

DISCUSSION PAPER SERIES

DP14171
(v. 3)

SUSTAINABLE INVESTING IN EQUILIBRIUM

Luboš Pástor, Robert F. Stambaugh and Lucian
Taylor

FINANCIAL ECONOMICS



SUSTAINABLE INVESTING IN EQUILIBRIUM

Luboš Pástor, Robert F. Stambaugh and Lucian Taylor

Discussion Paper DP14171
First Published 06 December 2019
This Revision 04 June 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Financial Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Luboš Pástor, Robert F. Stambaugh and Lucian Taylor

SUSTAINABLE INVESTING IN EQUILIBRIUM

Abstract

We model investing that considers environmental, social, and governance (ESG) criteria. In equilibrium, green assets have low expected returns because investors enjoy holding them and because green assets hedge climate risk. Green assets nevertheless outperform when positive shocks hit the ESG factor, which captures shifts in customers' tastes for green products and investors' tastes for green holdings. The ESG factor and the market portfolio price assets in a two-factor model. The ESG investment industry is largest when investors' ESG preferences differ most. Sustainable investing produces positive social impact by making firms greener and by shifting real investment toward green firms.

JEL Classification: G11, G12

Keywords: sustainable investing, ESG, Socially responsible investing, social impact

Luboš Pástor - lubos.pastor@chicagobooth.edu
University of Chicago and CEPR

Robert F. Stambaugh - stambaugh@wharton.upenn.edu
University of Pennsylvania

Lucian Taylor - luket@wharton.upenn.edu
University of Pennsylvania

Acknowledgements

We are grateful for comments from our discussants Bernard Dumas, Harrison Hong and Jacob Sagi, and also from Rui Albuquerque, Malcolm Baker, George Constantinides, Alex Edmans, Gene Fama, Sam Hartzmark, John Heaton, Ravi Jagannathan, Ralph Koijen, Yrjo Koskinen, Stavros Panageas, Raghu Rajan, Jeff Wurgler, Josef Zechner, conference participants at the 2020 Spring NBER Asset Pricing Meeting, the 2020 SFS Cavalcade, and the 2020 INSEAD Finance Symposium, and seminar participants at the University of Chicago, WU Vienna, and the National Bank of Slovakia.

Sustainable Investing in Equilibrium

Ľuboš Pástor,^{a, c, d, e} Robert F. Stambaugh,^{b, c} Lucian A. Taylor^{b*}

^a*University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637, USA*

^b*University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, USA*

^c*NBER, USA*

^d*CEPR, United Kingdom*

^e*National Bank of Slovakia*

June 4, 2020

Abstract

We model investing that considers environmental, social, and governance (ESG) criteria. In equilibrium, green assets have low expected returns because investors enjoy holding them and because green assets hedge climate risk. Green assets nevertheless outperform when positive shocks hit the ESG factor, which captures shifts in customers' tastes for green products and investors' tastes for green holdings. The ESG factor and the market portfolio price assets in a two-factor model. The ESG investment industry is largest when investors' ESG preferences differ most. Sustainable investing produces positive social impact by making firms greener and by shifting real investment toward green firms.

JEL classifications: G11, G12

Keywords: sustainable investing, socially responsible investing, ESG, social impact

* Corresponding author email: luket@wharton.upenn.edu. The views in this paper are the responsibility of the authors, not the institutions they are affiliated with. We are grateful for comments from our discussants Bernard Dumas, Harrison Hong and Jacob Sagi, and also from Rui Albuquerque, Malcolm Baker, George Constantinides, Alex Edmans, Gene Fama, Sam Hartzmark, John Heaton, Ravi Jagannathan, Ralph Koijen, Yrjo Koskinen, Stavros Panageas, Raghu Rajan, Jeff Wurgler, Josef Zechner, conference participants at the 2020 Spring NBER Asset Pricing Meeting, the 2020 SFS Cavalcade, and the 2020 INSEAD Finance Symposium, and seminar participants at the University of Chicago, WU Vienna, and the National Bank of Slovakia.

1. Introduction

Sustainable investing considers not only financial objectives but also environmental, social, and governance criteria. This investment approach initially gained popularity by imposing negative screens under the umbrella of socially responsible investing (SRI), but its scope has expanded significantly in recent years. Assets managed with an eye on sustainability have grown to tens of trillions of dollars and seem poised to grow further.¹ Given this rapid growth, the effects of sustainable investing on asset prices and corporate behavior are important to understand.

We analyze both financial and real effects of sustainable investing through the lens of an equilibrium model. The model features many heterogeneous firms and agents, yet it is highly tractable, yielding simple and intuitive expressions for the quantities of interest. The model illuminates the key channels through which agents' preferences for sustainability can move asset prices, tilt portfolio holdings, determine the size of the ESG investment industry, and cause real impact on society.

In the model, firms differ in the sustainability of their activities. “Green” firms generate positive externalities for society, “brown” firms impose negative externalities, and there are different shades of green and brown. Agents differ in their preferences for sustainability, or “ESG preferences,” which have multiple dimensions. First, agents derive utility from holdings of green firms and disutility from holdings of brown firms. Second, agents care about firms' aggregate social impact. In a model extension, agents additionally care about climate risk. Naturally, agents also care about financial wealth.

We show that agents' tastes for green holdings affect asset prices. Agents are willing to pay more for greener firms, thereby lowering the firms' costs of capital. Green assets have negative CAPM alphas, whereas brown assets have positive alphas. Consequently, agents with stronger ESG preferences, whose portfolios tilt more toward green assets and away from brown assets, earn lower expected returns. Yet such agents are not unhappy because they derive utility from their holdings.

The model implies three-fund separation, whereby each agent holds the market portfolio, the risk-free asset, and an “ESG portfolio” whose composition depends on assets' greenness. Agents with stronger than average tastes for green holdings deviate from the market largely

¹According to the 2018 Global Sustainable Investment Review, sustainable investing assets exceeded \$30 trillion globally at the start of 2018, a 34% increase in two years. As of November 2019, more than 2,600 organizations have become signatories to the United Nations Principles of Responsible Investment (PRI), with more than 500 new signatories in 2018/2019, according to the 2019 Annual Report of the PRI.

by overweighting green assets and underweighting brown ones. Agents with weaker ESG tastes deviate in the opposite direction, and agents with average tastes hold the market portfolio. If there is no dispersion in ESG tastes, all agents simply hold the market. Even if all agents derive a large amount of utility from green holdings, they nevertheless hold only the market if their ESG tastes are equally strong, because asset prices then fully adjust to reflect those tastes. For the ESG industry to exist, dispersion in ESG tastes is necessary.

We define the “ESG factor” as a scaled return on the ESG portfolio. We show that the ESG factor and the market portfolio together price assets in a two-factor model. Assets’ loadings on the ESG factor, their “ESG betas,” equal their ESG characteristics: green assets have positive ESG betas and brown assets have negative betas. A simple version of the ESG factor is a green-minus-brown portfolio return, where both green and brown portfolios are weighted by ESG characteristics. Assets’ CAPM alphas reflect exposure to the omitted, priced ESG factor. The factor has a negative premium that comes from investors’ ESG tastes.

We interpret the ESG factor as capturing unexpected changes in ESG concerns. These concerns can change in two ways: customers can shift their demands for goods of green providers, and investors can change their appreciation for green holdings. The ESG factor affects the relative performance of green and brown assets; its positive realizations boost green assets while hurting brown ones. If ESG concerns strengthen unexpectedly and sufficiently, green assets outperform brown ones despite having lower expected returns.

To assess the model’s quantitative implications, we calibrate a setting with two types of investors: those sharing equal concerns about ESG (“ESG investors”) and those having no concerns (“non-ESG investors”). Given their portfolios’ green tilts, ESG investors earn lower expected returns than non-ESG investors. The difference in expected returns increases with λ , the wealth share of ESG investors, and with Δ , the maximum certain return ESG investors are willing to forgo in exchange for investing in their desired portfolio instead of the market. Non-ESG investors earn an alpha that is positive and increasing in both λ and Δ . ESG investors earn a negative alpha whose magnitude is increasing in Δ , concave in λ , and greatest when the dispersion in ESG tastes is greatest (i.e., $\lambda = 0.5$).

Despite earning a negative alpha, ESG investors enjoy an “investor surplus”: they sacrifice less return than they are willing to in order to hold their desired portfolio. The reason is that equilibrium asset prices adjust to ESG tastes, thereby pushing the market portfolio toward the portfolio desired by ESG investors. Specifically, ESG tastes make green firms more valuable and brown firms less. The market portfolio thus moves closer to ESG investors’

desired portfolio, pushing those investors' negative alpha closer to zero. For example, when ESG investors have $\Delta = 4\%$, their alpha is at least -2% . We define investor surplus to be the difference between alpha and $-\Delta$. The surplus is always positive, ranging from $\Delta/2$ to Δ .

We measure the size of the ESG investment industry by the aggregate ESG dollar tilt away from the market portfolio. The ESG industry is largest when the dispersion in ESG tastes is greatest. In addition, the ESG industry's size is reduced by the price adjustment mentioned above. For example, suppose that the ESG industry reaches 24% of the stock market's value when Δ is 1%. Then doubling the strength of ESG tastes by raising Δ to 2% increases that maximum industry size by less than half, to 35% of the market's value.

Our model implies that sustainable investing leads to positive social impact. We define a firm's social impact as the product of the firm's greenness and its scale. We show that agents' tastes for green holdings increase firms' social impact through two channels. First, firms choose to become greener, because greener firms have higher market values. Second, real investment shifts from brown to green firms, due to shifts in firms' cost of capital (up for brown firms, down for green firms). We obtain positive aggregate social impact even if agents have no direct preference for it, shareholders do not engage with management, and managers simply maximize market value.

Finally, we extend the model by allowing climate to enter investors' utility. Expected returns then depend not only on market betas and investors' tastes but also on climate betas, which measure firms' exposures to climate shocks. Evidence suggests that brown assets have higher climate betas than green assets (e.g., Choi, Gao, and Jiang, 2019; Engle et al., 2020). This difference pushes up brown assets' expected returns in our model. The idea is that investors dislike unexpected deteriorations in the climate. If the climate worsens unexpectedly, brown assets lose value relative to green assets (e.g., due to new government regulation that penalizes brown firms). Because brown firms lose value in states of the world investors dislike, they are riskier, so they must offer higher expected returns. Brown stocks thus have positive CAPM alphas not only because of investors' distastes for brown holdings, but also because of brown stocks' larger exposures to climate risk.

Our theoretical treatment of climate risk is related to recent empirical work on the implications of such risk for asset prices. Hong, Li, and Xu (2019) analyze the response of food producers' stock prices to climate risks. Bolton and Kacperczyk (2019) conclude that investors demand compensation for exposure to carbon risk in the form of higher returns on carbon-intensive firms. Ilhan, Sautner, and Vilkov (2020) show that firms with higher carbon

emissions exhibit more tail risk and more variance risk. Engle et al. (2020) develop a procedure to dynamically hedge climate risk by constructing mimicking portfolios that hedge innovations in climate news series obtained by textual analysis of news sources. Bansal, Ochoa, and Kiku (2016) identify climate change as a long-run risk factor. Krueger, Sautner, and Starks (2019) find that institutional investors consider climate risks to be important investment risks.

Besides climate risk, other aspects of ESG-related risk have been studied. Hoepner et al. (2018) find that ESG engagement reduces firms’ downside risk as well as their exposures to a downside-risk factor. Luo and Balvers (2017) find a premium for boycott risk. We complement these studies with a theoretical contribution. We construct an ESG risk factor that is capable of pricing assets in a two-factor model, and we show that green and brown assets have opposite exposures to this factor.

Prior studies report, in various contexts, that green assets underperform brown assets. Hong and Kacperczyk (2009) find that “sin” stocks (i.e., stocks of public firms producing alcohol, tobacco, and gaming, which we would classify as brown) outperform non-sin stocks. They argue that social norms lead investors to demand compensation for holding sin stocks. Barber, Morse, and Yasuda (2020) find that venture capital funds that aim not only for financial return but also for social impact earn lower returns than other funds. They argue that investors derive nonpecuniary utility from investing in dual-objective funds. Baker et al. (2018) and Zerbib (2019) find that green bonds tend to be priced at a premium, offering lower yields than traditional bonds. Both studies argue that the premium is driven by investors’ environmental concerns. Similarly, Chava (2014) and El Ghouli et al. (2011) find that greener firms have a lower implied cost of capital. All of these results are consistent with our prediction that ESG tastes reduce green firms’ costs of capital.

Some studies find the opposite result, that green assets outperform brown, using alternative definitions of green and brown. Firms perform better if they are better-governed, judging by employee satisfaction (Edmans, 2011) or by strong shareholder rights (Gompers, Ishii, and Metrick, 2003), or if they have higher ESG ratings in the 1992–2004 period (Kempf and Osthoff, 2007). These results are also consistent with our model as long as ESG tastes strengthen unexpectedly over the sample period. We do not mean to imply that we can always declare empirical success for our model. The model clearly predicts that green assets underperform brown over a sufficiently long period—a period long enough so that unexpected changes in ESG tastes average to zero. We simply explain why it is difficult to distinguish *ex ante* versus *ex post* effects of ESG concerns by looking at realized returns over periods during which ESG tastes shift. Disentangling alphas from ESG taste shifts is a major challenge for

empirical work in this area.

Our model is also related to previous theoretical studies of sustainable investing. Heinkel, Kraus, and Zechner (2001) build an equilibrium model in which exclusionary ethical investing affects firm investment. They consider two types of investors, one of which refuses to hold shares in polluting firms. The resulting reduction in risk sharing increases the cost of capital of polluting firms, depressing their investment. Albuquerque, Koskinen, and Zhang (2019) construct a model in which a firm’s socially responsible investments increase customer loyalty, giving the firm more pricing power. This power makes the firm less risky and thus more valuable. Unlike these models, ours features neither a lack of risk sharing nor pricing power; instead, the main force is investors’ tastes for holding green assets.

Fama and French (2007) argue that tastes for holding green assets can affect prices. Baker et al. (2018) build a model featuring two types of investors with mean-variance preferences, where one type also has tastes for green assets. Their model predicts that green assets have lower expected returns and more concentrated ownership, and they find support for these predictions in the universe of green bonds. Pedersen, Fitzgibbons, and Pomorski (2019) consider the same two types of mean-variance investors but also add a third type that is unaware of firms’ ESG scores. This lack of awareness is costly if firms’ ESG scores predict their profits. The authors show that stocks with higher ESG scores can have either higher or lower expected returns, depending on the wealth of the third type of investors. They obtain four-fund separation and derive the ESG-efficient frontier characterizing the tradeoff between the ESG score and the Sharpe ratio.

While the models in these studies share some features with ours, we offer novel insights. We show that an ESG factor, along with the market portfolio, prices assets in a two-factor model. Positive realizations of this factor, which result from shifts in customers’ and investors’ tastes, can result in green assets outperforming brown. The size of the ESG investment industry, as well as investors’ alphas, crucially depend on the dispersion in investors’ ESG tastes. ESG investors earn an investor surplus. We have a continuum of investors with multiple dimensions of ESG preferences. Including climate in those preferences results in the pricing of climate risk. Finally, ESG investing has positive social impact.

Positive social impact also emerges from the model of Oehmke and Opp (2020), but through a different channel. Key ingredients to generating impact in their model are financing constraints and coordination among agents. Our model does not include those ingredients, but it produces social impact nevertheless, through tastes for green holdings. To emphasize these tastes, we do not model shareholder engagement with management,

which is another channel through which ESG investing can potentially increase market value (e.g., Dimson, Karakas, and Li, 2015). In our model, value-maximizing managers make their firms greener voluntarily, without pressure from shareholders, because greener firms command higher market values.²

Our assumption that some investors derive nonpecuniary benefits from green holdings has considerable empirical support in the mutual fund literature. Mutual fund flows respond to ESG-salient information, such as Morningstar sustainability ratings (Hartzmark and Sussman, 2019) and environmental disasters (Bialkowski and Starks, 2016). Flows to SRI mutual funds are less volatile than flows to non-SRI funds (Bollen, 2007) and less responsive to negative past performance (Renneboog, ter Horst, and Zhang, 2011). Investors in SRI funds also indicate willingness to forgo financial performance to accommodate their social preferences (Riedl and Smeets, 2017).

This paper is organized as follows. Section 2 presents our baseline model. Section 3 discusses the ESG factor. Section 4 explores the model’s quantitative implications. Section 5 extends the baseline model by letting agents care about the climate, showing that climate risk commands a premium. Section 6 examines social impact. Section 7 concludes.

2. Model

The model considers a single period, from time 0 to time 1, in which there are N firms, $n = 1, \dots, N$. Let \tilde{r}_n denote the return on firm n ’s shares in excess of the riskless rate, r_f , and let \tilde{r} be the $N \times 1$ vector whose n th element is \tilde{r}_n . We assume \tilde{r} is normally distributed:

$$\tilde{r} = \mu + \tilde{\epsilon} , \tag{1}$$

where μ contains equilibrium expected excess returns and $\tilde{\epsilon} \sim N(0, \Sigma)$. In addition to financial payoffs, firms produce social impact. Each firm n has an observable “ESG characteristic” g_n , which can be positive (for “green” firms) or negative (for “brown” firms). Firms with $g_n > 0$ have positive social impact, meaning they generate positive externalities (e.g., cleaning up the environment). Firms with $g_n < 0$ have negative social impact, meaning they generate negative externalities (e.g., polluting the environment). In Section 6, we model firms’ social impact in greater detail.

²Theoretical work on sustainable investing also includes Friedman and Heinle (2016), Gollier and Pouget (2014), and Luo and Balvers (2017). Bank and Insam (2017) do not mention sustainable investing, but they model investors with preferences for other stock characteristics. Empirical work on sustainable investing includes Geczy, Stambaugh, and Levin (2005), Hong and Kostovetsky (2012), and Cheng, Hong, and Shue (2016), among others. For recent experimental work, see Humphrey et al. (2020). For surveys of the early literature, see Bauer, Koedijk, and Otten (2005) and Renneboog, ter Horst, and Zhang (2008).

There is a continuum of agents who trade firms' shares and the riskless asset. The riskless asset is in zero net supply, whereas each firm's stock is in positive net supply. Let X_i denote an $N \times 1$ vector whose n th element is the fraction of agent i 's wealth invested in stock n . Agent i 's wealth at time 1 is $\tilde{W}_{1i} = W_{0i}(1 + r_f + X_i' \tilde{r})$, where W_{0i} is the agent's initial wealth. Besides liking wealth, agents also derive utility from holding green stocks and disutility from holding brown stocks.³ Each agent i has exponential utility

$$V(\tilde{W}_{1i}, X_i) = -e^{-A_i \tilde{W}_{1i} - b_i' X_i}, \quad (2)$$

where A_i is the agent's absolute risk aversion and b_i is an $N \times 1$ vector of nonpecuniary benefits that the agent derives from her stock holdings. Holding the riskless asset brings no such benefit. The benefit vector has agent-specific and firm-specific components:

$$b_i = d_i g, \quad (3)$$

where g is an $N \times 1$ vector whose n th element is g_n and $d_i \geq 0$ is a scalar measuring the degree of agent i 's "ESG taste." Agent i thus derives a nonpecuniary benefit of $d_i g_n$ from holding stock n . Agents with higher values of d_i have stronger tastes for the ESG characteristics of their holdings. In addition to having ESG tastes, agents care about firms' aggregate social impact, but that component of preferences does not affect agents' portfolio choice or asset prices. Therefore, we postpone the discussion of that component until Section 6.3.

2.1. Expected returns

Due to their infinitesimal size, agents take asset prices (and thus the return distribution) as given when choosing their optimal portfolios at time 0. To derive the first-order condition for X_i , we compute the expectation of agent i 's utility in Eq. (2) and differentiate it with respect to X_i . As we show in the Appendix, agent i 's portfolio weights on the N stocks are

$$X_i = \frac{1}{a_i} \Sigma^{-1} \left(\mu + \frac{1}{a_i} b_i \right), \quad (4)$$

where $a_i \equiv A_i W_{0i}$ is agent i 's relative risk aversion. For tractability, we assume that $a_i = a$ for all agents. We define ω_i to be the ratio of agent i 's initial wealth to total initial wealth: $\omega_i \equiv W_{0i}/W_0$, where $W_0 = \int_i W_{0i} di$. Because we assume a zero aggregate position in the riskless asset, market clearing requires that w_m , the $N \times 1$ vector of weights in the market portfolio of stocks, satisfies

$$\begin{aligned} w_m &= \int_i \omega_i X_i di \\ &= \frac{1}{a} \Sigma^{-1} \mu + \frac{\bar{d}}{a^2} \Sigma^{-1} g, \end{aligned} \quad (5)$$

³We frame the discussion in terms of green and brown stocks, but our main ideas apply more broadly to any set of green and brown assets, such as bonds and private equity investments.

where $\bar{d} \equiv \int_i \omega_i d_i di \geq 0$ is the wealth-weighted mean of ESG tastes d_i across agents and $\iota' w_m = 1$, with ι denoting an $N \times 1$ vector of ones. Note that $\bar{d} > 0$ unless the mass of agents who care about ESG is zero. Solving for μ gives

$$\mu = a \Sigma w_m - \frac{\bar{d}}{a} g . \quad (6)$$

Premultiplying by w'_m gives the market equity premium, $\mu_m = w'_m \mu$:

$$\mu_m = a \sigma_m^2 - \frac{\bar{d}}{a} w'_m g , \quad (7)$$

where $\sigma_m^2 = w'_m \Sigma w_m$ is the variance of the market return. In general, the equity premium depends on the average of ESG tastes, \bar{d} , through $w'_m g$, which is the overall “greenness” of the market portfolio. If the market is net green ($w'_m g > 0$) then stronger ESG tastes (i.e., larger \bar{d}) reduce the equity premium. If the market is net brown ($w'_m g < 0$), stronger ESG tastes increase the premium as investors demand compensation for this brownness. For simplicity, we assume that the market portfolio is ESG-neutral,

$$w'_m g = 0 , \quad (8)$$

which implies that the equity premium in Eq. (7) is independent of agents’ ESG tastes. Equivalently, we could view g as being defined so that agents derive utility (disutility) from holdings that are greener (brownier) than the market. Eqs. (7) and (8) imply $a = \mu_m / \sigma_m^2$. Combining this with Eq. (6) and noting that the vector of market betas is $\beta_m = (1/\sigma_m^2) \Sigma w_m$, we obtain our first proposition.

Proposition 1. *Expected excess returns in equilibrium are given by*

$$\mu = \mu_m \beta_m - \frac{\bar{d}}{a} g . \quad (9)$$

We see that expected excess returns deviate from their CAPM values, $\mu_m \beta_m$, due to ESG tastes for holding green stocks.

Corollary 1. *If $\bar{d} > 0$, the expected return on stock n is decreasing in g_n .*

As long as the mass of agents who care about sustainability is nonzero, \bar{d} is positive, and expected returns are decreasing in ESG characteristics. Given their ESG tastes, agents are willing to pay more for greener firms, thereby lowering the firms’ expected returns. Because the vector of stocks’ CAPM alphas is defined as $\alpha \equiv \mu - \mu_m \beta_m$, Eq. (9) yields the following corollary.

Corollary 2. *The CAPM alpha of stock n is given by*

$$\alpha_n = -\frac{\bar{d}}{a}g_n . \quad (10)$$

If $\bar{d} > 0$, green stocks have negative alphas, and brown stocks have positive alphas. Moreover, greener stocks have lower alphas.

As long as some agents care about sustainability, Eq. (10) implies that the alphas of stocks with $g_n > 0$ are negative, the alphas of stocks with $g_n < 0$ are positive, and α_n is decreasing with g_n . Furthermore, the negative relation between α_n and g_n is stronger when risk aversion, a , is lower and when the average ESG taste, \bar{d} , is higher.⁴

Proposition 2. *The mean and variance of the excess return on agent i 's portfolio are*

$$\mathbb{E}(\tilde{r}_i) = \mu_m - \delta_i \left(\frac{\bar{d}}{a^3} g' \Sigma^{-1} g \right) \quad (11)$$

$$\text{Var}(\tilde{r}_i) = \sigma_m^2 + \delta_i^2 \left(\frac{1}{a^4} g' \Sigma^{-1} g \right) , \quad (12)$$

where $\delta_i \equiv d_i - \bar{d}$.

Both equations are derived in the Appendix. Agents with $\delta_i > 0$ accept below-market expected returns in exchange for satisfying their stronger tastes for holding green stocks. As a result, agents whose tastes for green holdings are weaker ($\delta_i < 0$) enjoy above-market expected returns. In departing from market holdings, all agents with $\delta_i \neq 0$ incur higher return volatility than that of the market portfolio.

Corollary 3. *If $\bar{d} > 0$ and $g \neq 0$, agents with larger δ_i earn lower expected returns.*

Under the conditions of this corollary, the term in parentheses in Eq. (11) is strictly positive. Therefore, agents with stronger ESG tastes (i.e., larger δ_i) earn lower expected returns. The effect of δ_i on $\mathbb{E}(\tilde{r}_i)$ is stronger when the average ESG taste is stronger (i.e., when \bar{d} is larger), when risk aversion a is smaller, and when $g' \Sigma^{-1} g$ is larger.

The low expected returns earned by ESG-sensitive agents do not imply these agents are unhappy. As we show in the Appendix, agent i 's expected utility in equilibrium is given by

$$\mathbb{E} \left\{ V(\tilde{W}_{1i}) \right\} = \bar{V} e^{-\frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g} , \quad (13)$$

where \bar{V} is the expected utility if the agent has $\delta_i = 0$. Expected utility is increasing in δ_i^2 (note from Eq. (2) that $\bar{V} < 0$), so the more an agent's ESG taste d_i deviates from the average in either direction, the more ESG preferences contribute to the agent's utility.

⁴Proposition 1 and its corollaries continue to hold if agents disagree on stocks' ESG characteristics, g_n . In that case, the results hold with g_n replaced by the wealth-weighted average of agents' perceived values of g_n , adjusted for any covariance between those perceived values and ESG tastes. See the Appendix.

2.2. Portfolio tilts and the ESG portfolio

Substituting for μ from Eq. (9) into Eq. (4), we obtain an agent's portfolio weights:

Proposition 3. *Agent i 's equilibrium portfolio weights on the N stocks are given by*

$$X_i = w_m + \left(\delta_i/a^2\right) \Sigma^{-1}g . \quad (14)$$

Proposition 3 implies three-fund separation, as each agent's overall portfolio can be implemented with three assets: the riskless asset, the market portfolio, and an "ESG portfolio" whose weights are proportional to $\Sigma^{-1}g$. The fraction of agent i 's wealth in the riskless asset, $1 - \iota'X_i = -(\delta_i/a^2) \iota'\Sigma^{-1}g$, can be positive or negative. The agent's remaining wealth is invested in stocks. Specifically, the agent allocates a fraction ϕ_i of his remaining wealth to the ESG portfolio and a fraction $1 - \phi_i$ to the market portfolio. To see this, note that the $N \times 1$ vector of weights within agent i 's stock portfolio, w_i , equals the right-hand side of Eq. (14) multiplied by $1/(\iota'X_i)$, giving

$$\begin{aligned} w_i &= \frac{1}{\iota'(w_m + (\delta_i/a^2) \Sigma^{-1}g)} \left(w_m + (\delta_i/a^2) \Sigma^{-1}g \right) \\ &= (1 - \phi_i)w_m + \phi_i w_g , \end{aligned} \quad (15)$$

with the fraction of agent i 's stock portfolio invested in the ESG portfolio given by

$$\phi_i = \frac{(\delta_i/a^2) \iota'\Sigma^{-1}g}{1 + (\delta_i/a^2) \iota'\Sigma^{-1}g} , \quad (16)$$

and the $N \times 1$ vector of weights in the ESG portfolio given by

$$w_g = \frac{1}{\iota'\Sigma^{-1}g} \Sigma^{-1}g . \quad (17)$$

In the special case where $\iota'\Sigma^{-1}g = 0$, no agent holds the riskless asset, and the ESG portfolio is a zero-cost position, with⁵

$$w_g = \Sigma^{-1}g , \quad (18)$$

and $w_i = X_i$, so that

$$w_i = w_m + \phi_i w_g , \quad (19)$$

with ϕ_i then defined as

$$\phi_i = \delta_i/a^2 . \quad (20)$$

⁵In the special case considered in Section 4, in which $\iota'\Sigma^{-1}g = 0$ and Σ has a two-factor structure, w_g is proportional to g , so that the ESG portfolio goes long green stocks and short brown stocks.

Denote the ESG portfolio's greenness as

$$g_g = w'_g g . \quad (21)$$

From Eqs. (17) and (18), g_g is nonzero as long as $g \neq 0$. Also, g_g is negative if $\iota' \Sigma^{-1} g < 0$, but it is otherwise positive. From Eqs. (16) through (20), we see that ϕ_i has the same sign as the product of g_g and δ_i , if the denominator of ϕ_i in Eq. (16) is positive. From Eq. (14), this last condition obtains if agent i invests a positive fraction of her wealth in stocks, so that $\iota' X_i > 0$.

Therefore, for an agent with positive wealth in stocks and $\delta_i > 0$, ϕ_i is positive (negative) if g_g is positive (negative). That is, such an agent in general tilts away from the market portfolio in the direction of greenness, in that she tilts toward the ESG portfolio when it is green and away from it when it is brown. In contrast, agents with $\delta_i < 0$ tilt away from the ESG portfolio when it is green and toward it when it is brown. From Eq. (10), the ESG portfolio's CAPM alpha is

$$\alpha_g = -\frac{\bar{d}}{a} g_g , \quad (22)$$

whose sign is opposite that of g_g . Therefore, for the same agents described above, those with positive (negative) values of δ_i have ESG-portfolio tilts that produce negative (positive) alphas for their overall portfolios.

The ESG tilt is zero (i.e., $\phi_i = 0$) for agents with average ESG concerns, i.e., for whom $d_i = \bar{d}$ and thus $\delta_i = 0$. Those agents hold the market portfolio. In contrast, agents who are indifferent to ESG, for whom $d_i = 0$ and thus $\delta_i < 0$, tilt away from the market portfolio as explained above. It is suboptimal to say, "I don't care about ESG, so I just hold the market." In a world with ESG concerns, agents indifferent to ESG should tilt away from the market portfolio; otherwise they are not optimizing. The market portfolio is optimal for agents with average concerns about ESG but not for those indifferent to ESG.

If all agents have identical ESG concerns, so that $\delta_i = 0$ for all i , then Eqs. (16) and (20) imply a zero ESG tilt for each agent. We thus have the following corollary.

Corollary 4. *If there is no dispersion in ESG tastes across agents, then all agents hold the market portfolio.*

For example, all agents hold the market portfolio when none of them have ESG concerns, as in the familiar CAPM. All agents also hold the market, however, when they have strong but equal ESG tastes. The reason is that stock prices then fully adjust to reflect those tastes, again making the market everybody's optimal choice. Dispersion in ESG tastes is necessary for an ESG investment industry to exist.

2.3. Two-factor pricing with the ESG portfolio

The excess return on the ESG portfolio is $\tilde{r}_g = w'_g \tilde{r}$. From Eqs. (8) and (17) and $\beta_m = (1/\sigma_m^2)\Sigma w_m$, the ESG portfolio's market beta is zero (i.e., $w'_g \beta_m = 0$). Premultiplying both sides of Eq. (9) by w'_g gives the expected excess return on the ESG portfolio as

$$\mu_g = -\frac{\bar{d}}{a} g_g , \quad (23)$$

the same as its alpha in Eq. (22). The variance of the ESG portfolio's return is

$$\sigma_g^2 = \left(\frac{1}{\iota' \Sigma^{-1} g} \right)^2 g' \Sigma^{-1} g = \left(\frac{1}{\iota' \Sigma^{-1} g} \right) g_g , \quad (24)$$

and the covariance of its return with the N assets is

$$\text{Cov}(\tilde{r}, \tilde{r}_g) = \left(\frac{1}{\iota' \Sigma^{-1} g} \right) g . \quad (25)$$

Define the vector of simple betas with respect to \tilde{r}_g as $\beta_g = (1/\sigma_g^2)\text{Cov}(\tilde{r}, \tilde{r}_g)$. From Eqs. (24) and (25),

$$\beta_g = \frac{1}{g_g} g . \quad (26)$$

By combining Eqs. (9), (23), and (26), we relate expected returns to betas on the market and the ESG portfolio:

Proposition 4. *Expected excess returns in equilibrium are given by*

$$\mu = \mu_m \beta_m + \mu_g \beta_g . \quad (27)$$

As noted earlier, the ESG portfolio is zero-beta, so that $\text{Cov}(\tilde{r}_g, \tilde{r}_m) = 0$. Thus β_m and β_g are also the slope coefficients in the multivariate regression of \tilde{r} on \tilde{r}_m and \tilde{r}_g . Therefore, using Eq. (27), we have a two-factor asset pricing model:

Proposition 5. *Excess returns obey the regression model*

$$\tilde{r} = \beta_m \tilde{r}_m + \beta_g \tilde{r}_g + \tilde{\nu} , \quad (28)$$

in which $E(\tilde{\nu} | \tilde{r}_m, \tilde{r}_g, \beta_m, \beta_g) = 0$ and all assets have zero two-factor alphas, equivalent to zero intercepts in the above regression.

From Eqs. (9), (10) and (27), the vector of CAPM alphas is given by

$$\alpha = \beta_g \mu_g \quad (29)$$

$$= -(\bar{d}/a) g . \quad (30)$$

Eqs. (29) and (30) allow alternative interpretations of α . On one hand, Eq. (29) offers a risk-based interpretation: The elements of β_g represent exposures to the risky return \tilde{r}_g , and μ_g is the expected return accompanying a unit of that risk. In other words, one can attribute assets' non-zero CAPM alphas to that omitted priced risk factor. On the other hand, the only reason that investors expose themselves to the risk in \tilde{r}_g is that they have non-average tastes for green and brown holdings. While the popular risk-based interpretation of factor pricing models is mechanically valid, we see here an example of how that interpretation can miss the underlying economics. The latter are evident in Eq. (30), which reveals that the sources of α are tastes for known characteristics, g , not aversion to an additional fundamental risk. In Section 5, however, we extend our model to include an example of such risk, climate shocks, and we discuss how g can also reflect exposures to that risk.

3. The ESG factor

We next introduce an empirically identifiable ESG factor, closely related to the ESG portfolio, that maintains two-factor pricing. After discussing approaches for constructing the ESG factor, we analyze its underlying economic sources of risk. The latter analysis provides insights into ex post versus ex ante performance of green stocks relative to brown.

3.1. Constructing the ESG factor

We define the ESG factor as

$$\tilde{f}_g = (1/g_g) \tilde{r}_g , \quad (31)$$

so that the traded factor \tilde{f}_g is simply the excess return on a position in the ESG portfolio, either long or short, levered or de-levered, depending on the sign and value of g_g . Using Eq. (26), we can then rewrite the two-factor model in Eq. (28) as

$$\tilde{r} = \beta_m \tilde{r}_m + g \tilde{f}_g + \tilde{\nu} . \quad (32)$$

Assets' loadings on the ESG factor, their ESG betas, are simply their ESG characteristics, g . A higher-than-expected realization of \tilde{f}_g boosts the returns on green stocks and depresses those on brown ones. From Eqs. (23) and (31), the ESG factor's premium is negative:

$$E \{ \tilde{f}_g \} = -\bar{d}/a . \quad (33)$$

One approach to constructing the ESG factor is to run a cross-sectional regression of market-adjusted excess stock returns, $\tilde{r}^e \equiv \tilde{r} - \beta_m \tilde{r}_m$, on the stocks' ESG characteristics, g ,

with no intercept. The slope from that regression is

$$\hat{f}_g = \frac{g' \tilde{r}^e}{g'g} , \quad (34)$$

which from Eq. (32) has mean-zero estimation error $\hat{f}_g - \tilde{f}_g = g' \tilde{\nu} / g'g$. As N grows large, the probability limit of this estimation error is zero as long as the covariance matrix of $\tilde{\nu}$ has bounded eigenvalues and the cross-sectional second moment of the elements of g is bounded below by a positive value. The ESG factor is thus essentially just a g -weighted average of market-adjusted stock returns. To obtain the time series of the ESG factor's realizations in practice, one can run a series of such cross-sectional regressions, period by period.

A simpler version of \hat{f}_g arises if we add the assumptions of $g'\iota = 0$ and $g'\beta_m = 0$; that is, if we assume that not only the value-weighted average of g_n 's but also their equal- and beta-weighted averages are zero (recall from Eq. (8) that $g'w_m = 0$). In this case, $g'\tilde{r}^e$ in Eq. (34) equals $g'\tilde{r}$, so \hat{f}_g is just a scaled (by $g'g$) excess return on a zero-cost portfolio whose weights are proportional to g . Simplifying further, \hat{f}_g is proportional to the difference between returns on green-stock and brown-stock portfolios:

$$\hat{f}_g \propto \tilde{r}_{green} - \tilde{r}_{brown} , \quad (35)$$

with the weights in the green (brown) portfolio proportional to the positive (negative) elements of g . A popular approach to constructing traded factors (e.g., Fama and French, 1993) is to have them be excess returns on long-short portfolios whose stock weights sum to zero. Our model provides a formal justification for such an approach in the context of ESG investing. However, unlike in Fama and French (1993), stocks in our ESG factor are weighted by their g_n 's rather than by their market capitalizations.

3.2. Sources of ESG factor risk

In this subsection, we extend our model from Section 2 to identify potential sources of risk in the ESG factor. The strength of ESG concerns can change over time, both for investors in firms' shares and for the customers who buy the firms' goods and services. If ESG concerns strengthen, customers may shift their demands for goods and services to greener providers (the "customer" channel), and investors may derive more utility from holding the stocks of greener firms (the "investor" channel). Both channels contribute to the ESG factor's risk in our framework.

To model the customer channel, we need to model firm profits. Let \tilde{u}_n denote the financial payoff (profit in our one-period setting) that firm n produces at time 1, for each dollar

invested in the firm's stock at time 0. We assume a simple two-factor structure for the $N \times 1$ vector of these payoffs of the form

$$\tilde{u} - E_0\{\tilde{u}\} = \tilde{z}_m \beta_m + \tilde{z}_g g + \tilde{\zeta} , \quad (36)$$

where $E_0\{ \}$ denotes expectation as of time 0, the random quantities \tilde{z}_m , \tilde{z}_g , and $\tilde{\zeta}$ have zero means and are mutually uncorrelated, $\beta'_m g = 0$, and the elements of $\tilde{\zeta}$ have identical variances and are uncorrelated with each other. The shock \tilde{z}_m can be viewed as a macro output factor, with firms' sensitivities to that pervasive shock being proportional to their stocks' market betas. The shock \tilde{z}_g represents the effect on firms' payoffs of unanticipated ESG-related shifts in customers' demands. These shifts can result not only from changes in consumers' tastes but also from revisions of government policy. For example, pro-environmental regulations may subsidize green products, leading to more customer demand, or handicap brown products, leading to less demand. A positive \tilde{z}_g shock increases the payoffs of green firms but hurts those of brown firms.

To model the investor channel, we assume that the average ESG taste \bar{d} shifts unpredictably from time 0 to time 1. We therefore need to price stocks not only at time 0, as we have done so far, but also at time 1, after the preference shift in \bar{d} occurs. To make this possible in our simple framework, we split time 1 into two times, 1^- and 1^+ , that are close to each other. We calculate prices p_1 as of time 1^- , by which time ESG tastes have shifted and all risk associated with \tilde{u} has been realized. Stockholders receive \tilde{u} at time 1^+ . During the instant between times 1^- and 1^+ , these payoffs are riskless. For economy of notation, we assume the risk-free rate $r_f = 0$.

There are two generations of agents, Gen-0 and Gen-1. Gen-0 agents live from time 0 to time 1^- ; Gen-1 agents live from time 1^- to 1^+ . Gen-1 agents have identical tastes of $d_i = \bar{d}_1$, a condition that gives them finite utility, given the absence of both risk and position constraints during their lifespan. Neither a nor g change across generations. At time 1^- , Gen-0 agents sell stocks to Gen-1 agents at prices p_1 , which depend on Gen-1 ESG tastes \bar{d}_1 and the financial payoff \tilde{u} . This simple setting maintains single-period payoff uncertainty while also allowing risk stemming from shifts in ESG tastes to enter via both channels described earlier.

Given that the payoff \tilde{u}_n is known at the time when the price $p_{1,n}$ is computed, $p_{1,n}$ is equal to \tilde{u}_n discounted at the expected return implied by Eq. (9) with $\beta_{m,n}$ set to zero:

$$p_{1,n} = \frac{\tilde{u}_n}{1 - \frac{g_n}{a} \bar{d}_1} \approx \tilde{u}_n + \frac{g_n}{a} \bar{d}_1 . \quad (37)$$

The approximation above holds well for typical discount rates, which are not too far from

zero.⁶ Representing it as an equality for all assets gives

$$p_1 = \tilde{u} + \frac{1}{a} \bar{d}_1 g, \quad (38)$$

which is the vector of payoffs to Gen-0 agents. Its expected value at time 0 equals

$$\mathbb{E}_0\{p_1\} = \mathbb{E}_0\{\tilde{u}\} + \frac{1}{a} \mathbb{E}_0\{\bar{d}_1\} g. \quad (39)$$

Note that $p_1 - \mathbb{E}_0\{p_1\}$ equals the vector of unexpected returns for Gen-0 agents, because \tilde{u}_n is the firm's payoff per dollar invested in its stock at time 0. From Eqs. (36) through (39), these unexpected returns are given by

$$\tilde{r} - \mathbb{E}_0\{\tilde{r}\} = \beta_m \tilde{z}_m + g \tilde{f}_g^e + \tilde{\zeta} \quad (40)$$

with

$$\tilde{f}_g^e = \tilde{z}_g + \frac{1}{a} [\bar{d}_1 - \mathbb{E}_0\{\bar{d}_1\}]. \quad (41)$$

As shown in the Appendix, when N is large,

$$\tilde{f}_g^e \approx \tilde{f}_g - \mathbb{E}_0\{\tilde{f}_g\}, \quad (42)$$

where \tilde{f}_g is the ESG factor defined in Eq. (31).

Eq. (41) therefore identifies the two sources of risk in the ESG factor discussed earlier: \tilde{z}_g represents the customer channel while the other term represents the investor channel. While the customer channel follows closely from the structure assumed in Eq. (36), the investor channel emerges from the equilibrium dependence of stock prices on \bar{d} .

The elements of $\tilde{f}_g^e g$ in Eq. (40) drive a wedge between expected and realized returns for Gen-0 agents. Suppose that ESG concerns strengthen unexpectedly, so that $\tilde{f}_g^e > 0$. A firm's unexpected return in Eq. (40) is then expected to be positive for green firms (for which $\tilde{f}_g^e g_n > 0$) and negative for brown firms (for which $\tilde{f}_g^e g_n < 0$), because the expected values of \tilde{z}_m and $\tilde{\zeta}$ are both zero. In other words, if $\tilde{\epsilon}_n$ denotes the unexpected return for stock n , $\mathbb{E}\{\tilde{\epsilon}_n | \tilde{f}_g^e > 0, g_n > 0\} > 0$, and $\mathbb{E}\{\tilde{\epsilon}_n | \tilde{f}_g^e > 0, g_n < 0\} < 0$. We thus have the following proposition.

Proposition 6. *Green (brown) stocks perform better (worse) than expected if ESG concerns strengthen unexpectedly via either the customer channel or the investor channel.*

⁶Let $\rho_1 \equiv \tilde{u}_n - 1$ and $\rho_2 \equiv \frac{g_n}{a} \bar{d}_1$. The approximation in Eq. (37) follows from $\frac{1+\rho_1}{1-\rho_2} = \frac{(1+\rho_1)(1+\rho_2)}{1-\rho_2^2} \approx (1+\rho_1)(1+\rho_2) \approx 1 + \rho_1 + \rho_2$, where we assume that the second-order terms ρ_2^2 and $\rho_1\rho_2$ are small enough to be neglected. This assumption seems plausible because the magnitudes of ρ_1 and ρ_2 are comparable to discount rates. We rely on this approximation in the remainder of Section 3.

As noted earlier, green stocks have lower expected returns than brown stocks. A positive realization of \tilde{f}_g^e , however, boosts the realized performance of green stocks while hurting that of brown stocks. If one computes average returns over a sample period when ESG concerns strengthen more than investors expected, so that the average of \tilde{f}_g^e over that period is strongly positive, then green stocks outperform brown stocks, contrary to what is expected. Furthermore, if ESG concerns strengthen via the investor channel, making \bar{d} increase, then green stocks' alphas are more negative at the end of the period than the beginning (see Corollary 2). In this case, past outperformance of green stocks makes it especially likely they will underperform in the future.

To empirically distinguish alphas from unexpected shocks, one could use proxies for shifts in ESG tastes. Proxies for shifts in investors' tastes could come from investor surveys or from the flows in and out of ESG-tilted funds. Proxies for shifts in customers' ESG tastes could come from consumer surveys or from data on firm revenues or profitability. With such proxies, one could test whether green stocks outperform brown ones when either type of ESG tastes strengthens unexpectedly. In addition, one could attempt to separate the effects of investors' and customers' tastes, because only shifts in investors' tastes make green stocks' future alphas more negative after green stocks outperform.

4. Quantitative implications

To explore the model's quantitative implications, we consider a special case with two types of agents: ESG investors, for whom $d_i = d > 0$, and non-ESG investors, for whom $d_i = 0$. ESG investors thus enjoy nonpecuniary benefits dg , whereas non-ESG investors receive no benefits (see Eq. (3)). Let λ denote the fraction of total wealth belonging to ESG investors, so that $1 - \lambda$ is the corresponding fraction for non-ESG investors.

We further simplify the two-factor setting in Eq. (32) by assuming that ν in Eq. (32) has a scalar covariance matrix, $\eta^2 I_N$, where I_N is the identity matrix. The covariance matrix of \tilde{r} is therefore of the form

$$\Sigma = \sigma_m^2 \beta_m \beta_m' + \sigma_f^2 g g' + \eta^2 I_N. \quad (43)$$

Recall that $w_m' g = 0$, which here implies $\iota' g = 0$. We assume that $\beta_m' g = 0$, so that ESG characteristics are orthogonal to market betas. We also assume equal market weights across stocks, $w_m = (1/N)\iota$. Without loss of generality, we set $(g' g)/N = 1$. In all calculations, we take limits as the number of stocks, N , grows large.

4.1. Parameter values

In this simple setting there are only four parameters whose numerical values are relevant to the initial set of results we present: λ , a , σ_m , and Δ (defined below). We vary λ over its entire $[0, 1]$ range. We set $\sigma_m = 0.20$, roughly the historical standard deviation of the market portfolio's excess return. Following Eq. (7), we then set $a = \mu_m/\sigma_m^2$ with $\mu_m = 0.08$, roughly the market's historical mean excess return.⁷

Rather than calibrating d , we translate it to a more easily interpreted quantity, Δ . We define Δ as the maximum rate of return that an ESG investor is willing to sacrifice, for certain, to invest in her desired portfolio rather than in the market portfolio. The sacrifice is greatest when there are no other ESG investors, i.e., when $\lambda \approx 0$, because that is when the ESG investor's portfolio most differs from the market portfolio. Specifically, we define $\Delta \equiv r_{esg}^* - r_m^*$, where r_{esg}^* is the ESG investor's certainty equivalent excess return when investing in her optimal portfolio, and r_m^* is the same investor's corresponding certainty equivalent if forced to hold the market portfolio instead. Both certainty equivalents are computed for $\lambda = 0$. In this setting,

$$\Delta = \frac{d^2 g_g}{2a^3}, \quad (45)$$

as shown in the Appendix, along with the expressions for r_{esg}^* and r_m^* . Note that Δ is larger under stronger ESG tastes (larger d), lower risk aversion (smaller a), and a greener ESG portfolio (larger g_g). We consider four values of Δ : 1%, 2%, 3%, and 4% per year.

4.2. ESG versus non-ESG expected portfolio returns

The difference in expected excess returns on the portfolios of the two investor types is

$$\mathbb{E}\{\tilde{r}_{esg}\} - \mathbb{E}\{\tilde{r}_{non}\} = -2\lambda\Delta, \quad (46)$$

as shown in the Appendix. Fig. 1 plots this difference as λ goes from zero to one. The difference is zero at $\lambda = 0$, but it declines linearly as λ increases. At $\lambda = 1$, ESG tastes are fully reflected in prices, and the difference reaches its largest magnitude. In that scenario, the difference is -2% when $\Delta = 0.01$, but it is -8% when $\Delta = 0.04$. ESG investors thus

⁷Identifying σ_m^2 in Eq. (43) as the market variance is justified for large N . If we instead denote that variance as simply σ^2 , note that the implied variance of the market, $w'_m \Sigma w_m$, is

$$\sigma_m^2 = \frac{1}{N^2} \iota' (\sigma^2 \beta_m \beta'_m + \sigma_f^2 g g' + \eta^2 I_N) \iota = \sigma^2 + \frac{\eta^2}{N}, \quad (44)$$

noting $w'_m \beta_m = (1/N) \iota' \beta_m = 1$ and recalling $\iota' g = 0$, so we simply set $\sigma^2 = \sigma_m^2$, the limit as N grows large.

earn significantly lower returns than non-ESG investors when the former account for a larger fraction of wealth (larger λ) and when they have stronger ESG demands (larger Δ). In both [Fig. 1] scenarios, ESG tastes exert large effects on asset prices, hurting ESG investors' returns.

The certainty equivalent returns of the two types, r_{esg}^* for ESG investors and r_{non}^* for non-ESG investors, are both increasing in Δ , but r_{esg}^* decreases with λ whereas r_{non}^* increases with λ , as we show in the Appendix. As λ increases, stock prices are affected more by ESG investors' tastes, so these investors must pay more for the green stocks they desire. The resulting drop in r_{esg}^* need not imply, however, that an ESG investor is made less happy by an increased presence of ESG investors. With the latter, there is also greater social impact of ESG investing, as we discuss in Section 6. The additional utility that the ESG investor derives from the greater social impact, as in Eq. (68), can exceed the drop in utility corresponding to the lower r_{esg}^* . Non-ESG investors, on the other hand, do prefer to be lonely in their ESG tastes. A non-ESG investor is happiest when all other investors are ESG ($\lambda = 1$), because that scenario maximizes deviations of prices from pecuniary fundamentals, which the non-ESG investor exploits to her advantage. This investor's preference for loneliness in ESG tastes is even stronger if she derives utility from social impact, because that impact is maximized when $\lambda = 1$.

4.3. Correlation between the ESG return and the market return

The correlation between the return on an ESG investor's portfolio and the return on the market portfolio is derived in the Appendix:

$$\rho(\tilde{r}_{esg}, \tilde{r}_m) = \frac{\sigma_m}{\sqrt{\sigma_m^2 + \frac{2\Delta}{a}(1-\lambda)^2}}. \quad (47)$$

Fig. 2 plots the value of $\rho(\tilde{r}_{esg}, \tilde{r}_m)$ as λ goes from zero to one. The correlation takes its lowest value at $\lambda = 0$. For $\Delta = 0.01$, that value is nearly 0.9, whereas for $\Delta = 0.04$, it is just over 0.7. As Δ increases, indicating that ESG investors feel increasingly strongly about ESG, those investors' portfolios become increasingly different from the market portfolio in terms of $\rho(\tilde{r}_{esg}, \tilde{r}_m)$, and this effect is strongest when $\lambda = 0$. However, as λ approaches one, so does $\rho(\tilde{r}_{esg}, \tilde{r}_m)$. When ESG investors hold an increasingly large fraction of wealth, market prices adjust to their preferences, and all portfolios converge to the market portfolio. [Fig. 2]

4.4. Alphas and the investor surplus

The alphas of the ESG and non-ESG investors' portfolios are derived in the Appendix:

$$\alpha_{esg} = -2\lambda(1-\lambda)\Delta \quad (48)$$

$$\alpha_{non} = 2\lambda^2\Delta . \quad (49)$$

Panel A of Fig. 3 plots α_{esg} as λ goes from zero to one. ESG investors earn zero alpha at both extremes of λ . Their portfolio differs most from the market portfolio when $\lambda = 0$, but all stocks have zero alphas in that scenario, because there is no impact of ESG investors on prices. At the other extreme, when $\lambda = 1$, many stocks have non-zero alphas, due to the price impacts of ESG investors, but ESG investors hold the market, so again they earn zero alpha. Otherwise, ESG investors earn negative alpha, which is greatest in magnitude when $\lambda = 0.5$. At that peak, $\alpha_{esg} = -0.5\%$ when $\Delta = 0.01$, but $\alpha_{esg} = -2\%$ when $\Delta = 0.04$. [Fig. 3]

Interestingly, these worst-case alphas are substantially smaller in magnitude than the corresponding Δ 's. For example, when ESG investors are willing to give up 2% certain return to hold their portfolio rather than the market (i.e., $\Delta = 0.02$), their worst-case alpha is only -1% . The reason is that equilibrium stock prices adjust to ESG demands. These demands push the market portfolio toward the portfolio desired by ESG investors, thereby bringing those investors' negative alphas closer to zero. Through this adjustment of market prices, ESG investors earn an "investor surplus" in that they do not have to give up as much return as they are willing to in order to hold their desired portfolio.

The magnitude of this investor surplus is easy to read off Panel B of Fig. 3, which plots α_{esg} as a function of Δ . For any given value of λ , investor surplus is the difference between the corresponding solid line and the dashed line, which has a slope of -1 . The surplus increases with Δ because the stronger the ESG investors feel about greenness, the more they move market prices. The relation between the surplus and λ is richer. Formally, investor surplus $\mathcal{I} \equiv \alpha_{esg} + \Delta$ follows quickly from Eq. (48):

$$\mathcal{I} = \Delta[1 - 2\lambda(1 - \lambda)] . \quad (50)$$

Because $0 \leq \lambda \leq 1$, the value in brackets is always between 0.5 and 1, so \mathcal{I} is always between $\Delta/2$ and Δ . It reaches its smallest value of $\Delta/2$ when $\lambda = 0.5$ and its largest value of Δ when $\lambda = 0$ or 1. For example, when $\Delta = 0.02$, \mathcal{I} ranges from 1% to 2% depending on λ .

Fig. 4 plots α_{non} as a function of λ and Δ . Like ESG investors, non-ESG investors earn zero alpha when $\lambda = 0$ or $\Delta = 0$. However, α_{non} increases in both λ or Δ . This alpha can be as large as 8% when $\lambda = 1$ and $\Delta = 0.04$. A non-ESG investor earns the highest alpha when all other investors are ESG (i.e., $\lambda = 1$) and when those investors' ESG tastes are strong (i.e., Δ is large) because the price impact of ESG tastes is then particularly large. By overweighting brown stocks, whose alphas are positive and large, and underweighting green stocks, whose alphas are negative and large, the non-ESG investor earns a large positive [Fig. 4]

alpha.

Given our assumptions, the differences between the alphas plotted in Figures 3 and 4 are equal to the differences in expected returns plotted in Fig. 1. Specifically, from Eqs. (46) through (49), $\alpha_{esg} - \alpha_{non} = E\{\tilde{r}_{esg}\} - E\{\tilde{r}_{non}\}$.

4.5. Size of the ESG investment industry

We define the size of the ESG investment industry by the aggregate amount of ESG-driven investment that deviates from the market portfolio, divided by the stock market's total value. In general, this aggregate ESG tilt is given by

$$T = \int_{i:d_i>0} \omega_i T_i di , \quad (51)$$

where

$$T_i = \frac{1}{2} \iota' |w_i - w_m| . \quad (52)$$

The aggregate ESG tilt, T , is a wealth-weighted average of agent-specific tilts, T_i , across all agents who care at least to some extent about ESG (i.e., $d_i > 0$). Each T_i is one half of the sum of the absolute values of the N elements of agent i 's ESG tilt, $|w_i - w_m|$. We compute absolute values of portfolio tilts because ESG-motivated investors both over- and under-weight stocks relative to the market. We divide by two because we do not want to double-count: for each dollar that an agent moves into a green stock, she must move a dollar out of another stock.

With two types of agents, the expression for T simplifies to

$$T = \frac{1}{N} \lambda(1 - \lambda) \sqrt{\frac{\Delta}{2a\sigma_f^2}} \iota' |g| , \quad (53)$$

as we show in the Appendix. The aggregate tilt depends on the absolute values of the elements of g . To evaluate $\iota' |g|$ in this quantitative exercise, we further assume that the elements of g are normally distributed across stocks, in addition to the previous assumptions that these elements have zero mean and unit variance (recall $\iota' g = 0$ and $(g'g)/N = 1$). Then $\iota' |g| = NE(|g_n|) = N\sqrt{2/\pi}$. Therefore,

$$T = \lambda(1 - \lambda) \sqrt{\frac{\Delta}{a\pi\sigma_f^2}} . \quad (54)$$

For this analysis we now need to specify the value of one additional parameter, σ_f^2 , the standard deviation of the ESG factor. We set $\sigma_f = (0.2)\sigma_m$, but the effect of this parameter

is easily gauged from Eq. (54). The more volatile is the ESG factor, the more reluctant both ESG and non-ESG investors are to tilt away from the market and thereby expose themselves to the ESG factor’s risk. Higher risk aversion (greater a) also makes them more reluctant to do so, but they tilt more when ESG investors have stronger tastes (greater Δ).

Fig. 5 plots T for different values of λ and Δ . In Panel A, λ goes from zero to one. At both $\lambda = 0$ and $\lambda = 1$, we have $T = 0$ because all investors hold the market portfolio. Again, we see that dispersion in ESG tastes is needed for an ESG investment industry to exist. The maximum value of T in Eq. (54) always occurs at $\lambda = 0.5$, the maximum of $\lambda(1 - \lambda)$. In Panel B, Δ goes from 0 to 0.04. Larger values of Δ produce larger values of T . This relation between Δ and T is concave (see also Eq. (54)). For example, the ESG industry peaks at 35% of the stock market’s value when $\Delta = 0.02$, but doubling the strength of ESG tastes (raising Δ to 0.04) increases that maximum industry size by less than half, to 50% of the market’s value. We see that the price impact of ESG tastes weakens their impact on the size of the ESG investment industry. [Fig. 5]

5. Climate risk

Sustainable investing is motivated in part by concerns about climate change. Many experts expect climate change to impair quality of life, lowering utility of the typical individual beyond what is captured by climate’s effect on wealth. Unanticipated climate changes present investors with an additional source of risk, which is non-traded and only partially insurable.⁸ This section extends our model from Section 2 to include climate risk.

Let \tilde{C} denote climate at time 1, which is unknown at time 0. We modify the utility function for individual i in Eq. (2) to include \tilde{C} as follows:

$$V(\tilde{W}_{1i}, X_i, \tilde{C}) = -e^{-A_i\tilde{W}_{1i} - b'_i X_i - c_i \tilde{C}} . \quad (55)$$

Let $\bar{c} \equiv \int_i \omega_i c_i di$, the wealth-weighted mean of climate sensitivity across agents. We assume $\bar{c} > 0$, so that agents dislike low realizations of \tilde{C} , on average. We also assume \tilde{C} is normally distributed, and without loss of generality we set $E\{\tilde{C}\} = 0$ and $\text{Var}\{\tilde{C}\} = 1$. Besides replacing Eq. (2) by Eq. (55), we maintain all other assumptions from Section 2.

In principle, “climate” can be interpreted broadly, for example, as “social climate.” However, for shocks to climate to affect asset prices, these shocks must enter the average agent’s

⁸In that sense, climate risk is related to “background risk” analyzed in prior work. Research into the risk associated with non-marketable assets originates with Mayers (1972). Examples of non-traded systematic risk factors include human capital (Fama and Schwert, 1977), liquidity (Pástor and Stambaugh, 2003), and innovation-induced displacement (Garleanu, Kogan, and Panageas, 2012).

utility, in that $\bar{c} > 0$. This assumption is nontrivial because individuals' views on various social issues, such as guns and abortion, are quite heterogeneous in practice. We emphasize the narrow interpretation of climate ("E" in ESG), for which the assumption is likely to hold. Indeed, the correlations in ESG ratings across rating agencies are higher for the "E" ratings than for the "S" and "G" ratings (e.g., Berg, Koelbel, and Rigobon, 2019).

5.1. Expected returns and portfolio holdings

Climate risk affects equilibrium stock returns, as shown in the Appendix.

Proposition 7. *Expected excess returns in equilibrium are given by*

$$\mu = \mu_m \beta_m - \frac{\bar{d}}{a} g + \bar{c} (1 - \rho_{mC}^2) \psi, \quad (56)$$

where ψ is the $N \times 1$ vector of "climate betas" (slope coefficients on \tilde{C} in a multivariate regression of $\tilde{\epsilon}$ on both $\tilde{\epsilon}_m$ and \tilde{C}), and ρ_{mC} is the correlation between $\tilde{\epsilon}_m$ and \tilde{C} .

Expected returns depend on climate betas, ψ , which represent firms' exposures to non-market climate risk. To understand the regression defining ψ , recall that $\tilde{\epsilon}$ is an $N \times 1$ vector of unexpected stock returns from Eq. (1) and $\tilde{\epsilon}_m$ is the unexpected market return. A firm's climate beta is its loading on \tilde{C} after controlling for the market return. Climate betas, ψ_n , are likely to be related to ESG characteristics, g_n , as we argue in Section 5.2.

Compared to Eq. (9), expected excess returns contain an additional component given by the last term on the right-hand side of Eq. (56). Stock n 's climate beta, ψ_n , enters expected return positively. Thus, a stock with a negative ψ_n , which provides investors with a climate-risk hedge, has a lower expected return than it would in the absence of climate risk. Vice versa, a stock with a positive ψ_n , which performs particularly poorly when the climate worsens unexpectedly, has a higher expected return.

Proposition 8. *Agent i 's equilibrium portfolio weights on the N stocks are given by*

$$X_i = w_m + \frac{\delta_i}{a^2} (\Sigma^{-1} g) - \frac{\gamma_i}{a} (\Sigma^{-1} \sigma_{\epsilon C}), \quad (57)$$

where $\gamma_i \equiv c_i - \bar{c}$ and $\sigma_{\epsilon C}$ is an $N \times 1$ vector of covariances between $\tilde{\epsilon}_n$ and \tilde{C} .

Eq. (57), which we prove in the Appendix, implies four-fund separation. The first three funds are the same as in Proposition 3; the fourth one is a climate-hedging portfolio whose weights are proportional to $\Sigma^{-1} \sigma_{\epsilon C}$. Agents with $\gamma_i > 0$, whose climate sensitivity is above average, short the hedging portfolio, whereas agents with $\gamma_i < 0$ go long.

The climate-hedging portfolio, $\Sigma^{-1}\sigma_{\tilde{C}}$, is a natural mimicking portfolio for \tilde{C} . To see this, note that the N elements of $\Sigma^{-1}\sigma_{\tilde{C}}$ are the slope coefficients from the multiple regression of \tilde{C} on $\tilde{\epsilon}$. Therefore, the return on the hedging portfolio has the highest correlation with \tilde{C} among all portfolios of the N stocks. Investors in our model hold this maximum-correlation portfolio, to various degrees determined by their γ_i , to hedge climate risk. The climate-hedging portfolio can tilt toward either green stocks or brown stocks, depending on how returns on each type relate to climate shocks. The latter issue is addressed next.

5.2. Green stocks as climate hedges

Ultimately the issue of whether green stocks or brown stocks are better climate hedges is an empirical question, because sensible economic arguments can be made either way. The argument that green stocks should hedge climate risk can be motivated through both channels described in Section 3.

First, consider the customer channel. Unexpected worsening of the climate can heighten consumers' climate concerns, prompting greater demands for goods and services of greener providers. These demands can arise not only from consumers' preferences but also from government regulation. Negative climate shocks can prompt government regulations that favor green providers or penalize brown ones. For example, the new regulations could subsidize green products and tax, or even prohibit, brown ones. Half of the institutional investors surveyed by Krueger, Sautner, and Starks (2019) state that climate risks related to regulation have already started to materialize.

Second, consider the investor channel. Unexpected worsening of the climate can strengthen investors' preference for green holdings (i.e., increase \bar{d}), possibly as a result of stronger public pressure on institutional investors to divest from brown assets. For example, Choi, Gao, and Jiang (2019) show that retail investors sell carbon-intensive firms in extremely warm months, consistent with \bar{d} rising in such months.

Climate shocks are thus likely to correlate negatively with both components of the ESG factor in Eq. (41). Green stocks, which have positive exposures to this factor, are likely to have negative exposures to \tilde{C} . These arguments imply a negative correlation between g_n and ψ_n across firms.

One can also argue that the better hedges of climate risk are brown stocks, not green. Baker, Hollifield, and Osambela (2020) assume negative climate shocks result from positive shocks to the output of brown firms. The latter shocks translate to positive unexpected returns on those firms' stocks, thereby making brown stocks climate hedges. As noted

earlier, whether brown stocks or green stocks better hedge climate risk ultimately rests on empirical evidence.

The evidence suggests the better climate hedges are green stocks. For example, Choi, Gao, and Jiang (2019) show that green firms, as measured by low carbon emissions, outperform brown firms during months with abnormally warm weather, which the authors argue alerts investors to climate change. Engle et al. (2020) report that green firms, as measured by high E-Scores from Sustainalytics, outperform brown firms in periods with negative climate news. Both studies thus show that a high-minus-low g_n stock portfolio is a good hedge against climate risk, indicating that g_n is negatively correlated with ψ_n across firms.

In the special case where this negative correlation is perfect, so that

$$\psi_n = -\xi g_n , \quad (58)$$

where $\xi > 0$ is a constant, Eq. (56) simplifies to

$$\mu = \mu_m \beta_m - \left[\frac{\bar{d}}{a} + \bar{c} (1 - \rho_{mC}^2) \xi \right] g . \quad (59)$$

Stock n 's CAPM alpha is then given by

$$\alpha_n = - \left[\frac{\bar{d}}{a} + \bar{c} (1 - \rho_{mC}^2) \xi \right] g_n . \quad (60)$$

Both terms inside the brackets are positive, so the negative relation between α_n and g_n is stronger than in Corollary 2. Greener stocks now have lower CAPM alphas not only because of investors' tastes for green holdings, but also because of greener stocks' ability to better hedge climate risk. Climate risk thus represents another reason to expect green stocks to underperform brown ones over the long run. For the same reason, green stocks have a lower cost of capital than brown stocks relative to the CAPM.

In this special case, two-factor pricing from Section 3 continues to hold. Each stock has a zero alpha in the two-factor model in Eq. (32). The ESG factor is still defined as in Eq. (31), but its premium is reduced by $\bar{c}(1 - \rho_{mC}^2) \xi$, as compared to Eq. (33). This reduction reflects compensation for climate risk. The compensation is negative because greener firms are better hedges against this risk. The ESG factor's premium thus has one taste-based component, $-\bar{d}/a$, and one risk-based component, $-\bar{c}(1 - \rho_{mC}^2) \xi$.

6. Social impact

Does sustainable investing produce real social impact? This section explores how firms respond to the asset pricing effects from Section 2. We extend our baseline model from that

section to include firms' choices of investment and ESG characteristics.

We define the social impact of firm n as

$$S_n \equiv g_n K_n , \quad (61)$$

where K_n is the firm's operating capital. Social impact captures the firm's total externalities, which depend on both the nature of the firm's operations (g_n) and their scale (K_n). We consider two scenarios. In Section 6.1, we let the firm's manager choose K_n while taking g_n as given. In Section 6.2, we allow the manager to choose both K_n and g_n . Throughout, the manager maximizes the firm's market value at time 0.

The extra assumptions we make here change none of the previous sections' predictions. Since investors are infinitesimally small, they still take asset prices and firms' ESG characteristics as given, even though firms now choose those characteristics. Firms' choices of K_n and g_n affect their market values, which are consistent with the expected returns derived in Section 2.

6.1. *Green firms invest more, brown firms less*

The firm is initially endowed with operating capital $K_{0,n} > 0$. The firm's manager chooses how much additional capital, ΔK_n , to buy, while taking the firm's ESG characteristic, g_n , as given. The firm's capital investment produces a time-zero cash flow of $-\Delta K_n - \frac{\kappa_n}{2}(\Delta K_n)^2$, where $\kappa_n > 0$ controls capital-adjustment costs. The firm uses capital to produce an expected gross cash flow at time 1 equal to $\Pi_n K_n$, where Π_n is a positive quantity denoting one plus the firm's gross profitability.

The optimal amount of additional capital is derived in the Appendix:

$$\Delta K_n(\bar{d}) = \frac{1}{\kappa_n} \left[\frac{\Pi_n}{1 + r_f + \mu_m \beta_{m,n} - \frac{\bar{d}}{a} g_n} - 1 \right] . \quad (62)$$

This value is increasing in g_n , indicating that greener firms invest more, *ceteris paribus*.

For any firm n , agents' ESG tastes induce social impact equal to the difference between the firm's actual social impact and its hypothetical impact if agents did not care about ESG:

$$S_n(\bar{d}) - S_n(0) = g_n \left(\Delta K_n(\bar{d}) - \Delta K_n(0) \right) . \quad (63)$$

We prove the following proposition in the Appendix.

Proposition 9. *Firm n 's ESG-induced social impact is positive:*

$$S_n(\bar{d}) - S_n(0) = \frac{\bar{d} g_n^2 \Pi_n}{a \kappa_n \left(1 + r_f + \mu_m \beta_{m,n} - \frac{\bar{d}}{a} g_n\right) (1 + r_f + \mu_m \beta_{m,n})} > 0, \quad (64)$$

as long as $\bar{d} > 0$ and $g_n \neq 0$. Moreover, this impact is increasing in \bar{d} , decreasing in a , increasing in Π_n , decreasing in κ_n , and decreasing in $\beta_{m,n}$.

The intuition behind this result builds on Eq. (62), which shows that ESG tastes lead green firms to invest more and brown firms to invest less. That result relates to Corollary 1, which states that ESG tastes reduce green firms' expected returns and hence their costs of capital. Green firms' lower costs of capital increase their projects' NPVs, so green firms invest more. And vice versa, ESG tastes increase brown firms' costs of capital, reducing their investment. As a result, ESG tastes tilt investment from brown to green firms, which increases social impact for both types of firms.

The comparative statics are also intuitive. Social impact is larger when ESG tastes are stronger (i.e., when \bar{d} is larger) because stronger tastes move asset prices more. The impact is also larger when risk aversion is weaker (i.e., a is smaller) because less risk-averse agents tilt their portfolios more to accommodate their tastes, again resulting in larger price effects (see Propositions 1 and 3). The impact is larger when capital is less costly to adjust (i.e., when κ_n is smaller) because more investment reallocation takes place. The impact is also larger when firms are more productive (i.e., when Π_n is larger) because a given change in the cost of capital has a larger effect on investment. Finally, the impact is larger for firms with smaller market betas because such firms have a lower cost of capital to begin with, so the ESG-induced change in their cost of capital is relatively larger.

In our model, investors' ESG tastes tilt real investment from brown to green firms because those tastes generate alphas, which affect the cost of capital, which in turn affects investment. There is considerable empirical support for this mechanism. Baker and Wurgler (2012) survey studies that find a negative relation between corporate investment and alpha. Most of these studies interpret alpha as mispricing, whereas our study's ESG-induced alphas do not reflect mispricing. We expect ESG-induced alphas to have an especially strong effect on investment. Whereas mispricing is transient, firms' ESG traits are highly persistent, which makes ESG-induced alphas highly persistent. Van Binsbergen and Opp (2019) show that when alphas are more persistent, they have stronger effects on investment.

6.2. Firms become greener

We now extend the framework from Section 6.1 by allowing firm n 's manager to choose not only K_n but also g_n . The firm is initially endowed with an ESG characteristic $g_{0,n}$. The manager chooses both ΔK_n and Δg_n , the change in the firm's ESG characteristic. For example, a coal power producer can increase its g_n by installing scrubbers. Adjusting g_n is costly: it reduces the firm's time-one cash flow by a fraction $\frac{\chi_n}{2}(\Delta g_n)^2$, where $\chi_n > 0$ controls ESG-adjustment costs.

We prove the following proposition in the Appendix.

Proposition 10. *Firm n 's value-maximizing choices of ESG adjustment and investment are*

$$\Delta g_n(\bar{d}) \approx \frac{\bar{d}}{a\chi_n} \quad (65)$$

$$\Delta K_n(\bar{d}) = \frac{1}{\kappa_n} \left[\frac{\Pi_n \left(1 - \frac{\chi_n}{2}(\Delta g_n(\bar{d}))^2\right)}{1 + r_f + \mu_m \beta_{m,n} - \frac{\bar{d}}{a} g_n(\bar{d})} - 1 \right], \quad (66)$$

where $g_n(\bar{d}) = g_{0,n} + \Delta g_n(\bar{d})$, and the approximation uses $\log(1+x) \approx x$ for small x .

Both choices are intuitive given the results from Section 2. When $\bar{d} > 0$, expected returns decrease in g_n (Corollary 1), so firms' market values increase in g_n . Managers who wish to maximize market value therefore make their firms greener (i.e., $\Delta g_n > 0$). This effect is especially strong when risk aversion a is low because ESG characteristics then have large effects on market values. Firms also adjust g_n by more when doing so is less costly.

As in Section 6.1, ESG tastes lead green firms to invest more and brown firms to invest less. The denominator in Eq. (66) shows that ESG tastes reduce the costs of capital for green firms, which increases their projects' NPV and hence investment. And vice versa, ESG tastes increase brown firms' costs of capital, reducing their projects' NPV and investment. In addition, ESG tastes affect expected cash flows in the numerator of Eq. (66). Stronger ESG tastes induce all firms, green and brown, to adjust their g_n by more, which reduces their expected cash flows, and hence also their investment.

Agents' ESG tastes now increase social impact not only by tilting investment from brown to green firms, as before, but also by making firms greener:

$$S_n(\bar{d}) - S_n(0) = g_{n,0} \left(\Delta K_n(\bar{d}) - \Delta K_n(0) \right) + K_n(\bar{d}) \Delta g_n(\bar{d}). \quad (67)$$

The first term reflects the investment effect analogous to Eq. (63). As discussed previously, when firms cannot change their g_n 's, $\Delta K_n(\bar{d}) - \Delta K_n(0)$ is positive for green firms and

negative for brown firms, making this term positive for both types of firms. When firms can change their g_n 's, the first term in Eq. (67) is still generally positive. The second term reflects firms' capital becoming greener. This term is also positive since $\Delta g_n(\bar{d}) > 0$, as implied by Eq. (65).

Fig. 6 plots the ESG-induced social impact across firms with different initial ESG characteristics. We see that all firms have positive social impact. The two colored regions indicate the two sources of social impact from Eq. (67). The second source, from firms becoming greener, is roughly equal across firms (top green region). The first source, from tilting investment toward green firms, is zero for an ESG-neutral firm, but it is large for very green or very brown firms, which experience the largest shifts in investment (bottom blue region). Due to this non-monotonicity, the overall social impact induced by ESG-motivated investors is largest for firms with extreme ESG characteristics, but it is strictly positive even for ESG-neutral firms.

[Fig. 6]

The aggregate social impact induced by ESG investors, denoted $S(\bar{d}) - S(0)$, is the sum of $S_n(\bar{d}) - S_n(0)$ across firms n . This sum can be computed off the curve in Fig. 6. Since this curve is convex in $g_{0,n}$, $S(\bar{d}) - S(0)$ is greater when there is more dispersion in ESG characteristics across firms. A larger dispersion in $g_{0,n}$ deepens the cost-of-capital differentials between green and brown firms, leading to larger investment differentials. With green firms investing more and brown firms investing less, aggregate social impact increases.

Fig. 7 illustrates how aggregate social impact varies with the strength of ESG preferences. We assume firms differ only in their initial ESG characteristics $g_{n,0}$, which are uniformly distributed with mean zero. The figure shows that $S(\bar{d}) - S(0)$ increases as ESG preferences strengthen, which is intuitive. We also see that both sources of social impact from Eq. (67) grow larger as ESG preferences strengthen. These results hold whether ESG preferences strengthen because there are more ESG investors (Panel A), or because the ESG investors have stronger ESG tastes (Panel B).

[Fig. 7]

We have made the standard assumption that managers maximize the firm's market value. This assumption makes sense if, for example, managers wish to maximize the value of their stock-based compensation. Alternatively, a manager could maximize shareholder welfare, which depends not just on market value but also on the firm's ESG characteristics (e.g., Hart and Zingales, 2017). Such behavior could result from shareholders engaging actively with the firm, so that managers run the firm as shareholders desire (e.g., Dyck et al., 2019), or from shareholders appointing managers whose preferences match their own. Our model arguably provides a lower bound on social impact. Extending the model so that managers additionally

care about their firms’ ESG characteristics should produce Δg_n values (and hence social impact) even larger than we currently predict. Put differently, we show that ESG-motivated investors generate social impact even without direct engagement by shareholders, and even if managers do not care directly about firms’ ESG characteristics. Even a “selfish” manager who cares only about market value behaves in a way that increases social impact.

6.3. Preferences for aggregate social impact

As noted in Section 2, agents derive utility not only from their holdings, X_i , but also from firms’ aggregate social impact, $S = \sum_{n=1}^N S_n$. We assume each agent i ’s utility is increasing in S :

$$U(\tilde{W}_{1i}, X_i, S) = V(\tilde{W}_{1i}, X_i) + h_i(S) , \quad (68)$$

where $h'_i(S) > 0$ and V is the original utility function from Eq. (2). (The additive specification is not needed; our results are identical if S enters utility multiplicatively.)

Proposition 11. *If agents derive utility also from aggregate social impact (Eq. (68)), all of our results in Propositions 1 through 10 and Corollaries 1 through 4 continue to hold.*

The inclusion of S in the utility function does not affect any of our prior results. The reason is that infinitesimally small agents take stock prices, and hence S , as given when choosing their portfolios. When an agent tilts toward green stocks, she generates a positive externality on other agents via the $h_i(S)$ term in their utility.⁹ Being infinitesimal, though, she does not internalize any of this effect. As the preference for S does not affect portfolio choice, it does not affect equilibrium asset prices, real investment, or S . In the model of Oehmke and Opp (2020), agents’ preference for social impact does lead to impact because agents are assumed to coordinate. In our model, agents cannot coordinate. Social impact is caused by the inclusion of X_i , not S , in the utility function in Eq. (68).

To provide more intuition for the roles of X_i and S in the utility function, consider why people vote in elections. Many individuals vote because they derive utility directly from doing so, analogous to investors deriving utility from their holdings (X_i) in our setting. This utility from voting can have various sources; for example, some people enjoy participating in a democracy, others feel a warm glow from voting for their favorite candidate, and some might like to tell friends they have exercised their patriotic duty. Each individual’s utility may also depend on the election outcome (S), but that by itself is not why an individual

⁹In the presence of externalities, the competitive market solution generally differs from the social planner’s solution. For an example of a social planner’s solution in a different setting, with interesting implications for ESG mandates, see Hong, Wang, and Yang (2020).

votes. If there are a large number of voters, the individual sees her vote as having no effect on that component of her utility. Just as utility from voting produces an aggregate social good (a healthy democracy), investors' utility from their portfolio holdings generates aggregate social impact.

More research is clearly needed on the real effects of sustainable investing. For example, what is the relative importance of the investment channel (ΔK_n) and the “become-greener” channel (Δg_n)? What if agents care about both social impact and climate, and the effect of the former on the latter is uncertain? How would social impact change if we combined the asset pricing effects we examine with direct engagement by large shareholders? We leave these questions for future work.

7. Conclusion

We analyze both financial and real effects of sustainable investing in a highly tractable equilibrium model. The model produces a number of empirical implications regarding asset prices, portfolio holdings, the size of the ESG investment industry, climate risk, and the social impact of sustainable investing. We review those implications below.

First, ESG preferences move asset prices. Stocks of greener firms have lower ex ante CAPM alphas, especially when risk aversion is low and the average ESG preference is strong. Green stocks have negative alphas, whereas brown stocks have positive alphas. Green stocks' negative alphas stem from two sources: investors' tastes for green holdings and such stocks' ability to hedge climate risk. Green and brown stocks have opposite exposures to an ESG risk factor, which captures unexpected changes in ESG concerns of customers and investors. If either kind of ESG concerns strengthen unexpectedly over a given period of time, green stocks can outperform brown stocks over that period, despite having lower alphas. Stocks are priced by a two-factor asset pricing model, where the factors are the market portfolio and the ESG factor. A simple version of the ESG factor is a green-minus-brown portfolio return, where both green and brown portfolios are weighted by ESG characteristics.

Second, portfolio holdings exhibit three-fund separation. Investors with stronger than average ESG tastes hold portfolios that have a green tilt away from the market portfolio, whereas investors with weaker than average ESG tastes take a brown tilt. These tilts are larger when risk aversion is lower. Investors with stronger ESG tastes earn lower expected returns, especially when risk aversion is low and the average ESG taste is high. Yet these investors give up less return than they are willing to in order to hold their desired portfolio. In the model extension that adds climate risk, we obtain four-fund separation, with the

fourth fund representing a climate-hedging portfolio with a green tilt.

Third, the size of the ESG investment industry—the aggregate dollar amount of ESG-driven investment that deviates from the market portfolio—is increasing in the dispersion of investors’ ESG preferences. With no dispersion there is no ESG industry, because everyone holds the market.

Finally, sustainable investing generates positive social impact, in two ways. First, it leads firms to become greener. Second, it induces more real investment by green firms and less investment by brown firms.

While the model’s predictions for alphas have been examined empirically by prior studies, most of its other predictions remain untested, presenting opportunities for future empirical work. One challenge is that our model aims to describe the world of the present and the future, but not necessarily the world of the past. Although the “sin” aspects of investing have been recognized for decades, the emphasis on ESG criteria is a recent phenomenon. How the model fits in various time periods is another question for empirical work.




Fig. 1. ESG versus non-ESG expected portfolio return. This figure plots the expected excess return on the portfolio of ESG investors minus the corresponding value for non-ESG investors. Results are plotted against λ , the fraction of wealth belonging to ESG investors, and for different values of Δ , the maximum certain return an ESG investor would sacrifice to invest in her optimal portfolio instead of the market portfolio.





Fig. 2. Correlation of ESG investor’s portfolio return with the market return. The figure plots the correlation between the returns on the ESG investor’s portfolio and the market portfolio. Results are plotted against λ , the fraction of wealth belonging to ESG investors, and for different values of Δ , the maximum certain return an ESG investor would sacrifice to invest in her optimal portfolio instead of the market portfolio.

Panel A. The Role of λ



Panel B. The Role of Δ





Fig. 3. Alphas of ESG investors. This figure plots the alpha for the portfolio held by ESG investors as a function of λ , the fraction of wealth belonging to ESG investors, and Δ , the maximum certain return an ESG investor would sacrifice to invest in her optimal portfolio instead of the market portfolio. Panel A plots the ESG alpha as a function of λ for four different values of Δ ; Panel B flips the roles of λ and Δ . The dashed line in Panel B has a slope of -1 . The differences between the solid lines and the dashed line represent investor surplus.

Panel A. The Role of λ



Panel B. The Role of Δ





Fig. 4. Alphas of non-ESG investors. This figure plots the alpha for the portfolio held by non-ESG investors as a function of λ , the fraction of wealth belonging to ESG investors, and Δ , the maximum certain return an ESG investor would sacrifice to invest in her optimal portfolio instead of the market portfolio. Panel A plots the ESG alpha as a function of λ for four different values of Δ ; Panel B flips the roles of λ and Δ .

Panel A. The Role of λ



Panel B. The Role of Δ




Fig. 5. Size of the ESG industry. The figure plots the aggregate dollar size of ESG investors’ deviations from the market portfolio (the ESG “tilt”), expressed as a fraction of the market’s total capitalization. In Panel A, results are plotted against λ , the fraction of wealth belonging to ESG investors, and for different values of Δ , the maximum certain return an ESG investor would sacrifice to invest in her optimal portfolio instead of the market portfolio. In Panel B, results are plotted against Δ and for different values of λ .





Fig. 6. Firm-level social impact. This figure plots $S_n(\bar{d}) - S_n(0)$, the social impact induced by ESG-motivated investors, for different firms n . The horizontal axis indicates the firm's initial ESG characteristic, $g_{n,0}$. The two regions indicate the components of $S_n(\bar{d}) - S_n(0)$ from Eq. (67). This figure uses the same parameters as the previous figures, with $\lambda = 0.5$ and $\Delta = 0.02$, as well as $r_f = 0.02$, $K_{0,n} = 1$, $\Pi_n = 1.2$, $\chi_n = 0.5$, and $\kappa_n = 1$. These parameter values produce $\bar{d} = 0.0113$, $\Delta g_n(\bar{d}) = 0.0113$, $\Delta K_n(0) = 0.0909$, and $\Delta K_n(\bar{d})$ ranging from 0.0813 to 0.1007.

Panel A. The Role of λ



Panel B. The Role of Δ




Fig. 7. Aggregate social impact. The figure plots $S(\bar{d}) - S(0)$, the aggregate social impact induced by ESG-motivated investors. We assume the firms' initial ESG characteristics $g_{0,n}$ are uniformly distributed in $[-\sqrt{3}, \sqrt{3}]$. (These endpoints maintain $g'_0 \iota = 0$ and $g'_0 g_0 = 1$.) The two colored regions indicate the components of $S_n(\bar{d}) - S_n(0)$ from Eq. (67), aggregated across firms. In Panel A, results are plotted against λ , the fraction of wealth belonging to ESG investors, assuming $\Delta = 0.02$. In Panel B, results are plotted against Δ , the maximum certain return an ESG investor would sacrifice to invest in her optimal portfolio instead of the market portfolio, assuming $\lambda = 0.5$. All remaining parameter values are the same as in Fig. 6.

Appendix. Proofs and Derivations

Derivation of Eq. (4):

To compute agent i 's expected utility, we rely on Eq. (2), the relation $\tilde{W}_{1i} = W_{0i}(1 + r_f + X_i'\tilde{r})$, and the fact that \tilde{r} is normally distributed, $\tilde{r} \sim N(\mu, \Sigma)$:

$$\begin{aligned}
\mathbb{E} \left\{ V(\tilde{W}_{1i}, X_i) \right\} &= \mathbb{E} \left\{ -e^{-A_i \tilde{W}_{1i} - b_i' X_i} \right\} \\
&= \mathbb{E} \left\{ -e^{-A_i [W_{0i}(1+r_f+X_i'\tilde{r})] - b_i' X_i} \right\} \\
&= -e^{-a_i(1+r_f)} \mathbb{E} \left\{ e^{-a_i X_i' [\tilde{r} + \frac{1}{a_i} b_i]} \right\} \\
&= -e^{-a_i(1+r_f)} e^{-a_i X_i' [E(\tilde{r}) + \frac{1}{a_i} b_i] + \frac{1}{2} a_i^2 X_i' \text{Var}(\tilde{r}) X_i} \\
&= -e^{-a_i(1+r_f)} e^{-a_i X_i' [\mu + \frac{1}{a_i} b_i] + \frac{1}{2} a_i^2 X_i' \Sigma X_i}
\end{aligned} \tag{A1}$$

where $a_i \equiv A_i W_{0i}$ is agent i 's relative risk aversion. Agents take μ and Σ as given. Differentiating with respect to X_i , we obtain the first-order condition

$$-a_i \left[\mu + \frac{1}{a_i} b_i \right] + \frac{1}{2} a_i^2 (2 \Sigma X_i) = 0, \tag{A2}$$

from which we obtain agent i 's portfolio weights

$$X_i = \frac{1}{a_i} \Sigma^{-1} \left(\mu + \frac{1}{a_i} b_i \right). \tag{A3}$$

Derivation of Eq. (5):

The n th element of agent i 's portfolio weight vector, X_i , is given by

$$X_{i,n} = \frac{W_{0i,n}}{W_{0i}} \tag{A4}$$

where $W_{0i,n}$ is the dollar amount invested by agent i in stock n . Let $W_{0,n} \equiv \int_i W_{0i,n} di$ denote the total amount invested in stock n by all agents. Then the n th element of the market-weight vector, w_m , is given by

$$w_{m,n} = \frac{W_{0,n}}{W_0} = \frac{1}{W_0} \int_i W_{0i,n} di = \frac{1}{W_0} \int_i W_{0i} X_{i,n} di = \int_i \frac{W_{0i}}{W_0} X_{i,n} di = \int_i \omega_i X_{i,n} di \tag{A5}$$

Note that $\sum_{n=1}^N w_{m,n} = 1$ because $\sum_{n=1}^N W_{0,n} = W_0$, which follows from the riskless asset being in zero net supply. Plugging in for X_i from Eq. (A3) and imposing $a_i = a$, we have

$$\begin{aligned}
x &= \int_i \omega_i X_i di \\
&= \int_i \omega_i \left[\frac{1}{a} \Sigma^{-1} \left(\mu + \frac{1}{a} b_i \right) \right] di \\
&= \frac{1}{a} \Sigma^{-1} \mu \left(\int_i \omega_i di \right) + \frac{1}{a^2} \Sigma^{-1} g \left(\int_i \omega_i d_i di \right) \\
&= \frac{1}{a} \Sigma^{-1} \mu + \frac{\bar{d}}{a^2} \Sigma^{-1} g.
\end{aligned} \tag{A6}$$

Proof of the statement in footnote 4:

Let \tilde{g}_{in} denote agent i 's perceived ESG characteristic of firm n , known by all agents. Eq. (3) changes to $b_i = d_i \tilde{g}_i$, where \tilde{g}_i is an agent-specific $N \times 1$ vector containing the values of \tilde{g}_{in} . Eq. (5) is unchanged, with g redefined as

$$\begin{aligned} g &= (1/\bar{d}) \int_i \omega_i d_i \tilde{g}_i di \\ &= E_\omega[\tilde{g}_i] + \text{Cov}_\omega(d_i/\bar{d}, \tilde{g}_i), \end{aligned} \quad (\text{A7})$$

where E_ω and Cov_ω denote the wealth-weighted expectation and covariance, respectively, across agents. The first term on the right-hand side of Eq. (A7) is an $N \times 1$ vector whose n th element is the wealth-weighted average of \tilde{g}_{in} across agents. The second term is a vector whose n th element is the wealth-weighted covariance between agents' scaled ESG tastes, d_i/\bar{d} , and perceived ESG characteristics, \tilde{g}_{in} . It seems plausible to assume that the second term is a zero vector, but we do not need to make that assumption. Since Eq. (5) is unchanged, Eqs. (6) through (9) are also unchanged.

Derivation of Eq. (11):

Agent i 's expected excess return is given by $E(\tilde{r}_i) = X_i' \mu$. We take μ from Eq. (9) and express X_i in terms of w_m by subtracting Eq. (5) from Eq. (4). Recalling the assumption $w_m' g = 0$ from Eq. (8), we obtain agent i 's expected excess return as

$$\begin{aligned} E(\tilde{r}_i) &= X_i' \mu \\ &= \left[w_m' + \frac{\delta_i}{a^2} g' \Sigma^{-1} \right] \left[\mu_m \beta_m - \frac{\bar{d}}{a} g \right] \\ &= \left[w_m' + \frac{\delta_i}{a^2} g' \Sigma^{-1} \right] \left[\frac{\mu_m}{\sigma_m^2} \Sigma w_m - \frac{\bar{d}}{a} g \right] \\ &= \mu_m - \frac{\bar{d}}{a} w_m' g + \frac{\delta_i \mu_m}{a^2 \sigma_m^2} g' w_m - \frac{\delta_i \bar{d}}{a^3} g' \Sigma^{-1} g \\ &= \mu_m - \frac{\delta_i \bar{d}}{a^3} g' \Sigma^{-1} g. \end{aligned} \quad (\text{A8})$$

Derivation of Eq. (12):

Recall that agent i 's excess portfolio return is $\tilde{r}_i = X_i' \tilde{r}$, where $\tilde{r} \sim N(\mu, \Sigma)$. Therefore,

$$\begin{aligned} \text{Var}(\tilde{r}_i) &= X_i' \Sigma X_i \\ &= \left[w_m' + \frac{\delta_i}{a^2} g' \Sigma^{-1} \right] \Sigma \left[w_m + \frac{\delta_i}{a^2} \Sigma^{-1} g \right] \\ &= w_m' \Sigma w_m + \frac{\delta_i}{a^2} g' \Sigma^{-1} \Sigma w_m + w_m' \Sigma \frac{\delta_i}{a^2} \Sigma^{-1} g + \frac{\delta_i^2}{a^4} g' \Sigma^{-1} \Sigma \Sigma^{-1} g \\ &= w_m' \Sigma w_m + \frac{\delta_i}{a^2} g' w_m + \frac{\delta_i}{a^2} w_m' g + \frac{\delta_i^2}{a^4} g' \Sigma^{-1} g. \end{aligned} \quad (\text{A9})$$

Recognizing that $w'_m \Sigma w_m = \sigma_m^2$ and $w'_m g = 0$, we have

$$\text{Var}(\tilde{r}_i) = \sigma_m^2 + \frac{\delta_i^2}{a^4} g' \Sigma^{-1} g, \quad (\text{A10})$$

which is Eq. (12). We see that $\text{Var}(\tilde{r}_i) > \sigma_m^2$ as long as $\delta_i \neq 0$.

Derivation of Eq. (13):

The second exponent in agent i 's expected utility in Eq. (A1) contains the terms $-aX'_i \mu$, $-X'_i b_i$, and $(a^2/2)X'_i \Sigma X_i$. The first of these is simply minus a times the expression in Eq. (A8). The second is given by

$$\begin{aligned} -X'_i b_i &= - \left[w'_m + \frac{\delta_i}{a^2} g' \Sigma^{-1} \right] [d_i g] \\ &= - \frac{d_i \delta_i}{a^2} g' \Sigma^{-1} g, \end{aligned} \quad (\text{A11})$$

and the third is given by

$$\begin{aligned} \frac{a^2}{2} X'_i \Sigma X_i &= \frac{a^2}{2} \left[w'_m + \frac{\delta_i}{a^2} g' \Sigma^{-1} \right] \Sigma \left[w_m + \frac{\delta_i}{a^2} \Sigma^{-1} g \right] \\ &= \frac{a^2}{2} \sigma_m^2 + \frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g, \end{aligned} \quad (\text{A12})$$

recalling $w'_m g = 0$ in both cases. Adding the three terms then gives

$$\begin{aligned} -aX'_i \mu - X'_i b_i + (a^2/2)X'_i \Sigma X_i &= -a\mu_m + \frac{\delta_i \bar{d}}{a^2} g' \Sigma^{-1} g - \frac{d_i \delta_i}{a^2} g' \Sigma^{-1} g + \frac{a^2}{2} \sigma_m^2 + \frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g \\ &= -a\mu_m + \frac{a^2}{2} \sigma_m^2 + \frac{1}{a^2} \left(\delta_i \bar{d} - d_i \delta_i + \frac{1}{2} \delta_i^2 \right) g' \Sigma^{-1} g \\ &= -a \left(\mu_m - \frac{a}{2} \sigma_m^2 \right) - \frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g. \end{aligned} \quad (\text{A13})$$

Substituting this exponent into Eq. (A1) gives

$$\begin{aligned} \text{E} \left\{ V(\tilde{W}_{1i}, X_i) \right\} &= -e^{-a(1+r_f)} e^{-a(\mu_m - \frac{a}{2} \sigma_m^2) - \frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g} \\ &= \left[-e^{-a(1+r_f)} e^{-a(\mu_m - \frac{a}{2} \sigma_m^2)} \right] e^{-\frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g} \\ &= \bar{V} e^{-\frac{\delta_i^2}{2a^2} g' \Sigma^{-1} g}, \end{aligned} \quad (\text{A14})$$

noting that the bracketed term is, \bar{V} , the agent's expected utility if $\delta_i = 0$.

Derivation of Eq. (42):

The assumptions below Eq. (36), along with Eq. (40), imply that the covariance matrix of \tilde{r} is of the form

$$\Sigma = B\Omega B' + \sigma_\zeta^2 I, \quad (\text{A15})$$

in which $B = [\beta_m \ g]$, and both $B'B$ and Ω are diagonal matrices:

$$B'B = \begin{bmatrix} \beta_m' \beta_m & 0 \\ 0 & g'g \end{bmatrix}, \quad \Omega = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}. \quad (\text{A16})$$

Inverting Σ using the Woodbury identity gives

$$\Sigma^{-1} = \frac{1}{\sigma_\zeta^2} I - \frac{1}{\sigma_\zeta^4} B\Theta B', \quad (\text{A17})$$

in which Θ is the diagonal matrix,

$$\Theta = \begin{bmatrix} \theta_m & 0 \\ 0 & \theta_g \end{bmatrix} = \left[\Omega^{-1} + \frac{1}{\sigma_\zeta^2} B'B \right]^{-1}, \quad (\text{A18})$$

and

$$\theta_g = \frac{\sigma_2^2 \sigma_\zeta^2}{\sigma_\zeta^2 + g'g \sigma_2^2}. \quad (\text{A19})$$

Post-multiplying the right-hand side of Eq. (A17) by g and recognizing that $g'\beta_m = 0$ gives

$$\Sigma^{-1}g = \left(\frac{1}{\sigma_\zeta^2} - \frac{\theta_g g'g}{\sigma_\zeta^4} \right) g = \left(\frac{1}{\sigma_\zeta^2 + g'g \sigma_2^2} \right) g. \quad (\text{A20})$$

If $l'g \neq 0$, then from Eqs. (17) through (21) and the surrounding discussion, $w_g = g/(l'g)$, and $g_g = (g'g)/(l'g)$. If $l'g = 0$, then w_g equals the right-hand side of Eq. (A20), and $g_g = (g'g)/(\sigma_\zeta^2 + g'g \sigma_2^2)$. In either case, $(1/g_g)w_g = g/(g'g)$. The definition of \tilde{f}_g in Eq. (31) implies that pre-multiplying the right-hand side of Eq. (40) by $(1/g_g)w_g'$, recalling that $w_g'\beta_m = 0$, gives

$$\tilde{f}_g - E_0\{f_g\} = \tilde{f}_g^e + \tilde{\xi}, \quad (\text{A21})$$

with $\tilde{\xi} = (g'\tilde{\zeta})/(g'g)$. The variance of $\tilde{\xi}$ is $\sigma_\zeta^2/(g'g)$, which goes to zero as $N \rightarrow \infty$ if the cross-sectional second moment of the elements of g , $(g'g)/N$, is bounded below by a positive value for all N .

Derivation of Eq. (45):

First note that Σ is of the same form as in Eq. (A15), with the relabelings $\eta^2 = \sigma_\zeta^2$, $\sigma_m = \sigma_1$, and $\sigma_f = \sigma_2$. Also, as there, $\beta_m'g = 0$. Therefore, using Eq. (A20),

$$\Sigma^{-1}g = \left(\frac{1}{\eta^2 + g'g \sigma_f^2} \right) g = \left(\frac{1}{\eta^2 + N \sigma_f^2} \right) g. \quad (\text{A22})$$

Using Eq. (A22) and noting $g'g = N$, observe that for large N ,

$$g'\Sigma^{-1}g = \left(\frac{1}{\eta^2 + N\sigma_f^2} \right) g'g = \frac{1}{\sigma_f^2}. \quad (\text{A23})$$

Also observe that, because $\iota'\Sigma^{-1}g = 0$ (recall $\iota'g = 0$), the ESG portfolio has zero cost and weights $w_g = \Sigma^{-1}g$ as in Eq. (18). By Eq. (A22), the ESG portfolio goes long green stocks and short brown stocks. The greenness of the ESG portfolio is given by

$$g_g = w'_g g = g'\Sigma^{-1}g = \frac{1}{\sigma_f^2}. \quad (\text{A24})$$

The portfolio weights for each type of investor follow directly from Eq. (14), with $\delta_i = (1-\lambda)d$ for an ESG investor and $\delta_i = -\lambda d$ for a non-ESG investor:

$$X_{esg} = w_m + (1-\lambda)\frac{d}{a^2}\Sigma^{-1}g \quad (\text{A25})$$

$$X_{non} = w_m - \lambda\frac{d}{a^2}\Sigma^{-1}g. \quad (\text{A26})$$

Therefore, using Eq. (A22), the ESG investor's portfolio weights in Eq. (A25) become

$$X_{esg} = \frac{1}{N}\iota + (1-\lambda)\frac{d}{a^2}\Sigma^{-1}g = \frac{1}{N}(\iota + hg), \quad (\text{A27})$$

with

$$h = \frac{(1-\lambda)d}{a^2(\eta^2/N + \sigma_f^2)}, \quad (\text{A28})$$

which, as N grows large, converges to

$$h = \frac{(1-\lambda)d}{a^2\sigma_f^2}. \quad (\text{A29})$$

With expected utility as given by Eq. (A1), an ESG investor's certainty equivalent excess return from holding her optimal portfolio is

$$\begin{aligned} r_{esg}^* &= X'_{esg} \left(\mu + \frac{d}{a}g \right) - \frac{a}{2}X'_{esg}\Sigma X_{esg} \\ &= \frac{1}{N}(\iota + hg)' \left(\mu_m\beta_m - \frac{\lambda d}{a}g + \frac{d}{a}g \right) - \frac{a}{2}X'_{esg}\Sigma X_{esg} \\ &= \mu_m + \frac{h(1-\lambda)d}{a} - \frac{a}{2}X'_{esg}\Sigma X_{esg}. \end{aligned} \quad (\text{A30})$$

Recall that δ_i for the ESG investor is $(1-\lambda)d$, and thus the variance of the ESG investor's portfolio return, using Eq. (12), is

$$X'_{esg}\Sigma X_{esg} = \sigma_m^2 + \frac{(1-\lambda)^2 d^2}{a^4 \sigma_f^2}. \quad (\text{A31})$$

Combining Eqs. (A29), (A30), and (A31), we then see

$$r_{esg}^* = \mu_m - \frac{a}{2}\sigma_m^2 + \frac{(1-\lambda)^2 d^2}{2a^3\sigma_f^2}. \quad (\text{A32})$$

If the ESG investor is instead constrained to hold the market portfolio, the resulting certainty equivalent excess return is given by

$$r_m^* = w'_m \mu - \frac{a}{2} w'_m \Sigma w_m = \mu_m - \frac{a}{2} \sigma_m^2. \quad (\text{A33})$$

The ESG investor's certainty-equivalent gain from investing as desired, versus investing in the market, is therefore

$$r_{esg}^* - r_m^* = \frac{(1-\lambda)^2 d^2}{2a^3\sigma_f^2}. \quad (\text{A34})$$

This difference in certainty equivalents is largest when $\lambda = 0$. That largest difference, Δ , is therefore

$$\Delta = \frac{d^2}{2a^3\sigma_f^2}, \quad (\text{A35})$$

and substituting for σ_f^2 using Eq. (A24) gives Eq. (45). The corresponding value of d is

$$d = \sqrt{2\Delta a^3\sigma_f^2}. \quad (\text{A36})$$

Derivation of the certainty equivalent excess return of a non-ESG investor (Section 4.2):

Proceeding as above, the non-ESG investor's portfolio weights in Eq. (A26) become

$$X_{non} = \frac{1}{N}\iota - \lambda \frac{d}{a^2} \Sigma^{-1} g = \frac{1}{N}(\iota + kg), \quad (\text{A37})$$

with

$$k = -\frac{\lambda d}{a^2(\eta^2/N + \sigma_f^2)}. \quad (\text{A38})$$

Similarly, the variance of the non-ESG investor's portfolio return for large N is

$$X'_{non} \Sigma X_{non} = \sigma_m^2 + k^2 \sigma_f^2, \quad (\text{A39})$$

and a non-ESG investor's certainty equivalent excess return from holding his optimal portfolio is

$$\begin{aligned} r_{non}^* &= X'_{non} \mu - \frac{a}{2} X'_{non} \Sigma X_{non} \\ &= \frac{1}{N}(\iota + kg)' (\mu_m \beta_m - \frac{\lambda d}{a} g) - \frac{a}{2} X'_{non} \Sigma X_{non} \\ &= \mu_m - \frac{a}{2} \sigma_m^2 + \frac{\lambda^2 d^2}{2a^3 \sigma_f^2} \\ &= r_m^* + \frac{\lambda^2 d^2}{2a^3 \sigma_f^2}. \end{aligned} \quad (\text{A40})$$

Derivation of Eq. (46):

From Eqs. (11) and (A23), the difference in expected excess returns earned by the two types of investors is

$$\mathbb{E}(\tilde{r}_{esg}) - \mathbb{E}(\tilde{r}_{non}) = -\frac{\lambda d^2}{a^3} g' \Sigma^{-1} g = -\frac{\lambda d^2}{a^3 \sigma_f^2}. \quad (\text{A41})$$

Substituting for d^2 from Eq. (A36) gives Eq. (46).

Derivation of Eq. (47):

The covariance between the ESG investor's return and the market return, using Eq. (A27) is

$$\begin{aligned} X'_{esg} \Sigma w_m &= \frac{1}{N^2} (\iota + hg)' (\sigma_m^2 \beta_m \beta_m' + \sigma_f^2 gg' + \eta^2 I_N) \iota \\ &= \sigma_m^2 + \frac{1}{N} \eta^2, \end{aligned} \quad (\text{A42})$$

which equals σ_m^2 for large N . Combining this result with Eq. (A31) gives the correlation between the ESG investor's return and the market return as

$$\begin{aligned} \rho(\tilde{r}_{esg}, \tilde{r}_m) &= \frac{X'_{esg} \Sigma w_m}{\sigma_m \sqrt{X'_{esg} \Sigma X_{esg}}} \\ &= \frac{\sigma_m}{\sqrt{\sigma_m^2 + \frac{(1-\lambda)^2 d^2}{a^4 \sigma_f^2}}} \end{aligned} \quad (\text{A43})$$

Substituting for d^2 from Eq. (A36) gives Eq. (47).

Derivations of Eqs. (48) and (49):

Let α denote the $N \times 1$ vector of alphas given by Eq. (10). The alpha of the ESG investor is given by

$$\begin{aligned} \alpha_{esg} &= X'_{esg} \alpha \\ &= \frac{1}{N} (\iota + hg)' \left(-\frac{\lambda d}{a} g \right) \\ &= -\lambda(1-\lambda) \frac{d^2}{a^3 \sigma_f^2}, \end{aligned} \quad (\text{A44})$$

using Eqs. (A27) and (A29). Substituting for d^2 from Eq. (A36) gives Eq. (48). The wealth-weighted average alpha must equal zero,

$$\lambda \alpha_{esg} + (1-\lambda) \alpha_{non} = 0, \quad (\text{A45})$$

and applying that identity gives Eq. (49).

Derivation of Eq. (53):

Because $\iota'X_{esg} = 1$, ESG investors' stock portfolio weights, w_i , are simply X_{esg} from Eq. (A27). Using Eqs. (A27) and (A29), along with $w_m = (1/N)\iota$, gives

$$\begin{aligned} T &= \frac{1}{2}\lambda\iota'|X_{esg} - w_m| \\ &= \frac{1}{2}\lambda\iota'\left|\frac{1}{N}(\iota + hg) - \frac{1}{N}\iota\right| \\ &= \frac{1}{2}\lambda\iota'\left|\frac{(1-\lambda)d}{Na^2\sigma_f^2}g\right| \end{aligned} \quad (\text{A46})$$

Substituting for d from Eq. (A36), we obtain Eq. (53).

Derivation of Eq. (56):

Modifying the earlier derivation of Eq. (4), we obtain

$$\begin{aligned} \mathbb{E}\{V(\tilde{W}_{1i}, X_i, \tilde{C})\} &= -e^{-a_i(1+r_f)}\mathbb{E}\left\{e^{-a_iX_i'[\bar{r} + \frac{1}{a_i}b_i] - c_i\tilde{C}}\right\} \\ &= -e^{-a_i(1+r_f)}e^{-a_iX_i'[E(\bar{r}) + \frac{1}{a_i}b_i] + \frac{1}{2}a_i^2X_i'Var(\bar{\epsilon})X_i + a_i c_i X_i' Cov(\bar{\epsilon}, \tilde{C}) + \frac{1}{2}c_i^2 Var(\tilde{C})} \\ &= -e^{-a_i(1+r_f)}e^{-a_iX_i'[\mu + \frac{1}{a_i}b_i] + \frac{1}{2}a_i^2X_i'\Sigma X_i + a_i c_i X_i'\sigma_{\epsilon C} + \frac{1}{2}c_i^2\sigma_{\tilde{C}}^2}, \end{aligned} \quad (\text{A47})$$

where $\sigma_{\epsilon C} \equiv Cov(\bar{\epsilon}, \tilde{C})$. Differentiating with respect to X_i gives the first-order condition

$$-a_i\left[\mu + \frac{1}{a_i}b_i\right] + a_i^2\Sigma X_i + a_i c_i \sigma_{\epsilon C} = 0 \quad (\text{A48})$$

from which we obtain agent i 's portfolio weights

$$X_i = \frac{1}{a_i}\Sigma^{-1}\left(\mu + \frac{1}{a_i}b_i - c_i\sigma_{\epsilon C}\right). \quad (\text{A49})$$

Again imposing the market-clearing condition and $a_i = a$ gives

$$\begin{aligned} w_m &= \int_i \omega_i X_i di \\ &= \frac{1}{a}\Sigma^{-1}\mu + \frac{\bar{d}}{a^2}\Sigma^{-1}g - \frac{\bar{c}}{a}\Sigma^{-1}\sigma_{\epsilon C}, \end{aligned} \quad (\text{A50})$$

which implies

$$\mu = a\Sigma w_m - \frac{\bar{d}}{a}g + \bar{c}\sigma_{\epsilon C}. \quad (\text{A51})$$

Premultiplying by w_m' , again imposing the assumption $w_m'g = 0$, gives

$$\mu_m = a\sigma_m^2 + \bar{c}\sigma_{mC}, \quad (\text{A52})$$

where $\sigma_{mC} \equiv \text{Cov}(\tilde{\epsilon}_m, \tilde{C}) = w'_m \sigma_{\epsilon C}$. Solving Eq. (A52) for a and substituting into the first term on the right-hand side of Eq. (A51) gives

$$\begin{aligned} \mu &= \frac{\mu_m - \bar{c}\sigma_{mC}}{\sigma_m^2} \Sigma w_m - \frac{\bar{d}}{a} g + \bar{c}\sigma_{\epsilon C} \\ &= (\mu_m - \bar{c}\sigma_{mC})\beta_m - \frac{\bar{d}}{a} g + \bar{c}\sigma_{\epsilon C} \\ &= \mu_m\beta_m - \frac{\bar{d}}{a} g + \bar{c} \left(\sigma_{\epsilon C} - \frac{\sigma_{mC}}{\sigma_m^2} \sigma_{\epsilon m} \right), \end{aligned} \quad (\text{A53})$$

noting $\beta_m = (1/\sigma_m^2)\sigma_{\epsilon m} = (1/\sigma_m^2)\Sigma w_m$. To see that the third term on the right-hand side of Eq. (A53) is the same as that in Eq. (56), first observe that in the multivariate regression of $\tilde{\epsilon}$ on $\tilde{\epsilon}_m$ and \tilde{C} , the $N \times 2$ matrix of slope coefficients is given by

$$[\sigma_{\epsilon m} \ \sigma_{\epsilon C}] \begin{bmatrix} \sigma_m^2 & \sigma_{mC} \\ \sigma_{mC} & \sigma_C^2 \end{bmatrix}^{-1} = \frac{1}{\sigma_m^2 \sigma_C^2 - \sigma_{mC}^2} [\sigma_C^2 \sigma_{\epsilon m} - \sigma_{mC} \sigma_{\epsilon C} \quad \sigma_m^2 \sigma_{\epsilon C} - \sigma_{mC} \sigma_{\epsilon m}],$$

so the second column is given by

$$\psi = \frac{\sigma_m^2 \sigma_{\epsilon C} - \sigma_{mC} \sigma_{\epsilon m}}{\sigma_m^2 \sigma_C^2 - \sigma_{mC}^2}. \quad (\text{A54})$$

Using Eq. (A54), we can rewrite the third term on the right-hand side of Eq. (A53) as

$$\begin{aligned} \bar{c} \left(\sigma_{\epsilon C} - \frac{\sigma_{mC}}{\sigma_m^2} \sigma_{\epsilon m} \right) &= \bar{c} \frac{\sigma_m^2 \sigma_C^2 - \sigma_{mC}^2}{\sigma_m^2} \psi \\ &= \bar{c} (1 - \rho_{mC}^2) \psi, \end{aligned} \quad (\text{A55})$$

recalling that $\sigma_C = 1$.

Derivation of Eq. (57):

Substituting for μ from Eq. (A53) into Eq. (A49) and setting $a_i = a$, we obtain

$$\begin{aligned} X_i &= \frac{1}{a} \Sigma^{-1} \left(\mu + \frac{1}{a} b_i - c_i \sigma_{\epsilon C} \right) \\ &= \frac{1}{a} \Sigma^{-1} \left(\mu_m \beta_m - \frac{\bar{d}}{a} g + \bar{c} \left(\sigma_{\epsilon C} - \frac{\sigma_{mC}}{\sigma_m^2} \sigma_{\epsilon m} \right) + \frac{1}{a} b_i - c_i \sigma_{\epsilon C} \right) \\ &= \frac{\mu_m}{a} \Sigma^{-1} \beta_m - \frac{1}{a} \Sigma^{-1} \bar{c} \frac{\sigma_{mC}}{\sigma_m^2} \sigma_{\epsilon m} + \frac{1}{a} \Sigma^{-1} \left(\frac{d_i}{a} g - \frac{\bar{d}}{a} g \right) - \frac{1}{a} \Sigma^{-1} (c_i - \bar{c}) \sigma_{\epsilon C} \\ &= \frac{\mu_m}{a} \Sigma^{-1} \beta_m - \frac{1}{a} \Sigma^{-1} (\bar{c} \sigma_{mC}) \frac{\sigma_{\epsilon m}}{\sigma_m^2} + \frac{1}{a} \Sigma^{-1} \frac{\delta_i}{a} g - \frac{c_i - \bar{c}}{a} \Sigma^{-1} \sigma_{\epsilon C} \end{aligned} \quad (\text{A56})$$

Noting from Eq. (A52) that $\bar{c}\sigma_{mC} = \mu_m - a\sigma_m^2$, and that $\beta_m = \frac{\sigma_{\epsilon m}}{\sigma_m^2} = \frac{\Sigma w_m}{\sigma_m^2}$, we have

$$X_i = \frac{\mu_m}{a} \Sigma^{-1} \beta_m - \frac{1}{a} \Sigma^{-1} (\mu_m - a\sigma_m^2) \beta_m + \frac{\delta_i}{a^2} \Sigma^{-1} g - \frac{c_i - \bar{c}}{a} \Sigma^{-1} \sigma_{\epsilon C}$$

$$\begin{aligned}
&= \sigma_m^2 \Sigma^{-1} \beta_m + \frac{\delta_i}{a^2} \Sigma^{-1} g - \frac{c_i - \bar{c}}{a} \Sigma^{-1} \sigma_{\epsilon C} \\
&= w_m + \frac{\delta_i}{a^2} \Sigma^{-1} g - \frac{c_i - \bar{c}}{a} \Sigma^{-1} \sigma_{\epsilon C},
\end{aligned} \tag{A57}$$

which is Eq. (57).

Derivation of Eq. (64):

The firm's value at time 0 is

$$v_n = -\Delta K_n - \frac{\kappa_n}{2} (\Delta K_n)^2 + \frac{\Pi_n (K_{0,n} + \Delta K_n)}{1 + r_f + \mu_m \beta_n - \frac{\bar{d}}{a} g_n}. \tag{A58}$$

The manager maximizes v_n by choosing ΔK_n . The first-order condition yields

$$\Delta K_n(\bar{d}) = \frac{1}{\kappa_n} \left[\frac{\Pi_n}{1 + r_f + \mu_m \beta_n - \frac{\bar{d}}{a} g_n} - 1 \right]. \tag{A59}$$

Substituting into Eq. (63) produces

$$\begin{aligned}
S_n(\bar{d}) - S_n(0) &= g_n \frac{1}{\kappa_n} \left[\frac{\Pi_n}{1 + r_f + \mu_m \beta_n - \frac{\bar{d}}{a} g_n} - \frac{\Pi_n}{1 + r_f + \mu_m \beta_n} \right] \\
&= g_n \frac{\Pi_n}{\kappa_n} \left[\frac{\frac{\bar{d}}{a} g_n}{(1 + r_f + \mu_m \beta_n - \frac{\bar{d}}{a} g_n)(1 + r_f + \mu_m \beta_n)} \right],
\end{aligned} \tag{A60}$$

which produces Eq. (64). Comparative statics for Π_n , β_n , and κ_n follow immediately from Eq. (64). For the comparative statics for \bar{d} and a , we define $\tilde{d} \equiv \bar{d}/a$ and compute

$$\begin{aligned}
\frac{\partial}{\partial \tilde{d}} (S_n(\bar{d}) - S_n(0)) &= \frac{g_n^2 \Pi_n}{\kappa_n (1 + r_f + \mu_m \beta_n)} \left[\frac{(1 + r_f + \mu_m \beta_n - \tilde{d} g_n) + \tilde{d} g_n}{(1 + r_f + \mu_m \beta_n - \tilde{d} g_n)^2} \right] \\
&= \frac{g_n^2 \Pi_n}{\kappa_n} \left[\frac{1}{(1 + r_f + \mu_m \beta_n - \tilde{d} g_n)^2} \right],
\end{aligned} \tag{A61}$$

which is positive if $g_n \neq 0$. Since $S_n(\bar{d}) - S_n(0)$ increases in \tilde{d} , it increases in \bar{d} and decreases in a .

Derivation of Eqs. (65) and (66):

The firm's value at time 0 is now

$$v_n = -\Delta K_n - \frac{\kappa_n}{2} (\Delta K_n)^2 + \frac{\Pi_n (K_{0,n} + \Delta K_n) \left(1 - \frac{\chi_n}{2} (\Delta g_n)^2\right)}{1 + r_f + \mu_m \beta_n - \frac{\bar{d}}{a} (g_{n,0} + \Delta g_n)}. \tag{A62}$$

The manager maximizes v_n by choosing Δg_n and ΔK_n . The choice of Δg_n depends only on the third term of Eq. (A62), and we can maximize its log. Using the approximation that

$\log(1+x) \approx x$ and ignoring terms without Δg_n , the choice of Δg_n simplifies to

$$\max_{\Delta g_n} -\frac{\chi_n}{2}(\Delta g_n)^2 + \frac{\bar{d}}{a}\Delta g_n . \quad (\text{A63})$$

The first-order condition delivers Eq. (65). Without taking logs, the first-order condition for ΔK_n is

$$-1 - \kappa_n \Delta K_n + \frac{\Pi_n \left(1 - \frac{\chi_n}{2}(\Delta g_n)^2\right)}{1 + r_f + \mu_m \beta_n - \frac{\bar{d}}{a}g_n} = 0 , \quad (\text{A64})$$

which delivers Eq. (66).

References

- Albuquerque, R., Koskinen, Y., Zhang, C., 2019. Corporate social responsibility and firm risk: theory and empirical evidence. *Management Science* 65, 4451–4469.
- Baker, M., Bergstresser, D., Serafeim, G., Wurgler, J., 2018. Financing the response to climate change: the pricing and ownership of U.S. green bonds. Unpublished working paper. Harvard University.
- Baker, M., Wurgler, J., 2012. Behavioral corporate finance: an updated survey. In: Constantinides, G., Harris, M., Stulz, R. (Eds.), *Handbook of the Economics of Finance*, Vol. 2. Elsevier, Amsterdam, pp. 357–424.
- Baker, S., Hollifield, B., Osambela, E., 2020. Asset prices and portfolios with externalities. Unpublished working paper. University of Virginia.
- Bank, M., Insam, F., 2017. Return contributions of factors and characteristics in an integrated asset pricing approach. Unpublished working paper. University of Innsbruck.
- Bansal, R., Ochoa, M., Kiku, D., 2016. Climate change and growth risks. Unpublished working paper. Duke University.
- Barber, B., Morse, A., Yasuda, A., 2020. Impact investing. *Journal of Financial Economics*, forthcoming.
- Bauer, R., Koedijk, K., Otten, R., 2005. International evidence on ethical mutual fund performance and investment style. *Journal of Banking and Finance* 29, 1751–1767.
- Berg, F., Koelbel, J., Rigobon, R., 2019. Aggregate confusion: the divergence of ESG ratings. Unpublished working paper. Massachusetts Institute of Technology.
- Bialkowski, J., Starks, L., 2016. SRI funds: investor demand, exogenous shocks and ESG profiles. Unpublished working paper. University of Canterbury.
- van Binsbergen, J., Opp, C., 2019. Real anomalies. *Journal of Finance* 74, 1659–1706.
- Bollen, N., 2007. Mutual fund attributes and investor behavior. *Journal of Financial and Quantitative Analysis* 42, 683–708.
- Bolton, P., Kacperczyk, M., 2019. Do investors care about carbon risk? Unpublished working paper. Columbia University.
- Chava, S., 2014. Environmental externalities and cost of capital. *Management Science* 60, 2223–2247.
- Cheng, I., Hong, H., Shue, K., 2016. Do managers do good with other peoples' money? Unpublished working paper. Dartmouth College.
- Cremers, M., Petajisto, A., 2009. How active is your fund manager? A new measure that predicts performance. *Review of Financial Studies* 22, 3329–3365.
- Dimson, E., Karakas, O., Li, X., 2015. Active ownership. *Review of Financial Studies* 28, 3225–3268.

- Dyck, A., Lins, K., Roth, L., Wagner, H., 2019. Do institutional investors drive corporate social responsibility? International evidence. *Journal of Financial Economics* 131, 693–714.
- Edmans, A., 2011. Does the stock market fully value intangibles? Employee satisfaction and equity prices. *Journal of Financial Economics* 101, 621–640.
- El Ghoul, S., Guedhami, O., Kwok, C., Mishra, D., 2011. Does corporate social responsibility affect the cost of capital? *Journal of Banking and Finance* 35, 2388–2406.
- Engle, R., Giglio, S., Kelly, B., Lee, H., Stroebel, J., 2020. Hedging climate change news. *Review of Financial Studies* 33, 1184–1216.
- Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3–56.
- Fama, E., French, K., 2007. Disagreement, tastes, and asset prices. *Journal of Financial Economics* 83, 667–689.
- Fama, E., Schwert, G., 1977. Human capital and capital market equilibrium. *Journal of Financial Economics* 4, 95–125.
- Friedman, H., Heinle, M., 2016. Taste, information, and asset prices: implications for the valuation of CSR. *Review of Accounting Studies* 21, 740–767.
- Garleanu, N., Kogan, L., Panageas, S., 2012. Displacement risk and asset returns. *Journal of Financial Economics* 105, 491–510.
- Geczy, C., Stambaugh, R., Levin, D., 2005. Investing in socially responsible mutual funds. Unpublished working paper. University of Pennsylvania.
- Gollier, C., Pouget, S., 2014. The “washing machine”: investment strategies and corporate behavior with socially responsible investors. Unpublished working paper. Toulouse School of Economics.
- Gompers, P., Ishii, J., Metrick, A., 2003. Corporate governance and equity prices. *Quarterly Journal of Economics* 118, 107–156.
- Hart, O., Zingales, L., 2017. Companies should maximize shareholder welfare not market value. *Journal of Law, Finance, and Accounting* 2, 247–274.
- Hartzmark, S., Sussman, A., 2019. Do investors value sustainability? A natural experiment examining ranking and fund flows. *Journal of Finance* 74, 2789–2837.
- Heinkel, R., Kraus, A., Zechner, J., 2001. The effect of green investment on corporate behavior. *Journal of Financial and Quantitative Analysis* 36, 431–449.
- Hoepner, A., Oikonomou, I., Sautner, Z., Starks, L., Zhou, X., 2018. ESG shareholder engagement and downside risk. Unpublished working paper. University College Dublin.
- Hong, H., Kacperczyk, M., 2009. The price of sin: the effects of social norms on markets. *Journal of Financial Economics* 93, 15–36.
- Hong, H., Kostovetsky, L., 2012. Red and blue investing: values and finance. *Journal of*

- Financial Economics 103, 1–19.
- Hong, H., Li, F., Xu, J., 2019. Climate risks and market efficiency. *Journal of Econometrics* 208, 265–281.
- Hong, H., Wang, N., Yang, J., 2020. Implications of mitigating disaster risks for sustainable investing and growth. Unpublished working paper. Columbia University.
- Humphrey, J., Kogan, S., Sagi, J., Starks, L., 2020. The asymmetry in responsible investing preferences. Unpublished working paper. University of Queensland.
- Ilhan, E., Sautner, Z., Vilkov, G., 2020. Carbon tail risk. Unpublished working paper. Frankfurt School of Finance and Management.
- Kempf, A., Osthoff, P., 2007. The effect of socially responsible investing on portfolio performance. *European Financial Management* 13, 908–922.
- Krueger, P., Sautner, Z., Starks, L., 2019. The importance of climate risks for institutional investors. *Review of Financial Studies*, forthcoming.
- Luo, A., Balvers, R., 2017. Social screens and systematic investor boycott risk. *Journal of Financial and Quantitative Analysis* 52, 365–399.
- Mayers, D., 1972. Nonmarketable assets and capital market equilibrium under uncertainty. In: Jensen, M. (Ed.), *Studies in the Theory of Capital Markets*. Praeger, New York, pp. 223–248.
- Oehmke, M., Opp., M., 2020. A theory of socially responsible investment. Unpublished working paper. London School of Economics.
- Pástor, Ľ., Stambaugh, R., 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111, 642–685.
- Pedersen, L., Fitzgibbons, S., Pomorski, L., 2019. Responsible investing: the ESG-efficient frontier. *Journal of Financial Economics*, forthcoming.
- Renneboog, L., Ter Horst, J., Zhang, C., 2008. Socially responsible investments: institutional aspects, performance, and investor behavior. *Journal of Banking and Finance* 32, 1723–1742.
- Renneboog, L., Ter Horst, J., Zhang, C., 2011. Is ethical money financially smart? Nonfinancial attributes and money flows of socially responsible investment funds. *Journal of Financial Intermediation* 20, 562–588.
- Riedl, A., Smeets, P., 2017. Why do investors hold socially responsible mutual funds? *Journal of Finance* 72, 2505–2550.
- Zerbib, O., 2019. The effect of pro-environmental preferences on bond prices: evidence from green bonds. *Journal of Banking and Finance* 98, 39–60.