Quantifying the Impact of Impact Investing^{*}

Andrew W. Lo^{\dagger} and Ruixun Zhang[‡]

First Draft: 18 October 2021 This Draft: 18 December 2022

Abstract

We propose a quantitative framework for assessing the financial impact of any form of impact investing, including socially responsible investing (SRI), environmental, social, and governance (ESG) objectives, and other non-financial 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 propose a quantitative measure for the financial reward, or cost, of impact investing compared to passive index benchmarks. We illustrate our approach with applications to biotech venture philanthropy, semiconductor R&D consortium, divesting from "sin" stocks, investing in ESG, and "meme" stock rallies such as GameStop in 2021.

Keywords: Impact Investing; Environmental, Social, and Governance Investing; Socially Responsible Investing; Venture Philanthropy; Investments.

JEL Classification: C10, C20, G11, G12

^{*}We thank Mike Chen (discussant), Nan Chen, Ron Kahn, Mark Kritzman, Bob Merton, Jonathan Parker, Lukasz Pomorski, Krishna Ramaswamy, Luke Taylor (discussant), Lintong Wu, and Haoxiang Zhu 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. We also thank seminar and conference participants at the 2022 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 NYU Shanghai, the 2021 China Society for Industrial and Applied Mathematics Annual Conference, the International Association for Quantitative Finance (IAQF) seminar, the Guanghua School of Management at PKU, Cardiff University, Chinese University of Hong Kong, UC Berkeley, Blackrock, and Vanguard. 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 above.

[†]Charles E. and Susan T. Harris Professor, MIT Sloan School of Management; Director, MIT Laboratory for Financial Engineering; Principal Investigator, MIT Computer Science and Artificial Intelligence Laboratory; External Faculty, Santa Fe Institute. alo-admin@mit.edu (email).

[‡]Assistant Professor and Boya Young Fellow, Peking University School of Mathematical Sciences, Center for Statistical Science, National Engineering Laboratory for Big Data Analysis and Applications, and Laboratory for Mathematical Economics and Quantitative Finance. zhangruixun@pku.edu.cn (email).

Contents

1	Introduction	1
2	Literature Review	6
3	The Framework 3.1 The No-Impact Baseline Case 3.2 Impact Factors and Induced Order Statistics 3.3 Defining an Impact Portfolio	9 9 11 11
4	Characterizing Excess Returns 4.1 Finite-Sample Distribution 4.2 Comparison with Conventional Order Statistics 4.3 Asymptotic Distribution 4.4 Interpreting Excess Return as Omitted Factors	 12 13 15 17 18
5	Impact Portfolio Construction5.1Properties of Arbitrary Impact Portfolios5.2Estimation of ρ and σ_{α} .5.3Treynor-Black Portfolios5.4Combining Impact and Passive Portfolios5.5Qualifications	 20 21 22 24 32
6	Applications to Five Impact Investments6.1Venture Philanthropy: The Cystic Fibrosis Foundation6.2R&D Consortium: Sematech6.3Divesting Sin Stocks6.4ESG Investing6.5The GameStop Phenomenon	 33 34 36 38 40 43
7	Discussion 7.1 Impact Measurement Error 7.2 Asset-Pricing Models, Impact Factor Selection, and Economic Foundations 7.3 Nonstationarities and Estimating ρ	45 45 46 47
8	Conclusion	48
Α	AppendixA.1Numerical Examples for the Performance of Impact PortfoliosA.2A Simple Execution Model for Meme Stock TradingA.3Proof of Proposition 1A.4Proof of Proposition 2A.5Proof of Proposition 3A.6Proof of Proposition 4A.7Proof of Proposition 5A.8Proof of Proposition 6A.9Proof of Proposition 7A.10Proof of Proposition 8	50 50 55 55 57 57 58 58 58 58 59

A.11 Proof of Proposition 9										•										59
A.12 Proof of Proposition 10		•	 •				•	 •			•			•					 •	60
A.13 Proof of Proposition A.0).	•	 •		 •	•			•	•	•		•	•			•	•	 •	60

1 Introduction

Impact investing—broadly defined as investments that consider not only financial objectives but also other 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"). Such "socially responsible" investing (SRI) initially involved imposing filters so that certain companies were excluded from investable universes (also referred to as "exclusionary investing"). 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 As of December 2022, 5,301 organizations representing over \$100 trillion in assets under management have become signatories to the United Nations Principles of Responsible Investment (UNPRI), and the first three of these principles include explicit commitments to ESG.²

The growth in popularity and assets under management of impact investing has also triggered a backlash. For example, on 4 August 2022, a letter signed by the attorneys general of nineteen states was sent to BlackRock's CEO, Laurence Fink, expressing concern over the asset manager's ESG policies and how they may impact their holdings of fossil-fuel energy companies:³ "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. Our states will not idly stand for our pensioners' retirements to be sacrificed for BlackRock's climate agenda." 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.⁴

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, which is the subject of this article.

- Principle 1: We will incorporate ESG issues into investment analysis and decision-making processes.
- Principle 2: We will be active owners and incorporate ESG issues into our ownership policies and practices.
- Principle 3: We will seek appropriate disclosure on ESG issues by the entities in which we invest.

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

²Specifically:

See https://www.unpri.org/about-us/what-are-the-principles-for-responsible-investment, accessed 15 December 2022.

³https://www.texasattorneygeneral.gov/sites/default/files/images/executive-management/BlackRock% 20Letter.pdf, accessed 15 December 2022.

⁴See 9 U.S. Code §1109 - Liability for breach of fiduciary duty.

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 non-superior risk/reward profile compared to 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 non-superiority 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 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.

But 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 multi-factor model for asset returns and define excess returns or "alpha" as non-zero 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 set-up where "alpha" is re-interpreted as omitted factors of which investors are either unaware or unable to access as easily as 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 or strategy. 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, and a large body of literature has developed many results for the distributions of various types of order statistics. 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 which, 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,⁵ this simple but profound result highlights the mechanism through which an impact factor's excess return is influenced by the induced ordering of alpha—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 "super-efficient 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

 $^{^{5}}$ For standardized returns with unit variances, the market beta is simply the correlation between security returns and market returns.

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 due to 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,⁶ different market conditions,⁷ and different asset-pricing models for alpha,⁸ 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, 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 Foundation, a leading venture philanthropy organization dedicated to treating and, eventually, curing cystic fibrosis. This example shows 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" (Falck and Heblich, 2007; Eichholtz, Kok, and Quigley, 2010). 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, ρ .

The second application involves the R&D consortium, another form of impact investing. We consider the example of 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 Merton's (1987) model of neglected stocks and segmented markets. Calibrating to Hong and Kacperczyk (2009) as an example, we estimate the cost of divestment to be 1.7%–3.3% in forgone alpha per annum, depending on the specific selection criteria. However, when calibrating to Blitz and Fabozzi (2017) where the authors estimate alpha

⁶There is a substantial literature documenting the divergence of ESG ratings for the same firms (Dortfleitner, Halbritter, and Nguyen, 2015; Semenova and Hassel, 2015; Berg, Kölbel, and Rigobon, 2022).

⁷See, for example, the "luxury-good effect" of Bansal, Wu, and Yaron (2022).

⁸See, for example, Geczy, Stambaugh, and Levin (2021) and Madhavan, Sobczyk, and Ang (2021).

by controlling for two new Fama and French (2015) quality factors—profitability and investment in addition to classic factors, we obtain a smaller but still non-trivial cost to institutional investors, 0.6%–1.3% 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, Wu, and Yaron, 2022). Our framework is sufficiently flexible to be calibrated to asset-pricing models that account for the underlying financial market variables that drive ESG returns, such as good vs. bad economic times (Bansal, Wu, and Yaron, 2022) and market-wide shifts in attention to climate change (Pástor, Stambaugh, and Taylor, 2022; Lo, Zhang, and Zhao, 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 recently developed equilibrium models of ESG investing.

Finally, we apply our framework to explain the January 2021 price spike in GameStop Corp. and other "meme" stocks such as AMC Entertainment Holdings and Blackberry, 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. Of course, manipulating the prices of publicly traded equities clearly violates both securities law and anti-trust regulation, and our analysis is not meant to condone or encourage such activities. However, measuring the magnitude of such investments and understanding their financial implications can better inform regulators and policymakers as to the scope and severity of this phenomenon so 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 avoiding any concerns that their clients are neither aware of nor willing to bear the consequences of the managers' impact criterion.

2 Literature Review

There is a growing literature theorizing the impact of SRI, ESG, and other non-financial objectives on asset pricing. Heinkel, Kraus, and Zechner (2001) build an equilibrium model in which exclusionary ethical investing leads to lower stock prices for polluting firms. Fama and French's (2007) taste model shows that, if investors prefer to invest in socially responsible companies, the expected return on such companies will be lower. Idzorek, Kaplan, and Ibbotson (2021) expand Fama and French's (2007) model into an equilibrium popularity asset pricing model (PAPM).⁹ Pástor, Stambaugh, and Taylor (2021) provide a model for ESG investing where investors' taste for green assets imply lower returns, and assets can be priced in a two-factor model that includes the ESG factor and the market portfolio. Moreover, they show that the ESG factor exists when there is a large dispersion in investors' ESG tastes.¹⁰

However, other attempts to incorporate ESG explicitly into an asset-pricing framework imply that impact investing may positively predict expected returns in certain situations. Pedersen, Fitzgibbons, and Pomorski (2021) show that when the market is populated by ESG-motivated, ESG-aware, and ESG-unaware investors, the optimal allocation satisfies four-fund separation and is characterized by an ESG-efficient frontier. In their framework, ESG may either yield benefits to expected returns because it provides information about firm fundamentals (as in our example above in which constraints contain information about returns), or incur costs because it affects investor preferences and constraints. Chen and Mussalli (2020) and Sorensen, Chen, and Mussalli (2021) expand traditional portfolio theory to incorporate sustainability considerations for practitioners.

While 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, Stambaugh, and Taylor (2021), and Pedersen, Fitzgibbons, and Pomorski (2021) highlight that the expected return of ESG investing depends on the mix of investors and preferences in the market.¹¹ But 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).¹² 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 **X** and security returns reflects the particular market condition and shift

⁹Their model implies that ESG investments can have better returns when aggregate tastes shift.

¹⁰Other theoretical work on sustainable investing includes Friedman and Heinle (2016), Luo and Balvers (2017), Albuquerque, Koskinen, and Zhang (2019), Berk and van Binsbergen (2021), Goldstein et al. (2021), and Zerbib (2022). Adler and Kritzman (2008) use a simulation framework to assess the cost of socially responsible investing.

¹¹Pedersen, Fitzgibbons, and Pomorski (2021) show that high-ESG stocks deliver high expected returns when the market has many ESG-unaware investors, and low expected returns when the market has many ESG-motivated investors. Pástor, Stambaugh, and Taylor (2021) show that dispersion in ESG preferences increases the size of ESG investments and lowers equilibrium ESG returns.

¹²For example, Bebchuk, Cohen, and Wang (2013) document the disappearance of a return premium associated with highly rated corporate governance during an earlier period, and Bansal, Wu, and Yaron (2022) find a "luxury-good effect," both of which suggest time-dependent performance of stocks as a function of their ESG ratings.

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, Fitzgibbons, and Pomorski (2021) is that ESG's information about firm fundamentals can yield benefits to its returns, while screening constraints will incur costs to ESG investing.¹³ Our model shows that, when securities have non-zero 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, in addition to the ESG-efficient frontier of Pedersen, Fitzgibbons, and Pomorski (2021), we are able to explicitly construct the optimal super-efficient portfolio from any \mathbf{X} and explicitly quantify its financial impact.

Theories of SRI and ESG investing are also accompanied by a vast empirical literature focused on measuring their returns across asset classes and regions, and how much of these returns can be explained by traditional asset pricing factors.¹⁴ On the one hand, several studies suggest that investments with ESG considerations may sacrifice returns in markets including stocks (Alessandrini and Jondeau, 2020), bonds (Baker et al., 2022), and venture capital funds (Barber, Morse, and Yasuda, 2021). This is also consistent with the literature that documents positive excess returns for sin stocks.¹⁵ Geczy, Stambaugh, and Levin (2021) show that the SRI cost to mutual funds is minimal compared to a CAPM-investor but may be substantial when investors allow for size, value, and momentum factors. Pástor, Stambaugh, and Taylor (2022) show that the high returns for green assets in recent years reflect unexpectedly strong increases in environmental concerns, not high expected returns.

On the other hand, recent empirical evidence from both academic research (Kempf and Osthoff, 2007; Madhavan, Sobczyk, and Ang, 2021; Bansal, Wu, and Yaron, 2022; Berg et al., 2022b) and industry advocates (Shing, 2021; Xiong, 2021) suggests that impact investing and, in particular, ESG measures, is associated with higher expected returns, at least under certain market conditions. Lindsey, Pruitt, and Schiller (2021) find that modifying optimal portfolio weights to achieve an ESG-investing tilt negligibly affects portfolio performance because ESG measures do not provide information about future stock performance beyond what is provided by other observable firm characteristics. This raises the possibility that impact investing need not always imply lower risk-adjusted returns.

¹³In 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, Fitzgibbons, and Pomorski, 2021, p. 573).

¹⁴See, for example, Galema, Plantinga, and Scholtens (2008), Renneboog, Ter Horst, and Zhang (2008), Blitz and Fabozzi (2017), and Madhavan, Sobczyk, and Ang (2021).

¹⁵See, for example, Fabozzi, Ma, and Oliphant (2008), Hong and Kacperczyk (2009), Statman and Glushkov (2009), and Fauver and McDonald IV (2014).

Moreover, there is substantial divergence among impact measures—such as those for ESG even when they purport to capture the same concepts (Dortfleitner, Halbritter, and Nguyen, 2015; Semenova and Hassel, 2015; Gibson, Krueger, and Schmidt, 2021; Berg, Kölbel, and Rigobon, 2022). In particular, Khan, Serafeim, and Yoon (2016) find that only firms with good ratings on *material* sustainability issues significantly outperform firms with poor ratings on these issues.

These inconsistencies in the impact investing literature 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 explains 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. Our framework provides a unified methodology to quantify the financial consequences of all forms of impact investing, including SRI and ESG.

Additional literature on SRI and ESG investing includes climate change and its impact on asset pricing (Bolton and Kacperczyk, 2021),¹⁶ preference toward sustainable investments (Bauer, Ruof, and Smeets, 2021), market responses to companies' eco-friendly behavior (Klassen and McLaughlin, 1996; Flammer, 2013, 2021; Krüger, 2015), transmission channels between ESG information and company valuation (Dunn, Fitzgibbons, and Pomorski, 2018; Giese et al., 2019), the real social impact generated by green investors (Dyck et al., 2019; Chen, Dong, and Lin, 2020), and implications for bank loans (Goss and Roberts, 2011). The empirical evidence of a causal relation between the sustainability classification and capital inflows for U.S. mutual funds (Hartzmark and Sussman, 2019) further highlights both the popularity and importance of impact investing today.

More generally, our framework is applicable to portfolios constructed on the basis of any characteristic, including both impact proxies such as ESG and SRI measures and 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. This includes a large literature devoted to identifying asset pricing factors,¹⁷ a vast econometrics literature focused on factor models,¹⁸ the literature on active portfolio management,¹⁹ and the literature on data-snooping biases and the high dimensionality of cross-sectional asset-pricing models.²⁰ 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.

¹⁹The fundamental law of active management by Grinold (1989, 1994) and Grinold and Kahn (1999, 2019).

¹⁶See also Hong, Karolyi, and Scheinkman (2020), Giglio, Kelly, and Stroebel (2021), and references therein.

¹⁷See, for example, Chen, Roll, and Ross (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, Etula, and Muir (2014), Hou, Xue, and Zhang (2015), and He, Kelly, and Manela (2017).

¹⁸See, for example, Fama and MacBeth (1973), Ferson and Harvey (1991), Shanken (1992), Lewellen, Nagel, and Shanken (2010), Connor, Hagmann, and Linton (2012), Bai and Zhou (2015), Gagliardini, Ossola, and Scaillet (2016, 2019), Gu, Kelly, and Xiu (2020), and Raponi, Robotti, and Zaffaroni (2020).

²⁰See, for example, Harvey, Liu, and Zhu (2016), Green, Hand, and Zhang (2017), Kozak, Nagel, and Santosh (2018), Feng, Giglio, and Xiu (2020), and Freyberger, Neuhierl, and Weber (2020).

3 The Framework

We consider a universe of N securities with returns R_{it} that satisfy the following linear multi-factor model:

$$R_{it} - R_{ft} = \alpha_i + \beta_{i1} \left(\Lambda_{1t} - R_{ft} \right) + \dots + \beta_{iK} \left(\Lambda_{Kt} - R_{ft} \right) + \epsilon_{it} \tag{1}$$

such that
$$\operatorname{E}[\epsilon_{it}|\Lambda_{kt}] = 0$$
, $k = 1, \dots, K$ (2)

where Λ_{kt} is the k-th factor return, k = 1, ..., K, R_{ft} is the risk-free rate, α_i and β_{ik} are the excess return and factor betas, respectively, and ϵ_{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 multi-factor model (1) is consistent with a number of asset-pricing models such as the Sharpe-Lintner Capital Asset-Pricing Model (CAPM) (Sharpe, 1964; Lintner, 1965), Merton's Intertemporal CAPM (Merton, 1973), Ross's Arbitrage Pricing Theory (APT) (Ross, 1976), and the Fama-French multi-factor models (Fama and French, 1993, 2015). In particular, all of these asset-pricing models imply that $\alpha_i = 0$, and that returns are simply the sum of the risk-free rate plus all the risk premia multiplied by the asset's corresponding risk exposures or β_{ik} 's.

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.²¹

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

$$\max_{\left\{\omega_1,\ldots,\omega_N\mid\sum_{i=1}^N\omega_i=1\right\}} \mathbb{E}[U(W)] \geq \max_{\left\{\omega_1^c,\ldots,\omega_N^c\mid\sum_{i\in\mathcal{S}}\omega_i^c=1 \text{ and } \omega_i^c=0 \text{ for } i\notin\mathcal{S}\right\}} \mathbb{E}[U(W^c)]$$
(3)

for any non-decreasing concave utility function $U(\cdot)$ where

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

²¹Proofs of all propositions are provided in the Appendix.

In addition, under certain fairly realistic conditions given in the Appendix, the loss in utility by restricting to the subset S is generally small, as long as the number of securities excluded by 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 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 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.²² In addition, although impact investing in this special case cannot improve returns, the under-performance is likely to be small (assuming no transactions costs or fees, of course).

However, suppose we allow for non-zero 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.²³ If we relax the condition that S is independent of the joint distribution of $\{R_i\}$, then Proposition 1 clearly does not hold. For example, suppose:

$$S = \{ i : \alpha_i > 0, i = 1, \dots, N \}.$$
(5)

Clearly in this case, it is possible for the risk-adjusted returns of the 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 information²⁴ 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 no impact. But 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 counter-example assumes the existence of mispriced or positive-alpha securities (5), but an equally valid equilibrium/no-arbitrage interpretation is that the α_i s are omitted factors

 $^{^{22}}$ In this baseline portfolio selection problem, investors do not have information on the impact of individual securities, or **X** defined in the next section. In other words, investors maximize their unconditional mean-variance utility, which corresponds to the "type-A (ESG-aware)" investors in Pedersen, Fitzgibbons, and Pomorski (2021).

 $^{^{23}}$ An alternative interpretation adopted by, for example, Adler and Kritzman (2008) is that certain investors have skills that yield alphas private to themselves.

²⁴By private information, we mean that the alphas are assumed to be unobservable by investors. Without the constraint, 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.

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, eventdriven 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 below—and offers a systematic and quantitative approach to measuring impact in either case.

3.2 Impact Factors and Induced Order Statistics

To measure the effects of impact investing on investment performance, we assume that the excess return of the *i*-the security, α_i , is not observable, whereas the impact factor, X_i , for that security is. Contrary to the usual asset-pricing set-up in which the α_i s are assumed to be fixed constants (and, in equilibrium or under no-arbitrage conditions, identically equal to 0), 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 re-order the bivariate vector $(X_i, \alpha_i)^T$, $i = 1, 2, \cdots, N$, according to the magnitudes of their first components:

$$\begin{pmatrix} X_{1:N} \\ \alpha_{[1:N]} \end{pmatrix}, \begin{pmatrix} X_{2:N} \\ \alpha_{[2:N]} \end{pmatrix}, \cdots, \begin{pmatrix} X_{N:N} \\ \alpha_{[N:N]} \end{pmatrix}$$
(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*,²⁵ where the order is *induced* by another variable **X**.

Because α is defined with respect to the multi-factor model (1), by studying the interaction between **X** and α , our framework allows for the existence of any set of predefined asset-pricing factors, and focuses on the incremental role of **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, **X**. For example, an investor may choose to invest in the top n_0 securities ranked by **X**, or form portfolios long the top decile and short the bottom decile. In general, we define an *impact portfolio* to be any portfolio $\mathcal{S}(\mathbf{X})$ formed as a function of the impact factor, **X**. With portfolio weights $\{\omega_i, i \in \mathcal{S}\}$, the return of the impact portfolio is given by:

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

²⁵The 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.

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 **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 multi-factor asset-pricing model (1), this corresponds to a wide range of impact investing contexts.²⁶ 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 (Harvey, Liu, and Zhu, 2016; Feng, Giglio, and Xiu, 2020; Hou, Xue, and Zhang, 2020). For the purposes of this study, we focus on the impact investing interpretation, but 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:

(A1) $(X_i, \alpha_i)^T, i = 1, 2, \dots, N$, are independently and identically distributed (IID) 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 α_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 α_i as a measure of uncertainty as to the degree of mispricings of securities in our investment universe.²⁷ 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 α_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.²⁸ Moreover, we can calibrate the degree of mispricings in our model through σ_{α}^2 —smaller values correspond to greater efficiency, and larger values correspond to lower efficiency and more active management opportunities.

²⁶For example, the equilibrium two-factor ESG model of Pástor, Stambaugh, and Taylor (2021) corresponds to a CAPM asset-pricing model with an impact factor \mathbf{X} representing ESG in our framework, and the climate model of Bolton and Kacperczyk (2021) corresponds to an impact factor \mathbf{X} representing various measures of carbon emission and a multi-factor asset-pricing model representing a suite of control variables.

²⁷Pástor and Stambaugh (1999) also model this randomness as a measure of mispricing uncertainty.

²⁸In fact, Grossman and Stiglitz (1980) have argued that the presence of occasional mispricings is a pre-requisite for achieving informationally efficient markets, otherwise, no one has any incentive to gather information and incorporate it into market prices.

However, our framework can also be interpreted from an equilibrium/no-arbitrage perspective, where non-zero α_i s are due to the presence of omitted factors that investors are either unaware of or unable to access directly. Under this interpretation, we will see below that the randomness in α_i is due to cross-sectional variability in security *i*'s omitted-factor betas. In this case, however, it is possible for μ_{α} to be non-zero 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 **X** and α 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 4.4), and general marginal distributions and dependent structures for $(X_i, \alpha_i)^T$ (see Lo et al. (2022)), at the expense of simplicity.

Regardless of the interpretation of α_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, **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 **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 (A1), the expected value of the *i*-th induced order statistic $\alpha_{[i:N]}, i = 1, 2, \dots, N$ is given by:

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

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

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

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

$$\sigma_{ij} \equiv \operatorname{Cov}\left(\alpha_{[i:N]}, \alpha_{[j:N]}\right) = \sigma_{\alpha}^2 \rho^2 \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right).$$
(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 (ρ) between **X** and individual securities' alphas, the standard deviation of individual alphas (σ_{α}), and a cross-sectionally standardized impact score (E [$X_{i:N}$]). The correlation ρ 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),²⁹ who provides a simple decomposition of alpha into the product of three terms: the information coefficient (the correlation ρ 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 $\mathbf{E}[X_{i:N}]$. Finally, our results also provide the variance and covariances of individual alphas, which is 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 $E[X_{i:N}]$ in (8), $Var(X_{i:N})$ in (9), and $Cov(X_{i:N}, X_{j:N})$ 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 *i*-th 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:

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

and

$$\operatorname{Var}\left(X_{i:N}\right) \approx \frac{p_i(1-p_i)}{N+2} {Q'_i}^2 + \frac{p_i(1-p_i)}{(N+2)^2} \left[2(1-2p_i)Q'_iQ''_i + p_i(1-p_i)\left(Q'_iQ'''_i + \frac{1}{2}{Q''_i}^2\right)\right]$$
(12)

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

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

 29 See also Grinold and Kahn (1999).

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

$$Q'_{i} = \left(\Phi^{-1}(p_{i})\right)' = \frac{1}{\phi\left(\Phi^{-1}(p_{i})\right)}$$
(14)

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

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

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

 Φ and ϕ 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 $E[X_{i:N}]$ 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.³⁰

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, ρ , between α and \mathbf{X} is positive, the expected value of the induced order statistic increases as the order *i* increases (see Figure 1a). The dispersion of the mean is larger when the correlation, ρ , or the dispersion of the unknown α , σ_{α} , is larger.

In addition, Figure 1b shows that the variances, $Var(X_{i:N})$, stay relatively constant across the ordered securities *i* and are primarily determined by ρ and σ_{α} . In fact, we will see in Section 4.3 that as the number of securities increases without bound, the variance converges to a constant across all *i*.

Finally, the covariances, $\text{Cov}(X_{i:N}, X_{j:N})$, are very close to zero except when *i* and *j* are close to 0 or 50, the two extremes. We will also see in Section 4.3 that as N increases without bound, the covariances approach zero, implying that induced order statistics are mutually independent in the limit.

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

 $[\]overline{{}^{30}\text{In fact, E}[X_{i:N}] \approx \Phi^{-1}(p_i) \text{ is a reasonable first-order approximation. For example, David and Nagaraja (2004, Sections 4.5 and 4.6) give the following bound: <math>\Phi^{-1}\left(\frac{i-1}{N}\right) \leq \mathrm{E}[X_{i:N}] \leq \Phi^{-1}\left(\frac{i}{N}\right).$



Figure 1: Mean, variance, and covariances of the induced order statistic, $\alpha_{[i:N]}$. The total number of securities is set to 50 for illustrative purposes. In (c) we set $\rho = 20\%$ and $\sigma_{\alpha} = 5\%$.

Proposition 4. Under Assumption (A1), 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:

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

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

$$\sigma_{ij} \equiv \operatorname{Cov}\left(\alpha_{[i:N]}, \alpha_{[j:N]}\right) = \rho^2 \operatorname{Cov}\left(\alpha_{i:N}, \alpha_{j:N}\right).$$
(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, ρ , is precisely the correlation between **X** and α .

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, $\alpha_{i:N}$, for a collection of N=50 securities. As the correlation, ρ , increases to 1, the expected excess return approaches the hypothetical value of sorting based on α .

This result highlights the role that induced ordering plays in distinguishing securities with



Figure 2: Comparison between the expected value of the induced order statistic, $\alpha_{[i:N]}$, and the order statistic, $\alpha_{i:N}$, with N = 50 and $\sigma_{\alpha} = 5\%$.

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 α) determines how much of the mean, variance, and covariances from a hypothetical sorting based on α can actually be achieved via the induced ordering of **X**.

4.3 Asymptotic Distribution

As the number of securities, N, increases without bound, the limiting joint distribution of the induced order statistics, $\alpha_{[i:N]}$, has been derived by Yang (1977) and does not require the normality assumption (A1), hence we can rely on this asymptotic approximation for large samples.

Proposition 5. Assuming $(X_1, \alpha_1)^T, \dots, (X_N, \alpha_N)^T$ are IID, for any sequence $1 < i_1 < \dots < i_n < N$ such that, as $N \to \infty$, $i_k/N \to \xi_k \in (0, 1)$ for $k = 1, \dots, n$, we have:

$$\lim_{N \to \infty} \mathbb{P}\left(\alpha_{[i_1:N]} < a_1, \cdots, \alpha_{[i_n:N]} < a_n\right) = \prod_{k=1}^n \mathbb{P}\left(\alpha_k < a_k | F_x(X_k) = \xi_k\right),\tag{21}$$

where $F_x(\cdot)$ is the marginal CDF of X_i .

Proposition 5 implies that the induced order statistics at distinct quantiles are asymptotically independent, consistent with the finite sample observations in Proposition 3 and Figure 1. Also, because the conditional distribution of jointly normal random vectors is still normal, we can characterize the first two moments of the induced order statistics asymptotically via the following result.

Proposition 6. Under Assumption (A1), as N increases without bound, the induced order statistics, $\alpha_{[i_k:N]}$ ($k = 1, \dots, n$), converge in distribution to independent Gaussian random variables with mean $\mu(\xi_k)$ and variance $\sigma^2(\xi_k)$, where

$$\mu(\xi_k) \equiv \rho(\sigma_\alpha/\sigma_x) \left[F_x^{-1}(\xi_k) - \mu_x \right] = \rho \sigma_\alpha \Phi^{-1}(\xi_k), \tag{22}$$

$$\sigma^2(\xi_k) \equiv \sigma_\alpha^2 (1 - \rho^2). \tag{23}$$

Note that the mean and variance here are consistent with the finite-sample results in Proposition 2 when k/N converges to ξ_k . The mean depends on the order k (shown in Figure 3), and its shape is very similar to the finite-sample case (Figure 1a). On the other hand, the variance, $\sigma^2(\xi_k)$, is a constant across all quantiles.



Figure 3: Asymptotic mean of the induced order statistic, $\alpha_{[i_k:N]}$, as $i_k/N \to \xi \in (0,1)$.

4.4 Interpreting Excess Return as Omitted Factors

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

Let security returns follow the K-factor asset-pricing model as specified in (1)–(2), but now assume there are no mispricings. However, suppose that investors only account for the first factor Λ_1 , without loss of generality, and are unaware of the remaing K-1 factors $\Lambda_2, \ldots, \Lambda_K$. We define:

$$\lambda_{ik} \equiv \beta_{ik} (\Lambda_k - R_f) \tag{24}$$

to be factor k's contribution to security i's return, for i = 1, ..., N and k = 2, ..., K, and

$$\lambda_i \equiv \sum_{k=2}^K \lambda_{ik} \tag{25}$$

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:

$$\alpha_i \equiv \mathbf{E}[\lambda_i] = \sum_{k=2}^K \beta_{ik} (\mathbf{E}[\Lambda_k] - R_f) .$$
(26)

To characterize the distribution of λ_i after ranking securities based on the impact factor **X**, we make the following assumption:

(A2) $(X_i, \lambda_i)^T, i = 1, 2, \dots, N$, are bivariate normal random vectors with their marginal distributions and paired correlations defined by:

$$\mu_x \equiv \mathrm{E}[X_i], \ \mu_\lambda \equiv \mathrm{E}[\lambda_i], \ \sigma_x^2 \equiv \mathrm{Var}(X_i), \ \sigma_\lambda^2 \equiv \mathrm{Var}(\lambda_i), \ \mathrm{and} \ \rho_{x,\lambda} \equiv \mathrm{Corr}(X_i,\lambda_i)$$
(27)

for i = 1, ..., N. In addition, for $i \neq j$, the correlations across different securities are defined by:

$$\rho_x \equiv \operatorname{Corr}(X_i, X_j), \ \rho_\lambda \equiv \operatorname{Corr}(\lambda_i, \lambda_j), \ \text{and} \ \tilde{\rho}_{x,\lambda} \equiv \operatorname{Corr}(X_i, \lambda_j).$$
(28)

Under this assumption, the cross-sectional randomness of λ_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 $\rho_{x,\lambda}$, $\tilde{\rho}_{x,\lambda}$, ρ_x , and ρ_{λ} .

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 **X**. We again assume without loss of generality that $\mu_x = 0$ and $\sigma_x = 1$, so that **X** is a standard normal random vector. But we allow for a non-zero risk premium μ_{λ} .

Proposition 7. Under Assumption (A2), define

$$\rho_{adj} \equiv \frac{\rho_{x,\lambda} - \tilde{\rho}_{x,\lambda}}{1 - \rho_x} \tag{29}$$

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

$$\mathbf{E}\left[\lambda_{[i:N]}\right] - \mu_{\lambda} = \rho_{adj}\sigma_{\lambda}\mathbf{E}\left[X_{i:N}\right]. \tag{30}$$

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

$$\operatorname{Var}\left(\lambda_{[i:N]}\right) = \sigma_{\lambda}^{2} \left(1 - \rho_{adj}^{2} + \rho_{adj}^{2} \operatorname{Var}\left(X_{i:N}\right)\right).$$
(31)

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

$$\operatorname{Cov}\left(\lambda_{[i:N]}, \lambda_{[j:N]}\right) = \sigma_{\lambda}^{2} \rho_{adj}^{2} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) + \left(\rho_{\lambda} - \rho_{x} \rho_{adj}^{2}\right).$$
(32)

Similar to our results in Proposition 2, the expected return from omitted factors in (30) is closely related to Grinold's (1994) decomposition of alpha. In particular, we provide a generalization to Grinold's (1994) results in the context of factor returns when the cross-sectional dependence is specified by Assumption (A2).

Proposition 7 characterizes the return from omitted factors for the *i*-th security ranked by \mathbf{X} . This result highlights an important implication when estimating the financial impact of impact investing. Given any definition of impact, \mathbf{X} , if the portfolio selected based on \mathbf{X} produces a non-zero excess return, \mathbf{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 \mathbf{X} , such as an "ESG factor" or a "carbon factor" (Bolton and Kacperczyk, 2021).

On the other hand, Proposition 7 also implies that, when forming a portfolio, if one uses a selection criteria that appears 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 \mathbf{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, Sobczyk, and Ang (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, followed by a special case—equal-weighted portfolios—which is also related to how to estimate ρ and σ_{α} empirically. 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 S:

$$\mathcal{S} \equiv \{i_1, i_2, \cdots, i_{n_0}\}\tag{33}$$

which is obtained from a rank-ordering of securities from the investment universe according to the impact factor, \mathbf{X} . The excess return of the portfolio is then given by:

$$\tilde{\alpha} \equiv \sum_{i \in \mathcal{S}} \omega_i \alpha_{[i:N]} \tag{34}$$

where $\{\omega_i : i \in S\}$ are arbitrary portfolio weights that sum to 1. Based on the distribution of the induced order statistics in Proposition 2, we have the following result for portfolio excess returns:

Proposition 8. Under Assumption (A1), the expected excess return of a portfolio S defined in (33) is:

$$\mathbf{E}\left[\tilde{\alpha}\right] = \sum_{i \in \mathcal{S}} \omega_i \mu_i = \rho \sigma_\alpha \sum_{i \in \mathcal{S}} \omega_i \mathbf{E}\left[X_{i:N}\right],\tag{35}$$

and the variance is:

$$\operatorname{Var}\left(\tilde{\alpha}\right) = \sum_{i \in \mathcal{S}} \omega_{i}^{2} \sigma_{i}^{2} + 2 \sum_{i < j \in \mathcal{S}} \omega_{i} \omega_{j} \sigma_{ij}$$
$$= \sigma_{\alpha}^{2} \left(1 - \rho^{2} + \rho^{2} \left(\sum_{i \in \mathcal{S}} \omega_{i}^{2} \operatorname{Var}\left(X_{i:N}\right) + 2 \sum_{i < j \in \mathcal{S}} \omega_{i} \omega_{j} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) \right) \right).$$
(36)

Proposition 8 quantifies the distribution of excess returns for any portfolio constructed according to the impact factor, **X**. Appendix A.1 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), Sortino ratio (Sortino and Van Der Meer, 1991; Sortino and Price, 1994), and information ratios (Treynor and Black, 1973); performance attribution (Brinson, Hood, and Beebower, 1986); and active portfolio management and enterprise risk management (Grinold and Kahn, 1999).

5.2 Estimation of ρ and σ_{α} .

Two key parameters that characterize the distribution of induced order statistics in Propositions 2 and 6 are ρ , the correlation between unobserved α and \mathbf{X} , and σ_{α} , the cross-sectional standard deviation of α_i . A special case of Proposition 8—equal-weighted portfolios—provides a way to estimate these parameters in practice. Consider an equal-weighted portfolio \mathcal{S} defined in (33) with portfolio weights $\omega_i = 1/n_0$. In this case, Proposition 8 implies that the expected value and variance of portfolio alphas are given by:

$$\mathbf{E}\left[\tilde{\alpha}\right] = \frac{\rho \sigma_{\alpha}}{n_0} \sum_{i \in \mathcal{S}} \mathbf{E}\left[X_{i:N}\right],\tag{37}$$

$$\operatorname{Var}\left(\tilde{\alpha}\right) = \sigma_{\alpha}^{2} \left(1 - \rho^{2} + \frac{\rho^{2}}{n_{0}^{2}} \left(\sum_{i \in \mathcal{S}} \operatorname{Var}\left(X_{i:N}\right) + 2 \sum_{i < j \in \mathcal{S}} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) \right) \right).$$
(38)

Empirical studies usually report excess returns from equal-weighted portfolios formed by ranking some stock characteristics such as the P/E ratio, book-to-value, or ESG score. As a result, the expected value and variance of the impact-portfolio alpha in (37)–(38) lead to a natural estimator of these two parameters based on historical data.

In particular, suppose one empirically measures the portfolio alpha and its variance, which can be substituted into (37)–(38) to yield a system of two equations with respect to ρ and σ_{α} , where parameters such as the number of securities in the portfolio (n_0) and the total number of securities in the universe (N) can be easily obtained. This leads, in principle, to a solution for ρ and σ_{α} .

On the other hand, if the variance in (38) is difficult to estimate empirically, one can still use (37) to calibrate $\rho \sigma_{\alpha}$, from which ρ can be solved based on assumptions or prior empirical estimates (such as Pástor and Stambaugh (1999)) for the spread in cross-sectional α .

In addition, it is worth emphasizing that the estimation of ρ depends implicitly on the frequency of historical data used to estimate impact-portfolio excess returns, $\tilde{\alpha}$. In theory, if the two terms in (37), $\tilde{\alpha}$ and σ_{α} , both scale linearly as the frequency varies, the estimates of ρ should stay invariant with respect to weekly, monthly, or annual returns. However, they may lead to different empirical estimates in practice, and therefore, the correlation estimated from this procedure should be interpreted in the same frequency space as the return data used.

We apply these methods to five empirical examples in Section $6.^{31}$

5.3 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.³² 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, Treynor and Black's (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, α_i , as its mean plus noise:

$$\alpha_i = \mu_i + \zeta_i \tag{39}$$

where $\{\zeta_i\}$ are independent random variables with zero means. We can then combine ζ_i with security *i*'s idiosyncratic error, ϵ_i . Because ζ_i and ϵ_i are independent, the combined idiosyncratic

³¹In addition, the framework in this article is used to empirically measure ρ and σ_{α} for a wide range of ESG metrics (Berg et al., 2022b) and environmental metrics including carbon emission (Lo, Zhang, and Zhao, 2022).

 $^{^{32}}$ When no alpha or mispricing estimates are available, the optimal portfolio is given by the traditional mean-variance analysis.

variance for security *i* is simply $\sigma_i^2 + \sigma(\epsilon_i)^2$, where σ_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 9. Under Assumption (A1), the Treynor-Black weight of security *i* is proportional to its expected alpha divided by its combined idiosyncratic variance:

$$\omega_i \propto \frac{\mu_i}{\sigma_i^2 + \sigma(\epsilon_i)^2}.\tag{40}$$

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 (40) can be further simplified to:

$$\omega_i \propto \frac{\rho \sigma_\alpha \Phi^{-1}(\xi_i)}{\sigma_\alpha^2 (1-\rho^2) + \sigma(\epsilon_i)^2} \propto \Phi^{-1}(\xi_i) \cdot Constant.$$
(41)

Proposition 9 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 (40), recall that the variance of the *i*-th induced order statistic, σ_i^2 , is approximately a constant when N is large (see Figure 1b and Proposition 6). The expected excess return, $\mu_i = \rho \sigma_{\alpha} \mathbb{E}[X_{i:N}]$, varies with respect to *i* only through the last term $\mathbb{E}[X_{i:N}]$. As a result, if each security's idiosyncratic volatility is the same, the Treynor-Black weights of security *i* in (40) depend only on their relative ranking in the universe of N securities, which is specified by the term $\Phi^{-1}(\xi_i)$ in (41).

For an illustrative example, consider a portfolio formed by the top n_0 securities ranked by \mathbf{X} , and let n_0 vary from 1 to 250. We assume for a moment that the idiosyncratic volatility is 15% for all securities. Figure 4 depicts the weights of this portfolio. As expected, securities that rank higher have higher weights. Based on Proposition 9, the weights in Figure 4 are determined only by the relative rank of the *i*-th security in the universe of N securities. In other words, changing the correlation, ρ , between $\boldsymbol{\alpha}$ and \mathbf{X} does not affect these weights.

The portfolio selected by ranking \mathbf{X} and applying the Treynor-Black weights in (40) 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:

$$\alpha_A = \sum_{k=1}^n \omega_k \mu_k,\tag{42}$$

$$\beta_A = \sum_{k=1}^n \omega_k \beta_k,\tag{43}$$

$$\sigma(\epsilon_A)^2 = \sum_{k=1}^n \omega_k^2 \left(\sigma_k^2 + \sigma(\epsilon_k)^2 \right).$$
(44)

These results—together with the explicit quantification of individual-security alphas in Propo-



Figure 4: Treynor-Black weights of the securities in the impact portfolio formed by top-ranking securities based on the impact factor, **X**, with (a) N=50; and (b) N=500.

sition 2 and 7, and the optimal Treynor-Black weights in Proposition 9—provide a complete characterization of the performance of optimal impact portfolios. In particular, the information ratio of the impact portfolio, defined as $E[\alpha_A]/\sigma(\epsilon_A)$, is proportional to the correlation, ρ , between unobserved alpha, α , and the impact factor, **X**. This is closely related to the fundamental law of active management (FLAM) by Grinold (1989),³³ which provides a simple approximation of the information ratio of an active portfolio by the product of information coefficient (ρ in our notation) and the *breadth* of a strategy.³⁴

Figure 5 contains the expected excess return, α_A , for two examples of the impact portfolio in a collection of N = 500 securities. Figure 5a depicts portfolios formed by selecting the top n_0 securities ranked by **X**. The expected value decreases as n_0 increases and more securities are included. Figure 5b depicts portfolios formed by dividing all securities into four quantiles based on the ordering of **X**. In both cases, Treynor-Black portfolios (solid line) achieve higher expected excess returns than the equal-weighted portfolios (dashed line).

5.4 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

 $^{^{33}{\}rm See}$ also (Grinold and Kahn, 1999, chapter 6) and (Grinold and Kahn, 2019, chapter 4 and 5) for recent developments.

³⁴In Grinold's (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. In addition, our framework can be regarded as a generalization of the initial FLAM because the impact portfolio only uses the rank information in **X**. For more discussion on impact portfolio analytics and its relationship with FLAM under more general distributional assumptions on (α, \mathbf{X}) , see Lo et al. (2022).



Figure 5: Expected excess return of the impact portfolio formed based on Treynor-Black weights, with N = 500 and $\sigma_{\alpha} = 5\%$. The expected excess returns of the corresponding equal-weighted portfolios are shown in dashed lines for comparison. (a) shows the case where the top-ranking securities are selected. (b) shows the case where all securities are divided into four segments based on ranking.

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 ω_A denote the weight of the impact portfolio, and $1 - \omega_A$ 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:

$$\omega_A = \left(\frac{\alpha_A}{\sigma(\epsilon_A)^2}\right) \left/ \left(\frac{\mathrm{E}[R_m] - R_f}{\sigma_m^2}\right) \right.$$
(45)

where $E[R_m]$ and σ_m^2 are the expected return and variance of the passive portfolio, respectively.

We illustrate the impact portfolio's alpha and its corresponding weight, ω_A , using a numerical example. Suppose the passive portfolio has an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. The idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities.³⁵ Consider again a collection of N = 500 securities. We divide them into 10 decile portfolios ranked by the impact factor, **X**. For several different values of ρ and σ_{α} , Panel A of Table 1 reports the expected excess return of the impact portfolio,³⁶ α_A in (42), for the top and bottom two decile portfolios. Panel B of Table 1 reports the weight, ω_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,³⁷ but even with such a modest range of α , the impact portfolios

³⁵This is an innocuous assumption and we show later via simulation that cross-sectional heterogeneity in idiosyncratic volatilities does not affect our conclusions.

³⁶We set the individual security weights of the impact portfolios to be always positive by convention from the perspective of long-only investors.

³⁷For example, using several factor models, Pástor and Stambaugh (1999) estimate σ_{α} to be between 0% to 10%

Table 1: Expected excess returns for the impact portfolios and their corresponding Treynor-Black weights when combined with a passive portfolio. We set N = 500 and assume that the passive portfolio has an annualized expected excess return of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. The idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities. Panel A reports α_A for the impact portfolio and Panel B reports the weight, ω_A , and the corresponding excess return of the impact portfolio combined with the passive portfolio.

Correlation ρ	Bottom	2^{nd}	9^{th}	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 A: Expected excess return of the impact portfolio (α_A)

Panel B: Impact portfolios combined with a passive portfolio

Correlation a		Weight	ω_A		Expected	Excess	s Return	n $\omega_A \alpha_A$				
Correlation ρ	Bottom	2^{nd}	9^{th}	Top	Bottom	2^{nd}	9^{th}	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,	$\sigma_{\alpha} = 2^{0}$	70							
$30\% \ (R^2 = 9\%)$	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,	$\sigma_{\alpha} = 5\%$	70							
$30\% \ (R^2 = 9\%)$	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%				

yield economically significant alpha. For example, α_A for the top decile is 1.1% when the correlation $\rho = 30\%$, and 0.4% when $\rho = 10\%$. Observe that ρ^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, σ_{α} , 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 (top half of Panel B). For example, when $\sigma_{\alpha} = 2\%$, ω_A for the top decile is 1.31 when the correlation $\rho = 30\%$, and 0.44 when $\rho = 10\%$. When $\sigma_{\alpha} = 5\%$, the impact portfolio has a weight of 3.23 for the top decile when $\rho = 30\%$. This implies a highly leveraged portfolio in which more than 200% of the passive portfolio is shorted. The corresponding gain in expected excess return for the combined portfolio is 8.9% when $\rho = 30\%$ 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 shows the performance of long-only combined portfolios, which corresponds to a constrained weight, ω_A , between 0 and 1. When the correlation, ρ , 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, ρ , is negative, only the bottom two deciles provide positive excess returns. However, securities in these deciles have the lowest impact factor **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, Figure 6 generalizes the long/short portfolios in the top half of Panel B, and displays two metrics for the combined portfolio that consists of the impact and passive portfolios, with two different levels of σ_{α} . In Figures 6a and 6b we consider $\sigma_{\alpha}=2\%$. In other words, most of the securities have an alpha within [-4%, 4%]. The weights of the active portfolio range from -1.5 to 1.5, depending on the correlation between α and **X** (Figure 6a). The expected excess return of the combined portfolio ranges from 0% to over 2.5% (Figure 6b).

In Figures 6c and 6d we consider $\sigma_{\alpha} = 5\%$. In other words, most of the securities have an alpha within [-10%, 10%]. This is not unimaginable in some highly volatile sectors such as biotech. The weights of the active portfolio can be as high as two, indicating a leveraged impact portfolio and a short position in the passive market portfolio (Figure 6c). In this case, the expected excess return of the impact portfolio can yield up to 14% (Figure 6d)!

Figure 7 is the long-only counterpart of Figure 6, which generalizes the long-only portfolios in the

with a Bayesian framework (depending on different priors) for 1,994 stocks.



Figure 6: Performance metrics for the combined long/short portfolio that consists of the impact portfolio with N = 500 and passive market portfolio with an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. The idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities. (a) and (b) show the Treynor-Black weight for the impact portfolio and the overall expected excess return, respectively, for $\sigma_{\alpha} = 2\%$. (c) and (d) show the same metrics for $\sigma_{\alpha} = 5\%$.



Figure 7: Performance metrics for the combined long-only portfolio that consists of the impact portfolio with N = 500 and passive market portfolio with an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. The idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities. (a) and (b) show the Treynor-Black weight for the impact portfolio and the overall expected excess return, respectively, for $\sigma_{\alpha} = 2\%$. (c) and (d) show the same metrics for $\sigma_{\alpha} = 5\%$.

bottom half of Panel B in Table 1. The top-right regions in Figures 7a and 7c show positive weights for the impact portfolio, because top-ranking securities yield positive expected excess returns when the correlation, ρ , is positive. Similarly, the bottom-left regions in Figures 7a and 7c also show positive weights because these securities, despite having low impact, have positive expected excess returns. As a result, the positive excess returns in the top-right regions in Figures 7b and 7d provide a measure of the financial benefit of forming long-only impact portfolios, while the bottomleft regions can be interpreted as the financial cost of not investing in the lowest impact deciles.

Finally, we show how our impact portfolio improves the efficient frontier and the Capital Market Line to achieve a "super-efficient frontier" under the assumption that α are mispricings. Under the alternate omitted-factor interpretation, the "super-efficiency" of the new frontier may be viewed as the result of additional risk premia not accessible to investors except through impact portfolio managers.

Proposition 10. Under Assumption (A1), the return of the final portfolio, P, that consists of the impact portfolio with Treynor-Black weights and the passive market portfolio is:

$$R_P - R_f = \omega_A R_A + (1 - \omega_A) R_m - R_f = \omega_A \alpha_A + (R_m - R_f) (\beta_A \omega_A + (1 - \omega_A)) + \omega_A \epsilon_A, \quad (46)$$

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

$$\mathbf{E}[R_P] - R_f = \omega_A \alpha_A + (\mathbf{E}[R_m] - R_f)(\beta_A \omega_A + (1 - \omega_A)), \tag{47}$$

$$\operatorname{Var}[R_P] = \operatorname{Var}[R_m](\beta_A \omega_A + (1 - \omega_A))^2 + \omega_A^2 \sigma(\epsilon_A)^2.$$
(48)

This forms a super-efficient frontier in comparison to the Capital Market Line associated with the passive portfolio.

Figure 8 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 Figures 8a and 8c, the idiosyncratic volatility is assumed to be a constant $\sigma(\epsilon_i) = 15\%$ for all securities. As the correlation, ρ , and variance in alpha, σ_{α}^2 , increase, the impact portfolios (defined as the top half of the securities ranked by **X**) are able to improve the original Capital Market Line for both long/short and long-only portfolios, leading to super-efficient 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 lognormal distribution:

$$\log\left(\sigma(\epsilon_i)\right) \sim \operatorname{Normal}(\mu_{\epsilon}, \sigma_{\epsilon}). \tag{49}$$

Calibrating to empirically plausible values in the literature (e.g., Kuntz (2020)), we perform simulations for $\log(\mu_{\epsilon}) = 15\%$ and $\sigma_{\epsilon} = 1$. Figures 8b and 8d confirm that even with such cross-sectional heterogeneity, the Capital Market Line is still improved, leading to super-efficient frontiers.



Figure 8: Super-efficient frontiers from the combined portfolio that consists of the impact portfolio with N = 500 and passive market portfolio with an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$. (a) and (b) allow short positions in the passive index, while (c) and (d) consider long-only portfolios. In (a) and (c) the idiosyncratic volatility is a constant $\sigma(\epsilon_i) = 15\%$ for all securities. In (b) and (d) we simulated idiosyncratic volatility based on (49) and apply a maximum leverage ratio of 3:1.

5.5 Qualifications

The Treynor-Black portfolios outlined in Proposition 9–10 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. In particular, we provide examples of both long/short and long-only portfolios in Section 5.4. In the meantime, we emphasize that the performance metrics associated with these portfolios should be taken as proof-of-concepts, 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 to traditional approaches. First, investors may follow traditional mean-variance optimization and integrate the level of impact as part of either the objective or constraint.³⁸ 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 (2) 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 estimated by regressing returns onto the impact factor \mathbf{X} . Compared to this approach, our framework corresponds more closely to how the industry actually constructs impact portfolios, which usually relies on ordering securities based on \mathbf{X} and assigning weights in each quartile or decile. In fact, our theory requires weaker assumptions because the results only depend on the rank information in \mathbf{X} , not its value. As a result, the optimal portfolio relies on quantiles of \mathbf{X} , making our approach robust against noise and outliers, as demonstrated empirically by Lo, Zhang, and Zhao (2022). In addition, the focus on rank in \mathbf{X} allows for the generalization of our theory to capture nonlinear dependence between security returns and the impact factor \mathbf{X} , as shown in a subsequent study (Lo et al., 2022).

The key insight of our framework lies in the ability to quantify the distribution of impactranked returns when mispricings exist for individual securities. In the omitted factor interpretation, it identifies impact factors \mathbf{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 tradeoff 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 "super-efficient frontier" either through mispricings or additional risk premia not accessible to

³⁸In 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 expected utility subject to a constraint that the portfolio ESG is higher than a preset threshold. See Pedersen, Fitzgibbons, and Pomorski (2021) and Pástor, Stambaugh, and Taylor (2021) for example.

investors except through impact portfolio managers.

The four examples of "super-efficient frontier" in Figure 8 are based on impact portfolios formed with the top half of the securities ranked according to \mathbf{X} . More generally, each investor can decide on the most suitable subset of securities depending on their desired level of impact. In the special case of an ESG metric, this process yields the "ESG-efficient frontier" of Pedersen, Fitzgibbons, and Pomorski (2021), which provides the highest attainable Sharpe ratio for each ESG level.

Our results add to Pedersen, Fitzgibbons, and Pomorski (2021) in two ways. First, we provide the explicit construction of the optimal impact portfolio, thanks to our ability to characterize impact-ranked returns for individual securities in Proposition 2-9, which yields an explicit measure of the degree to which **X** improves or worsens the efficient frontier. Second, 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, Fitzgibbons, and Pomorski, 2021, p. 573). Our framework shows that, when there are non-zero α that are otherwise inaccessible to investor (either interpreted as mispricings or omitted factors), simply using the impact factor, **X**, to form portfolios based on ranking or subsetting can also improve the efficient frontier, because of the information implicit in the selection criteria.

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).³⁹ To see this, observe that our results in Proposition 7 quantify the excess returns of impact-ranked securities and, in particular, identifies \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, while the standardized impact score, $\mathbf{E}[X_{i:N}]$, 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} , though the corresponding SDF is not uniquely determined.⁴⁰

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.

³⁹We thank the associate editor for pointing out this interesting connection.

⁴⁰The correspondence between return factor models and SDF factor models are initially due to 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 $E[R^e] = \beta' E[\lambda]$, one can always find a SDF, m, such that m satisfies $m = 1 + b'\lambda$ and $0 = E[mR^e]$, where b is given by the equation $E[\lambda] = -\text{Var}[\lambda]b$. Nonetheless, m is not unique because one can always add to m any random variable orthogonal to returns.
6.1 Venture Philanthropy: The Cystic Fibrosis Foundation

The concept of venture philanthropy (VP) was introduced by Letts, Ryan, and Grossman (1997), who suggested that nonprofit organizations could learn useful practices from venture capitalists, including due diligence, risk management, performance measurement, relationship management, investment duration and size management, and exit strategies. This approach has received a great deal of attention both within and outside the field (Grossman, Appleby, and Reimers, 2013), and has now been applied to education (Scott, 2009), community redevelopment (Van Slyke and Newman, 2006), and medical R&D (Scaife, 2008; Salzman, 2016), among other fields. 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 Cystic Fibrosis (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 α , 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 thirty thousand Americans. Over a period of 12 years, the CF Foundation invested \$150 million to fund CF drug development efforts at Vertex Pharmaceuticals, a Boston-based biotechnology firm. This work led to the identification and development of Kalydeco, the first FDA-approved treatment to address the underlying causes of CF. The Foundation's investment entitled them to receive royalties calculated as a percentage of future sales of successful CF drugs. 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.

From the financial perspective, a \$3.3 billion return from a \$150 million investment is the dream scenario for any investor, but it could seem like just one individual success story. If we consider CF Foundation's entire portfolio of VP efforts, they allocated a medical and research budget of \$87 million across more than 500 awards in 2012, and over \$160 million across more than 1,100 awards in 2016 (Kim and Lo, 2019). Apparently, from the portfolio perspective, the \$3.3 billion return is still very attractive after factoring in CF Foundation's investments in other projects, even assuming everything else did not produce any financial reward.⁴¹

If we assume for simplicity that the \$150 million investment was made upfront 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 α of this investment, we require an estimate of the cost of capital of Vertex

⁴¹In fact, since 2014, the CF Foundation has sold additional royalty interests, bringing their total investment returns to over \$4 billion since inception. However, for our purposes, we focus only on the single sale to Royalty Pharma for simplicity since it occurred at a single point in time.

during the 12-year investment period from 2002 to 2013, prior to the 2014 royalty sale. We consider a simple CAPM model:

$$R_i - R_f = \alpha_i + \beta_i (R_M - R_f) + \epsilon_i, \tag{50}$$

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. Figure 9 displays the 250-day rolling-window daily estimated beta of Vertex from 17 July 1992 to 30 December 2020, and the average value between 2 January 2001 and 31 December 2013 was 1.42. The average 5-year constant-maturity Treasury yield from January 2001 to December 2013 was 2.8%,⁴² and the annualized compound return of the 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 multi-year 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 α 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).



Figure 9: 250-day rolling-window estimated daily beta coefficients for Vertex Pharmaceuticals from 17 July 1992 to 30 December 2020.

Using this estimated range for the CF Foundation's α , we can estimate its correlation with a "rare disease impact investing" factor **X**. Making a few additional assumptions about auxiliary parameters, we can reverse-engineer the implied correlation, ρ , that is consistent with this performance range, which is [35%, 86%].⁴³ Our highly stylized calculations are not meant to yield a

⁴²See https://fred.stlouisfed.org/series/GS5/.

 $^{^{43}}$ Using a Bayesian framework, Pástor and Stambaugh (1999) estimate the posterior σ_{α} to be between 0% to 10%

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. But 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 *pre-requisite* 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.

More generally, most early-stage drug development programs have low probabilities of success, long time horizons, and large capital requirements (Fagnan et al., 2013), making them less attractive investments than alternatives in other industries like software, social media, telecommunications, etc. In recent years, new tools have emerged to quantify and diversify the risk in these investments (Fagnan et al., 2013; Thakor et al., 2017). Our impact framework provides a systematic approach for constructing impact portfolios and measuring their financial performance, and properly measuring and managing the risk of these investments is the first step towards encouraging more capital to be allocated to accelerate drug development and build greater social value.

6.2 R&D Consortium: Sematech

The second example of impact investing is Semiconductor Manufacturing Technology ("Sematech"), the semiconductor R&D consortium that was established after the U.S. Congress passed the National Cooperative Research Act (NCRA) in 1984 (Evan and Olk, 1990). In 1987, 14 U.S. semiconductor firms and the U.S. government formed Sematech to solve common manufacturing problems by leveraging shared R&D resources and research results. During its early years, the U.S. government financed about half of Sematech's roughly \$200 million annual budget. After 1996, the Sematech Board of Directors decided to end matching federal funding because it had become clear that the U.S. semiconductor industry had regained strength and market share, but Sematech continued to serve its membership.

⁽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 α 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%].

Previous research has studied the effects of Sematech on members' R&D spending, profitability, investment, and productivity (Irwin and Klenow, 1996a,b; Link, Teece, and Finan, 1996). In particular, it has been estimated that Sematech induced members to cut their overall R&D spending on the order of \$300 million per year, and raised members' profitability relative to non-members' (Irwin and Klenow, 1996a). This reflects more sharing and less duplication of research. In other words, more research has been accomplished per R&D dollar.

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 non-members, 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 (SIC 3674) from 1975 to 1999.⁴⁴ Our sample consists of 52 firms, including 11 of the original 14 members of Sematech.⁴⁵

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

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

$$R_{i,\text{post}} - R_{f,\text{post}} = \alpha_{i,\text{post}} + \beta_{i,\text{post}} (R_{M,\text{post}} - R_{f,\text{post}}) + \epsilon_i,$$
(51)

where $R_{i,\cdot}$ is the return of the *i*-th security, $R_{f,\cdot}$ is the risk-free rate estimated by the 5-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 to 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 **X** with $X_i = 1$ representing Sematech membership and 0 otherwise, this leads to an estimated correlation, ρ , of 25% with a range of [1.4%, 49%] between the Sematech impact factor and excess returns.⁴⁶

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 such as the size, value, and industry effects are accounted for. Nonetheless,

⁴⁴Sematech consists 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 to 1999 which covers roughly 12 years before and after the formation of Sematech. Returns are winsorized at 15% each side to reduce the impact of outliers.

⁴⁵We require at least 6 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.

⁴⁶We assume that the cross-sectional standard deviation of α is $\sigma_{\alpha} = 20\%$, and 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%].

Table 2: Estimated annualized CAPM excess returns based on a difference-in-difference (DiD) approach that compares Sematech member firms against non-members, before and after the formation of Sematech. Estimated coefficients are significant at the 1 percent (***), 5 percent (**), or 10 percent (*) levels based on Bootstrap confidence intervals.

Firm	Period	Excess Return	Excess Return	Excess Return
			(Post minus Pre)	(DiD)
Sematech Members	Pre	-2.67%	9.90%**	$10.22\%^{*}$
	Post	$7.23\%^{***}$		
Non-Members	Pre	-8.90%***	-0.33%	
	Post	$-9.22\%^{***}$		

our results are consistent with prior evidence on Sematech's far-reaching effects on members' R&D spending and profitability (Irwin and Klenow, 1996a,b; Link, Teece, and Finan, 1996). Our 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 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 to mutual or hedge funds, and receive less coverage from analysts. As a result, sin stocks seem to yield higher expected returns (Fabozzi, Ma, and Oliphant, 2008; Hong and Kacperczyk, 2009; Statman and Glushkov, 2009), an observation also shared by Fauver and McDonald IV (2014) on international stocks.

This empirical fact implies a negative correlation between a stock's excess return and an "antisin stock" factor. In terms of the super-efficient frontier shown in Figure 8, 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 United States from 1965 to 2006.⁴⁷ This corresponds to the following four-factor asset-pricing model:

$$R_i - R_f = \alpha_i + \beta_{i,1}(R_M - R_f) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}MOM + \epsilon_i,$$
(52)

⁴⁷Hong and Kacperczyk (2009) use monthly returns. The authors also run a cross-sectional regression controlling firm characteristics and get similar excess returns.

and an impact factor **X** representing whether a stock is considered sin. Hong and Kacperczyk's (2009) estimate can be used to calculate the implied correlation between α and **X** in our model, using results from Propositions 2 and 6 (see also discussions in Section 5.2).

Panel A of Table 3 summarizes these calibration results.⁴⁸ The implied correlation is 27% $(R^2 = 7.2\%)$, assuming a standard deviation of cross-sectional alpha of $\sigma_{\alpha} = 5\%$.⁴⁹ 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 anti-sin 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%).

Table 3: Estimated cost in excess return per annum for avoiding sin stocks, calibrated to prior empirical studies. Here we assume that the passive portfolio has an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$, and $\sigma_\alpha = 5\%$.

Impost	Weight of	Expected Excess Return			
Portfolio	Impact Portfolio	Impact Portfolio	Combined with Passive Portfolio		
1 01 010110	ω_A	$lpha_A$	$\omega_A lpha_A$		
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%		

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:

$$R_{i} - R_{f} = \alpha_{i} + \beta_{i,1}(R_{M} - R_{f}) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}MOM + \beta_{i,5}RMW + \beta_{i,6}CMA + \epsilon_{i}.$$
 (53)

 $^{^{48}}$ 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, (22).

⁴⁹We use an intermediate value based on Pástor and Stambaugh's (1999) estimate of σ_{α} (between 0% to 10% for 1,994 stocks). Results for other values of σ_{α} follow trivially. Note that different values for σ_{α} lead to different estimates of ρ , but not the final estimates of the cost of avoiding sin stocks because the expected alpha in (22), for example, is invariant of the product of the two ($\rho\sigma_{\alpha}$).

We also summarize the implied correlation and cost of divesting sin stocks based on Blitz and Fabozzi (2017) in Panel B of Table 3.⁵⁰ 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)–(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.

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.

Compared to sin stocks, the empirical evidence on ESG's excess returns is mixed. On the one hand, several studies find that portfolios or funds with high ESG scores tend to outperform otherwise comparable investments (Shing, 2021; Madhavan, Sobczyk, and Ang, 2021; Bansal, Wu, and Yaron, 2022; Berg et al., 2022b). On the other hand, others argue that the evidence that markets reward companies for being "good" is weak to non-existent (Alessandrini and Jondeau, 2020; Cornell and Damodaran, 2020; Lindsey, Pruitt, and Schiller, 2021), which is supported by recent evidence that green bonds—bonds whose proceeds are used for environmentally sensitive purposes—are indeed priced at a premium, implying a lower yield compared to otherwise equivalent bonds (Baker et al., 2022). Ardia et al. (2022) and Pástor, Stambaugh, and Taylor (2022) further show that the high returns for green assets in recent years reflect unexpectedly strong increases in environmental concerns, not high expected returns.⁵¹

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. (2022b) who construct ESG portfolios based on a range of aggregation methods to de-noise 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 U.S. using a simple average ESG score, which implies a much higher 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 vary with the specific ESG metric.

In contrast, Baker et al. (2022) study the U.S. bond market and report a yield difference of 6 basis points at issuance for green bonds below other ordinary bonds.⁵² This corresponds to a

 $^{^{50}}$ Blitz and Fabozzi (2017) use monthly returns. For U.S. data in 1963–2016, the authors report a non-significant monthly excess return of 0.10%. This number becomes negative when restricted to data after 1990.

⁵¹For additional discussion on this topic, see Renneboog, Ter Horst, and Zhang (2008), de Franco (2020), and Pedersen, Fitzgibbons, and Pomorski (2021).

 $^{^{52}2,083}$ green U.S. municipal bonds are used in the sample, compared to 643,299 ordinary bonds.

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%.⁵³ 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.

Table 4: Estimated ESG excess returns per annum, calibrated to prior empirical studies. Here we assume that the passive portfolio has an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$, and $\sigma_\alpha = 5\%$ for stock markets and $\sigma_\alpha = 1\%$ for bond markets.

	Weight of	Expected Excess Return				
Impact	Impact Portfolio	Impact Portfolio	Combined with Passive Portfolio			
Portiolio	ω_A	$lpha_A$	$\omega_A lpha_A$			
Panel A: B	Panel A: Berg et al. (2022b)					
Implied cor	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: B:	aker et al. (2022)					
Implied cor	relation $\rho = -2.0\%$	$(R^2 = 0.04\%).$				
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: B	ansal, Wu, and Yard	on (2022) ("good ti	mes")			
Implied cor	relation $\rho = 22\%$ (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: B	ansal, Wu, and Yar	on (2022) ("bad tin	nes")			
Implied cor	relation $\rho = -0.2\%$	$(R^2 = 0.0\%).$				
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, Stambaugh, and Taylor (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, Zhang, and Zhao (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%			

⁵³We 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% to 10% in the stock market). Similar to our results for sin stocks, different values for σ_{α} lead to different estimates of ρ , but not the final estimates of the ESG alpha.

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, Wu, and Yaron (2022) document a "luxury-good effect" for a ESG 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 \mathbf{X} representing the customized ESG score, and a four-factor asset-pricing model with an interaction term:

$$R_{i} - R_{f} = \alpha_{i} + \beta_{i,0}I + \beta_{i,1}(R_{M} - R_{f}) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}MOM + \gamma_{i,1}(R_{M} - R_{f}) \cdot I + \gamma_{i,2}SMB \cdot I + \gamma_{i,3}HML \cdot I + \gamma_{i,4}MOM \cdot I + \epsilon_{i},$$
(54)

where I is a dummy variable representing good economic times defined by the cyclically-adjusted real P/E (CAPE) 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, Wu, and Yaron (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, and sizable positive excess returns for the impact portfolios. However, during bad economic times, the monthly Fama-French fourfactor alpha in Bansal, Wu, and Yaron (2022) becomes -0.0026%, rendering all of our estimates of correlation and ESG alpha to be essentially zero.

Pástor, Stambaugh, and Taylor (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 \mathbf{X} in our framework, they document a highly significant monthly Fama-French three-factor alpha of 0.50% for a green-minus-brown portfolio,⁵⁴ which reduces to an insignificant level of 0.05% when excess returns are regressed on proxies of unexpected climate shocks:

$$R_i - R_f = \alpha_i + \beta_{i,1}(R_M - R_f) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}DCC + \beta_{i,5}DCC_{\text{Lag1}} + \epsilon_i, \quad (55)$$

where DCC stands for a "delta climate concern" variable constructed from Ardia et al.'s (2022) climate change concerns index, and DCC_{Lag1} its one-month lagged values. This implies a much smaller correlation of 5.5% (Panel E) compared to estimates from, for example, Berg et al. (2022b).

Similar, Panel F shows the estimates based on Lo, Zhang, and Zhao (2022), who use the Trucost Environmental data and find a similar effect to Pástor, Stambaugh, and Taylor's (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

⁵⁴Stocks ranked by the greenness scores in the top third minus those in the bottom third.

proxy variables for unexpected climate shocks:

$$R_{i} - R_{f} = \alpha_{i} + \beta_{i,1}(R_{M} - R_{f}) + \beta_{i,2}SMB + \beta_{i,3}HML + \beta_{i,4}RMW + \beta_{i,5}CMA + \beta_{i,6}DCC + \beta_{i,7}DCC_{\text{Lag1}} + \epsilon_{i},$$
(56)

which corresponds to a small implied correlation of 3.8% between excess returns and a low-carbon impact factor \mathbf{X} .⁵⁵

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. For example, Madhavan, Sobczyk, and Ang (2021) show that the security selection alpha by U.S. equity mutual fund managers is related to ESG scores, but only through the component correlated with existing style factors such as value, quality, and momentum. In contrast, no significant relationship was found with the idiosyncratic ESG components not related to style factors. In the context of mutual funds, Geczy, Stambaugh, and Levin (2021) show that the SRI cost depends on the investor's views about asset-pricing models and manager skills. In particular, the SRI cost is minimal compared to a CAPM-investor but may be substantial when investors allow for size, value, and momentum factors, as well as managerial skill.

6.5 The GameStop Phenomenon

In January 2021, the share price of Gamestop Corp. (GME)—a struggling videogame retailer that had recently announced a 30% decline in 2020Q3 net sales, due in part to 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. 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."⁵⁶ While it is difficult to determine the exact cause and motivation behind the early initiators,⁵⁷ the GME price spike is unlikely to have been driven by changes in the fundamentals of the company,⁵⁸ but rather caused by a combination of a grass-roots "David vs. Goliath" conflict between retail investors and hedge-fund shortsellers, and 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 shortselling pressure from institutional investors in late 2020 and into January 2021. These events

 $^{^{55}}$ 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.

 $^{^{56}}$ A literature quantifying the effect of Reddit activities on the GME mania has already emerged. See Umar et al. (2021), Pedersen (2022), and Lyócsa, Baumöhl, and Výrost (2022) for example.

⁵⁷For example, Hasso et al. (2021) profile retail investors participating in the frenzy using individual brokerage data, and find that GME investors had a history of investing in lottery-like stocks prior to investing in GME. This implies that individual investors may not purely engage in a digital protest against Wall Street.

⁵⁸In fact, GME's revenue has been declining year over year since 2017 and its earnings-per-share has been negative since 2018.

attracted substantial media attention due to 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 10, 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. Although it was clear that this was driven by short-term investor mania, strikingly, GME's stock price stayed at a high level ever since. As of December 2022, the initial \$1 is still worth over \$10.

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



Figure 10: Growth of one dollar invested in GameStop (GME), AMC Entertainment Holdings (AMC), and Blackberry (BB) in logarithmic scale.

The Gamestop example is likely driven by retail mania and herding, rather than a social welfareenhancing 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 towards meme stocks, while the latter can be driven by slow but persistent shifts in preferences towards green assets (or concerns towards climate change), as shown by Pástor, Stambaugh, and Taylor (2022) and Lo, Zhang, and Zhao (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 α , at least for a short period of time.⁵⁹ 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

⁵⁹What constitutes a "short" period of time is clearly subjective and context dependent—as of December 2, 2022, GME's close price was \$27.52, still nearly 10 times higher than what it was in October 2020.

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 anti-trust regulation,⁶⁰ 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 Appendix A.2, 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 α 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 discuss in this section, including measurement error in the impact variable **X** (Section 7.1), the selection of impact factors and its relation to the benchmark asset-pricing model (Section 7.2), and the implications of nonstationarity and impact horizon for measuring the correlation ρ between impact and returns (Section 7.3).

7.1 Impact Measurement Error

It should be emphasized that measurement errors can exist in the impact variable \mathbf{X} itself, just like any other stock-level characteristic. This is especially challenging for emerging concepts such as ESG, as documented recently by Berg, Kölbel, and Rigobon (2022). 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., 2022a) or aggregation methods designed for denoising (Berg et al., 2022b). Nonetheless, empirical estimates of ESG returns should be taken with the caveat that they are potentially affected by the measurement noise.

⁶⁰See the Securities Exchange Act of 1934, Section 9, and the Sherman Act and the Commodity Exchange Act.

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 (50)–(56). 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 **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 multi-factor models (Fama and French, 1993, 2015).⁶¹ 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 measurement errors discussed above. In the case of biotech venture philanthropy, curing the disease is a pre-requisite 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 Merton's (1987) model of neglected stocks and segmented markets or Fama and French's (2007) taste model, which predicts that non-socially 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, Stambaugh, and Taylor, 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 vs bad economic times (Bansal, Wu, and Yaron, 2022) and market-wide shifts in attention to climate change (Pástor, Stambaugh, and Taylor, 2022; Lo, Zhang, and Zhao, 2022). And finally, the causal mechanism for the meme-stock phenomenon is given by the price-impact model in Appendix A.2 or Pedersen's (2022) model of influencers and thought leaders in social networks.

 $^{^{61}}$ It is known that testing an asset-pricing model against an unspecified alternative results in poor power (see, for example, MacKinlay (1987) and Campbell, Lo, and MacKinlay (1997, pp. 261)). From the Bayesian perspective, we need at least two models to update their posteriors with data. Without a concrete alternative model, the likelihood $P(\text{data}|\text{not } H_0)$ is ill-defined.

7.3 Nonstationarities and Estimating ρ

The specific motivation underlying each impact investment also plays an important role in determining the key parameter ρ , the correlation between the impact factor **X** and returns. In particular, such motivation is often nonstationary, which implies that ρ 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, Bolton, and Samama, 2016), in contrast to today's 5,301 UNPRI signatories.⁶² 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 **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 due to general environmental concerns, this factor can have an impact on the company's returns, thereby inducing a risk premium and, consequently, a non-zero ρ .

Because impact investing is so often associated with non-pecuniary 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 ρ using decades of historical data is unlikely to be fruitful given the shift in these correlations over time.⁶³ A more useful approach is to either estimate ρ 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, ρ is assumed to be a fixed constant, but in practice, ρ is not only timevarying 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 of 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 timespan that it is likely to be observed. And when disclosing the financial impact of such an

⁶²https://www.unpri.org/signatories/signatory-resources/signatory-directory (accessed 2 Dec 2022).

⁶³For example, Pástor, Stambaugh, and Taylor (2022) and Lo, Zhang, and Zhao (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.

investment, it would be prudent to disclose the estimated correlation ρ over multiple horizons, in the same way that retail investors are now given 1-, 3-, and 5-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 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 α to measure the correlation between **X** and α , which are demonstrated empirically in two projects involving a wide range of ESG metrics (Berg et al., 2022b) and environmental metrics such as carbon emission (Lo, Zhang, and Zhao, 2022). Then why not just estimate α and stop there? The reality is that not all investors have access to good estimates of α , not to mention the alphas of more sophisticated impact portfolios. Our framework provides a much lower (one) dimensional quantity (ρ), compared to α , 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. And by comparing the properties of impact portfolios on the Treynor-Black super-efficient-frontier with those of non-impact 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 above.

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 were 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, Jiang, and Kim, 2015), long-term fundamentals (Bebchuk, Brav, and Jiang, 2015), and stock performance (Dimson, Karakaş, and Li, 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, Liu, and Zhu, 2016; Feng, Giglio, and Xiu, 2020; Hou, Xue, and Zhang, 2020). From this perspective, our model has defined a measure for the alpha of any factor, providing a unified framework for SRI, ESG, and beyond.

Our framework may also help inform regulators and policymakers on the most appropriate tools to encourage investments with more socially-aware goals. Not all types of impact investing are created equal. When these investments create positive excess returns, one must understand what drives the initial under-valuation 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 to 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. While 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 so as to identify more opportunities for doing well by doing good.

A Appendix

In this Appendix, we provide additional technical results and proofs for all the propositions.

A.1 Numerical Examples for the Performance of Impact Portfolios

To develop intuition for Proposition 8, consider a portfolio formed by selecting the top n_0 securities based on **X**. For a market with N = 50 securities, Figure A.1 displays the mean and variance of the excess return of portfolios formed in this way. As the number of securities in the portfolio, n_0 , increases, the excess return decreases because more securities with weaker alphas are included. At the same time, the variance of the portfolio also decreases thanks to the diversification from more securities.



Figure A.1: Distribution of portfolio excess return formed by the top n_0 securities ranked by the impact factor, **X**. The number of total securities, N, is set to be 50.

Another typical way of forming portfolios is to sort all securities in the universe into 10 deciles based on \mathbf{X} . Figure A.2a contains the expected excess returns of the 10 deciles, which has a similar shape to the expected excess returns of individual securities in Figure 1a.

Finally, we can also consider portfolios as N increases without bound. Suppose we divide the [0, 1] interval into L segments each of length 1/L, and pick M equally-spaced quantiles within each segment. Specifically, the *l*-th portfolio is formed by selecting the following quantiles:

$$\xi_{l,m} = \frac{l-1+\frac{m}{M+1}}{L}, \quad m = 1, 2, \cdots, M$$
 (A.1)

for $l = 1, 2, \dots, L$. Figure A.2b shows the expected excess returns of this portfolio when L = M = 10, which, not surprisingly, has a similar shape to Figure A.2a because the portfolio formed by (A.1) is the limit of the decile portfolio when N increases without bound.

A.2 A Simple Execution Model for Meme Stock Trading

Following Bertsimas and Lo's (1998) framework and notation, we assume that an investor seeks to purchase a total of \bar{S} shares of a particular security over a fixed time interval, [0, T]. The investor decides how to divide \bar{S} into smaller purchases distributed throughout the interval so as to maximize



Figure A.2: Expected excess return for decile portfolios formed by ranking the impact factor, **X**. In (a) the number of total securities, N, is set to be 50, and in (b) we show the case when N increases without bound.

the final price-impact of the security.⁶⁴ The answer depends, of course, on the degree to which a single purchase affects the market price, i.e., the "price impact" and the dynamics of future market prices. Given a particular price-impact function and a specification for the price dynamics, an optimal trading strategy that maximizes the price impact of acquiring \bar{S} in [0, T] may be obtained.

Specifically, denote by S_t the number of shares acquired in period t at price P_t , where $t = 1, 2, \dots, T$. Then the investor's objective of maximizing final price impact is given by:

$$\max_{\{S_t\}} \mathbb{E}[P_T] \tag{A.2}$$

subject to the constraint that the desired number of shares are acquired:

$$\sum_{t=1}^{T} S_t = \bar{S}.$$
(A.3)

We assume that the security price follows the bivariate stochastic process:

$$P_t = P_{t-1} + \theta S_t^z + \gamma F_t + \epsilon_t, \quad \theta > 0, \ z \in (0, 1]$$

$$F_t = \delta F_{t-1} + \eta_t, \quad \delta \in (-1, 1)$$
(A.4)

where ϵ_t and η_t are independent white noise processes with mean 0 and variance σ_{ϵ}^2 and σ_{η}^2 , respectively. The parameter θ specifies the magnitude of the price impact, which is assumed to follow a power law in S_t , where the parameter z specifies the "price sensitivity" of the security or, equivalently, the security's degree of illiquidity. The latter interpretation is motivated by Kyle's (1985) market microstructure model in which liquidity is measured by a loglinear-regression estimate of the log-volume required to move the price by one dollar. Sometimes referred to as "Kyle's lambda,"

⁶⁴Note that this is *not* the objective function considered by Bertsimas and Lo (1998)—the problem they pose is how to divide \bar{S} so as to maximize *cumulative profits*, which they solve via stochastic dynamic programming.

this measure is an inverse proxy of liquidity, with higher values of lambda implying lower liquidity and lower market depth. 65

The presence of F_t in the law of motion for P_t captures the potential impact of market conditions or private information about the security. For example, F_t can represent new business opportunities created by the company, the impact of popular sentiment—as in the case of GME, as well as any of the other factors mentioned above. In either case, the impact of F_t on trading profits and the time series properties of F_t both have important implications for the feasibility and profitability of price-impact investing. With these price dynamics, the following result completely characterizes the optimal price-impact strategy and its corresponding expected profit:

Proposition A.0. Under the price dynamics specified by (A.4), the strategy that maximizes the total price impact, (A.2), is given by:

$$S_1 = S_2 = \dots = S_T = \frac{S}{T},\tag{A.5}$$

and its corresponding expected profit is given by:

$$V^* = \left(\frac{\theta \bar{S}^z (T-1)}{2T^z} + \frac{\gamma \delta F_1 \left(1 - T \delta^{T-1} + (T-1) \delta^T\right)}{(1-\delta)^2 T}\right) \bar{S}.$$
 (A.6)

In fact, when z=1 and price impact is linear in trading quantities,⁶⁶ it does not matter how trades are allocated because the total impact from T trades is always equal to the impact of one single trade of size \bar{S} . However, when the price impact is a concave function in general (0 < z < 1), the optimal strategy is to simply divide the total order \bar{S} into T equal "waves," and trade them at regular intervals, as specified in (A.5).

The expression for V^* in (A.6) shows that the expected profit of price-impact investing depends on two factors: the market impact as parameterized by θ and z, and influences from other factors (sentiment, liquidity, private information, etc.) as parameterized by γ and the AR(1) coefficient governing these other factors (δ).

To illustrate the effect of these parameters on trading profit V^* , we simulate a universe of N = 500 securities where the parameters, θ , z, γ , and δ , are generated by four independent uniform distributions on [0, 1]. In the following analysis, we assume that the first realization of $X_1 = 1$, without loss of generality.

In Figure A.3a, we first show the relationship of the expected profit V^* with respect to market impact (θ). As θ increases, expected profit increases as well. This is quite intuitive because the stronger the market impact, the easier it is for short squeezers to induce price momentum and generate profits. If we consider a collection of securities each with a different θ , the correlation between their market-impact coefficients and expected profit is 37%, implying that sorting α based on θ will generate positive excess returns in the context of our impact-investing framework.

Figure A.3b displays the relation between expected profit V^* and sensitivity z. As power increases from 0 to 1, the expected profit decreases. This is because lower values of z correspond to more concave price-impact functions, for which each small trading segment has larger price impact.

⁶⁵See also Lillo, Farmer, and Mantegna (2003) and Almgren et al. (2005) for more detailed explorations of the power law of price impact in equity markets. When z=1, this reduces to the "linear price impact with information" specification from Bertsimas and Lo (1998).

⁶⁶See also Bertsimas and Lo (1998).



Figure A.3: The expected profit, V^* , of price-impact investing as a function of four parameters in (A.4), for a market with N = 500 securities with simulated parameters. Here we set $\theta = 1$, z = 1, $\gamma = 1$, $\delta = 10\%$, $\bar{S} = 1$, $F_1 = 1$, and T = 30 by default, and vary each parameter accordingly.

The correlation between z and expected profit is -63%. In other words, one can achieve positive excess returns by selecting securities based on the reverse ordering of sensitivity z.

Figure A.3c displays the relation between expected profit V^* and influences from other factors (γ) , which has a weak positive correlation of 9%. Finally, Figure A.3d displays the relation between AR coefficient (δ) and expected profit. The expected profit is larger when δ is larger. This is because we have assumed the first realization of F_t is positive, and higher autocorrelations imply stronger momentum. Indeed, the correlation between the AR coefficient, δ , and expected profit is 26% in this simulated market.

We summarize the results from Figure A.3 in Table A.1, and provide their implied α when applied to a collection of 500 securities simultaneously, each with different price dynamics as specified in (A.4). Panels A, B, and C show the expected excess returns if investors apply θ , z, γ , and δ , respectively, to rank securities, where the correlations with trading profits are obtained from our simple execution model. The expected α can be very high with a leveraged portfolio, driven by the high correlation between stock α and the price-impact investing factor in certain cases.

Table A.1: Estimated excess returns per annum for the price-impact investing factor, based on the optimal strategy's profit in (A.6) and its implied correlations with respect to various characteristics of individual securities. Here we assume $\sigma_{\alpha} = 5\%$ —an intermediate value based on Pástor and Stambaugh's (1999) estimated range of σ_{α} (between 0% to 10%), and the passive portfolio has an annualized risk premium of $E[R_m] - R_f = 6\%$ and volatility of $\sigma_m = 15\%$.

Imama at	Weight of	Expected Excess Return			
Portfolio	Impact Portfolio	Impact Portfolio	Combined with Passive Portfolio		
Portiollo	ω_A	$lpha_A$	$\omega_A lpha_A$		
Panel A: Ranking based on price impact (θ)					
Model-implied correlation with alpha: $\rho = 37\%$.					
Top Half	9.19	2.3%	21.4%		
Top Decile	4.09	3.4%	14.0%		
Top 2%	1.20	4.5%	5.4%		
Panel B: Ranking based on market sensitivity (z) ; reverse order					
Model-impl	ied correlation with	alpha: $\rho = 63\%$.			
Top Half	15.58	3.9%	61.5%		
Top Decile	6.95	5.8%	40.3%		
Top 2%	2.04	7.6%	15.5%		
Panel C: Ranking based on other factors (γ)					
Model-implied correlation with alpha: $\rho = 9\%$.					
Top Half	2.20	0.6%	1.2%		
Top Decile	0.98	0.8%	0.8%		
Top 2%	0.29	1.1%	0.3%		
Panel D: Ranking based on AR coefficient for other factors (δ)					
Model-implied correlation with alpha: $\rho = 26\%$.					
Top Half	6.43	1.6%	10.5%		
Top Decile	2.87	2.4%	6.9%		
Top 2%	0.84	3.2%	2.7%		

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, this example highlights the fundamental determinants for a price-impact investor's α : the correlation between each stock's trading profit and stock characteristics, e.g., market capitalization, liquidity, specific forms of market impact, attention from the general public, main shareholders, short interest, or anything correlated with stock returns. Higher correlations lead to higher alpha when following that particular characteristic to select target stock.

A.3 Proof of Proposition 1

The constraints on the right-hand side optimization problem of (3) is a subset of the left-hand side optimization problem. Therefore the inequality follows.

To give a bound on the utility loss between the unconstrained portfolio W and the constrained portfolio W^c , we consider an intermediate portfolio W^{c1} that is also constrained to the subset S, but with equal factor loadings as the unconstrained portfolio W. In other words, the portfolio weights for W^{c1} satisfy the following conditions:

$$\omega_i^{c1} = 0 \text{ for } i \notin \mathcal{S} \tag{A.7}$$

$$\sum_{i \in \mathcal{S}} \omega_i^{c1} = 1 \tag{A.8}$$

$$\sum_{i=1}^{N} \omega_i \beta_{ik} = \sum_{i \in \mathcal{S}} \omega_i^{c1} \beta_{ik} \text{ for } k = 1, \dots, K.$$
(A.9)

Because W^c maximizes the utility in (3),

$$E[U(W^{c})] = E\left[U\left(R_{f} + \sum_{k=1}^{K}\sum_{i\in\mathcal{S}}\omega_{i}^{c}\beta_{ik}\left(\Lambda_{k} - R_{f}\right) + \sum_{i\in\mathcal{S}}\omega_{i}^{c}\epsilon_{i}\right)\right]$$

$$\geq E\left[U\left(R_{f} + \sum_{k=1}^{K}\sum_{i\in\mathcal{S}}\omega_{i}^{c1}\beta_{ik}\left(\Lambda_{k} - R_{f}\right) + \sum_{i\in\mathcal{S}}\omega_{i}^{c1}\epsilon_{i}\right)\right]$$

$$= E\left[U\left(R_{f} + \sum_{k=1}^{K}\sum_{i=1}^{N}\omega_{i}\beta_{ik}\left(\Lambda_{k} - R_{f}\right) + \sum_{i\in\mathcal{S}}\omega_{i}^{c1}\epsilon_{i}\right)\right] = E[U(W^{c1})].$$
(A.10)

Now we consider the utility of the following two portfolios,

$$E[U(W)] = E\left[U\left(R_f + \sum_{k=1}^{K} \sum_{i=1}^{N} \omega_i \beta_{ik} \left(\Lambda_k - R_f\right) + \sum_{i=1}^{N} \omega_i \epsilon_i\right)\right],$$

$$E[U(W^{c1})] = E\left[U\left(R_f + \sum_{k=1}^{K} \sum_{i=1}^{N} \omega_i \beta_{ik} \left(\Lambda_k - R_f\right) + \sum_{i \in \mathcal{S}} \omega_i^{c1} \epsilon_i\right)\right].$$
(A.11)

Note that they only differ in the last term in the parenthesis, the idiosyncratic volatilities. Denote $A \equiv R_f + \sum_{k=1}^{K} \sum_{i=1}^{N} \omega_i \beta_{ik} (\Lambda_k - R_f)$ and $B \equiv \sum_{i=1}^{N} \omega_i \epsilon_i$ (or $\sum_{i \in S} \omega_i^{c1} \epsilon_i$). For any well-behaved utility function U, because E[B] = 0, we have:

$$E[U(A+B)] \approx E\left[U(A) + U'(A)B + \frac{1}{2}U''(A)B^2\right] = E[U(A)] + \frac{1}{2}E[U''(A)]Var[B^2]$$
(A.12)

by second-order Taylor expansion around B=0. Since E[U(W)] and $E[U(W^{c1})]$ differs only through the idiosyncratic volatility term, B, we have:

$$E[U(W)] - E[U(W^{c})] \le E[U(W)] - E[U(W^{c1})]$$

$$\approx \frac{1}{2}E[U''(A)] \left(\operatorname{Var}\left(\sum_{i=1}^{N} \omega_{i} \epsilon_{i}\right) - \operatorname{Var}\left(\sum_{i \in \mathcal{S}} \omega_{i}^{c1} \epsilon_{i}\right) \right).$$
(A.13)

When the number of securities, N, is large, suppose further that:

$$\omega_i \approx \frac{1}{N}, \ \omega_i^{c1} \approx \frac{1}{N-n}, \ \sigma(\epsilon_i) \approx \sigma_{\epsilon},$$
 (A.14)

where n is the number of securities excluded in S, and σ_{ϵ} is the common idiosyncratic volatility for all securities. We have:

$$\mathbf{E}[U(W)] - \mathbf{E}[U(W^c)] \le \frac{1}{2} \mathbf{E}[U''(A)] \left(\frac{\sigma_{\epsilon}^2}{N} - \frac{\sigma_{\epsilon}^2}{N-n}\right) = -\frac{1}{2} \mathbf{E}[U''(A)] \sigma_{\epsilon}^2 \left(\frac{n}{N(N-n)}\right).$$
(A.15)

When the number of securities excluded in S, n, is small relative to the total number of securities, N, the utility loss (A.15) is also small.

Finally, we observe that the assumptions in (A.14) are non-critical for our main conclusions here, and can be relaxed at the expense of simplicity of the mathematical derivation.

A.4 Proof of Proposition 2

Because **X** and α are jointly normal, we can express α_i with the following linear relationship:

$$\alpha_i = \mu_\alpha + \rho \frac{\sigma_\alpha}{\sigma_x} (X_i - \mu_x) + e_i, \qquad (A.16)$$

where e_i are normal random variables with $E[e_i] = 0$ and $Var(e_i) = \sigma_{\alpha}^2(1 - \rho^2)$, and the X_i and the e_i are mutually independent. Ordering securities based on X_i , we have:

$$\alpha_{[i:N]} = \mu_{\alpha} + \rho \frac{\sigma_{\alpha}}{\sigma_x} (X_{i:N} - \mu_x) + e_{[i]}, \qquad (A.17)$$

where $e_{[i]}$ denotes the particular e_i associated with $X_{i:N}$. Note that $X_{i:N}$ on the right-hand side are the usual order statistics, while $\alpha_{[i:N]}$ on the left-hand side are the *induced* order statistics. Because X_i and e_i are independent, the set of $X_{i:N}$ and the set of $e_{[i]}$ are also independent. Therefore, we can calculate the first two moments of α based on the relationship in (A.17):

$$\mathbf{E}\left[\alpha_{[i:N]}\right] = \mu_{\alpha} + \rho \frac{\sigma_{\alpha}}{\sigma_{x}} (\mathbf{E}[X_{i:N}] - \mu_{x}) + e_{[i]} = \mu_{\alpha} + \rho \sigma_{\alpha} \mathbf{E}[X_{i:N}],$$
(A.18)

$$\operatorname{Var}\left(\alpha_{[i:N]}\right) = \rho^{2} \frac{\sigma_{\alpha}^{2}}{\sigma_{x}^{2}} \operatorname{Var}\left(X_{i:N}\right) + \sigma_{\alpha}^{2} (1 - \rho^{2}) = \sigma_{\alpha}^{2} \left(1 - \rho^{2} + \rho^{2} \operatorname{Var}\left(X_{i:N}\right)\right), \quad (A.19)$$

$$\operatorname{Cov}\left(\alpha_{[i:N]}, \alpha_{[j:N]}\right) = \operatorname{Cov}\left(\rho \frac{\sigma_{\alpha}}{\sigma_{x}} X_{i:N}, \rho \frac{\sigma_{\alpha}}{\sigma_{x}} X_{j:N}\right) = \sigma_{\alpha}^{2} \rho^{2} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right).$$
(A.20)

See also David and Nagaraja (2004, Section 6.8).

A.5 Proof of Proposition 3

We first observe that $U_{i:N} \equiv \Phi(X_{i:N})$ maps the *i*-th normal order statistics to the *i*-th order statistics from a uniform distribution on [0, 1], where Φ is the cumulative distribution function of standard normal random variables. We define $Q \equiv \Phi^{-1}$ and write $X_{i:N} = Q(U_{i:N})$. We then expand $Q(U_{i:N})$ in a Taylor series around the expected value of $Q(U_{i:N})$:

$$E[Q(U_{i:N})] = \frac{i}{n+1} = p_i,$$
(A.21)

which gives:

$$X_{i:N} = Q(U_{i:N}) = Q(p_i) + (U_{i:N} - p_i) Q'(p_i) + \frac{1}{2} (U_{i:N} - p_i)^2 Q''(p_i) + \frac{1}{6} (U_{i:N} - p_i)^3 Q'''(p_i) + \cdots$$
(A.22)

Substituting (A.22) into the definition of $E[X_{i:N}]$, $Var(X_{i:N})$, and $Cov(X_{i:N}, X_{j:N})$, and rearranging the terms lead to (11)-(13) in Proposition 3. See also David and Nagaraja (2004, Section 4.6).

In particular, for standard normal random variables we have $Q'(p_i) = 1/\phi(Q)$ where ϕ is the density function for standard normal random variables. Therefore we can calculate:

$$Q''(p_i) = \frac{d(1/\phi(Q))}{d\Phi(Q)} = \frac{d(1/\phi(Q))}{dQ} \frac{dQ}{d\Phi(Q)} = \frac{Q}{\phi^2(Q)},$$
(A.23)

$$Q'''(p_i) = \frac{1+2Q^2}{\phi^3(Q)},\tag{A.24}$$

$$Q''''(p_i) = \frac{Q(7+6Q^2)}{\phi^4(Q)},\tag{A.25}$$

which completes the proof for (14)-(17).

A.6 Proof of Proposition 4

Because **X** and α are both normally distributed, we observe that $\frac{X_i - \mu_x}{\sigma_x}$ and $\frac{\alpha_i - \mu_\alpha}{\sigma_\alpha}$ both follow the standard normal distribution. Therefore,

$$\mathbf{E}\left[\frac{X_{i:N}-\mu_x}{\sigma_x}\right] = \mathbf{E}\left[\frac{\alpha_{i:N}-\mu_\alpha}{\sigma_\alpha}\right],\tag{A.26}$$

$$\operatorname{Var}\left(\frac{X_{i:N} - \mu_x}{\sigma_x}\right) = \operatorname{Var}\left(\frac{\alpha_{i:N} - \mu_\alpha}{\sigma_\alpha}\right),\tag{A.27}$$

$$\operatorname{Cov}\left(\frac{X_{i:N} - \mu_x}{\sigma_x}, \frac{X_{j:N} - \mu_x}{\sigma_x}\right) = \operatorname{Var}\left(\frac{\alpha_{i:N} - \mu_\alpha}{\sigma_\alpha}, \frac{\alpha_{j:N} - \mu_\alpha}{\sigma_\alpha}\right).$$
(A.28)

We have assumed, without loss of generality, that $\mu_{\alpha} = \mu_x = 0$ and $\sigma_x = 1$, which leads to:

$$\mathbf{E}[X_{i:N}] = \frac{\mathbf{E}[\alpha_{i:N}]}{\sigma_{\alpha}},\tag{A.29}$$

$$\operatorname{Var}\left(X_{i:N}\right) = \frac{\operatorname{Var}\left(\alpha_{i:N}\right)}{\sigma_{\alpha}^{2}},\tag{A.30}$$

$$\operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) = \frac{\operatorname{Cov}\left(\alpha_{i:N}, \alpha_{j:N}\right)}{\sigma_{\alpha}^{2}}.$$
(A.31)

This together with (8)-(10) gives:

$$\mu_i = \mathbf{E}\left[\alpha_{[i:N]}\right] = \rho \sigma_\alpha \mathbf{E}\left[X_{i:N}\right] = \rho \mathbf{E}[\alpha_{i:N}]. \tag{A.32}$$

$$\sigma_i^2 - \sigma_\alpha^2 = \sigma_\alpha^2 \rho^2 \left[\operatorname{Var} \left(X_{i:N} \right) - 1 \right] = \rho^2 \left[\operatorname{Var} \left(\alpha_{i:N} \right) - \sigma_\alpha^2 \right],$$
(A.33)

$$\sigma_{ij} \equiv \operatorname{Cov}\left(\alpha_{[i:N]}, \alpha_{[j:N]}\right) = \sigma_{\alpha}^{2} \rho^{2} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) = \rho^{2} \operatorname{Cov}\left(\alpha_{i:N}, \alpha_{j:N}\right).$$
(A.34)

A.7 Proof of Proposition 5

This proposition follows from Yang (1977). See also Lo and MacKinlay (1990) for an application in a different context.

A.8 Proof of Proposition 6

This follows from Proposition 5 by observing that $\Phi(\xi_k) = F_x(\xi_k \sigma_x + \mu_x)$. Alternatively, this result can be proved by taking the limit as $N \to \infty$ based on the finite-sample results in Proposition 2-3.

A.9 Proof of Proposition 7

For simplicity, we define $\boldsymbol{\lambda} \equiv [\lambda_1 \cdots \lambda_N]^T$, and observe that \mathbf{X} and $\boldsymbol{\lambda}$ can be rewritten as:

$$\mathbf{X} = \mu_x \mathbf{1} + \mathbf{C}_x \mathbf{N}_x$$

$$\boldsymbol{\lambda} = \mu_\lambda \mathbf{1} + \mathbf{C}_\lambda \mathbf{N}_\lambda$$
(A.35)

where $\mathbf{1} \equiv [1 \cdots 1]^T$ is a column vector of ones with size N, \mathbf{N}_x and \mathbf{N}_y are both N-dimensional standard normal random vectors with $\operatorname{Cov}(\mathbf{N}_x, \mathbf{N}_y) = \mathbf{\Sigma}$, and \mathbf{C}_x and \mathbf{C}_y are both $N \times N$ deterministic matrices. The specification in (A.35) completely characterizes the joint distribution of \mathbf{X} and $\boldsymbol{\lambda}$. In light of the parameterization in Assumption (A2), we have:

$$\mathbf{C}_{x} = \sqrt{1 - \rho_{x}} \sigma_{x} \mathbf{I} + \left(\sqrt{1 + (N - 1)\rho_{x}} - \sqrt{1 - \rho_{x}}\right) \sigma_{x} \mathbf{L}$$

$$\mathbf{C}_{\lambda} = \sqrt{1 - \rho_{\lambda}} \sigma_{\lambda} \mathbf{I} + \left(\sqrt{1 + (N - 1)\rho_{\lambda}} - \sqrt{1 - \rho_{\lambda}}\right) \sigma_{\lambda} \mathbf{L}$$

$$\mathbf{\Sigma} = \frac{\rho_{x\lambda} - \tilde{\rho}_{x\lambda}}{\sqrt{(1 - \rho_{x})(1 - \rho_{\lambda})}} \mathbf{I} + \left(\frac{\rho_{x\lambda} + (n - 1)\tilde{\rho}_{x\lambda}}{\sqrt{(1 + (n - 1)\rho_{x})(1 + (n - 1)\rho_{\lambda})}} - \frac{\rho_{x\lambda} - \tilde{\rho}_{x\lambda}}{\sqrt{(1 - \rho_{x})(1 - \rho_{\lambda})}}\right) \mathbf{L}$$
(A.36)

where **I** is the identity matrix and $\mathbf{L} \equiv \frac{\mathbf{1} \cdot \mathbf{1}^T}{N}$ is a matrix whose elements are all 1/N.

We now define a projection matrix:

$$\mathbf{P} \equiv \left(\mathbf{C}_{\lambda} \mathbf{\Sigma}^{T} \mathbf{C}_{x}^{T}\right) \left(\mathbf{C}_{x} \mathbf{C}_{x}^{T}\right)^{-1}$$

$$= \underbrace{\rho_{adj} \frac{\sigma_{\lambda}}{\sigma_{x}}}_{a} \mathbf{I} + \underbrace{\left(\frac{\rho_{x\lambda} + (n-1)\tilde{\rho}_{x\lambda}}{1 + (n-1)\rho_{x}} - \rho_{adj}\right) \frac{\sigma_{\lambda}}{\sigma_{x}}}_{b} \mathbf{L}$$

$$= a\mathbf{I} + b\mathbf{L}$$
(A.37)

and it is easy to show that:

$$\boldsymbol{\lambda} - \mathbf{P}\mathbf{X} \perp \mathbf{X}. \tag{A.38}$$

Therefore, when assuming $\mu_x = 0$ and $\sigma_x = 1$, we have:

$$E \left[\lambda_{[i:N]}\right] = E \left[(\boldsymbol{\lambda} - \mathbf{P}\mathbf{X})_{[i:N]} \right] + E \left[(\mathbf{P}\mathbf{X})_{[i:N]} \right]$$
$$= \mu_{\lambda} - (a+b)\mu_{x} + aE \left[X_{i:N} \right] + b\mu_{x}$$
$$= \mu_{\lambda} + \rho_{adj}\sigma_{\lambda}E \left[X_{i:N} \right],$$
(A.39)

which proves (30). The variances and covariances in (31)-(32) can be proven similarly following the same decomposition in (A.39). See also Lee and Viana (1999).

A.10 Proof of Proposition 8

The expected excess return follows directly from the distribution of alphas for single securities in Proposition 2. The variance also follows by rearranging terms:

$$\operatorname{Var}\left(\tilde{\alpha}\right) = \sum_{i \in P} \omega_{i}^{2} \sigma_{i}^{2} + 2 \sum_{i < j \in P} \omega_{i} \omega_{j} \sigma_{ij}$$
$$= \sigma_{\alpha}^{2} \left(1 - \rho^{2} + \rho^{2} \sum_{i \in P} \omega_{i}^{2} \operatorname{Var}\left(X_{i:N}\right) + 2\rho^{2} \sum_{i < j \in P} \omega_{i} \omega_{j} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) \right)$$
$$= \sigma_{\alpha}^{2} \left(1 - \rho^{2} + \rho^{2} \left(\sum_{i \in P} \omega_{i}^{2} \operatorname{Var}\left(X_{i:N}\right) + 2 \sum_{i < j \in P} \omega_{i} \omega_{j} \operatorname{Cov}\left(X_{i:N}, X_{j:N}\right) \right) \right).$$
(A.40)

A.11 Proof of Proposition 9

Because of the decomposition in (39), and the fact that ζ_i are independent of ϵ_i , the combined idiosyncratic variance for security *i* is simply $\sigma_i^2 + \sigma(\epsilon_i)^2$, where σ_i^2 is the variance of the *i*-th induced order statistic given in (9), and $\sigma(\epsilon_i)^2$ is the original idiosyncratic variance for security *i* given in (1). The classical result of Treynor and Black (1973) maintains that to maximize the Sharpe ratio of the portfolio, security weights should be proportional to the expected excess returns divided by the idiosyncratic variance, which proves (40).

(41) follows from plugging in results from Proposition 6 into (40).

A.12 Proof of Proposition 10

By definitions in (42)-(43), the return of the impact portfolio in excess of the risk-free rate can be written as:

$$R_A - R_f = \alpha_A + \beta_A (R_m - R_f) + \epsilon_A. \tag{A.41}$$

When combining with the passive market portfolio, the weight of the impact portfolio, ω_A , is given in (45). Therefore, the return of the combined portfolio, in excess of the risk-free rate, is

$$R_P - R_f = \omega_A (R_A - R_f) + (1 - \omega_A)(R_m - R_f)$$

= $\omega_A (\alpha_A + \beta_A (R_m - R_f) + \epsilon_A) + (1 - \omega_A)(R_m - R_f)$ (A.42)
= $\omega_A \alpha_A + (R_m - R_f)(\beta_A \omega_A + (1 - \omega_A)) + \omega_A \epsilon_A,$

which completes the proof of (46). (47) and (48) follow directly from simple calculations of the expected value and variance of R_P based on (A.42).

A.13 Proof of Proposition A.0

Based on the price process, (A.4), the investor's objective, (A.2), can be written as:

$$E[P_T] = E[P_{T-1} + \theta S_T^z + \gamma F_T + \epsilon_T]$$

= $E[P_{T-2} + \theta S_{T-1}^z + \gamma F_{T-1} + \epsilon_{T-1} + \theta S_T^z + \gamma F_T + \epsilon_T]$
= $P_0 + \theta (S_1^z + \dots + S_T^z) + \gamma (F_1 + \delta F_1 + \dots + \delta^{T-1} F_1]$
= $P_0 + \theta (S_1^z + \dots + S_T^z) + \frac{\gamma (1 - \delta^T) F_1}{1 - \delta}.$ (A.43)

Maximizing $E[P_T]$ over S_1, S_2, \ldots, S_T is the same as maximizing the middle term in (A.43):

$$(S_1^z + \dots + S_T^z). \tag{A.44}$$

When z=1, it does not matter how trades are allocated because (A.44) is always equal to \bar{S} . When 0 < z < 1, (A.44) is a concave function with respect to S_1, S_2, \ldots, S_T , and is maximized when $S_1=S_2=\cdots=S_T=\bar{S}/T$, which completes the proof of the optimal strategy, (A.5).

The optimal profit is simply the total value of the position subtracted by the average execution cost:

$$V^* = \mathbf{E}[P_T] \cdot \bar{S} - \mathbf{E}\left[\sum_{t=1}^T P_t S_t\right] = \left(\mathbf{E}[P_T] - \frac{1}{T} \sum_{t=1}^T \mathbf{E}[P_t]\right) \bar{S}.$$
 (A.45)

Based on a similar derivation to (A.43), it is easy to show that

$$E[P_t] = P_0 + \theta(S_1^z + \dots + S_t^z) + \frac{\gamma(1 - \delta^t)F_1}{1 - \delta} = P_0 + \theta t \frac{\bar{S}^z}{T^z} + \frac{\gamma(1 - \delta^t)F_1}{1 - \delta},$$
(A.46)

for t = 1, 2, ..., T. Substituting (A.46) into (A.45), we have

$$\begin{split} V^* &= \left(\mathbf{E}[P_T] - \frac{1}{T} \sum_{t=1}^T \mathbf{E}[P_t] \right) \bar{S} \\ &= \left(P_0 + \theta T \frac{\bar{S}^z}{T^z} + \frac{\gamma(1 - \delta^T)F_1}{1 - \delta} - \frac{1}{T} \sum_{t=1}^T \left(P_0 + \theta t \frac{\bar{S}^z}{T^z} + \frac{\gamma(1 - \delta^t)F_1}{1 - \delta} \right) \right) \bar{S} \\ &= \left(\theta T \frac{\bar{S}^z}{T^z} + \frac{\gamma(1 - \delta^T)F_1}{1 - \delta} - \frac{1}{T} \sum_{t=1}^T \left(\theta t \frac{\bar{S}^z}{T^z} + \frac{\gamma(1 - \delta^t)F_1}{1 - \delta} \right) \right) \bar{S} \\ &= \left(\theta \frac{\bar{S}^z}{T^z} \left(T - \frac{1}{T} \sum_{t=1}^T t \right) + \frac{\gamma F_1}{1 - \delta} \left((1 - \delta^T) - \frac{1}{T} \sum_{t=1}^T (1 - \delta^t) \right) \right) \bar{S} \end{split}$$
(A.47)
$$&= \left(\theta \frac{\bar{S}^z}{T^z} \left(T - \frac{1 + T}{2} \right) + \frac{\gamma F_1}{1 - \delta} \left(\frac{1}{T} \sum_{t=1}^T \delta^t - \delta^T \right) \right) \bar{S} \\ &= \left(\theta \frac{\bar{S}^z}{T^z} \left(\frac{T - 1}{2} \right) + \frac{\gamma \delta F_1}{1 - \delta} \left(\frac{1 - \delta^T}{T(1 - \delta)} - \delta^{T-1} \right) \right) \bar{S} \\ &= \left(\frac{\theta \bar{S}^z(T - 1)}{2T^z} + \frac{\gamma \delta F_1 \left(1 - T \delta^{T-1} + (T - 1) \delta^T \right)}{(1 - \delta)^2 T} \right) \bar{S} \end{split}$$

which completes the proof of (A.6).

References

- Adler, T., and M. Kritzman, 2008, The cost of socially responsible investing, The Journal of Portfolio Management 35, 52–56.
- Adrian, T., E. Etula, and T. Muir, 2014, Financial intermediaries and the cross-section of asset returns, *The Journal of Finance* 69, 2557–2596.
- Albuquerque, R., Y. Koskinen, and C. Zhang, 2019, Corporate social responsibility and firm risk: Theory and empirical evidence, *Management Science* 65, 4451–4469.
- Alessandrini, F., and E. Jondeau, 2020, ESG investing: From sin stocks to smart beta, The Journal of Portfolio Management 46, 75–94.
- Almgren, R., C. Thum, E. Hauptmann, and H. Li, 2005, Direct estimation of equity market impact, *Risk* 18, 58–62.
- Andersson, M., P. Bolton, and F. Samama, 2016, Governance and climate change: A success story in mobilizing investor support for corporate responses to climate change, *Journal of Applied Corporate Finance* 28, 29–33.
- Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht, 2022, Climate change concerns and the performance of green versus brown stocks, Working Paper, National Bank of Belgium.
- Bai, J., and G. Zhou, 2015, Fama-macbeth two-pass regressions: Improving risk premia estimates, Finance Research Letters 15, 31–40.
- Baker, M., D. Bergstresser, G. Serafeim, and J. Wurgler, 2022, The pricing and ownership of US green bonds, Annual Review of Financial Economics 14, 415–437.
- Bansal, R., D. Wu, and A. Yaron, 2022, Socially responsible investing in good and bad times, The Review of Financial Studies 35, 2067–2099.
- Barber, B. M., A. Morse, and A. Yasuda, 2021, Impact investing, *Journal of Financial Economics* 139, 162–185.
- Bauer, R., T. Ruof, and P. Smeets, 2021, Get real! individuals prefer more sustainable investments, The Review of Financial Studies 34, 3976–4043.
- Bebchuk, L. A., A. Brav, and W. Jiang, 2015, The long-term effects of hedge fund activism, Columbia Law Review 115, 1085–1156.
- Bebchuk, L. A., A. Cohen, and C. C. Wang, 2013, Learning and the disappearing association between governance and returns, *Journal of Financial Economics* 108, 323–348.
- Berg, F., J. F. Koelbel, A. Pavlova, and R. Rigobon, 2022a, ESG confusion and stock returns: Tackling the problem of noise, Available at SSRN 3941514.
- Berg, F., J. F. Kölbel, and R. Rigobon, 2022, Aggregate confusion: The divergence of ESG ratings, *Review of Finance* 26, 1315–1344.
- Berg, F., A. W. Lo, R. Rigobon, M. Singh, and R. Zhang, 2022b, Quantifying the impact of impact investing: An empirical analysis with six ESG metrics, MIT working paper.

- Berk, J., and J. H. van Binsbergen, 2021, The impact of impact investing, Available at SSRN 3909166.
- Bertsimas, D., and A. W. Lo, 1998, Optimal control of execution costs, Journal of Financial Markets 1, 1–50.
- Bhattacharya, P., 1974, Convergence of sample paths of normalized sums of induced order statistics, The Annals of Statistics 2, 1034–1039.
- Blitz, D., and F. J. Fabozzi, 2017, Sin stocks revisited: Resolving the sin stock anomaly, *The Journal of Portfolio Management* 44, 105–111.
- Bolton, P., and M. Kacperczyk, 2021, Do investors care about carbon risk?, Journal of Financial Economics 142, 517–549.
- Brav, A., W. Jiang, and H. Kim, 2015, The real effects of hedge fund activism: Productivity, asset allocation, and labor outcomes, *The Review of Financial Studies* 28, 2723–2769.
- Brinson, G. P., L. R. Hood, and G. L. Beebower, 1986, Determinants of portfolio performance, *Financial Analysts Journal* 42, 39–44.
- Brodie, J., I. Daubechies, C. De Mol, D. Giannone, and I. Loris, 2009, Sparse and stable Markowitz portfolios, *Proceedings of the National Academy of Sciences* 106, 12267–12272.
- Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The econometrics of financial markets* (Princeton University Press, Princeton, NJ).
- Chen, M., and G. Mussalli, 2020, An integrated approach to quantitative ESG investing, *The Journal of Portfolio Management* 46, 65–74.
- Chen, N.-F., R. Roll, and S. A. Ross, 1986, Economic forces and the stock market, *Journal of Business* 383–403.
- Chen, T., H. Dong, and C. Lin, 2020, Institutional shareholders and corporate social responsibility, Journal of Financial Economics 135, 483–504.
- Cochrane, J., 2009, Asset pricing: Revised edition (Princeton university press).
- Connor, G., M. Hagmann, and O. Linton, 2012, Efficient semiparametric estimation of the famafrench model and extensions, *Econometrica* 80, 713–754.
- Cornell, B., and A. Damodaran, 2020, Valuing ESG: Doing good or sounding good?, The Journal of Impact and ESG Investing 1, 76–93.
- David, H., 1973, Concomitants of order statistics, *Bulletin of the International Statistical Institute* 45, 295–300.
- David, H. A., and H. N. Nagaraja, 2004, Order Statistics (John Wiley & Sons).
- de Franco, C., 2020, ESG controversies and their impact on performance, *The Journal of Investing* 29, 33–45.
- Dimson, E., O. Karakaş, and X. Li, 2015, Active ownership, The Review of Financial Studies 28, 3225–3268.

- Dortfleitner, G., G. Halbritter, and M. Nguyen, 2015, Measuring the level and risk of corporate responsibility–an empirical comparison of different ESG ratings approaches, *The Journal of Asset Management* 17, 450–466.
- Dunn, J., S. Fitzgibbons, and L. Pomorski, 2018, Assessing risk through environmental, social and governance exposures, *Journal of Investment Management* 16, 4–17.
- Dybvig, P. H., and J. E. Ingersoll, 1982, Mean-variance theory in complete markets, *Journal of Business* 233–251.
- Dyck, A., K. V. Lins, L. Roth, and H. F. Wagner, 2019, Do institutional investors drive corporate social responsibility? international evidence, *Journal of Financial Economics* 131, 693–714.
- Eichholtz, P., N. Kok, and J. M. Quigley, 2010, Doing well by doing good? Green office buildings, American Economic Review 100, 2492–2509.
- Evan, W. M., and P. Olk, 1990, R&d consortia: A new us organizational form, MIT Sloan Management Review 31, 37.
- Fabozzi, F. J., K. Ma, and B. J. Oliphant, 2008, Sin stock returns, The Journal of Portfolio Management 35, 82–94.
- Fagnan, D. E., J. M. Fernandez, A. W. Lo, and R. M. Stein, 2013, Can financial engineering cure cancer?, American Economic Review 103, 406–11.
- Falck, O., and S. Heblich, 2007, Corporate social responsibility: Doing well by doing good, Business Horizons 50, 247–254.
- Fama, E. F., and K. R. French, 1993, Common risk factors in the returns on stocks and bonds, Journal of Financial Economics 33, 3–56.
- Fama, E. F., and K. R. French, 2007, Disagreement, tastes, and asset prices, Journal of Financial Economics 83, 667–689.
- Fama, E. F., and K. R. French, 2015, A five-factor asset pricing model, Journal of Financial Economics 116, 1–22.
- Fama, E. F., and J. D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, Journal of Political Economy 81, 607–636.
- Fauver, L., and M. B. McDonald IV, 2014, International variation in sin stocks and its effects on equity valuation, *Journal of Corporate Finance* 25, 173–187.
- Feng, G., S. Giglio, and D. Xiu, 2020, Taming the factor zoo: A test of new factors, *The Journal of Finance* 75, 1327–1370.
- Ferson, W. E., and C. R. Harvey, 1991, The variation of economic risk premiums, Journal of Political Economy 99, 385–415.
- Flammer, C., 2013, Corporate social responsibility and shareholder reaction: The environmental awareness of investors, Academy of Management Journal 56, 758–781.
- Flammer, C., 2021, Corporate green bonds, Journal of Financial Economics 142, 499–516.

- Freyberger, J., A. Neuhierl, and M. Weber, 2020, Dissecting characteristics nonparametrically, The Review of Financial Studies 33, 2326–2377.
- Friedman, H. L., and M. S. Heinle, 2016, Taste, information, and asset prices: Implications for the valuation of csr, *Review of Accounting Studies* 21, 740–767.
- Gagliardini, P., E. Ossola, and O. Scaillet, 2016, Time-varying risk premium in large cross-sectional equity data sets, *Econometrica* 84, 985–1046.
- Gagliardini, P., E. Ossola, and O. Scaillet, 2019, A diagnostic criterion for approximate factor structure, *Journal of Econometrics* 212, 503–521.
- Galema, R., A. Plantinga, and B. Scholtens, 2008, The stocks at stake: Return and risk in socially responsible investment, *Journal of Banking & Finance* 32, 2646–2654.
- Geczy, C. C., R. F. Stambaugh, and D. Levin, 2021, Investing in socially responsible mutual funds, The Review of Asset Pricing Studies 11, 309–351.
- Gibson, R., P. Krueger, and P. S. Schmidt, 2021, ESG rating disagreement and stock returns, *Financial Analysts Journal* 1–24.
- Giese, G., L.-E. Lee, D. Melas, Z. Nagy, and L. Nishikawa, 2019, Foundations of ESG investing: How ESG affects equity valuation, risk, and performance, *The Journal of Portfolio Management* 45, 69–83.
- Giglio, S., B. Kelly, and J. Stroebel, 2021, Climate finance, Annual Reviews of Financial Economics 13, 15–36.
- Goldstein, I., A. Kopytov, L. Shen, and H. Xiang, 2021, On ESG investing: Heterogeneous preferences, information, and asset prices, Available at SSRN 3823042.
- Goss, A., and G. S. Roberts, 2011, The impact of corporate social responsibility on the cost of bank loans, *Journal of Banking & Finance* 35, 1794–1810.
- Green, J., J. R. Hand, and X. F. Zhang, 2017, The characteristics that provide independent information about average us monthly stock returns, *The Review of Financial Studies* 30, 4389–4436.
- Grinold, R. C., 1989, The fundamental law of active management, *The Journal of Portfolio Management* 15, 30–37.
- Grinold, R. C., 1994, Alpha is volatility times IC times score, *Journal of Portfolio Management* 20, 9.
- Grinold, R. C., and R. N. Kahn, 1999, Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk, second edition (McGraw-Hill Education, New York).
- Grinold, R. C., and R. N. Kahn, 2019, Advances in Active Portfolio Management: New Developments in Quantitative Investing (McGraw Hill Professional).
- Grossman, A., S. Appleby, and C. Reimers, 2013, Venture philanthropy: Its evolution and its future, Harvard Business School, N9-313-111.

- Grossman, S. J., and J. E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.
- Gu, S., B. Kelly, and D. Xiu, 2020, Empirical asset pricing via machine learning, The Review of Financial Studies 33, 2223–2273.
- Hartzmark, S. M., and A. B. Sussman, 2019, Do investors value sustainability? a natural experiment examining ranking and fund flows, *The Journal of Finance* 74, 2789–2837.
- Harvey, C. R., Y. Liu, and H. Zhu, 2016, ... and the cross-section of expected returns, *The Review* of Financial Studies 29, 5–68.
- Hasso, T., D. Müller, M. Pelster, and S. Warkulat, 2021, Who participated in the gamestop frenzy? evidence from brokerage accounts, *Finance Research Letters* 102140.
- He, Z., B. Kelly, and A. Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.
- Heinkel, R., A. Kraus, and J. Zechner, 2001, The effect of green investment on corporate behavior, Journal of Financial and Quantitative Analysis 36, 431–449.
- Hong, H., and M. Kacperczyk, 2009, The price of sin: The effects of social norms on markets, Journal of Financial Economics 93, 15–36.
- Hong, H., G. A. Karolyi, and J. A. Scheinkman, 2020, Climate finance, The Review of Financial Studies 33, 1011–1023.
- Hou, K., C. Xue, and L. Zhang, 2015, Digesting anomalies: An investment approach, The Review of Financial Studies 28, 650–705.
- Hou, K., C. Xue, and L. Zhang, 2020, Replicating anomalies, The Review of Financial Studies 33, 2019–2133.
- Idzorek, T. M., P. D. Kaplan, and R. G. Ibbotson, 2021, The popularity asset pricing model, Available at SSRN 3451554.
- Irwin, D. A., and P. J. Klenow, 1996a, High-tech r&d subsidies estimating the effects of sematech, Journal of International Economics 40, 323–344.
- Irwin, D. A., and P. J. Klenow, 1996b, Sematech: purpose and performance, Proceedings of the National Academy of Sciences 93, 12739–12742.
- Jagannathan, R., and Z. Wang, 1996, The conditional capm and the cross-section of expected returns, *The Journal of Finance* 51, 3–53.
- Jegadeesh, N., and S. Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65–91.
- Kempf, A., and P. Osthoff, 2007, The effect of socially responsible investing on portfolio performance, European Financial Management 13, 908–922.
- Khan, M., G. Serafeim, and A. Yoon, 2016, Corporate sustainability: First evidence on materiality, *The Accounting Review* 91, 1697–1724.

- Kim, E., and A. W. Lo, 2019, Venture philanthropy: a case study of the cystic fibrosis foundation, Available at SSRN 3376673.
- Klassen, R. D., and C. P. McLaughlin, 1996, The impact of environmental management on firm performance, *Management science* 42, 1199–1214.
- Kozak, S., S. Nagel, and S. Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.
- Krüger, P., 2015, Corporate goodness and shareholder wealth, Journal of Financial Economics 115, 304–329.
- Kuntz, L.-C., 2020, Beta dispersion and market timing, Journal of Empirical Finance 59, 235–256.
- Kyle, A. S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- Lee, H.-M., and M. Viana, 1999, The joint covariance structure of ordered symmetrically dependent observations and their concomitants of order statistics, *Statistics & probability letters* 43, 411–414.
- Lettau, M., and S. Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, The Journal of Finance 56, 815–849.
- Letts, C. W., W. Ryan, and A. Grossman, 1997, Virtuous capital: What foundations can learn from venture capitalists, *Harvard Business Review* 75, 36–50.
- Lewellen, J., S. Nagel, and J. Shanken, 2010, A skeptical appraisal of asset pricing tests, Journal of Financial Economics 96, 175–194.
- Lillo, F., J. D. Farmer, and R. N. Mantegna, 2003, Master curve for price-impact function, *Nature* 421, 129–130.
- Lindsey, L. A., S. Pruitt, and C. Schiller, 2021, The cost of esg investing, *Available at SSRN* 3975077.
- Link, A. N., D. J. Teece, and W. F. Finan, 1996, Estimating the benefits from collaboration: The case of sematech, *Review of Industrial Organization* 11, 737–751.
- Lintner, J., 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13–37.
- Lo, A. W., 2004, The adaptive markets hypothesis, Journal of Portfolio Management 30, 15–29.
- Lo, A. W., 2017, Adaptive Markets: Financial Evolution at the Speed of Thought (Princeton University Press, Princeton, NJ).
- Lo, A. W., and A. C. MacKinlay, 1990, Data-snooping biases in tests of financial asset pricing models, *The Review of Financial Studies* 3, 431–467.
- Lo, A. W., and R. T. Thakor, 2019, Risk and reward in the orphan drug industry, The Journal of Portfolio Management 45, 30–45.
- Lo, A. W., L. Wu, R. Zhang, and C. Zhao, 2022, Optimal impact portfolios with general dependence and marginals, Available at SSRN 4177277.

- Lo, A. W., R. Zhang, and C. Zhao, 2022, Measuring and optimizing the risk and reward of green portfolios, *The Journal of Impact and ESG Investing* 3, 55–93.
- Luo, H. A., and R. J. Balvers, 2017, Social screens and systematic investor boycott risk, Journal of Financial and Quantitative Analysis 52, 365–399.
- Lyócsa, S., E. Baumöhl, and T. Výrost, 2022, Yolo trading: Riding with the herd during the gamestop episode, *Finance Research Letters* 46, 102359.
- MacKinlay, A. C., 1987, On multivariate tests of the capm, *Journal of Financial Economics* 18, 341–371.
- Madhavan, A., A. Sobczyk, and A. Ang, 2021, Toward ESG alpha: Analyzing ESG exposures through a factor lens, *Financial Analysts Journal* 77, 69–88.
- Merton, R., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- Merton, R. C., 1987, A simple model of capital market equilibrium with incomplete information, *The Journal of Finance* 42, 483–510.
- Pástor, L., and R. F. Stambaugh, 1999, Costs of equity capital and model mispricing, *The Journal of Finance* 54, 67–121.
- Pástor, L., and R. F. Stambaugh, 2003, Liquidity risk and expected stock returns, Journal of Political Economy 111, 642–685.
- Pástor, L., R. F. Stambaugh, and L. A. Taylor, 2021, Sustainable investing in equilibrium, Journal of Financial Economics 142, 550–571.
- Pástor, L., R. F. Stambaugh, and L. A. Taylor, 2022, Dissecting green returns, Journal of Financial Economics 146, 403–424.
- Pedersen, L. H., 2022, Game on: Social networks and markets, *Journal of Financial Economics* 146, 1097–1119.
- Pedersen, L. H., S. Fitzgibbons, and L. Pomorski, 2021, Responsible investing: The ESG-efficient frontier, *Journal of Financial Economics*.
- Raponi, V., C. Robotti, and P. Zaffaroni, 2020, Testing beta-pricing models using large crosssections, *The Review of Financial Studies* 33, 2796–2842.
- Renneboog, L., J. Ter Horst, and C. Zhang, 2008, Socially responsible investments: Institutional aspects, performance, and investor behavior, *Journal of Banking & Finance* 32, 1723–1742.
- Ross, S., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.
- Ross, S. A., 1978, A simple approach to the valuation of risky streams, *Journal of business* 453–475.
- Salzman, R., 2016, Venture philanthropy and gene therapy: lessons from adrenoleukodystrophy, Human Gene Therapy 27, 14–18.

- Scaife, W., 2008, Venturing into venture philanthropy: Is more sustainable health and medical research funding possible through venture philanthropy and social entrepreneurship?, Journal of Nonprofit & Public Sector Marketing 20, 245–260.
- Scott, J., 2009, The politics of venture philanthropy in charter school policy and advocacy, *Educa*tional Policy 23, 106–136.
- Semenova, N., and L. G. Hassel, 2015, On the validity of environmental performance metrics, Journal of Business Ethics 132, 249–258.
- Shanken, J., 1992, On the estimation of beta-pricing models, *The Review of Financial Studies* 5, 1–33.
- Sharp, P., and S. Hockfield, 2017, Convergence: The future of health, *Science* 355, 589–589.
- Sharpe, W., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, The Journal of Finance 19, 425–442.
- Sharpe, W. F., 1966, Mutual fund performance, The Journal of Business 39, 119–138.
- Shiller, R. J., 2005, *Irrational Exuberance*, second edition (Princeton University Press, Princeton, New Jersey).
- Shing, E., 2021, Does sustainable investing have to cost performance?, BNP Paribas Wealth Management, https://ifa.bnpparibas.com/asia/en/expert-voices/ sustainable-investing-and-performance.html, accessed on August 2, 2021.
- Sorensen, E., M. Chen, and G. Mussalli, 2021, The quantitative approach for sustainable investing, The Journal of Portfolio Management.
- Sortino, F. A., and L. N. Price, 1994, Performance measurement in a downside risk framework, *The Journal of Investing* 3, 59–64.
- Sortino, F. A., and R. Van Der Meer, 1991, Downside risk, Journal of Portfolio Management 17, 27.
- Statman, M., and D. Glushkov, 2009, The wages of social responsibility, *Financial Analysts Journal* 65, 33–46.
- Thakor, R. T., N. Anaya, Y. Zhang, C. Vilanilam, K. W. Siah, C. H. Wong, and A. W. Lo, 2017, Just how good an investment is the biopharmaceutical sector?, *Nature Biotechnology* 35, 1149–1157.
- Treynor, J. L., and F. Black, 1973, How to use security analysis to improve portfolio selection, *The Journal of Business* 46, 66–86.
- Tu, J., and G. Zhou, 2011, Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies, *Journal of Financial Economics* 99, 204–215.
- Umar, Z., M. Gubareva, I. Yousaf, and S. Ali, 2021, A tale of company fundamentals vs sentiment driven pricing: The case of gamestop, *Journal of Behavioral and Experimental Finance* 30, 100501.
- Van Slyke, D. M., and H. K. Newman, 2006, Venture philanthropy and social entrepreneurship in community redevelopment, *Nonprofit Management and Leadership* 16, 345–368.
- Xiong, J. X., 2021, The impact of ESG risk on stocks, *The Journal of Impact and ESG Investing* 1.
- Yang, S.-S., 1977, General distribution theory of the concomitants of order statistics, The Annals of Statistics 996–1002.
- Yogo, M., 2006, A consumption-based explanation of expected stock returns, The Journal of Finance 61, 539–580.
- Zerbib, O. D., 2022, A sustainable capital asset pricing model (S-CAPM): Evidence from environmental integration and sin stock exclusion, *Review of Finance* 26, 1345–1388.