
Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

# Quantifying the Impact of Impact Investing

Andrew W. Lo, Ruixun Zhang

informs
https://pubsonline.informs.org/journal/mnsc
MANAGEMENT SCIENCE
*Articles in Advance*, pp. 1–26
ISSN 0025-1909 (print), ISSN 1526-5501 (online)



# Quantifying the Impact of Impact Investing

**Andrew W. Lo,**[a,b] **Ruixun Zhang**[c,*]

[a] MIT Sloan School of Management, MIT Laboratory for Financial Engineering, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, Massachusetts 02142; [b] Santa Fe Institute, Santa Fe, New Mexico 87501; [c] Peking University School of Mathematical Sciences, Peking University National Engineering Laboratory for Big Data Analysis and Applications, Peking University Center for Statistical Science, Peking University Laboratory for Mathematical Economics and Quantitative Finance, Beijing 100871, China

*Corresponding author

**Contact:** alo-admin@mit.edu, https://orcid.org/0000-0003-2944-7773 (AWL); zhangruixun@pku.edu.cn, https://orcid.org/0000-0002-7670-8393 (RZ)

**Received:** April 17, 2022
**Revised:** December 18, 2022; June 30, 2023
**Accepted:** August 5, 2023
**Published Online in Articles in Advance:** December 7, 2023

https://doi.org/10.1287/mnsc.2022.01168



**Abstract.** We propose a quantitative framework for assessing the financial impact of any form of impact investing, including socially responsible investing; environmental, social, and governance (ESG) objectives; and other nonfinancial investment criteria. We derive conditions under which impact investing detracts from, improves on, or is neutral to the performance of traditional mean-variance optimal portfolios, which depends on whether the correlations between the impact factor and unobserved excess returns are negative, positive, or zero, respectively. Using Treynor–Black portfolios to maximize the risk-adjusted returns of impact portfolios, we derive an explicit and easily computable measure of the financial reward or cost of impact investing as compared with passive index benchmarks. We illustrate our approach with applications to biotech venture philanthropy, a semiconductor research and development consortium, divesting from "sin" stocks, ESG investments, and "meme" stock rallies such as GameStop in 2021.

**History:** Accepted by Agostino Capponi, finance.
**Funding:** Research funding from the National Key Research and Development Program of China [Grant 2022YFA1007900], the National Natural Science Foundation of China [Grant 12271013], Peking University's Fundamental Research Funds for the Central Universities, and the Massachusetts Institute of Technology Laboratory for Financial Engineering is acknowledged.
**Supplemental Material:** The online appendix and data files are available at https://doi.org/10.1287/mnsc.2022.01168.

**Keywords:** impact investing • environmental, social, and governance (ESG) investing • socially responsible investing • venture philanthropy • investments

## 1. Introduction

Impact investing—broadly defined as investments that consider not only financial objectives but also additional goals that support certain social priorities and agendas—has drawn an increasing amount of attention in recent years. This concept was first introduced through populist efforts to effect social change by encouraging institutional investors to divest from companies engaged in businesses viewed by critics as unethical, immoral, or otherwise objectionable (e.g., exploitation of child labor; tacit support of apartheid or religious persecution; or gambling, pornography, alcohol, tobacco, and firearms businesses (collectively known as "sin stocks")). However, its scope has expanded significantly to include investment products that use environmental, social, and governance (ESG) criteria, "green" bonds, and private equity funds seeking social impact alongside financial returns.[1]

The growth in popularity and assets under management of impact investing has also triggered a backlash. For example, on August 4, 2022, a letter signed by the attorneys general of 19 states was sent to BlackRock expressing concern over the asset manager's ESG policies and how they may impact their holdings of fossil-fuel energy companies:[2] "BlackRock's actions on a variety of governance objectives may violate multiple state laws. Mr. McCombe's letter asserts compliance with our fiduciary laws because BlackRock has a private motivation that differs from its public commitments and statements. This is likely insufficient to satisfy state laws requiring a sole focus on financial return." These are not minor concerns given that the legal penalty for violating one's fiduciary duty involves personally making up any losses suffered by the client and restoring to the client any profits made by the fiduciary's service provision to said client.[3]

So, how can we reconcile impact investing with fiduciary duty? The answer lies in developing a framework in which the financial impact of impact investing can be measured and disclosed, which is the subject of this article.

Conventional wisdom typically views impact investing as a standard portfolio selection problem with additional constraints related to the degree of social impact of the underlying securities, thereby implying a nonsuperior risk/reward profile compared with the unconstrained case. Given that the constrained portfolio contains a proper subset of securities of the unconstrained version, mathematical logic suggests that the constrained optimum is, at best, equal to the unconstrained optimum or more likely, inferior.

However, the nonsuperiority of constrained optima relies on a key assumption that is almost never explicitly stated; the constraint is assumed to be statistically independent of the securities' returns. In some cases, such an assumption is warranted; imagine constructing a subset of securities with Committee on Uniform Security Identification Procedures (CUSIP) identifiers that contain prime numbers. Clearly, such a constraint has no relation to the returns of any security; hence, imposing such a constraint can only reduce the risk-adjusted return of the optimized portfolio.

What if the constraint is *not* independent of the returns? For example, consider the constraint "invest only in those companies for which their stock prices will appreciate by more than 10% over the next 12 months." Apart from the infeasibility of imposing such a condition, it should be obvious that this constraint would, in fact, increase the risk-adjusted return of the optimized portfolio. Therefore, the answer to the question of what is the impact of impact investing rests entirely on whether and how the impact criteria are related to the performance characteristics of the securities being considered.

In this article, we develop a general framework to quantify the financial impact of impact investing. We formalize impact investing as the sorting and selection of an investment universe of $N$ securities based on an *impact factor*, $X_i$, for security $i$ so that higher values of $X_i$ correspond to greater impact (e.g., lower carbon emissions, greater sustainability, higher ESG score, etc.). As a result, other things equal, impact investors are assumed to prefer securities with higher values of $X_i$. This impact factor defines a rank ordering for all securities in the universe from which an impact portfolio can be constructed (i.e., the top decile of ESG-ranked securities or the bottom decile of carbon emissions-ranked securities). Therefore, the impact on investment performance is determined by the joint distribution of the vector $\mathbf{X} \equiv [X_1\ X_2\ \cdots\ X_N]^T$ of impact measures with the investment performance of individual securities.

To formalize this idea, we first propose a general linear multifactor model for asset returns and define excess returns or "alphas" as nonzero intercepts that we model as mean-zero random variables. This framework allows for the possibility of superior investment performance for individual securities but also includes the conventional case of equilibrium or no-arbitrage pricing if we set the variance of the alphas to zero. In fact, the implications from our model are broadly applicable to an equilibrium asset-pricing setup where "alpha" is reinterpreted as omitted factors of which investors are either unaware or unable to access as easily as professional portfolio managers. Such an agnostic approach to investment performance allows us to determine conditions under which impact investing does and does not change the risk/reward profile of a given investment product.

In particular, we derive—both in finite samples and asymptotically (as the number of securities increases without bound)—the distribution of individual alphas that have been ranked according to their impact factors $\mathbf{X}$. It is well known that ranked random variables—known as *order statistics*—have different distributions than their unranked versions. However, for our purposes, a more relevant strand of that literature focuses on *induced order statistics*, in which random variables are ranked not by their own values but by the values of other random variables (e.g., ranking the returns of a collection of mutual funds not by their returns but by the funds' market betas). We use properties of induced order statistics to derive the distribution of an impact portfolio's alphas ranked by an arbitrary impact score $\mathbf{X}$, allowing us to quantify the impact of impact investing.

Using this framework, we show that the expected alpha from the induced ordering is determined by three terms: the correlation between $\mathbf{X}$ and the individual securities' alphas, the standard deviation of individual securities' alphas, and a cross-sectionally standardized impact score that captures whether the impact factor of a security is above or below average. In addition, we provide an alternative characterization of the expected alpha from the induced ordering as a discounted version of the expected alpha from ordering securities based on alpha (i.e., via an all-knowing oracle that, in reality, is of course unattainable because alphas are unobservable). Much like the Sharpe–Lintner Capital Asset-Pricing Model (CAPM) (Sharpe 1964, Lintner 1965), which quantifies the expected return of individual securities through market beta,[4] this simple but profound result highlights the mechanism through which an impact factor's excess return is earned; it achieves a fraction of the maximum possible alpha with perfect knowledge, where the fraction is simply the correlation between $\mathbf{X}$ and the individual securities' alphas.

Using this insight, we quantify the alphas of portfolios formed based on the impact factor $\mathbf{X}$—including common heuristics of creating portfolios from the top or bottom impact-factor quantiles—and then apply the Treynor and Black (1973) framework to derive the optimal weights when forming both long/short and long-only portfolios to maximize Sharpe ratio. We show that such impact portfolios are associated with "superefficient frontiers" as long as the impact factor, $\mathbf{X}$, is positively correlated with the unobserved alphas of the individual securities. We also provide an equilibrium/no-arbitrage interpretation of our results in which excess returns arise from omitted factors that investors may not be aware of but to which impact portfolio managers have access. In this case, the excess returns are simply "excess" with respect to factors that investors observe and represent risk premia from specific impact factors.

The Treynor–Black portfolio allows us to construct a natural measure of the financial impact of impact investing; an impact factor has positive alpha when it is positively correlated with individual securities' unobserved alphas. On the other hand, an impact factor can impose a cost—also quantifiable in our framework—when it is negatively correlated with alphas and investors divest of the bottom-ranked securities (which have positive alphas on average because of the negative correlation with $\mathbf{X}$). This provides a possible explanation for the inconsistent and sometimes contradictory empirical findings on the effects of adopting impacting investing. The correlation between the impact factor and alpha is affected by different measures of impact,[5] different market conditions,[6] and different asset-pricing models for alpha,[7] all of which can influence the final estimate of the benefit or cost of impact investing.

To illustrate the practical relevance of our results, we apply our framework to five specific impact-investing contexts. The specific correlation—positive or negative—for each form of impact investing depends on the specific nature of the impact, the risks involved to achieve that impact, and its relationship with the underlying process of alpha generation.

The first example is biotech venture philanthropy (VP), a particular form of impact investing in biomedicine where nonprofit and mission-driven organizations fund initiatives to advance their objectives and potentially achieve returns that can be reinvested toward their mission. We take the case study by Kim and Lo (2019) about the Cystic Fibrosis (CF) Foundation to show that a significantly positive alpha can be achieved by advancing drug development for rare diseases, which illustrates the feasibility of "doing well by doing good." In this case, the challenges associated with early-stage drug development programs from the financial perspective—low probabilities of success, long time horizons, and large capital requirements as highlighted by Fagnan et al. (2013)—are more than offset by a positive correlation, $\rho$.

The second application involves the research and development (R&D) consortium, another form of impact investing. We consider the example of Semiconductor Manufacturing Technology ("Sematech")—a high-profile R&D consortium formed in 1987 and funded by the U.S. Government and 14 U.S. semiconductor firms—whose purpose was to solve common manufacturing problems by leveraging shared R&D resources. Using a simple difference-in-difference approach, we estimate that joining the consortium leads to an increased alpha in the stocks of member firms of around 10%, implying significant returns from investments in the consortium. In this case, the R&D consortium reduces R&D duplication and increases profitability for member firms (Irwin and Klenow 1996a), leading to a positive correlation between impact and returns.

The third application involves measuring the cost of divesting from sin stocks (stocks of companies engaged in businesses considered by some to be socially undesirable but that are documented to have positive alphas), which can be explained by the Merton (1987) model of neglected stocks and segmented markets. Calibrating to Hong and Kacperczyk (2009) and Blitz and Fabozzi (2017) as examples, we estimate the cost of divestment ranging from 0.6% to 3.3% in forgone alpha per annum. This example illustrates the dependence of the magnitude of estimated alpha on the specific asset-pricing model used, a well-known issue with all performance attribution exercises.

Fourth, we apply our framework to several ESG empirical studies. Correlations between the specific ESG measures in these studies and the unobserved alphas of individual securities determine the final estimate of the benefit (or cost) of ESG investing. They range from −0.05% for bonds (Baker et al. 2022) to 2.65% for equities in certain market conditions (Bansal et al. 2022). This underscores the importance of asset class, impact measures, and specific market conditions in determining the alpha of impact investing as well as the need to rationalize the highly dynamic impact of ESG on asset prices beyond equilibrium models of ESG investing.

Finally, we apply our framework to explain the January 2021 price spike in GameStop Corp. (GME) and other "meme" stocks, where a decentralized short squeeze that exploited the short positions of institutional investors caused their prices to increase sharply before crashing. Classifying such phenomena as impact investing may seem strange, but they share similarities with the other examples in how their returns are generated mechanically, and perhaps, a separate category called "price-impact investing" would be more appropriate. Our framework helps to quantify the financial impact of price-impact investing. Measuring the magnitude of such investments and understanding their financial implications can better inform regulators and policy makers as to the scope and severity of this phenomenon so that they can devote the appropriate sources to addressing it.

Overall, these five examples highlight the importance of choosing a baseline asset-pricing model and forming impact variables based on economic, institutional, and market rationales in order to establish sound and robust relationships between impact and returns. Our model offers a unified framework that can be calibrated to a wide range of settings and provides academic rigor for how to think about impact investing and quantify their financial consequences. Moreover, our framework provides a systematic and politically neutral approach for portfolio managers to disclose the financial consequences of impact investing, thereby addressing any concerns that their clients are either unaware of or unwilling to bear the consequences of their managers' impact criterion.

## 2. Literature Review

There is a growing literature theorizing the impact of SRI, ESG, and other nonfinancial objectives on asset pricing. The Fama and French (2007) taste model shows that if investors prefer to invest in socially responsible companies, the expected return on such companies will be lower. Pástor et al. (2021) provide a model for ESG investing where investors' taste for green assets implies lower returns, and assets can be priced in a two-factor model that includes the ESG factor and the market portfolio. Pedersen et al. (2021) derive an ESG-efficient frontier and show that ESG may either yield benefits to expected returns because it provides information about firm fundamentals (as in our example, in which constraints contain information about returns) or incur costs because it affects investor preferences and constraints.[8]

Although these studies share some of the same implications as our framework, we add to this literature in several novel ways. The equilibrium frameworks of Fama and French (2007), Pástor et al. (2021), and Pedersen et al. (2021) highlight that the expected return of ESG investing depends on the mix of investors and preferences in the market. However, impact investing is still an evolving concept, and their expected returns are dynamic and context dependent. It is possible that market prices are still adjusting to reach a new equilibrium that reflects these considerations (Cornell and Damodaran 2020).[9] Our unified econometric framework provides an explicit method to quantify the excess returns of any form of impact investment—including, but not limited to, the equilibrium setting—during different stages of this adaptive process. These results are, in turn, consistent with the equilibrium-based models when the correlation between $\mathbf{X}$ and security returns reflects the particular market condition and shift in preferences over time. From the adaptive markets (Lo 2004, 2017) perspective, this correlation could reinforce itself as the amount of assets under management for a given impact factor increases over time and eventually stabilizes as the size of the new sector reaches a steady state.

Our framework also differs from existing models in that we allow for the possibility of nonzero alphas or omitted factors in the equilibrium/no-arbitrage interpretation, which is particularly relevant for the highly adaptive and dynamic ESG investment industry. An important insight from Pedersen et al. (2021) is that ESG's information about firm fundamentals can yield benefits to its returns, whereas screening constraints will incur costs to ESG investing.[10] Our model shows that when securities have nonzero alphas that are otherwise inaccessible to investors, ESG investing can derive financial benefit from constraints too because of the information about returns implicit in these constraints. This effect is formalized statistically by the correlation between the impact factor, $\mathbf{X}$, and returns. As a result, we are able to explicitly construct the optimal superefficient portfolio from any $\mathbf{X}$ and explicitly quantify its financial impact.

The empirical literatures on measuring the returns of SRI and ESG are not always consistent with each other. These inconsistencies raise the question of what the real financial impact of impact investing is, which is precisely the motivation for our current contribution. Our framework explains not only how to measure the financial impact of impact investing but also why there is such a wide range of empirical estimates for the expected returns of SRI and ESG investing. It stems from the wide range of impact definitions, date ranges, asset classes, and asset-pricing models for alpha, each of which leads to a different specification that may have a potentially different correlation between the impact factor and asset returns.

More generally, our framework is applicable to portfolios constructed on the basis of any characteristic, including traditional factors such as value, size, momentum, and other variables. As such, our work is related to several strands of the asset-pricing and econometrics literature. These include a large literature devoted to identifying asset-pricing factors[11] and a vast econometrics literature focused on factor models.[12] In particular, we make use of the same statistical results on induced order statistics first applied to financial data by Lo and MacKinlay (1990), albeit in a very different context.

## 3. The Framework

We consider a universe of $N$ securities with returns $R_{it}$ that satisfy the following multifactor model:

$$R_{it} - R_{ft} = \alpha_i + \beta_{i1}(\Lambda_{1t} - R_{ft}) + \cdots + \beta_{iK}(\Lambda_{Kt} - R_{ft}) + \epsilon_{it} \quad (1)$$

$$\text{such that } \mathrm{E}[\epsilon_{it} \mid \Lambda_{kt}] = 0, \; k = 1, \ldots, K, \quad (2)$$

where $\Lambda_{kt}$ is the $k$th factor return; $k = 1, \ldots, K$. Here $R_{ft}$ is the risk-free rate, the excess return and factor betas are represented by $\alpha_i$ and $\beta_{ik}$, respectively, and $\epsilon_{it}$ is the idiosyncratic return. Because we consider only a static model in this article, we omit the subscript $t$ throughout for notational simplicity.

Under suitable restrictions on the parameters $\{\alpha_i, \beta_{ik}\}$ and the definitions of the factor returns $\{\Lambda_k\}$, the linear multifactor model (1) is consistent with a number of asset-pricing models, such as the Sharpe–Lintner CAPM (Sharpe 1964, Lintner 1965), Merton's intertemporal

CAPM (Merton 1973), Ross's arbitrage pricing theory (APT) (Ross 1976), and the Fama–French multifactor models (Fama and French 1993, 2015). In particular, all of these asset-pricing models imply that $\alpha_i = 0$.

However, to measure the impact of impact investing, we take no position as to whether any particular asset-pricing model holds. Nor do we make any assumption on investor belief structures. Instead, we derive the implications of impact investing on the statistical properties of impact-portfolio returns without constraining excess returns to be zero. These properties lead to a framework that is flexible enough to interpret impact from multiple perspectives.

### 3.1. The No-Impact Baseline Case

We begin by stating the near-trivial result that arbitrary portfolios formed according to criteria unrelated to the parameters of the return-generating processes $\{R_{it}\}$ are necessarily less than or equal to the investment performance of the mean-variance optimal portfolio.[13]

**Proposition 1.** *If asset returns satisfy (1) and (2) and $\alpha_1 = \cdots = \alpha_N = 0$, then any arbitrary subset $\mathcal{S} \subseteq \{1, \ldots, N\}$ formed independently of the joint distribution of returns, $\{R_{it}\}$, satisfies*

$$\begin{aligned} &\max_{\{\omega_1, \ldots, \omega_N | \sum_{i=1}^N \omega_i = 1\}} \mathrm{E}[U(W)] \\ &\geq \max_{\{\omega_1^c, \ldots, \omega_N^c | \sum_{i \in \mathcal{S}} \omega_i^c = 1 \text{ and } \omega_i^c = 0 \text{ for } i \notin \mathcal{S}\}} \mathrm{E}[U(W^c)] \end{aligned} \tag{3}$$

*for any nondecreasing concave utility function $U(\cdot)$, where*

$$W \equiv \sum_{i=1}^N \omega_i R_i \quad \text{and} \quad W^c \equiv \sum_{i \in \mathcal{S}} \omega_i^c R_i. \tag{4}$$

*In addition, under certain fairly realistic conditions given in the online appendix, the loss in utility by restricting to the subset $\mathcal{S}$ is generally small as long as the number of securities excluded by $\mathcal{S}$ is small relative to the total number of securities, $N$.*

This proposition confirms the common critique that skeptics often level against impact investing. If the constraint $\mathcal{S}$ has nothing to do with the characteristics of the underlying asset returns, $\{R_i\}$, then imposing such constraints can only reduce investment performance or at best, achieve the unconstrained optimum. In particular, the independence of $\mathcal{S}$ and $\{R_i\}$ implies that the excess returns, $\{\alpha_i\}$, are indistinguishable from $\{\epsilon_i\}$, in which case we are essentially assuming zero excess returns, so there is no possibility of generating any excess performance.[14] In addition, although impact investing in this special case cannot improve returns, the underperformance is likely to be small (assuming no transactions costs or fees, of course).

However, suppose we allow for nonzero alphas that are unobserved to investors; in other words, the unconstrained optimization problem in (3) does not have the ability to find securities with positive alphas.[15] If we relax the condition that $\mathcal{S}$ is independent of the joint distribution of $\{R_i\}$, then Proposition 1 clearly does not hold. For example, suppose that

$$\mathcal{S} = \{i : \alpha_i > 0, \ i = 1, \ldots, N\}. \tag{5}$$

Clearly, in this case, it is possible for the risk-adjusted returns of the $\mathcal{S}$ portfolio to beat those of the unconstrained portfolio given that the subset contains all positive-alpha securities and the complement contains the reverse. This conclusion may seem counterintuitive because the constrained portfolio is, by definition, a feasible solution in the unconstrained case. So, how can imposing the constraint ever improve performance? The answer lies in the fact that in the unconstrained case, information about the $\{\alpha_i\}$ is not available; the constraint contains private information[16] that can dramatically improve performance. Therefore, the constrained solution is actually not feasible in the unconstrained case.

So, the fundamental question of whether an impact investment has positive (or negative) financial impact reduces to the information content in the constraint (i.e., the relation between the constraint and the joint distribution of asset returns). No relation implies no information; hence, there is no impact. However, the presence of even the slightest amount of dependence between the constraint and returns implies the possibility of some degree of impact. We can quantify this degree by being explicit about the statistical relation between asset returns and the impact factor.

Of course, this counterexample assumes the existence of mispriced or positive-alpha securities (5), but an equally valid equilibrium/no-arbitrage interpretation is that the $\alpha_i$'s are omitted factors from the investor's linear-factor benchmark. Either investors are unaware of these factors, or they do not have the ability to access them (e.g., exotic betas from private equity, distressed debt, event-driven opportunities, etc.). Under this interpretation, impact investing can be viewed as providing investors with alternative risk premia.

Our framework accommodates both interpretations—as we describe—and offers a systematic and quantitative approach to measuring impact in either case.

### 3.2. Impact Factors and Induced Order Statistics

We assume that the excess return of the *i*th security, $\alpha_i$, is not observable, whereas the impact factor, $X_i$, for that security is. Contrary to the usual asset-pricing setup in which the $\alpha_i$'s are assumed to be fixed constants (and in equilibrium or under no-arbitrage conditions, identically equal to zero), we assume that they are random variables.

Impact investors select a portfolio based on the impact factor, $\mathbf{X}$, and the excess return of their portfolio is determined by the corresponding vector of excess returns of the individual securities in that portfolio,

$\boldsymbol{\alpha} \equiv [\alpha_1 \cdots \alpha_N]^T$. Specifically, suppose an investor ranks $N$ securities according to $\mathbf{X}$. Let us reorder the bivariate vector $(X_i, \alpha_i)^T, i = 1, 2, \ldots, N$, according to the magnitudes of their first components:

$$\begin{pmatrix} X_{1:N} \\ \alpha_{[1:N]} \end{pmatrix}, \begin{pmatrix} X_{2:N} \\ \alpha_{[2:N]} \end{pmatrix}, \cdots, \begin{pmatrix} X_{N:N} \\ \alpha_{[N:N]} \end{pmatrix}, \tag{6}$$

where $X_{1:N} < X_{2:N} < X_{N:N}$ and the notation $X_{i:N}$ denotes the $i$th order statistic from a total of $N$ random variables. The notation $\alpha_{[i:N]}$ represents the $i$th *induced order statistic*,[17] where the order is *induced* by another variable $\mathbf{X}$.

Because $\boldsymbol{\alpha}$ is defined with respect to the multifactor model (1), by studying the interaction between $\mathbf{X}$ and $\boldsymbol{\alpha}$, our framework allows for the existence of any set of predefined asset-pricing factors and focuses on the incremental role of $\mathbf{X}$ in determining asset returns.

### 3.3. Defining an Impact Portfolio

Impact investing essentially corresponds to the selection of securities based on the impact factor, $\mathbf{X}$. For example, an investor may choose to invest in the top $n_0$ securities ranked by $\mathbf{X}$ or form portfolios long the top decile and short the bottom decile. In general, we define an *impact portfolio* as any portfolio $\mathcal{S}(\mathbf{X})$ formed as a function of the impact factor, $\mathbf{X}$. With portfolio weights $\{\omega_i, i \in \mathcal{S}\}$, the return of the impact portfolio is given by

$$R_{\mathcal{S}} = \sum_{i \in \mathcal{S}} \omega_i R_i. \tag{7}$$

To characterize $R_{\mathcal{S}}$, we therefore need to quantify the distribution of the excess returns—or the induced order statistic $\alpha_{[i:N]}$—given certain assumptions on the joint distribution of $(\mathbf{X}, \boldsymbol{\alpha})$.

Note that $\mathbf{X}$ can represent a variety of characteristics related to metrics for climate change, sustainable farming, tobacco usage, gambling, biomedical R&D, and any other SRI or ESG considerations. Together with the generality of our multifactor asset-pricing model (1), this corresponds to a wide range of impact-investing contexts. Section 6 provides five concrete examples, highlighting that the specific economic, institutional, and market variables that matter in each case will depend on the specific context and time period in consideration (see discussions in Section 7).

In fact, our framework applies more generally to any characteristics of a security including, for example, the traditional value, size, and momentum factors as well as denizens of the "factor zoo" described in the recent literature. For the purposes of this study, we focus on the impact-investing interpretation, but we will discuss broader interpretations in Section 8.

## 4. Characterizing Excess Returns

To assess the impact of impact portfolios, we require the distribution of $\alpha_{[i:N]}$, which can be derived explicitly under the following assumption.

**Assumption 1.** *$(X_i, \alpha_i)^T, i = 1, 2, \ldots, N$, are independently and identically distributed bivariate normal random vectors with mean $(\mu_x, \mu_\alpha)^T$, variance $(\sigma_x^2, \sigma_\alpha^2)^T$, and correlation $\rho \in (-1, 1)$.*

The assumption that $\alpha_i$ is random is somewhat unconventional, so a few clarifying remarks are in order. This assumption was first used in Lo and MacKinlay (1990) to represent cross-sectional estimation errors of intercepts from CAPM regressions. However, in our current context, we interpret the randomness in $\alpha_i$ as a measure of uncertainty as to the degree of mispricings of securities in our investment universe.[18] This uncertainty can be interpreted from a Bayesian perspective as the degree of conviction that mispricings exist in the cross-section. Under this interpretation, we will make the auxiliary assumption—without loss of much generality—that all $\alpha_i$'s are mean 0 ($\mu_\alpha = 0$). This corresponds to centering the Bayesian prior on zero average deviations from equilibrium or no-arbitrage pricing in our investment universe, a reasonable and more realistic first approximation that still allows for mispricings, which of course, motivates a significant portion of the asset management industry's products and services.[19] Moreover, we can calibrate the degree of mispricings in our model through $\sigma_\alpha^2$; smaller values correspond to greater efficiency, and larger values correspond to lower efficiency and more active management opportunities.

However, our framework can also be interpreted from an equilibrium/no-arbitrage perspective, where nonzero $\alpha_i$'s are because of the presence of omitted factors that investors are either unaware of or unable to access directly. Under this interpretation, we will see that the randomness in $\alpha_i$ is because of cross-sectional variability in security $i$'s omitted-factor betas. In this case, however, it is possible for $\mu_\alpha$ to be nonzero to reflect the risk premia of the omitted factors.

We deliberately make the assumption of joint normality for $(X_i, \alpha_i)^T$ so as to capture the interaction between $\mathbf{X}$ and $\boldsymbol{\alpha}$ while still being able to derive explicit and easily interpretable results for induced excess returns. This assumption can be relaxed considerably to allow for cross-sectional dependence (see Section 4.3) and general marginal distributions and dependent structures for $(X_i, \alpha_i)^T$ (see Lo et al. 2022b) at the expense of simplicity.

Regardless of the interpretation of $\alpha_i$, the theory of induced order statistics allows us to completely characterize its statistical properties. We first present its finite-sample distribution followed by asymptotic results when the number of securities, $N$, increases without bound.

### 4.1. Finite-Sample Distribution

We first observe that the mean and standard deviation of the impact factor, $\mathbf{X}$, do not actually matter for the distribution of $\alpha_{[i:N]}$'s because it is only the relative order of

$X_i$'s that determines the order of $\alpha_{[i:N]}$'s. Therefore, we assume without loss of generality that $\mu_x = 0$ and $\sigma_x = 1$ so that $\mathbf{X}$ is a standard normal random vector. Then, the following result characterizes the finite-sample distributions of the induced order statistics $\{\alpha_{[i:N]}\}$.

**Proposition 2.** *Under Assumption 1, the expected value of the ith induced order statistic $\alpha_{[i:N]}, i = 1, 2, \ldots, N$ is given by*

$$\mu_i \equiv \mathrm{E}[\alpha_{[i:N]}] = \rho\sigma_\alpha \mathrm{E}[X_{i:N}]. \tag{8}$$

*The variance of the ith induced order statistic $\alpha_{[i:N]}, i = 1, 2, \ldots, N$ is given by*

$$\sigma_i^2 \equiv \mathrm{Var}(\alpha_{[i:N]}) = \sigma_\alpha^2(1 - \rho^2 + \rho^2 \mathrm{Var}(X_{i:N})). \tag{9}$$

*The covariance of the ith and jth induced order statistics, $\alpha_{[i:N]}$ and $\alpha_{[j:N]}$, for $i \neq j$ is given by*

$$\sigma_{ij} \equiv \mathrm{Cov}(\alpha_{[i:N]}, \alpha_{[j:N]}) = \sigma_\alpha^2\rho^2 \mathrm{Cov}(X_{i:N}, X_{j:N}). \tag{10}$$

Proposition 2 gives us the first two moments of the induced order statistics, $\alpha_{[i:N]}$'s. In particular, the expected alpha in (8) is determined by three terms: the correlation ($\rho$) between $\mathbf{X}$ and individual securities' alphas, the standard deviation of individual alphas ($\sigma_\alpha$), and a cross-sectionally standardized impact score ($\mathrm{E}[X_{i:N}]$). The correlation $\rho$ here plays a critical role in determining the expected alpha of both individual securities and impact portfolios (see Section 5). This resembles the CAPM's market beta—which quantifies security returns attributable to systematic market risk—because the market beta is simply the correlation between security returns and market returns when they are both standardized with unit variances.

If we view the impact factor $\mathbf{X}$ as a signal for predicting asset returns, the expected alpha in (8) is closely related to the results by Grinold (1994),[20] who provides a simple decomposition of alpha into the product of three terms: the information coefficient (the correlation $\rho$ in our notation), the volatility of residual returns, and a standardized score that measures the strength of the signal for each asset. In our context, it is the volatility of the unobserved alpha (not the volatility of residual returns) that determines the expected alpha of each asset. In addition, because we only use the rank information in $\mathbf{X}$, the standardized score can be quantified explicitly by $\mathrm{E}[X_{i:N}]$. Finally, our results also provide the variance and covariances of individual alphas, which are crucial for quantifying the uncertainty of these alphas in practice.

We note that all three quantities in (8)–(10) depend on the distribution of the order statistics of standard normal random variables. In fact, the terms $\mathrm{E}[X_{i:N}]$ in (8), $\mathrm{Var}(X_{i:N})$ in (9), and $\mathrm{Cov}(X_{i:N}, X_{j:N})$ in (10) can be explicitly evaluated by numerical integration over the density function of $X_{i:N}$ (see David and Nagaraja 2004, section 3.1 for example).

On the other hand, we can also explicitly evaluate the quantities in (8)–(10) based on the following approximation results.

**Proposition 3.** *Let $p_i \equiv \frac{i}{N+1}$ denote the relative position of the order i in the population of N securities. The expected value and variance of the ith order statistic of the standard normal random variable, $X_{i:N}$, can be approximated up to order $(N+2)^{-2}$, when N increases without bound, by*

$$\mathrm{E}[X_{i:N}] \approx \Phi^{-1}(p_i) + \frac{p_i(1-p_i)}{2(N+2)}Q_i'' + \frac{p_i(1-p_i)}{(N+2)^2}\left[\frac{1}{3}(1-2p_i)Q_i''' + \frac{1}{8}p_i(1-p_i)Q_i''''\right] \tag{11}$$

$$\mathrm{Var}(X_{i:N}) \approx \frac{p_i(1-p_i)}{N+2}Q_i'^2 + \frac{p_i(1-p_i)}{(N+2)^2}\left[2(1-2p_i)Q_i'Q_i'' + p_i(1-p_i)\left(Q_i'Q_i''' + \frac{1}{2}Q_i''^2\right)\right] \tag{12}$$

*for $i = 1, 2, \ldots N$. Additionally, their covariances can be approximated up to order $(N+2)^{-2}$, when N increases without bound, by*

$$\mathrm{Cov}(X_{i:N}, X_{j:N}) \approx \frac{p_i(1-p_j)}{N+2}Q_i'Q_j' + \frac{p_i(1-p_j)}{(N+2)^2}\left[(1-2p_i)Q_i''Q_j' + (1-2p_j)Q_i'Q_j'' + \frac{1}{2}p_i(1-p_i)Q_i'''Q_j' + \frac{1}{2}p_j(1-p_j)Q_i'Q_j''' + \frac{1}{2}p_i(1-p_j)Q_i''Q_j''\right] \tag{13}$$

*for $1 \le i < j \le N$. Here, $Q_i'$, $Q_i''$, $Q_i'''$, and $Q_i''''$ are the first four derivatives of $\Phi^{-1}(p_i)$:*

$$Q_i' = \left(\Phi^{-1}(p_i)\right)' = \frac{1}{\phi\left(\Phi^{-1}(p_i)\right)} \tag{14}$$

$$Q_i'' = \left(\Phi^{-1}(p_i)\right)'' = \frac{\Phi^{-1}(p_i)}{\phi\left(\Phi^{-1}(p_i)\right)^2} \tag{15}$$

$$Q_i''' = \left(\Phi^{-1}(p_i)\right)''' = \frac{1 + 2\left(\Phi^{-1}(p_i)\right)^2}{\phi\left(\Phi^{-1}(p_i)\right)^3} \tag{16}$$

$$Q_i'''' = \left(\Phi^{-1}(p_i)\right)'''' = \frac{\Phi^{-1}(p_i)(7 + 6\left(\Phi^{-1}(p_i)\right)^2)}{\phi\left(\Phi^{-1}(p_i)\right)^4}. \tag{17}$$

*$\Phi$ and $\phi$ are the cumulative distribution function (CDF) and density function of the standard normal distribution, respectively.*

Although the approximations in Proposition 3 may seem daunting, their first-order terms are fairly intuitive.

The first term in (11) is $\Phi^{-1}(p_i)$, which simply approximates $\mathrm{E}[X_{i:N}]$ by the inverse CDF applied to the relative rank, $p_i \equiv \frac{i}{N+1}$, of the $i$th order statistic, which is a well-known first-order approximation by itself.

Figure 1 displays the mean, variance, and covariances of the induced order statistic, $\alpha_{[i:N]}$, for a collection of $N = 50$ securities as given in Proposition 2 using the approximations in Proposition 3. When the correlation, $\rho$, between $\boldsymbol{\alpha}$ and $\mathbf{X}$ is positive, the expected value of the induced order statistic increases as the order $i$ increases (see Figure 1(a)). The dispersion of the mean is larger when the correlation, $\rho$, or the dispersion of the unknown $\boldsymbol{\alpha}$, $\sigma_\alpha$, is larger.

In addition, Figure 1(b) shows that the variances, $\mathrm{Var}(X_{i:N})$, stay relatively constant across the ordered securities $i$ and are primarily determined by $\rho$ and $\sigma_\alpha$. In fact, we will see in Online Appendix A.1 that as the number of securities increases without bound, the variance converges to a constant across all $i$.

Finally, the covariances, $\mathrm{Cov}(X_{i:N}, X_{j:N})$, are very close to zero except when $i$ and $j$ are close to the two extremes. We show in Online Appendix A.1 that as $N$ increases without bound, the covariances approach zero, implying that induced order statistics are mutually independent in the limit.

Online Appendix A.1 provides the asymptotic mean, variance, and covariances of induced order statistics that generalize Proposition 2.

### 4.2. Comparison with Conventional Order Statistics

To develop further intuition for the effect of induced ordering, we compare the distributions of induced order statistics with their conventional order statistics counterparts, $\alpha_{i:N}$. Note that this comparison is merely meant to be an illustrative thought experiment; $\boldsymbol{\alpha}$ is unobservable in practice. Hence, such rankings are not feasible in practice. Nonetheless, this provides a useful comparison

**Figure 1.** (Color online) Mean, Variance, and Covariances of the Induced Order Statistic, $\alpha_{[i:N]}$

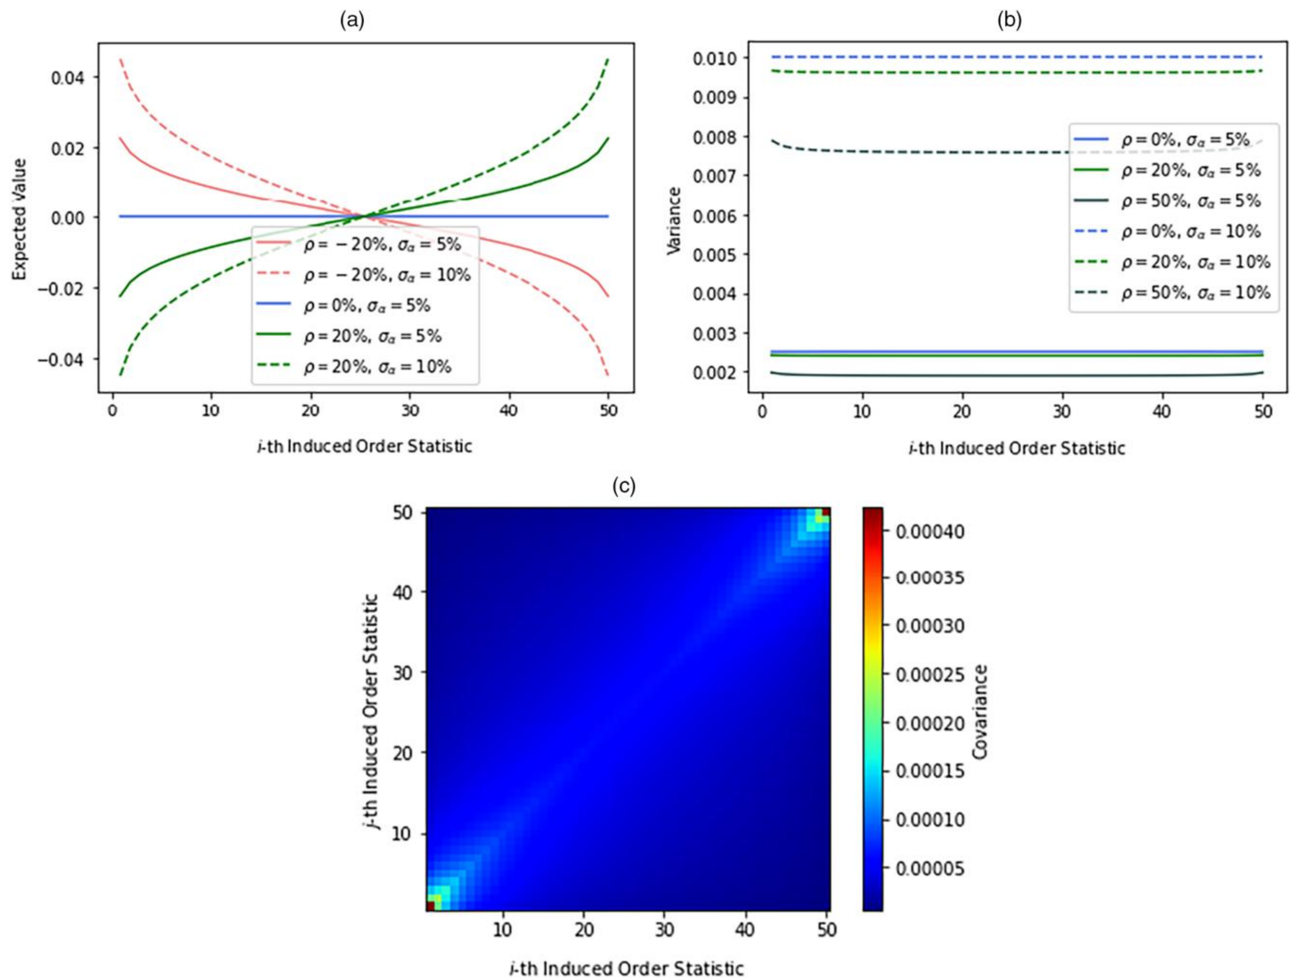


*Notes.* The total number of securities is set to 50 for illustrative purposes. In panel (c), we set $\rho = 20\%$ and $\sigma_\alpha = 5\%$. (a) Expected value. (b) Variance. (c) Covariances.

with what can be achieved by ordering based on the impact factor, $\mathbf{X}$.

**Proposition 4.** *Under Assumption 1, the first two moments of the induced order statistic, $\alpha_{[i:N]}$, are related to the order statistic, $\alpha_{i:N}$, by the following identities:*

$$\mu_i \equiv \mathrm{E}[\alpha_{[i:N]}] = \rho \mathrm{E}[\alpha_{i:N}] \tag{18}$$

$$\sigma_i^2 - \sigma_\alpha^2 \equiv \mathrm{Var}(\alpha_{[i:N]}) - \sigma_\alpha^2 = \rho^2[\mathrm{Var}(\alpha_{i:N}) - \sigma_\alpha^2] \tag{19}$$

$$\sigma_{ij} \equiv \mathrm{Cov}(\alpha_{[i:N]}, \alpha_{[j:N]}) = \rho^2 \mathrm{Cov}(\alpha_{i:N}, \alpha_{j:N}). \tag{20}$$

Proposition 4 tells us that the mean, variance, and covariances of the induced order statistics, $\alpha_{[i:N]}$, are essentially a discounted version of the corresponding moments of the conventional order statistics, $\alpha_{i:N}$. The discount factor, $\rho$, is precisely the correlation between $\mathbf{X}$ and $\boldsymbol{\alpha}$.

To visualize this effect, Figure 2 contains a comparison of the expected excess returns of the induced order statistic, $\alpha_{[i:N]}$, and the order statistic, $\alpha_{i:N}$. As the correlation, $\rho$, increases to one, the expected excess return approaches the hypothetical value of sorting based on $\boldsymbol{\alpha}$.

This result highlights the role that induced ordering plays in distinguishing securities with positive alpha from those with negative alpha. If alpha is fully observed by a hypothetical oracle, she can simply pick securities with the highest alphas to construct impact portfolios. In reality, the correlation between the sorting variable (in our case, the impact factor) and the target variable (in our case, the unobserved $\boldsymbol{\alpha}$) determines how much of the mean, variance, and covariances from a hypothetical sorting based on $\boldsymbol{\alpha}$ can actually be achieved via the induced ordering of $\mathbf{X}$.

### 4.3. Interpreting Excess Return as Omitted Factors

Having completely characterized the stochastic properties of the excess returns $\boldsymbol{\alpha}$ of securities ranked according to an arbitrary impact factor $\mathbf{X}$, we now provide an explicit derivation of the equilibrium/no-arbitrage interpretation of $\boldsymbol{\alpha}$ as risk premia associated with omitted factors.

**Figure 2.** (Color online) Comparison Between the Expected Value of the Induced Order Statistic, $\alpha_{[i:N]}$, and the Order Statistic, $\alpha_{i:N}$, with $N = 50$ and $\sigma_\alpha = 5\%$

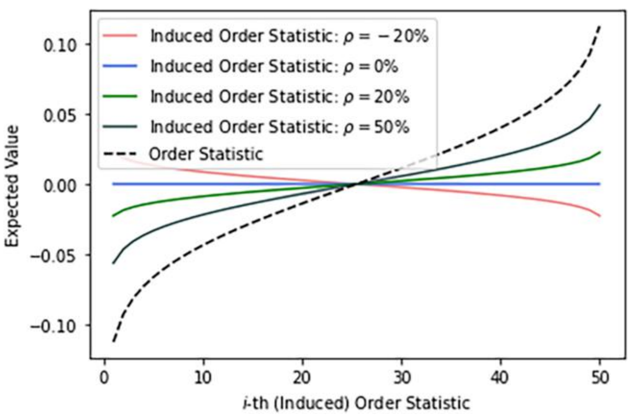


Let security returns follow the $K$-factor asset-pricing model as specified in (1) and (2), but now assume there are no mispricings. However, suppose that investors only account for the first factor $\Lambda_1$, without loss of generality, and are unaware of the remaining $K-1$ factors $\Lambda_2, \ldots, \Lambda_K$. We define $\lambda_{ik} \equiv \beta_{ik}(\Lambda_k - R_f)$ to be factor $k$'s contribution to security $i$'s return, for $i = 1, \ldots, N$ and $k = 2, \ldots, K$, and $\lambda_i \equiv \sum_{k=2}^{K} \lambda_{ik}$ to be the total net contribution of all the omitted factors to security $i$'s return. Given that investors are unaware of factors $2, \ldots, K$, the total excess expected returns for the securities in our universe appear to be alphas to such investors:

$$\alpha_i \equiv \mathrm{E}[\lambda_i] = \sum_{k=2}^{K} \beta_{ik}(\mathrm{E}[\Lambda_k] - R_f). \tag{21}$$

To characterize the distribution of $\lambda_i$ after ranking securities based on the impact factor $\mathbf{X}$, we make the following assumption.

**Assumption 2.** $(X_i, \lambda_i)^T, i = 1, 2, \ldots, N$, *are bivariate normal random vectors with their marginal distributions and paired correlations defined by*

$$\mu_x \equiv \mathrm{E}[X_i],\ \mu_\lambda \equiv \mathrm{E}[\lambda_i],\ \sigma_x^2 \equiv \mathrm{Var}(X_i),\ \sigma_\lambda^2 \equiv \mathrm{Var}(\lambda_i),\ \textit{and}\ \rho_{x,\lambda} \equiv \mathrm{Corr}(X_i, \lambda_i). \tag{22}$$

*In addition, for $i \neq j$, the correlations across different securities are defined by*

$$\rho_x \equiv \mathrm{Corr}(X_i, X_j),\ \rho_\lambda \equiv \mathrm{Corr}(\lambda_i, \lambda_j),\ \textit{and}\ \tilde{\rho}_{x,\lambda} \equiv \mathrm{Corr}(X_i, \lambda_j). \tag{23}$$

Under this assumption, the cross-sectional randomness of $\lambda_i$ can be interpreted as variations coming from both the factor values and the distribution of factor betas across companies in our universe. $(X_i, \lambda_i)^T$ can be correlated across securities, and their correlation structure is described by the four parameters $\rho_{x,\lambda}$, $\tilde{\rho}_{x,\lambda}$, $\rho_x$, and $\rho_\lambda$.

We can characterize the first two moments of $\lambda_{[i:N]}$. Recall that the notation $\lambda_{[i:N]}$ denotes the $i$th induced order statistic where the order is *induced* by the impact factor $\mathbf{X}$. We again assume without loss of generality that $\mu_x = 0$ and $\sigma_x = 1$ so that $\mathbf{X}$ is a standard normal random vector. However, we allow for a nonzero risk premium $\mu_\lambda$.

**Proposition 5.** *Under Assumption 2, define*

$$\rho_{adj} \equiv \frac{\rho_{x,\lambda} - \tilde{\rho}_{x,\lambda}}{1 - \rho_x} \tag{24}$$

*to be an adjusted correlation. The expected value of the ith induced order statistic $\lambda_{[i:N]}, i = 1, 2, \ldots, N$ is given by*

$$\mathrm{E}[\lambda_{[i:N]}] - \mu_\lambda = \rho_{adj}\sigma_\lambda \mathrm{E}[X_{i:N}]. \tag{25}$$

*The variance of the ith induced order statistic $\lambda_{[i:N]}, i = 1, 2, \ldots, N$ is given by*

$$\mathrm{Var}(\lambda_{[i:N]}) = \sigma_\lambda^2(1 - \rho_{adj}^2 + \rho_{adj}^2 \mathrm{Var}(X_{i:N})). \tag{26}$$

*The covariance of the ith and jth induced order statistics, $\lambda_{[i:N]}$ and $\lambda_{[j:N]}$, for $i \neq j$ is given by*

$$\mathrm{Cov}(\lambda_{[i:N]}, \lambda_{[j:N]}) = \sigma_\lambda^2 \rho_{adj}^2 \mathrm{Cov}(X_{i:N}, X_{j:N}) + (\rho_\lambda - \rho_x \rho_{adj}^2). \tag{27}$$

Proposition 5 characterizes the return from omitted factors for the *i*th security ranked by **X**. This result highlights an important implication when estimating the financial impact of impact investing. Given any definition of impact, **X**, if the portfolio selected based on **X** produces a nonzero excess return, **X** must be correlated with some factors not previously accounted for in the asset-pricing framework. This may imply the existence of a new factor that corresponds to the very definition of **X**, such as an "ESG factor" or a "carbon factor" (Bolton and Kacperczyk 2021).

On the other hand, Proposition 5 also implies that when forming a portfolio, if one uses selection criteria that appear independent of return characteristics such as market betas and factor loadings, it may still be correlated with omitted factor risk premiums, in which case the selection criteria will produce nonzero excess returns. In other words, what appears to be an "impact factor" (a selection criteria **X** based on a particular concept) may just be correlations with other omitted factors that are, in fact, unrelated to the impact concept one intends to capture. Therefore, impact estimates may be inaccurate and misleading without first properly accounting for all known factors.

This observation is supported empirically by both Blitz and Fabozzi (2017) in the case of estimating excess returns for sin stocks and Madhavan et al. (2021) for ESG scores, both of which we discuss in more detail in Section 6.

## 5. Impact-Portfolio Construction

Having quantified the distribution of the induced order statistics, $\alpha_{[i:N]}$, we can now construct portfolios based on the impact factor, **X**, and characterize the statistical properties of their excess returns. We first quantify the performance of arbitrary impact portfolios. We then use the Treynor and Black (1973) framework to derive the optimal weights for each security as well as the optimal way to combine an impact portfolio with any existing portfolio, such as the passive market index. The latter result follows directly from our ability to completely characterize the statistical properties of individual alphas in our framework.

### 5.1. Properties of Arbitrary Impact Portfolios

Consider an arbitrary impact portfolio of $n_0$ securities with indexes in $\mathcal{S}$:

$$\mathcal{S} \equiv \{i_1, i_2, \ldots, i_{n_0}\}, \tag{28}$$

which is obtained from a rank ordering of securities from the investment universe according to the impact factor, **X**. The excess return of the portfolio is then given by

$$\tilde{\alpha} \equiv \sum_{i \in \mathcal{S}} \omega_i \alpha_{[i:N]}, \tag{29}$$

where $\{\omega_i : i \in \mathcal{S}\}$ are arbitrary portfolio weights that sum to one. Based on the distribution of the induced order statistics in Proposition 2, we have the following result for portfolio excess returns.

**Proposition 6.** *Under Assumption 1, the expected excess return of a portfolio $\mathcal{S}$ defined in (28) is*

$$\mathrm{E}[\tilde{\alpha}] = \sum_{i \in \mathcal{S}} \omega_i \mu_i = \rho \sigma_\alpha \sum_{i \in \mathcal{S}} \omega_i \mathrm{E}[X_{i:N}], \tag{30}$$

*and the variance is*

$$\begin{aligned} \mathrm{Var}(\tilde{\alpha}) &= \sum_{i \in \mathcal{S}} \omega_i^2 \sigma_i^2 + 2 \sum_{i < j \in \mathcal{S}} \omega_i \omega_j \sigma_{ij} \\ &= \sigma_\alpha^2 \left( 1 - \rho^2 + \rho^2 \left( \sum_{i \in \mathcal{S}} \omega_i^2 \mathrm{Var}(X_{i:N}) \right.\right. \\ &\quad \left.\left. + 2 \sum_{i < j \in \mathcal{S}} \omega_i \omega_j \mathrm{Cov}(X_{i:N}, X_{j:N}) \right) \right). \end{aligned} \tag{31}$$

Proposition 6 quantifies the distribution of excess returns for any portfolio constructed according to the impact factor, **X**. Online Appendix A.3 provides several numerical examples where impact portfolios are formed based on top-ranking securities and decile portfolios. This result implies that the full range of tools and results from modern portfolio theory can be applied here, including the calculation of various performance metrics, such as the Sharpe ratio (Sharpe 1966), the Sortino ratio (Sortino and Van Der Meer 1991, Sortino and Price 1994), and information ratios (Treynor and Black 1973); performance attribution (Brinson et al. 1986); and active portfolio management and enterprise risk management (Grinold and Kahn 1999).

### 5.2. Treynor–Black Portfolios

A key advantage of our framework is the ability to characterize the alphas of arbitrary impact portfolios via induced order statistics. Given this representation, it is clear that equal-weighted portfolios are not optimal in terms of achieving the best risk-adjusted returns.

However, Treynor and Black (1973) provide a methodology that is designed to maximize a portfolio's

Sharpe ratio when investors have access to alpha forecasts for a certain subset of securities.[21] This can be interpreted as a temporary departure from equilibrium in the sense of mispricings that investors can exploit or differences in information across investors in the sense of omitted risk factors. In either case, the Treynor and Black (1973) methodology allows us to construct a portfolio that maximizes the Sharpe ratio, which can be directly applied in our case to derive optimal weights for securities selected by the impact factor.

To apply the Treynor–Black framework, we rewrite the excess return of the $i$th security, $\alpha_i$, as its mean plus noise:

$$\alpha_i = \mu_i + \zeta_i, \tag{32}$$

where $\{\zeta_i\}$ are independent random variables with zero means. We can then combine $\zeta_i$ with security $i$'s idiosyncratic error, $\epsilon_i$. Because $\zeta_i$ and $\epsilon_i$ are independent, the combined idiosyncratic variance for security $i$ is simply $\sigma_i^2 + \sigma(\epsilon_i)^2$, where $\sigma_i^2$ is given in (9).

Given any number of securities selected by $\mathbf{X}$, we can form an optimal portfolio based on the Treynor–Black weights, which we summarize in the following result.

**Proposition 7.** *Under Assumption 1, the Treynor–Black weight of security i is proportional to its expected alpha divided by its combined idiosyncratic variance*:

$$\omega_i \propto \frac{\mu_i}{\sigma_i^2 + \sigma(\epsilon_i)^2}. \tag{33}$$

*In addition, if the idiosyncratic volatility, $\sigma(\epsilon_i)$, is constant across securities i, as N increases without bound, the Treynor–Black weight of security i in (33) can be further simplified to*

$$\omega_i \propto \frac{\rho\sigma_\alpha\Phi^{-1}(\xi_i)}{\sigma_\alpha^2(1-\rho^2)+\sigma(\epsilon_i)^2} \propto \Phi^{-1}(\xi_i)\cdot \textit{Constant}. \tag{34}$$

Proposition 7 gives an explicit formula for the Treynor–Black weights that optimize the risk-adjusted returns of the impact portfolio, which can easily be implemented in practice. For further intuition behind (33), recall that the variance of the $i$th induced order statistic, $\sigma_i^2$, is approximately a constant when $N$ is large (see Figure 1(b)). The expected excess return, $\mu_i = \rho\sigma_\alpha \mathrm{E}[X_{i:N}]$, varies with respect to $i$ only through the last term $\mathrm{E}[X_{i:N}]$. As a result, if each security's idiosyncratic volatility is the same, the Treynor–Black weights of security $i$ in (33) depend only on their relative ranking in the universe of $N$ securities, which is specified by the term $\Phi^{-1}(\xi_i)$ in (34).

The portfolio selected by ranking $\mathbf{X}$ and applying the Treynor–Black weights in (33) is one specific example of an *impact portfolio* we defined in Section 3.3. Treynor and Black (1973) call this the "active management" portfolio, and its return characteristics are given by

$$\alpha_A = \sum_{k=1}^{n}\omega_k\mu_k, \tag{35}$$

$$\beta_A = \sum_{k=1}^{n}\omega_k\beta_k, \tag{36}$$

$$\sigma(\epsilon_A)^2 = \sum_{k=1}^{n}\omega_k^2\left(\sigma_k^2 + \sigma(\epsilon_k)^2\right). \tag{37}$$

These results—together with the explicit quantification of individual-security alphas in Propositions 2 and 5 and the optimal Treynor–Black weights in Proposition 7—provide a complete characterization of the performance of optimal impact portfolios. In particular, the information ratio of the impact portfolio, defined as $\mathrm{E}[\alpha_A]/\sigma(\epsilon_A)$, is proportional to the correlation, $\rho$, between the unobserved alpha, $\boldsymbol{\alpha}$, and the impact factor, $\mathbf{X}$. This is closely related to the fundamental law of active management (FLAM) by Grinold (1989),[22] which provides a simple approximation of the information ratio of an active portfolio by the product of information coefficient ($\rho$ in our notation) and the *breadth* of a strategy.[23] Online Appendix A.3 provides several numerical examples of the performance of Treynor–Black impact portfolios.

### 5.3. Combining Impact and Passive Portfolios

Once the relative weights of the securities within an impact portfolio are determined, one can combine the portfolio with any other portfolio. For example, we may form an impact portfolio by ranking a company's impact on global warming, which can be combined with other characteristics, such as sustainable farming, tobacco usage, and gaming, to form an overall "ESG" portfolio. We can also add the impact portfolio to the suite of portfolios mimicking more traditional asset-pricing factors, such as value, size, and momentum.

However, perhaps the most natural application is to consider combining the impact portfolio with a passive index fund, such as the market portfolio. Let $\omega_A$ denote the weight of the impact portfolio and $1-\omega_A$ denote the weight of a passive portfolio. To maximize the Sharpe ratio of the combined portfolio, the relative weight is determined by the impact portfolio's excess return and idiosyncratic volatility:

$$\omega_A = \left(\frac{\alpha_A}{\sigma(\epsilon_A)^2}\right)\Bigg/\left(\frac{\mathrm{E}[R_m]-R_f}{\sigma_m^2}\right), \tag{38}$$

where $\mathrm{E}[R_m]$ and $\sigma_m^2$ are the expected return and variance of the passive portfolio, respectively.

We illustrate the impact portfolio's alpha and its corresponding weight, $\omega_A$, using a numerical example. Suppose the passive portfolio has an annualized risk premium of $\mathrm{E}[R_m]-R_f = 6\%$ and volatility of $\sigma_m = 15\%$.

The idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities.[24] Consider again a collection of $N = 500$ securities. We divide them into 10 decile portfolios ranked by the impact factor, $\mathbf{X}$. For several different values of $\rho$ and $\sigma_\alpha$, panel A of Table 1 reports the expected excess return of the impact portfolio,[25] $\alpha_A$ in (35), for the top and bottom two decile portfolios. Panel B of Table 1 reports the weight, $\omega_A$, and the corresponding excess return of the impact portfolio combined with the passive portfolio.

We first consider the case in which the cross-sectional standard deviation $\sigma_\alpha = 2\%$. In other words, most securities' excess returns are within $[-2\sigma_\alpha, 2\sigma_\alpha] = [-4\%, 4\%]$. This is a fairly conservative assumption for U.S. equities,[26] but even with such a modest range of $\boldsymbol{\alpha}$, the impact portfolios yield economically significant alphas. For example, $\alpha_A$ for the top decile is 1.1% when the correlation $\rho = 30\%$ and 0.4% when $\rho = 10\%$. Observe that $\rho^2$ is simply the $R^2$ of the cross-sectional regression of $\boldsymbol{\alpha}$ on $\mathbf{X}$, so a 30% (10%) correlation implies that only 9% (1%) of the variation in $\boldsymbol{\alpha}$ is explained by $\mathbf{X}$, which is a fairly plausible assumption for a typical impact factor.

When the cross-sectional standard deviation, $\sigma_\alpha$, increases to 5%, the alpha of the top decile impact portfolio increases further to 2.8% with a 30% correlation and 0.9% with a 10% correlation.

When combined with the passive index to form long/short portfolios, the optimal portfolio contains significant weight from the impact portfolio, sometimes implying a highly leveraged position (the top half of panel B of Table 1). The corresponding gain in expected excess return for the combined portfolio is 8.9% when

**Table 1.** Expected Excess Returns for the Impact Portfolios and Their Corresponding Treynor–Black Weights When Combined with a Passive Portfolio

Panel A: Expected excess return of the impact portfolio ($\alpha_A$)

| Correlation $\rho$ | Bottom, % | Second, % | Ninth, % | Top, % |
|---|---|---|---|---|
| $\sigma_\alpha = 2\%$ | | | | |
| 30% ($R^2 = 9\%$) | −1.1 | −0.6 | 0.6 | 1.1 |
| 10% ($R^2 = 1\%$) | −0.4 | −0.2 | 0.2 | 0.4 |
| −10% ($R^2 = 1\%$) | 0.4 | 0.2 | −0.2 | −0.4 |
| −30% ($R^2 = 9\%$) | 1.1 | 0.6 | −0.6 | −1.1 |
| $\sigma_\alpha = 5\%$ | | | | |
| 30% ($R^2 = 9\%$) | −2.8 | −1.6 | 1.6 | 2.8 |
| 10% ($R^2 = 1\%$) | −0.9 | −0.5 | 0.5 | 0.9 |
| −10% ($R^2 = 1\%$) | 0.9 | 0.5 | −0.5 | −0.9 |
| −30% ($R^2 = 9\%$) | 2.8 | 1.6 | −1.6 | −2.8 |

Panel B: Impact portfolios combined with a passive portfolio

| | Weight $\omega_A$ | | | | Expected excess return $\omega_A\alpha_A$ | | | |
|---|---|---|---|---|---|---|---|---|
| Correlation $\rho$ | Bottom | Second | Ninth | Top | Bottom, % | Second, % | Ninth, % | Top, % |
| Long/short, $\sigma_\alpha = 2\%$ | | | | | | | | |
| 30% ($R^2 = 9\%$) | −1.31 | −0.78 | 0.78 | 1.31 | 1.4 | 0.5 | 0.5 | 1.4 |
| 10% ($R^2 = 1\%$) | −0.44 | −0.26 | 0.26 | 0.44 | 0.2 | 0.1 | 0.1 | 0.2 |
| −10% ($R^2 = 1\%$) | 0.44 | 0.26 | −0.26 | −0.44 | 0.2 | 0.1 | 0.1 | 0.2 |
| −30% ($R^2 = 9\%$) | 1.31 | 0.78 | −0.78 | −1.31 | 1.4 | 0.5 | 0.5 | 1.4 |
| Long/short, $\sigma_\alpha = 5\%$ | | | | | | | | |
| 30% ($R^2 = 9\%$) | −3.23 | −1.93 | 1.93 | 3.23 | 8.9 | 3.1 | 3.1 | 8.9 |
| 10% ($R^2 = 1\%$) | −1.08 | −0.64 | 0.64 | 1.08 | 1.0 | 0.3 | 0.3 | 1.0 |
| −10% ($R^2 = 1\%$) | 1.08 | 0.64 | −0.64 | −1.08 | 1.0 | 0.3 | 0.3 | 1.0 |
| −30% ($R^2 = 9\%$) | 3.23 | 1.93 | −1.93 | −3.23 | 8.9 | 3.1 | 3.1 | 8.9 |
| Long only, $\sigma_\alpha = 2\%$ | | | | | | | | |
| 30% ($R^2 = 9\%$) | 0 | 0 | 0.78 | 1.00 | 0 | 0 | 0.5 | 1.1 |
| 10% ($R^2 = 1\%$) | 0 | 0 | 0.26 | 0.44 | 0 | 0 | 0.1 | 0.2 |
| −10% ($R^2 = 1\%$) | 0.44 | 0.26 | 0 | 0 | 0.2 | 0.1 | 0 | 0 |
| −30% ($R^2 = 9\%$) | 1.00 | 0.78 | 0 | 0 | 1.1 | 0.5 | 0 | 0 |
| Long only, $\sigma_\alpha = 5\%$ | | | | | | | | |
| 30% ($R^2 = 9\%$) | 0 | 0 | 1.00 | 1.00 | 0 | 0 | 1.6 | 2.8 |
| 10% ($R^2 = 1\%$) | 0 | 0 | 0.64 | 1.00 | 0 | 0 | 0.3 | 0.9 |
| −10% ($R^2 = 1\%$) | 1.00 | 0.64 | 0 | 0 | 0.9 | 0.3 | 0 | 0 |
| −30% ($R^2 = 9\%$) | 1.00 | 1.00 | 0 | 0 | 2.8 | 1.6 | 0 | 0 |

*Notes.* We set $N = 500$ and assume that the passive portfolio has an annualized expected excess return of $\mathrm{E}[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. The idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities. Panel A reports $\alpha_A$ for the impact portfolio, and Panel B reports the weight, $\omega_A$, and the corresponding excess return of the impact portfolio combined with the passive portfolio.

$\rho = 30\%$ and 1.0% when $\rho = 10\%$. Of course, these calculations are meant only to be proofs of concept because we have not considered specific regulatory and institutional constraints for investors in practice. The leveraged portfolios with short positions in the passive index may be suitable only for certain types of hedge funds.

The bottom half of panel B of Table 1 shows the performance of long-only combined portfolios, which corresponds to a constrained weight, $\omega_A$, between zero and one. When the correlation, $\rho$, is positive, only the top two deciles provide positive alpha and therefore, yield positive weight in the combined portfolio. When $\sigma_\alpha = 5\%$, for example, this leads to a sizable expected excess return of 2.8% when $\rho = 30\%$ and 0.9% when $\rho = 10\%$. Section 6 provides two real-world examples with positive impact alpha in the context of venture philanthropy and R&D subsidies.

When the correlation, $\rho$, is negative, only the bottom two deciles provide positive excess returns. However, securities in these deciles have the lowest impact factor $\mathbf{X}$. Therefore, impact investors may get an estimate of the opportunity cost of not investing in these portfolios. Divesting from sin stocks provides a real-world example in this case, which we discuss further in Section 6.

More generally, Online Appendix A.4 provides additional examples that generalize results in Table 1.

Finally, we show how our impact portfolio improves the efficient frontier and the capital market line to achieve a "superefficient frontier" under the assumption that $\boldsymbol{\alpha}$ are mispricings. Under the alternate omitted-factor interpretation, the "superefficiency" of the new frontier may be viewed as the result of additional risk premia not accessible to investors except through impact-portfolio managers.

**Proposition 8.** *Under Assumption 1, the return of the final portfolio, P, that consists of the impact portfolio with Treynor–Black weights and the passive market portfolio is*

$$\begin{aligned} R_P - R_f &= \omega_A R_A + (1-\omega_A)R_m - R_f \\ &= \omega_A \alpha_A + (R_m - R_f)(\beta_A \omega_A + (1-\omega_A)) + \omega_A \epsilon_A, \end{aligned} \tag{39}$$

*where $R_m$ is the return of the passive portfolio. The expected value and variance of $R_P$ are*

$$E[R_P] - R_f = \omega_A \alpha_A + (E[R_m] - R_f)(\beta_A \omega_A + (1-\omega_A)), \tag{40}$$

$$\mathrm{Var}[R_P] = \mathrm{Var}[R_m](\beta_A \omega_A + (1-\omega_A))^2 + \omega_A^2 \sigma(\epsilon_A)^2. \tag{41}$$

*This forms a superefficient frontier in comparison with the capital market line associated with the passive portfolio.*

Figure 3 displays the passive portfolio as well as several combinations with impact portfolios in relation to the efficient frontier. We continue to assume that the passive portfolio has an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. In Figure 3, (a) and (c), the idiosyncratic volatility is assumed to be a constant $\sigma(\epsilon_i) = 15\%$ for all securities. As the correlation, $\rho$, and variance in alpha, $\sigma_\alpha^2$, increase, the impact portfolios (defined as the top half of the securities ranked by $\mathbf{X}$) are able to improve the original capital market line for both long/short and long-only portfolios, leading to superefficient frontiers.

The results in this section have so far relied on the assumption that the idiosyncratic volatility, $\sigma(\epsilon_i)$, is cross-sectionally constant. To check the robustness of our results, we simulate a collection of securities where the $i$th security's idiosyncratic volatility follows a log-normal distribution:

$$\log(\sigma(\epsilon_i)) \sim \mathrm{Normal}(\mu_\epsilon, \sigma_\epsilon). \tag{42}$$

Calibrating to empirically plausible values in the literature (e.g., Kuntz 2020), we perform simulations for $\log(\mu_\epsilon) = 15\%$ and $\sigma_\epsilon = 1$. Figure 3, (b) and (d) confirms that even with such cross-sectional heterogeneity, the capital market line is still improved, leading to superefficient frontiers.

## 5.4. Qualifications

The Treynor–Black portfolios outlined in Propositions 7 and 8 provide the optimal impact portfolios when investors face no constraints. In practice, investors may face additional regulatory and institutional constraints, such as limits on short positions and sector concentration. Because our framework quantifies the expected alpha for individual securities, it is easy to derive optimal impact portfolios accounting for these additional constraints as well. Section 5.3 provides examples of long-only portfolios. In the meantime, we emphasize that the performance metrics associated with these portfolios should be taken as proofs of concept, and real-world investors should account for their specific constraints and trading costs on top of our framework to derive proper benchmarks.

Our framework provides a methodology to construct more robust portfolios in practice compared with traditional approaches. First, investors may follow traditional mean-variance optimization and integrate the level of impact as part of either the objective or the constraint.[27] However, estimating the expected return and covariance matrix in traditional Markowitz portfolios is known to be challenging, which often leads to unstable portfolios (Brodie et al. 2009, Tu and Zhou 2011). Thanks to the ability to quantify individual security alphas based on a small (two) number of parameters, our approach provides a much more regularized approach to constructing optimal impact portfolios.

Second, investors may also build portfolios through a univariate regression, in which security alphas are

**Figure 3.** (Color online) Superefficient Frontiers from the Combined Portfolio That Consists of the Impact Portfolio and Passive Market Portfolio

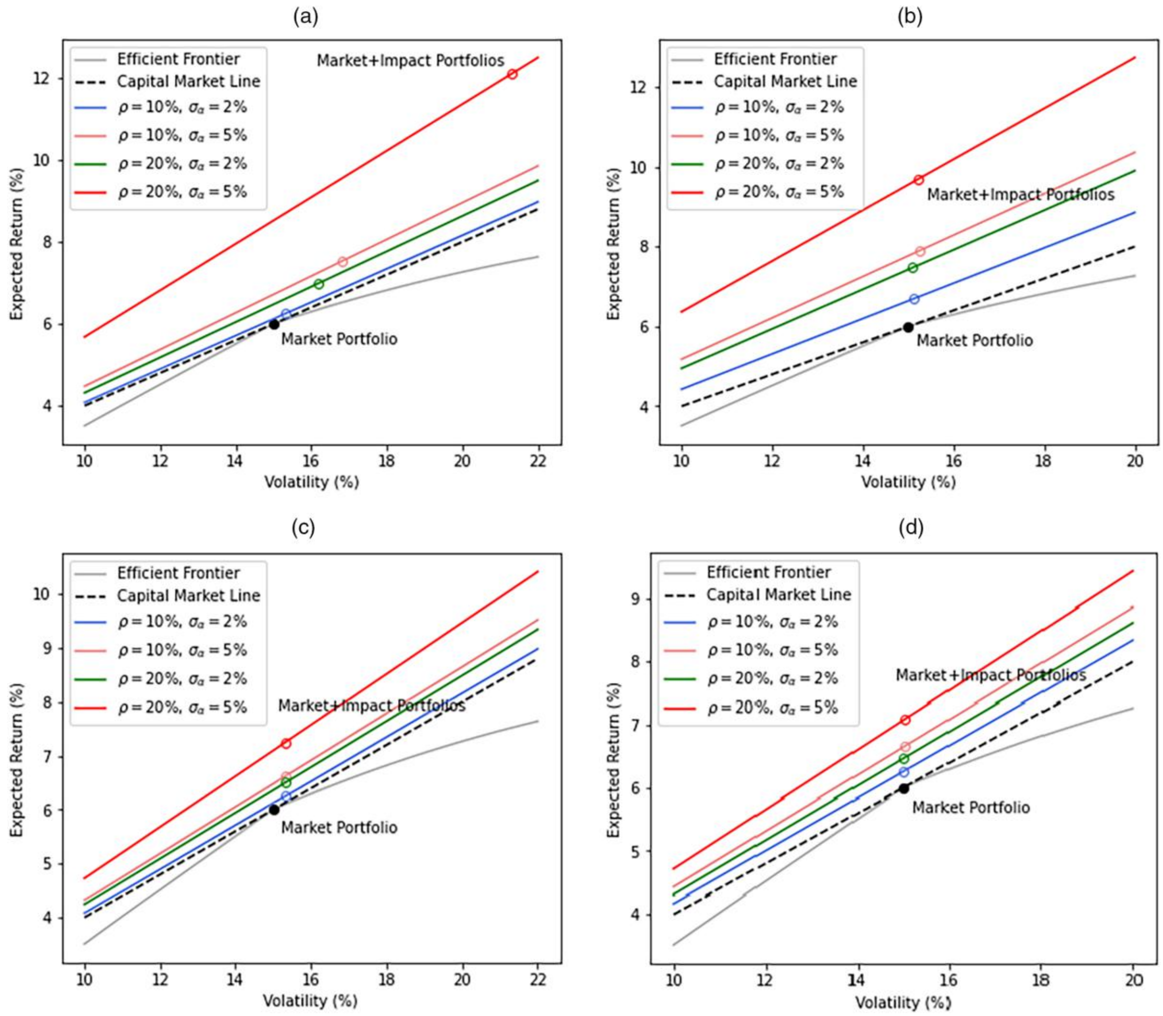


*Notes*. Panels (a) and (b) allow short positions in the passive index, whereas panels (c) and (d) consider long-only portfolios. In panels (a) and (c), the idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities. In panels (b) and (d), we simulated idiosyncratic volatility based on (42) and apply a maximum leverage ratio of 3:1. (a) Long/short, constant $\sigma(\epsilon_i)$. (b) Long/short, simulated $\sigma(\epsilon_i)$. (c) Long only, constant $\sigma(\epsilon_i)$. (d) Long only, simulated $\sigma(\epsilon_i)$.

estimated by regressing returns onto the impact factor **X**. Compared with this approach, our framework corresponds more closely to how the industry actually constructs impact portfolios, which usually relies on ordering securities based on **X** and assigning weights in each quartile or decile. In fact, our theory requires weaker assumptions because the results only depend on the rank of **X**, not its value. As a result, the optimal portfolio relies on quantiles of **X**, making our approach robust against noise and outliers as demonstrated empirically by Lo et al. (2022a). The focus on rank also allows for the generalization of our theory to capture nonlinear dependence between returns and the impact factor as shown in a subsequent study (Lo et al. 2022b).

The key insight of our framework lies in the ability to quantify the distribution of impact-ranked returns when mispricings exist for individual securities. In the omitted factor interpretation, it identifies impact factors **X** that explain the cross-section of security returns and furthermore, quantifies the excess returns associated with these omitted factors. If alphas are fully observed—either in the mispricing or omitted factor interpretation—then investors have complete information to trade-off returns and impact, and optimal impact portfolios can simply be constructed following Treynor and Black (1973). Our framework shows how to quantify the impact of impact investing when alphas are unknown in the real world and construct the optimal portfolio that achieves a "superefficient frontier" either through mispricings or additional risk premia not accessible to investors except through impact-portfolio managers.

When alphas are interpreted as omitted factors whose beta loadings are correlated with the impact factor, $\mathbf{X}$, our model also provides an approach to recovering the stochastic discount factor (SDF).[28] To see this, observe that our results in Proposition 5 quantify the excess returns of impact-ranked securities and in particular, identify $\mathbf{X}$ that explains the cross-section of returns. The correlation scaled by the standard deviation, $\rho_{adj}\sigma_\lambda$, can be treated as the risk premium for the omitted factor, whereas the standardized impact score, $\mathrm{E}[X_{i:N}]$, provides the factor loadings. Because a factor model in expected returns always corresponds to a factor model in SDF space, our framework essentially recovers the SDF with respect to this factor $\mathbf{X}$, although the corresponding SDF is not uniquely determined.[29]

## 6. Applications to Five Impact Investments

In this section, we apply our framework to five particular examples of impact investing: biotech venture philanthropy, semiconductor R&D consortium, divesting from sin stocks, ESG investing, and the GameStop short squeeze during January 2021.

### 6.1. Venture Philanthropy: The Cystic Fibrosis Foundation

The concept of VP was introduced by Letts et al. (1997), who suggested that nonprofit organizations could learn useful practices from venture capitalists. In particular, recent biomedical advances have created significant opportunities for a new generation of therapeutics (Sharp and Hockfield 2017). However, early-stage R&D efforts often face a dearth of funding given the high risk of failure and significant funding requirements. This has been particularly true for rare disease drug development, where market sizes are often too small to attract much attention and funding (Kim and Lo 2019).

We consider the example of the CF Foundation—profiled in the case study by Kim and Lo (2019)—and conclude that VP in biomedicine can produce significant positive excess returns. This example illustrates the possibility of an impact investment that is positively correlated with $\boldsymbol{\alpha}$ or an omitted factor that patient advocacy groups can more easily exploit than typical investors.

The CF Foundation is the world's leading philanthropic organization for CF, a rare genetic disease that currently affects more than 30,000 Americans. Over a period of 12 years, the CF Foundation invested \$150 million to fund CF drug development efforts at Vertex Pharmaceuticals, which led to the identification and development of Kalydeco, the first Food and Drug Administration-approved treatment to address the underlying causes of CF. In 2014, their rights to Vertex royalties were sold to an outside investment firm, New York City-based Royalty Pharma, for \$3.3 billion in cash.

If we assume for simplicity that the \$150 million investment was made up front and the \$3.3 billion sale occurred 12 years later, this implies a compound annual return of 29.4% over this period. To estimate the realized $\alpha$ of this investment, we require an estimate of the cost of capital of Vertex during the 12-year investment period from 2002 to 2013 prior to the 2014 royalty sale. We consider a simple CAPM model:

$$R_i - R_f = \alpha_i + \beta_i(R_M - R_f) + \epsilon_i, \tag{43}$$

where $R_i$ is the return of the $i$th security with CF Foundation being one of them, $R_f$ is the risk-free rate, and $R_M$ is the market return. The average beta of Vertex between January 2, 2001 and December 31, 2013 was 1.42 (see Online Appendix A.5). The average five-year constant-maturity treasury yield from January 2001 to December 2013 was 2.8%,[30] and the annualized compound return of the Center for Research in Security Prices (CRSP) value-weighted returns index with dividends during this period was 5.4%; hence, a simple CAPM estimate of the cost of capital is $1.42 \times (5.4\% - 2.8\%) + 2.82\% = 6.5\%$.

Of course, this crude estimate does not account for the illiquid nature of biomedical assets and the financing risks that their multiyear investment horizons pose. A cost of capital of 20% for privately held biotech investments is a commonly used industry benchmark. Therefore, a plausible range for the $\alpha$ of the CF Foundation's investment in Vertex is 9.4% (using a 20% cost of capital) to 22.9% (using a 6.5% cost of capital).

Using this estimated range for the CF Foundation's $\alpha$, we can estimate its correlation with a "rare disease impact investing" factor $\mathbf{X}$. Making a few additional assumptions about auxiliary parameters, we can reverse engineer the implied correlation, $\rho$, that is consistent with this performance range, which is $[35\%, 86\%]$.[31] Our highly stylized calculations are not meant to yield a rigorous estimate of the true alpha associated with drug development for rare diseases, and the plausible range of the true alpha is likely larger, potentially including zero. However, more systematic empirical analyses of the biopharma industry show that pharmaceutical companies have become increasingly profitable, with risk-adjusted returns outperforming the aggregate stock market in recent years (Thakor et al. 2017, Lo and Thakor 2019). The example of the CF Foundation provides additional intuition for how impact and investment performance need not be a zero-sum game in the presence of sufficient correlation between impact and performance.

However, there is a deeper message in this striking example, which is that in certain cases, impact is a *prerequisite* for performance. The CF Foundation's main objective—helping to create a disease-modifying drug for CF—was, in fact, the primary source of Vertex's outsized investment performance. The fact that the foundation focused on this one long-term goal—to the exclusion of shorter-term financial metrics and milestones—and was willing to continue investing in Vertex over multiple years despite business cycle fluctuations (including the 2008 Financial Crisis) contributed significantly to its success (both in impact and in financial returns). Indeed, many traditional venture capitalists have shied away from investing in projects with such high risks and long-term capital commitments. In other words, in this case, correlation may actually be *causation*; impact can sometimes be responsible for financial success.

### 6.2. An R&D Consortium: Sematech

The second example of impact investing is Sematech, an R&D consortium established in 1987 by 14 U.S. semiconductor firms and the U.S. Government to solve common manufacturing problems by leveraging shared R&D resources and research results. It was funded by a combination of the U.S. Government initially and member firms later on.

To measure the returns to Sematech as a form of impact investing, we follow Irwin and Klenow (1996a) and use a difference-in-difference approach to compare the excess return to Sematech member firms against those of nonmembers, controlling for past returns prior to the formation of Sematech. In particular, we collect monthly return data from CRSP for all U.S. firms whose principal business is semiconductors and related devices with Standard Industrial Classification (SIC) 3674 from 1975 to 1999.[32] Our sample consists of 52 firms, including 11 of the original 14 members of Sematech.[33]

For each firm, we divide the sample into pre-Sematech (January 1975 to August 1987) and post-Sematech (September 1987 to December 1999) periods and estimate two CAPM models:

$$R_{i,\text{pre}} - R_{f,\text{pre}} = \alpha_{i,\text{pre}} + \beta_{i,\text{pre}}(R_{M,\text{pre}} - R_{f,\text{pre}}) + \epsilon_i,$$
$$R_{i,\text{post}} - R_{f,\text{post}} = \alpha_{i,\text{post}} + \beta_{i,\text{post}}(R_{M,\text{post}} - R_{f,\text{post}}) + \epsilon_i, \tag{44}$$

where $R_{i,\cdot}$ is the return of the *i*th security, $R_{f,\cdot}$ is the risk-free rate estimated by the five-year constant-maturity treasury yield, and $R_{M,\cdot}$ is the market return estimated by the CRSP value-weighted returns index including dividends.

Table 2 summarizes the estimated annualized alphas. Sematech member firms have an average excess return of 7.23% after joining the R&D consortium compared with an average of −2.67% prior to Sematech's formation. In comparison, the average excess returns for non-member firms are around −9% in both periods. This leads to an overall difference-in-difference estimate of 10.22% lift in annualized excess returns for a firm that joined the R&D consortium, with a 90% bootstrap confidence interval of $[0.56\%, 19.88\%]$. If we consider an impact factor $\mathbf{X}$ with $X_i = 1$ representing Sematech membership and zero otherwise, this leads to an estimated correlation, $\rho$, of 25% with a range of $[1.4\%, 49\%]$ between the Sematech impact factor and excess returns.[34]

Our estimated excess returns based on the simple CAPM model should only be interpreted as a crude approximation to the true effect of Sematech, and the true effect may be different once other control variables are accounted for. In particular, firms' entry decisions into the consortium may be endogenously affected by other factors, such as the size and the profitability of the firm. We conduct a simple analysis in Online Appendix A.6 to mitigate this concern by comparing basic characteristics for Sematech versus non-Sematech firms and by showing that the impact from Sematech largely remains robust on samples that are constructed to match the distribution of Sematech firms.

Our results are also consistent with prior evidence on Sematech's far-reaching effects on members' R&D spending and profitability (Irwin and Klenow 1996a, b; Link et al. 1996). In particular, Irwin and Klenow (1996a) show that joining Sematech improves firms' profitability, which provides a potential channel through which excess returns are earned.

This example shows how the impact of R&D consortia can be measured in our framework, and the difference-in-difference approach yields a potentially causal estimate of the positive financial impact of the Sematech impact factor.

### 6.3. Divesting Sin Stocks

Another particular type of impact investing is avoiding or divesting sin stocks—stocks from companies involved in or associated with activities considered unethical or immoral. Although there may be a degree of subjectivity

**Table 2.** Estimated Annualized CAPM Excess Returns Based on a Difference-in-Difference Approach That Compares Sematech Member Firms with Nonmembers Before and After the Formation of Sematech

| Firm | Period | Excess return, % | Excess return (post minus pre), % | Excess return (difference in difference), % |
|---|---|---|---|---|
| Sematech members | Pre | −2.67 | 9.90** | 10.22* |
| | Post | 7.23*** | | |
| Nonmembers | Pre | −8.90*** | −0.33 | |
| | Post | −9.22*** | | |

*Estimated coefficients are significant at the 10% level based on bootstrap confidence intervals; **estimated coefficients are significant at the 5% level based on bootstrap confidence intervals; ***estimated coefficients are significant at the 1% level based on bootstrap confidence intervals.

involved in determining what is considered sinful, common examples include companies involved in producing, distributing, or otherwise supporting alcohol, tobacco, gambling, sex-related industries, and firearms. It has been found that sin stocks are less held by norm-constrained institutions, such as pension plans, as compared with mutual or hedge funds, and they receive less coverage from analysts. As a result, sin stocks seem to yield higher expected returns (Fabozzi et al. 2008, Hong and Kacperczyk 2009, Statman and Glushkov 2009).

This empirical fact implies a negative correlation between a stock's excess return and an "antisin stock" factor. In terms of the superefficient frontier shown in Figure 3, divesting from sin stocks likely yields a negative return and a lower efficient frontier. This leads to a natural definition of the cost to this specific impact factor.

We use Hong and Kacperczyk (2009) to calibrate our model and focus on tobacco, alcohol, and gambling as proxies for sin stocks. The authors report a monthly excess return of 0.26% for an equal-weighted portfolio long sin stocks and short their comparables by running a time series regression controlling for market, size, value, and momentum factors using equity data in the United States from 1965 to 2006. This corresponds to the following four-factor asset-pricing model,

$$R_i - R_f = \alpha_i + \beta_{i,1}(R_M - R_f) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}MOM + \epsilon_i, \tag{45}$$

and an impact factor $\mathbf{X}$ representing whether a stock is considered sin. The Hong and Kacperczyk (2009) estimate can be used to calculate the implied correlation between $\boldsymbol{\alpha}$ and $\mathbf{X}$ in our model using results from Proposition 2 and Proposition A.2 in the online appendix (see also the discussions in Online Appendix A.2).

Panel A of Table 3 summarizes these calibration results.[35] The implied correlation is 27% ($R^2 = 7.2\%$) assuming a standard deviation of cross-sectional alpha of $\sigma_\alpha = 5\%$.[36] This leads to a measure of the cost of avoiding sin stocks. If we form an impact portfolio based on the top half of all securities based on the antisin factor, it leaves an excess return of 1.7% per annum on the table. If we form a Treynor–Black portfolio based on the omitted sin stocks and the passive market portfolio, we could have achieved a leveraged alpha of 14.4% per annum with a (leveraged) weight of 8.58 for the sin stocks portfolio. On the other hand, if we form an impact portfolio by leaving out only the top decile (or top 2%) of the most sinful stocks, the opportunity cost is 2.5% (3.3%).

In fact, a few studies have tried to understand why sin stocks appear to show positive excess returns. In particular, Blitz and Fabozzi (2017) show that sin stocks indeed exhibit a significantly positive CAPM alpha, but this alpha disappears completely when controlling not only for classic factors, such as size, value, and momentum, but also for exposures to the two new Fama and French (2015) quality factors—profitability and investment. This corresponds to the following six-factor asset-pricing model:

$$R_i - R_f = \alpha_i + \beta_{i,1}(R_M - R_f) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}MOM + \beta_{i,5}RMW + \beta_{i,6}CMA + \epsilon_i. \tag{46}$$

We also summarize the implied correlation and cost of divesting sin stocks based on Blitz and Fabozzi (2017) in panel B of Table 3.[37] Both the correlation and sin stock excess returns decrease sharply based on this study.

This example highlights the fact that the measurement of excess returns of impact investing depends on the specific asset-pricing model used to estimate alpha. Our framework can be applied to any number of factors as specified in (1) and (2). Indeed, a factor may yield positive correlation with alpha under one asset-pricing model (implying a positive excess return) and may disappear or change sign after controlling for additional factors.

**Table 3.** Estimated Cost in Excess Return per Annum for Avoiding Sin Stocks Calibrated to Prior Empirical Studies

| | | Expected excess return, % | |
|---|---|---|---|
| Impact portfolio | Weight of impact portfolio $\omega_A$ | Impact portfolio $\alpha_A$ | Combined with passive portfolio $\omega_A\alpha_A$ |
| Panel A: Hong and Kacperczyk (2009) | | | |
| Implied correlation $\rho = 27\%$ ($R^2 = 7.2\%$) | | | |
| Top half | 8.58 | 1.7 | 14.4 |
| Top decile | 3.78 | 2.5 | 9.35 |
| Top 2% | 1.04 | 3.3 | 3.4 |
| Panel B: Blitz and Fabozzi (2017) | | | |
| Implied correlation $\rho = 10\%$ ($R^2 = 1.1\%$) | | | |
| Top half | 3.30 | 0.6 | 2.1 |
| Top decile | 1.45 | 1.0 | 1.4 |
| Top 2% | 0.40 | 1.3 | 0.5 |

*Note.* Here, we assume that the passive portfolio has an annualized risk premium of $\mathrm{E}[R_m] - R_f = 6\%$, volatility of $\sigma_m = 15\%$, and $\sigma_\alpha = 5\%$.

### 6.4. ESG Investing

More generally, SRI and ESG-aware investing have both drawn an increasing amount of attention in recent years. Our model provides a systematic framework to measure the financial impact of SRI and ESG—positive or negative—and construct optimal portfolios based on the correlation between the impact characteristic and excess returns.

As with the sin stocks in Section 6.3, we also calibrate our model with respect to several studies in Table 4. Panel A uses the estimates from Berg et al. (2022c), who construct ESG portfolios based on a range of aggregation methods to denoise ESG metrics from six different data vendors. They report an annualized Fama–French five-factor alpha of 4.5% for the top-bottom decile portfolio in the United States using a simple average ESG score, which implies a correlation of 26% ($R^2 = 6.7\%$) between stock alpha and the ESG impact factor. This is consistent with opinions from industry advocates of ESG, although the magnitude of excess returns in the literature varies with the specific ESG metric.

In contrast, Baker et al. (2022) study the U.S. bond market and report a yield difference of six basis points at issuance for green bonds below other ordinary bonds.[38] This corresponds to a plausible and economically meaningful 0.6% difference in value on a bond with a 10-year duration. Panel B of Table 4 shows the implied correlation of $-2\%$.[39] This result points to a negative to neutral ESG alpha in the bond market, in which case the excess return, $\omega_A\alpha_A$, in the last column should be interpreted as the cost to ESG investing in this particular market.

In addition, several recent studies provide additional insights into what economic state variables potentially drive the observed ESG returns. Using stock data from S&P 500 and Russell 3000 in 1993–2013, Bansal et al. (2022) document a "luxury-good effect" for an ESG

**Table 4.** Estimated ESG Excess Returns per Annum Calibrated to Prior Empirical Studies

| | | Expected excess return, % | |
|---|---|---|---|
| Impact portfolio | Weight of impact portfolio $\omega_A$ | Impact portfolio $\alpha_A$ | Combined with passive portfolio $\omega_A\alpha_A$ |
| **Panel A: Berg et al. (2022c)** | | | |
| Implied correlation $\rho = 26\%$ ($R^2 = 6.7\%$) | | | |
| Top half | 3.20 | 1.61 | 5.15 |
| Top decile | 1.41 | 2.38 | 3.35 |
| Top 2% | 0.39 | 3.14 | 1.24 |
| **Panel B: Baker et al. (2022)** | | | |
| Implied correlation $\rho = -2.0\%$ ($R^2 = 0.04\%$) | | | |
| Top half | −1.06 | −0.02 | 0.03 |
| Top decile | −0.47 | −0.04 | 0.02 |
| Top 2% | −0.13 | −0.05 | 0.00 |
| **Panel C: Bansal et al. (2022) ("good times")** | | | |
| Implied correlation $\rho = 22\%$ ($R^2 = 4.7\%$) | | | |
| Top half | 2.69 | 1.35 | 3.64 |
| Top decile | 1.19 | 1.99 | 2.37 |
| Top 2% | 0.33 | 2.65 | 0.88 |
| **Panel D: Bansal et al. (2022) ("bad times")** | | | |
| Implied correlation $\rho = -0.2\%$ ($R^2 = 0.0\%$) | | | |
| Top half | −0.02 | −0.01 | 0.00 |
| Top decile | −0.01 | −0.02 | 0.00 |
| Top 2% | −0.00 | −0.02 | 0.00 |
| **Panel E: Pástor et al. (2022)** | | | |
| Implied correlation $\rho = 5.5\%$ ($R^2 = 0.3\%$) | | | |
| Top half | 0.69 | 0.34 | 0.24 |
| Top decile | 0.30 | 0.51 | 0.15 |
| Top 2% | 0.08 | 0.68 | 0.06 |
| **Panel F: Lo et al. (2022)** | | | |
| Implied correlation $\rho = 3.8\%$ ($R^2 = 0.1\%$) | | | |
| Top half | 0.47 | 0.24 | 0.11 |
| Top decile | 0.21 | 0.35 | 0.07 |
| Top 2% | 0.06 | 0.46 | 0.03 |

*Note.* Here, we assume that the passive portfolio has an annualized risk premium of $E[R_m] - R_f = 6\%$, volatility of $\sigma_m = 15\%$, $\sigma_\alpha = 5\%$ for stock markets, and $\sigma_\alpha = 1\%$ for bond markets.

factor that combines analyst ratings, firm announcements, and realized incidents. Stocks with higher ESG ratings significantly outperform lower-ranked ones during good economic times but not during bad economic times, resembling the demand for luxury goods. Their analysis corresponds to an impact factor **X** representing the customized ESG score and a four-factor asset-pricing model with an interaction term:

$$R_i - R_f = \alpha_i + \beta_{i,0} I + \beta_{i,1}(R_M - R_f) + \beta_{i,2} SMB + \beta_{i,3} HML + \beta_{i,4} MOM + \gamma_{i,1}(R_M - R_f)\cdot I + \gamma_{i,2} SMB \cdot I + \gamma_{i,3} HML \cdot I + \gamma_{i,4} MOM \cdot I + \epsilon_i, \tag{47}$$

where $I$ is a dummy variable representing good economic times defined by the cyclically adjusted real P/E ratios from Shiller (2005). We report the implied correlation between stock alpha and the ESG factor based on their estimates for good and bad economic times in panels C and D of Table 4, respectively. During good economic times, Bansal et al. (2022) report a monthly Fama–French four-factor alpha of 0.315% for the top-bottom ESG portfolio. This implies a 22% correlation between stock alpha and the ESG factor as well as sizable positive excess returns for the impact portfolios. However, during bad economic times, the monthly Fama–French four-factor alpha in Bansal et al. (2022) becomes −0.0026%, rendering all of our estimates of correlation and ESG alpha essentially zero.

Pástor et al. (2022) provide another channel through which ESG returns can be explained. Based on the environmental score from MSCI ESG Ratings data, which corresponds to the impact factor **X** in our framework, they document a highly significant monthly Fama–French three-factor alpha of 0.50% for a green-minus-brown portfolio,[40] which reduces to an insignificant level of 0.05% when excess returns are regressed on proxies of unexpected climate shocks:

$$R_i - R_f = \alpha_i + \beta_{i,1}(R_M - R_f) + \beta_{i,2} SMB + \beta_{i,3} HML + \beta_{i,4} DCC + \beta_{i,5} DCC_{\text{Lag1}} + \epsilon_i, \tag{48}$$

where $DCC$ stands for a "delta climate concern" variable constructed from the Ardia et al. (2022) climate change concerns index and $DCC_{\text{Lag1}}$ is its one-month lagged values. This implies a much smaller correlation of 5.5% (panel E of Table 4) compared with estimates from, for example, Berg et al. (2022c).

Similarly, panel F of Table 4 shows the estimates based on Lo et al. (2022a), who use the Trucost Environmental data and find a similar effect to the Pástor et al. (2022) findings for a wide range of environmental measures including, in particular, carbon emissions. Based on the framework in this article, they estimate an annualized alpha of 2.38% for a long/short green portfolio. This alpha reduces to a statistically insignificant level of 0.3% per annum with the same proxy variables for unexpected climate shocks:

$$R_i - R_f = \alpha_i + \beta_{i,1}(R_M - R_f) + \beta_{i,2} SMB + \beta_{i,3} HML + \beta_{i,4} RMW + \beta_{i,5} CMA + \beta_{i,6} DCC + \beta_{i,7} DCC_{\text{Lag1}} + \epsilon_i, \tag{49}$$

which corresponds to a small implied correlation of 3.8% between excess returns and a low-carbon impact factor **X**.[41]

The five studies we highlight in Table 4 underscore the difficulty in measuring consistent excess returns of ESG, which depend on many factors, including the asset class, region, and time period. In addition, the specific choice of asset-pricing model also affects the empirical estimates of ESG alpha (see also Geczy et al. 2021, Madhavan et al. 2021).

### 6.5. The GameStop Phenomenon

In January 2021, the share price of GME—a struggling video game retailer that had recently announced a 30% decline in 2020Q3 net sales in part because of an 11% reduction in their store base—went from \$17.25 on January 5 to an all-time high of \$347.51 on January 27.[42] Although few investment professionals would consider GME an "impact investment," it is difficult to categorize it as anything else given the apparent origin of its meteoric price spike.

The key turning point for GME seemed to be growing interest among retail investors affiliated with the Reddit forum "r/WallStreetBets." Although it is difficult to determine the exact cause and motivation behind the early initiators, the GME price spike is unlikely to have been driven by changes in the fundamentals of the company,[43] but rather, it was caused by a combination of a grassroots "David versus Goliath" conflict between retail investors and hedge-fund short sellers and the trend followers taking advantage of this dynamic. Other stocks that seemed to be involved in this movement included AMC Entertainment Holdings (AMC) and Blackberry (BB), both of which were facing short-selling pressure from institutional investors. These events attracted substantial media attention because of the populist narrative that was playing out on social media at the time as well as the extraordinary price gyrations and wealth transfers involved. As shown in Figure 4, if an investor bought \$1 of GME at the beginning of October 2020, she would have gained over \$30 at the end of January 2021. Strikingly, GME's stock price stayed at a high level ever since. As of December 2022, the initial \$1 is still worth over \$10. This sustained price increase is no longer just attributable to temporary supply/demand imbalances; the increased value of GME equity gave the company substantial resources to improve its business, thereby allowing it to make positive changes in its

**Figure 4.** (Color online) Growth of One Dollar Invested in GME, AMC, and BB on a Semilogarithmic Scale

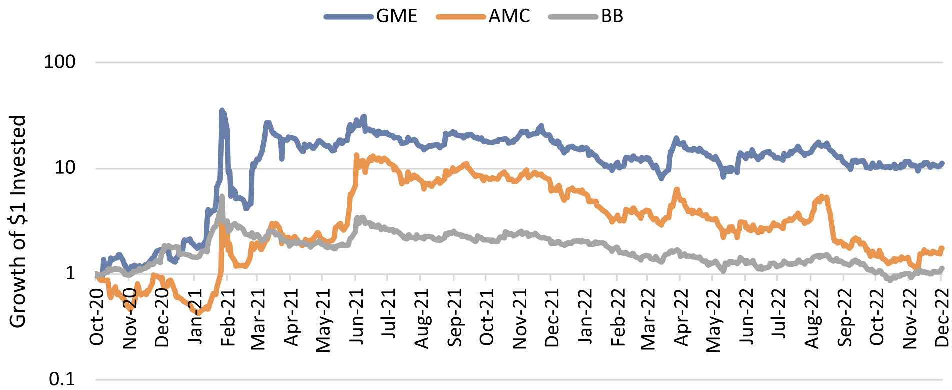


business that permanently altered its future prospects for the better.[44]

In a very real sense, WallStreetBets participants can be viewed as a general form of impact investors. Additionally, by most accounts, they have been highly successful in achieving the impact they desired (i.e., punishing the short sellers, pushing up the price of an underdog company bullied by elite institutional investors, and saving the company from bankruptcy). However, to distinguish this type of activity from traditional impact investing, we shall call the GME phenomenon "price-impact investing."

The GameStop example is likely driven by retail mania and herding rather than a social welfare-enhancing investment. It has a very different—perhaps opposite—underlying economic motivation from the previous impact-investing examples we gave. However, there are mechanical similarities between the GameStop mania and for example, ESG investing. The former can be driven by investor mania with fast and temporary shifts in preferences toward meme stocks, whereas the latter can be driven by slow but persistent shifts in preferences toward green assets (or concerns toward climate change), as shown by Lo et al. (2022a) and Pástor et al. (2022). In fact, some behavioral economists might argue that the distinction between manias and preference shifts is largely a matter of semantics.

In the case of GME, it is almost obvious in retrospect that the very act of investing can produce a positive $\alpha$, at least for a short period of time.[45] However, the same strategy may not work as well for other stocks. In general, all stocks can be affected by such price-impact investors in theory, but the degree to which each of them is susceptible depends on a number of factors, including its market capitalization, liquidity, price dynamics, main shareholders, amount of short interest, sentiment and attention from the general public, and so on. Moreover, manipulating the prices of publicly traded equities clearly violates both securities law and antitrust regulation.[46] Hence, there are significant ethical and legal ramifications of this type of price-impact investing that have yet to be fully explored. Nevertheless, our impact framework provides a means to measure the magnitude of such investments, which could be an important component of policy debates on whether and how to regulate this type of activity. Because this type of price-impact investing is very different from the other impact investments, in Online Appendix A.7, we demonstrate how to apply well-known market-microstructure models, such as Bertsimas and Lo (1998), to quantify the relation between short-term trading programs and market price reactions.

In practice, it is difficult to accurately calibrate the relevant parameters for each stock; hence, the expected profit of engaging in GME-like price-impact investing is correspondingly difficult to estimate. However, the fundamental determinant for a price-impact investor's $\alpha$ is the correlation between each stock's trading profit and its susceptibility to price manipulations as a function of stock characteristics (e.g., market capitalization, liquidity, and short interest). In fact, stocks like GME, AMC Entertainment Holdings, and Blackberry were the perfect targets for the short squeezes that occurred at the end of 2020 to early 2021 because of their highly publicized amounts of short interest from hedge funds and high customer concentration in the young people, both of which are arguably correlated features with short-squeeze profits.

## 7. Discussion

Our framework for assessing the financial consequences of impact investing has several limitations and potential extensions that we review in this section.

### 7.1. Impact Measurement Error

It should be emphasized that measurement errors can exist in the impact variable **X** itself, just like any other stock-level characteristic. This is especially challenging for emerging concepts, such as ESG, as documented recently by Berg et al. (2022a). One potential approach to dealing with the noise in ESG measurement is to combine multiple sources of data and get an improved signal with instrumented variable regression (Berg et al. 2022b) or aggregation methods designed for denoising (Berg et al. 2022c). Nonetheless, empirical estimates of ESG returns should be taken with the caveat that they are potentially affected by the measurement noise.

### 7.2. Asset-Pricing Models, Impact Factor Selection, and Economic Foundations

The examples in Section 6 underscore the fact that any quantitative assessment of impact investing depends critically on an asset-pricing model. In the same way that investment performance attribution requires a benchmark from which to measure deviations, our measures are always relative to a given benchmark asset-pricing model, such as those in (43)–(49). Therefore, to determine whether returns are attributable to a given impact variable or any other omitted variable, one can either include the proposed omitted variable in the factor model at the outset or directly test the omitted variable as $\mathbf{X}$ using our framework. In either case, a concrete alternative hypothesis is required as in any test of a given asset-pricing model, such as the CAPM (Sharpe 1964, Lintner 1965), the APT (Ross 1976), or the Fama–French multifactor models (Fama and French 1993, 2015).[47] In our example of sin stocks, Blitz and Fabozzi (2017) show that the initial excess returns documented by Hong and Kacperczyk (2009) reduce significantly after accounting for two concrete omitted quality factors—profitability and investment.

More generally, the selection of impact variables $\mathbf{X}$ should be on the basis of a priori economic, institutional, scientific, social, and/or other rationales, which not only helps to establish relationships that are more likely to be robust and causal but also mitigates the measurement errors discussed. In the case of biotech venture philanthropy, curing the disease is a prerequisite and therefore, likely a causal impact factor for financial performance. The excess returns are driven by well-documented challenges in early-stage drug development programs, such as the low probabilities of success, long time horizons, and large capital requirements (Fagnan et al. 2013). In the Sematech example, there exist well-documented economic channels through which the superior financial performance can be realized, which is reduced R&D duplication and increased profitability per R&D dollar for member firms (Irwin and Klenow 1996a). Our difference-in-difference approach also adds additional confidence to the causal nature of this impact factor. For sin stocks, the excess returns can be explained by the Merton (1987) model of neglected stocks and segmented markets or the Fama and French (2007) taste model, which predicts that nonsocially responsible companies that are out of favor by investors will earn higher expected returns. For ESG investing, although theories suggest that green assets should not earn higher returns in equilibrium, they can have prolonged periods of better returns when aggregate tastes are changing (Pástor et al. 2021). Our framework is flexible enough to be calibrated to asset-pricing models that account for the underlying financial market variables driving ESG returns, such as good versus bad economic times (Bansal et al. 2022) and market-wide shifts in attention to climate change (Lo et al. 2022a, Pástor et al. 2022). Finally, the causal mechanism for the meme stock phenomenon is given by the price-impact model in Online Appendix A.7 or the Pedersen (2022) model of influencers and thought leaders in social networks.

### 7.3. Nonstationarities and Estimating $\rho$

The specific motivation underlying each impact investment also plays an important role in determining the key parameter $\rho$, the correlation between the impact factor $\mathbf{X}$ and returns. In particular, such motivation is often nonstationary, which implies that $\rho$ may be time varying. The dynamics of such time variation are related to changes in the most pressing issues facing society at each point in history. For example, the idea that portfolio managers should include company-specific carbon risk exposures in their investment process was greeted with skepticism in a not-so-distant conference in 2010 (Andersson et al. 2016), in contrast to today's 5,301 United Nations Principles for Responsible Investment (UNPRI) signatories.[48] In this adaptive process (Lo 2004, 2017), the correlation between a climate impact factor and returns rises as the amount of assets under management and the number of products that are attempting to take advantage of a given $\mathbf{X}$ increase over time and eventually stabilizes as the size of the new sector reaches a steady state. Even in the absence of any direct physical relationship between a company's carbon emissions and its business prospects, if enough investors care about the company's carbon footprint because of general environmental concerns, this factor can have an impact on the company's returns, thereby inducing a risk premium and consequently, a nonzero $\rho$.

Because impact investing is so often associated with nonpecuniary preferences of investors and because behavioral economists and other social scientists have documented the fact that individual preferences change over time and across contexts, nonstationarity is especially relevant for such investments. Therefore, attempting to estimate $\rho$ using decades of historical data is unlikely to be fruitful given the shift in these correlations over time.[49] A more useful approach is either to estimate $\rho$ using a structural model based on the specific impact measure or to apply time series methods that are more robust to nonstationarities, such as rolling-window estimators, regime-switching models, or machine-learning techniques, which are often more adaptive than standard statistical estimators.

A separate but closely related issue involves the time horizon over which correlation is measured. In our theoretical analysis, $\rho$ is assumed to be a fixed constant, but in practice, $\rho$ is not only time varying but also horizon dependent. For some types of impact, the horizon is short, such as the price-impact example of Section 6.5, so correlations can be measured using daily or monthly returns. However, for other types of impact, the horizon can be multiple years, such as the case with the Cystic Fibrosis Foundation (Section 6.1). These two extremes reflect the specific mechanisms of impact. In the case of meme stocks, impact occurs trade by trade, whereas in the case of biomedical innovation, impact occurs as drug candidates reach clinical-trial milestones, which can take over a decade. Therefore, when applying the framework of Section 3 to a given context, a necessary prerequisite is a clear understanding of the nature of the impact being sought and over the specific time span that it is likely to be observed. Additionally, when disclosing the financial impact of such an investment, it would be prudent to disclose the estimated correlation $\rho$ over multiple horizons in the same way that retail investors are now given one-, three-, and five-year historical returns whenever available to evaluate current and potential future investments. Correlations over different horizons can reveal important nonstationarities as well as other structural and economic features about the impact that investors are hoping to obtain, and they can also help to set expectations as to when such impact might be realized.

## 8. Conclusion

In this article, we propose a new framework to quantify the financial value added/subtracted of impact investing. Using the theory of induced order statistics, we show that the correlation between the impact factor, $\mathbf{X}$, and the excess returns of individual securities determines the excess return of the impact portfolio. The impact factor provides a ranking and selection mechanism for portfolio construction, and its correlation with $\boldsymbol{\alpha}$ provides additional information that can be used to achieve better risk-adjusted returns as well as impact.

In practice, we require estimates of $\boldsymbol{\alpha}$ to measure the correlation between $\mathbf{X}$ and $\boldsymbol{\alpha}$, which is demonstrated empirically in two projects involving a wide range of ESG metrics (Berg et al. 2022c) and environmental metrics, such as carbon emission (Lo et al. 2022a). Then, why not just estimate $\boldsymbol{\alpha}$ and stop there? The reality is that not all investors have access to good estimates of $\boldsymbol{\alpha}$, not to mention the alphas of more sophisticated impact portfolios. Our framework provides a much lower (one-)dimensional quantity ($\rho$), compared with $\boldsymbol{\alpha}$, to be estimated either based on historical data or economic theories. It also provides a simple and unified quantity that asset managers can disclose to investors. In this sense, this correlation is analogous to the CAPM's market beta, which reduces to the correlation between individual security returns and market returns when they are both standardized to have unit variances.

The ability to quantify the distribution of alphas for impact-sorted securities allows us to form Treynor–Black portfolios to exploit the alphas optimally. This is particularly relevant for the investment management industry as it strives to bridge the gap between traditional investment products and the growing demand for impact investments. Regardless of the nature of the desired impact—whether it is biomedical innovation, promoting ESG, avoiding socially unsavory businesses, or attempting to achieve certain price objectives—our framework can be used to construct the most efficient way of investing in impact portfolios. Additionally, by comparing the properties of impact portfolios on the Treynor–Black superefficient frontier with those of nonimpact investments after accounting for investor-specific constraints, we have a concrete metric of the reward (or cost) of impact investing as demonstrated in the five examples.

In fact, an investment's alpha can itself be influenced by its impact as demonstrated in our example of venture philanthropy. If the Cystic Fibrosis Foundation was not able to achieve the impact to develop effective drugs for cystic fibrosis, it is unlikely that they could have generated any meaningful return. In this sense, there is an endogenous and likely a *causal* relationship between $\mathbf{X}$ and alpha. Another example of realizing alpha by achieving impact is activist investing, for which it has been empirically documented that activists may help their portfolio companies improve production efficiency (Brav et al. 2015), long-term fundamentals (Bebchuk et al. 2015), and stock performance (Dimson et al. 2015).

More broadly, our framework is relevant not just to impact proxies, such as SRI and ESG metrics, but applies to any characteristics that may be correlated with excess returns. This includes traditional factors, such as value, quality, size, and momentum, as well as hundreds of new factors and anomalies in the "Factor Zoo" discussed in the recent literature (Harvey et al. 2016, Feng et al. 2020, Hou et al. 2020). From this perspective, our model has defined a measure for the alpha of any factor.

Our framework may also help inform regulators and policy makers on the most appropriate tools to encourage investments with more socially-aware goals. Not all types of impact investing are created equal. When they create positive excess returns, one must understand what drives the initial undervaluation in the first place and what risks are preventing investors from participating in these opportunities. In the case of venture philanthropy in biomedical research and development, for example, it is crucial to develop new tools to mitigate risks from low probabilities of success, long time horizons, and large capital requirements (Fagnan et al. 2013, Thakor et al. 2017).

On the other hand, when impact investing incurs a cost to investors, at the very least it suggests the need for more explicit investor disclosures. It may also justify certain incentives and industrial policies, such as tax benefits and R&D grants to encourage the growth of these socially beneficial firms and organizations. One case in point is the area of green energy where for example, Baker et al. (2022) document a lower yield for green bonds compared with otherwise equivalent bonds. Governments around the world are designing policies to help grow industries, such as clean energy and electric vehicles. Even if they incur a cost in the short to medium term, as a society we need to invest in them if we value greater sustainability.

Indeed, our analysis underscores the fact that finance need not be a zero-sum game. Although impact investing does imply sacrificing excess returns in certain situations, in other situations it is, in fact, possible to achieve impact and attractive financial returns at the same time. We hope to apply our framework more broadly not only to allow portfolio managers to satisfy their fiduciary duties but also, to achieve the more ambitious goal of doing well by doing good on behalf of investors.

## Acknowledgments

The authors thank Agostino Capponi (editor), Mike Chen (discussant), Nan Chen, Ron Kahn, Mark Kritzman, Bob Merton, Jonathan Parker, Lukasz Pomorski, Krishna Ramaswamy, Luke Taylor (discussant), Lintong Wu, Haoxiang Zhu, and two anonymous reviewers for very helpful comments and discussion; Chaoyi Zhao for assistance in obtaining CRSP data for U.S. semiconductor companies; and Jayna Cummings for editorial assistance. The authors also thank seminar and conference participants at the 2022 American Finance Association (AFA) Annual Meetings (Boston), the 2022 Q Group Spring Seminar, the 2022 China International Conference in Finance, the 2022 INFORMS Annual Meetings, the 2022 Yale Initiative on Sustainable Finance Annual Symposium, the 2021 *Journal of Investment Management* Conference on Climate Change, the Seventh Annual Volatility Institute Conference at New York University Shanghai, the 2021 China Society for Industrial and Applied Mathematics Annual Conference, the International Association for Quantitative Finance seminar, the Guanghua School of Management at Peking University, Cardiff University, Chinese University of Hong Kong, China Europe International Business School, the University of California Berkeley, Blackrock, Vanguard, and the International Centre for Pension Management. The views and opinions expressed in this article are those of the authors only and do not necessarily represent the views and opinions of any other organizations, any of their affiliates or employees, or any of the individuals acknowledged.

## Endnotes

[1] Impact investing has been used to refer specifically to these investments (Barber et al. 2021), but we use the term "impact investing" more broadly in this article.

[2] See https://www.texasattorneygeneral.gov/sites/default/files/images/executive-management/BlackRock%20Letter.pdf (accessed December 15, 2022).

[3] See 9 U.S. Code §1109—Liability for breach of fiduciary duty.

[4] For standardized returns with unit variances, the market beta is simply the correlation between security returns and market returns.

[5] There is a substantial literature documenting the divergence of ESG ratings for the same firms (Dortfleitner et al. 2015, Semenova and Hassel 2015, Berg et al. 2022a).

[6] See, for example, the "luxury-good effect" of Bansal et al. (2022).

[7] See, for example, Geczy et al. (2021) and Madhavan et al. (2021).

[8] Other theoretical work on sustainable investing includes Heinkel et al. (2001), Adler and Kritzman (2008), Friedman and Heinle (2016), Luo and Balvers (2017), Albuquerque et al. (2019), Chen and Mussalli (2020), Berk and van Binsbergen (2021), Goldstein et al. (2021), Idzorek et al. (2021), Sorensen et al. (2021), and Zerbib (2022).

[9] For example, Bebchuk et al. (2013) document the disappearance of a return premium associated with highly rated corporate governance during an earlier period, and Bansal et al. (2022) find a "luxury-good effect," both of which suggest time-dependent performance of stocks as a function of their ESG ratings.

[10] In their framework, the standard mean-variance tangency portfolio has the highest Sharpe ratio among all portfolios, and restricting portfolios to have any ESG score other than that of the tangency portfolio must yield a lower Sharpe ratio (Pedersen et al. 2021, p. 573).

[11] See, for example, Chen et al. (1986), Fama and French (1993, 2015), Jegadeesh and Titman (1993), Jagannathan and Wang (1996), Lettau and Ludvigson (2001), Pástor and Stambaugh (2003), Yogo (2006), Adrian et al. (2014), Hou et al. (2015), and He et al. (2017).

[12] See, for example, Fama and MacBeth (1973), Ferson and Harvey (1991), Shanken (1992), Lewellen et al. (2010), Connor et al. (2012), Bai and Zhou (2015), Gagliardini et al. (2016, 2019), Feng et al. (2020), Gu et al. (2020), and Raponi et al. (2020).

[13] Proofs of all propositions are provided in the online appendix.

[14] In this baseline portfolio selection problem, investors do not have information on the impact of individual securities or $\mathbf{X}$ defined in the next section.

[15] An alternative interpretation adopted by, for example, Adler and Kritzman (2008) is that certain investors have skills that yield alphas private to themselves.

[16] By private information, we mean that the alphas are assumed to be unobservable by investors. Without the constraint, $\mathcal{S}$, investors have no way to select securities with positive alphas. In this sense, the constraint is, in fact, a mechanism for alpha selection and therefore, contains valuable information.

[17] The term was coined by Bhattacharya (1974) to distinguish between random variables ranked by their own realized values versus random variables ranked by the realizations of related random variables. These indirectly ranked statistics are also referred to as *concomitants* of the order statistic, $X_{i:N}$ (David 1973). Lo and MacKinlay (1990) applied these same statistical tools to quantify data-snooping biases in testing financial asset-pricing models.

[18] Pástor and Stambaugh (1999) also model this randomness as a measure of mispricing uncertainty.

[19] In fact, Grossman and Stiglitz (1980) have argued that the presence of occasional mispricings is a prerequisite for achieving informationally efficient markets; otherwise, no one has any incentive to gather information and incorporate it into market prices.

[20] See also Grinold and Kahn (1999).

[21] When no alpha or mispricing estimates are available, the optimal portfolio is given by the traditional mean-variance analysis.

[22] See also Grinold and Kahn (1999, chapter 6, 2019, chapters 4 and 5) for recent developments.

[23] In the Grinold (1989) framework, breadth is defined as the number of independent bets of a strategy in a given year. In our context, breadth is determined by the number of available assets in the universe. Our framework can also be regarded as a generalization of the FLAM because the impact portfolio only uses the rank information in **X**.

[24] This is an innocuous assumption, and we show later via simulation that cross-sectional heterogeneity in idiosyncratic volatilities does not affect our conclusions.

[25] We set the individual security weights of the impact portfolios to be always positive by convention from the perspective of long-only investors.

[26] For example, using several factor models, Pástor and Stambaugh (1999) estimate $\sigma_\alpha$ to be between 0% and 10% with a Bayesian framework (depending on different priors) for 1,994 stocks.

[27] In particular, investors can construct portfolios by maximizing the expected utility that adds a term for portfolio ESG on top of traditional mean variance of security returns. Alternatively, they can maximize the expected utility subject to a constraint that the portfolio ESG is higher than a preset threshold. See, for example, Pástor et al. (2021) and Pedersen et al. (2021).

[28] We thank the associate editor for pointing out this interesting connection.

[29] The correspondence between return factor models and SDF factor models is initially because of Ross (1978) and Dybvig and Ingersoll (1982). The example given by Cochrane (2009, section 6.3) illustrates this connection in its simplest form when asset returns are written in excess of the risk-free rate, which we denote as $R^e$ following Cochrane (2009). Given a return factor model $E[R^e] = \beta' \mathrm{E}[\lambda]$, one can always find an SDF, $m$, such that $m$ satisfies $m = 1 + b'\lambda$ and $0 = \mathrm{E}[mR^e]$, where $b$ is given by the equation $\mathrm{E}[\lambda] = -\mathrm{Var}[\lambda]b$. Nonetheless, $m$ is not unique because one can always add to $m$ any random variable orthogonal to returns.

[30] See https://fred.stlouisfed.org/series/GS5/.

[31] Using a Bayesian framework, Pástor and Stambaugh (1999) estimate the posterior $\sigma_\alpha$ to be between 0% and 10% (depending on different priors) for 1,994 stocks using several factor models. Given that private biopharma companies likely have larger mispricings than the average stock, we assume that the cross-sectional standard deviation of $\boldsymbol{\alpha}$ is $\sigma_\alpha = 10\%$, and the CF Foundation's investment in Vertex ranks at the top 1% of $N = 10{,}000$ securities based on a rare disease impact factor. If we assume, instead, that $\sigma_\alpha = 20\%$, the implied correlation range is $[18\%, 43\%]$.

[32] Sematech consisted of only U.S. companies during its initial years. In the late 1990s, members from Asia and Europe start to join Sematech in a limited capacity, and Sematech completed its first year of operations as a unified global consortium in 2000. Therefore, we choose our sample period to be 1975–1999, which covers roughly the 12 years before and after the formation of Sematech. Returns are winsorized at 15% each side to reduce the impact of outliers.

[33] We require at least six valid monthly returns both before and after the formation of Sematech for each firm to be included. The 11 Sematech members are AT&T Microelectronics, Advanced Micro Devices, International Business Machines, Hewlett-Packard, Intel, LSI Logic, Micron Technology, Motorola, National Semiconductor, Rockwell International, and Texas Instruments.

[34] We assume that the cross-sectional standard deviation of $\boldsymbol{\alpha}$ is $\sigma_\alpha = 20\%$ and that the universe of firms in the semiconductor industry is approximately 200. If we assume, instead, that $\sigma_\alpha = 10\%$, the average implied correlation is 51%, with a range of $[2.8\%, 99\%]$.

[35] Hong and Kacperczyk (2009) report 193 sin stocks in their selection, and Blitz and Fabozzi (2017) report that sin stocks are about 2.5% of the universe. We calibrate to these parameters when determining the quantiles of the induced order statistics in, for example, Equation (A.2) in the online appendix.

[36] We use an intermediate value based on the Pástor and Stambaugh (1999) estimate of $\sigma_\alpha$ (between 0% and 10% for 1,994 stocks). Results for other values of $\sigma_\alpha$ follow trivially. Note that different values for $\sigma_\alpha$ lead to different estimates of $\rho$ but not the final estimates of the cost of avoiding sin stocks because the expected alpha in Equation (A.2) in the online appendix, for example, is invariant of the product of the two ($\rho\sigma_\alpha$).

[37] Blitz and Fabozzi (2017) use monthly returns. For U.S. data in 1963–2016, the authors report a nonsignificant monthly excess return of 0.10%. This number becomes negative when restricted to data after 1990.

[38] In total, 2,083 green U.S. municipal bonds are used in the sample compared with 643,299 ordinary bonds.

[39] We assume the standard deviation of cross-sectional alpha is $\sigma_\alpha = 1\%$ because of smaller magnitudes for bond returns. This corresponds to a lower value in the range of estimates by Pástor and Stambaugh (1999) (between 0% and 10% in the stock market). Similar to our results for sin stocks, different values for $\sigma_\alpha$ lead to different estimates of $\rho$ but not the final estimates of the ESG alpha.

[40] This indicates stocks ranked by the greenness scores in the top third minus those in the bottom third.

[41] In particular, this estimate corresponds to a portfolio constructed from ranking stocks based on the negative value of the logarithm of scope 1 carbon emissions.

[42] Note that these prices are unadjusted for the subsequent four-to-one stock split that GME underwent on July 22, 2022.

[43] In fact, GME's revenue has been declining since 2017, and its earnings per share has been negative since 2018.

[44] For example, on July 6, 2021, GameStop announced an expansion of its fulfillment center through the leasing of a 530,000-square foot facility in Reno, Nevada. By September 2021, the company hired an additional 500 employees, and in November 2021, the company secured a $500 million global asset-based revolving credit facility, which was oversubscribed and which replaced its existing $420 million facility due in November 2022 (in addition to delivering enhanced liquidity, the new facility offered reduced borrowing costs, lighter covenants, and more flexibility).

[45] What constitutes a "short" period of time is clearly subjective and context dependent; as of December 2, 2022, GME's closing price was $27.52, more than 10 times higher than the $2.62 closing price on October 30, 2020 (note that these are split-adjusted prices).

[46] See the Securities Exchange Act of 1934, Section 9; the Sherman Act; and the Commodity Exchange Act.

[47] It is known that testing an asset-pricing model against an unspecified alternative results in poor power (see, for example, MacKinlay 1987, Campbell et al. 1997, p. 261). From the Bayesian perspective, we need at least two models to update their posteriors with data. Without a concrete alternative model, the likelihood $P(\text{data} \mid \text{not } H_0)$ is ill defined.

[48] See https://www.unpri.org/signatories/signatory-resources/signatory-directory (accessed December 2, 2022).

[49] For example, Lo et al. (2022a) and Pástor et al. (2022) show that it is important to account for shocks in preference shifts in order to get correlations that are more likely to persist in the future.

## References

Adler T, Kritzman M (2008) The cost of socially responsible investing. *J. Portfolio Management* 35(1):52–56.

Adrian T, Etula E, Muir T (2014) Financial intermediaries and the cross-section of asset returns. *J. Finance* 69(6):2557–2596.

Albuquerque R, Koskinen Y, Zhang C (2019) Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Sci.* 65(10):4451–4469.

Andersson M, Bolton P, Samama F (2016) Governance and climate change: A success story in mobilizing investor support for corporate responses to climate change. *J. Appl. Corporate Finance* 28(2):29–33.

Ardia D, Bluteau K, Boudt K, Inghelbrecht K (2022) Climate change concerns and the performance of green vs. brown stocks. *Management Sci.*, ePub ahead of print December 16, https://doi.org/10.1287/mnsc.2022.4636.

Bai J, Zhou G (2015) Fama–Macbeth two-pass regressions: Improving risk premia estimates. *Finance Res. Lett.* 15:31–40.

Baker M, Bergstresser D, Serafeim G, Wurgler J (2022) The pricing and ownership of US green bonds. *Annual Rev. Financial Econom.* 14:415–437.

Bansal R, Wu D, Yaron A (2022) Socially responsible investing in good and bad times. *Rev. Financial Stud.* 35(4):2067–2099.

Barber BM, Morse A, Yasuda A (2021) Impact investing. *J. Financial Econom.* 139(1):162–185.

Bebchuk LA, Brav A, Jiang W (2015) The long-term effects of hedge fund activism. *Columbia Law Rev.* 115:1085–1156.

Bebchuk LA, Cohen A, Wang CC (2013) Learning and the disappearing association between governance and returns. *J. Financial Econom.* 108(2):323–348.

Berg F, Kölbel JF, Rigobon R (2022a) Aggregate confusion: The divergence of ESG ratings. *Rev. Finance* 26(6):1315–1344.

Berg F, Koelbel JF, Pavlova A, Rigobon R (2022b) ESG confusion and stock returns: Tackling the problem of noise. Preprint, submitted October 12, http://dx.doi.org/10.2139/ssrn.3941514.

Berg F, Lo AW, Rigobon R, Singh M, Zhang R (2022c) Quantifying the impact of impact investing: An empirical analysis with six ESG metrics. MIT working paper, MIT, Cambridge, MA.

Berk J, van Binsbergen JH (2021) The impact of impact investing. Preprint, submitted August 21, http://dx.doi.org/10.2139/ssrn.3909166.

Bertsimas D, Lo AW (1998) Optimal control of execution costs. *J. Financial Marketing* 1(1):1–50.

Bhattacharya P (1974) Convergence of sample paths of normalized sums of induced order statistics. *Ann. Statist.* 2(5):1034–1039.

Blitz D, Fabozzi FJ (2017) Sin stocks revisited: Resolving the sin stock anomaly. *J. Portfolio Management* 44(1):105–111.

Bolton P, Kacperczyk M (2021) Do investors care about carbon risk? *J. Financial Econom.* 142(2):517–549.

Brav A, Jiang W, Kim H (2015) The real effects of hedge fund activism: Productivity, asset allocation, and labor outcomes. *Rev. Financial Stud.* 28(10):2723–2769.

Brinson GP, Hood LR, Beebower GL (1986) Determinants of portfolio performance. *Financial Anal. J.* 42(4):39–44.

Brodie J, Daubechies I, De Mol C, Giannone D, Loris I (2009) Sparse and stable Markowitz portfolios. *Proc. Natl. Acad. Sci. USA* 106(30):12267–12272.

Campbell JY, Lo AW, MacKinlay AC (1997) *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).

Chen M, Mussalli G (2020) An integrated approach to quantitative ESG investing. *J. Portfolio Management* 46(3):65–74.

Chen N-F, Roll R, Ross SA (1986) Economic forces and the stock market. *J. Bus.* 59:383–403.

Cochrane J (2009) *Asset Pricing*, revised ed. (Princeton University Press, Princeton, NJ).

Connor G, Hagmann M, Linton O (2012) Efficient semiparametric estimation of the Fama–French model and extensions. *Econometrica* 80(2):713–754.

Cornell B, Damodaran A (2020) Valuing ESG: Doing good or sounding good? *J. Impact ESG Investing* 1(1):76–93.

David H (1973) Concomitants of order statistics. *Bull. Internat. Statist. Inst.* 45(1):295–300.

David HA, Nagaraja HN (2004) *Order Statistics* (John Wiley & Sons, Hoboken, NJ).

Dimson E, Karakaş O, Li X (2015) Active ownership. *Rev. Financial Stud.* 28(12):3225–3268.

Dortfleitner G, Halbritter G, Nguyen M (2015) Measuring the level and risk of corporate responsibility—An empirical comparison of different ESG ratings approaches. *J. Asset Management* 17(7):450–466.

Dybvig PH, Ingersoll JE (1982) Mean-variance theory in complete markets. *J. Bus.* 55:233–251.

Fabozzi FJ, Ma K, Oliphant BJ (2008) Sin stock returns. *J. Portfolio Management* 35(1):82–94.

Fagnan DE, Fernandez JM, Lo AW, Stein RM (2013) Can financial engineering cure cancer? *Amer. Econom. Rev.* 103(3):406–411.

Fama EF, French KR (1993) Common risk factors in the returns on stocks and bonds. *J. Financial Econom.* 33(1):3–56.

Fama EF, French KR (2007) Disagreement, tastes, and asset prices. *J. Financial Econom.* 83(3):667–689.

Fama EF, French KR (2015) A five-factor asset pricing model. *J. Financial Econom.* 116(1):1–22.

Fama EF, MacBeth JD (1973) Risk, return, and equilibrium: Empirical tests. *J. Political Econom.* 81(3):607–636.

Feng G, Giglio S, Xiu D (2020) Taming the factor zoo: A test of new factors. *J. Finance* 75(3):1327–1370.

Ferson WE, Harvey CR (1991) The variation of economic risk premiums. *J. Political Econom.* 99(2):385–415.

Friedman HL, Heinle MS (2016) Taste, information, and asset prices: Implications for the valuation of CSR. *Rev. Accounting Stud.* 21(3):740–767.

Gagliardini P, Ossola E, Scaillet O (2016) Time-varying risk premium in large cross-sectional equity data sets. *Econometrica* 84(3):985–1046.

Gagliardini P, Ossola E, Scaillet O (2019) A diagnostic criterion for approximate factor structure. *J. Econometrics* 212(2):503–521.

Geczy CC, Stambaugh RF, Levin D (2021) Investing in socially responsible mutual funds. *Rev. Asset Pricing Stud.* 11(2):309–351.

Goldstein I, Kopytov A, Shen L, Xiang H (2021) On ESG investing: Heterogeneous preferences, information, and asset prices. Preprint, submitted April 3, http://dx.doi.org/10.2139/ssrn.3823042.

Grinold RC (1989) The fundamental law of active management. *J. Portfolio Management* 15(3):30–37.

Grinold RC (1994) Alpha is volatility times IC times score. *J. Portfolio Management* 20(4):9–126.

Grinold RC, Kahn RN (1999) *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*, 2nd ed. (McGraw-Hill Education, New York).

Grinold RC, Kahn RN (2019) *Advances in Active Portfolio Management: New Developments in Quantitative Investing* (McGraw Hill Professional, New York).

Grossman SJ, Stiglitz JE (1980) On the impossibility of informationally efficient markets. *Amer. Econom. Rev.* 70(3):393–408.

Gu S, Kelly B, Xiu D (2020) Empirical asset pricing via machine learning. *Rev. Financial Stud.* 33(5):2223–2273.

Harvey CR, Liu Y, Zhu H (2016) … and the cross-section of expected returns. *Rev. Financial Stud.* 29(1):5–68.

He Z, Kelly B, Manela A (2017) Intermediary asset pricing: New evidence from many asset classes. *J. Financial Econom.* 126(1):1–35.

Heinkel R, Kraus A, Zechner J (2001) The effect of green investment on corporate behavior. *J. Financial Quant. Anal.* 36(4):431–449.

Hong H, Kacperczyk M (2009) The price of sin: The effects of social norms on markets. *J. Financial Econom.* 93(1):15–36.

Hou K, Xue C, Zhang L (2015) Digesting anomalies: An investment approach. *Rev. Financial Stud.* 28(3):650–705.

Hou K, Xue C, Zhang L (2020) Replicating anomalies. *Rev. Financial Stud.* 33(5):2019–2133.

Idzorek TM, Kaplan PD, Ibbotson RG (2021) The popularity asset pricing model. Preprint, submitted October 25, http://dx.doi.org/10.2139/ssrn.3451554.

Irwin DA, Klenow PJ (1996a) High-tech R&D subsidies estimating the effects of Sematech. *J. Internat. Econom.* 40(3–4):323–344.

Irwin DA, Klenow PJ (1996b) Sematech: Purpose and performance. *Proc. Natl. Acad. Sci. USA* 93(23):12739–12742.

Jagannathan R, Wang Z (1996) The conditional CAPM and the cross-section of expected returns. *J. Finance* 51(1):3–53.

Jegadeesh N, Titman S (1993) Returns to buying winners and selling losers: Implications for stock market efficiency. *J. Finance* 48(1):65–91.

Kim E, Lo AW (2019) Venture philanthropy: A case study of the cystic fibrosis foundation. Preprint, submitted April 23, http://dx.doi.org/10.2139/ssrn.3376673.

Kuntz L-C (2020) Beta dispersion and market timing. *J. Empirical Finance* 59:235–256.

Lettau M, Ludvigson S (2001) Consumption, aggregate wealth, and expected stock returns. *J. Finance* 56(3):815–849.

Letts CW, Ryan W, Grossman A (1997) Virtuous capital: What foundations can learn from venture capitalists. *Harvard Bus. Rev.* 75:36–50.

Lewellen J, Nagel S, Shanken J (2010) A skeptical appraisal of asset pricing tests. *J. Financial Econom.* 96(2):175–194.

Link AN, Teece DJ, Finan WF (1996) Estimating the benefits from collaboration: The case of Sematech. *Rev. Indust. Organ.* 11(5):737–751.

Lintner J (1965) The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Rev. Econom. Statist.* 47(1):13–37.

Lo AW (2004) The adaptive markets hypothesis. *J. Portfolio Management* 30(5):15–29.

Lo AW (2017) *Adaptive Markets: Financial Evolution at the Speed of Thought* (Princeton University Press, Princeton, NJ).

Lo AW, MacKinlay AC (1990) Data-snooping biases in tests of financial asset pricing models. *Rev. Financial Stud.* 3(3):431–467.

Lo AW, Thakor RT (2019) Risk and reward in the orphan drug industry. *J. Portfolio Management* 45(5):30–45.

Lo AW, Zhang R, Zhao C (2022a) Measuring and optimizing the risk and reward of green portfolios. *J. Impact ESG Investing* 3(2):55–93.

Lo AW, Wu L, Zhang R, Zhao C (2022b) Optimal impact portfolios with general dependence and marginals. Preprint, submitted July 31, http://dx.doi.org/10.2139/ssrn.4177277.

Luo HA, Balvers RJ (2017) Social screens and systematic investor boycott risk. *J. Financial Quant. Anal.* 52(1):365–399.

MacKinlay AC (1987) On multivariate tests of the CAPM. *J. Financial Econom.* 18(2):341–371.

Madhavan A, Sobczyk A, Ang A (2021) Toward ESG alpha: Analyzing ESG exposures through a factor lens. *Financial Anal. J.* 77(1):69–88.

Merton R (1973) An intertemporal capital asset pricing model. *Econometrica* 41(5):867–887.

Merton RC (1987) A simple model of capital market equilibrium with incomplete information. *J. Finance* 42(3):483–510.

Pástor L, Stambaugh RF (1999) Costs of equity capital and model mispricing. *J. Finance* 54(1):67–121.

Pástor L, Stambaugh RF (2003) Liquidity risk and expected stock returns. *J. Political Econom.* 111(3):642–685.

Pástor L, Stambaugh RF, Taylor LA (2021) Sustainable investing in equilibrium. *J. Financial Econom.* 142(2):550–571.

Pástor L, Stambaugh RF, Taylor LA (2022) Dissecting green returns. *J. Financial Econom.* 146(2):403–424.

Pedersen LH (2022) Game on: Social networks and markets. *J. Financial Econom.* 146(3):1097–1119.

Pedersen LH, Fitzgibbons S, Pomorski L (2021) Responsible investing: The ESG-efficient frontier. *J. Financial Econom.* 142(2):572–597.

Raponi V, Robotti C, Zaffaroni P (2020) Testing beta-pricing models using large cross-sections. *Rev. Financial Stud.* 33(6):2796–2842.

Ross S (1976) The arbitrage theory of capital asset pricing. *J. Econom. Theory* 13(3):341–360.

Ross SA (1978) A simple approach to the valuation of risky streams. *J. Bus.* 51(3):453–475.

Semenova N, Hassel LG (2015) On the validity of environmental performance metrics. *J. Bus. Ethics* 132(2):249–258.

Shanken J (1992) On the estimation of beta-pricing models. *Rev. Financial Stud.* 5(1):1–33.

Sharp P, Hockfield S (2017) Convergence: The future of health. *Science* 355(6325):589.

Sharpe W (1964) Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance* 19(3):425–442.

Sharpe WF (1966) Mutual fund performance. *J. Bus.* 39(1):119–138.

Shiller RJ (2005) *Irrational Exuberance*, 2nd ed. (Princeton University Press, Princeton, NJ).

Sorensen E, Chen M, Mussalli G (2021) The quantitative approach for sustainable investing. *J. Portfolio Management* 47(8):38–49.

Sortino FA, Price LN (1994) Performance measurement in a downside risk framework. *J. Investing* 3(3):59–64.

Sortino FA, Van Der Meer R (1991) Downside risk. *J. Portfolio Management* 17(4):27–31.

Statman M, Glushkov D (2009) The wages of social responsibility. *Financial Anal. J.* 65(4):33–46.

Thakor RT, Anaya N, Zhang Y, Vilanilam C, Siah KW, Wong CH, Lo AW (2017) Just how good an investment is the biopharmaceutical sector? *Nature Biotechnology* 35(12):1149–1157.

Treynor JL, Black F (1973) How to use security analysis to improve portfolio selection. *J. Bus.* 46(1):66–86.

Tu J, Zhou G (2011) Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies. *J. Financial Econom.* 99(1):204–215.

Yogo M (2006) A consumption-based explanation of expected stock returns. *J. Finance* 61(2):539–580.

Zerbib OD (2022) A sustainable capital asset pricing model (S-CAPM): Evidence from environmental integration and sin stock exclusion. *Rev. Finance* 26(6):1345–1388.