

Dissecting Green Returns

Ľuboš Pástor

Robert F. Stambaugh

Lucian A. Taylor*

June 10, 2021

Abstract

Green assets delivered high returns in recent years. This performance reflects unexpectedly strong increases in environmental concerns, not high expected returns. German green bonds outperformed their higher-yielding non-green twins as the “greenium” widened, and U.S. green stocks outperformed brown as climate concerns strengthened. To show the latter, we construct a theoretically motivated green factor—a return spread between environmentally friendly and unfriendly stocks—and find that its positive performance disappears without climate-concern shocks. The factor lags those shocks, curiously, by about a month. A theory-driven two-factor model featuring the green factor explains much of the recent underperformance of value stocks.

JEL classifications: G12, G14

Keywords: Sustainable investing, ESG, environment, climate change, green bonds

*Pástor is at the University of Chicago Booth School of Business. Stambaugh and Taylor are at the Wharton School of the University of Pennsylvania. Pástor and Stambaugh are also at the NBER. Pástor is additionally at the National Bank of Slovakia and the CEPR. The views in this paper are the responsibility of the authors, not the institutions they are affiliated with. We thank Livia Amato for excellent research assistance. This research was funded in part by the Fama-Miller Center for Research in Finance at the University of Chicago Booth School of Business.

1. Introduction

The growth of sustainable investing is one of the most dramatic trends in the investment industry over the past decade. Today, sustainable strategies comprise one third of professionally managed U.S. assets (US SIF Foundation, 2020). Sustainable investing applies environmental, social, and governance (ESG) criteria, with environmental concerns playing the leading role. For example, 88% of the clients of BlackRock, the world’s largest asset manager, rank environment as “the priority most in focus” among ESG criteria (BlackRock, 2020). Investments considered environmentally friendly are often referred to as “green,” with “brown” denoting the opposite.

Asset managers often market sustainable investment products as offering superior risk-adjusted returns.¹ Past performance is a popular marketing tool, and indeed a number of studies report superior historical returns to sustainable strategies (e.g., Edmans, 2011, Nagy, Kassam, and Lee, 2016, and In, Park, and Monk, 2019). Of course, as the SEC generally requires of all marketed funds, managers must warn that past performance does not necessarily predict future performance. In this study we show why investors would be especially well advised to heed that warning when investing in green assets.

What does the past performance of green assets imply about their future performance? We address this question empirically, guided by the equilibrium model of Pástor, Stambaugh, and Taylor (2021, henceforth PST). The PST model predicts that green assets have lower *expected* returns than brown, due to investors’ tastes for green assets, yet green assets can have higher *realized* returns while agents’ tastes shift unexpectedly in the green direction. This wedge between expected and realized returns is central to our paper. As PST explain, green tastes can shift in two ways. First, investors’ preference for green assets can increase, directly driving up green asset prices. Second, consumers’ demands for green products can strengthen—for example, due to environmental regulations—driving up green firms’ profits and thus their stock prices. Similarly, investors’ preference for brown assets or consumers’ demand for brown products can decrease, again making green stocks outperform. We also leverage PST’s result that assets are priced by a two-factor model, where the factors are the market portfolio and the ESG factor. The ESG factor is the return on a portfolio that goes long green and short brown assets, where the assets are weighted by their greenness. The ESG factor’s expected return, as derived by PST, is negative.

¹For example, BlackRock believes that “integrating sustainability can help investors “build more resilient portfolios and achieve better long-term, risk-adjusted returns” (Fink, 2021). According to State Street, “ESG is a source of alpha that leads to positive portfolio performance” (Lester and He, 2018).

Our analysis focuses primarily on the U.S. stock market. We use individual stocks' environmental ratings from MSCI, a leading provider of ESG ratings. We construct a monthly "green factor," a return spread between green and brown stocks, following PST's procedure for constructing the ESG factor. Our sample begins in November 2012, when MSCI's data coverage increased sharply, and ends in December 2020. We find that the green factor earned a cumulative return spread of 35% over this period. The factor's average return is 31 basis points (bps) per month, with a t -statistic of 2.91. In short, green stocks significantly outperformed brown stocks in recent years.

Should green stocks' recent outperformance lead one to expect high green returns going forward? No, we argue. That outperformance likely reflects an unanticipated increase in environmental concerns. To reach this conclusion, we compute a measure of concerns about climate change, using the media index constructed by Ardia et al. (2021). We observe a steady increase in climate concerns during the last decade, with the level of our measure nearly doubling. We then isolate monthly shocks to climate concerns and find they exhibit a significant positive relation to the green factor. In other words, green stocks tend to outperform when there is bad news about climate change. If we set the climate shocks to zero, the green factor's estimated counterfactual performance becomes flat. That is, green stocks would not have outperformed brown without strengthened climate concerns.

In fact, green stocks might have underperformed brown, absent strengthened climate concerns. Such a possibility is suggested when we include capital flows into sustainable mutual funds as another measure of changes in climate concerns. The relation between fund flows and the green factor is more difficult to estimate reliably, due to endogeneity issues. Nevertheless, given our point estimates, when we zero out sustainable fund flows as well as the above climate-concern shocks, the resulting counterfactual performance of the green factor becomes negative, consistent with the factor having a negative expected return.

Our empirical explanation of green stocks' outperformance accords with the PST model. During a period when climate concerns strengthen sufficiently, the green factor delivers a positive return, as investors demand greener stocks or customers demand greener products. Outperformance caused by the strengthening of investor concerns is followed by lower expected performance of the green factor going forward. That is, a shift in the green factor's expected future performance relates inversely to its realized performance.

An inverse relation between realized returns and shifts in expected returns is not new in the stock return literature.² With stocks, a challenge to documenting this relation is

²For example, this inverse relation figures prominently in empirical analyses of the equity premium by

that expected stock returns are unobservable and generally hard to estimate. With bonds, however, we can see the relation more clearly. The inverse relation between a bond's realized return and the change in its yield to maturity is well understood, and the yield provides direct information about expected return, especially for buy-and-hold investors.

The case of German “twin” bonds illustrates this inverse relation in the context of climate concerns. Since 2020, the German government has issued green bonds, along with virtually identical non-green twins. The green bonds trade at lower yields, indicating lower expected returns compared to non-green bonds. The yield spread between the green and non-green twins, known as the “greenium,” reflects investors' willingness to accept a lower return in exchange for holding assets more aligned with their environmental values. Since issuance, the 10-year greenium experienced roughly a three-fold widening, presumably due to growing climate concerns. As a result, the green bond outperformed its non-green twin by a significant margin over the same period. However, this outperformance does not imply green outperformance going forward. Rather the opposite is clearly true, given the now wider greenium. This case study has a counterpart in the outperformance of the green factor in stocks. A downward shift in the green factor's expected future return is simply less easily documented, given that stocks offer no directly observable analog to the greenium.

Our main results relating climate shocks to green stock returns rely on the time series of the green factor. We also conduct a parallel analysis by running panel regressions on individual stocks, leading to two findings. First, there is a significantly positive relation between a stock's greenness and its average return. Second, that positive relation disappears when we interact the stock's greenness with climate-concern shocks, revealing that these shocks fully account for the superior performance of green stocks during the sample period. Both results echo our time-series evidence: despite having lower expected returns, green stocks outperform brown due to positive surprises over the sample period.

Green stocks' recent outperformance helps us understand the poor performance of value stocks in the 2010s, the worst decade on record for the HML factor of Fama and French (1993). We examine this performance through the lens of PST's two-factor model, with our green factor assuming the role of the ESG factor. We find that the two-factor model explains much of HML's recent underperformance. From November 2012 through December 2020, HML's monthly CAPM alpha is a marginally significant -71 bps, whereas HML's two-factor alpha is an insignificant -15 bps. In contrast, the green factor's alpha with respect to the Fama-French three-factor model is a significant 21 bps. The green factor and HML are negatively correlated, as value stocks are more often brown than green. Insofar as recent

Fama and French (2002) and Pástor and Stambaugh (2001, 2009).

average performance, however, the two-factor model explains HML's underperformance better than the three-factor model explains the green factor's outperformance. The two-factor model can also explain the momentum strategy's positive performance over the same period: momentum's monthly CAPM alpha is 66 bps, whereas its two-factor alpha is -6 bps.

As noted earlier, the green factor has a significantly positive relation to climate-concern shocks. Curiously, the factor reacts to those shocks with a nontrivial delay. While the factor has only a weak and insignificant positive relation to the current month's climate shock, it has a significantly positive relation to the previous month's shock. Consistent with these monthly results, at a weekly frequency we find climate shocks enter positively at lags of two to five weeks, most strongly at four weeks. It seems that stock prices are slow to incorporate relevant climate news. Ardia et al. (2021) find a positive contemporaneous daily relation between green stock returns and climate news. We confirm that result, but we also find that the relation flips to negative at lags of one and two days, offsetting the positive contemporaneous relation. This behavior is consistent with temporary price pressure from trading on same-day climate news. The bulk of the positive relation between green stock returns and climate-concern shocks evidently occurs with multi-week lags. Our results complement those of Hong, Li, and Xu (2019). They also find that stock prices are slow to react to climate-change risks, but they look at different assets (stocks in food industries across countries) and different climate shocks (trends in the risks of drought).

Our study relates to a large empirical literature investigating returns on green versus brown assets. One set of studies examine returns on an *ex ante* basis, using proxies for expected future returns. In the bond market, for example, Baker et al. (2018), Zerbib (2019), and Larcker and Watts (2020) analyze yields on green bonds versus brown. In the stock market, Chava (2014) and El Ghouli et al. (2011) compare implied costs of capital estimated for green firms versus brown. Most of these studies find lower *ex ante* returns on green assets, consistent with theory. A second, larger set of studies examine returns on an *ex post* basis, measuring realized green-versus-brown returns, generally for stocks. Examples include In, Park, and Monk (2019), Bolton and Kacperczyk (2020, 2021), Gorgen et al. (2020), and Hsu, Li, and Tsou (2020). We depart from all of these studies by focusing on the distinction between *ex ante* and *ex post* returns. In particular, we show why high green returns realized in recent years are likely to be misleading predictors of the future.

Our evidence on how climate shocks affect realized returns also relates to studies investigating the pricing of climate risk. Recent work examines that pricing in equities (e.g., Bolton and Kacperczyk, 2020, 2021, Hsu, Li, and Tsou, 2020, and Faccini, Matin, and Skiadopoulos, 2021), corporate bonds (Huynh and Xia, 2021, and Seltzer, Starks, and Zhu, 2021), municipi-

pal bonds (Painter, 2020, and Goldsmith-Pinkham et al., 2021), options (Ilhan, Sautner, and Vilkov, 2021), and real estate (Bernstein, Gustafson, and Lewis, 2019, Baldauf, Garlappi, and Yannelis, 2020, and Giglio et al., 2021). Engle et al. (2020) develop a procedure to dynamically hedge climate risk with the help of mimicking portfolios and textual analysis of news sources. Krueger, Sautner, and Starks (2020) document the importance of climate risk in a survey of institutional investors. For a survey of the climate finance literature, see Giglio, Kelly, and Stroebel (2020).

Our empirical analysis is guided by the theoretical model of PST, in which investors' tastes for green assets play a key role. Other models featuring tastes for green assets can be found in Fama and French (2007), Baker et al. (2018), Pedersen, Fitzgibbons, and Pomorski (2021), and Avramov et al. (2021). In some of these models, tastes are not the only force determining green assets' expected returns. For example, in the model of Pedersen et al., green returns are boosted by the presence of ESG-unaware investors. Their mechanism offers an alternative way to view positive green returns caused by increases in consumers' demands for green products. While these returns are viewed as unexpected in the PST setting, in the Pedersen et al. setting they are partly expected by investors who anticipate the taste shifts before market prices respond. In the model of Avramov et al., expected returns depend not only on green tastes but also on uncertainty about the firm's greenness. We rely on the PST model because it analyzes the effects of taste shifts and provides guidance in constructing the ESG factor.

Our results have important implications for research and practice. They underline the danger in using recent average returns to estimate expected returns. In particular, the recent outperformance of green assets does not imply high green returns going forward. In fact, if the outperformance resulted from increased demands by ESG investors, then green assets' expected returns are lower today than a decade ago. In the same spirit, value stocks' recent underperformance is less likely to continue, because value stocks tend to be brown and growth stocks green. From the corporate finance perspective, our findings imply that greener firms have lower costs of capital than their recent stock performance might suggest. This is good news for ESG investors, because one way they exert social impact is by decreasing green firms' cost of capital (e.g., Heinkel et al., 2001, PST).

This paper is organized as follows. Section 2 highlights the gap between expected and realized returns in the context of German twin bonds. Section 3 describes how we measure greenness in our main analysis, which is based on U.S. stocks. Section 4 discusses how we construct the green factor and measure its performance. Section 5 relates this performance to proxies for shifts in green tastes, such as climate news and flows into sustainable funds.

Section 6 analyzes the relation between returns and greenness at the stock level. Section 7 documents the delayed reaction of stock prices to climate news. Section 8 concludes.

2. German twin bonds

This paper emphasizes the difference between expected and realized returns on green assets. Quantifying this difference for stocks is challenging because expected stock returns are not directly observable. In this section, we illustrate this difference for bonds, whose expected returns are tightly linked to yields to maturity. Conveniently, bond yields are easily observable.

Since 2020, the government of Germany, the largest European economy, has been issuing green securities to finance its transition towards a low-carbon, sustainable economy.³ The first green security, a 10-year bond, was issued in September 2020 in the amount of 6.5 billion euros. The second green security, a 5-year note, followed two months later in the amount of 5 billion euros. Both securities have zero coupon rates. Germany plans to issue at least one green security per year, including a 30-year bond in May 2021 and a 10-year bond in September 2021. We refer to these securities as “green bonds.”

Each green bond is issued with the same characteristics as an existing conventional bond issued by the German government. Besides having the same issuer, the two bonds have the same maturity date, the same coupon rate, and the same coupon payment dates. This pairing creates “twin” bonds, which offer identical streams of cash flows with identical credit risk but different greenness. By comparing the prices of twin bonds, we can gain some insight into the value assigned to greenness by bond market investors.

Even though the twin bonds are paired very carefully, some differences between them remain. First, the issuance date of the green bond always comes after the initial issuance date of the conventional bond. For example, the green bond issued in September 2020 has a conventional twin issued in June 2020. Second, conventional bonds tend to be issued at larger volumes than their green twins. For example, in 2020, the issuance of conventional bonds was almost five times larger than that of the corresponding green bonds. Conventional bonds could thus in principle be more liquid than their green twins. However, the German Finance Agency has committed to play an active role in the secondary market for green bonds to make their liquidity comparable to that of conventional bonds.

³For more details, see <https://www.deutsche-finanzagentur.de/en/institutional-investors/federal-securities/green-federal-securities/>.

We obtain daily data on the first pair of twin bonds, downloading the end-of-day bond prices and mid-yields to maturity for the 10-year green bond (ISIN DE0001030708) and the 10-year non-green bond (DE0001102507) from Bloomberg. We download all available data since the first date of trading for the green bond, which is September 8, 2020, through the present date of April 12, 2021. Over this 7-month period, the two bonds' annual yields fluctuate between -67 and -27 bps. We plot these yields in the Appendix.⁴

Panel A of Figure 1 plots the time series of the difference between the yields of green and non-green bonds, also known as the green premium, or the “greenium” (e.g., Larcker and Watts, 2020). The greenium is always negative, ranging mostly between -5 and -2 bps per year.⁵ Therefore, for investors holding the bonds to maturity, the green bond always has a lower expected return than the non-green bond. This evidence is consistent with theories predicting that green assets offer lower expected returns than non-green assets (e.g., Pástor, Stambaugh, and Taylor, 2021).⁶

Given the lower yield of the green bond, one would expect it to deliver a lower return than its conventional twin. Instead the green bond delivered a higher return in our sample. We calculate bond returns as daily percentage changes in bond prices. The full-sample cumulative returns are negative, -1.47% for the green bond and -1.78% for the non-green bond, due to a rise in yields between September 2020 and April 2021. More interesting, the green bond outperforms its non-green twin. This outperformance accrues steadily through the sample, as shown by Panel B of Figure 1. The figure plots cumulative returns on a long-short portfolio, which goes long the green bond and short the non-green bond. The portfolio's average daily return of 0.2 bps is statistically significant ($t = 2.33$), and its cumulative return of 31 bps is substantial relative to German government bond yields.

Importantly, the positive average return of the long-short portfolio does not imply that the portfolio's expected return is positive. On the contrary, we know with certainty that the portfolio's expected return is negative if the bonds are held to maturity. For example, on September 8, 2020, the green bond's yield was -51.2 bps per year, whereas the non-green

⁴The Appendix is available on the authors' websites. It reports the results also for the second pair of twin bonds, which was first issued in November 2020. Those results are similar to those presented here. We prioritize the first twin pair due to its longer history.

⁵These greenium values are close to those estimated by prior studies in different settings. For example, Baker et al. (2018) estimate a greenium of about -6 bps in a sample of over 2,000 U.S. municipal and corporate green bonds, whereas Zerbib (2019) estimates -2 bps in a sample of over 1,000 supranational, sub-sovereign and agency, municipal, corporate, financial, and covered green bonds.

⁶This conclusion is reinforced by liquidity considerations. As noted earlier, the non-green bond has been issued at larger volume than its green twin. If this volume difference makes the conventional bond more liquid despite the aforementioned efforts of the German Finance Agency, then the resulting liquidity premium pushes the greenium up, and the expected return penalty associated with greenness is even larger.

bond's yield was -49.6 bps. Therefore, if both bonds are held to maturity, the green bond delivers a return 1.6 bps lower than the non-green bond. The green bond's expected return is lower also if the bonds are not held to maturity under a variety of plausible conditions, such as changes in the greenium being unpredictable. That condition is likely to hold, especially in efficient, or near-efficient, markets. Under that condition, the green bond's expected return is lower at the beginning of the sample, and the expected return of the long-short portfolio is negative. The cumulative value of this expected return is plotted by the dashed line in Panel B of Figure 1, which is gently downward-sloping.

How can we reconcile the higher realized return of the green bond with its lower expected return? The answer is that the greenium in Panel A grows increasingly negative between September 2020 and April 2021, deepening from -1.6 to -5.1 bps. This steady deepening is responsible for the steady outperformance of the long-short portfolio in Panel B. In the language of Pástor, Stambaugh, and Taylor (2021), if investors' tastes shift toward green assets, they push up the price of the green bond relative to the non-green bond. However, the green bond's outperformance is temporary, as it comes entirely at the expense of the bond's future return. Investors buying the bonds on September 8, 2020 and holding them to maturity expected to earn 1.6 bps less from the green bond, but those buying on April 12, 2021 expected to earn 5.1 bps less.

This example illustrates how time variation in expected returns drives a wedge between returns expected ex ante and those realized ex post. Even though the green bond's realized return is higher than that of the non-green bond, the green bond's expected long-term return is demonstrably lower. In other words, the expected return of the long-short portfolio is negative even though the portfolio's average realized return is positive (and significant at the 5% confidence level). Unlikely events do happen sometimes, and the outperformance of the German green bond in the first seven months of its existence is one of them.

3. Measuring stocks' greenness

We compute stock-level environmental scores based on MSCI ESG Ratings data, a successor to the MSCI KLD data used in many academic studies. Our data have a number of advantages. According to Eccles and Stroehle (2018), MSCI is the world's largest provider of ESG ratings. The MSCI ESG Ratings data are used by more than 1,700 clients, including pension funds, asset managers, consultants, advisers, banks, and insurers.⁷ MSCI covers more

⁷See <https://www.msci.com/our-solutions/esg-investing>, as of May 2021. In addition, MSCI has been voted 'Best firm for SRI research' in the Extel & SRI Connect Independent Research in Responsible

firms than other ESG raters, such as Asset4, KLD, RobecoSAM, Sustainalytics, and Vigeo Eiris (Berg et al., 2020). MSCI generates its ratings from corporate documents, government data, various journals, and news media. It updates those ratings at least annually. MSCI’s ESG research unit employs more than 200 analysts and incorporates artificial intelligence, machine learning, and natural language processing into its methodology.

The availability of industry-unadjusted granular data is another advantage of MSCI data for our purposes. With industry adjustment, a heavily-polluting firm is classified as green if it pollutes less than other firms in its heavily-polluting industry. That unappealing scenario does not arise in PST’s theory, in which there is no industry adjustment. We therefore use MSCI’s granular data to compute an environmental measure that is not industry adjusted. In contrast, MSCI’s composite ESG rating is industry-adjusted, as are ratings from other leading providers.

We use the MSCI variables “Environmental pillar score” (E_score) and “Environmental pillar weight” (E_weight). E_score is a number between 0 and 10 measuring the firm’s weighted-average score across 13 environmental issues related to climate change, natural resources, pollution and waste, and environmental opportunities. These scores are designed to measure a company’s resilience to long-term environmental risks. E_weight , which is typically constant across firms in the same industry, is a number between 0 and 100 measuring the importance of environmental issues relative to social and governance issues.⁸

We compute the unadjusted greenness score of firm i at the beginning of month t as

$$G_{i,t-1} = -(10 - E_score_{i,t-1}) \times E_weight_{i,t-1}/100 , \quad (1)$$

where $E_score_{i,t-1}$ and $E_weight_{i,t-1}$ are from company i ’s most recent MSCI ratings date before month t , looking back no more than 12 months. The quantity $10 - E_score_{i,t-1}$ measures how far the company is from a perfect environment score of 10. The product $(10 - E_score_{i,t-1}) \times E_weight_{i,t-1}$ measures how brown the firm is, specifically, the interaction of how badly the firm scores on environmental issues and how large the environmental impacts are for the industry’s typical firm (i.e., $E_weight_{i,t-1}$). The initial minus sign converts the measure from brownness to greenness.

Including E_weight in equation (1) is important for capturing a company’s greenness. For example, in 2019, Exxon Mobil and Best Buy had similar E_score values: 4.2 and 4.1,

Investment Survey in each year from 2015 through 2019 (<https://www.msci.com/zh/esg-ratings>).

⁸MSCI’s E, S, and G weights sum to 100. According to MSCI, “The weightings take into account both the contribution of the industry, relative to all other industries, to the negative or positive impact on the environment or society; and the timeline within which we expect that risk or opportunity for companies in the industry to materialize....” We follow MSCI in using the GICS sub-industry classification.

respectively. If we only used E_score , we would judge these companies to be similarly green. But E_weight is 48 for Exxon and only 11 for Best Buy, reflecting that oil and gas companies have larger environmental impacts than consumer retail companies. Exxon and Best Buy end up with $G_{i,t} = -2.78$ and -0.65 , respectively, indicating that Best Buy is much greener than Exxon. Similar to us, MSCI uses the interaction of E_score and E_weight when computing firms' composite ESG ratings.⁹

The environmental score we use in our analysis is

$$g_{i,t} = G_{i,t} - \bar{G}_t, \quad (2)$$

where \bar{G}_t is the value-weighted average of $G_{i,t}$ across all firms i . Since we subtract \bar{G}_t , $g_{i,t}$ measures the company's greenness relative to the market portfolio, as in PST.

We compute $g_{i,t}$ in the sample of stocks with non-missing MSCI data and CRSP share codes of 10 or 11. We merge CRSP and MSCI by using a combination of CUSIP, ticker, and company name. Our sample extends from November 2012 to December 2020. We begin in November 2012 because MSCI's coverage increases dramatically in October 2012, when MSCI began covering small U.S. stocks.¹⁰ Figure 2 plots the number of U.S. stocks with non-missing lagged MSCI ratings. This number increases sharply in November 2012, from roughly 500 to over 2,000. Our purchased MSCI data end in March 2020, but we extend our sample through December 2020 by looking back up to 12 months when computing $G_{i,t-1}$.

Table 1 shows industries ranked by their equal-weighted average $g_{i,t}$ scores at the end of 2019. The lowest-ranked industries include chemicals, oil and gas exploration and production, steel, mining (including coal), paper and forest products, and marine transport. It is reassuring that these industries, which are generally viewed as having negative environmental impacts, appear at the bottom of our ranking.

⁹MSCI's composite ESG rating is based on their "Weighted Average Key Issue" score, which equals $[E_score \times E_weight + S_score \times S_weight + G_score \times G_weight]/100$, where S and G refer to social and governance. So if MSCI used a formula like equation (1) to compute greenness not just on environmental but also on social and governance dimensions, then we could express MSCI's composite ESG score as 10 plus the sum of E, S, and G greenness.

¹⁰Before October 2012, MSCI covered only the largest 1,500 companies in the MSCI World Index, plus large companies in the UK and Australia MSCI indexes. In October 2012 MSCI added many smaller U.S. firms when it began covering also the MSCI U.S. Investible Market Index.

4. The green factor

4.1. Constructing the green factor

We construct the green factor by following the ESG factor methodology derived by PST. The factor is a return spread between environmentally friendly and unfriendly stocks. PST show that the factor's realizations can be estimated month by month by running cross-sectional regressions of market-adjusted excess stock returns on the stocks' greenness characteristics, with no intercept. The slope from one such regression, which represents the green factor's realization in month t , is given in equation (34) of PST as

$$\hat{f}_{gt} = \frac{g'_{t-1} \tilde{r}_t^e}{g'_{t-1} g_{t-1}}, \quad (3)$$

where g_{t-1} is the vector containing stocks' greenness characteristics, $g_{i,t-1}$, and $\tilde{r}_t^e \equiv \tilde{r}_t - \beta_{m,t-1} \tilde{r}_{mt}$ is the vector of stocks' market-adjusted excess returns. Specifically, \tilde{r}_t is the vector of stocks' returns in excess of the risk-free rate, \tilde{r}_{mt} is the market return in excess of the risk-free rate, and $\beta_{m,t-1}$ is the vector of stocks' market betas, which we estimate from rolling monthly regressions of individual stocks' excess returns on excess market returns using up to 60 months (and no less than 36 months) of data ending in month t . Since \tilde{r}_t and \tilde{r}_{mt} are excess returns, the green factor is the return on a zero-cost portfolio.

Zero-cost portfolios are commonly used as factors in the finance literature, but they are generally empirically motivated, and their construction details are somewhat arbitrary (e.g., Fama and French, 1993, 2015). In contrast, our green factor is theoretically motivated, and its construction methodology is derived analytically. The factor is essentially a portfolio of market-adjusted stock positions where each stock is weighted by its greenness, with green stocks receiving positive weights and brown stocks negative weights. PST show that when factor weights are based on greenness rather than market capitalization, assets are priced in equilibrium by two factors: the market portfolio and the green factor.

4.2. The green factor's performance

Figure 3 plots the green factor's cumulative return. Green stocks strongly outperformed brown in the 2010s, with a cumulative return difference of nearly 40% over our 8.2-year sample period. The factor averaged 31 bps per month (t -statistic: 2.91), as reported in the first column of Table 2. The table's remaining columns report results of regressing the green factor on various other return factors, including those in the three- and five-factor

models of Fama and French (1993, 2015), the momentum factor (UMD) as constructed by those authors, and the traded liquidity factor of Pástor and Stambaugh (2003). In all cases the green factor's alpha (regression constant) is economically and statistically significant, ranging from 18 to 37 bps per month, with t -statistics between 2.42 and 3.38.

Much of the above performance of green stocks stems from industry-level greenness. If we construct the green factor as in equation (3) but instead use industry-adjusted $g_{i,t-1}$ values, the resulting factor's performance is substantially lower, by more than half. Details appear in the Appendix.

The green factor's lowest alpha in Table 2 occurs in column 4, where we adjust for the three Fama-French factors and momentum. Its significant exposures to SMB, HML, and UMD indicate that the green factor tilts toward large stocks, growth stocks, and recent winners. Net of those exposures, the factor's alpha is 18 bps per month ($t = 2.46$).

4.3. Pricing value and momentum

During our sample period, the market-adjusted monthly alphas of HML and UMD are -71 bps and 66 bps, respectively, with t -statistics of -1.93 and 1.92, as shown in columns 1 and 3 of Table 3. The green factor's significant exposures to value and momentum, noted above, prompt us to ask a performance question in the reverse direction: To what extent can the green factor's strong performance account for the last decade's historic underperformance of value, or for the positive performance of momentum?

To address this question, we turn to the equilibrium setting of PST, in which expected returns obey a two-factor model that includes the market and an ESG factor. Here we assign the latter role to the green factor. HML's and UMD's alphas with respect to the two-factor model, which are shown in columns 2 and 4 of Table 3, are much smaller in magnitude than with just market adjustment. HML's alpha becomes -15 bps instead of -71 bps; UMD's alpha becomes -6 bps instead of 66 bps. The t -statistics shrink to -0.50 and -0.22 .

While the green factor's significant performance survives controlling for HML and UMD exposures, the reverse is not true. Nearly 80% of HML's negative alpha, and all of UMD's positive alpha, disappear after controlling for the green factor's strong performance. Recognizing the brown nature of value stocks, and the green nature of growth stocks, thus helps us understand why the value strategy experienced its worst decade ever in the 2010s.

5. Explaining the green factor's performance

What accounts for the green factor's strong performance over the last decade? After all, the factor's expected performance is negative, according to PST's model. As those authors explain, however, the green factor's realized performance can be positive in periods of unanticipated increases in demands for green firms' products and stocks (or decreases in demands for brown firms' products and stocks). These green demands can increase for various reasons, but a likely leading source is increased concerns about climate change.

In this section we investigate whether the green factor's positive performance can be explained by increases in climate concerns and green demands. We first describe how we proxy empirically for changes in (i) climate concerns, via media coverage, (ii) green-product demand, via firms' earnings news, and (iii) green-investment demand, via flows into sustainable funds. We then estimate the extent to which these shocks explain the green factor's realized performance. We find that climate concerns play the most important role, explaining virtually all of the green factor's positive performance over the sample period.

5.1. Measuring climate concerns

We measure concerns about climate change by using the Media Climate Change Concerns index (MCCC) of Ardia et al. (2021). This index, which is available from January 2003 through June 2018, is constructed by using data from eight major U.S. newspapers. It captures the number of climate news stories each day as well as their negativity and focus on risk. For each news article discussing climate change, Ardia et al. compute a "concern" measure that interacts two quantities: the fraction of total words related to risk and the scaled difference between negative and positive words. They aggregate this measure to the newspaper-day level by adding the concern values across stories. Next, they aggregate to the day level by averaging across newspapers, after adjusting for heterogeneity across newspapers. Finally, they take the square root of this daily measure because, as they put it, "One concerning article about climate change may increase concerns, but 20 concerning articles are unlikely to increase concerns 20 times more."

We measure the level of climate concern by using a distributed-lag model that assumes individuals' memory of climate news stories decays gradually over time. Let $MCCC_t$ be the MCCC index averaged across days in month t . We define the level of climate concerns at

the end of month t as

$$C_t = \sum_{\tau=0}^T \rho^\tau MCCC_{t-\tau} , \quad (4)$$

where $0 < \rho < 1$ measures how long climate news persists in investors' memories. We set the half-life of news stories to one year, which implies $\rho = 0.94$. The C_t series looks very similar for similar values of ρ . We set $T = 36$ months, because the MCCC index is relatively short-lived and its further lags have small effects on C_t (as $0.94^{36} \approx 0.1$).

Figure 4 plots the climate concern measure, C_t , between November 2012 and June 2018. The level of C_t nearly doubles during this period. Over the same period, the green factor's performance, also plotted in Figure 4, is strongly positive, cumulating to nearly 18%.

Green factor returns should respond to unanticipated changes in climate concerns. The change in climate concerns, defined as $\Delta C_t \equiv C_t - C_{t-1}$, follows from equation (4):

$$\Delta C_t = MCCC_t - \sum_{\tau=1}^T MCCC_{t-\tau} \rho^{\tau-1} (1 - \rho) - \rho^T MCCC_{t-1-T} . \quad (5)$$

We treat this change as unanticipated, given that its autocorrelation is insignificantly different from zero.

Ardia et al. compute unexpected changes in climate concerns as the prediction errors from rolling AR(1) models applied to the MCCC index. Our motivation for ΔC_t is different, but ΔC_t has a 94% correlation with the AR(1) error series. While we prefer our approach, we find very similar results if we use the latter series, as we report in the Appendix.

5.2. Other drivers of green demands

Increased climate concerns are likely to play a key role in boosting demands by consumers for green firms' products as well as demands by investors for green firms' stocks. Nonetheless, green demands can also arise from other sources. We allow for such sources by including additional proxies for green demand shifts. We first focus on the product-demand channel by constructing measures of firms' earnings news. We then turn to the investment-demand channel by measuring flows of capital into sustainable funds. Of course, all of the above forces driving green demands also drive brown demands, just in the opposite direction.

News about firms' profits affects the green factor's performance. Positive performance can reflect better earnings news for green firms than brown, due to effects not necessarily captured by our climate-concern measure. To allow for such effects, we compute two earnings-news measures using data from CRSP and I/B/E/S.

The first measure is based on the idea that a large portion of earnings news occurs on days when firms make earnings-related announcements (Beyer et al., 2010). We consider two types of announcements: those of quarterly earnings and voluntary forward guidance regarding future earnings. We compute stock returns in excess of the market during the three-trading-day windows centered on these announcement dates. We add the excess returns across unique events within a given stock-quarter. For about 70% of observations, no summation is required because the forward-guidance date coincides with the earnings-announcement date. Many firms never issue forward guidance, and some firms start or stop issuing forward guidance. We find that our announcement-return measure explains about 20% of the variance of quarterly stock-level returns (see the Appendix).

Our second measure captures news about long-term earnings. Such news can arrive gradually over time, in between the quarterly announcements. This second measure uses data on analysts' forecasts of each firm's long-run earnings growth rate. For firm i and quarter t , the measure equals the earliest mean analyst forecast in quarter $t + 1$ minus the latest mean forecast in quarter $t - 1$. Using forecasts from quarters $t - 1$ and $t + 1$ helps to capture all news arriving in quarter t . The measure may also include a small amount of information that arrives in quarters $t - 1$ or $t + 1$, but those inclusions are innocuous since they should not help explain returns in quarter t . We winsorize this measure at the 1% level. We find that this measure is significantly related to quarterly stock-level returns but explains less than 1% of their variance (see the Appendix).

Measuring the part of returns coming from earnings news is known to be difficult, and our measures surely miss important earnings news. Our first measure misses short-term news that arrives outside the three-day announcement windows. Analysts' long-term forecasts are only three- to five-year forecasts, so they exclude news about distant earnings. Any climate-related news affecting earnings more than five years in the future would elude our measures. Another limitation is that analysts' forecasts can differ from investors' forecasts.

Since the green factor is an aggregate time-series variable, we need to convert our firm-level earnings measures into aggregate ones. We do so by following the construction of our green factor in equation (3). For earnings measure X_{it} , we compute its aggregate green-minus-brown counterpart as $g'_t X_t / (g'_t g_t)$, where X_t is the vector containing X_{it} .

To proxy for shifts in investors' demand for green assets, we use flows into sustainable funds. From Morningstar's *2021 Sustainable Funds U.S. Landscape Report*, we obtain data on quarterly total flows into U.S. sustainable funds.¹¹ We scale these flows, which we refer

¹¹The data combine active and passive funds, equity and bond funds, open-end funds, and ETFs. Morn-

to as “ESG flows,” by the average total market capitalization of CRSP stocks during the quarter. ESG flows increased dramatically in 2013–2020, especially beginning in 2019.

5.3. Sources of the green factor’s performance

Table 4 relates monthly green-factor performance to our proxies for shifts in green demands by investors and consumers. We first regress the green factor on climate-concern shocks from both the current and previous month. We find that the green factor’s response to climate news is significantly positive, albeit mostly lagged. (See Section 7 for a deeper analysis of the lag.) In column 1 of Table 4, the coefficient on the current climate shock is positive but insignificant, whereas the coefficient on the previous month’s shock is nearly four times larger and has a t -statistic of 2.85. Aside from this curious lag, the green factor’s positive response to climate-concern shocks makes sense. An increase in climate concerns is likely to raise both the demand for products of green firms and the demand to hold those firms’ stocks; both effects push green stock prices higher relative to brown. Climate shocks explain 17% of the green factor’s monthly variance. Column 2 of Table 4 adds the earnings-news variables to the previous regression. Both variables enter positively, indicating better earnings news for greener stocks, but neither coefficient is estimated precisely enough to be statistically significant. The climate coefficients remain similar to those in column 1.

What if there had been no shocks to climate concerns or to green-versus-brown earnings? Figure 5 compares the green factor’s realized performance to its counterfactual performance in the absence of climate and earnings shocks. Using the regression estimated in column 2 of Table 4, we compute the counterfactual monthly green factor as the regression intercept plus the estimated residual, thereby assuming zero shocks to climate concerns and earnings. (Equivalently, the counterfactual equals the realized value minus the regressors times their respective coefficients.) The dashed line in Figure 5 plots the cumulative counterfactual return, and the solid line shows the cumulative realized return. We also plot a 95% confidence interval around the counterfactual, recognizing that the regression coefficients are estimated with error. To compute this interval, we repeatedly draw regression coefficients from their estimated sampling distribution, use those coefficients to compute simulated counterfactual returns, and then plot the simulated returns’ 95% confidence intervals.

The striking result in Figure 5 is that, absent climate-concern and earnings shocks, the green factor’s performance is essentially flat. Moreover, the counterfactual performance is

ingstar defines a sustainable fund as follows: “For a fund to be included in the sustainable funds universe, it must hold itself out to be a sustainable investment. In other words, ESG concerns must be central to its investment process and the fund’s intent should be apparent from a simple reading of its prospectus....”

reliably below the realized performance, as the latter lies well outside the 95% confidence interval. Because the climate-concern and earnings shocks are unanticipated, so too is essentially all of the green factor's positive performance over the period. That is, the positive performance does not imply higher expected returns on green stocks versus brown.

Column 3 of Table 4 adds sustainable fund flows to the previous regression. Reverse causation is a potential concern when regressing returns on contemporaneous flows. Instead of flows (or shifts in ESG tastes) causing returns, flows could be chasing same-period returns. We address this potential endogeneity by instrumenting for same-quarter sustainable flow using its previous-quarter value. The exclusion restriction plausibly holds, because flows cannot chase future realized returns. We find large first-stage t -statistics, indicating that the relevance condition holds and there is no problem with weak instruments.

In column 3 we also add sustainable funds' lagged total AUM, a proxy for the level of ESG tastes. This addition is motivated by PST's theoretical result that the expected green-factor return depends negatively on the average strength of ESG tastes (see equation (33) in PST), and the size of the ESG industry depends positively on those tastes (see Figure 5 in PST). We obtain annual sustainable fund AUM from the previously mentioned Morningstar report.¹² We scale ESG AUM by the total market capitalization of CRSP stocks.

The coefficients on both of the fund variables in column 3 exhibit their predicted signs. The estimated coefficient on flows is positive, consistent with higher flows indicating stronger green-stock demands and thus upward pressure on green-stock prices. The lagged AUM of sustainable funds gets a negative estimated coefficient, consistent with the above PST prediction. As with the earnings shocks, however, neither of the fund variables enters with enough precision to achieve statistical significance. One possible explanation is that our proxies for ESG flows and assets are noisy, given that they are derived from data on U.S. sustainable mutual funds and ETFs, omitting other ESG holdings. The coefficients on the climate and earnings variables are virtually unchanged from their values in column 2. These results suggest that our proxies for earnings news and ESG flows do not provide information about shifts in green demands beyond that captured by our climate-concern variable.

We also decompose the green factor's performance with respect to its sources. Figure

¹²We convert Morningstar's annual series to a quarterly series by using data on ESG flows and market returns, to approximate capital gains and losses. We estimate ESG funds' AUM at the end of quarter t , denoted $A\hat{U}M_t$ as

$$A\hat{U}M_t = \begin{cases} \text{True, known } AUM_t \text{ if } t \text{ is the year's last quarter} \\ A\hat{U}M_{t-1}(1 + R_t^{mkt}) + ESGFlow_t(1 + \frac{1}{2}R_t^{mkt}) \text{ otherwise,} \end{cases} \quad (6)$$

where R_t^{mkt} is the market return in quarter t . The fraction $\frac{1}{2}$ reflects that flows arrive throughout a quarter.

6 displays the decomposition based on the estimates in column 3 of Table 4. The solid line plots the green factor’s realized cumulative performance. The other lines show the factor’s counterfactual performance after “turning off” one or more shocks. First, we compute counterfactual green-factor returns assuming zero changes in climate concerns. These counterfactual returns equal realized green-factor returns minus the contemporaneous and lagged changes in climate concerns times their respective estimated coefficients. Next, we additionally assume zero earnings shocks. Finally, we assume zero ESG flows. To plot the latter line, we set ESG flows to zero in the regression model, and we set ESG assets to their counterfactual values in the absence of ESG flows. Counterfactual ESG assets grow by ESG funds’ imputed returns but not their flows.¹³

Figure 6 reveals that climate-concern shocks account for most of the gap between the actual and counterfactual performance plotted previously in Figure 5; the additional contribution from earnings shocks is modest. New in Figure 6 is the effect of zeroing out sustainable fund flows, in addition to the climate-concern and earnings shocks removed previously. Under that all-in scenario, the counterfactual performance of the green factor is substantially negative, as shown by the bottom dash-dot line. In other words, shutting down all of the identified ex post shocks to the green factor leaves it with negative estimated performance, consistent with its theoretically implied ex ante performance.

Next, we analyze sources of the green factor’s alpha. As shown earlier in Table 2, the green factor’s performance over our sample period remains significantly positive after controlling for various factors. In particular, even though the green factor tilts away from value stocks, the value factor’s unprecedented underperformance in the 2010s cannot explain the green factor’s positive performance. To assess similar robustness of the results in this section, we compute each month’s realized green factor net of its exposure to the three factors of Fama and French (1993) by taking the intercept plus the residual from the time-series regression of the green factor on the Fama-French factors. We then use this green-factor “alpha” instead of the green factor to repeat the analyses in Table 4 and Figure 5. The results appear in

¹³Define $ESGAssets_t$ to be AUM at end of quarter t , $ESGFlows_t$ to be ESG flows during t , and $ESGR_t$ to be return on ESG funds in t . We impute the value of $ESGR_t$ by taking the average of $ESGR1_t$ and $ESGR2_t$ from the following two equations, which differ only in the assumed timing of flows:

$$ESGR1_t = \frac{ESGAssets_t - ESGFlows_t}{ESGAssets_{t-1}} - 1 \quad (7)$$

$$ESGR2_t = \frac{ESGAssets_t}{ESGAssets_{t-1} + ESGFlows_t} - 1. \quad (8)$$

We set $ESG\tilde{Assets}_t = ESGAssets$ in 2009q4, then in subsequent quarters we grow ESG assets by the imputed $ESGR_t$. We then scale $ESG\tilde{Assets}_t$ by the size of CRSP, as before. We find that counterfactual ESG assets are roughly flat from 2009 to 2020.

Table 5 and Figure 7.

The results in Table 5 are very similar to those in Table 4. Climate-concern shocks enter positively for both the current and previous month, with the previous month's effect being much larger and statistically significant. Although none of the other variables enter with significance, their coefficients again have their predicted signs, with the minor exception of ESG flow, which enters with a small negative coefficient. Overall, the regression results using the green-factor alpha deliver a virtually identical message to those using the green factor. Not surprisingly, the same statement then applies to the counterfactual analysis displayed in Figure 7. That is, zeroing out climate and earnings shocks removes all of the otherwise substantial positive alpha. We find similar results if we repeat these analyses using Carhart (1997) four-factor alphas, Pástor-Stambaugh (2003) four-factor alphas, or Fama and French (2015) five-factor alphas (see the Appendix).

As explained earlier, our measure of climate concerns builds on that of Ardia et al. (2021). Those authors in turn acknowledge the prior work of Engle et al. (2020), who construct two measures of climate concerns, also based on media coverage. Ardia et al. discuss those alternative measures and explain that their measure adds risk as an another component of climate concerns. We rely on that more recent measure, but we also examine the robustness of our results to including the Engle et al. measures. We find that doing so does not change our conclusions. For example, we augment the independent variables in column 3 of Table 4 by including climate-concern shocks based on both of the Engle et al. measures for the current and previous month. The coefficients on all of those additional variables are statistically insignificant. In contrast, the shocks we construct based on the Ardia et al. measure still enter in that augmented regression as they do in the earlier regressions: the coefficient is positive and insignificant for the current month but much larger and significant for the previous month. See the Appendix for details.

6. Greenness and individual stock returns

All of our empirical analysis thus far is based on the green factor. That factor's positive performance is essentially equivalent to green stocks outperforming brown, given the factor's construction in equation (3). Next, to show that our conclusions do not hinge solely on the green factor's time series, we run panel regressions using individual stocks.

Table 6 reports regressions of individual stock returns in month t on various regressors. All regressions include time fixed effects and therefore capture cross-sectional variation in

returns. We begin in column 1 with a single regressor: the stock's greenness, $g_{i,t-1}$. The coefficient on greenness is significantly positive, indicating that greener stocks perform better over the sample period, consistent with the positive performance of the green factor.

Column 2 of Table 6 adds two regressors: $g_{i,t-1}$ interacted with the climate-concern shock in months t and $t - 1$. The coefficients on those regressors mirror the results in column 1 of Table 4: the coefficient on $g_{i,t-1}$ interacted with month t 's climate-concern shock is positive but insignificant, while the coefficient on the interaction with month $t-1$'s shock is larger and significant. As before, this result implies that green stocks outperform when climate concerns increase, albeit mostly with a lag. Notably, the coefficient on stand-alone greenness, $g_{i,t-1}$, turns slightly negative.¹⁴ That is, all of green stocks' outperformance indicated in column 1 disappears after we remove the effects of increased climate concerns. This result accords with that in Figure 6, in which the green factor's positive performance all but disappears after removing climate shocks.

Column 3 of Table 6 adds the two earnings-news variables for individual stocks. Not surprisingly, both are strongly related to individual stock returns. Even after controlling for these important drivers of individual stock returns, however, the climate-concern coefficients are little changed, and the coefficient on $g_{i,t-1}$ remains slightly negative. The same is true in column 4, which adds interactions between $g_{i,t-1}$ and the ESG flow and AUM variables defined earlier. Similar to before, we instrument for the interaction of $g_{i,t-1}$ and contemporaneous ESG flows by using $g_{i,t-1}$ interacted with lagged ESG flows. Consistent with the predicted signs, the flow interaction enters positively and the AUM interaction enters negatively, but neither is significant. These results mirror the time-series results in Table 4. Finally, adding book-to-market as a regressor in column 5 has negligible effects. Importantly, across columns 2 through 5, the climate-concern coefficients remain little changed, and the coefficient on $g_{i,t-1}$ remains slightly negative but insignificant.

Overall, the panel-regression analysis of individual stocks delivers the same message as the time-series analysis of the green factor: the outperformance of green stocks over brown is attributable entirely to climate-concern shocks.

¹⁴The sample shrinks as we move across the columns and add more regressors. These changes in sample are not responsible for the changes in g 's coefficient, however. In the Appendix, we hold the sample constant across the columns and show a similar pattern in g 's coefficients.

7. Timing of climate news and green-factor returns

In this section, we take a closer look at the strong relation between the green factor's returns and shocks to climate concerns. As shown previously in Section 5.3, this relation is positive but mostly asynchronous. Recall from Tables 4 and 5 that the green factor has a positive but insignificant relation to the current month's climate-concern shock, whereas its relation to the previous month's climate shock is significant and much stronger. To better understand this relation, we examine it at a higher data frequency.

First, we turn to the weekly frequency. The weekly change in climate concerns is well approximated by $MCCC_t$, the value of the MCCC index averaged across the days in week t .¹⁵ To create the weekly green-factor series, we compute a daily series by applying equation (3) at the daily frequency and then compounding those daily factor realizations within each week. A convenient by-product of this approach is the daily factor series, which we use later in this section. We estimate the slope coefficients in the time-series regression

$$GF_t = a + \sum_{\tau=0}^T \beta_{\tau} MCCC_{t-\tau} + e_t, \quad (9)$$

where GF_t is the green factor in week t and $T = 7$ weeks.

Figure 8 plots the estimated β_{τ} coefficients for lags $\tau = 0, 1, \dots, 7$ weeks, along with their 95% confidence intervals. The point estimate of β_0 is positive but far from significant. The estimates of $\beta_2, \beta_3, \beta_4$, and β_5 are all positive, though only β_4 is significant, marginally, at the 95% confidence level. The weekly relation is thus lagged, peaking at the lag of four weeks. This evidence is consistent with our prior evidence based on monthly data. Stock prices seem to incorporate shifts in climate concerns with a multi-week delay.

Second, we conduct the same analysis at the daily frequency. We measure the daily change in climate concerns by $MCCC_t$, the value of the MCCC index on day t , with an adjustment for non-trading days. On such days, stock returns are missing, but the MCCC index is available. Since news released during non-trading days gets into stock prices on the first subsequent trading day, we define $MCCC_t$ to be the sum of the MCCC index values since the end of the previous trading day. For example, $MCCC_t$ on a Monday is the sum of the raw MCCC values from Saturday, Sunday, and Monday. We then run the time-series regression in equation (9), except that t is now measured in days, with $T = 5$ days.

¹⁵The weekly value of ρ implied by the monthly value of 0.94 is 0.985. As a result, $1 - \rho \approx 0$ and the second term in equation (5) is small. With $T = 3 \times 52 = 156$ weeks, $\rho^T = 0.985^{156} \approx 0.09$, so the third term in equation (5) is also small, implying $\Delta C_t \approx MCCC_t$. We use the same approximation later at the daily frequency, where it is even more precise because the daily $\rho = 0.998$ and $T = 3 \times 365 = 1,095$ days.

Figure 9 plots the estimates of β_τ for lags of $\tau = 0, 1, \dots, 5$ days, along with their 95% confidence intervals. Unlike in weekly data, the positive estimate of β_0 is now statistically significant. This evidence is consistent with that of Ardia et al. (2021), who also report a positive contemporaneous daily relation between their MCCC index and green stock returns. This is not a given because Ardia et al. measure stocks' greenness differently, by using firms' self-reported greenhouse gas emissions as reported by the Asset4/Refinitiv database. Ardia et al.'s stock universe is also smaller than ours as they focus on the S&P 500 firms.

Figure 9 also shows negative estimates of β_1 and β_2 , though only that of β_1 is significant, and marginally so. The estimated magnitudes of β_1 and β_2 add up approximately to that of β_0 , indicating that the negative relations estimated at lags of one and two days offset the positive contemporaneous relation. These results explain why we do not observe significant contemporaneous relations in weekly and monthly data. A natural interpretation of the daily results is that trading driven by same-day climate news exerts temporary price pressure that is fully reversed within two days.

8. Conclusion

Realized returns are a popular proxy for expected returns in the empirical asset pricing literature. However, high realized returns do not always indicate high expected returns, especially if they are realized over a relatively short period. We offer the salient example of green assets over the past decade. We show that green assets' high recent returns are unexpected, reflecting news about environmental concerns rather than high expected returns. After constructing a theoretically motivated green factor from U.S. stock data, we show that the factor's recent outperformance vanishes after removing the effects of climate-concern shocks. Surprisingly, those shocks get reflected in the green factor's returns with a multi-week delay. We also find that a two-factor asset pricing model featuring the green factor absorbs much of the historic underperformance of value stocks in the 2010s.

Realized asset returns are notoriously noisy, and much of their volatility is viewed as inexplicable (e.g., Roll, 1984). We explain some of it for the green factor by linking the factor's returns to various proxies for shocks to green demands, such as a text-based measure of climate concerns, two measures of green-versus-brown earnings news, and sustainable fund flows. Additional proxies can be considered by future studies, with the aim of explaining a larger fraction of the green factor's returns. Future work can also apply our approach to other aspects of ESG investing, within various equity styles, and for other asset classes.

Our results contain a warning for investigations of climate-risk pricing. We find that green stocks typically outperform brown when climate concerns increase. This result echoes similar findings by Choi, Gao, and Jiang (2020), Engle et al. (2020), and Ardia et al. (2021). Equilibrium expected returns of stocks that are better hedges against adverse climate shocks include a negative hedging premium if the representative investor is averse to such shocks (e.g., PST). Empirically confirming a climate risk premium, however, must confront the large unanticipated positive component of green stock returns during the last decade. Without accounting for those unexpectedly high returns on stocks that appear to be relatively good climate hedges, one could be led astray. That is, one could infer that stocks providing better climate hedging have higher expected returns, not lower as theory predicts.




Figure 1. German twin bonds. Panel A plots the daily time series of the “greenium,” the difference between the yields of the German government’s 10-year green bond and its non-green twin, in annual terms. Panel B plots the performance of a portfolio that goes long the 10-year green bond and short its non-green twin. The solid line plots this long-short portfolio’s daily cumulative realized return. The dashed line plots the expected cumulative return as of the first day of trading of the green bond (September 8, 2020), absent a subsequent change in the greenium, which was -1.6 bps on that day.

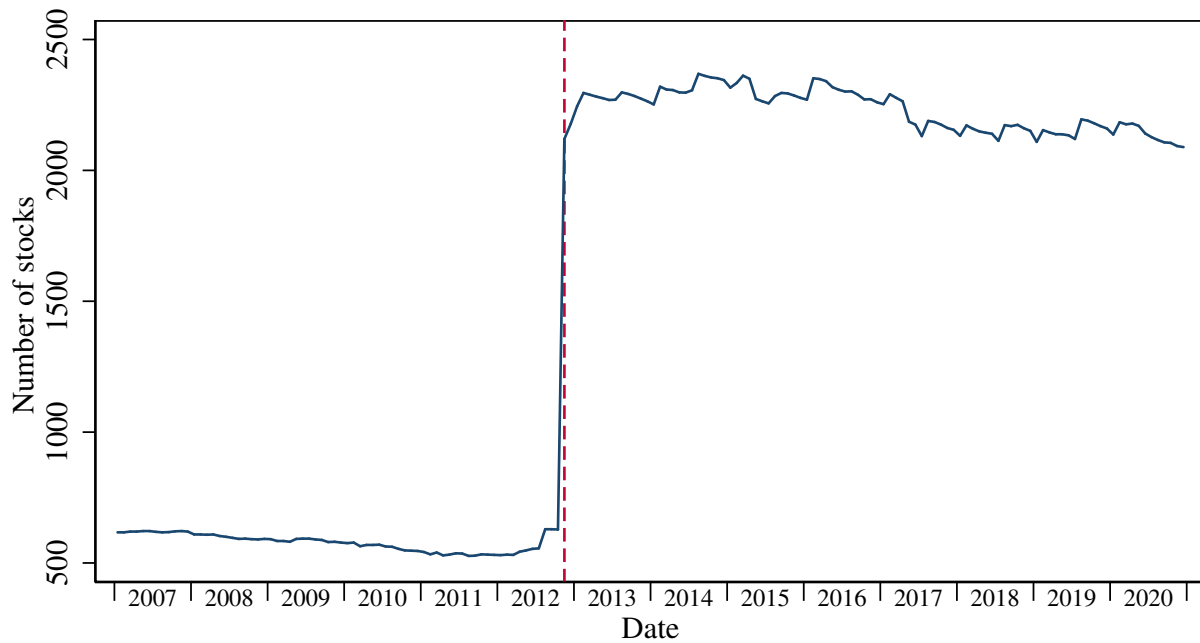


Figure 2. MSCI coverage. The figure plots the number of stocks in our sample with non-missing MSCI environmental scores at the beginning of the month. The dashed red line is at November 2012, where our sample begins. MSCI expanded its coverage in October 2012. We begin our sample in November 2012, as we require lagged environmental scores.




Figure 3. Cumulative return of the green factor. This figure plots the cumulative return of the green factor, which is computed from equation (3) on a monthly basis.




Figure 4. Climate concerns and the green factor. The level of climate concerns is computed as $C_t = \sum_{\tau=0}^{36} \rho^\tau MCCC_{t-\tau}$, where $MCCC_t$ is the monthly measure of Ardia et al. (2021), and $\rho = 0.94$.




Figure 5. Counterfactual green-factor returns. The solid line shows realized cumulative, compounded green-factor returns. The dashed line shows its counterfactual counterpart computed from column 2 of Table 4. The counterfactual monthly green-factor return equals its realized value minus the regressors times their respective regression slopes. Dotted lines indicate the counterfactual's 95% confidence interval. We compute confidence intervals using the following steps: (1) Estimate the regression from column 2 of Table 4 and store the estimated coefficients and their covariance matrix. (2) Repeat the following steps (2a)–(2c) 500 times: (2a) draw a new coefficient vector from a normal distribution with mean and variance saved in step (1); (2b) use the new coefficient to compute each period's counterfactual return; (2c) compute and store cumulative counterfactual returns. (3) Each month, compute the 2.5th and 97.5th percentiles of the counterfactual cumulative returns stored in step (2c).




Figure 6. Components of green-factor returns. The solid line plots the realized cumulative, compounded green-factor returns. The remaining lines show counterfactual green-factor returns computed using the model from column 3 of Table 4. To create the line “(1) No climate news,” we compute counterfactual monthly green-factor returns as their realized value minus the values of “ Δ Climate concerns (same month)” and “ Δ Climate concerns (prev. month)” times their respective regression coefficients. To create the line “(2) Also no CF news,” we use the previous counterfactual returns but also subtract “Earnings announcement returns” and “ Δ Earnings forecasts” times their respective regression coefficients. To create the line “(3) Also no ESG flows,” we use the previous counterfactual returns but also subtract “ESG flows” times its regression coefficient and [“ESG assets” minus counterfactual ESG assets] times the coefficient on “ESG assets,” where counterfactual ESG assets is computed as in footnote 13.

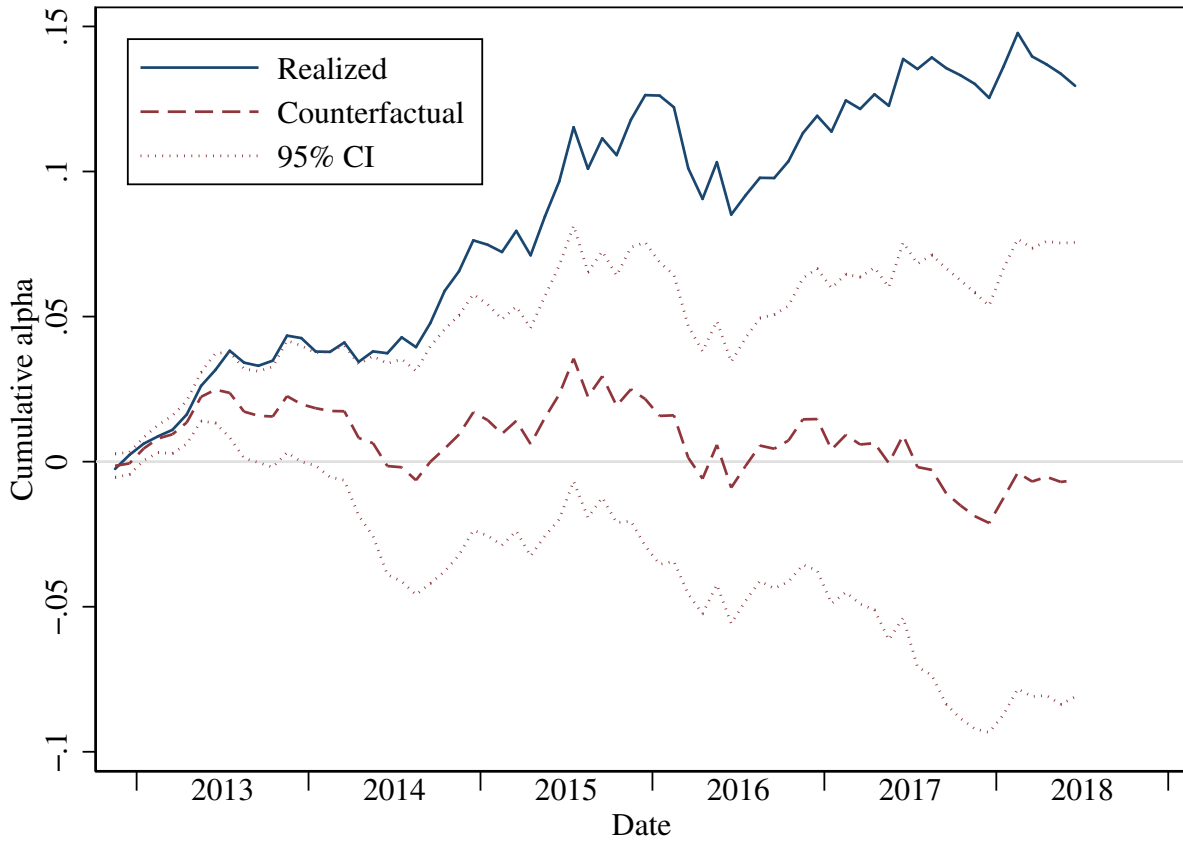


Figure 7. Counterfactual green-factor alpha. This is the same as Figure 5 but replaces the green factor with its Fama-French three-factor alphas, computed as in Table 5.




Figure 8. Weekly response of the green factor to climate news. This figure plots the β_τ coefficients from the weekly time-series regression in equation (9). The sample runs from November 2012 to June 2018. Dashed lines indicate 95% confidence intervals.




Figure 9. Daily response of the green factor to climate news. This figure plots the β_τ coefficients from the daily time-series regression in equation (9). The sample runs from November 2012 to June 2018. Dashed lines indicate 95% confidence intervals.

Table 1
Industries ranked by environmental scores

Average g is the environmental score averaged across firms within each MSCI industry at the end of 2019. MSCI uses the GICS sub-industry classification.

Rank	MSCI Industry	Average g	Rank	MSCI Industry	Average g
1	Asset Management & Custody Banks	0.870	33	Textiles, Apparel & Luxury Goods	-0.502
2	Professional Services	0.850	34	Auto Components	-0.505
3	Telecommunication Services	0.841	35	Property & Casualty Insurance	-0.506
4	Consumer Finance	0.837	36	Casinos & Gaming	-0.542
5	Health Care Equipment & Supplies	0.835	37	Real Estate Development	-0.548
6	Health Care Providers & Services	0.825	38	Semiconductors	-0.657
7	Life & Health Insurance	0.761	39	Electrical Equipment	-0.750
8	Interactive Media & Services	0.736	40	Construction & Farm Machinery	-0.758
9	Diversified Financials	0.732	41	Tobacco	-0.885
10	Media & Entertainment	0.704	42	Trading Companies & Distributors	-0.987
11	Diversified Consumer Services	0.614	43	Industrial Machinery	-1.040
12	Biotechnology	0.567	44	Containers & Packaging	-1.091
13	Pharmaceuticals	0.489	45	Energy Equipment & Services	-1.159
14	Multi-Line Insurance & Brokerage	0.405	46	Real Estate Management & Services	-1.198
15	Investment Banking & Brokerage	0.387	47	Airlines	-1.214
16	Banks	0.348	48	Hotels & Travel	-1.566
17	Restaurants	0.309	49	Building Products	-1.620
18	Construction & Engineering	0.125	50	Utilities	-1.903
19	Aerospace & Defense	0.097	51	Integrated Oil & Gas	-2.008
20	Commercial Services & Supplies	0.069	52	Food Products	-2.019
21	Air Freight & Logistics	-0.055	53	Beverages	-2.044
22	Household Durables	-0.116	54	Metals and Mining, Precious	-2.193
23	Software & Services	-0.130	55	Oil & Gas Refining, Marketing	-2.522
24	Electronic Equipment, Instruments	-0.170	56	Construction Materials	-2.556
25	Leisure Products	-0.173	57	Specialty Chemicals	-2.818
26	Automobiles	-0.215	58	Marine Transport	-2.828
27	Retail - Food & Staples	-0.251	59	Paper & Forest Products	-2.930
28	Retail - Consumer Discretionary	-0.269	60	Metals and Mining, Non-Precious	-2.947
29	Road & Rail Transport	-0.299	61	Steel	-2.955
30	Household & Personal Products	-0.300	62	Oil & Gas Exploration & Production	-3.010
31	Industrial Conglomerates	-0.364	63	Diversified Chemicals	-3.212
32	Technology Hardware, Storage	-0.391	64	Commodity Chemicals	-3.783

Table 2
Green-factor performance

We estimate monthly time-series regressions using data from November 2012 to December 2020. The dependent variable is the green factor. Mkt-Rf is the excess market return. SMB and HML are the size and value factors of Fama and French (1993). UMD is the momentum factor of Carhart (1997). LIQ is the traded liquidity factor of Pástor and Stambaugh (2003). RMW and CMA are the profitability and investment factors of Fama and French (2015). Returns are in percent per month. Robust t -statistics are in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.312 (2.91)	0.371 (3.38)	0.205 (2.48)	0.180 (2.46)	0.202 (2.42)	0.209 (2.60)
Mkt-RF		-0.0474 (-1.28)	0.00648 (0.17)	0.0374 (1.19)	0.00925 (0.29)	0.00552 (0.17)
SMB			-0.121 (-3.77)	-0.0977 (-3.54)	-0.107 (-2.24)	-0.169 (-4.02)
HML			-0.198 (-6.42)	-0.120 (-3.78)	-0.196 (-5.70)	-0.150 (-4.49)
UMD				0.127 (4.57)		
LIQ					-0.0229 (-0.38)	
RMW						-0.130 (-1.62)
CMA						-0.138 (-2.36)
Observations	98	98	98	98	98	98
R^2	0.000	0.034	0.408	0.519	0.411	0.464

Table 3
Pricing value and momentum in the green-factor model

We estimate monthly time-series regressions of either HML (in columns 1 and 2) or UMD (in columns 3 and 4) on the excess market return and the green factor by using data from November 2012 to December 2020. Returns are in percent per month. Robust t -statistics are in parentheses.

	Value		Momentum	
Constant	-0.709 (-1.93)	-0.151 (-0.50)	0.663 (1.92)	-0.0640 (-0.22)
Mkt-RF	0.139 (1.18)	0.0678 (0.70)	-0.368 (-3.75)	-0.275 (-3.14)
Green factor		-1.503 (-4.55)		1.960 (6.18)
Observations	98	98	98	98
R^2	0.041	0.345	0.173	0.487

Table 4
Sources of green-factor returns

We estimate monthly time-series regressions using data from November 2012 to June 2018. The dependent variable is the green factor. “ Δ Climate concerns” is a monthly change in the level of climate concerns, computed as in equation (5). The two earnings-news measures, “Earnings announcement returns” and “ Δ Earnings forecasts,” are described in Section 5.2. They correspond to the quarter that contains the given month. “ESG flows” equals the quarter’s dollar flow into ESG funds scaled by the average total CRSP market capitalization during the quarter that contains the given month, times 1000. We instrument for contemporaneous ESG flow by using its previous-quarter value. The first-stage t -statistic for lagged flows is 3.23. “ESG assets” equals total AUM in ESG funds scaled by the total CRSP market capitalization and measured at the beginning of the quarter containing the given month, times 1000. Robust t -statistics are in parentheses.

	(1)	(2)	(3)
Δ Climate concerns (same month)	0.00637 (0.95)	0.00328 (0.49)	0.00357 (0.54)
Δ Climate concerns (prev. month)	0.0235 (2.85)	0.0211 (2.52)	0.0212 (2.59)
Earnings announcement returns		0.558 (0.98)	0.509 (0.87)
Δ Earnings forecasts		0.227 (0.41)	0.260 (0.42)
ESG flows			0.0429 (0.46)
ESG assets			-0.00157 (-0.59)
Constant	0.000132 (0.11)	0.000289 (0.22)	0.00195 (0.37)
Observations	68	68	68
R^2	0.171	0.190	0.181

Table 5
Sources of green-factor alpha

This is the same as Table 4, except the dependent variable is the green factor's Fama-French three-factor alpha. We estimate these alphas in time-series regressions of the monthly green factor on the Fama-French factors, using data from November 2012 to June 2018. We set each month's alpha equal to the estimated intercept plus residual.

	(1)	(2)	(3)
Δ Climate concerns (same month)	0.00730 (1.34)	0.00583 (1.08)	0.00498 (0.85)
Δ Climate concerns (prev. month)	0.0183 (3.32)	0.0170 (3.03)	0.0168 (3.06)
Earnings announcement returns		0.219 (0.53)	0.307 (0.66)
Δ Earnings forecasts		0.184 (0.39)	0.0987 (0.21)
ESG flows			-0.0103 (-0.12)
ESG assets			-0.00111 (-0.48)
Constant	-0.000173 (-0.18)	-0.0000734 (-0.07)	0.00307 (0.70)
Observations	68	68	68
R^2	0.187	0.194	0.193

Table 6
Greenness and individual stock returns

This table shows results from panel regressions in which the dependent variable is stock i 's return in month t . $g_{i,t-1}$ is the stock's lagged greenness. ΔC_t is month t 's change in aggregate climate concerns, computed from equation (5). “[Earnings announcement ret.] $_{i,t}$ ” is the stock's sum of the three-trading-day excess returns (stock minus market) around earnings announcements and management earnings forecasts (if available) during the quarter containing month t . “[Δ Earnings forecast] $_{i,t}$ ” is the change in analysts' mean long-term earnings growth rate forecast for stock i during the quarter containing month t . “[ESG flows] $_t$ ” is the flow into ESG funds scaled by the total CRSP market capitalization during the quarter that contains month t , times 1000. We instrument for $g_{i,t-1} \times [\text{ESG flows}]_t$ by using $g_{i,t-1}$ times scaled ESG flows from the previous quarter. The first-stage t -statistics for the instrument in columns 4 and 5 are 3.61 and 5.27, respectively. “[ESG assets] $_{t-1}$ ” equals total ESG AUM scaled by CRSP at the end of the previous quarter, times 1000. We subtract from ESG assets a constant equal to the counterfactual ESG assets averaged across regression observations. Subtracting this constant affects the coefficient on $g_{i,t-1}$ but not the coefficient on $g_{i,t-1} \times [\text{ESG assets}]_{t-1}$. BE/ME is lagged at least six months. The sample begins in November 2012. All regressions include month fixed effects, cluster by month, and use robust standard errors.

	(1)	(2)	(3)	(4)	(5)
$g_{i,t-1}$	0.00213 (2.24)	-0.0000103 (-0.01)	-0.000267 (-0.27)	-0.00309 (-0.84)	-0.00416 (-0.85)
$g_{i,t-1} \times \Delta C_t$		0.00769 (1.15)	0.00802 (1.36)	0.00830 (1.31)	0.00806 (1.15)
$g_{i,t-1} \times \Delta C_{t-1}$		0.0166 (2.21)	0.0148 (2.24)	0.0159 (2.30)	0.0168 (2.29)
[Earnings announcement ret.] $_{i,t}$			0.320 (13.14)	0.320 (13.14)	0.315 (12.36)
[Δ Earnings forecast] $_{i,t}$			0.0592 (5.02)	0.0596 (5.08)	0.0587 (4.45)
$g_{i,t-1} \times [\text{ESG flows}]_t$				0.0753 (0.79)	0.0813 (0.77)
$g_{i,t-1} \times [\text{ESG assets}]_{t-1}$				-0.00160 (-0.58)	-0.000847 (-0.33)
$\ln(\text{BE/ME})_{i,t-1}$					-0.000741 (-0.52)
Observations	218,208	151,294	131,689	131,689	114,320

REFERENCES

- Ardia, David, Keven Bluteau, Kris Boudt, and Koen Inghelbrecht, 2021, Climate change concerns and the performance of green versus brown stocks, Working paper, National Bank of Belgium.
- Avramov, Doron, Si Cheng, Abraham Lioui, and Andrea Tarelli, 2021, Sustainable investing with ESG rating uncertainty, Working paper, IDC.
- Baker, Malcolm, Daniel Bergstresser, George Serafeim, and Jeffrey Wurgler, 2018, Financing the response to climate change: The pricing and ownership of U.S. green bonds, Working paper, Harvard University.
- Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis, 2020, Does climate change affect real estate prices? Only if you believe in it, *Review of Financial Studies* 33, 1256–1295.
- Berg, Florian, Julian F. Koelbel, and Roberto Rigobon, 2020, Aggregate confusion: The divergence of ESG ratings, Working paper, MIT.
- Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis, 2019, Disaster on the horizon: The price effect of sea level rise, *Journal of Financial Economics* 134, 253–272.
- Beyer, Anne, Daniel A. Cohen, Thomas Z. Lys, and Beverly R. Walther, 2010, The financial reporting environment: Review of the recent literature, *Journal of Accounting and Economics* 50, 296–343.
- BlackRock, 2020, Global Sustainable Investing Survey.
- Bolton, Patrick, and Marcin Kacperczyk, 2020, Do investors care about carbon risk? *Journal of Financial Economics*, forthcoming.
- Bolton, Patrick, and Marcin Kacperczyk, 2021, Global pricing of carbon-transition risk, NBER working paper.
- Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Chava, Sudheer, 2014, Environmental externalities and cost of capital, *Management Science* 60, 2223–2247.
- Choi, Darwin, Zhenyu Gao, and Wenxi Jiang, 2020, Attention to global warming, *Review of Financial Studies* 33, 1112–1145.
- Eccles, Robert G., and Judith C. Strohle, 2018, Exploring social origins in the construction of ESG measures, Working paper, University of Oxford.
- El Ghouli, Sadok, Omrane Guedhami, Chuck CY Kwok, and Dev R. Mishra, 2011, Does corporate social responsibility affect the cost of capital?, *Journal of Banking and Finance* 35, 2388–2406.
- Engle, Robert F., Stefano Giglio, Bryan T. Kelly, Heebum Lee, and Johannes Stroebel, 2020, Hedging climate change news, *Review of Financial Studies* 33, 1184–1216.

- Faccini, Renato, Rastin Matin, and George S. Skiadopoulos, 2021, Dissecting climate risks: Are they reflected in stock prices? Working paper, University of London.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F., and Kenneth R. French, 2002, The equity premium, *Journal of Finance* 57, 637–659.
- Edmans, Alex, Does the stock market fully value intangibles? Employee satisfaction and equity prices, *Journal of Financial Economics* 101, 621–640.
- Fama, Eugene F., and Kenneth R. French, 2007, Disagreement, tastes, and asset prices, *Journal of Financial Economics* 83, 667–689.
- Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.
- Fink, Larry, 2021, BlackRock CEO’s annual letter to shareholders. <https://www.blackrock.com/corporate/investor-relations/larry-fink-chairmans-letter>.
- Giglio, Stefano, Bryan T. Kelly, and Johannes Stroebe, 2020, Climate Finance, *Annual Review of Financial Economics*, forthcoming.
- Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebe, and Andreas Weber, 2021, Climate change and long-run discount rates: Evidence from real estate, *Review of Financial Studies*, forthcoming.
- Goldsmith-Pinkham, Paul S., Matthew Gustafson, Ryan Lewis, and Michael Schwert, 2021, Sea level rise exposure and municipal bond yields, Working paper, Yale University.
- Görgen, Maximilian, Andrea Jacob, Martin Nerlinger, Ryan Riordan, Martin Rohleder, and Marco Wilkens, 2020, Carbon risk, Working paper, University of Augsburg.
- Heinkel, Robert, Alan Kraus, and Josef Zechner, 2001, The effect of green investment on corporate behavior, *Journal of Financial and Quantitative Analysis* 36, 431–449.
- Hong, Harrison, Frank Weikai Li, and Jiangmin Xu, 2019, Climate risks and market efficiency, *Journal of Econometrics* 208, 265–281.
- Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, 2020, The pollution premium. Working paper, HKUST.
- Huynh, Thanh, and Ying Xia, 2021, Climate change news risk and corporate bond returns, *Journal of Financial and Quantitative Analysis*, forthcoming.
- Ilhan, Emirhan, Zacharias Sautner, and Grigory Vilkov, 2021, Carbon tail risk, *Review of Financial Studies* 34, 15401571.
- In, Soh Young, Ki Young Park, and Ashby Monk, 2019, Is ‘being green’ rewarded in the market? An empirical investigation of decarbonization and stock returns, Working paper, Stanford University.
- Krueger, Philipp, Zacharias Sautner, and Laura T. Starks, 2020, The importance of climate risks for institutional investors, *Review of Financial Studies* 33, 1067–1111.

- Larcker, David F., and Edward M. Watts, 2020, Where's the greenium? *Journal of Accounting and Economics* 69, 1–26.
- Lester, Anna, and Chen He, 2018, Harnessing ESG as an alpha source in active quantitative equities, State Street Global Advisors white paper.
- Nagy, Zoltán, Altaf Kassam, and Lenda-Eling Lee, 2016, Can ESG add alpha? An analysis of ESG tilt and momentum strategies, *Journal of Investing* 25 (2), 113–124.
- Painter, Marcus, 2020, An inconvenient cost: The effects of climate change on municipal bonds, *Journal of Financial Economics* 135, 468–482.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2001, The equity premium and structural breaks, *Journal of Finance* 56, 1207–1239.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2009, Predictive systems: Living with imperfect predictors, *Journal of Finance* 64, 1583–1628.
- Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics*, forthcoming.
- Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski, 2021, Responsible investing: The ESG-efficient frontier, *Journal of Financial Economics*, forthcoming.
- Roll, Richard, 1984, Orange juice and weather, *American Economic Review* 74, 861–880.
- Seltzer, Lee, Laura T. Starks, and Qifei Zhu, 2021, Climate regulatory risks and corporate bonds, Working paper, University of Texas.
- US SIF Foundation, 2020, Report on US Sustainable and Impact Investing Trends.
- Zerbib, Olivier David, 2019, The effect of pro-environmental preferences on bond prices: Evidence from green bonds, *Journal of Banking and Finance* 98, 39–60.