Journal of Financial Economics 142 (2021) 756–774

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

The benchmark inclusion subsidy☆

Anil K Kashyap[a,b,c], Natalia Kovrijnykh[d], Jian Li[e], Anna Pavlova[f,*]

[a] Booth School of Business, University of Chicago, National Bureau of Economic Research, United States
[b] Centre for Economic Policy Research, United Kingdom
[c] Bank of England, United Kingdom
[d] Department of Economics, Arizona State University, United States
[e] Department of Economics, University of Chicago, United States
[f] London Business School and Centre for Economic Policy Research, United Kingdom

ARTICLE INFO

Article history:
Received 2 April 2020
Revised 11 September 2020
Accepted 15 October 2020
Available online 30 April 2021

JEL classification:
G11
G12
G23
G32
G34

Keywords:
Project valuation
Investment
Mergers
Asset management
Benchmark

ABSTRACT

We argue that the pervasive practice of evaluating portfolio managers relative to a benchmark has real effects. Benchmarking generates additional, inelastic demand for assets inside the benchmark. This leads to a "benchmark inclusion subsidy:" a firm inside the benchmark values an investment project more than the one outside. The same wedge arises for valuing M&A, spinoffs, and IPOs. This overturns the proposition that an investment's value is independent of the entity considering it. We describe the characteristics that determine the subsidy, quantify its size (which could be large), and identify empirical work supporting our model's predictions.

1. Introduction

The asset management industry is estimated to control more than $85 trillion worldwide. Most of this money is managed against benchmarks. For instance, S&P Global reports that in 2017 just under $10 trillion were managed against the S&P 500 alone. Existing research related to benchmarking has largely focused on its asset pricing implications. Our contribution is to study the implications

☆ We thank the Editor (Bill Schwert) and two anonymous referees for their excellent comments. We have also benefited from comments by Jan Bena, Jules van Binsbergen, Andrea Buffa, Hui Chen, Simon Gervais, Denis Gromb, Marcin Kacperczyk, Steve Kaplan, Ralph Koijen, Felix Kubler, Shohini Kundu, Ailsa Roell, Jeremy Stein, David Thesmar, Dimitri Vayanos, Rob Vishny, Costas Xiouros, and seminar participants at ASU, Chicago Booth, LBS, NBIM, Philadelphia Fed, University of Zurich, AFA 2020 Annual Meeting, the 5th BI-SHOF conference, the 2019 Bank of England Research Forum on Macro-Finance, Cambridge Corporate Finance Theory Symposium 2019, the 2019 European Winter Finance conference, ISF 2019, the Macro Financial Modeling (MFM) Winter 2019 Meeting, Fall 2018 NBER Corporate Finance Program Meeting, the 2019 NBER Long-Term Asset Management Meeting, the 2019 Paul Woolley Centre conference, and the 2019 UBC Winter Finance conference. We thank Taisiya Sikorskaya for excellent research assistance. This research has been supported by a grant from the Alfred P. Sloan Foundation to the MFM project at the University of Chicago. The views expressed here are ours only and not necessarily those of the institutions with which we are affiliated, and all mistakes are our own.

* Corresponding author.
E-mail addresses: anil.kashyap@chicagobooth.edu (A.K. Kashyap), natalia.kovrijnykh@asu.edu (N. Kovrijnykh), lijian@uchicago.edu (J. Li), apavlova@london.edu (A. Pavlova).

https://doi.org/10.1016/j.jfineco.2021.04.021
0304-405X/

of benchmarking for corporate decisions such as investments, M&A, spinoffs, and IPOs. We argue that firms included in a benchmark are effectively subsidized by asset managers and so should evaluate investment opportunities differently.

A familiar proposition states that the value of an investment is independent of the entity considering it. We show that this result breaks down when some investors use portfolio (or fund) managers, whose performance is measured against a benchmark. Instead, the appropriate discount rate for an investment project differs across firms, depending on whether or not they are part of the benchmark. Specifically, when a firm adds risky cash flows, say because of an acquisition or by investing in a new project, the increase in the stockholder value is larger if the firm is inside the benchmark rather than outside. We call this the "benchmark inclusion subsidy." Because of this subsidy, a firm in the benchmark would accept cash flows with lower mean and/or larger variance than a non-benchmark firm would.

The reason for this result is that benchmarking induces portfolio managers to hold some shares of firms in their benchmarks regardless of those firms' cash-flow characteristics. So when a firm adds risky cash flows, the market demand for claims on these cash flows is higher if the firm is inside the benchmark rather than outside. The firm, therefore, should take this consideration into account in deciding on its investments, acquisitions and spinoffs. These implications differ from related research that identifies the effect of benchmarking on the value of securities associated with existing assets.

Here is how the model works. We introduce portfolio managers into a standard model, extended to allow for investment and production. Some investors manage their own portfolios and others use managers. A portfolio manager's compensation depends on performance relative to a benchmark portfolio. This assumption, critical for our argument, is in keeping with prevalent industry practices and evidence such as Ma et al. (2019). We show that portfolio managers optimally choose to hold a combination of the standard mean-variance portfolio and the benchmark portfolio, the latter appearing because of benchmarking in their compensation. Specifically, managers hold a fixed part of their portfolio in benchmark firms regardless of their stock prices and cash-flow characteristics,most importantly, irrespective of cash-flow variance. As a result, the equilibrium stock price of a firm in the benchmark is less adversely affected by the same cash-flow risk than that of an otherwise identical firm that is outside the benchmark.

Within our production economy, we show that the same argument applies to risky investments. Specifically, for a firm in the benchmark, risky investments are subsidized. The extra variance of its cash flows resulting from the project will be penalized less than that of an identical non-benchmark firm. This occurs because managers have an inelastic demand for a fixed number of shares of benchmark firms, regardless of their risk. This inelastic demand benefits riskier projects more. As cash-flow riskiness decreases, so does the benchmark inclusion subsidy, converging to zero for risk-free projects.

A testable implication of the model is that because of the subsidy, a firm in the benchmark is more likely to invest. Likewise, a benchmark firm would also accept projects with riskier cash flows.

To demonstrate these predictions in the most transparent way, we construct the simplest possible example. In it, there are three firms with uncorrelated cash flows. Absent portfolio managers, combining firms creates no value. With portfolio managers, a firm inside the benchmark creates value when it takes over a firm that is outside the benchmark.

We then turn to an extended model with many firms whose cash flows are correlated. Allowing for correlations brings out additional effects and predictions. In particular, we find that investments whose cash flows are positively correlated with those of the benchmark firms are valued more. As in the economy with uncorrelated cash flows, portfolio managers' excess demand for benchmark stocks raises those stocks' prices. Now, the stock prices of firms whose cash flows are positively correlated with those of the benchmark stocks also rise. This happens because, in seeking exposure to the benchmark's cash-flow risk, investors substitute away from expensive benchmark stocks into stocks that are correlated with them. This same reasoning means that investment projects or acquisitions that are positively correlated with the benchmark are valued more by all firms (relative to an economy without portfolio managers).

We show that the main mechanism that delivered the benchmark inclusion subsidy in the example generalizes. The extended model also features an additional channel, owing to the correlation of a project's cash flows with those of the firm's assets-in-place. We demonstrate that if this correlation is positive, the value of the new asset to the benchmark firm exceeds the asset's value were it to join the benchmark as a standalone.

We derive a closed-form expression for the benchmark inclusion subsidy, which turns out to be very simple, and study the variables that influence its size. We show that the higher the cash-flow risk of an investment, the larger the benchmark inclusion subsidy. Furthermore, the benchmark inclusion subsidy is increasing in the correlation of a project's cash flows with the existing assets; in particular, the subsidy is the largest for projects that are clones of a firm's existing assets. Finally, the size of the subsidy rises with the size of the asset management sector.

The ability to characterize the exact determinants of the subsidy allows us to predict the situations when benchmarking is the most and least important. To the best of our knowledge, other theories do not deliver such cross-sectional predictions. We are able to tie the size of the subsidy to characteristics of firms, investment opportunities, or potential acquisitions or divestitures.

For M&A, we stress that a merger of a firm in the benchmark with a firm outside creates value for stockholders. There are no cash-flow synergies, but there is a financial synergy that comes from bringing the target into the benchmark. This contrasts with textbook analyses arguing that such a merger creates no value. For IPOs, we highlight the implication that an IPO would have a higher value if the stock joins the benchmark. For spinoffs, we demon-

strate that if an asset divested by a benchmark firm lands outside the benchmark, it is stripped of the benchmark inclusion subsidy and hence its value drops.

It is worth noting that our model applies to both active and passive asset management. We show that the benchmark inclusion subsidy is larger when more managers are passive rather than active. While active managers in our model invest only a part of their portfolio in the benchmark, passive managers invest their entire portfolio in it, which amplifies the inelastic demand for stocks in the benchmark.

Our model also points to the importance of benchmark design for ESG (environmental, social, and governance) funds. As long as a company with a low ESG score remains part of the benchmark, it receives the benchmark inclusion subsidy. Dropping such "brown" companies from the benchmark and replacing them with "green" companies not only gives a cost-of-capital advantage to the green companies but also encourages other firms in the economy to mimic the green companies. The latter implication is due to our result that cash flows that are positively correlated with those of benchmark firms have higher values.

We attempt to quantify the size of the subsidy. Doing so requires us to pin down a number of model parameters whose values are not obvious. However, the primary determinant of the subsidy's size is the amount of the inelastic demand that the portfolio managers have for stocks inside the benchmark. According to the FactSet/LionShares database, in 2017, stocks belonging to the S&P 500 have institutional ownership of 83%, whereas institutional ownership for the overall market is 67%. We take this gap to be indicative of the inelastic demand for benchmark stocks. When we ask the model to account for this demand, as well as the relative size of firms inside and outside the benchmark, and of estimates of the spread in market beta's observed in the data, the implied subsidy is 94 basis points. We explore how alternative parameter choices change this number and conclude that the subsidy could easily be large enough to be important for the kind of corporate decisions that we have analyzed.

Finally, we briefly review existing empirical work that relates to the model's predictions. Past research confirms, to varying degrees, the predictions regarding the propensity to invest and engage in acquisitions for benchmark vs. non-benchmark firms, the factor structure of returns, as well as the size of the benchmark inclusion subsidy increasing in assets under management.

The remainder of the paper is organized as follows. In the next section, we explain how our perspective compares to previous work. Section 3 presents the example, and Section 4 analyzes the general model. Section 5 describes how we quantify the size of the subsidy. Section 6 reviews related empirical evidence. Section 7 presents our conclusions and suggestions for future areas of promising research. Omitted proofs and derivations are in the appendicies.

2. Related literature

Our paper revisits one of the main tenets of corporate finance that the value of an investment is independent of the entity considering it. This point is emphasized in leading textbooks such as Berk and DeMarzo (2017), Brealey et al. (2019), and Grinblatt and Titman (2002). There is a huge academic and applied literature on capital budgeting (see, e.g. Stulz, 1999, Jacobs and Shivdasani, 2012, Dessaint et al., 2018, for a range of representative studies). Finally, there is survey evidence (e.g. Graham and Harvey, 2001) on how managers make such decisions. In these analyses, asset management plays no role.

Our analysis shows that the textbook approach's failure to properly account for the effects of asset management has important consequences for real corporate decisions. The motivation for taking asset management seriously comes from the index additions and deletions literature. Harris and Gurel (1986) and Shleifer (1986) were the first to document that when stocks are added to the S&P 500 index, their prices rise. Subsequent papers have shown that firms that are deleted experience a decline in price. The findings have been confirmed across many studies and for many markets, so that financial economists consider these patterns to be stylized facts.[1] The estimated magnitudes of the index effect vary across studies, and typically most of the effect is permanent. For example, Chen et al. (2004) find the cumulative abnormal returns of stocks added to the S&P 500 during 1989–2000, measured over two months post announcement, to be 6.2%.[2]

Several theories have been used to interpret the index effect. The first is the investor awareness theory of Merton (1987). Merton posits that some investors become aware of and invest in a stock only when it gets included in a popular index. It is unclear why investor awareness declines for index deletions, although there is evidence of a decrease in analyst coverage. The second theory posits that index inclusions convey information about a firm's improved prospects. This theory has difficulty explaining the presence of index effects around mechanical index recompositions (see, e.g., Boyer, 2011, among others). The third theory is that index inclusion leads to improved liquidity, and this in turn boosts stock prices. This theory, however, does not explain increased correlations with other index stocks (documented in, e.g., Barberis et al., 2005 and Boyer, 2011).

The final theory can be broadly described as the price pressure theory, proposed by Scholes (1972). Scholes' prediction is that prices of included stocks should rise temporarily to compensate liquidity providers, but should revert back as investors find substitutes for these stocks. Subsequent literature has argued that the price pressure effects could be (more) permanent, driven by changing compositions of investors. Our model is broadly consistent with the price pressure view. Our benchmarked managers put permanent upward pressure on prices of stocks as long

[1] See, e.g, Beneish and Whaley (1996), Lynch and Mendenhall (1997), Wurgler and Zhuravskaya (2002), Chen et al. (2004), Petajisto (2011), and Hacibedel (2018).

[2] Calomiris et al. (2019) find an index effect for emerging market corporate bonds. They trace the rise of the JP Morgan Corporate Emerging Market Bond Index and show how firms in countries that became eligible change issuance patterns (to qualify for index inclusion) and pay lower yields on qualifying bond issues.

as they are in the benchmark. Despite any overpricing, benchmarking creates a fixed demand for these stocks by a particular clientele, the portfoio managers. Holding a substitute stock is costly for a manager because this entails a (risky) deviation from her benchmark.

The index effect literature only considers the average effect of index inclusion. Our theory has a host of cross-sectional predictions that one could potentially test. For example, stocks with larger cash-flow variance should experience a larger index effect. We also stress that what matters for our channel is whether a stock is in the benchmark, not the index. One can separate competing theories by studying stocks that are in the index but not the benchmark (e.g., "sin" stocks, as analyzed in Hong and Kacperczyk, 2009).

Our work is also related to a theoretical literature in asset pricing that explores the effects of delegated asset management and benchmarking on stock returns and their comovement. The first paper in this line of research is Brennan (1993), who derives a two-factor CAPM in an economy with portfolio managers. Cuoco and Kaniel (2011), Basak and Pavlova (2013), Buffa et al. (2014), and Buffa and Hodor (2018) show how benchmarking creates additional demand for stocks included in the benchmark index, generating an index effect. Basak and Pavlova also derive excess comovement of index stocks, and Buffa and Hodor introduce heterogeneous benchmarks and investigate the resulting rich structure of asset-price comovements within and across benchmarks. Greenwood (2005) considers a model with passive indexers and arbitrageurs (who are like our direct investors) and shows that an index reconstitution not only lifts prices of stocks added to the index but also those of non-index stocks that are positively correlated with them. Kaniel and Kondor (2013) study equilibrium trading strategies in an environment in which fund managers face a convex flow-performance relation. This literature considers pure exchange economies and focuses on the effects on *existing* assets, while ours is a production economy, in which we study the effects of benchmarking on *new* assets that are created through investment and other corporate decisions.

Kashyap et al. (2020) study the design of optimal (linear) contracts for portfolio managers and the resulting welfare implications. In contrast, this paper takes the contracts as exogenous and looks at their implications for corporate decisions.

Our paper is perhaps most closely related to Stein (1996). He also studies capital budgeting in situations where the CAPM does not correctly describe expected stock returns. He assumes, however, that the deviations are temporary and arise because of investor irrationality.[3] If market participants fail to appreciate risk and allow a firm to issue mispriced equity, he explains why rational managers may want to issue equity and invest even if the CAPM-based valuation of a project is negative. In our case, the price differences are not due to irrationality. Instead, they arise because of fundamental differences in demand from different types of investors. In Stein's setup the horizon that managers use for making decisions is critical, and those that are short-term oriented will potentially respond to mispricing if it is big enough. In our model, all managers of firms in the benchmark should account for the subsidy (for as long as the firm remains in the benchmark).[4]

Stein's paper led to a number of follow-up studies that look at other potential behavioral effects that could be associated with inclusion in a benchmark (see Baker and Wurgler, 2013 for a survey). These papers contain much of the empirical work that we cite in favor of our model. While we share several predictions with Stein (1996) there are some notable differences. For instance, Stein's model connects managerial time horizons and financial constraints to capital budgeting decisions. Our model has nothing to say about these considerations. However, we also have many implications that are distinct from his. For instance, our closed-form expression for the benchmark inclusion subsidy generates a number of predictions about which factors should lead firms in the benchmark to make different decisions than ones outside. On the whole, we see the behavioral theories and ours complementing each other.

3. Example

To illustrate the main mechanism, we begin with a simple example with three firms with uncorrelated cash-flows. We first consider an economy populated by identical investors in these firms who manage their own portfolios. We then modify the economy by introducing another group of investors who hire portfolio managers to run their portfolios. Portfolio managers' performance is evaluated relative to a benchmark. We show that the presence of portfolio managers changes some familiar corporate valuation principles.

3.1. Baseline economy

Consider the following environment. There are two periods, $t = 0, 1$. Investment opportunities are represented by three risky assets denoted by 1, 2, and y, and one risk-free bond. The risky assets are claims to cash flows D_i realized at $t = 1$, where $D_i \sim N(\mu_i, \sigma_i^2)$, $i = 1, 2, y$, and these cash flows are uncorrelated. We think of these assets as stocks of all-equity firms. There is a risk-free bond that pays an interest rate that is normalized to zero. Each of the risky assets is available in a fixed supply that is normalized to one. The bond is in infinite net supply. Let S_i denote the price of asset $i = 1, 2, y$.

There is measure one of identical agents who invest their own funds. Each investor has a constant absolute risk aversion (CARA) utility function over final wealth W, $U(W) = -e^{-\gamma W}$, where $\gamma > 0$ is the coefficient of absolute risk aversion. All investors are endowed with one share of

[3] van Binsbergen and Opp (2019) study the effects of other anomalies that are associated with mispricing on corporate decisions. However, they do not take a stand on the source of these anomalies.

[4] This stands in contrast to the recent literature on mistakes that managers make in project valuation. For example, Krüger et al. (2015) document that diversified firms appear to make investment decisions in non-core businesses by using the discount rate from their core business.

each stock and no bonds. At $t = 0$, each investor chooses a portfolio of stocks $x = (x_1, x_2, x_y)^\top$ and the bond holdings to maximize his utility, with $W(x) = \sum_{i=1,2,y} S_i + x_i(D_i - S_i)$.

As is well-known in this kind of setup, the demand x_i for risky asset i and the corresponding equilibrium price S_i will be

$$x_i = \frac{\mu_i - S_i}{\gamma \sigma_i^2}, \tag{1}$$

$$S_i = \mu_i - \gamma \sigma_i^2 \tag{2}$$

for $i = 1, 2, y$, where the second equation follows from setting the number of shares demanded equal to the supply (which is 1).[5]

When firms $i \in \{1, 2\}$ and y merge into a single firm, the demand for the combined firm's stock and the corresponding equilibrium stock price are

$$x_i' = \frac{\mu_i + \mu_y - S_i'}{\gamma(\sigma_i^2 + \sigma_y^2)}, \tag{3}$$

$$S_i' = \mu_i + \mu_y - \gamma(\sigma_i^2 + \sigma_y^2) = S_i + S_y. \tag{4}$$

Notice that the combined value of either firm is exactly equal to the sum of its initial value plus the value of y. This is the familiar result that says that the owner of a firm does not determine its value. Instead, the value arises from the cash flows and risks associated with those cash flows, which are the same regardless of who owns them.

3.2. Adding portfolio managers

Now we extend the example by assuming that some investors hire portfolio managers to manage their portfolios. There are now three types of agents in the economy, the same investors as before who manage their own portfolios and whom we refer to as "direct" investors from now on (constituting a fraction λ_D of the population), portfolio (or fund) managers (a fraction λ_M), and fund investors who hire those managers (a fraction λ_F).[6] All agents have the same preferences (as in the prior example).

Fund investors can buy the bond directly, but cannot trade stocks. They delegate the selection of their portfolios to portfolio managers. The managers receive compensation w from shareholders. This compensation has three parts: one is a linear payout based on absolute performance of the portfolio x, the second piece depends on the performance relative to the benchmark portfolio, and the third is independent of performance.[7] For concreteness, suppose that the benchmark is simply the stock of firm 1. Then

$$w = ar_x + b(r_x - r_{\mathbf{b}}) + c = (a+b)r_x - br_{\mathbf{b}} + c, \tag{5}$$

where $a \geq 0$, $b > 0$ and c are constants, $r_x = \sum_{i=1,2,y} x_i(D_i - S_i)$ and $r_{\mathbf{b}} = D_1 - S_1$. For simplicity, we assume that a, b, and c are set exogenously.[8]

A direct investor's demand for asset i continues to be

$$x_i^D = \frac{\mu_i - S_i}{\gamma \sigma_i^2}, \quad i = 1, 2, y. \tag{6}$$

A manager's demands for stocks inside and outside the benchmark are

$$x_1^M = \frac{1}{a+b}\frac{\mu_1 - S_1}{\gamma \sigma_1^2} + \frac{b}{a+b}, \tag{7}$$

$$x_i^M = \frac{1}{a+b}\frac{\mu_i - S_i}{\gamma \sigma_i^2}, \quad i = 2, y. \tag{8}$$

Managers' portfolio choices differ from those of the direct investors in two ways. First, they hold a scaled version of the same mean-variance portfolio as the one held by the direct investors. The reason for the scaling is that, as we can see from the first term in (5), for each share that the manager holds, she gets a fraction $a + b$ of the total return. Thus the manager scales her asset holdings by $1/(a+b)$ relative to those of a direct investor.

Second, and more importantly, the portfolio managers are penalized by b for underperforming the benchmark. Because of this penalty, the manager always holds $b/(a+b)$ shares of stock 1 (or more generally whatever is in the benchmark), regardless of its risk-return characteristics. This consideration explains the second term in (7). This inelastic demand for the benchmark will be critical for all of our results. In particular, the managers' incentive to hold the benchmark portfolio (regardless of the risk characteristics of its constituents) creates an asymmetry between stocks in the benchmark and all other stocks.

The second implication is very general and extends beyond our model with CARA preferences. Having a relative performance component as part of her compensation exposes the manager to an additional source of risk, fluctuations in the benchmark, which she optimally decides to hedge. The manager would, therefore, hold a hedging portfolio that is (perfectly) correlated with the benchmark, i.e., the benchmark itself.

Given the demands, we can now solve for the equilibrium prices. Using the market-clearing conditions, $\lambda_M x_i^M + \lambda_D x_i^D = 1$, $i = 1, 2, y$, we find

$$S_1 = \mu_1 - \gamma \Lambda \sigma_1^2 \left(1 - \lambda_M \frac{b}{a+b}\right), \tag{9}$$

$$S_2 = \mu_2 - \gamma \Lambda \sigma_2^2, \tag{10}$$

$$S_y = \mu_y - \gamma \Lambda \sigma_y^2, \tag{11}$$

where $\Lambda = [\lambda_M/(a+b) + \lambda_D]^{-1}$ modifies the market's effective risk aversion.

[5] We omit derivations for this simple example, but the analysis of our main model contains all proofs for the general case.

[6] For simplicity, we assume that each fund investor employs one manager, so that $\lambda_M = \lambda_F$. (The generalization to $\lambda_M \neq \lambda_F$ is straightforward.) Furthermore, $\lambda_D + \lambda_M + \lambda_F = 1$.

[7] This part captures features such as a fee linked to initial assets under management.

[8] Kashyap et al. (2020) endogenize optimal linear contracts for portfolio managers and show that benchmarking emerges as part of the optimal contract.

For concreteness, suppose that $\mu_1 = \mu_2$ and $\sigma_1 = \sigma_2$ so that the return and risks of stocks 1 and 2 are identical. Our first noteworthy finding is that the share price of firm 1 that is inside the benchmark is higher than that of its twin that is not. This happens because portfolio managers automatically tilt their demand towards the benchmark, effectively reducing the supply of this stock by $b/(a+b)$. The lower the supply of the stock (all else equal), the higher must be its equilibrium price. Another way to understand the result is that the managers' inelastic demand for the benchmark means that the adverse effects of variance that typically reduce the demand for any stock, are less relevant for the assets in the benchmark.[9]

Next, consider potential mergers. Suppose first that firm y is acquired by the non-benchmark firm (firm 2). The new demands of direct investors and portfolio managers for the stock of firm 2 are

$$x_2'^D = \frac{\mu_2 + \mu_y - S_2'}{\gamma\left(\sigma_2^2 + \sigma_y^2\right)}, \tag{12}$$

$$x_2'^M = \frac{1}{a+b}\frac{\mu_2 + \mu_y - S_2'}{\gamma\left(\sigma_2^2 + \sigma_y^2\right)}. \tag{13}$$

The new equilibrium price of firm2ı́s stock is

$$S_2' = \mu_2 + \mu_y - \gamma\Lambda\left(\sigma_2^2 + \sigma_y^2\right) = S_2 + S_y. \tag{14}$$

As before, the combined value of firm 2, continues to be the sum of the initial value plus the value of y.

Suppose instead that y is acquired by firm 1, which is in the benchmark. Re-normalizing the combined number of shares of firm 1 to one, the demands for the stock of the combined firm are

$$x_1'^D = \frac{\mu_1 + \mu_y - S_1'}{\gamma\left(\sigma_1^2 + \sigma_y^2\right)}, \tag{15}$$

$$x_1'^M = \frac{1}{a+b}\frac{\mu_1 + \mu_y - S_1'}{\gamma\left(\sigma_1^2 + \sigma_y^2\right)} + \frac{b}{a+b}. \tag{16}$$

Our next major finding is that there is a benchmark inclusion subsidy. Specifically, the new price of firm1ı́s shares is

$$\begin{aligned} S_1' &= \mu_1 + \mu_y - \gamma\Lambda\left(\sigma_1^2 + \sigma_y^2\right)\left(1 - \lambda_M\frac{b}{a+b}\right) \\ &= S_1 + S_y + \gamma\Lambda\sigma_y^2\lambda_M\frac{b}{a+b}, \end{aligned} \tag{17}$$

which is strictly larger than the sum of S_1 and S_y. So when a firm inside the benchmark acquires y (which had been outside the benchmark), the combined value exceeds the sum of the initial value plus the value of y.[10] We refer to the increment as the *benchmark inclusion subsidy*. This subsidy exists because portfolio managers' demand for the benchmark is partially divorced from the risk and return characteristics of the benchmark, and thus, this kind of acquisition raises the value of the target firm. You can see this by noting that the last term in (17) is proportional to the variance of y, σ_y^2. This is because when y is acquired by firm 1, a portion of managers' demand for y becomes inelastic and independent of its variance. Hence, the market penalizes the variance of y's cash flows less when they are inside firm 1 rather than firm 2.

In contrast, notice that if firm y had started out inside the benchmark, then S_1' would be exactly equal to the sum of prices of stocks 1 and y. In that case, the inelastic demand for the stock would already have been embedded in its price before the merger. So, the extra value of acquisition that accrues to firm 1 relative to firm 2 arises from the increase in the price of y when it becomes part of the benchmark.

To put this more formally, let S_y' denote the price of y if it were inside the benchmark. Then $S_y' - S_y = \gamma\Lambda\sigma_y^2\lambda_M b/(a+b)$, which is precisely the extra term in Eq. (17). Notice that this is directly related to the "index effect" estimated in the literature, which is the percentage change in a firm's stock price when it joins the benchmark, and in our model is

$$\frac{S_y' - S_y}{S_y} = \frac{\gamma\Lambda\sigma_y^2}{S_y}\lambda_M\frac{b}{a+b}, \tag{18}$$

where S_y is given by (11).

Thus, in this simple example, the benchmark inclusion subsidy reduces to the index effect for the target firm. As we will show in the general model in Section 4, if we allow for any correlation between the acquirer's and the target's cash flows, the benchmark inclusion subsidy will have an additional term accounting for the correlation. When the correlation is positive, the subsidy exceeds the index effect for the target firm.

It is worth noting that our model predicts that the index effect is larger for firms with riskier cash flows. This can be seen from Eq. (18), where the index effect is increasing in σ_y^2, even after controlling for the stock price before the inclusion. The literature so far has focused on estimating the average index effect. In contrast, our model makes cross-sectional implications about how the index effect varies with firms' risk characteristics.

Finally, the impact of portfolio managers can also work in the other direction, reducing valuations of spinoffs and divestitures. If y had been part of a firm inside the benchmark and is sold to a non-benchmark firm, the value of y would drop when it is transferred.

In the next section, we consider a richer version of the setup that allows us to analyze several additional questions. Based just on this extremely simplified example, however, we already have seen two empirical predictions. First, consistent with the existing literature on index inclusions, we see that there should be an increase in a

[9] Notice that in this model the asymmetry between benchmark and non-benchmark stocks cannot be arbitraged away. The direct investors are unrestricted in their portfolio choice and therefore can engage in any arbitrage activity. However, as the managers *permanently* reduce the supply of the benchmark stock, direct investors simply reduce their holdings of the benchmark stock and hold more of the non-benchmark stock. As long as managers represent a meaningful fraction of the market (i.e., λ_M is non-negligible), there are always differences in prices of stocks inside and outside the benchmark.

[10] This result does not depend on the firm being entirely equity financed. We assume no debt financing here just for simplicity.

firm's share price when it is added to the benchmark. We view this as a necessary condition for the existence of the benchmark inclusion subsidy. In our framework, the stock price increase would remain present for as long as the firm is part of the benchmark.

The other, more novel, prediction is related to acquisitions (and spinoffs). If a firm that has not previously been part of the benchmark is acquired by a benchmark firm, its value should go up purely from moving into the benchmark. This breaks the usual valuation result which presumes that an asset purchase that does not alter any cash flows (of either the target or acquirer) should not create any value. Alternatively, if a firm were spun-off so that it moves out of the benchmark, its value should drop even though its cash flows are unchanged. Furthermore, these kinds of transactions also have non-trivial welfare implications, see Section 4.6 and Appendix C.

Our results in this paper depend on the compensation contract having a non-zero value of b. There is both direct empirical evidence and strong intuitive reasons for why this assumption should hold. For instance, U.S. mutual funds are required to include a "Statement of Additional Information" in the prospectus that describes how portfolio managers are compensated. Ma et al. (2019) collect this information for active U.S. equity mutual funds and find more than three quarters of the funds explicitly base compensation on performance relative to a benchmark. To see why these results are expected, consider any portfolio manager that runs multiple funds with different characteristics, for instance, a bond fund and an equity fund. To compensate the portfolio managers of each fund, the simple returns cannot be meaningfully compared because of the differences in risk. However, if each fund's performance is adjusted for a benchmark for its type, then the relative performances can be compared. So, it is hardly surprising that the use of benchmarks is so pervasive and our assumption concerning b is not controversial.

4. The general model

We now generalize the example studied in Section 3 in several directions. All results from the previous section hold in this richer model. To analyze a new implication for investment, we will assume that y is not traded initially. So our economy becomes a production economy in which y is interpreted as a potential project.

We will only describe elements of the environment that differ from those described in the previous section. There are n risky stocks, whose total cash flows $D = (D_1, \dots, D_n)^\top$ are jointly normally distributed, $D \sim N(\mu, \Sigma)$, where $\mu = (\mu_1, \dots, \mu_n)^\top$, $\Sigma_{ii} = \text{Var}(D_i) = \sigma_i^2$, and $\Sigma_{ij} = \text{Cov}(D_i, D_j) = \rho_{ij}\sigma_i\sigma_j$. We assume that the matrix Σ is invertible. Stock prices are denoted by $S = (S_1, \dots, S_n)^\top$. For simplicity of exposition and for easier comparison to Section 3, we normalize the total number of shares of each asset to one. However, for generality, all of our proofs in Appendix A are written for the case when asset i's total number of shares is $\bar{x}_i$.

Some stocks are part of a benchmark. We order them so that all shares of the first k stocks are in the benchmark, and none of the remaining $n-k$ stocks are included. Thus, the ith element of the benchmark portfolio equals the total number of shares of asset i times $\mathbf{1}_i$, where $\mathbf{1}_i = 1$ if $i \in \{1, \dots, k\}$ and $\mathbf{1}_i = 0$ if $i \in \{k+1, \dots, n\}$. Denote further $\mathbf{1_b} = (\mathbf{1}_1, \dots, \mathbf{1}_n)^\top = (\underbrace{1, \dots, 1}_{k}, \underbrace{0, \dots, 0}_{n-k})^\top$.

We follow the convention in the literature (see, e.g., Buffa et al., 2014) by defining $r_x = x^\top(D - S)$ to be the performance of portfolio $x = (x_1, \dots, x_n)^\top$ and $r_{\mathbf{b}} = \mathbf{1}_{\mathbf{b}}^\top(D - S)$ to be the performance of the benchmark portfolio. Then the compensation of a portfolio manager with contract (a, b, c) is $w = ar_x + b(r_x - r_{\mathbf{b}}) + c$.[11]

Denote by $x^D = (x_1^D, \dots, x_n^D)^\top$ and $x^M = (x_1^M, \dots, x_n^M)^\top$ the optimal portfolio choices of a direct investor and a portfolio manager, respectively.

Lemma 1 (Portfolio Choice). *Given asset prices S, the demands of a direct investor and a portfolio manager are given by*

$$x^D = \Sigma^{-1}\frac{\mu - S}{\gamma}, \tag{19}$$

$$x^M = \frac{1}{a+b}\Sigma^{-1}\frac{\mu - S}{\gamma} + \frac{b}{a+b}\mathbf{1_b}. \tag{20}$$

The demands generalize those from the example exactly as would be expected. In particular, the direct investors opt for the mean-variance portfolio and the portfolio managers choose a linear combination of that portfolio and the benchmark. The fact that part of the managers' portfolio is invested in the benchmark regardless of prices or other characteristics of these assets will again be crucial for our results below.

An extreme form of our manager is a passive manager—someone who faces a very high b, which incentivizes her to hold just the benchmark portfolio and severely punishes any deviations from it. All of our results hold in this special case, which we discuss in Section 4.5.

Using (19) - (20) and the market-clearing condition $\lambda_M x^M + \lambda_D x^D = \mathbb{1} \equiv (1, \dots, 1)^\top$, we have the following result:

Lemma 2 (Asset Prices). *The equilibrium asset prices are*

$$S = \mu - \gamma\Lambda\Sigma\left(\mathbb{1} - \lambda_M\frac{b}{a+b}\mathbf{1_b}\right). \tag{21}$$

Equation (21) is a generalization of Eqs. (9) - (10). As before, the price of a benchmark firm is higher than it would be for an otherwise identical non-benchmark firm. The reason is that, as Lemma 1 shows, managers demand a larger amount of the stock in the benchmark.

4.1. Investment

Suppose there is a project with cash flows $D_y \sim N(\mu_y, \sigma_y^2)$, and $\text{Corr}(D_i, D_y) = \rho_{iy}$ for $i = 1, \dots, n$. Investing in this project requires spending I. If firm i (whose cash

[11] In Appendix B we repeat all of the analysis for the case where a portfolio manager's compensation is tied to the per-dollar return on the benchmark rather than the per-share return (performance) and confirm that our key results continue to hold.

flows are D_i) invests, its cash flows in period 1 become $D_i + D_y$. Let $S^{(i)} = \left(S_1^{(i)}, \ldots, S_n^{(i)}\right)^\top$ denote the stock prices if firm i invests in the project. The firm finances investment by issuing equity.[12] That is, we assume that if firm i invests in the project, it issues δ_i additional shares to finance it, where $\delta_i S_i^{(i)} = I$. If firm i is in the benchmark, then the additional shares enter the benchmark.

To proceed, suppose firm i (and only firm i) invests in the project. Then the new cash flows are $D^{(i)} = D + (0, \ldots, 0, \underbrace{D_y}_{i}, 0, \ldots, 0)^\top$, distributed according to $N\left(\mu^{(i)}, \Sigma^{(i)}\right)$, where $\mu^{(i)} = \mu + (0, \ldots, 0, \underbrace{\mu_y}_{i}, 0, \ldots, 0)^\top$ and

$$\Sigma^{(i)} = \Sigma + \begin{pmatrix} 0 & \rho_{1y}\sigma_1\sigma_y & 0 & & \\ & & \vdots & & \\ \rho_{1y}\sigma_1\sigma_y & \ldots & \sigma_y^2 + 2\rho_{iy}\sigma_i\sigma_y & \ldots & \rho_{ny}\sigma_n\sigma_y \\ & & \vdots & & \\ 0 & & \rho_{ny}\sigma_n\sigma_y & & 0 \end{pmatrix}. \tag{22}$$

Denote $I^{(i)} = (0, \ldots, 0, \underbrace{I}_{i}, 0, \ldots, 0)^\top$.

Lemma 3 (Post-Investment Asset Prices). *The equilibrium stock prices when firm i invests in the project are given by*

$$S^{(i)} = \mu^{(i)} - I^{(i)} - \gamma\Lambda\Sigma^{(i)}\left(\mathbb{1} - \lambda_M \frac{b}{a+b}\mathbf{1_b}\right). \tag{23}$$

The change in the stockholder value of the investing firm i, $\Delta S_i \equiv S_i^{(i)} - S_i$, is

$$\Delta S_i = \mu_y - I - \gamma\Lambda\left(\sigma_y^2 + \rho_{iy}\sigma_i\sigma_y\right)\left(1 - \lambda_M \frac{b}{a+b}\mathbf{1}_i\right) - \gamma\Lambda\sum_{j=1}^{n}\rho_{jy}\sigma_j\sigma_y\left(1 - \lambda_M\frac{b}{a+b}\mathbf{1}_j\right). \tag{24}$$

The first two terms in the first line of (24) are the expected cash flows of the project net of the cost of investment, and the remaining terms reflect the penalty for risk. It is evident from (24) that this penalty differs if i is part of the benchmark, so it will turn out to be subject to the benchmark inclusion subsidy that we have already seen in the example in Section 3.

Notice that the terms on second line of (24) are the same regardless of the identity of the investing firm. When any firm invests in a project positively correlated with the benchmark, this firm's cash flows become more correlated with the benchmark. The presence of managers makes stocks that covary positively with the benchmark more expensive relative to what they would have been in the economy with only direct investors.[13] This happens because the investors substitute from the expensive stocks in the benchmark to stocks that are correlated with it. A similar logic applies to new investment projects that are correlated with the benchmark: relative to the economy without portfolio managers, firms have a higher valuation for projects whose cash flows are positively correlated with cash flows of benchmark firms. This force, not present in Section 3, applies to all firms contemplating investment, including the non-benchmark ones. We comment further on how it affects firms' investment incentives at the end of this subsection.

We are now ready to derive the benchmark inclusion subsidy in this generalized setting. Consider incentives of firm i to invest in a project. It will do so if its stockholder value goes up as a result of the investment, that is, if $\Delta S_i > 0$. Consider two firms, i_{IN} and i_{OUT}, one in the benchmark and the other is not (i.e., $i_{\text{IN}} \le k$ and $i_{\text{OUT}} > k$). Suppose that their cash flows with and without the project are identical; specifically, $\sigma_{i_{\text{IN}}} = \sigma_{i_{\text{OUT}}} = \sigma$ and $\rho_{i_{\text{IN}}y} = \rho_{i_{\text{OUT}}y} = \rho_y$. The difference in the incremental stockholder value created by the investment for the two firms is

$$\Delta S_{i_{\text{IN}}} - \Delta S_{i_{\text{OUT}}} = \gamma\Lambda\left(\sigma_y^2 + \rho_y\sigma\sigma_y\right)\lambda_M\frac{b}{a+b}. \tag{25}$$

The right hand side of (25) is the analytical expression for the benchmark inclusion subsidy.

Assumption 1. $\sigma_y^2 + \rho_y\sigma\sigma_y > 0$.

So long as Assumption 1 holds, the benchmark inclusion subsidy is positive, and the increase in the stockholder value for the firm in the benchmark is larger than that for the firm outside the benchmark.

In practice one would expect Assumption 1 to hold for most investments. A typical project that a firm undertakes is similar to its existing activities. Even if a project is diversifying, it is still typically positively correlated with the firm's original cash flows.

The more general structure that we consider in this section allows us to fully characterize the benchmark inclusion subsidy in (25) and to derive additional implications relative to Section 3. Notice that the subsidy is the sum of two terms. The first term, $\gamma\Lambda\sigma_y^2\lambda_M b/(a+b)$, is the one that we have already seen in Section 3. It essentially captures the "index effect" for project y, since the investment effectively moves y's cash flows in the benchmark.[14] The second term, $\gamma\Lambda\rho_y\sigma\sigma_y\lambda_M b/(a+b)$, is new. (Recall that $\rho_y = 0$ in the example in Section 3 with uncorrelated cash flows, so this term was missing in the example.) It is proportional to the covariance between the existing and new cash flows, $\rho_y\sigma\sigma_y$, and so we refer to this term as the "covariance subsidy." Intuitively, when the existing and new cash flows are positively correlated, the co-

[12] As in the example, the results that follow hold even if the firm uses some debt financing. For instance, the size of the benchmark inclusion subsidy is literally identical even if the firm has risk-free debt.

[13] See also Lemma 4 in Appendix A that formally derives the two-factor CAPM in our environment, where the stock's expected returns depend on the covariance with the market portfolio as well as covariance with the benchmark.

[14] Technically speaking, if firm y were to join the index as a standalone firm and firm k gets pushed out from the index as a result, then the index effect (in absolute terms) for firm y would be $S_y' - S_y = \gamma\Lambda(\sigma_y^2 - \rho_{ky}\sigma_k\sigma_y)\lambda_M b/(a+b)$. (We discuss this in more detail in the context of IPOs in Section 4.3.) Since in our setup the effective "index inclusion" of y happens as a result of investment or merger, no other firm is removed from the index, and hence the expression for the benchmark inclusion subsidy does not include the $(-\rho_{ky}\sigma_k\sigma_y)$ term.

variance increases the overall variance of post-investment cash flows. And because the cash-flow variance is penalized less for firms that are inside the benchmark, the subsidy increases in the covariance. If ρ_y is positive, the covariance subsidy is positive and hence, the benchmark inclusion subsidy exceeds the index effect. The covariance subsidy is the largest when $\rho_y = 1$, so that y is a clone of the existing assets. Moreover, assuming that the correlation ρ_y is large enough and the variance of existing cash flows exceeds that of the new cash flows, i.e., $\sigma > \sigma_y$ (both are empirically reasonable assumptions), the covariance subsidy exceeds the index effect.

Keeping in mind that the benchmark inclusion subsidy arises from taking a difference-in-differences, we can further explain the terms that comprise it. The managers subsidize the variance of a benchmark firm's post-investment cash flow, which is $\sigma_i^2 + \sigma_y^2 + 2\rho_{iy}\sigma_i\sigma_y$. The first term, σ_i^2, washes out of the first difference given by Eq. (24) because it is present for the benchmark firm pre- and post-investment. Furthermore, the subsidy includes only one covariance term $\rho_{iy}\sigma_i\sigma_y$, not two. This is because any firm, either inside the benchmark or not, receives a subsidy for the covariance with the benchmark (one can see this from the second line of (24), which is the same for all firms). That is, projects with a positive covariance with the benchmark are more valuable, even if a non-benchmark firm undertakes them. This is because prices of benchmark stocks are inflated by the inelastic demand from portfolio managers, which leads direct investors (and managers through their mean-variance portfolios) to substitute into assets that provide exposure to the benchmark without being in the benchmark itself. Consequently, of the two covariances that enter the extra variance, one is subsidized regardless of which firm invests and the other is subsidized only when the investing firm is a benchmark firm. Hence, one of the two covariances drops out from the difference-in-differences.

The presence of the benchmark inclusion subsidy translates into different investment rules for firms inside and outside the benchmark. We formalize this result in Proposition 1.

Proposition 1 (Project Valuation). *A firm in the benchmark is more likely to invest in a project than a firm outside the benchmark if and only if Assumption 1 holds. More precisely, all else equal, a firm in the benchmark accepts projects with a lower mean μ_y, larger variance σ_y^2, and/or larger correlation ρ_y than an otherwise identical firm outside the benchmark if and only if Assumption 1 holds.*

Proposition 1 is at odds with the textbook treatment of valuation taught in basic corporate finance courses. The usual rule states that a project's value is independent of which firm undertakes it, and is simply given by the project's cash flows discounted at the project-specific (not firm-specific) cost of capital.[15] In a world with benchmarking, the proper discount rate for a project depends on whether the firm undertaking it belongs to the benchmark.

The reason why a project is worth more to a firm in the benchmark than to one outside it is because when the project is adopted by the benchmark firm, it is incrementally financed by portfolio managers regardless of its variance. So, the additional overall cash-flow variance that the project generates is penalized less for a firm inside the benchmark.

To further understand the importance of the variance, consider a special case where the project is risk free, i.e., $\sigma_y^2 = 0$. Then Assumption 1 fails and we can see that the project would be priced identically by all firms.

Remark 1 (Risk-Free Projects). If $\sigma_y^2 = 0$, then a firm's valuation of project y is independent of whether this firm is included in the benchmark or not.

We can build further intuition about the model by considering what happens when the inequality in Assumption 1 is reversed. This happens if the project is sufficiently negatively correlated with the cash flows of assets in place, that is, if $\rho_y \leq -\sigma_y/\sigma$. In this case, a project is a hedge because it reduces the variance of a firm's cash flow. Firms inside the benchmark benefit less from this reduction because their cash-flow variance is subsidized by managers and they lose some of that subsidy.[16] Consequently, a benchmark firm will value such a project less than a non-benchmark firm would.

Finally, as we mentioned earlier, our model also implies that projects correlated with assets inside and outside the benchmark are valued differently (by firms both inside and outside the benchmark). Projects that are positively correlated with the benchmark provide alternative cheaper exposure to the benchmark firms' cash flows. This is reflected in Eq. (24), which shows that for any firm, investing in a project that is positively correlated with a component of the benchmark is more beneficial than if the project has the same degree of correlation with an asset outside of the benchmark.

To illustrate this insight graphically, Fig. 1 plots the change in the stockholder value ΔS_i given by (24) as a function of ρ_{jy}, where the solid line corresponds to some asset j inside the benchmark, and the dashed line to some j outside the benchmark. In the figure, for concreteness, the investing firm i is in the benchmark. If i were outside the benchmark, these lines would shift down in parallel.

The figure shows that the change in stockholder value, ΔS_i, is decreasing in the correlation coefficient ρ_{jy}. If j is in the benchmark, then the downward sloping line is flatter. Moreover, the solid line is above the dashed line for positive correlations and below for negative correlations. This is because positive (negative) correlation of the project with an asset in the benchmark is penalized (rewarded) less than the same correlation with an asset outside the benchmark.

Remark 2 (Relevance for ESG investing). Because of the impact of benchmarks on investment decisions, our model also speaks to the ongoing debate over the efficacy of ESG

[15] See for example Jacobs and Shivdasani (2012) or Berk and DeMarzo (2017), chapter 19.

[16] Notice that when $-\sigma_y/\sigma \leq \rho_y \leq -\sigma_y/(2\sigma)$, although the investment reduces the firm's cash-flow variance, the benchmark inclusion subsidy is positive.

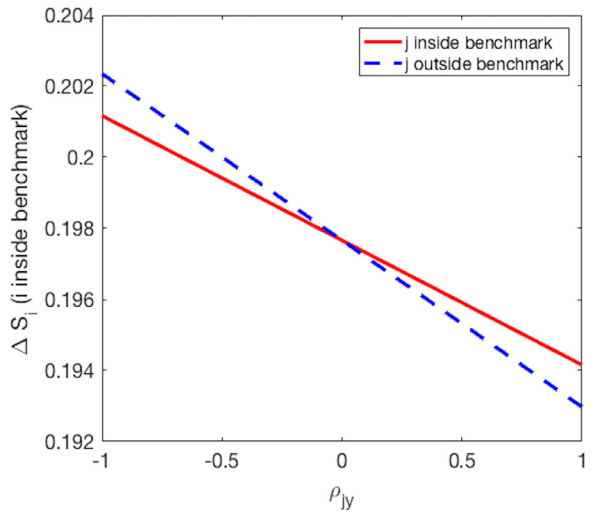

Fig. 1. Change in the stockholder value, ΔS_i, as a function of correlations of project y's cash flows with cash flows of assets inside and outside the benchmark, ρ_{jy} for some $j \leq k$ and some $j > k$. Parameter values: $n = 5$, $k = 3$, $\mu_y = 1.2$, $I = 1$, $\sigma_j = 0.15$, $\sigma_y = 0.1$, $\rho_{jy} = 0$ unless it is plotted on the horizontal axis, $\rho_{j\ell} = 0$, $\ell \neq j$, $j = 1, \ldots, n$, $\gamma = 2$, $\lambda_M = 0.3$, $a = 0.008$, $b = 0.042$. Investing firm: $i = 1$. Solid line: $j = 2$. Dashed line: $j = 4$.

(environmental, social, and governance) benchmarks. First, the model suggests a direct effect: an exclusion of companies with poor ESG characteristics in the benchmark denies such companies the benchmark inclusion subsidy, and therefore these companies would not be able to grow as much. Second, there is an indirect effect: having a company with poor ESG characteristics (e.g., highly polluting) in the benchmark encourages all firms to make investments in projects whose cash flows are positively correlated with those of the poor ESG company in the benchmark (i.e., highly polluting). In contrast, replacing such company in the benchmark with a firm with a higher ESG score (e.g., with low emissions) would encourage firms outside the benchmark to mimic such a firm instead.

Discussions about benchmarking often revolve around the possibility that it leads to more correlation in risk exposures for the people hiring portfolio managers. Our model points to an additional source of potential correlation generated by benchmarking. Because benchmarking leads to higher valuations of stocks that are correlated with the benchmark, it induces firms—both inside and outside the benchmark—to take on more fundamental risk that is correlated with the benchmark (relative to the economy without benchmarking). Thus, our model predicts that cash flows in the economy with managers endogenously become more homogeneous/correlated with each other.

In practice, the criterion for equity benchmark inclusion is often based on firms' market capitalization. In our model, firms that are already in the benchmark benefit from the benchmark inclusion subsidy and invest more than their twins outside the benchmark. Firms in the benchmark, therefore, grow faster than their peers outside. This would make the benchmark composition persistent.

4.1.1. Cost of capital

Notice that the benchmark inclusion subsidy is directly related to the project-specific cost of capital for an investment. Define firm i's cost of capital, applied to the investment project y, r_i, by $\Delta S_i = -I + \mu_y/(1 + r_i)$, where ΔS_i, defined in (24), measures the increase in stockholder value from adopting the project. Then the difference in the cost of capital for firms outside and inside the benchmark is

$$r_{out} - r_{in} = \frac{\mu_y}{\Delta S_{out} + I} - \frac{\mu_y}{\Delta S_{in} + I} = \frac{\mu_y(\Delta S_{in} - \Delta S_{out})}{(\Delta S_{in} + I)(\Delta S_{out} + I)}. \tag{26}$$

Notice that unlike (25), which is measured in dollars, (26) is unit-free (or can be expressed in percent). For that reason, we will use (26) for the quantitative analysis in Section 5.

4.2. Mergers and acquisitions

As we have seen in the example considered in Section 3, the model can also be used to think about mergers and acquisitions.

Proposition 2 (Mergers and Acquisitions). *Suppose firm i considers acquiring firm y that is outside the benchmark, and $\sigma_y^2 + \rho_{iy}\sigma_i\sigma_y > 0$. Then firm i is more likely to acquire y if firm i is inside the benchmark than if it is outside. More precisely, all else equal, a firm in the benchmark will acquire an asset with a lower mean μ_y, larger variance σ_y^2, and/or larger correlation ρ_y than a firm outside the benchmark if and only if $\sigma_y^2 + \rho_{iy}\sigma_i\sigma_y > 0$ holds.*

The logic behind this statement is identical to the reasoning that leads to the bias in investment. If a benchmark firm acquires y, it gets the benchmark inclusion subsidy. Again, this result is in contrast to the direct wisdom about the role of financing synergies in the evaluation of potential acquisitions. For example, if a firm has unused debt capacity, it might choose to use more debt financing than otherwise to buy another firm. The usual view is that the discount rate used to value the cash flows of the target firm should not be altered by the availability of the extra debt funding. The case for not adjusting the discount rate is that the same additional debt funding could have been used for any other potential acquisition. So, it would be a mistake to say that any particular target company is a more attractive firm to acquire just because some low-risk debt could be issued to finance the purchase.

In our setup, there is a more fundamental synergy that is responsible for lower financing costs. Because the managers will want to purchase part of any stock that is issued to undertake the transaction, those savings should be accounted for. The size of the subsidy will depend on the parameters that appear in Assumption 1. Thus, for example, all else equal, the higher the correlation of the cash flows of the target firm with the acquiring benchmark firm, the larger the financing advantage associated with that acquisition. Conversely, a hedging acquisition by a firm in the benchmark, where the target firm's cash flows are negatively correlated with acquirer's, always comes with a lower subsidy.

Proposition 2 works in reverse for spinoffs and divestitures. Specifically, assuming that the condition $\sigma_y^2 + \rho_{iy}\sigma_i\sigma_y > 0$ is satisfied, a division y is worth more if it is part of a firm inside the benchmark than if it is spun off and trades as a separate entity outside the benchmark or is sold to a firm outside the benchmark.

Finally, if firms i and j are both in the benchmark and $\sigma_i > \sigma_j$, then the subsidy will be bigger for i than for j. The difference arises because the subsidy for the variance of the cash flows added in the merger depends on the variance of existing cash flows.

4.3. IPOs and incentives to join the benchmark

Suppose y is now a standalone firm, which is held privately by direct investors and is considering an IPO. We demonstrate that y's incentive to go public depends on whether it will be included in the benchmark.

We consider two scenarios. In the first scenario, when firm y becomes public and gets included in the benchmark, no other firm leaves the benchmark. Most of the best known stock indices in the world have a fixed number of firms. Thus in the second scenario, if y joins the benchmark, then firm k is removed, so that the number of firms in the benchmark remains constant.

Proposition 3 (IPOs and Benchmarks). *Consider a privately-held firm y contemplating an IPO.*

(i) *It is more profitable for firm y to proceed with an IPO if it gets included in the benchmark and no other firm leaves the benchmark, than if it is does not get included in the benchmark.*

(ii) *It is more profitable for firm y to proceed with an IPO if it gets included in the benchmark and firm k is removed from the benchmark, than if it does not get included in the benchmark, if and only if* $\sigma_y^2 - \rho_{ky}\sigma_k\sigma_y > 0$.

The argument for the result in part (*i*) is the same as for other results in the paper—firm y gets the benchmark inclusion subsidy if it joins the benchmark. In part (*ii*) where y pushes another firm out of the benchmark, there is an additional consideration, as firm y loses part of the benchmark subsidy coming from its correlation with that firm. In other words, when firm y is included in the benchmark and firm k is pushed out, firm y's correlation with the benchmark increases by σ_y^2 because firm y is correlated with itself (and it enters the benchmark), and is reduced by $\rho_{ky}\sigma_k\sigma_y$ because firm k drops out of the benchmark. The net subsidy is therefore proportional to $\sigma_y^2 - \rho_{ky}\sigma_k\sigma_y$.

One could apply the above argument to any firm, not just a newly listed one. A firm is worth more inside the benchmark rather than outside. So, there is an added benefit to any corporate action that results in the firm's benchmark inclusion—for example, aimed at increasing the firm's size or meeting other criteria for benchmark inclusion. The costs of taking such action of course have to be outweighed by the benchmark inclusion subsidy, but a clear empirical prediction emerging from this discussion is that firms with good prospects of benchmark inclusion have an incentive to alter their behavior in order to gain membership in the benchmark. Similarly, benchmark firms that are close to the threshold for exclusion, have incentives to engage in potentially costly corporate actions that ensure that they retain their benchmark membership.

Remark 3 (Subsidy and Stock Return Volatility). While past research (e.g. Basak and Pavlova, 2013) have found that benchmarking alters volatilities of stock returns, that channel is absent in our main model. One question that might arise is whether accounting for this effect would change the analysis. In Appendix D, we extend our model to three periods and adopt a discrete-time square-root specification for the dividend process from Buffa et al. (2014), which delivers nontrivial second-moment implications in the presence of benchmarking.[17] We show that volatilities and correlations of per-share returns of stocks inside the benchmark are higher than those outside the benchmark. We also show that if benchmark inclusion increases stock return volatility, this force in isolation dampens the benchmark inclusion subsidy. Intuitively, since investors are risk averse, the higher the stock return volatility, the lower the stock price. Therefore, in the presence of portfolio managers, stock prices are inflated but not as much as in our baseline model. Following the same logic, when two firms adopt the same project, the firm inside the benchmark also experiences a larger increase in stock return volatility. As a result, this channel (in isolation) contributes to a smaller subsidy.[18] However, volatility also affects the subsidy in other ways, for instance, because it also changes the level of stock prices, so the total effect of volatility on the subsidy is difficult to sign.

4.4. Comparative statics with respect to λ_M

In this subsection we analyze the benchmark inclusion subsidy as a function of the size of the asset management sector. Consider (25), and rewrite it recognizing that $\Lambda = [\lambda_M/(a+b) + \lambda_D]^{-1}$ and $\lambda_D = 1 - 2\lambda_M$:

$$\Delta S_{i_{\text{IN}}} - \Delta S_{i_{\text{OUT}}} = \gamma \left[1 + \frac{1 - 2\lambda_M}{\lambda_M}(a+b)\right]^{-1} b\left(\sigma_y^2 + \rho_y \sigma \sigma_y\right). \tag{27}$$

Notice that this expression is strictly increasing in λ_M. This means that the effects described in this paper related to the difference in valuations by a firm inside the benchmark relative to a firm outside the benchmark become larger as the size of the asset management sector increases. Similarly, the difference in the cost of capital for new investments of a firm outside the benchmark relative to one inside it becomes larger as λ_M rises.

[17] While the model is tractable and delivers closed-form expressions for the quantities of interest, the complexity of the expressions rises significantly relative to our main model.

[18] In practice, however, stocks inside the benchmark (e.g., the S&P 500) have higher liquidity, which lowers their return volatility. This effect works in the opposite direction to the one we identified above.

4.5. Passive asset management

As we mentioned earlier, a limiting case of our setup when $b \to \infty$ corresponds to having passive asset management. In this case, it is easy to see that managers hold only the benchmark portfolio, i.e, $x^M = \mathbf{1_b}$. A generalization of our model would be to include both active and passive managers. If we denote the fractions of them in the economy by λ_M^A and λ_M^P, then the equilibrium stock prices would be

$$S = \mu - \gamma \Lambda \Sigma \left(\mathbb{1} - \left[\lambda_M^A \frac{b}{a+b} + \lambda_M^P \right] \mathbf{1_b} \right), \tag{28}$$

where $\Lambda = \left[\lambda_M^A/(a+b) + \lambda_D \right]^{-1}$.

All of our results extend to this case. Passive managers hold benchmark stocks irrespective of their characteristics, and they invest nothing in the mean-variance portfolio. Therefore, with passive managers the benchmark inclusion subsidy becomes even larger.[19] In particular, instead of the expression in (25), the subsidy becomes $\gamma \Lambda (\sigma_y^2 + \rho_y \sigma \sigma_y) \left[\lambda_M^A b/(a+b) + \lambda_M^P \right]$. The following proposition shows that this value is increasing if more managers are passive.

Proposition 4 (Passive Asset Managment). *The benchmark inclusion subsidy is larger when λ_M^P/λ_M^A is larger, holding $\lambda_M^P + \lambda_M^A$ constant.*

An analogous comparative-statics statement applies to the difference in the cost of capital for new investments of a firm outside vs. inside the benchmark—this difference is increasing in λ_M^P/λ_M^A holding $\lambda_M^P + \lambda_M^A$ constant.

One could argue that an increase in passive management comes not only from a (relative) decrease in active managers, but also from a reduction in agents who previously did not participate in the market, but now invest through passive funds. It is straightforward to incorporate such a dimension in our above extension. Specifically, denote the measure of the non-participating group (or "outsiders") by λ_O, and let $\lambda_O + \lambda_M^P + \lambda_M^A + \lambda_F + \lambda_D = 1$. As everywhere above, we maintain the assumption that the measure of managers of each type is equal to the measure of investors of the same type. Replacing λ_M^P with $(1 - \lambda_D - \lambda_O)/2 - \lambda_M^A$ in the above expressions, it is straightforward to show that more participation in the risky-asset market through passive investment increases the subsidy.[20]

4.6. Welfare considerations

A natural question that arises from our analysis is what are the welfare implications of the benchmark inclusion subsidy? A short answer is that welfare effects of the benchmark inclusion subsidy are ambiguous. The details of our analysis are in Appendix C, here we provide a brief summary.

[19] In a recent paper, Buss and Sundaresan (2020) propose an interesting complementary channel, based on information acquisition, through which passive ownership affects firms' investment.

[20] As the measure of non-participants λ_O decreases, half of them are moved to passive fund managers and half are moved to passive fund investors.

First, note that due to the subsidy, agents who are initially endowed with a large (small) amount of the benchmark firm's stock benefit (lose) from the subsidized increase in the price of this stock following an investment or merger. However, this effect is purely distributive, representing movement along a Pareto frontier, so it washes away in the aggregate.

Second, the question of whether the aggregate social welfare increases or decreases as a result of the subsidy is much more subtle. Appendix C illustrates the issues using a simple version of our model with uncorrelated assets and mergers considered in Section 3. We show that when a non-benchmark firm 2 acquires firm y, the agents' utilities (and hence the aggregate welfare) remain unchanged. However, when the benchmark firm 1 makes the acquisition, part of cash-flow risk is reallocated from the direct investors and managers to fund investors. The impact of changing the risk-sharing in this way depends on the parameters of the compensation contract and, depending on those parameters, this could raise or lower overall welfare.

The above results are derived for mergers, where the total cash flows in the economy are kept unchanged. With investment, total cash flows change as a result of an investment, and a benchmark firm will invest in projects that a non-benchmark firm will not. This will introduce an additional consideration into the analysis that further clouds the ability to determine the aggregate effect on welfare.

5. Quantitative analysis

One obvious question is whether the benchmark inclusion subsidy is quantitatively relevant for corporate actions? While the CARA-normal framework is widely used because of the tractable formulas that it delivers, it is well-understood that using it for quantitative analysis is challenging. Recognizing these issues, here is an indicative calculation of the potential effect of the subsidy on the cost of capital.

We intend to measure the difference in the cost of capital for a project undertaken by a firm in the benchmark and a firm outside the benchmark, as computed in Eq. (26). In practice, criteria for benchmark inclusion are typically based on market capitalization, and we will construct our benchmark that way too. From Ken French's website, we obtain 10 portfolios formed on size, which consist of U.S. stocks sorted into the NYSE size deciles. We set $n = 10$, with each asset representing a corresponding size portfolio. In what follows, the properties of each of the 10 assets are computed using equal weights for the constituent stocks that the asset represents. We assume the benchmark is the S&P 500, which includes only a subset of U.S. stocks. The three largest NYSE deciles contain about 500 stocks, and we therefore set $k = 3$, i.e., the first three assets are in the benchmark and the remaining ones are outside. The ten assets differ not only in terms of their market capitalization, but also in terms of their market betas. We therefore specify the cash flow process so that it will generate both types of heterogeneity. In particular, we assume that

$$D_i = \alpha_i^0 \left[\bar{\mu} + \alpha_i^1 Z + \left(1 - \alpha_i^1\right) \varepsilon_i \right], \qquad i = 1, \ldots, 10, \tag{29}$$

Table 1
Parameter values.

Parameter	Value	Explanation
α^0	(1, 0.1936, 0.1024, 0.0605, 0.0382, 0.0251, 0.0165, 0.0104, 0.0053, 0.0011)	For the largest asset, α_1^0 is normalized to be 1; the other 9 are chosen to match data for estimates of size ratios of deciles 2–10 relative to decile 1: (0.1914, 0.1008, 0.0594, 0.0373, 0.0245, 0.0161, 0.0101, 0.0051, 0.0011)
α^1	(0.1, 0.1174, 0.1224, 0.1269, 0.1317, 0.1342, 0.1386, 0.1442, 0.1510, 0.1504)	α_1^1 is set to be 0.1; the other 9 are chosen to match estimates of market betas for deciles 2–10: (1.0914, 1.1407, 1.1852, 1.2339, 1.2600, 1.3047, 1.3634, 1.4344, 1.4280)
$\gamma\Lambda$	0.2789	Chosen to match the market risk premium of 7.7%
σ_z^2	316.33	Chosen to match the standard deviation of return on the market of 19.7%
σ_ε^2	0.0957	Chosen to match the average idiosyncratic risk of the average asset (across the size portfolios), 3.5%
$\lambda_M/(a+b)$	0.302	Chosen to match the equally-weighted institutional ownership of all stocks in the market, 67%
$b\zeta$	1.9	Chosen to match the equally-weighted institutional ownership of benchmark stocks, 83%

where $Z \sim N(0, \sigma_z^2)$ is a common factor and $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$ is an idiosyncratic shock. The parameter α_i^0 captures asset i's size and α_i^1 governs its exposure to systematic vs. idiosyncratic risk. We define $\alpha^0 = (\alpha_1^0, \ldots, \alpha_{10}^0)$ and $\alpha^1 = (\alpha_1^1, \ldots, \alpha_{10}^1)$.

Table 1 reports the parameter values that we use along with the intuition for how they were chosen. Of course, all the parameters interact to determine the model outcomes, so the rationale listed in the table for how each is chosen is subject to that caveat.

The size of the subsidy depends only on the relative sizes of the various assets. So we normalize the size of the market capitalization of the largest asset to be 1, i.e. $\alpha_1^0 = 1$. We then calibrate the remaining α_i^0, $i = 2, \ldots, 10$, so that the relative market capitalizations of the other nine assets (compared to the first asset) match the December 2018 data on their relative market capitalizations.

To infer the relative importance of aggregate and idiosyncratic risk we proceed as follows. We set $\alpha_1^1 = 0.1$.[21] We then calibrate the remaining α_i^1, $i = 2, \ldots, 10$, so that the market betas of assets 2 - 10 in the model match the estimated market betas of the nine portfolios (based on monthly data from 1926 to 2018). The regressions that are used to estimate the betas also give us estimates of each asset's idiosyncratic risk, which on average is 3.5%. We use this average to choose the variance of the idiosyncratic risk in (29).

We infer the market's effective risk aversion, $\gamma\Lambda$, and the standard deviation of common factor, σ_z, in (29) by choosing values that will match the market risk premium and the volatility of returns on the market portfolio of 7.7% and 19.7%, respectively, that are reported by Brealey et al. (2019), Chapter 7 (for the period 1900–2017). In the CARA-normal setup, the level of the cash-flow mean $\bar{\mu}$ in (29) does not make a difference for the estimated cost of capital.

A critical determinant of the subsidy is the level of inelastic demand from the portfolio managers. To quantify this we turn to data on institutional ownership from the FactSet/LionShares database, available from WRDS for 2000–2017. Consistent with our model, in the data the institutional ownership of stocks in the benchmark (the S&P 500) exceeds that of stocks outside the benchmark. The (equally-weighted) ownership of the stocks in the benchmark (deciles 1–3) as of December 2017 is 83% while that of the whole market is 67%. In the model, these two quantities correspond to $\frac{1}{k}\sum_{i=1}^{k} x_i^M$ and $\frac{1}{n}\sum_{i=1}^{n} x_i^M$, respectively (recall that we normalized the total supply of each stock to one).

A well-known issue with the CARA utility is that it implies that there are no wealth effects. The standard solution to this problem in this kind of exercise is to adjust the risk aversion of agents so that an agent with higher wealth (e.g., a large fund manager) has a lower risk aversion and therefore buys more risky assets. This approach appropriately adjusts the first term in the manager's demand for risky assets in (20), the mean-variance portfolio, but not the second term, the inelastic demand due to benchmarking. The latter does not scale with wealth (or assets under management) simply because, as is standard in CARA models, we define relative performance as a per-share measure. In Appendix B, we define relative performance as a per-dollar return and show that it implies that the inelastic demand does scale with assets under management, as it would in practice. We introduce this type of wealth effect by multiplying the inelastic demand by a parameter ζ.[22] Then the manager's demand and the asset prices become

$$x^M = \frac{1}{a+b}\Sigma^{-1}\frac{\mu - S}{\gamma} + \frac{b\zeta}{a+b}\mathbf{1_b}, \tag{30}$$

$$S = \mu - \gamma\Lambda\Sigma\left(\mathbb{1} - \lambda_M\frac{b\zeta}{a+b}\mathbf{1_b}\right). \tag{31}$$

In light of this, we set $b\zeta$ and $\lambda_M/(a+b)$ so that we match the data on institutional ownership.

We imagine a hypothetical firm, which does not belong to the benchmark, but has the same characteristics (size and beta) as firms represented by asset 3, i.e., the bottom of the S&P 500. We think of this experiment as contrasting similar firms, where one is in the S&P 500 and its twin is not. We assume that the investment project under consideration is a clone of the investing firm, $\rho_y = 1$, as would

[21] For the calculations, choosing any number between 0 and 1 would be make no difference. This normalization simply pins down the level of σ_ε^2.

[22] The parameter ζ corresponds to $W_0^M/\sum_i \mathbf{1}_i\tilde{x}_i S_i$ in Appendix B, see e.g., equation (B.1), which is the value of assets under management divided by the benchmark market capitalization.

Table 2
The benchmark inclusion subsidy and institutional ownership. This table presents a sensitivity analysis of the difference in the cost of capital for new investments (in basis points) for identical firms outside vs. inside the benchmark, as institutional ownership varies. Each of the identical firms has the characteristics of asset 3. The top 3 size deciles of stocks are in the benchmark.

		Institutional ownership of the market		
		59%	67%	75%
Institutional ownership of the benchmark	75%	67	35	0
	83%	133	**94**	51
	91%	260	215	159

be the case if a firm opened a new plant, and that the size of the project μ_y/σ_y is 1% of the firm's expected cash flows. Under these assumptions, the difference in the cost of capital for new investments for the hypothetical outside firm and the one inside the benchmark as measured by (26) is 94 basis points.[23] Our model implies that if the asset management sector is larger (the institutional ownership of stocks in the benchmark is higher), the subsidy will be larger. In the data, we see a steady increase in the institutional ownership of U.S. stocks over the past three decades. Table 2 explores how our differences in the cost of capital estimates vary with the institutional ownership of the benchmark and the market. Our baseline estimate is at the center of the table. The difference in the cost of capital is the highest when the difference between the institutional ownership of the benchmark and the market is the highest (because this gap reflects the managers' inelastic demand). The other limiting case is when the institutional ownership of the benchmark and the market are the same, in which case the subsidy is zero. This can happen in our model if the benchmark and the market portfolios coincide or if the manager's contract does not include a relative performance component ($b = 0$).

In unreported analysis, we perform similar experiments assuming that the hypothetical firm is a twin of asset 1 or asset 2 instead of asset 3. In those experiments, the difference in the cost of capital is higher than our baseline estimate, and it is larger for bigger firms. This is because bigger firms have larger cash flow volatility, and from the previous analysis we know that larger volatility firms enjoy a higher benchmark inclusion subsidy. Conversely, if the project under consideration is not a clone of the investing firm's assets, then the subsidy will be smaller. This follows from the earlier discussion in Section 4.1 about the magnitude of the covariance subsidy.

As a robustness check, we explore how the subsidy changes if the benchmark is broader than the S&P 500 and closer to the market portfolio. To do so, we consider the Russell 1000 as the benchmark instead of the S&P 500. It turns out that the (equally-weighted) institutional ownership of the Russell 1000 stocks is the same as that for the S&P 500, which is 83%.[24] We recalibrate the model and target the same moments as in the original calibration, except the top 5 size deciles of stocks are now in the benchmark instead of top 3 (k=5). We find that in the Russell 1000 calibration, the difference in the cost of capital for the firm just outside the benchmark relative to that just inside is lower, amounting to 60 basis points in the baseline case (compared to 94 basis points in the S&P 500 calibration). The reason is that the smallest firm included in the Russell index is smaller than that for the S&P 500 index—has a lower cash-flow variance—and the lower the variance, the lower the subsidy.

In our calibration we have assumed there is only one benchmark. The key driver of the results, however, is the empirical difference in the level of institutional ownership between the S&P 500 firms and the non-S&P 500 firms. That gap reflects the inelastic demand from portfolio managers created by membership in the S&P 500. Our calibration strategy is valid even if there are other benchmarks that firms belong to. In particular if there were a second benchmark, the total market index, to which all firms belonged, then our calibration would be literally identical. All that would change is the interpretation of the results. In this case, we would have the incremental subsidy that the smallest S&P 500 firm receives relative to a twin firm that is in the total market index but not in the S&P 500.

We recognize that there is a fair bit of uncertainty about many of the objects that we would need to know in order to accurately estimate the subsidy. We consider this quantification to merely be indicative and to show that it is reasonable to think the subsidy could be quantitatively relevant for corporate decisions.

6. Related empirical evidence

We now turn to the empirical evidence that is related to the predictions of our model. In keeping with the presentation in Section 4, we organize the discussion around three predictions of the model. The first implication of our model is that upon inclusion in a benchmark, there should be an increase in a firm's stock price. This prediction is not original, but our theory does imply that there is an important distinction between a benchmark and an index. The second one is that firms inside the benchmark should be more prone to invest and to engage in mergers. Third, the subsidy should be higher when there are more assets under management.

6.1. Benchmark effect

Consistent with the empirical evidence, our model generates an index effect. Stock price changes are symmetric for index additions and deletions and the effect persists for as long as the stock is in the index. We also have a more subtle prediction: the share price response should depend on becoming part of a benchmark and not just because of being added to an index. In most cases, separat-

[23] This number is in line with Calomiris et al. (2019), who estimate the yield discount for emerging markets investment grade corporate issues eligible for index inclusion to be between 76 to 99 basis points in recent data.

[24] The estimate is as of December 2017. We again use data on institutional ownership from the FactSet/LionShares database, available from WRDS for 2000–2017.

ing the effect of being in the index and benchmark is challenging. One exception arises for firms that operate in so-called "sin" industries such as alcohol, tobacco, and gaming. Large firms in these industries would be included in indices such as the S&P 500, Russell 1000 or FTSE 100, but are deemed odious by some investors and hence excuded from their benchmarks. Hong and Kacperczyk (2009) study the returns of these sin firms and find that sin firms earn higher expected returns than comparable firms by about 26 basis points per month—i.e. roughly 3% per year. They also estimate the differences in the levels of stock prices for sin stocks and a matched sample of non-sin stocks and find that sin firms are valued between 15% to 20% lower.

Hong and Kacperczyk's results about sin stocks have also subsequently been confirmed in several studies.[25] As noted in Remark 2, these studies also have implications for the construction of ESG benchmarks.

6.2. Changes in corporate actions following benchmark inclusion

There are some papers that attempt to assess whether the model predictions regarding investment and mergers hold for benchmark firms versus non-benchmark firms. This is challenging because ideally one wants to control for both the selection into the benchmark and all the other factors that influence these kinds of expenditures.[26]

There are five papers that we are aware of that attempt to measure these effects and all find some evidence in favor of our model's predictions. Massa et al. (2005) compare 222 firms that were added to the S&P 500 with a control set of firms that were not and find that inclusion is associated with higher levels of equity issuance and more investment, with a substantial portion of the investment coming via increased mergers. Vijh and Yang (2008) also document that firms in the S&P 500 undertake significantly more acquisitions, in line with our model's predictions. These findings hold after they account (as much as they can) for observable differences in target and acquirer characteristics, though it is hard to know whether the controls are truly adequate.

Third, there are a pair of papers that investigate how inclusion in the Morgan Stanley Capital International (MSCI) All Country World Index, a benchmark for many international equity funds, alters the behavior of firms that join the index. Luong et al. (2017) show that higher foreign ownership shares that are related to index inclusion are associated with higher levels of research and development (as measured by patents). It is common to use research and development indicators as a proxy for corporate risk-taking (see, e.g. Coles et al., 2006 and Gormley et al., 2013). This evidence is consistent with the model prediction that a firm inside the benchmark would accept riskier projects than an otherwise identical firm outside the benchmark would. Kacperczyk et al. (2019) also show that firm investment rates rise upon being included in the MSCI benchmark. They find that measures of corporate governance do not seem to change when ownership changes, casting doubt on that potential explanation for the effect. This result is also consistent with the predictions of our theory, but a skeptic might argue that index inclusions are not entirely exogenous events.

Finally, further evidence comes from Bena et al. (2017). They study differences in investment and employment for firms across 30 countries between 2001 and 2010. Their basic regression relates capital expenditures relative to assets (or the number of employees) to institutional ownership by foreign investors and a host of firm-level controls (including sales, Tobin's Q, and cash holdings). Like the previous two papers, they instrument for the ownership variable using additions to the MSCI ACWI index. They find a large, statistically significant effect of the benchmark additions on investment, employment and R&D. Importantly, the results are also present when they restrict the analysis to firms that are close to the cutoff for inclusion in the index and when they estimate the effects of inclusion using a difference-in-difference experimental design.

6.3. Variation in the subsidy size

A number of well-known empirical results can be reinterpreted as evidence that the benchmark inclusion subsidy increases with the size of the asset management sector, λ_M. There are two types of studies that are relevant. One set studies index effects, which we reviewed in Section 2. By and large, this literature finds that the index effect (and hence the subsidy) has been rising over time (e.g., Chen et al., 2004). Since the asset management sector has also been growing, this evidence is consistent with our predictions. Much of this evidence, however, is for the S&P 500 index inclusions, and there is a confounding problem that there may be changes to cash flows or stock liquidity associated with the S&P 500 membership (see Denis et al., 2003).

The second set of papers exploit the cutoff between the Russell 1000 and Russell 2000 indexes to identify price pressure from index funds and ETFs following these indices.[27] The Russell 1000/2000 cutoff separates firms with very similar market values and stock liquidity. Chang et al. (2015) study firms that transition between the Russell 1000 and 2000. As they note, when a firm's fortunes improve, the company will move from the Russell

[25] See Fabozzi et al. (2008), Statman and Glushkov (2009), and Kim and Venkatachalam (2011) for evidence of superior performance of sin stocks, as well as Blitz and Fabozzi (2017) who caution that the performance of sin stocks can be explained by two new quality factors.

[26] For example, consider the evidence in Gutiérrez and Philippon (2017). They document that higher institutional ownership accompanies lower investment. They stress, however, that it is difficult to establish causality without any plausibly exogenous movement in ownership. Their preferred interpretation is that investment is crowded out by higher payouts, though they admit this could be due to a preference by institutional firms to find firms with high payouts. In Gutiérrez and Philippon (2018) they attempt to isolate variation in payout variation that can be ascribed to ownership structure. They find that controlling for cash flows (and other firm-specific variables), the higher ownership induced payouts are associated with lower investment. Of course, the sources and uses of funds accounting identity may also lead to this kind of pattern in the data.

[27] Our model implies that the price pressure comes not just from passive funds and ETFs, but, importantly, also from active funds, which control a larger fraction of assets under management benchmarked to these indices.

2000 to the Russell 1000, but when this happens, the institutional demand for the stock will drop (because more managers are benchmarked against the Russell 2000). The authors find that despite the improved fundamentals, their share price drops by about 5% from the rebalancing. Conversely, firms that fall into the Russell 2000 see price increases of about 5%.

Finally and most importantly, Pavlova and Sikorskaya (2020) test a version of our theory that accounts for the fact that there are multiple benchmarks.[28] One of their objectives is to show that stocks with higher inelastic demand from fund managers have lower expected returns. Pavlova and Sikorskaya's measure of funds' inelastic demand for a stock is the cumulative weight of the stock in all benchmarks, weighted by assets under management following each benchmark. They refer to it as the benchmarking intensity (BMI), and construct it using a dataset of 33 U.S. equity indices. Exploiting the discontinuity in the BMI of stocks at the Russell 1000/2000 cutoff, they show that an increase in a stock's BMI leads to lower long-run returns. This is direct evidence in favor of our theory.

7. Concluding remarks

We have seen that for firms that are part of a benchmark, the inelastic demand for their shares by portfolio managers lowers their cost of capital for investments, mergers, and IPO decisions. We have specific cross-sectional predictions for the size of this effect. While there is empirical evidence that speaks to some of these predictions, there are others that have yet to be tested. One obvious direction for future work would be to fill in these gaps.

For instance, there are many claims by practitioners (e.g. McKinsey on Finance, 2004) that a strong motive for undertaking an IPO is to become part of a benchmark. We believe no one has tested this hypothesis. Despite the practitioner attention, this implication is not part of the very long list of commonly cited reasons by economists that are usually considered.[29] So there would be some novelty value to confirming the model prediction.

More importantly, international differences create variation in IPO incentives that would make it possible to cleanly uncover the predicted effect. Specifically, not only do different exchanges have different requirements about how many shares have to be floated, but the relevance of benchmarks also varies across markets. So, the ease of qualifying for a public listing across markets will differ from the size of the subsidy implied by our theory. This should make it feasible to test the theory.

It would also be interesting to test the model's predictions about how the presence of benchmarks can alter the incentives regarding mergers. We saw that the benchmark inclusion subsidy is larger (smaller) for targets whose cash flows are more positively (negatively) correlated with the acquiring firm. While there is a large literature studying merger patterns, we believe this somewhat unusual prediction of our theory has not been investigated. Furthermore, our model may suggest an alternative explanation for the rise in the industry concentration ("monopolization") in the U.S. over the past 15 years.[30] A significant driving force behind this phenomenon is mergers. According to our model, firms in prominent benchmarks (e.g., the S&P 500) value targets outside the benchmark above their standalone values, and the valuation gap increases with the size of the asset management sector. Perhaps the rapid growth of the asset management industry over the last 15 years has contributed to the increased merger activity.

Another model prediction is that firms, both inside and outside the benchmark, are rewarded for taking on more fundamental risk that is correlated with the benchmark (relative to the economy without benchmarking). Thus, our model predicts that cash flows in the economy with portfolio managers endogenously become more homogeneous/correlated with each other. This force could eventually subtly change business-cycle dynamics. However, this effect will take time to play out, so finding an empirical strategy to identify it will be challenging.

Appendix A

In the main text, for simplicity of exposition we normalize the total supply of shares of each asset to one. Here, to show that this is innocuous, we suppose that stock i has a total supply of $\bar{x}_i$ shares. The per-share cash flow of asset i is then $D_i/\bar{x}_i$.

Proof of Lemma 1. Denote by $\hat{x}_i^\ell$ the fraction of shares of asset i that agent $\ell \in \{D, M\}$ holds, i.e., $\hat{x}_i^\ell = x_i^\ell/\bar{x}_i$. Let $\hat{x}^\ell = \left(\hat{x}_1^\ell, \ldots, \hat{x}_n^\ell\right)^\top$, $\ell \in \{D, M\}$, and $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_n)^\top$. Then the maximization problem of a direct investor with $\hat{x}^D = z$ is the same as that of a portfolio manager with $(a+b)\hat{x}^M - b\mathbf{1_b} = z$ and can be written as $\max_z -E\exp\{-\gamma(z(D - \bar{x}\cdot S))\}$, where $\bar{x}\cdot S = (\bar{x}_1 S_1, \ldots, \bar{x}_n S_n)^\top$. It is well known that when asset returns are normally distributed, the optimization of an agent with CARA preferences is equivalent to the following mean-variance problem:

$$\max_z z^\top(\mu - \bar{x}\cdot S) - \frac{\gamma}{2} z^\top \Sigma z. \tag{32}$$

The optimal solution is $z = \Sigma^{-1}(\mu - \bar{x}\cdot S)/\gamma$. Thus we have

$$\hat{x}^D = \Sigma^{-1}\frac{\mu - \bar{x}\cdot S}{\gamma}, \tag{33}$$

$$\hat{x}^M = \frac{1}{a+b}\Sigma^{-1}\frac{\mu - \bar{x}\cdot S}{\gamma} + \frac{b}{a+b}\mathbf{1_b}. \tag{34}$$

When $\bar{x} = \mathbb{1} \equiv (1, \ldots, 1)^\top$, $\hat{x}^\ell = x^\ell$ for $\ell \in \{D, M\}$ and we have Eqs. (19) and (20). □

[28] See also Cremers et al. (2012) for a multi-benchmark model that explains the cross-section of mutual fund returns.

[29] For example, Celikyurt et al. (2010) observe that "in theory, an IPO creates liquidity for the firm's shares, provides an infusion of capital to fund growth, allows insiders to cash out, provides cheaper and ongoing access to capital, facilitates the sale of the company, gives founders the ability to diversify their risk, allows venture capitalists and other early stage investors to exit their investment, and increases the transparency of the firm by subjecting it to capital market discipline."

[30] See David Leonhardt, "The Monopolization of America," *The New York Times*, November 25, 2018, https://www.nytimes.com/2018/11/25/opinion/monopolies-in-the-us.html.

Proof of Lemma 2. Using the market-clearing condition $\lambda_M \hat{x}^M + \lambda_D \hat{x}^D = \mathbb{1}$, we have the vector of the total share value of the firms

$$\bar{x} \cdot S = \mu - \gamma \Lambda \Sigma \left(\mathbb{1} - \lambda_M \frac{b}{a+b} \mathbf{1}_{\mathbf{b}} \right), \tag{35}$$

This gives us (21) when $\bar{x} = \mathbb{1}$. □

Proof of Lemma 3. Suppose firm i adopts the project. Then the total number of shares of asset i becomes $\bar{x}_i^{(i)} = \bar{x}_i + \delta_i$, where $\delta_i S_i^{(i)} = I$, and $\bar{x}_j^{(i)} = \bar{x}_j$ for $j \neq i$.

From (35) we know that

$$\bar{x}^{(i)} \cdot S^{(i)} = \mu^{(i)} - \gamma \Lambda \Sigma^{(i)} \left(\mathbb{1} - \lambda_M \frac{b}{a+b} \mathbf{1}_{\mathbf{b}} \right). \tag{36}$$

Finally, using the definition of $\bar{x}^{(i)}$, we have

$$\bar{x} \cdot S^{(i)} = \mu^{(i)} - I^{(i)} - \gamma \Lambda \Sigma^{(i)} \left(\mathbb{1} - \lambda_M \frac{b}{a+b} \mathbf{1}_{\mathbf{b}} \right), \tag{37}$$

which simplifies to (23) when $\bar{x} = \mathbb{1}$.

The *ith* element of $\bar{x} \cdot S$ is

$$\bar{x}_i S_i = \mu_i - \gamma \Lambda \sum_{j=1}^{n} \rho_{ij} \sigma_i \sigma_j \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_j \right) \tag{38}$$

and the *ith* element of $\bar{x} \cdot S^{(i)}$ is

$$\begin{aligned} \bar{x}_i S_i^{(i)} = {} & \mu_i + \mu_y - I - \gamma \Lambda \left(\sigma_y^2 + \rho_{iy} \sigma_i \sigma_y \right) \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_i \right) \\ & - \gamma \Lambda \sum_{j=1}^{n} \left[\rho_{ij} \sigma_i \sigma_j + \rho_{jy} \sigma_j \sigma_y \right] \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_j \right). \end{aligned} \tag{39}$$

Subtracting (38) from (39) yields

$$\begin{aligned} \bar{x}_i \Delta S_i = {} & \mu_y - I - \gamma \Lambda \left(\sigma_y^2 + \rho_{iy} \sigma_i \sigma_y \right) \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_i \right) \\ & - \gamma \Lambda \sum_{j=1}^{n} \rho_{jy} \sigma_j \sigma_y \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_j \right), \end{aligned} \tag{40}$$

which is (24) when $\bar{x}_i = 1$. For $i_{\text{IN}} \in \{1, \ldots, k\}$ and $i_{\text{OUT}} \in \{k+1, \ldots, n\}$ we have

$$\bar{x}_{i_{\text{IN}}} \Delta S_{i_{\text{IN}}} - \bar{x}_{i_{\text{OUT}}} \Delta S_{i_{\text{OUT}}} = \gamma \Lambda (\sigma_y^2 + \rho \sigma \sigma_y) \lambda_M \frac{b}{a+b}. \tag{41}$$

□

Proof of Proposition 1. The result follows immediately from (25) (or its analog (41)). □

Proof of Proposition 2. The only essential difference with the proof of Lemma 3 that implies Proposition 1 is that when firm y is traded before the merger, then (38) becomes

$$\bar{x}_i S_i = \mu_i - \gamma \Lambda \left[\sum_{j=1}^{n} \rho_{ij} \sigma_i \sigma_j \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_j \right) + \rho_{iy} \sigma_i \sigma_y \right]. \tag{42}$$

Subtracting this from (39) (and removing the explicit cost of investment), obtain

$$\begin{aligned} \bar{x}_i \Delta S_i = {} & \mu_y - \gamma \Lambda \left[\left(\sigma_y^2 + \rho_{iy} \sigma_i \sigma_y \right) \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_i \right) \right. \\ & \left. - \rho_{iy} \sigma_i \sigma_y \right] - \gamma \Lambda \sum_{j=1}^{n} \rho_{jy} \sigma_j \sigma_y \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_j \right) \\ = {} & \mu_y - \gamma \Lambda \sigma_y^2 + \gamma \Lambda \left(\sigma_y^2 + \rho_{iy} \sigma_i \sigma_y \right) \lambda_M \frac{b}{a+b} \mathbf{1}_i \\ & - \gamma \Lambda \sum_{j=1}^{n} \rho_{jy} \sigma_j \sigma_y \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_j \right). \end{aligned} \tag{43}$$

Thus (41) in this case is

$$\bar{x}_{i_{\text{IN}}} \Delta S_{i_{\text{IN}}} - \bar{x}_{i_{\text{OUT}}} \Delta S_{i_{\text{OUT}}} = \gamma \Lambda \left(\sigma_y^2 + \rho_{i_{\text{IN}} y} \sigma_{i_{\text{IN}}} \sigma_y \right) \lambda_M \frac{b}{a+b}, \tag{44}$$

and $\bar{x}_{i_{\text{IN}}} \Delta S_{i_{\text{IN}}} > \bar{x}_{i_{\text{OUT}}} \Delta S_{i_{\text{OUT}}} \iff \sigma_y^2 + \rho_{i_{\text{IN}} y} \sigma_{i_{\text{IN}}} \sigma_y > 0$. Notice that unlike in Proposition 1, we do not need to assume that $\sigma_{i_{\text{IN}}} = \sigma_{i_{\text{OUT}}}$ and $\rho_{i_{\text{IN}} y} = \rho_{i_{\text{OUT}} y}$. □

Proof of Proposition 3.

(*i*) Suppose firm y issues $\bar{x}_y$ shares when it goes public (in the main text we normalized $\bar{x}_y$ to one). The stock price of firm y if it is not included in the benchmark is

$$\bar{x}_y S_y^{\text{OUT}} = \mu_y - \gamma \Lambda \left[\sigma_y^2 + \sum_{i=1}^{n} \rho_{iy} \sigma_i \sigma_y \left(1 - \lambda_M \frac{b}{a+b} \mathbf{1}_i \right) \right]. \tag{45}$$

The price of firm y if it enters the benchmark and no other firm leaves it, is

$$\begin{aligned} \bar{x}_y S_y^{\text{IN}} = {} & \mu_y - \gamma \Lambda \left\{ \sigma_y^2 \left[1 - \lambda_M \frac{b}{a+b} \right] \right. \\ & \left. + \sum_{i=1}^{n} \rho_{iy} \sigma_i \sigma_y \left[1 - \lambda_M \frac{b}{a+b} \mathbf{1}_i \right] \right\}. \end{aligned} \tag{46}$$

Taking the difference, $\bar{x}_y \left(S_y^{\text{IN}} - S_y^{\text{OUT}} \right) = \gamma \Lambda \sigma_y^2 \lambda_M b/a + b > 0$.

(*ii*) The price of firm y if it replaces firm k in the benchmark is

$$\begin{aligned} \bar{x}_y \hat{S}_y^{\text{IN}} = {} & \mu_y - \gamma \Lambda \left\{ \left[\sigma_y^2 + \sum_{i=1}^{k-1} \rho_{iy} \sigma_i \sigma_y \right] \right. \\ & \left. \times \left(1 - \lambda_M \frac{b}{a+b} \right) + \sum_{i=k}^{n} \rho_{iy} \sigma_i \sigma_y \right\}. \end{aligned} \tag{47}$$

Taking the difference, $\bar{x}_y \left(\hat{S}_y^{\text{IN}} - S_y^{\text{OUT}} \right) = \gamma \Lambda \left(\sigma_y^2 - \rho_{ky} \sigma_k \sigma_y \right) \lambda_M b/(a+b) > 0$. Thus $\hat{S}_y^{\text{IN}} > S_y^{\text{OUT}} \iff \sigma_y^2 - \rho_{ky} \sigma_k \sigma_y > 0$.

□

Proof of Proposition 4. The subsidy is

$$\Delta S_{i_{\text{IN}}} - \Delta S_{i_{\text{OUT}}} = \gamma \Lambda (\sigma_y^2 + \rho_y \sigma \sigma_y) \left[\lambda_M^A \frac{b}{a+b} + \lambda_M^P \right]. \tag{48}$$

Using $\Lambda = \left[\lambda_M^A/(a+b) + \lambda_D\right]^{-1}$ and expressing $\lambda_M^A = (1-\lambda_D)/2 - \lambda_M^P$, the subsidy becomes

$$\begin{aligned}\Delta S_{i_{\text{IN}}} - \Delta S_{i_{\text{OUT}}} &= \gamma(\sigma_y^2 + \rho_y \sigma \sigma_y)\frac{\frac{b}{a+b}((1-\lambda_D)/2 - \lambda_M^P) + \lambda_M^P}{\frac{1}{a+b}((1-\lambda_D)/2 - \lambda_M^P) + \lambda_D} \\ &= \gamma(\sigma_y^2 + \rho_y \sigma \sigma_y)\frac{\frac{b}{a+b}(1-\lambda_D)/2 + \left(1 - \frac{b}{a+b}\right)\lambda_M^P}{\frac{1}{a+b}(1-\lambda_D)/2 + \lambda_D - \frac{1}{a+b}\lambda_M^P}.\end{aligned} \tag{49}$$

The numerator is increasing and the denominator is decreasing in λ_M^P while holding λ_D constant. □

Lemma 4 below demonstrates, the standard CAPM does not hold in our environment. The standard CAPM applies only in the special case in which no managers are present ($\lambda_M = 0$ and $\lambda_D = 1$). Otherwise, the stocks' expected returns depend on *two factors*, the usual market portfolio and the benchmark.[31]

Lemma 4 (Two-Factor CAPM). *Asset returns* $R_i = D_i/S_i$, $i = 1, \ldots, n$, *can be characterized by*[32]

$$E(R_i) - 1 = \beta_i^{\mathbf{m}} \kappa_{\mathbf{m}} - \beta_i^{\mathbf{b}} \kappa_{\mathbf{b}}, \ \ i = 1, \ldots, n, \tag{50}$$

where

$$\beta_i^{\mathbf{m}} = \frac{Cov(R_i, R_{\mathbf{m}})}{Var(R_{\mathbf{m}})}, \ \ \beta_i^{\mathbf{b}} = \frac{Cov(R_i, R_{\mathbf{b}})}{Var(R_{\mathbf{b}})}, \ \ i = 1, \ldots, n, \tag{51}$$

and $\kappa_{\mathbf{m}} > 0$ *and* $\kappa_{\mathbf{b}} > 0$ *are the market and benchmark risk premia, and* $R_{\mathbf{m}}$ *and* $R_{\mathbf{b}}$ *are the market and benchmark returns, respectively,*

$$R_{\mathbf{m}} = \sum_{j=1}^{n} \omega_j^{\mathbf{m}} \frac{D_j}{\bar{x}_j S_j} = \frac{\sum_{j=1}^{n} D_j}{\sum_{j=1}^{n} \bar{x}_j S_j}, \tag{52}$$

$$R_{\mathbf{b}} = \sum_{j=1}^{n} \omega_j^{\mathbf{b}} \frac{D_j}{\bar{x}_j S_j} = \frac{\sum_{j=1}^{k} D_j}{\sum_{j=1}^{k} \bar{x}_j S_j}, \tag{53}$$

where $\omega_i^{\mathbf{m}} = \bar{x}_i S_i / \sum_{j=1}^{n} \bar{x}_j S_j$, $i = 1, \ldots, n$, *denote the market portfolio weights and* $\omega_i^{\mathbf{b}} = \mathbf{1}_i \bar{x}_i S_i / \sum_{j=1}^{n} \mathbf{1}_j \bar{x}_j S_j$, $i = 1, \ldots, n$, *denote the benchmark portfolio weights.*

Proof of Lemma 4. To show (50), recall that $E(R_i) = \mu_i/(\bar{x}_i S_i)$. Take the ith row of (21), divide both sides by $\bar{x}_i S_i$ and rearrange terms to get

$$\begin{aligned}E(R_i) - 1 &= \gamma \Lambda \sum_{j=1}^{n} \bar{x}_j S_j \mathrm{Cov}(R_i, R_{\mathbf{m}}) \\ &\quad - \gamma \Lambda \lambda_M \frac{b}{a+b} \sum_{j=1}^{k} \bar{x}_j S_j \mathrm{Cov}(R_i, R_{\mathbf{b}}) \\ &= \frac{\mathrm{Cov}(R_i, R_{\mathbf{m}})}{\mathrm{Var}(R_{\mathbf{m}})} \mathrm{Var}(R_{\mathbf{m}}) \gamma \Lambda \sum_{j=1}^{n} \bar{x}_j S_j \\ &\quad - \frac{\mathrm{Cov}(R_i, R_{\mathbf{b}})}{\mathrm{Var}(R_{\mathbf{b}})} \mathrm{Var}(R_{\mathbf{b}}) \gamma \Lambda \lambda_M \frac{b}{a+b} \sum_{j=1}^{k} \bar{x}_j S_j \\ &= \beta_j^{\mathbf{m}} \kappa_{\mathbf{m}} - \beta_j^{\mathbf{b}} \kappa_{\mathbf{b}},\end{aligned} \tag{54}$$

where $\kappa_{\mathbf{m}} = \mathrm{Var}(R_{\mathbf{m}}) \gamma \Lambda \sum_{j=1}^{n} \bar{x}_j S_j, \kappa_{\mathbf{b}} = \mathrm{Var}(R_{\mathbf{b}}) \gamma \Lambda \sum_{j=1}^{k} \bar{x}_j S_j \lambda_M b/(a+b)$. □

[31] This result has been obtained in Brennan (1993).

[32] The left-hand side of Eq. (50) contains the return in excess of the (gross) return on the risk-free bond, where the latter is normalized to one in our model.

References

Baker, M., Wurgler, J., 2013. Behavioral corporate finance: an updated survey. Handb. Econ. Finance 2, 357–424.

Barberis, N., Shleifer, A., Wurgler, J., 2005. Comovement. J. Financ. Econ. 75 (2), 283–317.

Basak, S., Pavlova, A., 2013. Asset prices and institutional investors. Am. Econ. Rev. 103 (5), 1728–1758.

Bena, J., Ferreira, M., Matos, P., Pires, P., 2017. Are foreign investors locusts? The long-term effects of foreign institutional ownership. J. Financ. Econ. 126 (2), 122–146.

Beneish, M., Whaley, R., 1996. An anatomy of the "S&P game": the effects of changing the rules. J. Finance 51, 1909–1930.

Berk, J., DeMarzo, P., 2017. Corporate Finance, fourth ed. Pearson Education.

Blitz, D., Fabozzi, F.J., 2017. Sin stocks revisited: resolving the sin stock anomaly. J. Portf. Manage. 44, 105–111.

Boyer, B.H., 2011. Style-related comovement: fundamentals or labels? J. Finance 66, 307–332.

Brealey, R., Myers, S., Allen, F., 2019. Principles of Corporate Finance, 13th ed. McGraw-Hill Irwin.

Brennan, M., 1993. Agency and Asset Pricing. Working Paper. UCLA.

Buffa, A., Vayanos, D., Woolley, P., 2014. Asset Management Contracts and Equilibrium Prices. Working Paper 20480. NBER.

Buffa, A.M., Hodor, I., 2018. Institutional Investors, Heterogeneous Benchmarks and the Comovement of Asset Prices. Working Paper.

Buss, A., Sundaresan, S., 2020. More Risk, More Information: How Passive Ownership Can Improve Informational Efficiency. Working Paper.

Calomiris, C., Larrain, M., Schmukler, S., Williams, T., 2019. The Search for Yield in Large International Corporate Bonds: Investor Behavior and Firm Responses. NBER Working Paper 25979.

Celikyurt, U., Sevilir, M., Shivdasani, A., 2010. Going public to acquire? The acquisition motive in IPOs. J. Financ. Econ. 96, 345–363.

Chang, Y.-C., Hong, H., Liskovich, I., 2015. Regression discontinuity and the price effects of stock market indexing. Rev. Financ. Stud. 28, 212–246.

Chen, H., Noronha, G., Singal, V., 2004. The price response to S&P 500 index additions and deletions: evidence of asymmetry and a new explanation. J. Finance 59, 1901–1929.

Coles, J., Daniel, N., Naveen, L., 2006. Managerial incentives and risk-taking. J. Financ. Econ. 79, 431–468.

Cremers, M., Petajisto, A., Zitzewitz, E., 2012. Should benchmark indices have alpha? Revisiting performance evaluation. Crit. Finance Rev. 2, 1–48.

Cuoco, D., Kaniel, R., 2011. Equilibrium prices in the presence of delegated portfolio management. J. Financ. Econ. 101 (2), 264–296.

Denis, D.K., McConnell, J.J., Ovtchinnikov, A.V., Yu, Y., 2003. S&P 500 index additions and earnings expectations. J. Finance 58 (5), 1821–1840.

Dessaint, O., Olivier, J., Otto, C., Thesmar, D., 2018. CAPM-Based Company (Mis)valuations. Working Paper No. 3050928. Rotman School of Management.

Fabozzi, F.J., Ma, K.C., Oliphant, B.J., 2008. Sin stock returns. J. Portf. Manage. 35 (1), 82–94.

Gormley, T., Matsa, D., Milbourn, T., 2013. Ceo compensation and corporate risk: evidence from a natural experiment. J. Account. Econ. 56, 79–101.

Graham, J., Harvey, C., 2001. The theory and practice of corporate finance: evidence from the field. J. Financ. Econ. 60, 187–243.

Greenwood, R., 2005. Short- and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage. J. Financ. Econ. 75 (3), 607–649.

Grinblatt, M., Titman, S., 2002. Financial Markets and Corporate Strategy, second ed. McGraw-Hill Irwin.

Gutiérrez, G., Philippon, T., 2017. Investmentless growth: an empirical investigation. Brookings Pap. Econ. Act. Fall 2017, 89–190.

Gutiérrez, G., Philippon, T., 2018. Ownership, concentration, and investment. AEA Pap. Proc. 108, 432–437.

Hacibedel, B., 2018. Index Changes in Emerging Markets. Working Paper. Swedish Institute for Financial Research.

Harris, L., Gurel, E., 1986. Price and volume effects associated with changes in the S&P 500 list: new evidence for the existence of price pressures. J. Finance 41, 815–829.

Hong, H., Kacperczyk, M., 2009. The price of sin: the effects of social norms on markets. J. Financ. Econ. 93, 15–36.

Jacobs, M., Shivdasani, A., 2012. Do you know your cost of capital? Harv. Bus. Rev.. July-August issue, https://hbr.org/2012/07/do-you-know-your-cost-of-capital.
Kacperczyk, M., Sundaresan, S., Wang, T., 2019. Do Foreign Institutional Investors Improve Market Effciency? Working Paper. Imperial College London.
Kaniel, R., Kondor, P., 2013. The delegated lucas tree. Rev. Financ. Stud. 26, 929–984.
Kashyap, A., Kovrijnykh, N., Li, J., Pavlova, A., 2020. Is There Too Much Benchmarking in Asset Management? Working Paper. University of Chicago.
Kim, I., Venkatachalam, M., 2011. Are sin stocks paying the price for accounting sins? J. Account. Auditing Finance 26, 415–442.
Krüger, P., Landier, A., Thesmar, D., 2015. The WACC fallacy: the real effects of using a unique discount rate. J. Finance 70, 1253–1285.
Luong, H., Moshirian, F., Nguyen, L., Tian, X., Zhang, B., 2017. How do foreign institutional investors enhance firm innovation? J. Financ. Quant. Anal. 52, 1449–1490.
Lynch, A., Mendenhall, R., 1997. New evidence on stock price effects associated with changes in the S&P 500 index. J. Bus. 70, 351–383.
Ma, L., Tang, Y., Gómez, J.-P., 2019. Portfolio manager compensation in the U.S. mutual fund industry. J. Finance 72, 587–638.
Massa, M.M., Peyer, U., Tong, Z., 2005. Limits of Arbitrage and Corporate Financial Policy. Working Paper. INSEAD.
McKinsey on Finance, 2004. What is stock index membership worth? Perspect. Corp. Finance Strategy 10, 14–16.
Merton, R., 1987. A simple model of capital market equilibrium with incomplete information. J. Finance 17, 483–510.
Pavlova, A., Sikorskaya, T., 2020. Benchmarking Intensity. Working Paper. London Business School.
Petajisto, A., 2011. The index premium and its hidden cost for index funds. J. Empir. Finance 18, 271–288.
Scholes, M., 1972. The market for securities: substitution versus price pressure and the effects of information on share prices. J. Bus. 45, 179–211.
Shleifer, A., 1986. Do demand curves for stocks slope down? J. Finance 41, 579–590.
Statman, M., Glushkov, D., 2009. The wages of social responsibility. Financ. Anal. J. 65, 33–46.
Stein, J., 1996. Rational capital budgeting in an irrational world. J. Bus. 69, 429–455.
Stulz, R., 1999. What's Wrong with Modern Capital Budgeting? Financial Practice and Education Fall/Winter.
van Binsbergen, J.H., Opp, C.C., 2019. Real anomalies. J. Finance 74, 1659–1706.
Vijh, A., Yang, K., 2008. The Acquisition Performance of S&P 500 Firms. Working Paper. University of Iowa.
Wurgler, J., Zhuravskaya, E., 2002. Does arbitrage flatten demand curves for stocks? J. Bus. 75, 583–608.