

The Impact of Impact Investing*

Jonathan B. Berk

Stanford University and NBER

Jules H. van Binsbergen

University of Pennsylvania and NBER

August 27, 2021

Abstract

We evaluate the quantitative effect of ESG divestitures on the cost of capital of affected firms. We derive a simple expression for the change in the cost of capital as a function of three inputs: (1) the fraction of socially conscious capital, (2) the fraction of targeted firms in the economy and (3) the correlation between the targeted firms and the rest of the stock market. Given the current state of ESG investment we find that the impact on the cost of capital is too small to meaningfully affect real investment decisions. We empirically corroborate these small estimates by studying firm changes in ESG status. When firms are either included or excluded from the leading socially conscious US index (FTSE USA 4Good) we find no detectable effect on the cost of capital. 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.

*We thank Greg Buchak, Peter DeMarzo, Arvind Krishnamurthy, Christian Opp, Amit Seru and seminar participants at Stanford GSB for their helpful comments.

Impact 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 stock, 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 divestiture 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 by determining whether or not they have materially affected the cost of capital and, if not, whether they 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 assumptions that underlie the standard model in financial economics, the Capital Asset Pricing Model (CAPM), 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

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

where MRP is the historical market risk premium. We estimate, using the holdings of the largest socially conscious index fund over the past 5 years, that f is 48.5% and ρ is 0.97. Further, we estimate that currently socially conscious wealth makes up about 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 0.35 basis points. Given the uncertainty in the capital budgeting process, one third of a basis point cannot meaningfully impact firms' investment strategies.

We then show that the empirical data is consistent with this theoretical prediction. Even with the growth in the popularity of impact investing in the last 10 years, we find no detectable difference in the cost of capital between firms that are targeted for their social or environmental costs and firms that are not. Specifically, we study the effect of a firm either being included or excluded from the FTSE USA 4 Good index. The stocks in this index are a strict subset of the FTSE USA index, and so the inclusion and exclusion events are driven purely by changes in the social status of the firm. The index is also widely used — the world's largest socially responsible index fund, the Vanguard FTSE Social Index Fund, tracks this index. 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. In line with our theory, we find that the effect of a change in social status of a firm is very small and not statistically different from zero.

Finally, we consider the question of what it would take for a divestment strategy to successfully impact firm investment. Using the most optimistic estimates, we show that 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 question the effectiveness of disinvestment. On the other hand, the set of companies that are targeted all socially conscious funds comprise only 18% of the market, so a more effective strategy might be to do the opposite. Instead of divesting, socially conscious investors could purchase the stock in these companies and effect change through the proxy process or by gaining a majority stake and replacing upper management. To successfully implement this strategy would require less than 50% participation.

The reason divestiture 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 enough relative to the supply of investable capital and stocks 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 for most investors to choose to divest. In this case, because the shareholders that must hold the targeted 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.

1 Background

This paper fits into a sizable literature that studies impact investing.¹ In this paper, our focus is on divestment. One of the earliest paper that studies the effect of a divestment strategy is Teoh, Welch and Wazzan (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 industries.² 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.³ The advantage of our empirical approach is that it does not require us to take a stand 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.

Other factors to consider in interpreting the results of these studies is that in the sample period (60’s to early 2000’s) there was very little organized pressure to divest

¹Brest, Gilson and Wolfson (2018), in an article intended for a law review audience, survey the entire impact investing landscape.

²See Hong and Kacperczyk (2009), Chava (2014), El Ghoul, Guedhami, Kwok and Mishra (2011), Fabozzi, Ma and Oliphant (2008), Statman and Glushkov (2016), Barber, Morse and Yasuda (2021), Derwall and Verwijmeren (2007) and Salaber (2013) and references therein.

³See Blitz and Fabozzi (2017).

(other than in the case of South Africa) and studies using more recent data have failed to find the effect (Mollet and Ziegler (2014) and Trinks, Scholtens, Mulder and Dam (2018)). The effect is also not consistent across localities, as documented in Durand, Koh and Tan (2013), who attribute the differences to “cultural norms,” and Feng, Wang and Huang (2015).⁴

The first paper that theoretically models how a divestiture strategy affects corporate behavior is Heinkel, Kraus and Zechner (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 so that a significant fraction of firms would 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.⁵

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 this 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. Based on the graphs in their article, this estimate is likely very sensitive to the choice of sin industries, in particular, coal and mining. It is also very large — it seems somewhat implausible that investors who do not care about the social consequences of their investments would willingly pass over such a large return.

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, Fitzgibbons and Pomorski (2020) derive a model

⁴Renneboog, Ter Horst and Zhang (2008) provide review of the evidence in the literature on the effectiveness of impact investing.

⁵In a followup paper, Barnea, Heinkel and Kraus (2005), endogenize firm investment and study how total investment is affected by divestiture.

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, Stambaugh and Taylor (2020) also derive a model featuring agents with ESG preferences and study the security market line in this environment. That paper does not study the effect on the cost of capital of introducing ESG investors into the economy. Avramov, Cheng, Lioui and Tarelli (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 disinvesting, in socially costly firms and using their rights of control to alter corporate policy. Dimson, Karakaş and Li (2015) provide empirical evidence supporting the effectiveness of this alternative strategy.

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 (as is often assumed in existing papers that model this behavior). Another possibility is that investors use the strategy to signal “good behavior.” Riedl and Smeets (2017) provide evidence that supports this latter hypothesis.

A related line of research studies the effect on prices of demand shocks. In a neoclassical environment, since the supply of stocks is not fixed, one would expect little effect. In fact, studies have documented large price effects (see Koijen and Yogo (2019) and Koijen, Richmond and Yogo (2019)). Conceivably the effect on the cost of capital could operate through these demand effects. However, even if one takes the documented effects at face value, since it is not clear that they are permanent, it is not clear that they affect the long term cost of capital. Furthermore, the extent to which the measured effects reflect estimation error is subject to debate, mainly because it is difficult to understand how so much value could be left on the table in markets as competitive as U.S. stock markets.

2 The Model

To assess the impact of divestment, we consider the question in the work horse economy in financial economics, the single period Capital Asset Pricing Model (CAPM) developed by Sharpe (1964), Mossin (1966), Lintner (1965) and Treynor (1961). To derive that equilibrium, assume all investors are mean-variance maximizers,⁶ 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 the investor's risk appetite.

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. Thus, our results represent an upper bound on the effect of ESG investors, because including other margins of adjustment in our computations would weaken the effects of equity divestment initiatives on real investment decisions.

We assume that a set of investors exist who have ESG preferences in that they will only hold clean stocks, for simplicity, *ESG investors*. We will 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 will 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 will 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

⁶Berk (1997) provides the necessary and sufficient conditions that imply mean-variance maximization.

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 $R_D = \frac{D_D}{V_D} - 1$ respectively 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:

$$\begin{aligned} \sigma_D^2 &\equiv \text{var}(D_D) \\ \sigma_E^2 &\equiv \text{var}(D_E) \\ \sigma_{ED} &\equiv \text{cov}(D_E, D_D) \\ \sigma^2 &\equiv \text{var}(D) = \text{var}(D_D + D_E) = \sigma_D^2 + \sigma_E^2 + 2\sigma_{DE} = \text{var}(R) \\ \rho &\equiv \frac{\sigma_{ED}}{\sigma_D\sigma_E} = \frac{\text{cov}(R_E, R_D)}{\text{std}(R_D)\text{std}(R_E)}. \end{aligned}$$

Notice that because of the price normalization, the cashflow variance of the market is the same as the return variance. Because of the nature of correlations, the cashflow 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.⁷ So, the ESG investors' equilibrium allocation is e fraction of the clean

⁷If 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.

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)),

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

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 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:

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

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 $\beta_E^m \equiv \frac{\Sigma_E}{V_E}$:

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

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

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

2.2 Equilibrium without ESG investors

To assess the effect of ESG investors, we now derive the equilibrium when all investors are identical. 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 return of the two portfolios are now given

by

$$\begin{aligned}
V_E^* &= \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2} \\
V_D^* &= \frac{\bar{D}_D - 2k\Sigma_D\sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2} \\
\bar{R}_E^* &= \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_E}{V_E^*}\right) \\
\bar{R}_D^* &= \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_D}{V_D^*}\right).
\end{aligned}$$

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

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

where $\beta_D^* = \frac{\Sigma_D}{V_D^*}$ and $\beta_E^* = \frac{\Sigma_E}{V_E^*}$ 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. Equation (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:

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

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 (33) 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 a first order 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 change prices (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 (41) are observable, we can use this approximation to infer to what extent impact investing 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 index, to identify the clean portfolio. The FTSE USA 4 Good 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-capitalisation weighted index representing the performance of US large and mid cap stocks. Thus, the Vanguard FTSE Social Index Fund index fund invests exclusively in the United States, and with a current AUM of \$12 Billion, is the largest social index fund in the world.

The evolution of the combined total AUM invested in the three share classes of the




Figure 1: Combined total AUM invested in the three share classes of the Vanguard FTSE Social Index Fund in Million of dollars.

Vanguard FTSE Social Index Fund is plotted in Figure 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 4.1), in the last 5 years of our sample (December 2015-December 2020), the dirty portfolio comprised a little under half of US market capitalization: $V_D = 48.5\%$. 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.⁸ The measured correlation between the clean (FTSE USA 4 Good) and the dirty (the other stocks in the FTSE USA) portfolio over this 5-year period is $\rho = 0.97$. We use the historical market risk premium of 6%. The only remaining quantity to identify is the fraction of wealth controlled by ESG investors. We use the fraction of mutual fund wealth invested in ESG mutual funds, 2%, assuming this fraction is representative of all capital investment.⁹ Using these parameters, the effect on the cost of capital is

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

A difference of one third of a basis point cannot meaningfully affect the capital budgeting

⁸As 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.

⁹We 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 funds.

decision and so an effect of this size cannot affect real investment decision making.

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 good 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 to be clean), we estimate V_D to be 18% and ρ to be 0.8. These two changes partially offset each other resulting in an estimated change in the cost of capital of

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

The possibility exists that the mutual fund sector might not be representative of impact investing as a whole because impact investors might be more concentrated in other sectors. We therefore looked for other sources to calibrate the fraction of impact investors. One source is an estimate reported in the *Report on US Sustainable and Impact Investing Trends* (2020) published by SIF, an organization that represents sustainable investors. The report surveys investors asking them whether they consider any ESG criteria in their investments 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. Nevertheless, using this estimate gives

$$\Delta\bar{R} - \Delta\bar{R}^* = 6\% \times 48.5\% \left(\frac{0.33}{1 - 0.33} \right) (1 - 0.97^2) = 8.5 \text{ b.p.}$$

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

One might ask how many ESG investors would it take to materially affect the firm's real investment decision. To answer that question we calibrate the model as follows. We use the return volatility of the clean portfolio of $\sigma_{RE} = 15\%$, based on the monthly return volatility of the FTSE 4 Good 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 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 $\sigma_{RD} = 15\%$. The correlation between these two portfolios over the same 5-year period is $\rho = 0.97$. We use a risk free rate of 2%, MRP of 6% and $V_D = 48.5\%$, 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 and solving for the equilibrium in the economy without ESG investors. Taking the exogenous parameters as given, Figure 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 Figure 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 impact investors, that is, (40). The dashed curve in Figure 2 is (10), the approximation to (40). 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 86% of investors choose to hold only clean stocks.

We can also infer what the impact would be if the largest investor in the world decided to divest all dirty stocks. For example, the largest investor in the world, Blackrock, manages about \$8 trillion, 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 3.7 b.p.

As the difference between the red and black dashed curves in Figure 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 ESG investors need to be such a large fraction of investors.




Figure 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, (40). The dashed black curve is the first order approximation, that is, (10).

To first order, the last three variables are not under the control of ESG investors. Using current estimates we have

$$6\% \times 48.5\%(1 - 0.97^2) = 0.00172.$$

This calculation constrains the effectiveness of impact investing. For (41) to reach 1%, we need $\frac{\gamma}{1-\gamma} > 5.8$, implying that $\gamma > 85\%$. There are two reasons for this. First, the correlation coefficient of 0.97 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 force non-ESG investors to hold only dirty stocks, which is what happens when γ approaches one.

4 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. As we explained in Section 1, there are a number of studies that have looked at this question, although there is no consensus conclusion from those studies on the effect of ESG investors on the cost of capital. One possible explanation for this lack of consensus is that those studies rely on risk models to measure abnormal performance. Because many dirty stocks are concentrated in particular industries, the results might reflect risk differences that are uncontrolled for by the models.

To properly assess the effect of ESG investors on the cost of capital, we adopt a different approach. We focus on changes in ESG classifications, that is, the effect on the cost of capital when a firm either becomes clean or becomes 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.

We use the same two 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 two 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 index, when FTSE Russell changes the constituents of the FTSE USA 4 Good 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 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 index likely also reflects the investment decisions of other ESG investors in the economy. Finally, the fact that the FTSE USA 4 Good index is a strict 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.

4.1 Data and Descriptive Statistics: FTSE Indices

We obtained data from FTSE on the constituents and their portfolio weights for both indices starting in the early 2000s until the end of 2020.¹⁰ The number of constituents over this time sample for both indices is plotted in Figure 3. 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.

In Figure 4 we plot the total market capitalization of both 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 index was \$18.3 Trillion which equals 55% of the FTSE USA stock market capitalization, and a little under half of the total US stock market capitalization. This implies that the dirty stocks in the FTSE USA index make up 45% of the index at that point in time. At the beginning of the sample (June 2001), the stock market capitalization of the FTSE USA index is \$10.8 Trillion, that of the FTSE USA 4 Good index is 5.2 Trillion (48% of the FTSE USA index, implying a portfolio weight of dirty stocks of 52%) and the total U.S. stock market capitalization was \$14.7 Trillion. In our calibration, we have used the

¹⁰We thank FTSE for generously sharing their data with us.




Figure 3: Number of constituents of the FTSE USA and the FTSE USA 4 Good Indices.

average of the portfolio weight of dirty stocks in the FTSE USA over the last five years of our data set (December 2015-December 2020), which equals 48.5%.




Figure 4: Total Market Capitalization of the FTSE USA and the FTSE USA 4 Good Indices in \$ Trillions.

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 the index. First define the dummy that describes whether firm i is in the FTSE USA index at time t :

$$I_{i,t} = \begin{cases} 1 & \text{if firm } i \text{ is in the FTSE USA index at time } t \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

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

$$I_{i,t}^{4G} = \begin{cases} 1 & \text{if firm } i \text{ is in the FTSE USA 4 Good index at time } t \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

The inclusion and exclusion events can then be described as:

$$\Delta I_{it} \equiv I_{i,t} - I_{i,t-1} \quad (13)$$

$$\Delta I_{it}^{4G} \equiv I_{i,t}^{4G} - I_{i,t-1}^{4G}. \quad (14)$$

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 both indices.

Table 1: Inclusions and Exclusions

	No. of events
$\Delta I_{it} = 1$	872
$\Delta I_{it} = -1$	795
$\Delta I_{it}^{4G} = 1$	411
$\Delta I_{it}^{4G} = -1$	385
$\Delta I_{it} = 1 \ \& \ \Delta I_{it}^{4G} = 1$	54
$\Delta I_{it} = -1 \ \& \ \Delta I_{it}^{4G} = -1$	200

The table shows that there are many inclusion and exclusion events in our sample. The total number of firms that have been in the FTSE USA and the FTSE USA 4 Good index at some point during our sample period (2001-2020) equals 1,474 and 589 respectively. Based on the number of inclusion and exclusion events it is clear that both indices have experienced turnover over the sample period. The table also assesses when inclusions and exclusions coincide between the two indices. 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.

To further explore the relation between inclusion in the two indices, we plot in Figure 5 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 each of the 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 represent 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.

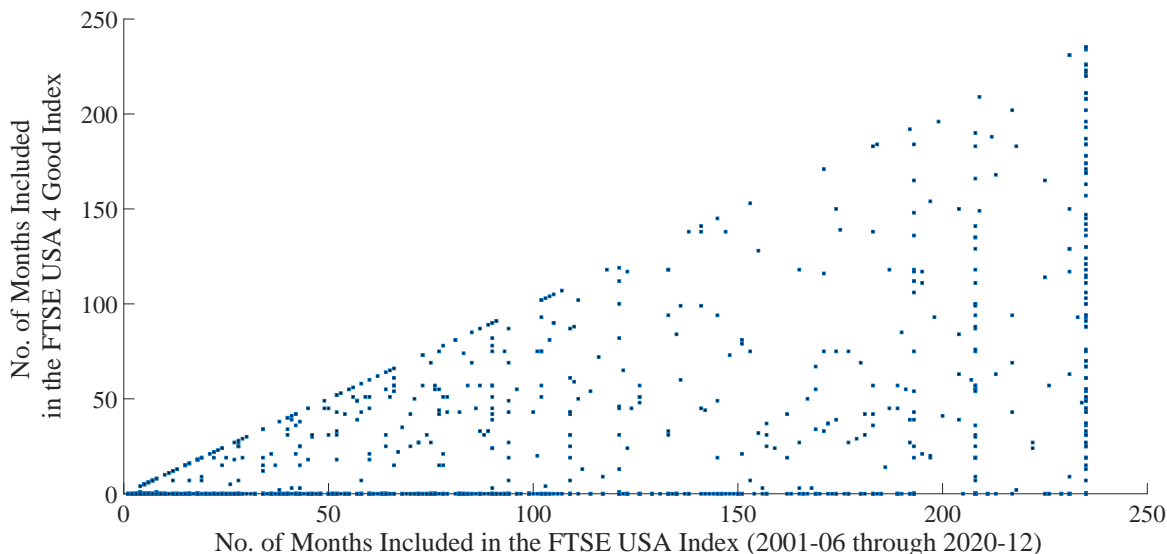


Figure 5: Number of total months included in the FTSE USA and the FTSE USA 4 Good Indices.

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 the FTSE index data provides ISINs, and the CRSP database provides CUSIPs, we use the S&P Global Market Intelligence database to provide the mapping to merge the two. Because CUSIPs change over time and companies can issue multiple securities, we use the CRSP identifier, PERMCO, to identify the single security associated with each holding.

4.2 Results Gauging Price and Return Effects

Our theory predicts that the effect on returns of ESG investing will be small, or, equivalently, we should not observe a large effect on the cost of capital when a stock experiences either an inclusion or an exclusion event. In this subsection we test this implication of the theory.

Estimating differences in expected returns, given the limited data, is difficult. On the other hand, under the rational expectations hypothesis, changes in expected returns are reflected in price changes. Because a price change reflects the capitalized value of a change in expected returns, it is potentially easier to identify. For that reason we will adopt a two prong approach and test for both changes in average returns following an inclusion and exclusion event, as well as price changes contemporaneous with the events.

Our main regression specification is the following:

$$R_{it} = c + \gamma I_{it} + \delta \Delta I_{it} + \gamma_{4G} I_{it}^{4G} + \delta_{4G} \Delta I_{it}^{4G} + \varepsilon_{it} \quad (15)$$

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. 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 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 average returns following the instantaneous price appreciation.

Table 2 reports the results. In line with our theory, we find small estimates of both coefficients. The point estimate of the instantaneous price appreciation is 0.24% with a t-statistic of 0.62 (specification (1) in Table 2). Although the point estimate is not statistically different from zero, it is in line with the order of magnitude estimate predicted by the theory (since this is the capitalized value of a change in the cost of capital, the inferred change in the cost of capital itself is smaller than this estimate). Given the difficulty in estimating expected return differences, it is perhaps not too surprising that the coefficient on the ΔI_{it} dummy is not precisely estimated and has the wrong sign.

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. The events are associated with an instantaneous price change of 0.81% and a lower average return of 10 basis points per month (both statistically significant) suggesting a role for liquidity. Specification (3) in Table 2 explores this further by measuring the prices effects in the months following the event. What is evident is that the price effect quickly reverses suggesting that the price change is associated with the price pressure from participants seeking to match the index, rather than a permanent change in liquidity due to the fact the stock is included in the index.

In the remaining columns of Table 2 we explore various alternative specifications 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 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 index, suggest that the overall price effect of inclusion and exclusion events is actually smaller than the estimate in specification (1). The overall conclusion is that the effect on the cost of capital of a change in ESG status of a firm is so small (if it exists at all) that it cannot meaningfully influence the firm's investment decision.

5 Conclusion

In this paper we have evaluated, both theoretically and empirically, the quantitative impact of socially conscious investing. We conclude that at current levels impact investing is unlikely to have a large impact on the long-term cost of capital of targeted firms. A substantial increase in the amount of socially conscious capital is required for the strategy to affect corporate policy. Given the current levels of socially conscious capital, a more effective strategy to put that capital to use is to follow a policy of engagement. By

Dependent Variable: R_{it}				
	(1)	(2)	(3)	(4)
I_{it}	-0.0010 [-2.93]	-0.0011 [-3.03]	-0.0014 [-3.90]	-0.0013 [-3.74]
I_{it}^{4G}	0.0006 [1.00]	0.0005 [0.91]	0.0012 [2.11]	0.0012 [2.08]
ΔI_{it}	0.0081 [2.14]		0.0073 [1.86]	
$\Delta I_{i,t-1}$		-0.0008 [-0.20]	0.0004 [0.09]	-0.0006 [-0.15]
$\Delta I_{i,t-2}$			0.0085 [2.11]	0.0083 [2.07]
$\Delta I_{i,t-3}$			-0.0046 [-1.36]	-0.0047 [-1.38]
$\Delta I_{i,t-4}$			-0.0079 [-1.97]	-0.0080 [-1.99]
ΔI_{it}^{4G}	0.0024 [0.62]		0.0022 [0.46]	
$\Delta I_{i,t-1}^{4G}$		0.0051 [1.12]	0.0031 [0.66]	0.0026 [0.57]
$\Delta I_{i,t-2}^{4G}$			-0.0126 [-2.21]	-0.0126 [-2.22]
$\Delta I_{i,t-3}^{4G}$			-0.0022 [-0.52]	-0.0024 [-0.57]
$\Delta I_{i,t-4}^{4G}$			-0.0085 [-1.71]	-0.0084 [-1.70]
Constant	0.0098 [62.65]	0.0101 [63.84]	0.0107 [67.66]	0.0107 [67.63]

Table 2: Return and Price Effects of Index Inclusions: FTSE USA vs. FTSE USA 4 Good. The table reports regression results of the type presented in Equation 15. 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 clustered by yearmonth.

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.

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

$$\beta_E \equiv \frac{\text{cov}(R_E, R_{mv})}{\text{var}(R_{mv})} = \frac{1 - \gamma + b}{V_E} \left(\frac{\Sigma_E \sigma^2 - e \sigma_E^2}{\sigma^2 - 2e \Sigma_E \sigma^2 + e^2 \sigma_E^2} \right), \quad (16)$$

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

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

Similarly, the beta of the dirty portfolio, β_D is

$$\beta_D \equiv \frac{1 - \gamma + b}{V_D} \left(\frac{\Sigma_D \sigma^2 - e \rho \sigma_D \sigma_E}{\sigma^2 - 2e \Sigma_E \sigma^2 + e^2 \sigma_E^2} \right) \quad (17)$$

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

$$\begin{aligned} r - E[R_E] + 2k \left(1 - \frac{b}{\gamma}\right) \text{var}[R_E] &= 0 \\ 1 + r &= \frac{\bar{D}_E}{V_E} - 2k \left(1 - \frac{b}{\gamma}\right) \frac{\sigma_E^2}{V_E^2}. \end{aligned} \quad (18)$$

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

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

$$1 + r = \frac{\bar{D}_D + (1-e)\bar{D}_E}{1-\gamma+b} - 2k \frac{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2}{(1-\gamma)(1-\gamma+b)}. \quad (19)$$

Equations (18) and (19) jointly determine r and b .

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

$$E[R_E] - r = \beta_E(E[\omega_D R_D + \omega_E R_E] - r). \quad (20)$$

Substituting (16) and (3) into this expression gives:

$$\frac{\bar{D}_E}{V_E} - (1+r) = \frac{1-\gamma+b}{V_E} \left(\frac{\Sigma_E\sigma^2 - e\sigma_E^2}{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2} \right) \left(\frac{\bar{D}_D + (1-e)\bar{D}_E}{1-\gamma+b} - (1+r) \right). \quad (21)$$

Substituting (19) into the right side of (21) and rearranging terms gives,

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

Substituting (18) into the left hand side of this expression and simplifying provides,

$$\gamma - b = V_E\gamma \left(1 + \frac{\rho\sigma_D}{\sigma_E} \right). \quad (22)$$

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

$$e = \gamma \left(1 + \frac{\rho\sigma_D}{\sigma_E} \right) \quad (23)$$

Substituting (22) into (18) gives r in terms of primitives:

$$1 + r = \frac{\bar{D}_E}{V_E} - 2k \left(1 + \frac{\rho\sigma_D}{\sigma_E} \right) \frac{\sigma_E^2}{V_E}. \quad (24)$$

Finally, substituting (23) into (19) provides b in terms of primitives:

$$\gamma - b = 1 - \frac{\bar{D}_D + (1 - e)\bar{D}_E}{1 + r} + 2k \frac{\sigma^2 - 2e\Sigma_E\sigma^2 + e^2\sigma_E^2}{(1 - \gamma)(1 + r)}. \quad (25)$$

Using the expressions for the risk free rate and equilibrium holdings, (23), (24) and (25), we can solve (19) for V_E :

$$\begin{aligned} V_E &= \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_E + \bar{D}_D - 2k\sigma^2 \left(1 + \left(\frac{\gamma}{1-\gamma}\right) (1 - \rho^2) \frac{\sigma_D^2}{\sigma^2}\right)} \\ &= \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_E + \bar{D}_D - 2k\sigma^2 (1 + \Gamma)}, \end{aligned} \quad (26)$$

where Γ is defined by (6). Rearranging terms in (26) gives

$$\begin{aligned} \bar{1} + \bar{R}_E &= \frac{\bar{D}_E}{V_E} = \bar{D}_E + \bar{D}_D - 2k\sigma^2 (1 + \Gamma) + 2k\sigma^2 \frac{\Sigma_E}{V_E} \\ \bar{R}_E &= \bar{R} + 2k\sigma^2 (\beta_E^m - (1 + \Gamma)) \end{aligned} \quad (27)$$

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

$$\begin{aligned} V_D &= 1 - V_E \\ &= \frac{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)}{\bar{D}_E + \bar{D}_D - 2k\sigma^2 (1 + \Gamma)} \end{aligned} \quad (28)$$

so

$$\begin{aligned} \bar{R}_D &= \bar{R} - 2k\sigma^2 \left(1 + \Gamma - \frac{\Sigma_D + \Gamma}{V_D}\right) \\ &= \bar{R} + 2k\sigma^2 \left(\beta_D^m - \left(1 - \Gamma \frac{V_E}{V_D}\right)\right) \end{aligned} \quad (29)$$

where $\beta_D^m \equiv \frac{\Sigma_D}{V_D}$ is the *market* beta of the dirty portfolio. Using (27) and (29), the difference in the expected return of dirty and clean stocks, $\Delta\bar{R}$, is

$$\begin{aligned} \Delta\bar{R} &\equiv R_D - R_E = 2k\sigma^2 \left(\beta_D^m - \beta_E^m + \frac{\Gamma}{V_D}\right) \\ &= 2k\sigma^2 (\bar{D}_E + \bar{D}_D - 2k\sigma^2(1 + \Gamma)) \left(\frac{\Sigma_D + \Gamma}{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)} - \frac{\Sigma_E}{\bar{D}_E - 2k\sigma^2 \Sigma_E}\right) \end{aligned} \quad (30)$$

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

$$E[R_D^*] - r^* = \beta_D(E[R] - r^*), \quad (31)$$

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.¹¹ Now

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

so

$$\beta_D^* = \frac{\Sigma_D \sigma^2}{V_D^*}. \quad (32)$$

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

$$r^* = \bar{R} - 2k\sigma^2, \quad (33)$$

¹¹Because 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.

where $\bar{R} \equiv E[R]$. Substituting these expressions and β_D^* into the pricing equation (31) gives the cost of capital, $\bar{R}_D^* \equiv E[R_D^*]$,

$$\bar{R}_D^* = \bar{R} - 2k\sigma^2 + 2k\sigma^2\beta_D \quad (34)$$

$$= \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_D}{V_D^*}\right). \quad (35)$$

Following similar logic gives

$$\bar{R}_E^* = \bar{R} - 2k\sigma^2 \left(1 - \frac{\Sigma_E}{V_E^*}\right). \quad (36)$$

To get expressions for V_D^* and V_E^* in terms of primitives, substitute $\bar{R}_D^* = \frac{\bar{D}_D}{V_D^*} - 1$ into the pricing relation (31),

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

Using (33) and (32) gives

$$\begin{aligned} \frac{\bar{D}_D}{V_D^*} - (1 + \bar{R} - 2k\sigma^2) &= \frac{\Sigma_D}{V_D^*} (2k\sigma^2) \\ V_D^* &= \frac{\bar{D}_D - 2k\Sigma_D\sigma^2}{1 + \bar{R} - 2k\sigma^2} = \frac{\bar{D}_D - 2k\Sigma_D\sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2}. \end{aligned} \quad (37)$$

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

$$V_E^* = \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{1 + \bar{R} - 2k\sigma^2} = \frac{\bar{D}_E - 2k\Sigma_E\sigma^2}{\bar{D}_D + \bar{D}_E - 2k\sigma^2}. \quad (38)$$

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

$$\begin{aligned} \Delta\bar{R}^* &\equiv \bar{R}_D^* - \bar{R}_E^* = 2k\sigma^2 \left(\frac{\Sigma_D}{V_D^*} - \frac{\Sigma_E}{V_E^*} \right) = 2k\sigma^2 (\beta_D^* - \beta_E^*) \\ &= 2k\sigma^2 (D_D + \bar{D}_E - 2k\sigma^2) \left(\frac{\Sigma_D}{\bar{D}_D - 2k\Sigma_D\sigma^2} - \frac{\Sigma_E}{\bar{D}_E - 2k\Sigma_E\sigma^2} \right). \end{aligned} \quad (39)$$

C Difference in the Cost of Capital with and without ESG Investors

Equation (30) 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. Equation (39) 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:

$$\begin{aligned}
\Delta \bar{R} - \Delta \bar{R}^* &= 2k\sigma^2 (\bar{D}_E + \bar{D}_D - 2k\sigma^2) \left(\frac{\Sigma_D + \Gamma}{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)} - \frac{\Sigma_D}{\bar{D}_D - 2k\sigma^2 \Sigma_D} \right) \\
&\quad - (2k\sigma^2)^2 \Gamma \left(\frac{\Sigma_D + \Gamma}{\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma)} - \frac{\Sigma_E}{\bar{D}_E - 2k\sigma^2 \Sigma_E} \right) \\
&= 2k\sigma^2 \Gamma (\bar{D}_E + \bar{D}_D - 2k\sigma^2) \left(\frac{\Sigma_D}{(\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma))(\bar{D}_D - 2k\sigma^2 \Sigma_D)} \right) \\
&\quad - (2k\sigma^2)^2 \Gamma \left(\frac{\Sigma_D \bar{D}_E - \Sigma_E \bar{D}_D + \Sigma_E \Gamma}{(\bar{D}_D - 2k\sigma^2 (\Sigma_D + \Gamma))(\bar{D}_E - 2k\sigma^2 \Sigma_E)} \right) \\
&= 2k\sigma^2 \Gamma \beta_D^m \left(\frac{(1 + \bar{R} - 2k\sigma^2)}{(1 + \bar{R} - 2k\sigma^2(1 + \Gamma)) V_D (1 + \bar{R}_D - 2k\sigma^2 \beta_D^m)} \right) \\
&\quad - (2k\sigma^2)^2 \Gamma \left(\frac{\beta_D^m (1 + \bar{R}_E) - \beta_E^m (1 + \bar{R}_D) + \beta_E^m \frac{\Gamma}{V_D}}{\left(1 + \bar{R}_D - 2k\sigma^2 \left(\beta_D^m + \frac{\Gamma}{V_D}\right)\right) (1 + \bar{R}_E - 2k\sigma^2 \beta_E^m)} \right) \\
&= 2k\sigma^2 V_D \Gamma_R \beta_D^m \left(\frac{(1 + \bar{R} - 2k\sigma^2)}{(1 + \bar{R} - 2k\sigma^2(1 + \Gamma_R V_D^2)) (1 + \bar{R}_D - 2k\sigma^2 \beta_D^m)} \right) \left(\frac{\sigma_{R_D}^2}{\sigma^2} \right) \quad (40) \\
&\quad - (2k\sigma^2 V_D)^2 \Gamma_R \left(\frac{\beta_D^m (1 + \bar{R}_E) - \beta_E^m (1 + \bar{R}_D) + \beta_E^m \Gamma_R V_D}{(1 + \bar{R}_D - 2k\sigma^2 (\beta_D^m + \Gamma_R V_D)) (1 + \bar{R}_E - 2k\sigma^2 \beta_E^m)} \right) \left(\frac{\sigma_{R_D}^2}{\sigma^2} \right)
\end{aligned}$$

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 (40) are approximately equal to 1. Then note from (33), 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 (40) is so small

we can ignore it. Hence, we can approximate (40) with

$$\Delta\bar{R} - \Delta\bar{R}^* \approx 2k\sigma^2 V_D \Gamma_R = \text{MRP} \times V_D \times \left(\frac{\gamma}{1-\gamma} \right) (1 - \rho^2). \quad (41)$$

References

- Avramov, Doron, Si Cheng, Abraham Lioui, and Andrea Tarelli, “Sustainable investing with ESG rating uncertainty,” *Available at SSRN*, 2021, 3711218.
- Barber, Brad M, Adair Morse, and Ayako Yasuda, “Impact investing,” *Journal of Financial Economics*, 2021, 139 (1), 162–185.
- Barnea, Amir, Robert Heinkel, and Alan Kraus, “Green investors and corporate investment,” *Structural Change and Economic Dynamics*, 2005, 16 (3), 332–346.
- Berk, Jonathan B., “Necessary Conditions for the CAPM,” *Journal of Economic Theory*, 1997, 73 (1), 245–257.
- Blitz, David and Frank J Fabozzi, “Sin stocks revisited: Resolving the sin stock anomaly,” *The Journal of Portfolio Management*, 2017, 44 (1), 105–111.
- Brest, Paul Andrew, Ronald J. Gilson, and Mark A. Wolfson, “How Investors Can (and Can’t) Create Social Value,” 2018.
- Chava, Sudheer, “Environmental externalities and cost of capital,” *Management science*, 2014, 60 (9), 2223–2247.
- Dam, Lammertjan and Bert Scholtens, “Toward a theory of responsible investing: On the economic foundations of corporate social responsibility,” *Resource and Energy Economics*, 2015, 41, 103–121.
- Derwall, Jeroen and Patrick Verwijmeren, “Corporate social responsibility and the cost of equity capital,” *Erasmus University Working Paper*, 2007.
- Dimson, Elroy, Oğuzhan Karakaş, and Xi Li, “Active ownership,” *The Review of Financial Studies*, 2015, 28 (12), 3225–3268.
- Durand, Robert B, SzeKee Koh, and Paul LiJian Tan, “The price of sin in the Pacific-Basin,” *Pacific-Basin Finance Journal*, 2013, 21 (1), 899–913.
- Fabozzi, Frank J, KC Ma, and Becky J Oliphant, “Sin stock returns,” *The Journal of Portfolio Management*, 2008, 35 (1), 82–94.
- Feng, Zhi-Yuan, Ming-Long Wang, and Hua-Wei Huang, “Equity financing and social responsibility: further international evidence,” *The international journal of accounting*, 2015, 50 (3), 247–280.

- Ghoul, Sadok El, Omrane Guedhami, Chuck CY Kwok, and Dev R Mishra**, “Does corporate social responsibility affect the cost of capital?,” *Journal of Banking & Finance*, 2011, *35* (9), 2388–2406.
- Heinkel, Robert, Alan Kraus, and Josef Zechner**, “The effect of green investment on corporate behavior,” *Journal of financial and quantitative analysis*, 2001, *36* (4), 431–449.
- Hong, Harrison and Marcin Kacperczyk**, “The price of sin: The effects of social norms on markets,” *Journal of Financial Economics*, 2009, *93* (1), 15–36.
- Koijen, Ralph SJ and Motohiro Yogo**, “A demand system approach to asset pricing,” *Journal of Political Economy*, 2019, *127* (4), 1475–1515.
- Koijen, RS, Robert J Richmond, and Motohiro Yogo**, “Which investors matter for global equity valuations and expected returns?,” *Available at SSRN*, 2019.
- Lintner, John**, “The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets,” *The Review of Economics and Statistics*, 1965, *47* (1), 13–37.
- Luo, H. Arthur and Ronald J. Balvers**, “Social Screens and Systematic Investor Boycott Risk,” *Journal of Financial and Quantitative Analysis*, 2017, *52* (1), 365–399.
- Mollet, Janick Christian and Andreas Ziegler**, “Socially responsible investing and stock performance: New empirical evidence for the US and European stock markets,” *Review of Financial Economics*, 2014, *23* (4), 208–216.
- Mossin, Jan**, “Equilibrium in a Capital Asset Market,” *Econometrica*, 1966, *34* (4), 768–783.
- Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor**, “Sustainable investing in equilibrium,” *Journal of Financial Economics*, 2020.
- Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski**, “Responsible investing: The ESG-efficient frontier,” *Journal of Financial Economics*, 2020.
- Renneboog, Luc, Jenke Ter Horst, and Chendi Zhang**, “Socially responsible investments: Institutional aspects, performance, and investor behavior,” *Journal of banking & finance*, 2008, *32* (9), 1723–1742.
Report on US Sustainable and Impact Investing Trends

Report on US Sustainable and Impact Investing Trends, Technical Report, US SIF Foundation 2020.

Riedl, Arno and Paul Smeets, “Why Do Investors Hold Socially Responsible Mutual Funds?,” *The Journal of Finance*, 2017, 72 (6), 2505–2550.

Salaber, Julie, “Religion and returns in Europe,” *European Journal of Political Economy*, 2013, 32, 149–160.

Sharpe, William F, “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk,” *Journal of Finance*, 1964, 19 (3), 425–442.

Statman, Meir and Denys Glushkov, “Classifying and measuring the performance of socially responsible mutual funds,” *The Journal of Portfolio Management*, 2016, 42 (2), 140–151.

Teoh, Siew Hong, Ivo Welch, and C Paul Wazzan, “The effect of socially activist investment policies on the financial markets: Evidence from the South African boycott,” *The Journal of Business*, 1999, 72 (1), 35–89.

Treynor, Jack, “Toward a Theory of the Market Value of Risky Assets,” 1961.

Trinks, Arjan, Bert Scholtens, Machiel Mulder, and Lammertjan Dam, “Fossil fuel divestment and portfolio performance,” *Ecological economics*, 2018, 146, 740–748.