



The impact of impact investing[☆]

Jonathan B. Berk^{a,b,*}, Jules H. van Binsbergen^{c,b,d}

^a Stanford University, United States of America

^b NBER, United States of America

^c University of Pennsylvania, United States of America

^d CEPR

ARTICLE INFO

Keywords:

ESG

Cost of Capital

Real Investment

Divestment

ABSTRACT

The change in the cost of capital that results from a divestiture strategy can be closely approximated by a simple function of three parameters: (1) the fraction of socially conscious capital, (2) the fraction of targeted firms in the economy and (3) the return correlation between the targeted firms and the rest of the stock market. When calibrated to current data, we demonstrate that the impact on the cost of capital is too small to meaningfully affect real investment decisions. We then derive the conditions that would be required for the strategy to have a meaningful impact. We empirically corroborate our theoretical results by studying firm changes in ESG status and are unable to detect an impact of ESG divestiture strategies on the cost of capital of treated firms. Our results suggest that to have impact, instead of divesting, socially conscious investors should invest and exercise their rights of control to change corporate policy.

Socially conscious investing aims to reduce social and environmental costs in society through investment choice. Divestment is an important tool investors employ to achieve this goal. Broadly defined, divestment is the idea that when investors refuse to invest in companies that impose social and environmental costs, society benefits. As plausible as this argument appears to be at first blush, it is missing the mechanism linking investor action and the effect on society. What is not clear is *how* divestment affects corporate strategy. Why would investors' choice to divest achieve the stated goal of reducing social costs in society?

When an investor chooses to sell a stock like a tobacco company, she necessarily sells it to another investor. Since this transaction simply exchanges one investor for another, it cannot directly impact how the company does business. There are, however, two possible ways the divestment decision can indirectly affect the company's business strategy. One way is for the new owners to exercise their rights of control. But since the impact of divesting is to effectively swap shareholders who care about social and environmental costs for other shareholders, it is hard to see why the new shareholders would be any more disposed to exercising their control rights for the good of society than the old ones.

Alternatively, when a socially conscious shareholder sells her stake to another shareholder, the new shareholder needs to be induced into buying shares in the company. This inducement comes in the form of a lower price. The lower price implies a higher cost of capital which affects the company's future real investment strategy by lowering the number of positive net present value (NPV) investment opportunities, thus lowering the company's growth rate. The opposite is true when a socially conscious shareholder chooses to buy. The extra demand for clean companies increases the price of those companies, lowering the cost of capital and thus increasing their growth rates. In the long term, socially desirable companies become a larger fraction of the economy at the expense of socially undesirable companies and the social and environmental costs on society are reduced. Therefore, for this mechanism to have impact, it is essential that the divestment strategy results in a large enough change in the cost of capital to materially affect the firm's investment opportunity set.

Our objective in this paper is to evaluate the impact of divestiture initiatives in secondary markets by determining whether or not they have materially affected the cost of capital and, if not, whether they

[☆] Dimitris Papanikolaou was the editor for this article. We thank Jiayi Cheng and Jiarui (Andy) Su for outstanding research assistance and Greg Brown, Greg Buchak, Peter DeMarzo, Tyler Muir, Arvind Krishnamurthy, Christian Opp, Dimitris Papanikolaou, Monika Piazzesi, Amit Seru and seminar participants at the Jackson Hole Finance Group, Stanford GSB, Amsterdam Business School, Erasmus University Rotterdam, Washington University, EDHEC, NYU Law, ICPM, Inquire Europe, the University of Bristol, the Guanghua School of Management, the University of Zurich, and the GMU LEC Research Roundtable for their helpful comments. The research associated with this writing was partially funded by the Law & Economics Center at the George Mason University Antonin Scalia Law School. The opinions expressed here are the authors' alone and do not represent the official position of the Law & Economics Center, George Mason University.

are likely to do so in the future. We begin by first calculating, given current market conditions, the predicted impact of socially responsible investing on the cost of capital. Under the common assumption of mean–variance preferences, we demonstrate that the change in the cost of capital can be closely approximated by a simple formula. If socially conscious investors choose to target a set of companies that make up f fraction of the economy and the correlation between these companies and the rest of the market is ρ , then a divestment strategy will lead to a change in the cost of capital of

$$\text{MRP} \times \left(\frac{\text{Socially Conscious Investor Wealth}}{\text{Rest of Investor Wealth}} \right) \times f \times (1 - \rho^2),$$

where MRP is the market risk premium. We estimate, using ESG mutual fund data, that f is 27% and ρ is 0.93. Further, we estimate that currently socially conscious wealth makes up less than 2% of stock market wealth in the United States. Using a market risk premium of 6% leads to an estimated change in the cost of capital of 0.44 basis points. Given the uncertainty in the capital budgeting process, half a basis point cannot meaningfully impact firms' investment strategies.

We then show that the data is consistent with this theoretical prediction. Even with the growth in the popularity of impact investing in the last 10 years, we find that the difference in the cost of capital between firms that are targeted for their social or environmental costs and firms that are not, is small and not statistically different from zero. Specifically, we study the effect of a firm either being included or excluded from the FTSE USA 4 Good Select Index. The stocks in this index are a subset of the FTSE USA Index, and so the inclusion and exclusion events are driven by changes in the ESG status of the firm. In addition, this index is replicated by the world's largest socially responsible index fund, the Vanguard FTSE Social Index Fund. The advantage of this approach over existing approaches is that our results do not rely on the assumption that the risk of the firm is correctly measured.

We also investigate the sensitivity of our results to the specific assumptions in our model. First, some readers are concerned that our baseline model does not feature investors who effectively do not participate in the price setting process because they are constrained (or mandated) to a passive buy-and-hold strategy. We show that at current ESG participation rates, to have a meaningful effect of 1% on the cost of capital, 98% of investors would need to be constrained in this way. Second, we do not explicitly model investor preference over social outcomes. Instead, we only consider the extreme case where socially conscious investors refuse to hold *any* firm that imposes social costs. In that sense, our results can be viewed as an upper bound on the effect of investors with “Warm Glow” preferences who trade off the financial benefits of investing in these firms against the social costs. Finally, we have calibrated the model to other indices like the MSCI ESG Aware Index, and have found quantitatively similar results.

We then consider the question of what it would take for a divestment strategy to successfully impact firm investment. To effect a more than 1% change in the cost of capital, impact investors would need to make up more than 80% of all investable wealth. Given the low likelihood of achieving such a high participation rate, the results in this paper raise the question of whether divestment in secondary markets is the most effective strategy to achieve social good.

The reason divestiture in secondary markets has so little impact is that stocks are highly substitutable, and socially costly stocks make up less than half of the economy. It therefore does not take much of a price change to induce an investor who does not care about the social costs to hold more of a stock than they otherwise would. To put this in the language of modern finance, when socially responsible investors divest, they must induce other investors to move away from their fully diversified portfolio. But because the fraction of stocks that are subject to divestment is small relative to the supply of investable capital, and diversified stock portfolios are highly correlated with each other, the new portfolio is only slightly less diversified than the old one. So the

new investors do not demand much of an increase in their expected return. Thus, the effect on the cost of capital is small. The only way to materially affect this basic trade-off is if a large majority of investors choose to divest. In this case, because the shareholders that must hold the targeted (dirty) stocks comprise only a small minority of shareholders, they must be induced to hold a more highly concentrated portfolio which materially affects their overall diversification. In equilibrium, this causes a larger price impact.

Contrast this strategy with an investment in the primary market. If an investor chooses to fund an otherwise unfunded negative NPV investment opportunity, the reaction of other investors is irrelevant. By funding the negative NPV investment, the investor can achieve social good even when all other investors are indifferent to social outcomes. Notice that this insight is limited to negative NPV primary market investments only. If an investor chooses to “divest” in primary markets by choosing not to fund a *positive* NPV investment opportunity (that has social costs) she will have little effect if other investors exist who chase positive NPV projects.

By the same token, if the same agent engages in a negative NPV investment in the secondary market by overpaying for a stock, this extra value accrues to the seller of the stock and so this action has no direct effect on the firm's investment strategies. Similarly, if the agent decides to sell her stake in a company that engages in socially undesirable investments at a discount, the extra value will go to the person purchasing the stake. These secondary market actions can only indirectly affect social outcomes by changing the terms of trade. The extent to which the terms of trade are affected depends crucially on the reaction of other investors. If many other investors are willing to step in (because they are attracted to the positive NPV opportunity), the effect on the price will be minimal resulting in little impact on the terms of trade. This distinction between primary and secondary markets is crucial because it implies that a much higher degree of participation is required to achieve social change using the secondary market rather than the primary market.

In summary, because of the reaction of other investors, at current levels of participation, the impact of socially motivated divestment in the secondary market is small. Furthermore, the degree of participation that would be required to have meaningful impact is implausibly high. Our conclusion is that to have impact, socially conscious investors should use the primary market (what people often term “impact investing”) rather than the secondary market. Alternatively, instead of divesting, investors could invest and use their stakes to exercise their control rights. By taking a substantive equity position in firms that impose social costs, socially conscious investors could achieve social good through the proxy process or by gaining a majority stake and replacing upper management. We estimate that the set of companies that are targeted by all socially conscious funds comprise only 27% of the market, so successfully implementing this strategy would require a substantially smaller fraction of investable wealth than the fraction of wealth required for a divestiture strategy to succeed.

1. Background

This paper fits into a sizable literature that studies socially conscious investing.¹ In this paper, our focus is on divestment in secondary markets. One of the earliest papers that studies the effect of a divestment strategy is Teoh et al. (1999) who fail to detect any effect of divestment in South African stocks during apartheid. Later papers generally documented an excess return for holding “sin” stocks, usually companies in the tobacco, alcohol, fossil fuels, weapons and gaming industries, suggesting that divestment might have affected the returns in these

¹ Brest et al. (2018), in an article intended for a law review audience, survey the entire impact investing landscape. There is also a practitioner literature on the subject, see Cornell (2020) and the references therein.

industries.² More recent work has found contradictory results, that is, that “clean” stocks outperform “sin” stocks.³ The main issue with this line of research is that the stocks of interest are highly concentrated in a few industries that likely have risk-based reasons for returns different to the market. Thus the “sin” premium could well be attributable to risk not correctly adjusted for in these papers.⁴ There are two advantages of our empirical approach relative to this literature. First, our approach does not require us to take a stance on the risk model. All we require is the much weaker assumption that a decision by the firm to change its social policy is not correlated with changes in the firm’s riskiness. Second, we use the price response to a policy change as a measure of the effect on the cost of capital, which is easier to detect than the concomitant change in the expected return.

Other confounding factors to consider are that in the early sample period (1960’s to early 2000’s) there was very little organized pressure to divest (other than in the case of South Africa) while studies using more recent data have failed to find the effect (Mollet and Ziegler, 2014; Trinks et al., 2018). The effect is also not consistent across localities, as documented in Durand et al. (2013), who attribute the differences to “cultural norms”, and Feng et al. (2015).⁵ More fundamentally, given the level of uncertainty in returns, detecting changes in expected returns is challenging given the short time span over which socially conscious criteria have come to effect investment choices.

The first paper that theoretically models how a divestiture strategy affects corporate behavior is Heinkel et al. (2001). The paper models an equilibrium where the investment behavior of ESG investors raises the cost of capital of polluting firms and lowers the cost of capital for green firms. Firms then endogenously choose to become green by paying a cost. In equilibrium this cost leaves the marginal firm indifferent. The model provides a rich set of predictions and insights but, as is often the case with rich models, has the disadvantage that it does not yield a simple characterization of the equilibrium. Although the authors do not formally calibrate their model, they illustrate their equilibrium with a numerical example. Unfortunately, because the parameter choices in the numerical example were selected to best illustrate the tradeoffs in the model, the example appears to have left subsequent researchers with the impression that the effect on the cost of capital is large enough to induce a significant fraction of firms to choose to pay the cost to become green. Of course, at the time the paper was written, impact investing was in its infancy and so the data to properly calibrate their model likely did not exist. As we show in this paper, with the benefit of an extra twenty years of data, when properly calibrated to current market conditions, the effect on the cost of capital is too small to be consequential.⁶

Like us, Luo and Balvers (2017) study the theoretical effect of divestment in a single period mean–variance environment. They derive a boycott factor risk premium and show that it is positive. They do not calibrate the size of the premium. In the empirical section of their paper, they estimate the boycott factor risk premium to be 16% per annum. Our results suggest that this large estimated premium is unlikely to represent the effect of ESG preferences broadly and is likely more characteristic of the sample period and the choice of sin industries.

A number of other papers derive equilibrium models that feature ESG investors, although they do not focus on the question of

whether the actions of ESG investors achieve socially desirable outcomes. Pedersen et al. (2020) derive a model that includes investors whose preferences depend on ESG scores. The focus of that paper is characterizing the equilibrium and then using it to provide insight on optimal portfolio choice. In their setting, optimal portfolios can be spanned by four funds, thereby providing a practical methodology to choose optimal portfolios in a world that features ESG investors. The paper does not address the question of whether the actions of ESG investors actually achieve social good. Dam and Scholtens (2015) derive an equilibrium model of firm behavior, but since the cost of capital is assumed fixed at the risk free rate in their model, it is not obvious why firms care about investor preferences in that model. Presumably the authors have in mind that investors exercise their unmodeled control rights. Pástor et al. (2020) also derive a model featuring agents with ESG preferences and study the security market line in this environment. That paper does not assess the magnitude of the effect on the cost of capital of introducing ESG investors into the economy. Avramov et al. (2021) study the effect of ESG uncertainty on the ESG profile of firms in a similar model.

One implication of our paper is that investors are likely to be more effective in reducing social costs by investing, rather than divesting, in socially costly firms and using their rights of control to alter corporate policy. Dimson et al. (2015) provide empirical evidence supporting the effectiveness of this alternative strategy. Broccardo et al. (2020) explicitly model the choice between investing and divesting. Although their focus is slightly different from ours (Broccardo et al., 2020 study the individual decision to follow a particular investment strategy taking as given what other investors do) they reach the same conclusion we do. They show that the incentives of individuals to join an exit strategy are not necessarily aligned with social incentives, while they are aligned when investors invest and use their rights of control.

An important conclusion of our paper is that the primary market is likely to be a more effective place to achieve beneficial social outcomes. But with that said, competition in the primary market is an important determinant of the effectiveness of the strategy, a point emphasized in Oehmke and Opp (2023) who study the effectiveness of primary investment in a market with perfect investor competition so capital for positive NPV investment is in infinite supply. To achieve social good, investing in firms with no social externalities is suboptimal (because those projects would be financed anyway)—to achieve social change in equilibrium it is important for socially conscious investors to fund firms that are imposing social costs, that is, to fund more socially desirable (negative NPV) investments at firms that otherwise would impose higher social costs.

A recent paper that studies the effect of improving social outcomes by direct investment in the primary market is Dangi et al. (2023). The relative cost of capital between green and brown firms is unaffected by the social preferences of agents in that model. However, firms issue shares on the primary market to finance investment. Because firm managers in that model anticipate agents’ investment choices, socially conscious agents’ investment decisions in the primary market affect corporate policy. Consistent with this evidence, Krueger et al. (2020) survey institutional investors and find that they consider engagement rather than divestment to be a more effective approach to address climate risks.

Given the ineffectiveness of divestiture, a natural question that arises is why investors engage in the strategy at all. One possibility is that investors either do not realize that the strategy is ineffective or derive utility from the strategy without regard to its effectiveness. Another possibility is that investors use the strategy to signal “good behavior”. Riedl and Smeets (2017) provide evidence that supports this latter hypothesis.

Our work is closely related to the literature that studies index inclusion effects.⁷ At first blush our results might seem inconsistent with

² See Hong and Kacperczyk (2009), Chava (2014), Ghoul et al. (2011), Fabozzi et al. (2008), Statman and Glushkov (2016), Barber et al. (2021), Derwall and Verwijmeren (2007) and Salaber (2013) and references therein.

³ See, for example, Pastor et al. (2021) who argue that this empirical finding is just an ex-post realization rather than indicative of the true expected return.

⁴ See Blitz and Fabozzi (2017).

⁵ Renneboog et al. (2008) provide a review of the evidence in the literature on the effectiveness of impact investing.

⁶ In a followup paper, Barnea et al. (2005), endogenize firm investment and study how total investment is affected by divestiture.

⁷ See Shleifer (1986) and Harris and Gurel (1986).

this literature that documents seemingly large price changes resulting from demand effects. In fact, we show that our model is also able to explain the observed index inclusion effects. Greenwood and Sammon (2022) show that in the most recent decade, the price inclusion effect into the S&P500 Index is 0.3%.⁸ In Section 3 we calibrate our model to this number and show that the demand elasticities implied by our model are fully consistent with the index inclusion effects observed in the data. More importantly, the cost of capital is a long-term rate that is unaffected by short term price effects. Only the permanent component of a price change affects the cost of capital while many of the demand effects documented in the literature have been shown to be largely temporary — most of this effect mean reverts within a short period of time (as Harris and Gurel (1986) document).

2. The model

To assess the impact of divestment, we start with a frictionless economy in which the workhorse model in financial economics, the Capital Asset Pricing Model (CAPM) holds.⁹ We assume that investors are homogeneous with mean–variance preferences.¹⁰ That is, if R_p is the return of their portfolios and σ_p^2 is the variance of this return, then all investors select their portfolios by solving:

$$\max E[R_p] - k\sigma_p^2,$$

where k measures investors' risk appetites.

Although the CAPM applies to all investments, in this paper we will focus on the cost of capital implied by equity markets, which is where most of the ESG focus has been. Of course, firms have several other sources of financing, such as public debt markets, banks, as well as internally generated funds. Including other margins of adjustment in our computations would weaken the effects of equity divestment initiatives on real investment decisions. Thus, our results represent an upper bound on the effect of ESG investors in the equity market.

To model the divestiture strategy, we then allow a subset of investors to have ESG preferences in that they will only hold clean stocks, for simplicity, *ESG investors*.¹¹ We divide the stocks in the economy into two sets, the set of *clean* stocks that satisfy ESG requirements that ESG investors have, and the rest. For convenience we refer to the set of remaining stocks as *dirty* stocks, although it should be emphasized that the choice of this term is not supposed to be pejorative. Rather, the term is shorthand for the complement of the clean stocks.

In this single period economy, all investors are endowed with a share of the market portfolio. They trade at the beginning of the period and consume the liquidating dividend that stocks pay out at the end of the period. The cumulative dividend payout of all stocks is denoted D . We adopt the price normalization that the price of holding all stocks is 1, implying that the return of the market portfolio is $R = D - 1$. The market portfolio consists of two portfolios, the *clean* portfolio and the *dirty* portfolio. Let D_E be the cumulative liquidating dividend paid to all investors holding shares in the clean portfolio. Similarly, D_D is the cumulative liquidating dividend paid to all investors holding shares in the dirty portfolio, implying $D_E + D_D = D$. Denote the value at the beginning of the period of the clean stock portfolio (the total amount of wealth invested in the portfolio) as V_E and the value of the dirty

stock portfolio as V_D , implying, under the price normalization, that $V_E + V_D = 1$. Thus, under the price normalization, V_E can be interpreted as the fraction of market value that the clean stock portfolio makes up, and V_D is the fraction of market value that the dirty stock portfolio makes up. The return of these two portfolios is therefore $R_E = \frac{D_E}{V_E} - 1$ and $R_D = \frac{D_D}{V_D} - 1$ and the price normalization implies that the market return is

$$R = D_E + D_D - 1,$$

with expectation

$$\bar{R} = E[R] = \bar{D}_E + \bar{D}_D - 1,$$

where we adopt the notation throughout the paper that a bar over a random variable denotes its expectation.

Next, define the following second moment primitives:

$$\sigma_D^2 \equiv \text{var}(D_D)$$

$$\sigma_E^2 \equiv \text{var}(D_E)$$

$$\sigma_{ED} \equiv \text{cov}(D_E, D_D)$$

$$\sigma^2 \equiv \text{var}(D) = \text{var}(D_D + D_E) = \sigma_D^2 + \sigma_E^2 + 2\sigma_{ED} = \text{var}(R)$$

$$\rho \equiv \frac{\sigma_{ED}}{\sigma_D \sigma_E} = \frac{\text{cov}(R_E, R_D)}{\text{std}(R_D) \text{std}(R_E)}.$$

Notice that because of the price normalization, the cash flow variance of the market is the same as the return variance. Because of the nature of correlations, the cash flow correlation between dirty and clean stocks is the same as the return correlation. Finally, we assume that a risk free asset in zero net supply exists with return r which will be determined in equilibrium.

2.1. Equilibrium with ESG investors

We begin by solving for the equilibrium. The total wealth of ESG investors is denoted γ , implying under the price normalization, that the total wealth of the rest of investors is $1 - \gamma$. Because ESG investors will not hold dirty stocks, in equilibrium they hold the tangency portfolio of clean stocks on the constrained mean–variance frontier that includes only clean stocks. This tangency portfolio defines the clean stock portfolio, and the dirty portfolio is then defined as the portfolio that combined with the clean portfolio gives the market portfolio.¹² So, the ESG investors' equilibrium allocation is e fraction of the clean portfolio given by their budget constraint:

$$\gamma = V_E e + b,$$

where b is the total wealth ESG investors invest in the risk free asset (implying that other investors invest $-b$ in the risk free asset). Solving for e gives

$$e = \frac{\gamma - b}{V_E}. \quad (1)$$

Because ESG investors hold a constrained portfolio, the portfolio that prices assets is the portfolio held by the other investors. This portfolio, that is, the mean–variance efficient portfolio, is the tangency portfolio on the unconstrained mean–variance efficient frontier of risky assets. The rest of investors' wealth invested in risky assets is given by

$$V_D + (1 - e)V_E,$$

which implies that their portfolio weights are (using (1)):

$$\omega_D = \frac{V_D}{V_D + (1 - e)V_E} = \frac{V_D}{1 - \gamma + b}$$

⁸ Bennett et al. (2020) also investigate the impact on firms of joining the S&P 500 Index from 1997 to 2017. They find that the positive announcement effect on the stock price of index inclusion over this sample period has disappeared and even argue that the long-run impact of index inclusion has become negative.

⁹ The CAPM was developed by Sharpe (1964), Mossin (1966), Lintner (1965) and Treynor (1961).

¹⁰ Berk (1997) provides the necessary and sufficient conditions that imply mean–variance maximization.

¹¹ These preferences represent the extreme case — investors refuse to hold any stock that does social harm rather than just tilting away from such stocks.

¹² If the market portfolio is on the mean–variance efficient frontier of risky assets, then the dirty portfolio will contain only dirty stocks, and both portfolios will consist of the renormalized market weights of the clean and dirty sets. Otherwise, the dirty portfolio will also contain clean stocks.

in the dirty portfolio and

$$\omega_E = \frac{(1-e)V_E}{V_D + (1-e)V_E} = \frac{V_E - \gamma + b}{1 - \gamma + b}$$

in the clean portfolio. Thus the return of the mean-variance efficient portfolio, R_{mv} is

$$R_{mv} = \omega_D R_D + \omega_E R_E = \frac{D_D + D_E}{1 - \gamma + b} - 1 - \frac{D_E}{V_E} \left(\frac{\gamma - b}{1 - \gamma + b} \right) \quad (2)$$

$$= \frac{D_D + (1-e)D_E}{1 - \gamma + b} - 1. \quad (3)$$

Let $\Sigma_E \equiv \frac{\sigma_E^2 + \rho \sigma_D \sigma_E}{\sigma^2}$ and $\Sigma_D \equiv \frac{\sigma_D^2 + \rho \sigma_D \sigma_E}{\sigma^2}$ (the cash flow betas of the two portfolios), implying that $\Sigma_E + \Sigma_D = 1$. In [Appendix A](#) we derive the equilibrium returns in this economy. There we show that the return beta (with respect to the mean-variance efficient portfolio), β_E , of the clean portfolio is

$$\beta_E = \frac{1 - \gamma + b}{V_E} \left(\frac{\Sigma_E \sigma^2 - e \sigma_E^2}{\sigma^2 - 2e \Sigma_E \sigma^2 + e^2 \sigma_E^2} \right), \quad (4)$$

and the beta of the dirty portfolio, β_D is

$$\beta_D = \frac{1 - \gamma + b}{V_D} \left(\frac{\Sigma_D \sigma^2 - e \rho \sigma_D \sigma_E}{\sigma^2 - 2e \Sigma_E \sigma^2 + e^2 \sigma_E^2} \right). \quad (5)$$

Using these expressions and the expressions derived in [Appendix A](#) for the risk free rate, we derive, in the appendix, equilibrium prices using the security market line defined by the above mean-variance efficient portfolio:

$$V_E = \frac{\bar{D}_E - 2k\sigma^2 \Sigma_E}{\bar{D}_E + \bar{D}_D - 2k\sigma^2 (1 + \Gamma)}$$

$$V_D = \frac{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)}{\bar{D}_E + \bar{D}_D - 2k\sigma^2 (1 + \Gamma)}$$

where

$$\Gamma \equiv \left(\frac{\gamma}{1 - \gamma} \right) (1 - \rho^2) \frac{\sigma_D^2}{\sigma^2}. \quad (6)$$

Using these prices we then compute expected returns as a function of the market betas of the two portfolios, $\beta_D^m \equiv \frac{\Sigma_D}{V_D}$ and $\beta_E^m \equiv \frac{\Sigma_E}{V_E}$:

$$\bar{R}_E = \bar{R} + 2k\sigma^2 (\beta_E^m - (1 + \Gamma))$$

$$\bar{R}_D = \bar{R} + 2k\sigma^2 \left(\beta_D^m - \left(1 - \Gamma \frac{V_E}{V_D} \right) \right).$$

The difference in the expected return of dirty and clean stocks, $\Delta \bar{R}$, is then given by

$$\Delta \bar{R} \equiv \bar{R}_D - \bar{R}_E = 2k\sigma^2 \left(\beta_D^m - \beta_E^m + \frac{\Gamma}{V_D} \right). \quad (7)$$

2.2. Equilibrium without ESG investors

To assess the effect of ESG investors, we now derive the equilibrium when all investors are identical and do not have ESG preferences. To differentiate this equilibrium from the equilibrium with ESG investors, we will mark the equilibrium variables with asterisks. In this standard CAPM equilibrium, all investors choose to hold the market portfolio, which is mean-variance efficient. As we show in [Appendix B](#), the value and expected returns of the two portfolios are now given by

$$V_E^* = \frac{\bar{D}_E - 2k \Sigma_E \sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2}$$

$$V_D^* = \frac{\bar{D}_D - 2k \Sigma_D \sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2}$$

$$\bar{R}_E^* = \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_E}{V_E^*} \right)$$

$$\bar{R}_D^* = \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_D}{V_D^*} \right).$$

So the difference in the cost of capital between clean and dirty stocks is

$$\Delta \bar{R}^* \equiv \bar{R}_D^* - \bar{R}_E^* = 2k\sigma^2 (\beta_D^* - \beta_E^*), \quad (8)$$

where $\beta_D^* \equiv \frac{\Sigma_D}{V_D^*}$ and $\beta_E^* \equiv \frac{\Sigma_E}{V_E^*}$ are the (market) betas of the portfolios.

3. Effect on the cost of capital of ESG investors

To study the effect of introducing ESG investors, we will assume that all investors are initially identical and hold the market portfolio. A portion of them then acquire ESG preferences and trade to the ESG equilibrium. Eq. (7) is the difference in the cost of capital between clean and dirty stocks after a portion of investors acquire ESG preferences and (8) is the difference in the cost of capital before the existence of ESG investors. The difference between the two is therefore the effect of ESG investors on the cost of capital.

Although the market betas of the portfolios are not the same in the two economies, we show in [Appendix C](#) that this difference is second order, and so the difference in the cost of capital due to the presence of ESG investors is approximated by the difference between (7) and (8) assuming the betas are the same:

$$\begin{aligned} \Delta \bar{R} - \Delta \bar{R}^* &\approx 2k\sigma^2 \frac{\Gamma}{V_D} \\ &= 2k\sigma^2 V_D \left(\frac{\gamma}{1 - \gamma} \right) (1 - \rho^2) \frac{\sigma_D^2}{V_D^2 \sigma^2} \\ &= 2k\sigma^2 V_D \left(\frac{\gamma}{1 - \gamma} \right) (1 - \rho^2) \frac{\sigma_{R_D}^2}{\sigma^2} \\ &\approx 2k\sigma^2 V_D \left(\frac{\gamma}{1 - \gamma} \right) (1 - \rho^2), \end{aligned} \quad (9)$$

where $\sigma_{R_D}^2 \equiv \text{var}(R_D)$ is the return variance of the dirty portfolio and we have assumed that the variance of the market return and the return of the dirty portfolio is approximately the same. Now, in a standard CAPM equilibrium, $2k\sigma^2$ equals the market risk premium (see (B.3) in [Appendix B](#)), so if we assume that risk preferences have not changed over time then we can set $2k\sigma^2$ equal to the historical market risk premium (MRP). Under this assumption, (9) becomes

$$\Delta \bar{R} - \Delta \bar{R}^* \approx \text{MRP} \times V_D \times \left(\frac{\gamma}{1 - \gamma} \right) (1 - \rho^2). \quad (10)$$

In [Appendix C](#) we derive the exact relation. We show there why the above formula is an approximation of that relation and we will presently demonstrate that the approximation is very accurate.

The above approximation is quite informative on how impact investing affects the cost of capital. For impact investing to materially affect the cost of capital, three conditions need to be met. First, dirty stocks cannot be easily substituted for clean stocks. In any risk-based model, the degree of substitutability is measured by the correlation between clean and dirty stocks, ρ , so the term $1 - \rho^2$ measures the degree to which clean and dirty stocks are not substitutable. Second, impact investors must make up a significant fraction of investors, so the term $\frac{\gamma}{1 - \gamma}$, which is the ratio of the total wealth of investors with ESG preferences to other investors, measures the influence of ESG investors. Finally, because the non-ESG investors have limited wealth and must hold the dirty stocks in equilibrium, the greater the fraction of the economy that dirty stocks make up, the greater the price impact. This effect is measured by V_D , the fraction of the economy that dirty stocks make up. The product of these terms multiplied by a measure of investors' risk appetite, as implied by the historical market risk premium, gives the change in the cost of capital brought about by the choice of ESG investors to divest themselves of dirty stocks.

Because the quantities in (10) are observable, we can use this approximation to infer to what extent divestiture has altered the cost of capital in the United States. We use the Vanguard FTSE Social Index Fund (VFTSX) which replicates the FTSE USA 4 Good Select Index, to identify the clean portfolio. The Vanguard FTSE Social Index Fund




Fig. 1. Combined total AUM invested in the three share classes of the Vanguard FTSE Social Index Fund in Millions of dollars.

invests exclusively in the United States, and with a 2021 AUM of \$12 Billion, has the longest track record of any social index fund and is the largest social index fund in the world.¹³ The FTSE USA 4 Good Select Index consists of the subset of stocks in the FTSE USA Index that pass an ESG screening procedure. The FTSE USA index is a market-capitalization weighted index representing the performance of US large and mid cap stocks. The evolution of the combined total AUM invested in the three share classes of the Vanguard FTSE Social Index Fund is plotted in Fig. 1.

Because this index contains clean stocks in their (normalized) market weights, we assume that the dirty portfolio consists of only dirty stocks in their (normalized) market weights. Under these assumptions (further details are provided in Section 5.1), in the last 5 years of our sample (December 2015–December 2020), the dirty portfolio comprised a little more than a quarter of US market capitalization: $V_D = 27\%$. We assume here that the fraction of clean and dirty stocks is the same for the small and mid cap stocks that are not in the FTSE USA Index.¹⁴ The measured correlation between the clean (FTSE USA 4 Good Select) and the dirty (the other stocks in FTSE USA) portfolio over this 5-year period is $\rho = 0.93$. We use the historical market risk premium of 6%. The only remaining quantity to identify is the fraction of wealth controlled by ESG investors. The fraction of mutual fund wealth invested in ESG mutual funds in 2021 is less than 1%.¹⁵ To account for the possibility of Morningstar misclassification error, we double this estimate to 2%, and assume this fraction is representative of all capital investment in the economy. Using these parameters, the effect on the cost of capital is

$$\Delta \bar{R} - \Delta \bar{R}^* \approx 6\% \times 27\% \left(\frac{0.02}{1 - 0.93^2} \right) (1 - 0.93^2) = 0.44 \text{ b.p.}$$

A difference of half a basis point cannot meaningfully affect the capital budgeting decision and so an effect of this size cannot affect real investment decision making.

The FTSE USA 4 Good Select Index appears to be specifically designed by FTSE for Vanguard. It is related to (but less stringent than) the FTSE USA 4 Good Index that is marketed to most investors. If we instead use the FTSE USA 4 Good Index, over the same time period, the dirty portfolio comprised a little under half of US market capitalization:

¹³ VFTSX has been in existence since 2000. In 2016, Blackrock launched a competitor ETF, the iShares ESG Aware MSCI USA ETF which eclipsed the Vanguard fund in AUM in 2022 for a brief period but has since lost almost half of its AUM.

¹⁴ As the FTSE USA Index represents over 80% of the total U.S. stock market capitalization, these stocks represent less than 20% of the US stock market.

¹⁵ We identified the universe of ESG mutual funds by using Morningstar classifications and the names of the funds. We then hand checked the prospectuses of the top 20 ESG funds by AUM.

$V_D = 48.5\%$. The measured correlation between the clean (FTSE USA 4 Good) and the dirty (the other stocks in FTSE USA) portfolio over the same time period is $\rho = 0.97$. Using these parameters provides

$$\Delta \bar{R} - \Delta \bar{R}^* \approx 6\% \times 48.5\% \left(\frac{0.02}{1 - 0.97^2} \right) (1 - 0.97^2) = 0.35 \text{ b.p.}$$

Even though the FTSE USA 4 Good Index classifies more firms as dirty which increases the effect on the cost of capital, because the dirty portfolio now comprises a greater fraction of the economy, it is more correlated with the clean portfolio which more than offsets the effect resulting in the same conclusion — the difference cannot meaningfully impact the capital budgeting decision.

We have also explored the holdings of all socially conscious mutual funds to gauge the fraction of dirty stocks in the economy. Such ESG funds may choose not to hold certain stocks because they do not find them to be attractive investment opportunities, regardless of the stocks' ESG status. For this reason, the intersection of holdings across all ESG funds will underestimate the fraction of clean stocks and overestimate the fraction of dirty stocks. If we focus on the union of all stocks held by ESG funds (which will include stocks that not all ESG investors will agree on being clean), we estimate V_D to be 18% and ρ to be 0.8. These two changes once again offset each other resulting in an estimated change in the cost of capital of

$$\Delta \bar{R} - \Delta \bar{R}^* \approx 6\% \times 18\% \left(\frac{0.02}{1 - 0.8^2} \right) (1 - 0.8^2) = 0.79 \text{ b.p.}$$

The possibility exists that the mutual fund sector might not be representative of clean investing as a whole because clean investors might be more concentrated in other sectors. We therefore looked for other sources to calibrate the fraction of clean investors. One source is an estimate reported in SIF (2020), a report published by The Forum for Sustainable and Responsible Investment, an organization that represents sustainable investors. The report surveys investors asking them whether they consider any ESG criteria in their investment decisions. Based on this survey, the report estimates the total wealth of U.S.-domiciled assets using sustainable investing strategies to be \$17.1 trillion, or about 33% of U.S. market wealth. This number is almost certainly an overestimate. For example, the report concludes that of these assets, \$16.5 Trillion is controlled by money managers, which would imply (based on the size of the money management industry) that over 90% of all managed money uses ESG criteria. It appears that in answering the survey, when a representative of an organization states that it takes impact into account, the authors of the survey then assume that *all* the capital that that organization manages is subject to ESG criteria, regardless of whether a particular fund in the organization actually uses ESG criteria to manage money. In addition, the survey also includes strategies other than divestiture. Using this estimate (with the other original parameter estimates) gives

$$\Delta \bar{R} - \Delta \bar{R}^* = 6\% \times 27\% \left(\frac{0.33}{1 - 0.93} \right) (1 - 0.93^2) = 10.6 \text{ b.p.}$$

Using what is undoubtedly an overestimate of the fraction of wealth managed under ESG criteria, the effect on the cost of capital is still too small to meaningfully impact the firm's real investment decision making.

Some readers have argued that instead of the FTSE indices, we should have focused our analysis on MSCI indices. The stated motivation was that at the time of the writing of our paper, the iShares ESG Aware MSCI USA ETF which replicates the MSCI ESG Aware Index had shown large growth in AUM. Indeed the AUM of the iShares ESG Aware MSCI USA fund eclipsed the AUM of the Vanguard FTSE Social Index Fund for a brief period in 2022. Since that time, however, the iShares ETF has lost a large amount of its AUM unlike the Vanguard fund. Therefore, the Vanguard FTSE Social Index quickly regained its position as the largest ESG index fund. Regardless, it is interesting to explore what would have happened if we had used MSCI indices to calibrate our model. As it turns out, the ESG Aware MSCI USA as well as its parent index the MSCI USA indices are very similar to the two corresponding FTSE index funds. The correlation between the iShares ESG Aware Index and the FTSE USA 4 Good Select Index is very high, with a correlation of 0.997. Also the number of stocks in both indices is similar with an implied V_D close to 50%. Because the estimated correlation between clean and dirty stocks as defined by MSCI is 0.99 our calibration exercise yields, once again, a change in the cost of capital of less than a basis point.

4. What would it take?

In this section we consider the question of what it would take for a divestment strategy to meaningfully impact real investment policy. We begin by first considering how many ESG investors it would take to impact the cost of capital by 1%. To answer this question we calibrate the model as follows. We use the return volatility of the clean portfolio of $\sigma_{R_E} = 15.8\%$, based on the monthly return volatility of the FTSE 4 Good Select Index over the past 5 years multiplied by $\sqrt{12}$. We identify the set of dirty stocks by taking the stocks missing in the FTSE 4 Good Select that are nevertheless in the FTSE USA stock index, and then directly estimate the return volatility of this portfolio over the same sample period providing $\sigma_{R_D} = 15.2\%$. The correlation between these two portfolios over the same 5-year period is $\rho = 0.93$. We use a risk free rate of 2%, MRP of 6% and $V_D = 27\%$, as before. We then use these moments to infer the values of the exogenous parameters. That is, we infer the cash flow standard deviations by multiplying the return standard deviations by the value of the clean and dirty stocks respectively. We then infer the expected liquidating dividends by using the above values of MRP, V_D and V_E in the equilibrium in the economy without ESG investors. Taking these exogenous parameters as given, Fig. 2 plots the equilibrium $\Delta \bar{R} - \Delta \bar{R}^*$, the effect on the cost of capital of ESG investors, as a function of the fraction of wealth ESG investors comprise.

The red curve in Fig. 2 plots the exact relation derived in Appendix C of $\Delta \bar{R} - \Delta \bar{R}^*$ as a function of γ , the fraction of wealth controlled by clean investors, that is, (C.1). The dashed curve in Fig. 2 is (10), the approximation to (C.1). Even when ESG investors make up 50% of wealth, the impact on the cost of capital is less than 20 b.p. To impact the cost of capital by at least 1% requires that at least 84% of investors choose to hold only clean stocks.

We can also infer what the impact would be if a large investor decided to divest all dirty stocks. For example, the largest investor in the world, Blackrock, managed about \$8 trillion at the end of our sample period, which, if you assume is all invested in domestic stocks (it is not currently, the portfolio contains a large bond allocation and a large international allocation) is about 17% of the market. So if an investor the size of Blackrock were to shift all their capital into clean U.S. stocks (and none of Blackrock's investors reacted by withdrawing

funds, a very unlikely scenario), the fraction of clean shareholders would rise from 2% to at most 19%. At 19%, the impact on the cost of capital is just 4.6 b.p.

As the difference between the red and black dashed curves in Fig. 2 makes clear, the approximation is very accurate at current values of the parameters, implying that the four variables in the approximation – the fraction of ESG investors, the fraction of dirty stocks, the risk premium demanded by investors and the correlation between clean and dirty stocks – are the primary determinants of the impact of divestiture on the cost of capital. This explains why, for a divestiture strategy to work, ESG investors need to be such a large fraction of investors. To first order, the last three variables are not under the control of ESG investors.¹⁶ Using current estimates we have

$$6\% \times 27\% (1 - 0.93^2) = 0.0021.$$

This calculation constrains the effectiveness of the divestiture strategy. For (10) to reach 1%, we need $\frac{\gamma}{1-\gamma} > 4.6$, implying that $\gamma > 82\%$. There are two reasons for this. The correlation coefficient of 0.93 implies that clean and dirty stocks are close substitutes. This and the fact that most stocks are clean implies that to induce non-ESG investors to hold dirty stocks does not require much price adjustment. The only way to get a modest impact is to effectively induce non-ESG investors to hold only dirty stocks, which is what happens when γ approaches one.

Another point where readers have criticized our model is that it does not include investors who derive utility from not holding dirty stocks, so called “warm glow preferences”. Although we assume in the model that all investors have mean-variance preferences, because ESG investors restrict their investment universe to only clean stocks, one can view these investors as having extreme warm glow preferences. In a model in which they have a preference for holding clean stocks and thus face a tradeoff between these preferences and economic returns, the constraint that they cannot hold dirty stocks in any amount would be relaxed, thereby reducing their impact on the cost of capital. So one can view our results as an upper bound on the effectiveness of a divestiture strategy when some investors have warm glow preferences.

Another source of criticism of our results is that we do not account for the role of passive investing. In a recent paper (Cheng et al., 2023) examine the role of “passive” investing and conclude that the effects studied in our article are magnified as the fraction of passive investors increases. This inference is critically dependent on the definition of passive. In our setting all investors are passive in the sense that they hold the mean-variance efficient portfolio. To have the effect (Cheng et al., 2023) identify, investors need to be passive in a much stronger sense. Instead of choosing to follow a passive strategy (as our investors do), they need to be *constrained* to follow a passive strategy. Because these investors never trade, it is straightforward to compute how such investors would affect the inferences from our model. One can simply recompute the equilibrium inferences assuming these investors were not part of the economy. So if constrained passive investors make up θ fraction of the economy, then the effective fraction of ESG investors, γ' is given by

$$\gamma' \equiv \gamma / (1 - \theta).$$

In their calibration, Cheng et al. (2023) assume that 70% of investors are constrained passive, and that ESG investors make up 20% of all investors. They conclude, under those assumptions, that the effect on the cost of capital is 27 b.p. Although the model in Cheng et al. (2023) is very different to the one in this article, the two models are remarkably consistent. Using Fig. 2, to get a 27 b.p. effect on the cost of capital would require $\gamma' = 0.6$, implying that with $\theta = 70\%$, ESG

¹⁶ Of course, as the fraction of ESG investors changes, in general equilibrium, prices change which necessarily has a second order effect on these variables, as is modeled in Fig. 2.




Fig. 2. Effect on the cost of capital of introducing ESG investors into the economy: The curves plot the change in the cost of capital, $\Delta \bar{R} - \Delta \bar{R}^*$, as a function of the fraction of wealth ESG investors comprise. The red curve is the exact effect on the cost of capital, that is, (C.1). The dashed black curve is the approximation, that is, (10).

investors would make up $0.6 \times 0.3 \approx 18\%$ of all investors, which is line with what Cheng et al. (2023) assumed. Although their effect is an order of magnitude larger than what we find, the fact that it is still small is an indication on how difficult it is for this strategy to have impact. After all, assuming that 70% of investors are constrained passive is tantamount to assuming that all passive investors are constrained. That is very unlikely to be true. Conversely, given current ESG participation ($\gamma = 2\%$), one could also compute what fraction θ of investors would have to be constrained passive to achieve a 1% change in the cost of capital. Using Fig. 2, to have a 1% effect on the cost of capital, $\gamma' = 82\%$, implying that $\theta = 1 - \frac{0.02}{0.82} = 97.6\%$ of investors would need to be constrained to hold passive portfolios.

Finally, readers have rightly pointed out that we do not provide any argument for what change in the cost of capital would be significant enough to change real investment policy. Clearly, a change less than 10 b.p. is inconsequential given the uncertainty in the capital budgeting process. We implicitly assumed that one would need at least a change of 1%, but we fully acknowledge that that cut off is arbitrary. The cut off also depends on the duration of the project. Small changes in the cost of capital can have a large effect on NPV calculations for long duration projects. Ultimately, this is an empirical question. However it makes little sense investing time and effort in attempting to answer this question without first demonstrating that the effect of the divestment strategy on the cost of capital is more than just a few basis points.

The model we use in this article is the standard model in finance and has been used in the theoretical literature studying divestment almost exclusively. Nevertheless, readers of this article have criticized our conclusion by arguing that the model in this paper cannot explain the index inclusion effect. Because of the close relation between that effect and the divestment strategy we study, these readers question the applicability of the model to answer the question at hand. However, this argument is flawed. When properly calibrated, the model in this paper is entirely consistent with the documented index inclusion effect.

Greenwood and Sammon (2022) show that, in the most recent decade, the price inclusion effect into the S&P500 index is 0.3%. Assuming an average duration of 40 years (based on a dividend yield of 2.5%), a price change of 30 b.p. implies a change in the expected return of 0.75 b.p. The model-predicted difference in the expected return between S&P500 (Indexed) and non-S&P500 (Non-Indexed) stocks in equilibria with and without index investors can be approximated by:

$$\Delta \bar{R} - \Delta \bar{R}^* \approx \text{MRP} \times V_{NI} \times \left(\frac{\gamma}{1-\gamma} \right) (1-\rho^2) \frac{\sigma_{R_{NI}}^2}{\sigma^2}. \quad (11)$$

As before, MRP is the market risk premium, γ is the fraction of investor wealth in S&P 500 index funds (or potentially funds that closely track

the S&P 500), V_{NI} is the market capitalization of stocks not in the S&P 500 and ρ is the correlation between S&P 500 stocks and non-S&P 500 stocks. The final term, the ratio of the return variance of stocks not in the S&P 500 and the market variance, is required because, unlike in the case of clean and dirty stocks, small and large stocks have substantially different return variances.

To calibrate our model we use the data made available by Ken French on size sorted portfolios. In the past decade, the top three size deciles comprise a little over 500 stocks (521 on average with a minimum of 505 and a maximum of 554). They therefore form a very close approximation of the S&P 500 Index. Indeed, the return correlation of these three top deciles with the S&P500 is 0.997 over this sample period. The three top deciles represent an average of 86.6% in market capitalization, implying $V_{NI} = 13.4\%$. Further, the standard deviation of the stocks in the bottom seven deciles is $\sigma_{R_{NI}} = 23.7\%$, the market standard deviation is $\sigma = 18.4\%$ and the correlation between the return of the top three deciles and the bottom seven deciles is $\rho = 0.89$. Fig. 3 plots the effect of the cost of capital as a function of γ (the fraction of investor wealth in the S&P500) at these calibrated numbers. To get an effect on the expected return of 0.75 b.p. per annum the model implies that 2.7% of investors restrict their investment universe to only S&P500 stocks. Cutting the duration estimate in half to twenty years, yields a change in the expected return of 1.5 b.p. which requires just over 5% of investors to restrict their investment universe to only S&P500 stocks. The implication that between 2% and 6% of investors restrict their portfolios to S&P500 stocks is eminently plausible.

In summary, our model's predictions are consistent with the observed S&P 500 inclusion effects. Historically the observed index inclusion effects have been of larger magnitude. Greenwood and Sammon (2022) attributed the decline in the inclusion effect to lower bid ask spreads and increased market efficiency. This is consistent with the original finding in Harris and Gurel (1986) who document that nearly all of the detected effect mean reverted in two weeks indicating that a large part of the original effect can be attributed to liquidity effects which are absent in our model. Given that the topic of our paper is on how to best induce firms to change their real investment behavior, short-term price pressure effects that mean revert in two weeks are not our focus. That said, to generate the price effect of 3% that was originally documented (implying a cost of capital effect of about 7 b.p.), Fig. 3 shows that even in this case one would only need about 20% of investors to restrict their investment universe to large stocks.

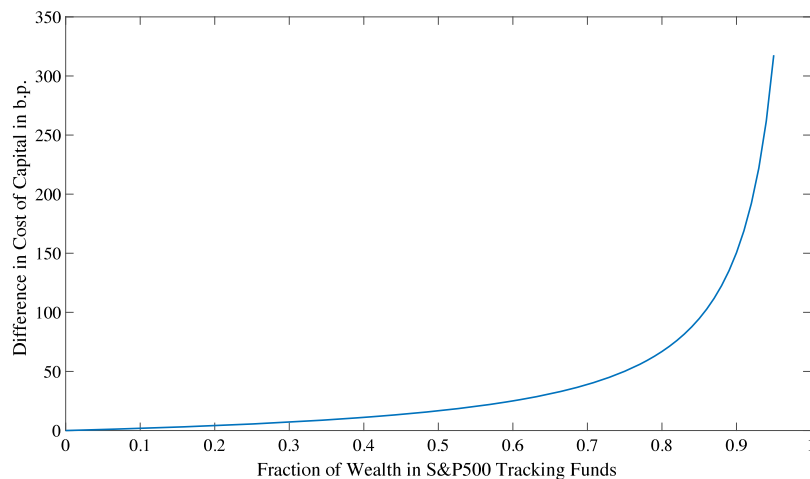


Fig. 3. S&P500 index inclusion effect: The curve plots the change in the expected return, $\Delta \bar{R} - \Delta \bar{R}^*$, as a function of γ the fraction of wealth S&P 500 investors comprise, that is, (11).

5. Empirical evidence

The theory we developed in the previous sections suggests that the observed effect of ESG investors on the cost of capital should be small. In this section we evaluate this prediction empirically. To assess the effect of ESG investors on the cost of capital we focus on changes in ESG classifications, that is, the effect on the cost of capital when a firm is classified as clean or classified as dirty. According to our theory, when ESG investors react to the change in status we should observe a change in the cost of capital. Our objective in this section is to measure the magnitude of this change.

One immediate concern is that because average realized returns are a noisy measure of the expected return (cost of capital), observing a change in this expectation is difficult. Consequently we will also measure the concurrent change in the firm's price. Because prices are directly observable and, under the rational expectations assumption, should reflect the capitalized value of the change in the cost of capital, in principle the price change should be easier to detect. Note that if the actions of investors actually do materially affect the cost of capital, clean (dirty) firms would have better (worse) growth opportunities, which would lead to an even larger price impact.

We use the same stock indices published by FTSE Russell to identify changes in the status of companies. In the words of FTSE Russell:

"The FTSE 4 Good indices are designed to measure the performance of companies that have demonstrated strong Environmental, Social and Governance (ESG) practices. Transparent management and clearly-defined ESG criteria make FTSE 4 Good Indexes suitable tools to be used by a wide variety of market participants when creating or assessing sustainable investment products".

There are a number of reasons why we choose to focus on these indices to identify changes in the ESG status of stocks. Because the Vanguard-FTSE Social Index fund chooses to replicate the FTSE USA 4 Good Select Index, when FTSE Russell changes the constituents of the FTSE USA 4 Good Select Index, there is an immediate redeployment of capital associated with the rebalancing activity of the largest index fund in the space. But more importantly, Vanguard's choice to replicate the FTSE USA 4 Good Select Index is a business decision that presumably reflects their belief that this index most effectively captures the space of large clean US companies. Thus a change in the FTSE USA 4 Good Select Index likely also reflects the investment decisions of other ESG investors in the economy. Finally, the fact that the FTSE USA 4 Good Select Index is a subset of the FTSE USA Index implies that any stock not in the 4 Good Index must not have satisfied ESG criteria. If, instead, we used the holdings of other impact funds, we could not differentiate between stocks that are not included because they do not satisfy ESG criteria, or because they do not represent good investment opportunities.

5.1. Data and descriptive statistics: FTSE indices

To assess the impact of ESG classification changes, we use both the FTSE USA 4 Good Select and the FTSE USA 4 Good Index. Both indices are subsets of the FTSE USA index. The FTSE USA 4 Good Select Index has as its primary advantage the fact that the Vanguard index fund tracks it, so we can be assured that a significant amount of capital does react to an index change. However, FTSE Russell markets the FTSE USA 4 Good Index more intensively, so it is likely that there are investors that react to changes in this index as well. We therefore report our results for both indices.

We obtained data from FTSE on the constituents and their portfolio weights for all three indices starting in the early 2000s until the end of 2020.¹⁷ The number of constituents over this time sample for all three indices is plotted in Fig. 4. For the FTSE USA Index the number varies between 446 and 745. Both extremes occur early in the sample after which the number of constituents stabilizes. The average over the sample is 616. The number of constituents for the FTSE USA 4 Good Index varies between 130 and 274 with an average of 189 and the number of constituents for the FTSE USA 4 Good Select Index varies between 286 and 506 with an average of 403.

In Fig. 5 we plot the total market capitalization of all indices. The total market capitalization of the FTSE USA Index at the end of our sample (December 2020) is \$33.1 Trillion representing 83% of the total stock market capitalization of the United States (\$40 Trillion). At that time, the market capitalization of the FTSE USA 4 Good (Select) Index was \$18.3 Trillion (\$28.2 Trillion) which equals 55% (85%) of the FTSE USA stock market capitalization. This implies that the dirty stocks in the FTSE USA Index make up 45% (15%) of the index at that point in time.

Because we are interested in documenting the effects on prices and returns of index inclusions, we next document how often firms are added and excluded from each index. To facilitate this exercise, we define dummy variables that capture inclusions and exclusions in each index. First define the dummy that describes whether firm i is in the FTSE USA index at time t :

$$I_{i,t} = \begin{cases} 1 & \text{if firm } i \text{ is in the FTSE USA Index at time } t \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

¹⁷ We thank FTSE for generously sharing their data with us.




Fig. 4. FTSE index constituents: Number of constituents of the FTSE USA, the FTSE USA 4 Good and the FTSE USA 4 Good Select Indices.




Fig. 5. FTSE index market capitalization: Total market capitalization of the FTSE USA, the FTSE USA 4 Good and the FTSE USA 4 Good Select Indices in \$ Trillions.

Table 1

Inclusions and exclusions: Incidences of dummy variable realizations for the FTSE USA Index, the FTSE USA 4 Good Index and the FTSE USA 4 Good Select Index.

	Number of events	
	4 Good (2001–2020)	4 Good Select (2005–2020)
$\Delta I_{it} = 1$	872	508
$\Delta I_{it} = -1$	795	653
$\Delta I_{it}^{4G} = 1$	411	812
$\Delta I_{it}^{4G} = -1$	385	824
$\Delta I_{it} = 1 \text{ \& } \Delta I_{it}^{4G} = 1$	54	118
$\Delta I_{it} = -1 \text{ \& } \Delta I_{it}^{4G} = -1$	200	386

The corresponding dummy variable for the FTSE USA 4 Good (Select) index is defined as:

$$I_{i,t}^{4G} = \begin{cases} 1 & \text{if firm } i \text{ is in the FTSE USA 4 Good (Select) Index at time } t \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

The inclusion and exclusion events can then be described as:

$$\Delta I_{it} \equiv I_{i,t} - I_{i,t-1} \quad (14)$$

$$\Delta I_{it}^{4G} \equiv I_{i,t}^{4G} - I_{i,t-1}^{4G}. \quad (15)$$

The latter two variables take the value 1 for an inclusion, -1 for an exclusion, and 0 otherwise. In Table 1 we summarize the inclusion and exclusion events in our sample for all three indices.

The table shows that there are many inclusion and exclusion events for all three indices. The second column in the table compares inclusion and exclusion events for the FTSE USA and the FTSE USA 4 Good Index for the sample period that both indices are available: 2001–2020. The total number of firms that have been in the FTSE USA and the FTSE USA 4 Good Index at some point during this sample period equals 1474 and 589 respectively. Over this sample period, the FTSE USA had 872 inclusion events and 795 exclusion events, whereas the FTSE 4 Good Index had 411 inclusion and 385 exclusion events. Only a small fraction of inclusion events coincide: only 54 of the 411 FTSE USA 4 Good inclusion events coincide with the companies' inclusion in the FTSE USA Index. This coinciding fraction is higher for index exclusion events, $\frac{200}{385} = 52\%$. This is to be expected, as exclusion from the FTSE USA Index implies exclusion from the FTSE USA 4 Good Index as membership of the FTSE USA Index is a necessary condition for inclusion in the FTSE USA 4 Good Index.

The third column in the table repeats the analysis presented in the second column, but now comparing the FTSE US Index with the FTSE 4 Good Select Index over the sample period that both indices are available: 2005–2020. Over this sample period, the FTSE USA had 508 inclusion events and 653 exclusion events, whereas the FTSE 4 Good Select Index had 812 inclusion and 824 exclusion events. The number of inclusion events that coincided between the FTSE US and the FTSE US 4 Good Select Index equals 118, and the number of coinciding exclusion events equals 386. Note that the FTSE USA 4 Good Select had substantially higher turnover than the FTSE USA 4 Good Index.




Fig. 6. Inclusion months: Number of total months included in the FTSE USA and the FTSE USA 4 Good Indices.

To further explore the relation between inclusion in the two indices, we plot in Fig. 6 for all the stocks that were ever included in the FTSE USA Index over the June 2001–December 2020 sample period, the total number of inclusion months for the FTSE USA and the FTSE USA 4 Good indices. The scatter plot exhibits a triangular shape, which is to be expected given that (as mentioned above), membership of the FTSE USA Index is a necessary condition for inclusion in the FTSE USA 4 Good Index. Observations on the vertical line on the right side of the graph represent companies that were included in the FTSE USA Index for the full sample period. The horizontal line on the bottom represents companies that were never included in the FTSE 4 Good Index, and were thus deemed “dirty” firms for the whole sample period. Companies on the diagonal represent firms that were included in both indices for the same number of periods, presumably because they were added and deleted at the same time. The graph shows that there are many companies below the diagonal, indicating that for many firms the inclusion and exclusion decision into the FTSE 4 Good Index show no clear coincidence with the inclusion into the FTSE USA Index. They therefore represent a change in status — an inclusion event in the FTSE 4 Good Index indicates a firm transitioning from dirty to clean and an exclusion event indicates a firm transitioning from clean to dirty.

When we repeat the analysis for the FTSE 4 Good Select Index (over the available data sample period of November 2005 through December 2020), a similar pattern emerges, though there are a small number of deviations in the scatter plot above the 45-degree line, resulting from the fact that in rare cases firms stay a few months longer in the FTSE 4 Good Select Index than in the FTSE US Index (see Fig. 7).¹⁸

To gauge the effect of inclusion and exclusion events on prices and returns, we need to construct a merged data set of FTSE index data with return data from the Center for Research in Security Prices (CRSP). While such a merge appears straightforward, it turns out to be quite involved. The reason is that different security identifiers are used in different contexts. The main security level identifiers used by FTSE Russell are the ISIN and the SEDOL numbers, whereas the CRSP database provides CUSIPs and (firm-level) PERMCO numbers. To accomplish the merge, we use the S&P Global Market Intelligence database provided on WRDS which has an identifiers database linking company ISINs and CUSIPs. Because CUSIPs change over time and companies can issue multiple securities, we use PERMCO to identify the single security associated with each holding.

5.2. Results gauging price and return effects

Our theory predicts that when a stock experiences either an inclusion or an exclusion event the effect on the cost of capital will be small. In this subsection we test this implication of the theory.

Our main regression specification is the following:

$$R_{it} = c + \gamma I_{it} + \delta \Delta I_{it} + \gamma_{4G} I_{it}^{4G} + \delta_{4G} \Delta I_{it}^{4G} + \varepsilon_{it} \quad (16)$$

where R_{it} is the monthly stock return including dividends on all stocks in the CRSP database. The coefficient, γ_{4G} , on the dummy variable I_{it}^{4G} measures the average return difference between clean and dirty FTSE USA stocks and is therefore an estimate of the effect of ESG investors on the cost of capital of the average stock in the FTSE USA Index. The coefficient, δ_{4G} , on the dummy variable ΔI_{it}^{4G} measures the instantaneous price reaction of an inclusion or exclusion event (recall that this dummy is 1 in the month of inclusion and -1 in the month of exclusion) and thus measures the capitalized value of the implied change in the cost of capital (the associated percentage price change). Notice that although the dummy variables I_{it} and ΔI_{it} are primarily included to control for inclusion in the FTSE USA Index, they also measure the effect of other inclusion effects unrelated to social investing such as liquidity.

If inclusion in the 4 Good (Select) Index has a measurable effect on the cost of capital, we would expect to find a positive coefficient on ΔI_{it}^{4G} suggesting an instantaneous price appreciation (depreciation) upon inclusion (exclusion). Similarly, we would expect a negative coefficient on I_{it}^{4G} implying lower (higher) average returns following the instantaneous price appreciation (depreciation).

Table 2 reports the results of the test of (16). In line with our theory, we find small estimates of δ_{4G} . The point estimate of the instantaneous price change of the 4 Good Index is 0.22% (first data column in Table 2). Assuming an average stock duration of 40 years, the implied change in the cost of capital is 0.55 b.p.¹⁹ This duration estimate is based on aggregate stock market data. Of course, there is large cross sectional variation in stock specific duration estimates, but it would be hard to argue for a stock duration less than 10 years, implying an upper bound on the implied change in the cost of capital of 2.2 b.p. The estimated instantaneous price change using the 4 Good Select Index is 0.48% (second data column in Table 2), implying a cost of capital impact for the average firm of 1.2 b.p., with an upper bound of 6 b.p. Although the point estimates are not statistically different from zero, they are in line with the order of magnitude predicted by the theory.

Given the extremely small estimated effect on the cost of capital, measuring the change directly using average realized returns is more challenging. In line with this, the estimate of γ_{4G} in both specifications in Table 2 is not significantly different from zero and has a sign that is both inconsistent with the theory and the prior estimate using the price change. Given the uncertainty in estimating the effect and the closeness of the coefficients to zero, there is essentially a near 50% probability of estimating the wrong sign.

¹⁸ It is unclear why FTSE Russell departed from their policies in these rare cases, but we conjecture that it is related to Vanguard’s optimal tracking strategies.

¹⁹ van Binsbergen (2024) estimates stock durations implied by dividend yields.

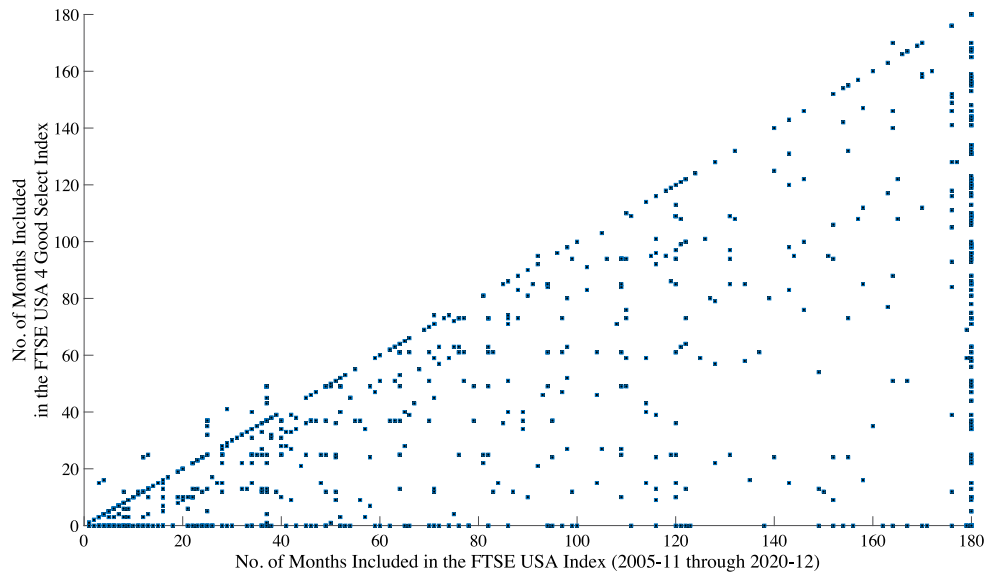


Fig. 7. Inclusion months: Number of total months included in the FTSE USA and the FTSE USA 4 Good Select Indices.

Table 2

Return and price effects of index inclusions: FTSE USA vs. FTSE USA 4 Good (Select). The table reports estimation results of the regression specified in (16). The dummies I_{it} and I_{it}^{4G} equal 1 for all months that a stock is in the index. The variables ΔI_{it} and ΔI_{it}^{4G} equal 1 in the month of inclusion, -1 in the month of exclusion, and 0 otherwise. Standard errors are double clustered by firm and yearmonth.

	(1) 4Good	(2) 4Good Select
I_{it}	−0.0010 (−0.65)	−0.0009 (−0.50)
$I_{it}^{4G(Select)}$	0.0004 (0.28)	0.0001 (0.02)
ΔI_{it}	0.0131 (1.56)	0.0117 (1.34)
$\Delta I_{it}^{4G(Select)}$	0.0022 (0.39)	0.0048 (0.43)
Constant	0.0098*** (2.60)	0.0101*** (2.65)
Observations	1 376 791	1 365 665
R^2	0.00	0.00

t-statistics in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

One concern is that because investors might be slow to update their portfolios in response to changes in ESG classifications, the above estimates will underestimate the magnitude of the effects. To address this concern we explore, in Table 3, the effect of including additional lags of ΔI_{it}^{4G} . The main insight remains the same: there seems little to no evidence that inclusion in the FTSE USA 4 Good (Select) Index has any meaningful price or return effects. The cumulative effect of all the lags of ΔI_{it}^{4G} , that is, the cumulative price appreciation within 4 months after inclusion into the FTSE US 4 Good (Select) Index is smaller than the original estimate, suggesting that the magnitude of the measured price effect of inclusion and exclusion events is not dampened by investor sluggishness. The overall conclusion is that the effect on the cost of capital of a change in ESG status of a firm is so small that it cannot meaningfully influence the firm's investment decisions.

In line with prior work on the effect of index inclusions in general, we do find modest effects associated with inclusions and exclusions from the FTSE USA Index itself (third row in Table 2). The events are associated with an instantaneous price change of 1.31% (4Good specification) and 1.17% (4Good Select specification) and a lower average return of 10 b.p. (4Good specification) and 9 b.p. (4Good Select specification) consistent with what Greenwood and Sammon (2022)

find for the S&P 500 Index. Although the signs are all consistent, none of these estimates are statistically significantly different from zero. Specification (2) and (5) in Table 3 measure the price effects in the months following the events. Although the estimates are not statistically significantly different from zero, what appears to be evident is that the price effect reverses in the months following the inclusion/exclusion event which is consistent with the original evidence on the S&P 500 inclusion effect documented in Harris and Gurel (1986). The subsequent reversal suggests that the price change is associated with the short-term price pressure from participants seeking to match the index, rather than a permanent change in the cost of capital.

5.3. Changes in risk

One advantage of our empirical approach is that we do not need to take a stand on the underlying risk model in the economy. Instead we make the much weaker assumption that there is no correlation between the inclusion event and changes to the riskiness of the firm. In this section we test this assumption by exploring whether inclusion in the 4Good (or the 4Good Select) Index leads to changes in firm risk as measured by the CAPM. To evaluate whether being included in the index has an effect on the systematic risk of the firm we use daily stock return data to compute for each stock and each month the realized (Scholes–Williams) beta in that month. That is, we regress daily stock returns in excess of the risk free rate on daily market returns and lagged market returns in excess of the risk free rate (mktrf) for each month, sum the coefficients and use these estimated betas, denoted by $\hat{\beta}_{it}$ as left-hand-side variables in our regression framework:

$$\hat{\beta}_{it} = c + \gamma I_{it} + \gamma_{4G} I_{it}^{4G} + \epsilon_{it}. \quad (17)$$

To better evaluate whether the inclusion event itself causes the riskiness of the firm to change rather than high or low beta firms having a higher or lower probability of being selected for one of the indices, we also include a set of stock fixed effects. The results are summarized in Table 4.

The table shows that being included in the FTSE4 Good US Select Index has no detectable influence on the CAPM beta of firms. The coefficient is -0.0169 with a t-stat of -0.96 . For the FTSE 4 Good index we find a small, though statistically significant, effect: the coefficient is -0.06 with a t-statistic of -2.85 , which for an equity risk premium of 6% corresponds to a change in the annual cost of capital of roughly

Table 3

Return and price effects of index inclusions: FTSE USA vs. FTSE USA 4Good (Select). The table reports estimation results of the regression specified in (16). The dummies I_{it} and I_{it}^{4G} equal 1 for all months that a stock is in the index. The variables ΔI_{it} and ΔI_{it}^{4G} equal 1 in the month of inclusion, -1 in the month of exclusion, and 0 otherwise. Standard errors are double clustered by firm and yearmonth.

	(1) 4Good	(2) 4Good	(3) 4Good	(4) 4Good Select	(5) 4Good Select	(6) 4Good Select
I_{it}	−0.0010 (−0.65)	−0.0013 (−0.89)	−0.0013 (−0.84)	−0.0009 (−0.50)	−0.0007 (−0.34)	−0.0006 (−0.31)
$I_{it}^{4G(Select)}$	0.0003 (0.21)	0.0011 (0.68)	0.0010 (0.64)	−0.0000 (−0.01)	−0.0013 (−0.36)	−0.0012 (−0.36)
ΔI_{it}		0.0119 (1.41)			0.0092 (1.03)	
$\Delta I_{i,t-1}$	0.0012 (0.10)	0.0007 (0.05)	0.0006 (0.05)	0.0053 (0.46)	0.0025 (0.18)	0.0024 (0.18)
$\Delta I_{i,t-2}$		0.0106 (0.82)	0.0105 (0.82)		0.0057 (0.41)	0.0056 (0.40)
$\Delta I_{i,t-3}$		−0.0056 (−0.68)	−0.0057 (−0.69)		−0.0065 (−0.79)	−0.0062 (−0.76)
$\Delta I_{i,t-4}$		−0.0097 (−0.66)	−0.0099 (−0.67)		−0.0118 (−0.87)	−0.0119 (−0.87)
$\Delta I_{it}^{4G(Select)}$		0.0027 (0.47)			0.0063 (0.56)	
$\Delta I_{i,t-1}^{4G(Select)}$	0.0058 (0.89)	0.0035 (0.52)	0.0035 (0.52)	−0.0069 (−0.56)	−0.0056 (−0.44)	−0.0057 (−0.44)
$\Delta I_{i,t-2}^{4G(Select)}$		−0.0124 (−1.11)	−0.0124 (−1.11)		0.0103 (1.34)	0.0102 (1.33)
$\Delta I_{i,t-3}^{4G(Select)}$		0.0016 (0.23)	0.0014 (0.21)		−0.0030 (−0.48)	−0.0032 (−0.51)
$\Delta I_{i,t-4}^{4G(Select)}$		−0.0130* (−1.91)	−0.0130* (−1.91)		−0.0028 (−0.22)	−0.0029 (−0.22)
Constant	0.0101*** (2.64)	0.0107*** (2.85)	0.0107*** (2.85)	0.0104*** (2.73)	0.0107*** (2.78)	0.0107*** (2.78)
Observations	1 368 499	1 327 152	1 327 152	1 346 087	1 271 890	1 271 891
R^2	0.00	0.00	0.00	0.00	0.00	0.00

t-statistics in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4

CAPM Betas for included and excluded firms. The table reports regression results of firm-level estimated CAPM betas (left hand side) on the inclusion dummies. Stock fixed effects are included. Standard errors are double clustered by yearmonth and firm.

	(1) 4Good	(2) 4Good Select
I_{it}	−0.0501* (−1.86)	−0.0459 (−1.64)
$I_{it}^{4G(Select)}$	−0.0603*** (−2.85)	−0.0169 (−0.96)
Constant	0.9484*** (58.05)	0.9468*** (57.83)
Observations	1 386 980	1 375 132
Stock Fixed Effects	Y	Y
R^2	0.08	0.08

t-statistics in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

36bps. The sign of the coefficient further strengthens the arguments presented in this paper. If inclusion in the 4Good Index slightly lowers systematic firm risk, this should lead to an instantaneous risk-based increase in the price upon inclusion. The measured instantaneous price inclusion effect in the previous sections was already very small, and thus becomes even smaller after adjusting for this effect of risk.

6. Conclusion

In this paper we have evaluated, both theoretically and empirically, the quantitative impact of socially conscious investing in the secondary market. We conclude that at current participation levels the divestiture strategy is unlikely to have a large impact on the long-term cost of capital of targeted firms. An unrealistically large increase in the level

of participation would be required for the strategy to affect corporate policy.

Given the current levels of socially conscious capital, our conclusion is that using primary markets would be a more effective strategy to effect social change. Investors should either directly invest in firms by funding negative NPV projects (that achieve social good) or follow a policy of engagement. By purchasing the stock in targeted companies rather than selling the stock, socially conscious investors could potentially have greater impact by exercising their rights of control through the proxy process or by gaining a majority stake and replacing upper management.

CRedit authorship contribution statement

Jonathan B. Berk: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Jules H. van Binsbergen:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors have nothing to disclose.

Appendix A. Derivation of the equilibrium with ESG investors

The beta (with respect to the mean–variance efficient portfolio), β_E , of the clean portfolio is

$$\beta_E \equiv \frac{cov(R_E, R_{mv})}{var(R_{mv})} = \frac{1 - \gamma + b}{V_E} \left(\frac{\Sigma_E \sigma^2 - e \sigma_E^2}{\sigma^2 - 2e \Sigma_E \sigma^2 + e^2 \sigma_E^2} \right), \quad (A.1)$$

where we have used

$$\begin{aligned} \text{cov}(R_E, R_{mv}) &= \text{cov}\left(\frac{D_E}{V_E}, \frac{D_D + (1-e)D_E}{1-\gamma+b}\right) = \frac{(1-e)\sigma_E^2 + \rho\sigma_D\sigma_E}{V_E(1-\gamma+b)} \\ &= \frac{\Sigma_E\sigma^2 - e\sigma_E^2}{V_E(1-\gamma+b)} \end{aligned}$$

and

$$\begin{aligned} \text{var}(R_{mv}) &= \frac{\sigma_D^2 + (1-e)^2\sigma_E^2 + 2(1-e)\rho\sigma_D\sigma_E}{(1-\gamma+b)^2} \\ &= \frac{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2}{(1-\gamma+b)^2}. \end{aligned}$$

Similarly, the beta of the dirty portfolio, β_D is

$$\beta_D \equiv \frac{1-\gamma+b}{V_D} \left(\frac{\Sigma_D\sigma^2 - e\rho\sigma_D\sigma_E}{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2} \right) \quad (\text{A.2})$$

where

$$\text{cov}(R_D, R_{mv}) = \frac{\Sigma_D\sigma^2 - e\rho\sigma_D\sigma_E}{V_D(1-\gamma+b)}.$$

To solve for r and b we need to solve each group's maximization problem. ESG investors pick b to maximize

$$\frac{b}{\gamma}r + \left(1 - \frac{b}{\gamma}\right)E[R_E] - k\left(1 - \frac{b}{\gamma}\right)^2 \text{var}[R_E],$$

Taking the derivative and setting it equal to zero gives

$$r - E[R_E] + 2k\left(1 - \frac{b}{\gamma}\right)\text{var}[R_E] = 0$$

$$1+r = \frac{\bar{D}_E}{V_E} - 2k\left(1 - \frac{b}{\gamma}\right)\frac{\sigma_E^2}{V_E^2}. \quad (\text{A.3})$$

Other investors pick b to maximize

$$-\frac{b}{1-\gamma}r + \left(1 + \frac{b}{1-\gamma}\right)E[R_{mv}] - k\left(1 + \frac{b}{1-\gamma}\right)^2 \text{var}[R_{mv}].$$

Taking the derivative and setting it equal to zero gives

$$r - E[\omega_D R_D + \omega_E R_E] + 2k\left(1 + \frac{b}{1-\gamma}\right)\text{var}(\omega_D R_D + \omega_E R_E) = 0$$

$$1+r = \frac{\bar{D}_D + (1-e)\bar{D}_E}{1-\gamma+b} - 2k\frac{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2}{(1-\gamma)(1-\gamma+b)}. \quad (\text{A.4})$$

Eqs. (A.3) and (A.4) jointly determine r and b .

The return of the clean portfolio is given by the pricing equation

$$E[R_E] - r = \beta_E(E[\omega_D R_D + \omega_E R_E] - r). \quad (\text{A.5})$$

Substituting (A.1) and (3) into this expression gives:

$$\begin{aligned} \frac{\bar{D}_E}{V_E} - (1+r) &= \frac{1-\gamma+b}{V_E} \left(\frac{\Sigma_E\sigma^2 - e\sigma_E^2}{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2} \right) \\ &\times \left(\frac{\bar{D}_D + (1-e)\bar{D}_E}{1-\gamma+b} - (1+r) \right). \end{aligned} \quad (\text{A.6})$$

Substituting (A.4) into the right side of (A.6) and rearranging terms gives,

$$\begin{aligned} \bar{D}_E - (1+r)V_E &= 2k(1-\gamma+b) \left(\frac{\Sigma_E\sigma^2 - e\sigma_E^2}{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2} \right) \\ &\times \left(\frac{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2}{(1-\gamma)(1-\gamma+b)} \right). \end{aligned}$$

Substituting (A.3) into the left hand side of this expression and simplifying provides,

$$\gamma - b = V_E \gamma \left(1 + \frac{\rho\sigma_D}{\sigma_E} \right). \quad (\text{A.7})$$

Substituting this expression into (1) gives e in terms of primitives:

$$e = \gamma \left(1 + \frac{\rho\sigma_D}{\sigma_E} \right) \quad (\text{A.8})$$

Substituting (A.7) into (A.3) gives r in terms of primitives:

$$1+r = \frac{\bar{D}_E}{V_E} - 2k \left(1 + \frac{\rho\sigma_D}{\sigma_E} \right) \frac{\sigma_E^2}{V_E}. \quad (\text{A.9})$$

Finally, substituting (A.8) into (A.4) provides b in terms of primitives:

$$\gamma - b = 1 - \frac{\bar{D}_D + (1-e)\bar{D}_E}{1+r} + 2k \frac{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2}{(1-\gamma)(1+r)}. \quad (\text{A.10})$$

Using the expressions for the risk free rate and equilibrium holdings, (A.8), (A.9) and (A.10), we can solve (A.4) for V_E :

$$\begin{aligned} V_E &= \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_E + \bar{D}_D - 2k\sigma^2 \left(1 + \left(\frac{\gamma}{1-\gamma} \right) (1-\rho^2) \frac{\sigma_D^2}{\sigma^2} \right)} \\ &= \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_E + \bar{D}_D - 2k\sigma^2(1+\Gamma)}, \end{aligned} \quad (\text{A.11})$$

where Γ is defined by (6). Rearranging terms in (A.11) gives

$$\begin{aligned} \bar{1} + \bar{R}_E &= \frac{\bar{D}_E}{V_E} = \bar{D}_E + \bar{D}_D - 2k\sigma^2(1+\Gamma) + 2k\sigma^2 \frac{\Sigma_E}{V_E} \\ \bar{R}_E &= \bar{R} + 2k\sigma^2 (\beta_E^m - (1+\Gamma)) \end{aligned} \quad (\text{A.12})$$

where $\beta_E^m \equiv \frac{\Sigma_E}{V_E}$ is the *market* beta of the clean portfolio. Using the price normalization

$$\begin{aligned} V_D &= 1 - V_E \\ &= \frac{\bar{D}_D - 2k\sigma^2(\Sigma_D + \Gamma)}{\bar{D}_E + \bar{D}_D - 2k\sigma^2(1+\Gamma)} \end{aligned} \quad (\text{A.13})$$

so

$$\begin{aligned} \bar{R}_D &= \bar{R} - 2k\sigma^2 \left(1 + \Gamma - \frac{\Sigma_D + \Gamma}{V_D} \right) \\ &= \bar{R} + 2k\sigma^2 \left(\beta_D^m - \left(1 - \Gamma \frac{V_E}{V_D} \right) \right) \end{aligned} \quad (\text{A.14})$$

where $\beta_D^m \equiv \frac{\Sigma_D}{V_D}$ is the *market* beta of the dirty portfolio. Using (A.12) and (A.14), the difference in the expected return of dirty and clean stocks, $\Delta\bar{R}$, is

$$\begin{aligned} \Delta\bar{R} &\equiv R_D - R_E = 2k\sigma^2 \left(\beta_D^m - \beta_E^m + \frac{\Gamma}{V_D} \right) \\ &= 2k\sigma^2 (\bar{D}_E + \bar{D}_D - 2k\sigma^2(1+\Gamma)) \\ &\times \left(\frac{\Sigma_D + \Gamma}{\bar{D}_D - 2k\sigma^2(\Sigma_D + \Gamma)} - \frac{\Sigma_E}{\bar{D}_E - 2k\sigma^2\Sigma_E} \right) \end{aligned} \quad (\text{A.15})$$

Appendix B. Derivation of the equilibrium without ESG investors

To assess the effect of ESG investors, we now derive the equilibrium when all investors are identical. In this standard CAPM equilibrium, all investors choose to hold the market portfolio, which is mean-variance efficient. This implies that the expected return of the dirty portfolio is given by the CAPM pricing relation

$$E[R_D^*] - r^* = \beta_D(E[R] - r^*), \quad (\text{B.1})$$

where

$$\beta_D^* = \frac{\text{cov}(R_D^*, R)}{\text{var}(R)} = \frac{\text{cov}(R_D^*, R)}{\sigma^2},$$

and asterisks denote equilibrium variables in the economy with identical investors.²⁰ Now

²⁰ Because of the price normalization, the return of the market portfolio does not depend on the preferences of investors implying that market variables do not require asterisks.

$$\begin{aligned} \text{cov}(R_D^*, R) &= \text{cov}(R_D^*, V_D^* R_D^* + V_E^* R_E^*) \\ &= \frac{\text{cov}(D_D, D_D + D_E)}{V_D^*} \\ &= \frac{\sigma_D^2 + \sigma_{DE}}{V_D^*}, \end{aligned}$$

so

$$\beta_D^* = \frac{\Sigma_D \sigma^2}{V_D^*}. \quad (\text{B.2})$$

To solve for r^* investors maximize

$$\alpha^* r^* + (1 - \alpha^*) E[R] - k(1 - \alpha^*)^2 \sigma^2.$$

Taking the derivative, setting it equal to zero and then setting $\alpha^* = 0$ gives

$$r^* - E[R] + 2k\sigma^2 = 0,$$

implying

$$r^* = \bar{R} - 2k\sigma^2, \quad (\text{B.3})$$

where $\bar{R} \equiv E[R]$. Substituting these expressions and β_D^* into the pricing Eq. (B.1) gives the cost of capital, $\bar{R}_D^* \equiv E[R_D^*]$,

$$\bar{R}_D^* = \bar{R} - 2k\sigma^2 + 2k\sigma^2 \beta_D \quad (\text{B.4})$$

$$= \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_D}{V_D^*}\right). \quad (\text{B.5})$$

Following similar logic gives

$$\bar{R}_E^* = \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_E}{V_E^*}\right). \quad (\text{B.6})$$

To get expressions for V_D^* and V_E^* in terms of primitives, substitute

$$\bar{R}_D^* = \frac{\bar{D}_D}{V_D^*} - 1 \text{ into the pricing relation (B.1),}$$

$$\frac{\bar{D}_D}{V_D^*} - (1 + r^*) = \beta_D (\bar{R} - r^*).$$

Using (B.3) and (B.2) gives

$$\frac{\bar{D}_D}{V_D^*} - (1 + \bar{R} - 2k\sigma^2) = \frac{\Sigma_D}{V_D^*} (2k\sigma^2)$$

$$V_D^* = \frac{\bar{D}_D - 2k\Sigma_D\sigma^2}{1 + \bar{R} - 2k\sigma^2} = \frac{\bar{D}_D - 2k\Sigma_D\sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2}. \quad (\text{B.7})$$

Because $V_E^* = 1 - V_D^*$, we have

$$V_E^* = \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{1 + \bar{R} - 2k\sigma^2} = \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2}. \quad (\text{B.8})$$

Finally, the difference in the cost of capital between clean and dirty stocks is

$$\begin{aligned} \Delta \bar{R}^* &\equiv \bar{R}_D^* - \bar{R}_E^* = 2k\sigma^2 \left(\frac{\Sigma_D}{V_D^*} - \frac{\Sigma_E}{V_E^*} \right) = 2k\sigma^2 (\beta_D^* - \beta_E^*) \\ &= 2k\sigma^2 (\bar{D}_D + \bar{D}_E - 2k\sigma^2) \left(\frac{\Sigma_D}{\bar{D}_D - 2k\Sigma_D\sigma^2} \right. \\ &\quad \left. - \frac{\Sigma_E}{\bar{D}_E - 2k\Sigma_E\sigma^2} \right). \end{aligned} \quad (\text{B.9})$$

Appendix C. Difference in the cost of capital with and without ESG investors

Eq. (A.15) is the difference in the cost of capital between clean and dirty stocks after a portion of investors acquire ESG preferences and trade to a new equilibrium. Eq. (B.9) is the difference in the cost of capital before the existence of ESG investors. The difference between the two is therefore the effect of ESG investors on the cost of capital:

$$\begin{aligned} \Delta \bar{R} - \Delta \bar{R}^* &= 2k\sigma^2 (\bar{D}_E + \bar{D}_D - 2k\sigma^2) \\ &\quad \times \left(\frac{\Sigma_D + \Gamma}{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)} - \frac{\Sigma_D}{\bar{D}_D - 2k\sigma^2 \Sigma_D} \right) \\ &\quad - (2k\sigma^2)^2 \Gamma \left(\frac{\Sigma_D + \Gamma}{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)} - \frac{\Sigma_E}{\bar{D}_E - 2k\sigma^2 \Sigma_E} \right) \\ &= 2k\sigma^2 \Gamma (\bar{D}_E + \bar{D}_D - 2k\sigma^2) \\ &\quad \times \left(\frac{\Sigma_D}{(\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma))(\bar{D}_D - 2k\sigma^2 \Sigma_D)} \right) \\ &\quad - (2k\sigma^2)^2 \Gamma \left(\frac{\Sigma_D \bar{D}_E - \Sigma_E \bar{D}_D + \Sigma_E \Gamma}{(\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma))(\bar{D}_E - 2k\sigma^2 \Sigma_E)} \right) \\ &= 2k\sigma^2 \Gamma \beta_D^m \left(\frac{(1 + \bar{R} - 2k\sigma^2)}{(1 + \bar{R} - 2k\sigma^2(1 + \Gamma)) V_D(1 + \bar{R}_D - 2k\sigma^2 \beta_D^m)} \right) \\ &\quad - (2k\sigma^2)^2 \Gamma \\ &\quad \times \left(\frac{\beta_D^m(1 + \bar{R}_E) - \beta_E^m(1 + \bar{R}_D) + \beta_E^m \frac{\Gamma}{V_D}}{(1 + \bar{R}_D - 2k\sigma^2 (\beta_D^m + \frac{\Gamma}{V_D})) (1 + \bar{R}_E - 2k\sigma^2 \beta_E^m)} \right) \\ &= 2k\sigma^2 V_D \Gamma_R \beta_D^m \\ &\quad \times \left(\frac{(1 + \bar{R} - 2k\sigma^2)}{(1 + \bar{R} - 2k\sigma^2(1 + \Gamma_R V_D^*)) (1 + \bar{R}_D - 2k\sigma^2 \beta_D^m)} \right) \\ &\quad \times \left(\frac{\sigma_{R_D}^2}{\sigma^2} \right) \\ &\quad - (2k\sigma^2 V_D)^2 \Gamma_R \\ &\quad \times \left(\frac{\beta_D^m(1 + \bar{R}_E) - \beta_E^m(1 + \bar{R}_D) + \beta_E^m \Gamma_R V_D}{(1 + \bar{R}_D - 2k\sigma^2 (\beta_D^m + \Gamma_R V_D)) (1 + \bar{R}_E - 2k\sigma^2 \beta_E^m)} \right) \\ &\quad \times \left(\frac{\sigma_{R_D}^2}{\sigma^2} \right) \end{aligned} \quad (\text{C.1})$$

where $\Gamma_R \equiv \frac{\Gamma}{V_D^2} \left(\frac{\sigma^2}{\sigma_{R_D}^2} \right) = \left(\frac{\gamma}{1-\gamma} \right) (1 - \rho^2)$ and $\sigma_{R_D}^2 = \text{var}(R_D)$. Notice that

all the terms in the large parentheses in (C.1) are approximately equal to 1. Then note from (B.3), that prior to the advent of impact investing,

$$2k = \frac{R - r^*}{\sigma^2} = \frac{\text{MRP}}{\sigma^2},$$

where MRP is the market risk premium. If we assume that risk preferences have not changed over time then $2k\sigma^2$ equals the historical market risk premium implying that $2k\sigma^2 V_D$ is on the order of about 1%. That implies that the second term in (C.1) is so small we can ignore it. Hence, we can approximate (C.1) with

$$\Delta \bar{R} - \Delta \bar{R}^* \approx 2k\sigma^2 V_D \Gamma_R = \text{MRP} \times V_D \times \left(\frac{\gamma}{1-\gamma} \right) (1 - \rho^2). \quad (\text{C.2})$$

References

- Avramov, Doron, Cheng, Si, Lioui, Abraham, Tarelli, Andrea, 2021. Sustainable investing with ESG rating uncertainty. Available at SSRN, 3711218.
- Barber, Brad M., Morse, Adair, Yasuda, Ayako, 2021. Impact investing. *J. Financ. Econ.* 139 (1), 162–185.
- Barnea, Amir, Heinkel, Robert, Kraus, Alan, 2005. Green investors and corporate investment. *Struct. Change Econ. Dyn.* 16 (3), 332–346.

- Bennett, Benjamin, Stulz, René M., Wang, Zexi, 2020. Does Joining the S & P 500 Index Hurt Firms? Technical Report, National Bureau of Economic Research.
- Berk, Jonathan B., 1997. Necessary conditions for the CAPM. *J. Econom. Theory* 73 (1), 245–257.
- Blitz, David, Fabozzi, Frank J., 2017. Sin stocks revisited: Resolving the sin stock anomaly. *J. Portf. Manag.* 44 (1), 105–111.
- Brest, Paul Andrew, Gilson, Ronald J., Wolfson, Mark A., 2018. How investors can (and can't) create social value.
- Broccardo, Eleonora, Hart, Oliver D., Zingales, Luigi, 2020. Exit vs. Voice. Working Paper 27710, National Bureau of Economic Research.
- Chava, Sudheer, 2014. Environmental externalities and cost of capital. *Manage. Sci.* 60 (9), 2223–2247.
- Cheng, Gong, Jondeau, Eric, Mojon, Benoit, Vayanos, Dimitri, 2023. The Impact of Green Investors on Stock Prices. Technical Report, Bank for International Settlements.
- Cornell, Bradford, 2020. ESG investing: Conceptual issues. *J. Wealth Manag.* 23 (3), 61–69.
- Dam, Lammertjan, Scholtens, Bert, 2015. Toward a theory of responsible investing: On the economic foundations of corporate social responsibility. *Resour. Energy Econ.* 41, 103–121.
- Dangl, Thomas, Halling, Michael, Yu, Jin, Zechner, Josef, 2023. Social preferences and corporate investment. Available at SSRN 4470567.
- Derwall, Jeroen, Verwijmeren, Patrick, 2007. Corporate Social Responsibility and the Cost of Equity Capital. Erasmus University Working Paper.
- Dimson, Elroy, Karakas, Oğuzhan, Li, Xi, 2015. Active ownership. *Rev. Financ. Stud.* 28 (12), 3225–3268.
- Durand, Robert B., Koh, SzeKee, Tan, Paul LiJian, 2013. The price of sin in the Pacific-Basin. *Pac.-Basin Finance J.* 21 (1), 899–913.
- Fabozzi, Frank J., Ma, K.C., Oliphant, Becky J., 2008. Sin stock returns. *J. Portf. Manag.* 35 (1), 82–94.
- Feng, Zhi-Yuan, Wang, Ming-Long, Huang, Hua-Wei, 2015. Equity financing and social responsibility: further international evidence. *Int. J. Account.* 50 (3), 247–280.
- Ghoul, Sadok El, Guedhami, Omrane, Kwok, Chuck C.Y., Mishra, Dev R., 2011. Does corporate social responsibility affect the cost of capital? *J. Bank. Financ.* 35 (9), 2388–2406.
- Greenwood, Robin, Sammon, Marco C., 2022. The Disappearing Index Effect. Technical Report, National Bureau of Economic Research.
- Harris, Lawrence, Gurel, Eitan, 1986. Price and volume effects associated with changes in the S & P 500 list: New evidence for the existence of price pressures. *J. Finance* 41 (4), 815–829.
- Heinkel, Robert, Kraus, Alan, Zechner, Josef, 2001. The effect of green investment on corporate behavior. *J. Financ. Quant. Anal.* 36 (4), 431–449.
- Hong, Harrison, Kacperczyk, Marcin, 2009. The price of sin: The effects of social norms on markets. *J. Financ. Econ.* 93 (1), 15–36.
- Krueger, Philipp, Sautner, Zacharias, Starks, Laura T., 2020. The importance of climate risks for institutional investors. *Rev. Financ. Stud.* 33 (3), 1067–1111.
- Lintner, John, 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Rev. Econ. Stat.* 47 (1), 13–37.
- Luo, H. Arthur, Balvers, Ronald J., 2017. Social screens and systematic investor boycott risk. *J. Financ. Quant. Anal.* 52 (1), 365–399.
- Mollet, Janick Christian, Ziegler, Andreas, 2014. Socially responsible investing and stock performance: New empirical evidence for the US and European stock markets. *Rev. Financ. Econ.* 23 (4), 208–216.
- Mossin, Jan, 1966. Equilibrium in a capital asset market. *Econometrica* 34 (4), 768–783.
- Oehmke, Martin, Opp, Marcus M., 2023. A Theory of Socially Responsible Investment. Swedish House of Finance Research Paper, (20-2).
- Pástor, Luboš, Stambaugh, Robert F., Taylor, Lucian A., 2020. Sustainable investing in equilibrium. *J. Financ. Econ.*
- Pastor, Lubos, Stambaugh, Robert F., Taylor, Lucian A., 2021. Dissecting Green Returns. Technical Report, National Bureau of Economic Research.
- Pedersen, Lasse Heje, Fitzgibbons, Shaun, Pomorski, Lukasz, 2020. Responsible investing: The ESG-efficient frontier. *J. Financ. Econ.*
- Renneboog, Luc, Horst, Jenke Ter, Zhang, Chendi, 2008. Socially responsible investments: Institutional aspects, performance, and investor behavior. *J. Bank. Financ.* 32 (9), 1723–1742.
- Riedl, Arno, Smeets, Paul, 2017. Why do investors hold socially responsible mutual funds? *J. Finance* 72 (6), 2505–2550.
- Salaber, Julie, 2013. Religion and returns in Europe. *Eur. J. Political Econ.* 32, 149–160.
- Sharpe, William F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance* 19 (3), 425–442.
- Shleifer, Andrei, 1986. Do demand curves for stocks slope down? *J. Finance* 41 (3), 579–590.
- SIF, 2020. Report on US Sustainable and Impact Investing Trends. Technical Report, The Forum for Sustainable and Responsible Investment.
- Statman, Meir, Glushkov, Denys, 2016. Classifying and measuring the performance of socially responsible mutual funds. *J. Portf. Manag.* 42 (2), 140–151.
- Teoh, Siew Hong, Welch, Ivo, Wazzan, C. Paul, 1999. The effect of socially activist investment policies on the financial markets: Evidence from the South African boycott. *J. Bus.* 72 (1), 35–89.
- Treynor, Jack, 1961. Toward a theory of the market value of risky assets.
- Trinks, Arjan, Scholtens, Bert, Mulder, Machiel, Dam, Lammertjan, 2018. Fossil fuel divestment and portfolio performance. *Ecol. Econom.* 146, 740–748.
- van Binsbergen, Jules H., 2024. Duration-based stock valuation: reassessing stock market performance and volatility. *J. Financ. Econ. Forthcoming*.