Mean variance efficiency of corporate risk taking decisions

Liuren Wu,* Yang Xu,†

Abstract

This paper introduces the concept of cross-sectional mean-variance efficiency in corporate risk-taking decisions. We construct mean-variance ratio forecasts on the return on asset of US public companies and examine how much the forecast can explain the cross-sectional risk-taking variation with a common proportionality coefficient. Estimation shows that cross-sectional mean-variance efficiency explains a large proportion of corporate risk-taking variation at the firm level and almost completely explains the variation at the industry level. Over time, corporations target mean-variance efficient allocations to maintain a stable financial leverage. When a firm deviates from mean-variance efficiency, it rebalances toward it at a slow speed.

JEL Classification: G11, G12, G32

Keywords: Mean-variance efficiency, corporate risk taking, mean-variance ratio, financial leverage, value maximization, risk rebalancing decision, error-correction model

We thank Michael Anthropelos, John Chung-Yen Chu, Andrew Grant, Armen Hovakimian, Dimitrios Kyriazis, Francisco Peñaranda, Artur Rodrigues, George Skiadopoulo, Ivo Welch, Chuanqian Zhang, Dexin Zhou, and seminar participants at Baruch College, University of Piraeu, the 2022 Asian Finance Association Annual Conference, the 2022 European Financial Management Association Annual Meeting at the Campus Bio-Medico University, the 2022 Financial Management Association European Conference at the Emlyon Business School, the 41st Eurasia Business and Economics Society Conference, and the Xiangjiang River Forum in Economics, Finance and Management for Young Scholars for their comments and suggestions. An earlier version of the paper is circulated under the title "Capital structure as an investment decision."

*Zicklin School of Business, Baruch College, The City University of New York, One Bernard Baruch Way, New York, NY 10010, USA. E-mail: liuren.wu@baruch.cuny.edu.

†Memorial University of Newfoundland, Faculty of Business Administration, Morrissey Rd, St. John's, NL A1B 3X9, Canada. E-mail: y.xu@mun.ca.

1. Introduction

Mean-variance analysis has been the cornerstone of modern investment theory since Markowitz (1952). This paper examines the mean-variance efficiency of corporate risk-taking decisions. Different firms take on different business activities to lever their domain expertise and make unique contributions to the investment frontier. Meanwhile, a firm can adjust its risk taking through financing decisions, e.g., holding cash to reduce the risk or levering up its risk level through debt. This paper characterizes the risk-taking decision of a firm with the ratio of the firm's total asset to its total equity. The ratio is one when the firm does not take on any debt and the riskiness of the total asset adjusts for the effect of the firm's cash position. When the firm takes on debt and the ratio is greater than one, the firm's risk target becomes higher than the risk level of its asset investment. Under the classic mean-variance framework and with some simplifying assumptions, we show that the optimal risk taking of a firm is proportional to the mean-variance ratio forecast on its return on asset, with the proportionality coefficient reflecting the firm's risk preference and proportional financing cost per unit risk.

A large literature examines the mean-variance efficiency of market portfolios based on the two-fund spanning principle of the mean-variance frontiers.[1] Different from this literature, this paper does not examine the mean-variance efficiency of a portfolio, but rather examines whether corporations follow the mean-variance principle in setting their risk targets via financing decisions. The information on corporate financing decisions is publicly available for publicly traded companies. This information provides a unique and valuable angle for examining the mean-variance efficiency, not from its implication on the performance of a large portfolio, but from the direct evidence on the actual risk taking decisions across a large universe of companies. Still, if each company were to claim that it follows mean-variance principle but with its own risk preference and hence different proportionality coefficient, we would not be able to examine the validity of the claim no matter what the actual allocation is. In this paper, based on the principle of market integration

[1] In this literature, a portfolio is considered mean-variance efficient with respect to a given set of assets if it is not possible to form another portfolio of those assets with the same expected return but a lower variance. See, for example, MacKinlay and Richardson (1991), Green and Hollifield (1992), Engel and Rodrigues (1993), Levy and Roll (2010), Amengual and Sentana (2010), Sentana (2009).

and cross-sectional consistency, we introduce the concept of cross-sectional mean-variance efficiency, under which a single common risk preference coefficient determines the cross-sectional risk allocation variation across all companies at a given point in time.

To examine the cross-sectional mean-variance efficiency of corporate risk taking decisions, we collect data over the past 40 years on US publicly traded companies. At each date, we measure a company's risk allocation weight as the ratio of total asset to common equity, and construct the mean-variance ratio forecast based on the company's return on asset history. We perform a cross-sectional regression of the risk allocation weight against the mean-variance ratio forecasts. The slope coefficient represents the estimate for the common proportionality coefficient at that date under the condition of cross-sectional mean-variance efficiency. Historical regression estimation shows that a common proportionality coefficient can explain a large proportion of the cross-sectional risk-taking variation. The regression R^2 estimates average at 68%, and vary within a narrow range over time from 54% to 81% within the 5-95th percentile range.

We construct the mean-variance ratio forecasts on each company based on the company's historical return on asset history. The forecasts necessarily contain forecasting errors. The errors-in-variable issue can bias the slope coefficient estimate toward zero. The forecasting bias between our estimate and a firm's internal forecast can create a distortion between our risk taking prediction and the firm's actual target. A firm's actual risk taking level can also deviate from its mean-variance efficient target for a long period of time due to high adjustment cost. To mitigate these issues, we aggregate the firm-level mean-variance forecasts and the risk-taking decisions to the industry level. We expect the aggregation to smooth out some of the forecasting errors and biases, and to average out some of the temporary deviations from the mean-variance targets for individual firms. Indeed, we find that the cross-industry variation of the mean-variance ratio forecasts almost completely explains the cross-industry risk taking variation. When aggregating to the 4-digit GICS industry level, the cross-industry regressions produce extremely high R^2 estimates that average at 92% and vary within a narrow range from 88% to 95% within the 5-95th percentile range. Aggregating to the 2-digit GICS sector level further increases the average regression R^2 to 95%. The average slope

coefficient estimates also become higher with the aggregation.

It is quite remarkable that a common proportionality coefficient can explain a large proportion of the risk taking variation across individual firms, and almost all the risk-taking variation across industries. The cross-sectional regressions highlight the dominant role played by the mean-variance principle in setting corporate risk taking decisions.

In deriving the mean-variance efficient risk allocation, we assume that the financing cost increases with the levered risk level. Under this assumption, the proportionality coefficient reflects not only the risk preference but also the proportional financing cost per unit risk taken. As a motivation for firm creation and expansion, we assume that firms have better financing access than individual investors of the firm and larger firms have better financing access than smaller firms. To test this hypothesis, we use the ratio of interest expense to total debt to approximate a firm's average financing cost, and we regress it against the firm's levered risk level and a relative firm size measure that we construct based on total sales. As expected, financing cost increases with risk. Conditional on the same risk level, the financing cost indeed declines with firm size. Therefore, conditional on the same mean-variance ratio forecast, we expect a larger firm to have a smaller financing cost per unit risk and a larger risk allocation. Regressing the cross-sectional variation of firm risk allocation decisions against the firm size measure indeed generates a positive and statistically significant average slope coefficient estimate; nevertheless, compared to the mean-variance ratio forecasts, the relative contribution of the firm size to the cross-sectional risk taking variation is very small. In a standardized bivariate regression with the mean-variance ratio forecasts and the firm size measure, the average contribution from the mean-variance ratio is 82%, whereas the average contribution from the firm size is merely 8%. The average R^2 of the bivariate regression is 68%, not visibly higher than that from the univariate regression on the mean-variance ratio forecasts.

We also examine the relative contributions of other firm characteristics that the literature has considered for explaining the financing decision variation, including tangibility, profitability, and market-to-book asset ratio. Their contributions are all much smaller than that from the mean-variance ratio. A multivariate

cross-sectional regression including the mean-variance ratio and the other firm characteristics only raises the average adjusted R^2 estimates to 70%, merely 2 percentage point higher than the univariate regression on the mean-variance ratio alone.

By regressing the risk taking decisions against the mean-variance ratio forecasts, we decompose the risk taking decision into two components, a component predicted by cross-sectional mean-variance efficiency and a regression residual capturing the firm's risk deviation from mean-variance efficiency. We examine the linkage between firm valuation and the firm's deviation from mean-variance efficiency. We measure a company's relative value with a simple construct of the Tobin (1969)'s Q, and we perform a cross-sectional local linear regression of the firm value against the firm deviation from mean-variance efficiency. The average estimated relation exhibits a hump shape, where the firm value reaches the maximum at zero deviation and declines with increasing deviation of either direction. The hump-shaped relation provides support that cross-sectional mean-variance efficient corporate risk allocation maximizes firm value.

The investment literature has long realized that the unconditional, long-run performance of an investment depends not only on conditional mean-variance efficiency at each period, but also on the time variation, or market timing, of the mean-variance efficient allocation (Hansen and Richard (1987)). The literature has examined the unconditional performance of various market-timing strategies in the securities market, e.g., constant leverage (Ackermann, Pohl, and Schmedders (2016) and Maurer, To, and Tran (2018)), constant volatility target (Moreira and Muir (2017) and Penãranda and Wu (2022)), constant mean return targets (Penãranda and Wu (2022)), and strategies that maximize some unconditional performance metric (Ferson and Siegel (2001) and Peñaranda (2016)). We find that the time variation of the mean-variance efficient corporate risk allocation best matches a constant leverage strategy, potentially because of the high adjustment cost for financial leverage. Meanwhile, we also find that the risk allocation negatively co-vary with the credit spread between BAA corporate bonds and US Treasuries. The spread can be regarded as a proxy for the variation of market sentiment and financing cost per unit risk, the reciprocal of which determines the proportionality coefficient in the mean-variance allocation.

When the risk allocation of an individual firm deviates from mean-variance efficiency, the firm rebalances its risk position in the future toward the mean-variance efficient level. We buid a cross-sectional error-correction model to examine the corporate dynamic risk rebalancing behavior. Model estimation shows that when the risk allocations of individual firms deviate from mean-variance efficiency, the firms on average rebalance their risk allocation over time toward mean-variance efficiency. Nevertheless, due to the high adjustment cost, this risk rebalancing process is very slow, with an average adjustment speed of merely 4.39% per year. The adjustment speed tends to be higher for companies with higher asset return variance estimates.

When we consider different specifications to explain the corporate risk allocation, the error-correction model also provides an out-of-sample test on whether a specification explains mean-variance efficient risk allocation variation or variations away from mean-variance efficiency. A specification that generates better prediction of the mean-variance efficient variation is expected to have stronger forecasting powers on future risk allocation changes toward to the predicted risk level. Estimation shows that although profitability and market-to-book ratio increase the explanatory power on the cross-sectional risk taking variation, they actually reduce the forecasting power on future leverage changes. Therefore, their contributions are less for the mean-variance efficient risk allocation variation, but more for variations away from mean-variance efficiency.

Our analysis shows that the classic mean-variance framework can be applied successfully to examine the behavior of corporate risk taking decisions. Different from this mean-variance investment perspective, the classic capital structure literature tends to focus on the cost-benefit tradeoff of the debt component (Kraus and Litzenberger (1973) and Scott (1976)), or the adverse selection costs of security issuance (Myers (1984) and Myers and Majluf (1984)). Future research can embed these effects into the classic mean-variance framework for a more integrated analysis. Furthermore, several researchers attempt to examine the optimal capital structure and dynamic rebalancing decisions within the classic structural model setting.[2] These models tend to start with some risk-neutral firm value process, which is a valid market-clearing assumption

[2]See, for example, Leland (1994), Fischer, Hernkel, and Zechner (1989), Goldstein, Ju, and Leland (2001), Titman and Tsyplakov (2007), Strebulaev (2007), and Lambrecht and Myers (2017).

under equilibrium, but can nevertheless prevent researchers from directly examining the effect of firm-specific risk-return tradeoff behaviors. This paper builds on the classic literature on the nature of firms (Hart (1995) and Hart and Moore (1998)), and directly considers the optimal risk-taking behavior of individuals firm, much as in classic investment analysis on individual investors.

2. Mean-variance efficiency of corporate risk taking decisions

Given the particular business a company embarks on, the company can adjust its risk taking through financing decisions: Holding cash when the business is deemed too risky and levering up with debt when the business is below the target risk level. We propose to examine the optimal corporate risk taking decision from the classic mean-variance allocation framework.

To examine the optimal risk taking decision of an individual firm, one cannot start in aggregation with the efficient market hypothesis (Fama (1970)), under which all productions can be organized in the market, leaving no room for the creation and growth of a firm, and all investors can raise financing with equal ease, leaving no relevance for firm-level financing decisions (Modigliani and Miller (1958)). For the market to be nearly as efficient, paradoxically, individual investors and individual firms cannot sit idle, but must work hard to optimize their production and investment decisions. In particular, the literature on the nature of firms, e.g., Coase (1937), Alchian and Demsetz (1972), Benjamin Klein and Alchian (1978), Williamson (1985), and Hart (1995), strives to identify realistic features of the market that motivate the creation and expansion of a firm as a coordinated production process. The coordinated production process can create value via several mechanisms, from reducing negotiating, transacting, contracting, and information costs to increasing market pricing powers. We examine the optimal risk taking decision of an individual firm from the same foundations that motivate the creation and expansion of a firm.

2.1. Two premises that motivate the creation and expansion of a firm

Among the many rationales that have been proposed to motivate the creation and expansion of a firm, we distill two main tenets as the starting assumptions of our analysis.

Assumption 1. A firm must be unique to be a worthy existence. A firm is created, not for gaining exposures to the currently existing market, but quite the opposite, for generating new, independent investment opportunities with superior risk-return tradeoffs, so that the firm creation can expand the investment opportunity and add value to the existing market. A firm with nothing but an exposure to the market portfolio has no worth for existence and shall not be created to begin with.

This premise states one of the most fundamental motivations for the creation of a firm. If a firm could only expect to generate returns proportional to its exposure to the existing market, as dictated by the classic capital asset pricing model built on the efficient market assumption, there would be no reason for the firm to be created to begin with. Investors could simply invest in the existing market portfolio to obtain the same expected return, but without the additional idiosyncratic risk of the firm.

When a firm makes presentations to investors, the emphasis is rarely on how close the firm is to the market portfolio, which can become a synonym to being mediocre. Instead, the emphasis is mostly on how "unique" the firm is to create value and expand the investment frontier. Investors are willing to invest in the firm only when they think the firm can add value and can expand the investment opportunity.

Assumption 2. Size matters for financing power. Everything else equal, the financing cost of a firm can be much lower than the financing cost faced by the average individual investor of the firm. Furthermore, the financing cost tends to decline as the firm size grows.

Market pricing power has long been recognized as an important motivation for firm creation and expansion. A large firm, once becoming dominant in a market, can gain pricing power not only for selling its products at higher market prices, but also for obtaining production factors at lower costs. Anti-trust laws are

put in place to check a firm's pricing power at the expense of smaller competitors and consumers. Cross-firm coordination, or collusion, has also been deemed illegal in most countries. Both types of regulations underpin the practical significance of such market pricing power. Financing cost can be regarded as the cost of a production factor and a firm's pricing power extends to its power in securing cheap financing.

This second premise directly contradicts one of the key foundations of the irrelevance argument that whatever leverage an investor wants to achieve, the investor can achieve it with equal ease and cost either in the market place individually or through the firm, thus making the firm-level financing decision irrelevant. In reality, investors and their associated firms do not have equal access to capital. At the very basic level, most banks create tiered pricing for customers who have different levels of wealth under management. The customers at the higher tier with more wealth under management enjoy a higher savings rate, a lower financing rate, and easier access to capital and services. A firm, as a coordinated group of investors, becomes a larger business customer for the banks, and can enjoy much easier access to capital than the firm's average investors as individuals. Furthermore, as a firm grows larger, not only does it gain pricing power on its production factors and its products, but it also gains pricing power on financing, giving the firm better access to cheaper financing.

An implication of this second premise is that a firm should take on the appropriate financial leverage through the firm's coordinated channel, instead of leaving the financing to its individual investors.

2.2. Mean-variance efficient corporate risk taking decisions

Building on the two premises that we distill from the long list of rationales motivating the creation and expansion of a firm, we derive an individual firm's optimal risk taking decision from the perspective of the classic mean-variance analysis. We measure an firm's time-t risk-taking decision, w_t, as the ratio of its total asset to its common equity at that time, i.e., the proportion of investor wealth allocated to the risky project of the firm.

In classic investment decisions, the allocation weight can be either less than one, in which case a fraction $(1-w_t)$ of the investor wealth is held in cash or a riskfree asset, or greater than one, in which case the investor borrows money to lever up the risky investment. Applying the investment concept to a firm's total asset, the allocation weight w_t can be no smaller than one because a firm's total asset cannot be smaller than the firm's equity. When a firm's investment is deemed too risky, the firm can hold cash reserves. It can also engage in corporate hedging to remove risk exposures that are not part of its core expertise. Both the cash reserves and hedging investments are included in the total asset, so that the riskiness of the total asset is adjusted and the firm's allocation weight to the total asset is no smaller than one by construction.

A firm takes on financial leverage when its total asset investment is more than its equity level, in which case the allocation weight w_t is greater than one. Traditional leverage measures (such as debt ratios) tend to center around the debt amount. Our allocation weight measure w_t captures the financial leverage from the perspective of risk allocation.

We formulate a firm's optimal risk taking decision within the classic mean-variance framework of Markowitz (1952), where the firm chooses the optimal allocation weight w_t to maximize the return-risk tradeoff on its asset investment,

$$\max_{w_t} \quad w_t(\mu_t - c_t(w_t)) - \frac{1}{2}\lambda_t w_t^2 \sigma_t^2, \tag{1}$$

where μ_t and σ_t^2 denote the conditional mean and variance on the firm's return on asset, λ_t captures the risk preference, or relative risk aversion, of the firm, and $c_t(w_t)$ denotes the firm's cost of borrowing. The borrowing cost can vary depending on how much risk the firm takes on and is therefore an implicit function of the optimal risk taking decision. We assume that the implicit function $c_t(w_t)$ satisfies the usual technical conditions to guarantee the existence of a solution for the optimization problem in equation (1).

We set up the optimization problem in (1) by assuming that the management team of the firm represents the interest of its equity investors. Classic moral hazard problems between managers and investors (e.g.,

Jensen and Meckling (1976) and Jensen (1986)) can drive a wedge between the optimal risk target for the firm and the risk level the management team decides to take. While outside the scope of this paper, one can in principle examine the impacts of the moral hazard problem on the optimal risk taking decision based on the deviations between the two. Furthermore, we assume that the investment opportunity as captured by μ_t and σ_t^2 are static conditional forecasts. Under the classic setting of Merton (1971), utility maximization can also generate an intertemporal hedging demand when the investment opportunity varies stochastically and exhibits nonzero correlations with the investment returns.

The first-order condition for the optimal risk taking decision in (1) can be written as,

$$\mu_t - c_t'(w_t)w_t - c_t(w_t) = \lambda_t w_t \sigma_t^2, \tag{2}$$

from which we can write the optimal leverage target, $\overline{w}_t$, as an implicit solution to the following equation,

$$\overline{w}_t = \lambda_t^{-1} \frac{\mu_t}{\sigma_t^2} - \frac{w_t(\overline{w}_t) + c_t'(\overline{w}_t)\overline{w}_t}{w_t \sigma_t^2}. \tag{3}$$

The following proposition summarizes the implication of this implicit solution.

Proposition 1. *Under the premise that a firm targets its optimal risk taking to maximize its risk-return tradeoff on its investment as in (1), a major determinant of the optimal risk allocation target is the firm's expected mean-variance ratio μ_t/σ_t^2 on its investment. Conditional on the same mean-variance ratio, the optimal risk target declines with the level of the financing cost and the speed at which it increases with the risk taking.*

To convert the implicit solution in (3) into an explicit solution, we make the following linear approximation of the financing cost function:

Assumption 3. Financing cost increases linearly with leverage and project risk. The financing cost of a firm

is proportional to its leverage-adjusted risk level:

$$c_t(w_t) = \frac{1}{2}\xi_t \sigma_t^2 w_t, \tag{4}$$

where we multiply the asset return variance by the risk allocation weight w_t to capture the levered risk level. The proportionality coefficient ξ_t can vary both with market financing conditions over time and across firms at a given point in time reflecting the financing power variations. In particular, our second assumption identifies firm size as a differentiating factor in the financing cost function. If larger firms enjoy cheaper financing for the same risk level, the proportionality coefficient ξ_t should decline with increasing firm size.

With the linear approximation of the financing cost function assumption in (4), we can aggregate the impact of business risk and financing cost, and derive the optimal risk taking decision explicitly.

Proposition 2. *Under the premise that a firm targets its risk allocation to maximize its risk-return tradeoff as in (1), and the firm's financing cost is proportional to its leverage-adjusted risk level as in (4), the optimal risk allocation target $\overline{w}_t$ is proportional to the mean-variance ratio of the firm's asset return,*

$$\overline{w}_t = \frac{1}{\lambda_t + \xi_t} \frac{\mu_t}{\sigma_t^2}, \tag{5}$$

where the proportionality coefficient is inversely proportional to the risk aversion of the firm's investors λ_t and the firm's proportional financial cost ξ_t.

2.3. Cross-sectional mean-variance efficiency

Proposition 2 states that a firm's optimal risk allocation increases with the firm's mean-variance ratio on its investment; nevertheless, if each firm claims to be mean-variance efficient as defined by the proposition, but all with different risk preferences, one would not be able to validate the claim empirically regardless of the risk allocation level. To generate empirical implications for cross-sectional comparisons, we introduce the concept of cross-sectional mean-variance efficiency.

Definition. Cross-sectional mean-variance efficiency. The cross-sectional variations of corporate risk allocation are cross-sectionally mean-variance efficient if the variations can be explained by mean-variance optimization with a common risk preference coefficient λ_t.

The concept of cross-sectional mean-variance efficiency can be regarded as a condition for cross-sectional consistency. With market integration, we expect the same cohort of investors to invest in the different firms at a given point in time, leading to the same risk preference across the firms.

By imposing cross-sectional consistency on risk preference, the propositions in the previous subsection generate the following empirically testable implication on the cross-sectional variation of the mean-variance efficient risk allocation across different firms.

Proposition 3. *At a given point in time, when the financing powers of different companies are similar as captured by a similar financing coefficient ξ_t, the cross-sectional variation of the mean-variance efficient corporate risk allocation is driven by the corresponding variation in the conditional mean-variance ratio forecasts,*

$$\overline{w}_{t,i} = \omega_t \frac{\mu_{t,i}}{\sigma^2_{t,i}}, \quad for \quad i = 1, 2, \cdots, N_t \tag{6}$$

where a single proportionality coefficient ω_t links the cross-sectional variation of the risk allocation weights to the corresponding variation of the mean-variance ratio forecasts across the N_t firms at time t. The single coefficient is jointly determined by the common financing power ξ_t and the common risk preference λ_t,

$$\omega_t = \frac{1}{\lambda_t + \xi_t}. \tag{7}$$

When the financing power varies across companies, the variation introduces another source of variation in the corporate risk allocation. The cheaper a firm can secure its financing at a given risk level and accordingly the smaller the coefficient ξ, the larger is its optimal risk allocation. Combining Assumption 2 with the optimal risk-taking decision in Proposition 2, we obtain the following corollary on the effect of firm size on the firm's mean-variance efficient risk allocation.

Corollary 1. *Under the premise that the proportional financing cost per unit risk declines with firm size, conditional on the same mean-variance ratio forecast, the mean-variance efficient risk allocation increases with firm size.*

We can further decompose the mean-variance ratio into two components, (1) the information ratio that measures the risk-adjusted performance $\mathcal{S}_t = \mu_t/\sigma_t$, and (2) the risk level as measured by the return volatility σ_t. The optimal risk taking decision is to increase with the information ratio $\mathcal{S}_t$, but decline with the investment risk σ_t. When the information ratio is similar across a set of firms, the cross-sectional mean-variance efficient risk allocation variation is mainly to achieve risk parity.

Corollary 2. *At a given point in time, when a set of firms have similar information ratio $\mathcal{S}_t$ on investment and similar proportional financing cost ξ, their mean-variance efficient risk allocation is proportional to the reciprocal of their asset return volatility,*

$$\overline{w}_{t,i} = \frac{\zeta_t}{\sigma_{t,i}}, \tag{8}$$

where the common proportionality coefficient is determined jointly by the common risk preference, common proportional financing cost, and the information ratio,

$$\zeta_t = \frac{\mathcal{S}_t}{\lambda_t + \xi_t}. \tag{9}$$

Alternatively speaking, these firms choose their risk allocation to target the same leveraged risk level, ζ_t,

$$\overline{w}_{t,i}\sigma_{t,i} = \zeta_t, \tag{10}$$

for all i.

Corollary 2 on corporate risk taking decisions follows largely the same risk parity argument in the investment literature (Qian (2011)). Several empirical studies also show that corporate financial leverages tend to vary inversely with various measures of the company's operating risk, e.g., Titman and Wessels

(1988), Rajan and Zingales (1995), Fama and French (2002), Faulkender and Petersen (2005), Klasa, Ortiz-Molina, Serfling, and Srinivasan (2017) Lambrinoudakis, Skiadopoulos, and Gkionis (2019), Lemmon, Roberts, and Zender (2008), and Im, Kang, and Shon (2020).

2.4. Dynamic risk rebalancing decision

As the firm's business evolves over time, the firm's expectation about its future prospect of risk and profitability can vary. As such, a firm's optimal risk taking is fundamentally a dynamic decision that necessitates dynamic rebalancing. To capture this dynamic rebalancing process, this section extends the static decision in (1) to a dynamic setting. We consider the optimal risk taking decision at time t, given a starting risk position at time $t-1$. In particular, we choose the optimal risk rebalancing amount $x_t = w_t - w_{t-1}$ to maximize the firm's risk-return tradeoff, subject to an adjustment cost, $a_t(|x_t|)$, which tends to increase with the adjustment amount $|x_t|$. Formally,

$$\max_{x_t} \quad (w_{t-1}+x_t)(\mu_t - c_t(w_{t-1}+x_t)) - \frac{1}{2}\lambda_t(w_{t-1}+x_t)^2\sigma_t^2 - |x_t|a(|x_t|) \tag{11}$$

For an explicit solution, we maintain the linear approximation of the financing cost structure in (4), and we further assume a proportional adjustment cost structure.

Assumption 4. The adjustment cost is linear in the adjustment amount. We approximate the adjustment cost with a linear structure,

$$a(|x_t|) = \frac{1}{2}\eta_t|x_t|, \tag{12}$$

with η_t capturing the proportional adjustment cost at time t.

With the above assumption, we can derive the following optimal rebalancing decision.

Proposition 4. *Large transaction cost dictates that a firm's risk allocation can persistently deviate from its mean-variance efficient target. A firm rebalances its risk allocation toward its mean-variance efficient target*

over time,

$$x_t = \kappa_t \left(\overline{w}_t - w_{t-1} \right), \tag{13}$$

where $\overline{w}_t$ *denotes the mean-variance efficient risk allocation target at time t in the absence of adjustment cost as in (5). The speed of adjustment* κ_t *declines with the proportional adjustment cost* η_t*, but increases with the riskiness of the investment* σ_t^2*,*

$$\kappa_t = \frac{1}{1 + \frac{\eta_t}{(\xi_t + \lambda_t)\sigma_t^2}}. \tag{14}$$

A corporation can in principle adjust its risk allocation by altering the equity amount, the debt amount, or the investment itself, but with different adjustment costs for different channels at different situations (Leary and Roberts (2005)). The proportional adjustment cost structure in (12) represents a stylized approximation, which results in the optimal dynamic adjustment amount in (13) proportional to the deviation from the optimal target. In practice, both equity and debt issuance incur significant fixed costs. As a result, the actual capital structure rebalancing actions are very much discrete, and only happen when the actual risk position deviates from the optimal target by a large amount or when new capital needs become imminent.

3. Data collection and summary behaviors

We collect data on US publicly traded companies. We retrieve quarterly financial reporting information and the Global Industry Classification Standard (GICS) classification for each company from Compustat. We sample the data monthly at the end of each month. At each sampling date, we select the universe of companies based on the following filtering criteria: (i) The company is public traded with common shares listed on one of the three main US stock exchanges, AMEX, NYSE, or NASDAQ, with valid GICS classification information. (ii) The company has strictly positive market capitalization, common book equity, total book asset, and total sales, with the ratio of market capitalization to book equity ratio no larger than 10. (iii) The company has a minimum of 12 quarters of recent history in income statements.

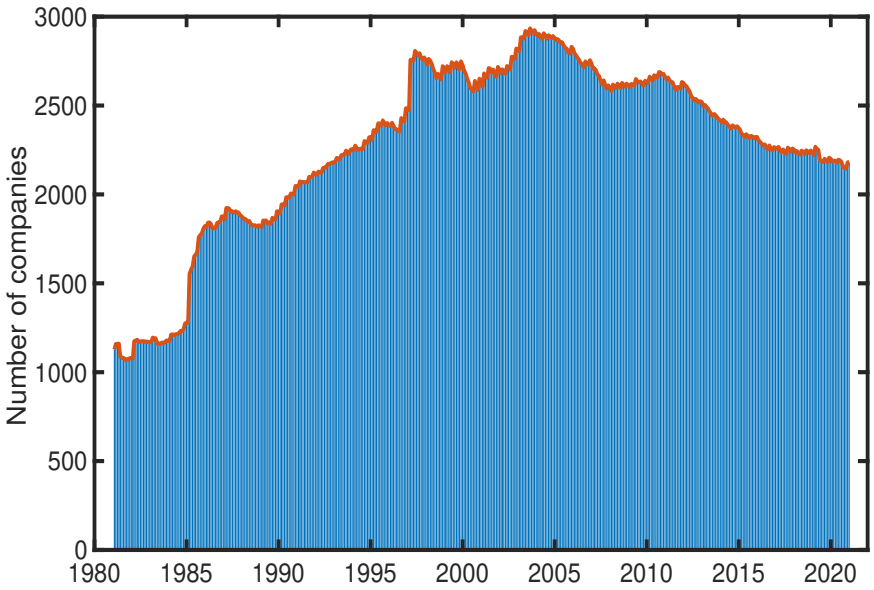


Figure 1
Number of selected companies per month
The bar chart shows the number of selected companies at each month from January 1981 to December 2020.

We use the company's total book asset and common equity to construct the risk allocation measure, and we use total sales to construct a relative company size measure. The strict positivity filter on total asset, common equity, and total sales guarantees that the risk allocation and relative size measures are well-defined. We match the book equity value with the market capitalization at the end of the corresponding quarter. The filter on the upper limit of the market-to-book ratio, as in Baker and Wurgler (2002), guarantees that the company's book common equity is not extremely small relative to its market valuation. These companies tend to have high valuation purely because of future prospects that are difficult to capture based on the company's historical performance. Finally, we require a minimum of three years (12 quarters) of recent history in the company's return on asset to generate the mean-variance ratio forecast.

We perform our analysis from January 1981 to December 2020, for a total of 480 months. Figure 1 plots the number of selected companies at each month over the 40 years span. The size of the cross section ranges from 1,069 names (in August 1981) to 2,934 names (in August 2003), with an average of 2,273 names per sampling date.

3.1. Risk allocation

We construct the risk allocation weight $w_{t,i}$ for each company i at each date t as the ratio of the company's total book asset (AT) to its book common equity (CE). We use book value to define the risk allocation to avoid temporary distortions induced by market pricing variations. Table 1 reports the time-series averages of the cross-sectional summary statistics on the risk allocation weight in panel A. On average, the allocation weight varies from 1.19 at the 5th percentile to 13.39 at the 95th percentile. To mitigate the impact of extreme observations, we winsorize the data cross-sectionally at the 5-95% level for all our cross-sectional analysis. The mean and standard deviation statistics are computed on the winsorized sample at each date. The winsorized mean risk allocation weight is 3.56. The winsorized cross-sectional standard deviation is at a similar level at 3.35. The average cross-sectional distribution shows positive skewness and excess kurtosis.

3.2. Return on asset behaviors

We generate the mean-variance ratio forecast on each company based on its return on asset history. First, we construct the company's return on asset at each quarter as the ratio of the company's reported quarterly operating income after depreciation (OIADP) to the company's total book asset at the beginning of the quarter. Then, we construct the historical mean and volatility estimators using the company's return on asset history over the past five years, with a minimum requirement of 12 quarters.

Table 1 reports the time-series averages of the cross-sectional statistics on the historical estimators in panel B. The mean estimators ("Mean") are reported in annualized percentage points. Although all companies are created in anticipation of making a positive return in the long run, any single company can have losses over any period of time. The mean return on asset in our sample averages around 8% per year, but the average value at the 5th percentile is negative at -8.18%. The cross-sectional distributions of the mean estimators are reasonably symmetric, with small average skewness and excess kurtosis estimates.

Potentially due to accounting smoothing, some companies report extremely smooth quarterly return

Table 1
Summary statistics of firm characteristics

			Percentiles					Tails	
	Mean	SD	5	25	50	75	95	skew	kurt
	A. Risk allocation								
w	3.56	3.35	1.19	1.58	2.24	3.62	13.39	1.99	2.91
	B. Return-on-asset behaviors								
Mean	8.00	7.73	-8.18	3.24	7.83	12.86	23.04	-0.05	-0.07
Volatility	6.35	3.76	1.09	3.65	5.68	8.41	15.09	0.71	-0.08
IR	1.36	0.85	-0.91	1.16	1.77	1.93	1.99	-1.61	1.44
IRF	1.36	0.35	0.69	1.14	1.45	1.59	1.82	-0.94	1.26
MVF	0.40	0.43	0.08	0.15	0.24	0.40	1.62	2.11	3.55
	C. Other firm characteristics								
Size	0.63	1.06	0.01	0.05	0.17	0.62	4.10	2.29	4.30
Tangibility	0.51	0.12	0.27	0.44	0.52	0.59	0.73	-0.17	-0.40
Profitability	7.40	9.12	-12.32	2.25	7.07	13.07	25.16	-0.11	0.00
MB ratio	2.13	1.32	0.66	1.15	1.70	2.72	5.51	1.18	0.57

Entries report the time-series averages of the cross-sectional summary statistics on company risk allocation weights (w) in panel A, return on asset behaviors in panel B, and other firm characteristics in panel C. The risk allocation weight is constructed as the ratio of total book asset to book common equity. The return on asset statistics include the annualized percentage mean return on asset ("Mean"), annualized percentage return volatility ("Volatility"), the annualized information ratio ("IR"), its cross-sectionally smoothed forecast ("IRF"), and the mean-variance ratio forecast ("MVF") constructed as the ratio of the information ratio forecast to the volatility estimate. Firm characteristics in panel C include (i) a company relative size measure ("Size"), constructed as the total sales of the company divided by the corresponding industry sector average at the 2-digit GICS level, (ii) asset tangibility ("Tangibility"), constructed as the ratio of net property, plant, and equipment to total assets, (iii) profitability, measured by the return on asset of the most recent quarter in annualized percentages, and (iv) the market-to-book ("MB") ratio, constructed as the ratio of market capitalization to book common equity. The summary statistics include the mean, standard deviation ("SD"), percentile values, skewness ("skew") and excess kurtosis ("kurt"), all computed on the cross-sectional samples winsorized at 5th-95th percentiles.

on asset numbers. To avoid degeneration in our mean-variance ratio forecasts, we construct the volatility estimators ("Volatility") without demeaning the quarterly return series. The estimates are reported in annualized percentage points. On average, the volatility estimates vary from a low of 1.09% at the 5th percentile to 15.09% at the 95th percentile. The cross-sectional distributions of the volatility estimates show positive skewness but little excess kurtosis on average.

We construct the annualized information ratio ("IR") as the ratio of annualized mean return to annualized volatility. The estimates vary from an average of -0.91 at the 5th percentile to an average of 1.99 at the 95th percentile. The information ratio estimates can become negative due to the negative mean return estimates. The average cross-sectional distribution shows negative skewness and positive excess kurtosis.

To generate a robust mean-variance ratio forecast from the historical estimates, we assume that companies from the same industry group share similar risk-adjusted performance. With this assumption, we perform cross-sectional smooth averaging on the information ratio estimates within each industry group at the 4-digit GICS level. We treat the industry-smoothed information ratio estimate as our information ratio forecast ("IRF"). Table 1 shows that, through the industry smoothing, the average information ratio forecast at the 5th percentile becomes positive at 0.69. As no companies are created to lose money, it makes sense to generate a positive information ratio forecast, even for companies with a negative average return over the past five years.

To construct the mean-variance ratio forecasts ("MVF"), we set a lower bound on the information ratio forecast to be no lower than 0.5, and we divide the information ratio forecast by the volatility estimate in percentage points. The mean-variance ratio forecast varies from an average of 0.08 at the 5th percentile to an average of 1.62 at the 95th percentile. The average cross-sectional distribution shows positive skewness and excess kurtosis.

3.3. Firm characteristics

We also construct a set of commonly identified firm characteristics that can affect corporate financing decisions. Through a comprehensive examination, Frank and Goyal (2009) conclude that the most reliable explanatory variables include the industry average, the market-to-book ratio, asset tangibility, firm profits, and size. They find that firms with larger size, more tangible assets, but lower profits and lower market-to-book ratios tend to have higher leverage. We construct similar measures and examine their relative contribution to the corporate risk taking variation.

We use total sales to represent the company size, and we scale the total sales of a company by the corresponding industry average at the 2-digit GICS level to capture the relative pricing power of the company relative to its industry peers. We measure the asset tangibility as the ratio of net property, plant, and equipment to total assets; we measure the profitability of a company with its return on asset in the last quarter; and we construct the market-to-book ratio as the ratio of market capitalization to book common equity.

Table 1 reports the summary statistics of these firm characteristics in panel C. The relative size varies from an average of 1% of the industry average at the 5th percentile to 4.1 times the industry average at the 95th percentile. The tangibility measure averages at 0.51, and varies within a narrow range from 0.27 at the 5th percentile to 0.73 at the 95th percentile. The return on asset as a profitability proxy has an annualized average of 7.4%, and varies from -12.32% at the 5th percentile to 25.16% at the 95th percentile. Compared to the statistics on the five-year average return in panel B, the quarterly return on asset is much more volatile and varies over a much wider range. Finally, the market-to-book ratio has an average of 2.13 and a median of 1.7, suggesting that the market capitalization tends to be higher than the book common equity value on average. Among the four variables, the cross-sectional distributions of size and market-to-book ratio show positive skewness and excess kurtosis, but the skewness and excess kurtosis estimates are small for asset tangibility and profitability.

4. Cross-sectional mean-variance efficiency in corporate risk taking decisions

This section examines the cross-sectional mean-variance efficiency of the risk taking decisions of US publicly traded companies over the past 40 years. First, we examine how much of the cross-sectional risk allocation variation can be explained by the mean-variance ratio forecast with a single proportionality coefficient at each date. Second, we aggregate the risk allocation weights and the mean-variance ratio forecasts to industry level to examine the cross-sectional mean-variance efficiency at the industry level. Third, we test the empirical validity of our assumption that larger firms have lower financing cost at the same risk level. We examine the implication of this assumption on the risk taking decision. Fourth, we examine the contributions of other firm characteristics in explaining the risk allocation variation. Finally, we establish a linkage between mean-variance efficient risk allocation and firm value maximization.

4.1. Cross-sectional mean-variance efficiency at firm level

When different firms have similar proportional financing cost, corporate risk taking decisions are cross-sectionally mean-variance efficient if their cross-sectional variations can be explained by the variation in the mean-variance ratio forecasts with a single proportionality coefficient at each date. To examine the degree of cross-sectional mean-variance efficiency at the firm level, we perform a cross-sectional regression of the risk allocation weight $w_{t,i}$ on the mean-variance ratio forecast $\mathrm{MVF}_{t,i}$ at each date t,

$$w_{t,i} = \alpha_t + \omega_t \mathrm{MVF}_{t,i} + e_{t,i}. \tag{15}$$

According to Proposition 3, the null value of the intercept α_t is zero, and the slope ω_t captures the single proportionality coefficient that defines the cross-sectional mean-variance efficiency. The R^2 of the regression represents an estimate on the degree of the cross-sectional mean-variance efficiency at the firm level.

Table 2 summarizes in panel A the results from the firm-level cross-sectional regressions. The statistics include the time-series averages ("Average"), the Newey and West (1987) t-statistics ("NW") on the average,

Table 2
Cross-sectional mean-variance efficiency of corporate risk taking decisions

			Percentiles				
	Average	NW	5	25	50	75	95
	A. Firm level regression						
α	0.97	6.44	0.33	0.60	1.02	1.24	1.65
ω	6.77	10.83	3.90	5.57	6.58	8.52	9.15
R^2	0.68	31.58	0.54	0.64	0.67	0.72	0.81
	B. Industry level regression						
α	0.79	4.69	0.07	0.34	0.84	1.13	1.48
ω	7.40	10.19	4.08	5.98	7.06	9.37	10.06
R^2	0.92	140.72	0.88	0.91	0.92	0.94	0.95
	C. Sector level regression						
α	0.70	4.39	-0.02	0.29	0.72	1.04	1.47
ω	7.51	10.37	4.01	6.13	7.25	9.39	10.12
R^2	0.95	85.10	0.84	0.93	0.96	0.98	0.99

Entries report the summary statistics of the intercept (α), the slope (ω), and the R^2 estimates from the cross-sectional regressions of the risk allocation weights against the mean-variance ratio forecasts. The statistics include the time-series sample average, the Newey-West t-value (NW) on the average, and the percentile values. The three panels perform the regressions on the firm level (A), the 4-digit GICS industry level (B), and the 2-digit GICS sector level (C), respectively.

and the percentile values of the coefficients and regression R^2 estimates. The average slope estimate is strongly positive and highly statistically significant. The average intercept estimate is also positive and statistically significant, but with a much smaller t-value.

The cross-sectional regressions generate high R^2 estimates, ranging from 54% at the 5th percentile to 81% at the 95th percentile, with a sample average of 68%. The high R^2 estimates highlight the degree to which the corporate risk taking decisions comply with cross-sectional mean-variance efficiency.

4.2. Cross-sectional mean-variance efficiency at industry level

The firm level cross-sectional regression generates high explanatory powers overall, but large deviations remain. The deviations can come from several sources. First, our simple history-based mean-variance ratio forecasts can certainly contain forecasting errors. The errors-in-variable issue not only reduces the explanatory power, but also biases the slope coefficient toward zero. Second, the financing power of different firms can be different, especially across firms with different sizes. Third, as a company's business evolve over time and as the adjustment cost can be high, the risk position of a company can deviate from its mean-variance efficient target at any given point in time.

To mitigate the impacts of forecasting errors and risk deviations at the individual firm level, this section examines the cross-sectional mean-variance efficiency at the industry level. Different industries take on different types of businesses, which can naturally have different risk-return tradeoff prospects. Aggregating to the industry level highlights the cross-industry differences in risk-return behaviors, while smoothing out the forecasting errors, the financing cost differences, and the risk deviations of individual firms.

Based on the GICS classification, we consider smooth averaging both at the 4-digit industry group level and at the broader 2-digit industry sector level. As each date, for each industry group, we perform simple equal-weighted averaging on the risk allocation weights to smooth out the temporary deviations and financing power differences in individual firms, and we also perform equal-weighted averaging on the mean-variance ratio forecasts to smooth out the individual forecasting noise. Then, we repeat the cross-sectional regressions at both the industry level and the sector level. Table 2 summarizes the results in panels B and C, respectively. The GICS classification includes 26 industry groups and 11 sectors. The number of names varies across sectors and industry groups. The regressions exclude groups with fewer than 10 companies.

By aggregating the firm level estimates to the industry and the sector level, the average intercept estimates become smaller, but the slope estimates become larger, potentially due to the mitigation of the error-in-variable issue. The Newey and West (1987) t-statistics on the average intercept estimate become

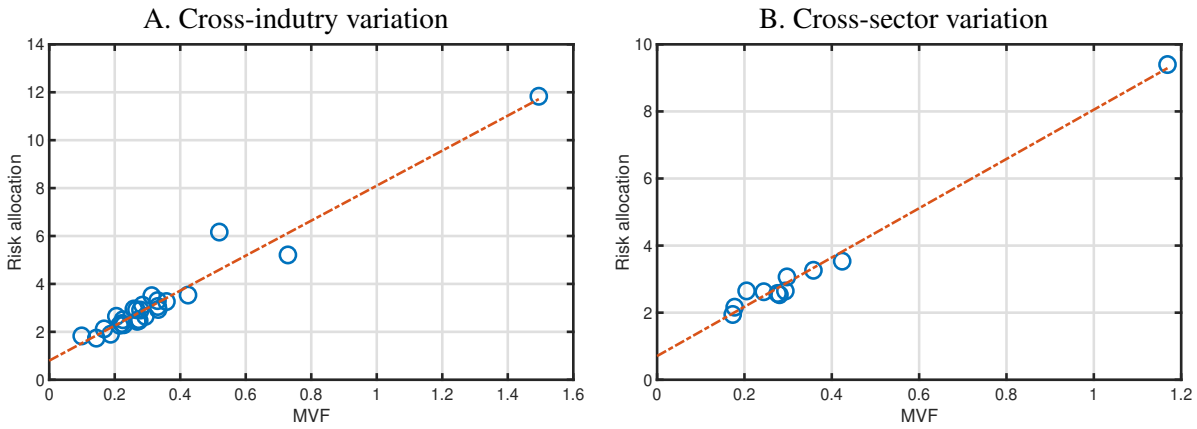


Figure 2
Scatter plots of risk allocation against mean-variance ratio forecasts
We compute the average risk allocation weights and mean-variance ratio forecasts within each 4-digit GICS industry and each 2-digit GICS sector. The two panels show the scatter plots on the time-series averages of the cross-industry (panel A) and cross-sector variation (panel B). The dash-dotted line in each panel represents a linear fit.

smaller, but the Newey and West (1987) t-statistics on the average slope coefficient estimates remain at about the same level despite the much smaller sample size.

With the aggregation, the regression R^2 estimates become much higher. At the 4-digit GICS industry level, the regression R^2 estimates average at 92%. Over time, the R^2 estimates of the cross-industry regressions vary within a very narrow range between 88% at the 5th percentile to 95% at the 95th percentile. At the 2-digit sector level, the average R^2 becomes even higher at 95%, and the R^2 estimates vary over time within a similarly narrow range. Therefore, cross-industry risk-taking variation can almost completely be explained by the corresponding cross-industry variation in the mean-variance ratio forecasts. Corporate risk taking decisions are nearly completely cross-sectionally mean-variance efficient at the industry level.

Figure 2 plots the scatter points of the time-series averages of the average risk allocation weight in each industry (panel A) and sector (panel B) against the corresponding average mean-variance ratio forecast. The dash-dotted line in each panel represents the linear fit. In both panels, the average scatter points fall closely onto the linear line.

In panel A, the industry with the largest risk allocation is banking with an average of 11.8, followed

by financial services with an average of 6.2 and insurance with an average of 5.2. These three industries constitute the finance sector, which shows up in panel B as the sector with the largest risk allocation at an average of 9.4. The average risk allocation weights for all other industries and sectors are much smaller. Historically, many empirical studies exclude financials because they are "special" relative to other industries. The scatter plot shows that they indeed look special in terms of the much higher leverage level relative to other industries, but the much higher leverages are dictated by their high mean-variance ratio forecasts and remain mean-variance consistent with the risk allocations of other industries.

In related literature, Gornall and Strebulaev (2018) explain the leverage difference between banks and their borrowers via a model of supply chain. In such a theoretical setting, bank seniority dramatically reduces bank asset volatility. Our analysis of the historical return-on-asset series on financial companies confirms their theoretical prediction. Among the 11 sectors, the financials sector has indeed the lowest average return volatility at 2.53%, compared to the grand average of 6.35% reported in Table 1. The sector's average information ratio estimate is also the highest at 1.72, compared to the grand average of 1.36. The high information ratio and the low volatility contribute to the high mean-variance ratio forecast and accordingly the high financial leverage for the sector.

Industry average has been found to be a strong explanatory variable for the cross-sectional variation of financial leverage (e.g., Leary and Roberts (2014) and Frank and Goyal (2009)), but for reasons not well understood (Welch (2012)). By applying classic mean-variance theory to corporate risk taking decisions, we show that the cross-industry variation in corporate risk taking is almost completely consistent with mean-variance efficient risk allocation.

4.3. The effect of firm size on financing cost and optimal risk taking

As a motivation for the creation and expansion of a firm and firm-level financing decision, we assume that a firm can enjoy easier access to cheaper financing than can its individual investors, and the larger a firm becomes relative to its peers, the better access the firm enjoys to cheaper financing. A lower financing cost

per unit risk allows a larger firm to target a higher optimal risk allocation weight (Corollary 1).

To examine our assumption on firm size and its implication on the optimal risk allocation, we construct a relative company size measure (RS) based on the company's total sales. We also construct a financing cost measure (FC) as the ratio of total interest expense to total debt. According to our assumptions, for a given company, the financing cost shall increase with increased risk; conditional on the same risk level, the financing cost shall become lower for larger companies. To test this hypothesis, we perform the following cross-sectional regression at each date,

$$\mathrm{FC}_{t,i} = \alpha_t + \beta_t \left(w_{t,i} \sigma^2_{t,i} \right) + \gamma_t \mathrm{RS}_{t,i} + e_{t,i}, \tag{16}$$

where the slope coefficient β_t measures how the financing cost responds to the leverage-adjusted risk level and the coefficient γ_t measures how the financing cost responds to firm size conditional on the same risk level. Under our hypothesis, β_t should be positive whereas γ_t should be negative.

Table 3 summarizes in panel A the results from the cross-sectional regression, including the time-series averages of the coefficients and R^2 estimates, and the Newey and West (1987) t-statistics on the time-series average. The time-series average of the β_t coefficient is positive and strongly statistically significant, confirming that financing cost increases with risk. On the other hand, the time-series average of the γ_t coefficient is negative and statistically significant, validating our hypothesis that financing cost declines with increasing firm size.

Panel B of Table 3 considers an alternative construction of the financing cost measure as the ratio of interest expense to long-term debt. Repeating the regressions with the alternative financing cost construction generates similar results. The time-series average of the β_t coefficient remains positive and statistically significant, whereas the time-series average of the γ_t coefficient remains negative and statistically significant.

If financing cost declines with firm size, the optimal risk allocation shall increases with it conditional on the same mean-variance ratio forecast. To gauge the relative contribution of firm size to the risk taking

Table 3
Size effect on company financing cost

	α_t	β_t	γ_t	R^2
	A. FC = Interest Expense/(Total Debt)			
Average	7.78	12.67	-0.05	0.15
NW	11.63	6.52	-2.24	0.37
	B. FC = Interest Expense/(Long-term Debt)			
Average	11.35	50.19	-0.10	0.12
NW	11.83	7.10	-2.27	0.51

At each date, we perform a cross-sectional regression of a financing cost measure (FC) on the leverage-adjusted risk level and the relative company size measure. Entries report the time-series average ("Average") of the regression coefficients and R^2 estimates, and the Newey-West t-statistics ("NW") on the time-series average. The financing cost measure is constructed as the ratio of total interest expense to total debt in panel A, and the ratio of interest expense to long-term debt in panel B, both in percentage points.

variation, we cross-sectionally regress the risk allocation weights at each date on the relative size measure, both with and without controlling the mean-variance ratio forecast. To make the regression coefficients more comparable relative to one another, we cross-sectionally standardize both the risk allocation weight and the explanatory variables by removing the cross-sectional sample average and dividing the de-meaned observations by the sample standard deviation. All variables are cross-sectionally winsorized before the standardization. The slope of the standardized univariate regression reflects the cross-sectional correlation between the risk allocation and the explanatory variable, and the coefficient squared is just the regression R^2. For the standardized bivariate regression, the two slope coefficients reflect the relative contribution and direction of the two explanatory variables to the risk allocation variation.

Table 4 summarizes the standardized regression results, with panel A for the univariate regression and panel B for the bivariate regression. For each coefficient, we report the time-series average of the coefficient and the Newey and West (1987) t-statistics (NW) on the average. The last column in panel B reports the average adjusted R^2 estimates of the bivariate regression. In the first row of the statistics, we repeat the univariate regression on the mean-variance ratio forecast, this time standardized, and use this as a benchmark

Table 4
Marginal contributions to the cross-sectional risk allocation variation

	A. Univariate		B. Bivariate				
			MVF		Other		R^2
	Average	NW	Average	NW	Average	NW	
MVF	0.82	64.03	–	–	–	–	0.68
Size	0.08	3.95	0.82	61.25	0.08	5.39	0.68
Tangibility	-0.02	-0.51	0.82	63.67	-0.02	-2.98	0.67
Profitability	-0.21	-8.55	0.82	60.64	0.00	0.27	0.68
MB ratio	-0.14	-7.96	0.87	90.60	0.15	14.28	0.70

At each date, we perform a standardized univariate cross-sectional regression of the risk allocation weight against each of the explanatory variables, and we also perform a standardized bivariate cross-sectional regression with the mean-variance ratio forecast (MVF) as the common variable. Entries report the time-series averages of the slope coefficient estimates and the Newey-West t-statistics (NW) on the average slope. The last column reports the average adjusted R^2 from the bivariate regressions.

for other regressions. The average cross-sectional correlation between the mean-variance ratio forecast and the leverage ratio is 82%.

The second row summarizes the contribution of the relative company size measure. From the univariate regression, the relative size measure shows a positive average cross-sectional correlation with the risk allocation weight at an average of 8%. The average correlation estimate is magnitudes lower than the 82% benchmark from the mean-variance ratio forecast; nevertheless, the Newey-West t-statistic shows that the average contribution of the relative size measure is statistically significant. The positive contribution remains after we control for the mean-variance ratio forecast in the standardized bivariate regression. The slope estimate on the relative size measure in the bivariate regression also averages at 8% while the slope estimate on the mean-variance ratio forecast averages at 82%. Because of the much smaller contribution from the relative size measure, the bivariate regression does not visibly increase the average adjusted R^2 estimate. The average adjusted R^2 estimate remains around 68%. Therefore, although firm size is helpful in explaining the financing cost difference across firms, its additional explanatory power on the risk allocation variation is small. The cross-firm difference in financing power, or its impact on risk allocation, is

much smaller than the cross-firm difference in the mean-variance ratio forecast and its impact on the risk allocation.

4.4. Contributions from other firm characteristics

The literature has identified several firm characteristics explaining the cross-sectional financial leverage variation. To examine their relative contribution to corporate risk taking decision, we perform similar standardized univariate and bivariate cross-sectional regressions on these variables. The remaining rows of Table 4 summarize the regression results on these variables.

Frank and Goyal (2003, 2009) and DeAngelo and Roll (2015), among others, find that the asset tangibility ratio can positively explain the cross-sectional capital structure variation. The general argument is that more tangible assets can increase the value of collateral and reduce the financing cost. We do not find this effect to be significant in our sample. The slope estimates average around -2% under both the standardized univariate and bivariate specifications. The average slope estimate is not statistically significant in the univariate regression and only becomes statistically significant under the bivariate setting.

By itself, the profitability measure contributes strongly negatively to the risk allocation variation. The correlation estimates average at -21%, with strong statistical significance. The literature has documented similarly negative relations between profitability and the leverage ratio.[3] In principle, holding everything else equal, a high return on asset should contribute positively to the mean-variance ratio forecast and accordingly positively to the optimal risk allocation. Nevertheless, due to the large adjustment cost, the risk rebalancing decision can be slow and discrete (Strebulaev (2007)). In between, a company's risk allocation can move away from its target level. One channel is through the cumulation of wealth via operating profits. When a company has been making more money than it distributes, its retained earnings account increases, which increases its total asset and common equity and reduces the risk allocation. Therefore, the wealth

[3] See, for example, Hovakimian, Opler, and Titman (2001), Flannery and Rangan (2006), Frank and Goyal (2009), Graham and Leary (2011), Abel (2018), and Chen, Harford, and Kamara (2018).

cumulation effect of operating profits can generate a negative contribution to the cross-sectional risk allocation variation.[4] In our sample, once controlling for the mean-variance ratio effect in the bivariate regression, we obtain a near-zero loading on the profitability measure, potentially due to the cancellation between the positive informational effect of profitability on future mean-variance ratio forecast and the negative wealth cumulation effect on the actual risk allocation.

By itself, the market-to-book equity ratio contributes negatively to the risk allocation variation. The cross-sectional correlation estimates average at -14%, with strong statistical significance. Once we control for the mean-variance ratio forecast, the bivariate regression generates an average positive slope on the market-to-book ratio at 15%, also with high statistical significance. Adding the market-to-book ratio to the standardized bivariate regression increases the contribution of the mean-variance ratio forecast from an average of 82% to an average of 87%. The bivariate regression also generates the highest average adjusted-R^2 estimate at 70%, two percentage points higher than from the benchmark univariate regression.

The literature uses the market-to-book ratio mainly as a relative value measure. The classic market timing argument (e.g., Fischer and Merton (1984) and Lucas and McDonald (1990)) suggests that when a company is overvalued, the company tends to increase its net stock issuance to benefit from the temporary mispricing opportunity. The stock overvaluation can also have a spillover effect to the debt market, allowing the company to also raise more debt at a cheaper cost.[5] Meanwhile, the catering theory suggests that optimistic market valuations can encourage firms to take actions to confirm such expectations by raising more capital and making more investments (Polk and Sapienza (2009)). As our regressions are on the risk allocation weights instead of net debt or equity issuances, the net contributions of these different considerations are not clear. In our sample, once controlled for the mean-variance ratio forecast, the market-to-book ratio on average has a net positive contribution to the cross-sectional risk allocation variation.

To capture the joint contributions of the mean-variance ratio forecasts and the list of firm characteristics,

[4] Danis, Rettl, and Whited (2014) find that the contributions of profitability can indeed become positive when the firms are at or close to their optimal leverage level, but otherwise can become negative due to variations away from the target.

[5] See, for example, Dong, Hirshleifer, and Teoh (2012), Dong, Loncarski, Horst, and Veld (2012), and Hu, Sy, and Wu (2020) for empirical evidence supporting the market timing argument on misvalued companies.

Table 5
Explaining corporate risk allocation variation in a multivariate setting

	Average	NW	Percentiles				
			5	25	50	75	95
MVF	0.86	70.94	0.78	0.83	0.86	0.88	0.91
Size	0.06	5.28	-0.00	0.03	0.07	0.09	0.11
Tangibility	-0.02	-3.42	-0.05	-0.03	-0.02	-0.01	0.02
Profitability	-0.06	-11.03	-0.09	-0.07	-0.06	-0.05	-0.03
MB ratio	0.15	16.07	0.07	0.13	0.16	0.18	0.20
R^2	0.70	40.75	0.57	0.67	0.69	0.75	0.79

At each date, we perform a standardized cross-sectional regression of the risk allocation against the mean-variance ratio forecast (MVF), relative company size, asset tangibility, profitability, and market-to-book (MB) ratio. Entries reports the time-series averages ("Average") of the coefficient and R^2 estimates, the Newey-West t-values (NW) on the average, and percentile values of the estimates.

we also perform a standardized multivariate cross-sectional regression at each date that includes all five explanatory variables. Table 5 reports the time-series averages of the slope coefficients and adjusted-R^2 estimates, the Newey and West (1987) t-statistics on the average, and percentile values of the estimates. The multivariate regression has an average adjusted R^2 estimate of 70%, two percentage points higher than that from the benchmark univariate regression on the mean-variance ratio forecast.

Under the multivariate setting, the relative company size contributes an average of 6% with strong statistical significance. The average contribution from asset tangibility remains small at -2%. The average contribution from the profitability measure is negative at -6% with strong statistical significance, while the average contribution from the market-to-book ratio is strongly positive at 15%. Above all, the contribution from the mean-variance ratio forecast remains magnitudes larger than contributions from all other variables. The slope estimates on the mean-variance ratio forecasts average at 86%, and vary within a narrow range from 78% at the 5th percentile to 91% at the 95th percentile. The Newey-West t-statistic on the average is also multiples larger than those on any other variables. With or without controlling for the other firm characteristics, the mean-variance ratio forecast remains the principal driver of the cross-sectional risk taking variation.

4.5. Mean-variance efficiency and firm value maximization

Our analysis shows that corporate risk taking decisions are broadly consistent with cross-sectional mean-variance efficiency under a single risk preference coefficient at each date. In this section, we establish a linkage between cross-sectional mean-variance efficiency and firm valuation maximization.

We start with a series of contemporaneous cross-sectional regression specifications to predict the cross-sectional variation of the mean-variance efficient risk taking target,

$$w_{t,i} = \alpha_t + \omega_t \mathrm{MVF}_{t,i} + c_t^k X_{t,i}^k + e_{t,i}^k. \tag{17}$$

We use the mean-variance ratio forecast MVF as the benchmark explanatory variable, and generate a benchmark prediction by itself. We then add each of the other explanatory variables X^k into the regression together with MVF to predict the risk allocation variation in a bivariate specification. We use the regression residual as a deviation measure from the predicted variation. If the predicted variation reflects the mean-variance efficient risk targets, the residual represents the deviation from cross-sectional mean-variance efficiency. We examine how company valuation varies with this deviation.

We measure the relative value of a company with a simple construct of the Tobin (1969)'s Q ratio, as the ratio of total book asset minus common book equity plus market capitalization to total book asset. We perform local linear cross-sectional regressions of the Q ratio against the residual from each specification k in (17) at each date t. The local linear regression uses a standard Gaussian kernel with three times the default bandwidth choice for a smooth estimated relation. To make the different specifications comparable to each other, we cross-sectionally standardize the regression residuals and the Tobin's Q ratio, and estimate the relation on the standardized series. Figure 3 plots the time-series averages of the estimated relations for each specification. The solid line denotes the benchmark univariate specification that predicts the risk allocation with the mean-variance ratio forecast. The other lines represent the other bivariate specifications.

The solid line from the benchmark prediction shows a clear hump-shaped average relation, where the

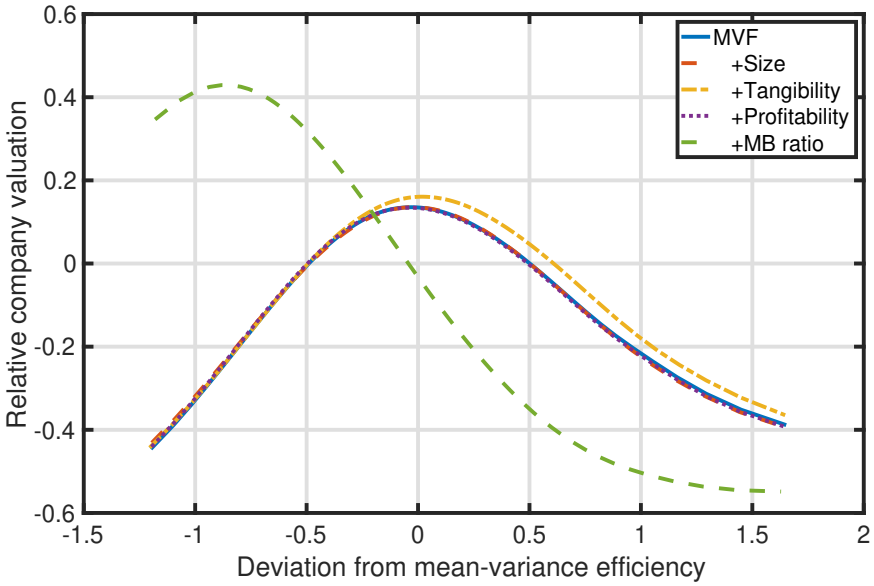


Figure 3

Deviations from mean-variance efficiency and relative company valuation

Each line denotes the time-series average of the local linear relation between Tobin's Q ratio and deviations from the risk allocations predicted by a particular specification. The solid line denotes the benchmark univariate specification that predicts the risk allocation with the mean-variance ratio forecast. The other lines denote bivariate specifications that include the mean-variance ratio forecast and one other variable.

hump peaks around zero deviation from the predicted risk allocation. The relation shows that cross-sectional mean-variance efficiency in corporate risk allocation maximizes firm value.

The two dotted lines represent specifications that include the relative company size and the probability measure, respectively. They are indistinguishable from the solid line, suggesting that these two variables add very little to the firm value maximization. The specification with asset tangibility (the dash-dotted line) shows a slightly higher peak around zero deviation. All these lines tell the same story that cross-sectional mean-variance efficiency in corporate risk allocation maximizes firm value.

The dashed line denotes the specification that includes the market-to-book ratio as an additional explanatory variable. The estimated relation deviates strongly from the benchmark hump shape. With the market-to-book ratio in the specification, the firm value is no longer maximized at zero deviation from the specification. The distortion highlights the complex nature of the interaction between market valuation and the risk taking decision. On the one hand, managers can try to time the market based on the relative valuation of its equity and debt, or try to reformulate its expectation to confirm with market valuation. On the other hand, the risk allocation the management takes also has a direct impact on market valuation of the company. In a sense, the market-to-book ratio is less a determinant of the mean-variance efficient risk allocation, but more a reflection of the mean-variance efficiency of the risk allocation decision.

To quantify the strength of each specification in generating value-maximizing risk allocation, at each date, we perform a series of standardized cross-sectional regressions of the Tobin's Q ratio, $Q_{t,i}$, against the negative of the absolute deviation from each specification, $\left|e^k_{t,i}\right|$,

$$Q_{t,i} = -\phi^k_{t,h}\left|e^k_{t,i}\right| + \varepsilon^k_{t+h,i}\,. \tag{18}$$

We cross-sectional standardize both the Q ratio and the absolute deviation so that the coefficient estimate $\phi_{t,h}$ measures the negative of the cross-sectional correlation between the absolute deviation from the predicted risk allocation target and the relative company valuation. It represents a standardized measure on how fast

Table 6
Link company relative valuation to deviations from mean-variance efficient risk allocation

	Average slope	NW slope	NW outperformance
Benchmark	Explain risk allocation variation with MVF		
MVF	0.232	26.05	
X^k	Explain risk allocation variation with MVF + X^k		
Size	0.231	28.65	-0.37
Tangibility	0.234	24.84	1.00
Profitability	0.232	24.11	-0.36
MB ratio	0.112	11.33	-15.32

At each date, we perform a series of standardized cross-sectional regressions of company relative value against the negative of the absolute deviation of the risk allocation from its predicted level. We predict the risk allocation first with the mean-variance ratio forecast (MVF) alone as a benchmark, and then with a list of other variables together with the MVF in a bivariate regression. Entries report the time-series averages of the slope coefficient estimate and the Newey-West (NW) t-statistics on the average. The last column reports the Newey-West (NW) t-statistics on the average slope coefficient difference between each bivariate specification and the univariate benchmark.

firm value declines with increasing deviation from cross-sectional mean-variance efficiency as predicted by a particular specification. The better the specification is in identifying the value-maximizing risk allocation, the stronger the firm value declines with the deviation, and the stronger will be the coefficient estimate.

Table 6 reports the time-series averages of the coefficient estimates and the Newey and West (1987) t-statistics on the average for each specification. The last column of the table reports the Newey-West t-statistics on the average slope coefficient difference between each bivariate specification and the benchmark univariate specification. A positive average slope difference indicates that the bivariate specification outperforms the benchmark in identifying the value-maximizing risk allocation.

The average slope estimate on deviations from the benchmark specification is positively and strongly statistically significant. Company relative value declines strongly with increasing deviation from the mean-variance ratio prediction of the risk allocation. Adding company relative size, asset tangibility, or profitability to the specification does not generate significant performance differences. The t-statistics on the average slope difference are weakly negative for size and profitability, and weakly positive for asset tangibility, but

none with statistical significance. Adding the market-to-book ratio generates significantly lower average coefficient estimate, suggesting that the market-to-book ratio does not help identify the value-maximizing risk allocation, even though it helps explain some of the actual risk allocation variation.

5. Time-series variation of corporate risk taking decisions

The previous section shows that cross-sectional mean-variance efficiency can explain a large proportion of the corporate risk taking variation at the firm level and can almost completely explain the variation at the industry level at a given date. This section examines the time-series variation of the mean-variance efficient risk allocation and the dynamic rebalancing behavior when a firm's risk allocation deviates from mean-variance efficiency.

5.1. Time-series variation of mean-variance efficient risk allocation

The existence of a common proportionality coefficient determines the cross-sectional mean-variance efficiency at a given point in time. How the proportionality coefficient varies over time dictates the unconditional long-run performance of the investment. The importance of this time-series dimension, or the market-timing decision, has long been recognized since Hansen and Richard (1987), who explain that a portfolio that is conditionally mean-variance efficient period by period is not necessarily unconditionally efficient.

In the securities market, the literature has examined the unconditional performance of various market-timing strategies, e.g., constant leverage target (Ackermann, Pohl, and Schmedders (2016) and Maurer, To, and Tran (2018)), constant volatility target (Moreira and Muir (2017) and Penãranda and Wu (2022)), constant mean return target (Penãranda and Wu (2022)), and strategies that maximize some unconditional performance metric (Ferson and Siegel (2001) and Peñaranda (2016)). Penãranda and Wu (2022) show that constant risk target tends to generate better unconditional performance than constant mean return target under most market conditions. Ingersoll, Spiegel, Goetzmann, and Welch (2007) discuss the potential gaming

of the unconditional Sharpe ratio at the expense of increasing downside risk.

For corporate risk taking decisions, the time-series variation can be different from that in the securities market due to the high adjustment cost. We analyze the time-series variation of the proportionality coefficient estimates ω_t and examine which market timing strategy best matches the time variation of the corporate risk taking behavior. In particular, by targeting a constant market average level on different measures, we can derive the associated market-timing strategy in terms of the variation of the proportionality coefficient,

$$\begin{aligned} \omega_t &\sim \frac{1}{\mathrm{MVF}_{t,m}} && \text{Constant average leverage} \\ &\sim \frac{1}{\mathrm{IRF}_{t,m}} && \text{Constant average levered volatility} \\ &\sim \frac{1}{\mathrm{IRF}^2_{t,m}} && \text{Constant average levered mean return} \end{aligned} \tag{19}$$

where $\mathrm{MVF}_{t,m}$, $\mathrm{IRF}_{t,m}$, and $\mathrm{IRF}^2_{t,m}$ represent the corresponding market averages of the mean-variance ratio, the information ratio, and the squared information ratio forecasts. See the Appendix for the derivation.

We use common equity as the weight to construct the market weighted averages of the mean-variance ratio forecast, the information ratio forecast, and the squared information forecast. Based on these market average estimates, we construct the time series that proxy the three market-timing strategies that target constant average leverage, constant average levered volatility, and constant average levered mean return, respectively, all with common equity as the weight. We measure the time-series correlation of the different market-timing strategies with the proportionality coefficient ω_t estimates.

Table 7 reports in panel A the time-series correlation estimates. The three columns correspond to the three sets of proportionality coefficients estimated from the cross-sectional regressions at the firm level, the 4-digit GICS industry level, and the 2-digit GICS sector level, respectively. The time-series variations of the proportionality coefficients show strongly positive correlations with the constant leverage market timing strategy, but negative correlations with the market-timing strategies that target either constant volatility or constant mean return. Targeting stable financial leverage makes practical sense for corporate risk taking as corporate leverage adjustment tends to incur much higher cost than that in the securities market.

Table 7
Correlations of the proportionality coefficients with different market-timing strategies

Market-timing strategies	Proportionality coefficient ω estimates at					
	Firm level		Industry level		Sector level	
	A. Time series correlation					
Constant leverage	0.878		0.864		0.875	
Constant volatility	-0.646		-0.652		-0.627	
Constant mean	-0.580		-0.585		-0.553	
1/Credit spread	0.346		0.369		0.371	
	B. Bivariate regression					
Intercept	-3.945	(-16.35)	-5.053	(-17.37)	-5.051	(-18.35)
Constant leverage	4.185	(42.08)	4.769	(39.78)	4.819	(42.47)
1/Credit spread	2.869	(10.28)	3.774	(11.22)	3.767	(11.83)
R^2	0.813		0.799		0.819	

We construct different market timing strategies. Panel A reports the time-series correlation between the proportionality coefficient estimates ω_t and these strategies. The three columns correspond to the three sets of proportionality coefficients estimated from cross-sectional regressions at the firm level, the 4-digit GICS industry level, and the 2-digit GICS sector level, respectively. Panel B reports the results from bivariate regressions of the proportionality coefficients against the constant leverage strategy and the reciprocal of the credit spread. In parentheses are the t-values of the corresponding estimates.

Our mean-variance setup in (1) does not target any risk or return level, but starts with a risk-return trade-off coefficient λ_t capturing the market risk preference. With a linear approximation of the financing cost function, the explicit solution has a proportionality coefficient determined by the reciprocal of the summation of the risk preference coefficient λ_t and the proportional financing cost coefficient per unit risk η_t. Thus, according to the theoretical derivation, the mean-variance efficient risk allocation can vary over time due fluctuations in investor risk preference and the market's financing condition. We construct a credit spread between Moody's seasoned BAA corporate bond yield and the 10-year constant maturity Treasury yield, both obtained from the Federal Reserve Bank of St. Louis. As the credit spread tends to increase when the corporate credit market becomes tight and when bond investors become more averse to credit risk, we can use the credit spread time series to proxy variations in the market's risk preference and the market's financing condition. In the last role of the panel, we use the reciprocal of the credit spread to proxy the proportionality coefficient variation in our structural setting and examine how it correlates with the time-series variation of the actual proportionality coefficient estimates. The correlation estimates are positive for all three sets of proportionality coefficients, suggesting that risk preference and market financing condition also play a role in the market timing of corporate risk allocation decisions.

In panel B, we perform a bivariate regression of the proportionality coefficients against the constant leverage market timing strategy ($1/\mathrm{MVF}_{t,m}$) and the reciprocal of the credit spread. Entries report the regression coefficient estimates and the t-values of the estimates (in parentheses). The regressions generate positive and strongly statistically significant slope coefficients on both variables. The regressions generate high R^2 estimates between 79.9% and 81.3%.

5.2. Dynamic risk rebalancing decisions

When a firm's risk allocation deviates from mean-variance efficiency, Proposition 4 shows that the company's future risk allocation dynamically adjusts toward the mean-variance efficient level. We propose a

cross-sectional error-correction model to capture the risk rebalancing behavior,

$$w_{t+1,i} - w_{t,i} = \delta_{t+1} + \kappa_{t+1}\left(\overline{w}_{t,i} - w_{t,i}\right) + \varepsilon_{t+1,i}, \tag{20}$$

where the slope coefficient κ_{t+1} captures the rebalancing speed over the next year per each unit deviation from the mean-variance efficient risk allocation $(\overline{w}_{t,i} - w_{t,i})$. We start by generating the mean-variance efficient allocation $\overline{w}_{t,i}$ with the univariate cross-sectional regression on the mean-variance ratio forecasts,

$$w_{t,i} = \alpha_t + \omega_t \mathrm{MVF}_{t,i} + e_{t,i}, \tag{21}$$

where the negative of the residual represents the deviation from cross-sectional mean-variance efficiency $-e_{t,i} = (\overline{w}_{t,i} - w_{t,i})$.

We estimate the cross-sectional error-correction model in (20) on the full universe of companies at each date. Table 8 reports in panel A the time-series statistics of the rebalancing speed estimates. The estimates have a time-series average of 4.39% and a median of 4.76% per year, suggesting that the risk rebalancing process is very slow. The estimates show high serial persistence with a monthly autocorrelation estimate of 0.93. Nevertheless, even after we adjust for the serial persistence, the Newey and West (1987) t-value for the sample average is still very high at 7.12, highlighting the strong statistical significance of the risk rebalancing behavior toward mean-variance efficiency.

Proposition 4 further predicts that the rebalancing speed declines with the proportional adjustment cost η_t but increases with the variance σ_t^2 of the investment. We do not have a measure quantifying the adjustment cost variation, but we have constructed historical variance estimators $\sigma_{t,i}^2$ in constructing the mean-variance ratio forecast. We examine how the adjustment speed varies with the variance estimator. At each date, we divide the company universe into five subsamples based on the quintile sorting of the variance estimators. We repeat the estimation of the error-correlation model for each subsample. Panel B of Table 8 reports the summary statistics on the risk rebalancing speed estimates for each variance quintile. In line with the

Table 8
Risk rebalancing speeds

Statistics	A. Full-sample estimates	B. Estimates for each variance quintile				
		1	2	3	4	5
Mean	4.39	3.82	3.73	5.32	6.68	7.82
Median	4.76	4.19	4.25	5.96	7.28	8.03
SD	3.88	4.29	5.17	5.83	6.05	4.92
Auto	0.93	0.93	0.92	0.89	0.89	0.87
NW	7.12	5.94	4.54	6.56	6.64	10.87

At each date, we estimate cross-sectional error-correction model to obtain the risk rebalancing speed. Panel A reports the statistics on the full-sample estimates. Panel B divides the sample into five subsamples based on the quintile sorting of asset return variance estimators and repeats the estimation for each sub-sample. The statistics include the time-series mean, median, standard deviation (SD), monthly autocorrelation (Auto), and the Newey-West t-statistics (NW) on the average.

prediction of the Proposition, the average and median rebalancing speed estimates increase with the variance quintile. The time series of the estimates become somewhat less persistent at the higher variance quintile. The Newey-West t-value on the sample average is the highest for the fifth variance quintile.

In estimating the rebalancing speeds in Table 8, we use the mean-variance ratio forecast alone to generate the mean-variance efficient risk allocation prediction through the benchmark specification in (21). When we consider different specifications to explain the cross-sectional risk allocation variation, the error-correction model can also be used as an out-of-sample test on whether a specification explains the mean-variance efficient risk allocation or variations away from mean-variance efficiency. As future risk rebalancing tends to move toward mean-variance efficiency, the closer a specification is to capture the mean-variance efficient risk allocation, the higher should be the error-correction speed toward its risk allocation prediction.

In addition to the benchmark univariate specification, we consider a series of bivariate specifications to explain the cross-sectional risk allocation variation as in (17),

$$w_{t,i} = \alpha_t + \omega_t \mathrm{MVF}_{t,i} + c_t^k X_{t,i}^k + e_{t,i}^k. \tag{22}$$

For each specification k, we use the following cross-sectionally standardized error-correction regression as an out-of-sample test of its predictability on future risk allocation adjustments,

$$w_{t+h,i} - w_{t,i} = -\rho^k_{t,h} e^k_{t,i} + \varepsilon^k_{t+h,i}, \tag{23}$$

where we predict future risk allocation changes $(w_{t+h,i} - w_{t,i})$ over future horizon h with the regression residual $e^k_{t,i}$ from the kth specification. We cross-sectionally standardize both the risk allocation change $(w_{t+h,i} - w_{t,i})$ and the residual $e^k_{t,i}$ so that the slope coefficient from the standardized regression $\rho^k_{t,h}$ captures the forecasting correlation between the current risk deviation and the future allocation change.

In this two-step procedure, the R^2 from the contemporaneous regression in (22) captures how much the specification can explain the cross-sectional variation of the observed risk allocation, whereas the slope estimate ρ from the standardized error-correction regression in (23), or equivalent the forecasting R^2 of the regression, captures whether the specification in (22) explains the mean-variance efficient risk allocation that future allocation rebalances to. A higher explanatory R^2 from the first regression does not necessarily lead to a higher error-correction forecasting correlation ρ in the second regression if the explained variation does not represent the mean-variance efficient risk allocation.

For each specification, we estimate the error-correction model on risk allocation changes over the next $h = 1$ to 5 years. Table 9 reports in Panel A the time-series averages of the forecasting correlation estimates across different forecasting horizons. Panel B reports the Newey and West (1987) t-statistics on the average differences in the forecasting correlation estimates between each bivariate specification and the univariate benchmark. A positive t-statistic on a bivariate specification suggests that this specification on average generates higher forecasting correlations with future risk allocation adjustments and accordingly outperforms the univariate benchmark specification in explaining the mean-variance efficient risk allocation.

The first row of statistics show that deviations from the benchmark specification can strongly predict future risk allocation changes from one to five years. The average forecasting correlation increases with the

Table 9
Error-correction forecasting correlation on future risk allocation changes

	A. Average forecasting correlations					B. NW t-statistics on outperformance				
h	1	2	3	4	5	1	2	3	4	5
Benchmark	Explain risk allocation variation with MVF									
MVF	0.152	0.239	0.291	0.319	0.345					
X_k	Explain risk allocation variation with MVF + X_k									
Size	0.154	0.243	0.296	0.325	0.352	1.71	2.39	2.62	2.88	3.15
Tangibility	0.158	0.244	0.296	0.323	0.347	2.09	1.57	1.34	1.00	0.29
Profitability	0.151	0.239	0.291	0.319	0.345	-0.67	-1.21	-2.15	-2.39	-2.19
MB ratio	0.141	0.228	0.280	0.309	0.336	-4.19	-4.70	-5.54	-5.47	-4.92

At each date, we perform a series of standardized cross-sectional forecasting regressions of future risk allocation changes over horizons from one to five years against the current deviations from the predicted levels. We predict the risk allocation level first with the mean-variance ratio forecast (MVF) alone as a benchmark, and then with a list of other variables together with the MVF in a bivariate regression. Panel A reports the time-series averages of the cross-sectional forecasting correlation estimates obtained from the error-correction forecasting regressions. Panel B reports the Newey-West (NW) t-statistics on the average forecasting correlation difference between each bivariate specification and the univariate benchmark.

forecasting horizon from 15.2% at one-year horizon to 34.5% at the five-year horizon.

When we add the relative company size measure to the specification in a bivariate regression, the size variable contributes positively not only to the cross-sectional variation of the risk allocation, but also helps enhance the forecasting relation with future risk allocation changes. The average forecasting correlation estimates are higher across all horizons. Although the average increases in the forecasting correlation estimates are small, the Newey-West t-statistics in panel B suggest that they are statistically significant, especially over longer horizons.

The asset tangibility ratio also contributes to higher error-correction forecasting correlation estimates. The enhancement is statistically significant at short forecasting horizons, but becomes less significant at longer horizons.

The contribution of profitability to the risk allocation variation is very small. The error-correction estimation shows that this small contribution is not toward a better risk allocation prediction. The forecasting

correlation estimates are lower than that from the benchmark specification. The Newey-West t-statistics on the forecasting correlation differences are negative across all forecasting horizons, and become more statistically significant at longer horizons. This negative contribution to the forecasting power supports the literature explanation of the wealth cumulation effect, which can move the risk allocation away from mean-variance efficiency in between rebalancing actions.

Adding the market-to-book ratio to the bivariate specification generates the largest R^2 increase in explaining the cross-sectional risk allocation variation; yet, deviations from this specification generates the lowest forecasting correlation on future risk allocation changes. The Newey-West t-statistics on the forecasting correlation differences from the benchmark specification are strongly negative across all forecasting horizons. Therefore, although the market-to-book ratio can explain the market timing actions and its effects on the risk taking variation, these explained variations are mostly away from the mean-variance efficient allocation. The variable's significantly negative contribution to the forecasting correlation matches its significantly negative contribution to firm value maximization. Both suggest that the market-to-book ratio is not a determinant of the mean-variance efficient risk allocation, even though it does affect a firm's actual risk allocation decision.

6. Concluding remarks

Mean-variance analysis has been the cornerstone of modern investment theory. This paper examines the mean-variance efficiency in US corporate risk taking decisions. We introduce the concept of cross-sectional mean-variance consistency in corporate risk taking decisions, characterized by the condition that the cross-sectional corporate risk taking variations can be explained by the mean-variance ratio forecasts with a common risk preference coefficient at a given point in time. We find that cross-sectional mean-variance efficiency can explain a large proportion of the corporate risk taking variation at the firm level and can almost completely explain the corporate risk taking variation at the industry level. Over time, due to the high adjustment cost, corporations tend to market time the mean-variance efficient allocation to main a stable financial

leverage. When the risk allocation of an individual firm deviates from mean-variance efficiency, its future risk allocation tends to rebalance toward to the mean-variance efficient level, but at a slow speed.

The successful application of mean-variance analysis to corporate risk taking decisions opens ample ground for future research. First, to accurately characterize the investment behavior, one must have an accurate assessment on the firm's risk-return prospect. As a first effort, we generate a company's mean-variance ratio forecast based on its historical return on asset behaviors. Future research can explore different approaches, potentially including analysts forecasts, to enhance the accuracy of the risk-return forecasts. Second, we assume that managers are fully aligned with equity investors for both incentives and the information set in making the investment decision. In practice, the decision-making managers can be overly optimistic or pessimistic about the future investment prospects and can accordingly choose a higher or lower risk allocation (Heaton (2002), Hackbarth (2009), and Ben-David, Graham, and Harvey (2013)). The managers may also have different incentives from the equity investors (Jensen and Meckling (1976)) and decide to over- or under-take risks based on how well they think they can reach their performance targets. Future research can use cross-sectional mean-variance efficiency as a benchmark and examine the sources of deviations from mean-variance efficiency in terms of information sets and incentives. Third, corporate hedging decisions can be examined within the same mean-variance framework. To maximize its risk-return tradeoff, a firm should focus on its domain expertise, lever up its comparative advantage, while hedging away exposures outside of its expertise. What this means is that not all firms should hedge the same risk. Whether a source of risk should be hedged depends crucially on whether the risk source is part of the firm's domain expertise and whether the hedging decision can enhance the firm's risk-return tradeoff.

Appendix. Market timing on mean-variance efficient risk allocations

We start with cross-sectionally mean-variance efficient risk allocations at a given point in time t across a universe of N_t companies, characterized by the following risk allocation relation,

$$\overline{w}_{t,i} = \omega_t \mathrm{MVF}_{t,i} \tag{24}$$

where the cross-sectional risk allocation variation is explained by the corresponding mean-variance forecast variation with a single proportionality coefficient ω_t. A market-timing strategy determines the time-series variation of the proportionality coefficient ω_t to achieve a certain unconditional objective.

A.1. Targeting a constant average leverage level

With $\overline{w}_{t,i}$ denoting the leverage level of a particular company at time t, we need to aggregate the leverage level across all companies to generate a market average leverage level that we can target on. A common aggregation is to define the market leverage as the ratio of the sum of total assets to the sum of common book equity across the universe. This definition of market leverage can be represented as the weighted average leverage of all companies, with common equity as the weighting,

$$\overline{w}_{t,m} = \frac{\sum_i \mathrm{TA}_{t,i}}{\sum_i \mathrm{CE}_{t,i}} = \frac{\sum_i \mathrm{CE}_t \overline{w}_{t,i}}{\sum_i \mathrm{CE}_{t,i}}, \tag{25}$$

where the sum is over the universe of companies at time t. We can apply the same weighed average to mean-variance efficient risk allocation relation in (24) to obtain the relation at the market level,

$$\overline{w}_{t,m} = \omega_t \mathrm{MVF}_{t,m}\,, \tag{26}$$

where the market mean-variance ratio $\mathrm{MVF}_{t,m}$ is constructed with the same common equity weighting,

$$\mathrm{MVF}_{t,m} = \frac{\sum_i \mathrm{CE}_t \mathrm{MVF}_{t,i}}{\sum_i \mathrm{CE}_{t,i}}. \tag{27}$$

To target a constant market average leverage level, $\overline{w}_{t,m} = c$, we can plug in the constant target to equation (26) to obtain the corresponding market-timing strategy as,

$$\omega_t = c \frac{1}{\mathrm{MVF}_{t,m}}. \tag{28}$$

If one intends to target a constant market leverage level with a different cross-company weighting scheme (say equal weighting), the market-timing strategy remains in the same form as in (28), but with the market mean-variance ratio $\mathrm{MVF}_{t,m}$ constructed in the same corresponding weighting scheme. This principle also applies to the other constant targeting strategies in the next two subsections.

A.2. Targeting a constant average volatility level

Given the allocation weight $\overline{w}_{t,i}$ for a company, the levered equity return volatility is given by $\sigma^E_{t,i} = \overline{w}_{t,i}\sigma_t$. Plugging in the cross-sectional mean-variance efficiency equation in (24), we have

$$\sigma^E_{t,i} = \overline{w}_{t,i}\sigma_t = \omega_t \mathrm{MVF}_{t,i}\sigma_t = \omega_t \mathrm{IRF}_{t,i}. \tag{29}$$

That is, the mean-variance efficient levered volatility variation is proportional to the variation of the information ratio forecast with the same proportionality coefficient ω_t.

To target a constant average equity return volatility, we can apply weighted average directly to (29) to obtain,

$$\sigma^E_{t,m} = \omega_t \mathrm{IRF}_{t,m}, \tag{30}$$

where we apply the same weighting to the equity return volatility and the information ratio forecast. A constant market average equity volatility target $\sigma^E_{t,m} = c$ leads to the following market-timing strategy

$$\omega_t = c\frac{1}{\text{IRF}_{t,m}}. \tag{31}$$

A.3. Targeting a constant average mean equity return level

Given the allocation weight $\overline{w}_{t,i}$ for a company, the levered mean return on equity is given by $\mu^E_{t,i} = \overline{w}_{t,i}\mu_t$. Plugging in the cross-sectional mean-variance efficiency equation in (24), we have

$$\mu^E_{t,i} = \overline{w}_{t,i}\mu_t = \omega_t \text{MVF}_{t,i}\mu_t = \omega_t \text{IRF}^2_{t,i}. \tag{32}$$

That is, the mean-variance efficient expected equity mean return variation is proportional to the variation of the squared information ratio forecast with the same proportionality coefficient ω_t.

Applying weighted average to (32) to obtain the relation at the market average level, we have

$$\mu^E_{t,m} = \omega_t \text{IRF}^2_{t,m}, \tag{33}$$

where we apply the same weighting to the equity mean return and the squared information ratio forecast. A constant market average mean equity return target $\mu^E_{t,m} = c$ implies the following market-timing strategy,

$$\omega_t = c\frac{1}{\text{IRF}^2_{t,m}}. \tag{34}$$

References

Abel, A. B. 2018. Optimal Debt and Profitability in the Trade-Off Theory. *Journal of Finance* 73(1):95–143.

Ackermann, F., W. Pohl, and K. Schmedders. 2016. Optimal and Naive Diversification in Currency Markets. *Management Science* 63(10):3347–3360.

Alchian, A. A, and H. Demsetz. 1972. Production, Information Costs, and Economic Organization. *American Economic Review* 62(5):777–795.

Amengual, D, and E. Sentana. 2010. A Comparison of Mean-variance Efficiency Tests. *Journal of Econometrics* 154(1):16–34.

Baker, M, and J. Wurgler. 2002. Market Timing and Capital Structure. *Journal of Finance* 57(1):1–32.

Ben-David, I., J. R. Graham, and C. Harvey. 2013. Managerial Miscalibration. *Quarterly Journal of Economics* 128(4):1547–1584.

Benjamin Klein, R. G. C, and A. A. Alchian. 1978. Vertical Integration, Appropriable Rents, and the Competitive Contracting Process. *Journal of Law and Economics* 21(2):297–326.

Chen, Z., J. Harford, and A. Kamara. 2018. Operating Leverage, Profitability, and Capital Structure. *Journal of Financial and Quantitative Analysis* 54(1):369–392.

Coase, R. H. 1937. The Nature of the Firm. *Economica* 4(16):396–405.

Danis, A., D. A. Rettl, and T. M. Whited. 2014. Refinancing, Profitability, and Capital Structure. *Journal of Financial Economics* 114(3):424–443.

DeAngelo, H, and R. Roll. 2015. How Stable Are Corporate Capital Structures?. *Journal of Finance* 70(1):373–418.

Dong, M., D. Hirshleifer, and S. Teoh. 2012. Overvalued Equity and Financing Decisions. *Review of Financial Studies* 25(12):3645–3683.

Dong, M., I. Loncarski., J. Horst, and C. Veld. 2012. What Drives Security Issuance Decisions: Market Timing, Pecking Order, or Both?. *Financial Management* 41(3):637–663.

Engel, C. M, and A. P. Rodrigues. 1993. Tests of Mean-Variance Efficiency of International Equity Markets. *Oxford Economic Papers* 45(3):403–421.

Fama, E. F. 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance* 11:404–419.

Fama, E. F, and K. R. French. 2002. The Equity Premium. *Journal of Finance* 57(2):637–659.

Faulkender, M, and M. A. Petersen. 2005. Does the Source of Capital Affect Capital Structure?. *Review of Financial Studies* 19(1):45–79.

Ferson, W. E, and A. F. Siegel. 2001. The Efficient Use of Conditioning Information in Portfolios. *Journal of Finance* 56(3):967–982.

Fischer, E. O., R. Hernkel, and J. Zechner. 1989. Dynamic Capital Structure Choice: Theory and Tests. *Journal of Finance* 44(1):19–40.

Fischer, S, and R. C. Merton. 1984. Macroeconomics and Finance: The Role of the Stock Market. *Carnegie-Rochester Conference Series on Public Policy* 21(1):57–108.

Flannery, M. J, and K. P. Rangan. 2006. Partial Adjustment toward Target Capital Structures. *Journal of Financial Economics* 79(3):469–506.

Frank, M. Z, and V. K. Goyal. 2003. Testing the Pecking Order Theory of Capital Structure. *Journal of Financial Economics* 67(2):217–248.

Frank, M. Z, and V. K. Goyal. 2009. Capital Structure Decisions: Which Factors Are Reliably Important?. *Financial Management* 38(1):1–37.

Goldstein, R., N. Ju, and H. Leland. 2001. An EBIT-Based Model of Dynamic Capital Structure. *Journal of Business* 74(4):483–512.

Gornall, W, and I. A. Strebulaev. 2018. Financing as a Supply Chain: The Capital Structure of Banks and Borrowers. *Journal of Financial Economics* 129(3):510–530.

Graham, J. R, and M. T. Leary. 2011. A Review of Empirical Capital Structure Research and Directions for the Future. *Annual Review of Financial Economics* 3(1):309–345.

Green, R. C, and B. Hollifield. 1992. When Will Mean-Variance Efficient Portfolios Be Well Diversified?. *Journal of Finance* 47(5):1785–1809.

Hackbarth, D. 2009. Determinants of Corporate Borrowing: A Behavioral Perspective. *Journal of Corporate Finance* 15(4):389–411.

Hansen, L. P, and S. F. Richard. 1987. The Role Of Conditioning Information in Deducing Testable Restrictions Implied By Dynamic Asset Pricing Models. *Econometrica* 55(3):587–614.

Hart, O. 1995. *Firms, Contracts, and Financial Structure*. Clarendon Press Oxford.

Hart, O, and J. Moore. 1998. Default and Renegotiation: A Dynamic Model of Debt. *Quarterly Journal of Economics* 113(1):1–41.

Heaton, J. B. 2002. Managerial Optimism and Corporate Finance. *Financial Management* 31(2):33–45.

Hovakimian, A., T. Opler, and S. Titman. 2001. The Debt-Equity Choice. *Journal of Financial and Quantitative Analysis* 36(1):1–24.

Hu, X., M. Sy, and L. Wu. 2020. A Factor Model of Company Valuation. Working paper, RMIT and Baruch College.

Im, H. J., Y. Kang, and J. Shon. 2020. How Does Uncertainty Influence Target Capital Structure. *Journal of Corporate Finance* 64:1–24.

Ingersoll, J., M. Spiegel., W. Goetzmann, and I. Welch. 2007. Portfolio Performance Manipulation and Manipulation-proof Performance Measures. *Review of Financial Studies* 20(5):1503–1546.

Jensen, M. C. 1986. Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers. *American Economic Review* 76(2):323–329.

Jensen, M. C, and W. H. Meckling. 1976. Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure. *Journal of Financial Economics* 3(4):305–360.

Klasa, S., H. Ortiz-Molina., M. Serfling, and S. Srinivasan. 2017. Protection of Trade Secretes and Capital Structure Decisions. *Journal of Financial Economics* 128(2):266–286.

Kraus, A, and R. H. Litzenberger. 1973. A State-Preference Model of Optimal Financial Leverage. *Journal of Finance* 28(4):911–922.

Lambrecht, B. M, and S. C. Myers. 2017. The Dynamics of Investment, Payout and Debt. *Review of Financial Studies* 30(11):3759–3800.

Lambrinoudakis, C., G. Skiadopoulos, and K. Gkionis. 2019. Capital Structure and Financial Flexibility: Expectations of Future Shocks. *Journal of Banking and Finance* 104(1):1–18.

Leary, M. T, and M. R. Roberts. 2005. Do Firms Rebalance Their Capital Structures?. *Journal of Finance* 60(6):2575–2619.

Leary, M. T, and M. R. Roberts. 2014. Do Peer Firms Affect Corporate Financial Policy?. *Journal of Finance* 69(1):139–178.

Leland, H. E. 1994. Risky Debt, Bond Covenants and Optimal Capital Structure. *Journal of Finance* 49(4):1213–1252.

Lemmon, M. L., M. R. Roberts, and J. F. Zender. 2008. Back to the Beginning: Persistence and the Cross-Section of Corporate Capital Structure. *Journal of Finance* 63(4):1575–1608.

Levy, M, and R. Roll. 2010. The Market Portfolio May Be Mean/Variance Efficient After All: The Market Portfolio. *Review of Financial Studies* 23(16):2462–2491.

Lucas, D. J, and R. L. McDonald. 1990. Equity Issues and Stock Price Dynamics. *Journal of Finance* 45(4):1019–1043.

MacKinlay, A. C, and M. P. Richardson. 1991. Using Generalized Method of Moments to Test Mean-Variance Efficiency. *Journal of Finance* 46(2):511–527.

Markowitz, H. M. 1952. Portfolio Selection. *Journal of Finance* 7(1):77–91.

Maurer, T. A., T.-D. To, and N.-K. Tran. 2018. Optimal Factor Strategy in FX Markets. Working paper, Washington University in St. Louis and University of New South Wales and Virginia Tech.

Merton, R. C. 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory* 3(4):373–413.

Modigliani, F, and M. H. Miller. 1958. The Cost of Capital, Corporation Finance and the Theory of Investment. *American Economic Review* 48(3):261–297.

Moreira, A, and T. Muir. 2017. Volatility-Managed Portfolios. *Journal of Finance* 72(4):1611–1644.

Myers, S. C. 1984. The Capital Structure Puzzle. *Journal of Finance* 39(3):574–592.

Myers, S. C, and N. S. Majluf. 1984. Corporate Financing and Investment Decisions When Firms Have Information That Investors Do Not Have. *Journal of Financial Economics* 13(2):187–221.

Newey, W. K, and K. D. West. 1987. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55(3):703–708.

Penãranda, F, and L. Wu. 2022. Targets, Predictability, and Performance. *Management Science* 68(2):1537–1555.

Peñaranda, F. 2016. Understanding Portfolio Efficiency with Conditioning Information. *Journal of Financial and Quantitative Analysis* 51(3):985–1011.

Polk, C, and P. Sapienza. 2009. The Stock Market and Corporate Investment: A Test of Catering Theory. *Review of Financial Studies* 22(1):187–217.

Qian, E. 2011. Risk Parity and Diversification. *Journal of Investing* 20(1):119–127.

Rajan, R. G, and L. Zingales. 1995. What Do We Know about Capital Structure? Some Evidence from International Data. *Journal of Finance* 50(5):1421–1460.

Scott, J. H. 1976. A Theory of Optimal Capital Structure. *Bell Journal of Economics* 7(1):33–54.

Sentana, E. 2009. The Econometrics of Mean-variance Efficiency Tests: A Survey. *Econometrics Journal* 12(3):C65–C101.

Strebulaev, I. A. 2007. Do Tests of Capital Structure Theory Mean What They Say?. *Journal of Finance* 62(4):1747–1787.

Titman, S, and S. Tsyplakov. 2007. A Dynamic Model of Optimal Capital Structure. *Review of Finance* 11(3):401–451.

Titman, S, and R. Wessels. 1988. The Determinants of Capital Structure Choice. *Journal of Finance* 43(1):1–19.

Tobin, J. 1969. A General Equilibrium Approach To Monetary Theory. *Journal of Money, Credit and Banking,* 1(1):15–29.

Welch, I. 2012. A Critique of Recent Quantitative and Deep-Structure Modeling in Capital Structure Research and Beyond. *Critical Finance Review* 2(1):131–172.

Williamson, O. E. 1985. *The Economic Institutions of Capitalism: Firms, Markets, and Relational Contracting*. The Free Press New York.