, if the top-down approach allocates 35% of a portfolio to U.S. equities, there are many ways to get that exposure. Supplementing it with factor exposures can help align the bottom-up implementation with the top-down expectations.

To align the implementation with the expectation, two different risk model approaches can be useful. One is to treat the top-down country/asset class indexes as explanatory variables. Regressing each manager's returns or each security's returns on the explanatory variables give betas (the slope coefficient from ordinary least squares) to the top-down categories. The managers or securities can then be chosen to match weighted average exposures to the exposures from the top-down analysis. This approach is also amenable to considering alpha.

Another returns-based approach is to treat the benchmark as the single explanatory variable and to select managers such that the beta to the benchmark is statistically indistinguishable from one. This approach is computationally intensive, but can result in more stable exposure to the investor's benchmark and increase confidence in generating alpha relative to the benchmark. Jacobsen and Ma (2020) detail an approach to building a portfolio with this approach. A simple example is to consider an investor with a benchmark that is 60% equities and 40% bonds. The asset allocator can search (computationally) for the combination of managers from a roster whose combined beta to the 60/40 benchmark is one while maximizing the t -value of the intercept term. This gives a beta similar to the benchmark and the maximized t -value is equivalent to maximizing the confidence that the alpha will be positive (or the least negative).

6 Conclusion

To sufficiently meet investors' needs and achieve the outcomes they want, portfolios must consider both financial and climate risks. There are two approaches to incorporating climate risks into the top-down construction of portfolios. One is the conventional approach of modifying risk and return assumptions to account for climate change. Another approach is to incorporate climate risk from a risk-budgeting perspective. Instead of modifying risks and returns, a portfolio manager can take a conventional risk model (a covariance matrix) and adjust the risk targets of assets where the adjustments are informed by climate considerations. This is the approach we detail in this paper by walking through the risk-budgeting framework and then how to use principal component analysis to aggregate climate information into risk-budget adjustments.

In addition to the top-down approach to building a portfolio, we explain the three different ways to integrate climate and other considerations into the

bottom-up portfolio implementation. One is the holdings-based approach to match the features of the top-down allocation. Another is the returns-based approach where each top-down asset or index is treated as an explanatory variable in explaining manager or security returns. The final approach is to use the top-down allocation as a benchmark, treating it as the sole explanatory variable. The merit of this latter approach is that there is more flexibility in incorporating a wide range of assets in the portfolio implementation and it has a statistically more reliable relationship with the actual benchmark that matters to the investor.

Appendix: Two-Variable PCA

In this Appendix we walk through some of the technical details of converting the two variables (GDP and temperature) into a single score based on PCA. First, we convert these raw values into ranks in Exhibit A.1. This is for convenience and to control for outliers since PCA can handle variables of different characteristics. However, PCA is sensitive to the scales of variables and will assign more importance to the variable with the greatest variation even if the variation only stems from the scale of measurement and is not really all that meaningful. The per capita GDP data in

rank form loses information about the distance between countries, but adding variables to capture that information could be useful.

Just for convenience, we rank the variables in the order that should be related to climate risk resilience. Richer countries have more resources so a higher per capita GDP should be associated with more resilience. Hotter countries are likely more vulnerable to climate change, we rank the variable in descending order where Canada (the coldest) has the highest ranked value.

The variables are next standardized by calculating their z -scores (number of standard deviations from their mean). PCA is then applied to the standardized variables. Technically, the variables did not need to be z -scored, but it does aid in the interpretation of the results. PCA changes the axes of the data to orient it in a way that the major axis (the first principal component) points in the direction of maximum variation of the data. It also gives us a way to combine the variables into one score. For a brief overview of PCA, especially in this two-variable case, please see the Appendix for details.

Though it is not necessary now that all the variables are on the same scale, we standardize the variables by calculating their z -scores. This aids

Exhibit A.1 Ranked per capita GDP (ascending) and average annual temperature (descending order).

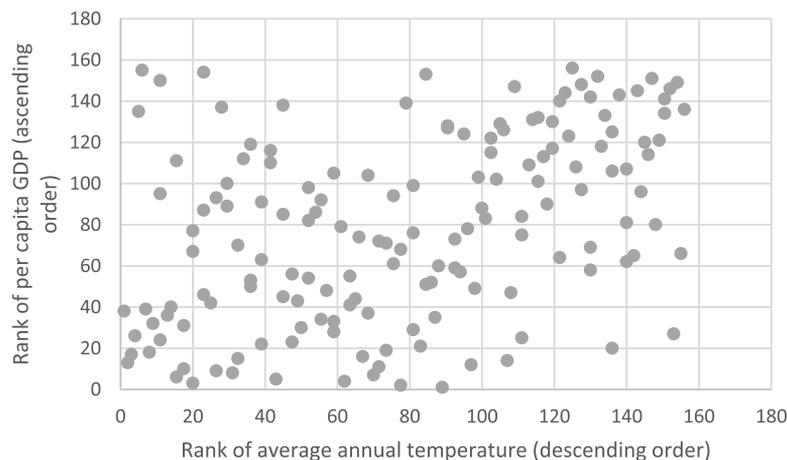
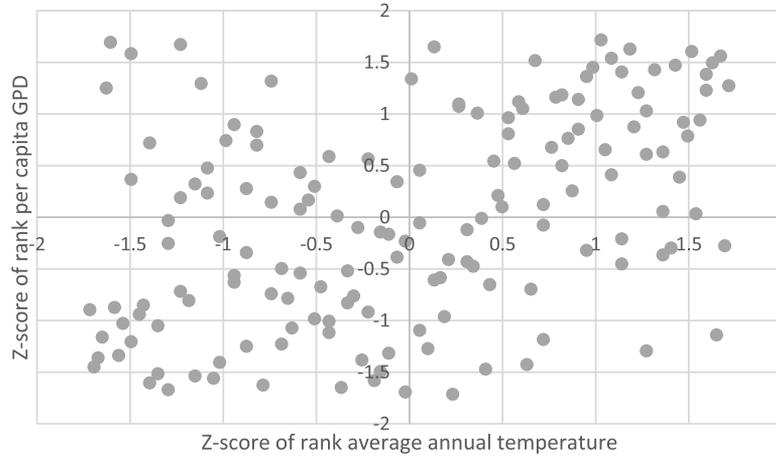


Exhibit A.2 z-Score of ranked per capita GDP and ranked average annual temperature.

in the interpretation of the output from PCA. To calculate the z -score for the rank average annual temperature, we calculate the average rank and the standard deviation of the ranks. If each observation was unique, this average and standard deviation would be the same for all variables because they are all ranks and we have observations on all countries. Due to rounding of values, some ranks are the same for different countries and the average and standard deviation of each variable may differ. Each country's rank value is replaced with how many standard deviations away it is from the mean. The same process is used for the per capita GDP rank variable. Visually, you can probably see that the direction of maximum variation goes from the lower left of Exhibit A.2 to the upper right. PCA finds the exact vector that confirms that eye-conometric result.

Because we are working with two standardized variables, it is very simple to calculate the principal components. There are two principal components and for two variables that have non-zero correlation, they will always be $(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ and $(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$. If the correlation is zero, then the axes do not change and the first principal component is in the direction of the variable with the greatest variation.

Generally, to find the principal components we need to find the eigen values and then the eigen vectors. Alternatively, we could perform singular value decomposition, but the distinction here is not important. The eigen values come from solving a system of equation $\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$, where \mathbf{A} is the correlation matrix between the variables, \mathbf{v} is the eigen vector, and λ is the eigen value. The eigen value in the two-dimensional case is one plus the correlation or one minus the correlation (there are two eigen values). The eigen vectors are not unique, so it is standard to impose the condition that the vectors are standardized to have a length of one. Only in this very special case where we have two standardized variables, the first eigen vector is $(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ if the correlation between the two variables is positive, while the second is $(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$. If the variables have a negative correlation, the first principal component is $(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$, while the second is $(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$. The first principal component is plotted in Exhibit 4.

References

- Bruno, G., Esakia, M., and Goltz, F. (2021). "Honey, I Shrunk the ESG Alpha: Risk-Adjusting ESG Portfolio Returns," Scientific Beta Working Paper.
- Jacobsen, B., Cheng, E., and Lee, W. (2021). "Climate Change and Asset Allocation: A Distinction That

- Makes a Difference,” *The Journal of Portfolio Management* Multi-Asset Special Issue, jpm.2021.1.218; <https://doi.org/10.3905/jpm.2021.1.218>.
- Jacobsen, B., Ma, C., and Lee, W. (2019). “The Alpha, Beta, and Sigma of ESG: Better Beta, Additional Alpha?” *The Journal of Portfolio Management* September, jpm.2019.1.091; <https://doi.org/10.3905/jpm.2019.1.091>.
- Jacobsen, B. and Lee, W. (2020). “Risk Parity Optimality Even with Negative Sharpe Ratio Assets,” *The Journal of Portfolio Management* Multi-Asset Special Issue, jpm.2020.1.151; <https://doi.org/10.3905/jpm.2020.1.151>.
- Jacobsen, B. and Ma, C. (2020). “Alpha Alchemy: Diversifying without Diluting Alpha,” *The Journal of Wealth Management* **23**(2), 75–87; <https://doi.org/10.3905/jwm.2020.1.106>.
- Luis, J. and Alvarez, C. (2021). “The Economic Geography of Global Warming,” National Bureau of Economic Research Working Paper 28466.
- Madhavan, A., Sobczyk, A., and Ang, A. (2020). “Towards ESG Alpha: Analyzing ESG Exposures through a Factor Lens,” *Financial Analysts Journal* **77**(1), 69–88.
- Mejia, S., Mrkaic, M., Novta, N., Pugacheva, E., and Topalova, P. (2018). “The Effects of Weather Shocks on Economic Activity: What are the Channels of Impact?” IMF Working Paper WP/18/144. (<https://www.imf.org/en/Publications/WP/Issues/2018/06/22/The-Effects-of-Weather-Shocks-on-Economic-Activity-What-are-the-Channels-of-Impact-45970>).
- Pastor, L., Stambaugh, R., and Taylor, L. (2021). “Dissecting Green Returns,” National Bureau of Economic Research Working Paper.
- Sharpe, W. (1964). “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk,” *Journal of Finance* **19**(3), 425–442.

Keywords: Risk budgeting, climate change, ESG, asset allocation, manager selection