

# Resolving the Zero-Beta Rate Puzzle

Rong Wang \*

May, 2026

## Abstract

This paper resolves a long-standing zero-beta rate puzzle—the empirical finding that estimated zero-beta rates remain persistently high across factor models. I show that this apparent robustness arises from pervasive model misspecification rather than reflecting a genuinely high risk-free rate. When a factor model fails to perfectly price assets, the zero-beta rate is no longer unique within the model itself, and the standard estimator—based on the minimum-variance zero-beta portfolio—converges toward the mean return of the global minimum-variance portfolio as model misspecification increases. To quantify this mechanism, I introduce a new investment-based measure of model misspecification: the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios. This measure captures the economic magnitude of pricing errors and links model misspecification to empirically observable investment opportunities. Studying a comprehensive set of classical and modern factor models, I find substantial misspecification, explaining why all models yield similarly elevated zero-beta rates. Simulation analyses confirm that realistic degrees of misspecification can fully reproduce the empirical puzzle even when the true risk-free rate is low.

---

\*Fuqua School of Business, Duke University; [rong.wang1996@duke.edu](mailto:rong.wang1996@duke.edu). I am deeply grateful for the invaluable guidance and continuous support from Anna Cieslak, Paymon Khorrami, Ravi Bansal, Nuno Clara, Campbell Harvey, and David Hsieh. I sincerely thank Manuel Adelino, Francesco Beraldi, Alon Brav, Doug Breeden, Simon Gervais, John Graham, Ben Hébert, Jinge Liu, Hao Pang, Manju Puri, Adriano Rampini, David Robinson, Yinan Su, Arthur Taburet, Joao Thereze, Vish Viswanathan, Melanie Wallskog, Ming Yang, and participants at the 2025 AFA Ph.D. Poster Session, 2026 WFA, Fuqua Brownbag Seminars, Fuqua Interdisciplinary Seminars for their helpful comments.

# 1. Introduction

The zero-beta rate, defined as the expected return of a portfolio orthogonal to the stochastic discount factor (SDF), has long occupied a central place in theoretical and empirical asset pricing. Since the seminal contributions of [Black \(1972\)](#) and [Black, Jensen, and Scholes \(1972\)](#), the notion of a zero-beta portfolio—one uncorrelated with the factor portfolio—has provided the foundation for asset pricing models that operate without a risk-free asset. In such settings, the zero-beta rate serves as a substitute for the unobservable risk-free rate, representing the expected return on a portfolio that bears no systematic risk.

Within the CAPM framework, empirical estimates of this rate have persistently appeared high, often far exceeding Treasury bill yields or other risk-free proxies. This pattern corresponds to a flat security market line with a large intercept and motivates the extensive literature on the beta anomaly, wherein low-beta stocks tend to earn higher risk-adjusted returns than high-beta stocks. A common interpretation attributes this discrepancy to the inadequacy of a single-factor model: the market factor alone cannot capture the cross-section of expected returns, and additional factors are required. Yet, [Di Tella, Hébert, Kurlat, and Wang \(2026\)](#) demonstrate that even when a rich set of cross-sectional risk factors is included, the estimated zero-beta rate remains well above the level of Treasury yields. I confirm this finding: across a wide range of prominent factor models—from traditional Fama–French models with pre-specified factors to more recent machine-learning models with latent factors—the estimated zero-beta rate remains strikingly high. I refer to this persistent pattern as the *zero-beta rate puzzle*, which continues to challenge our understanding of factor models and risk pricing mechanisms in financial markets.

The connection between zero-beta rate, equity risk premium, and convenience yield deepens the significance of the zero-beta rate puzzle. Safe asset yields, such as Treasury yields, are typically lower than the frictionless risk-free rate due to a “convenience yield” that reflects their non-pecuniary benefits in providing liquidity, collateral, hedging, and regulatory services ([Bansal and Coleman, 1996](#); [Krishnamurthy and Vissing-Jorgensen, 2012](#); [Nagel, 2016](#); [Acharya and Laarits, 2025](#); [Cieslak, Li, and Pflueger, 2025](#); etc.). If the true unobserved risk-free rate is indeed high, then the market risk premium must be low and the convenience yield must be high. Conversely, a low risk-free rate implies a high market risk premium and a small convenience yield. The empirical robustness of zero-beta rate estimates across models appears to suggest a high risk-free rate ([Di Tella et al., 2026](#)).

This paper challenges the interpretation of the robustness of zero-beta rate estimates. Rather than viewing the stability of zero-beta rate estimates across diverse models as evidence of economic validity, I argue that such robustness reflects a common statistical problem:

model misspecification. To demonstrate the relationship between model misspecification and zero-beta rate estimation, I proceed in four steps.

First, I discuss the conceptual relationship among factor model misspecification, the zero-beta rate, and the risk-free rate. To clarify terminology at the outset: the risk-free rate is the expected return on an asset that delivers a certain payoff in every future state of the world. It is a fundamental economic concept, not tied to any specific asset pricing model. The zero-beta rate, by contrast, is defined within a particular factor model or, equivalently, with respect to a given stochastic discount factor (SDF). It represents the expected return of any portfolio orthogonal to the factor space.

Under a correctly specified factor model—one that perfectly prices all risky assets—all zero-beta portfolios share the same expected return, and the zero-beta rate implied by that model is unique. However, this object need not coincide with the true risk-free rate. The reason lies in market incompleteness. When no risk-free asset is traded, there exist infinitely many admissible stochastic discount factors—and thus infinitely many factor models—that can all price the same set of risky returns. Each model implies its own internally consistent “risk-free rate”. This idea can also be understood through mean–variance analysis: any efficient portfolio on the mean–variance frontier (except the global minimum-variance portfolio) defines an SDF, or equivalently a factor model, that perfectly prices all risky assets. Each such model implies a unique, model-specific zero-beta rate, which need not coincide with the true risk-free rate. In short, the risk-free rate is fundamentally unidentified from risky returns alone (Hansen and Jagannathan, 1991; Cochrane, 2009). Hence, attempts to infer the true risk-free rate from factor models are inherently limited—even if the model appears to perfectly price all risky assets. This is the first layer of non-uniqueness, arising from multiple admissible SDFs in an incomplete market.

I next highlight a distinct and empirical source of non-uniqueness in the zero-beta rate. In practice, empirical factor models are typically misspecified and do not perfectly price the cross section of returns. When a model fails to perfectly price the cross section, pricing errors generate multiple zero-beta portfolios with distinct expected returns, making the zero-beta rate non-unique even within the model itself. This insight builds on Roll (1980), who studies orthogonal portfolios in the context of the CAPM. This is the second layer of non-uniqueness, arising from model misspecification.

Among the infinitely many possible zero-beta portfolios, the empirical literature conventionally selects a particular one—the unit-investment, minimum-variance zero-beta portfolio—and interprets its expected return as the zero-beta rate. This convention originates from Black (1972) and Black et al. (1972), where, under a perfectly specified factor model, the minimum-variance zero-beta portfolio lies on the mean–variance frontier. However, when

the model is misspecified, this choice loses theoretical justification and introduces systematic estimation bias. Specifically, when the factor model corresponds to an inefficient portfolio in mean–variance space, greater inefficiency—manifested as higher variance for a given mean or lower mean for a given variance—tends to push the estimated zero-beta rate upward. In the limit, severe misspecification drives the estimate toward the mean return of the global minimum-variance (GMV, hereafter) portfolio. Intuitively, I show that the zero-beta rate can be expressed as the mean return of the GMV portfolio, adjusted for the component attributable to factor exposures embedded in the GMV portfolio. When the model is severely misspecified, this adjustment term vanishes, and the zero-beta rate converges to the mean return of the GMV portfolio.

These analytical results reveal that the robustness of zero-beta rate estimates and their proximity to the mean return of the GMV portfolio likely arise because most factor models share similarly large degrees of misspecification. I use the term “model misspecification” broadly, without imposing assumptions on its source—whether omitted factors, weak factors, genuine mispricing, or pure measurement errors. Any deviation from perfect pricing constitutes misspecification. In mean–variance terms, this implies that the factor space fails to span or intersect the efficient frontier. Hence, throughout the analysis, I treat portfolio inefficiency and factor model misspecification as equivalent concepts.

Second, to quantify model misspecification empirically, I propose a general and economically grounded measure based on the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios. Under a correctly specified factor model, such portfolios should yield zero expected returns, resulting in a Sharpe ratio of zero. Any positive Sharpe ratio therefore reflects systematic pricing errors that can be potentially exploited, implying that the greater the Sharpe ratio of these portfolios, the greater the degree of model misspecification. This measure offers several advantages over traditional diagnostics.

First, the proposed measure does not require estimation of SDF parameters or factor risk premia, thereby avoiding well-known issues associated with misspecification and weak identification (Kan and Zhang, 1999b; Gospodinov et al., 2014; Kleibergen and Zhan, 2020; etc.). In contrast, the Hansen–Jagannathan distance (HJD) is inherently parameter-dependent, as it relies on estimating SDF parameters and computing the minimum distance between a proposed SDF and the set of admissible SDFs that correctly price test assets. Second, the approach provides an investment-based diagnostic in the spirit of GRS, but without requiring specification or estimation of a risk-free rate or zero-beta rate, which are central inputs in GRS-type tests (Shanken, 1986; Gibbons et al., 1989). Third, the underlying zero-investment, zero-beta portfolios exhibit modest turnover and limited short positions, making them economically more implementable than GRS-based maximum Sharpe ratio

portfolios, which often rely on extreme leverage and shorting (Fama and French, 2015). Finally, the framework is naturally suited to out-of-sample evaluation and the incorporation of transaction costs, making it a practical and flexible tool for assessing economic model misspecification in realistic settings.

Empirically, I examine a broad range of factor models—Fama–French (FF), principal component analysis (PCA), instrumented PCA (IPCA; Kelly, Pruitt, and Su, 2019), and conditional autoencoder (AE; Gu, Kelly, and Xiu, 2021) models. Across all specifications, the evidence reveals substantial misspecification: the maximum Sharpe ratios of zero-investment, zero-beta portfolios are economically large, exceeding 3.0 in-sample and 1.0 out-of-sample on an annualized basis. This implies that no prominent factor model is close to achieving mean–variance efficiency required for a unique zero-beta rate within the model. Moreover, the similar Sharpe ratios across models suggest comparable degrees of misspecification—helping explain why all models tend to yield similar, biased zero-beta rates close to the GMV mean return.

Third, to demonstrate how estimated zero-beta rates align with the mean return of the GMV portfolio, I examine two distinct universes of test assets that differ in their GMV mean returns. I find that the estimated zero-beta rates are systematically higher in the asset universe with a higher GMV mean return. Moreover, these estimates remain remarkably stable across different model specifications and lie close to the GMV mean returns within their respective asset groups. This pattern indicates that the estimated zero-beta rates primarily reflect the mean return of the GMV portfolio rather than the true risk-free rate, providing empirical support for my analytical result that substantial model misspecification biases zero-beta rate estimates toward the GMV mean return.

Fourth, I use simulation analysis to quantify the magnitude of this bias and demonstrate that the observed empirical patterns can arise purely from model misspecification. The first simulation constructs returns from a fully specified “true” twelve-factor model encompassing standard sources of systematic risk—market, size, value, profitability, investment, momentum, mispricing, volatility, and liquidity factors. In this setting, the true SDF and risk-free rate are known. I then deliberately introduce misspecification by omitting one or more factors from the estimated model and re-estimating the zero-beta rate using standard procedures. The results show a strong, monotonic relationship: as the number of omitted factors increases, the estimated zero-beta rate rises and converges towards the GMV mean return in the simulated world. This indicates that empirically observed biases can be fully explained by factor model misspecification alone, even when the true risk-free rate is low.

The previous simulation exercise only generates model misspecification from missing factors. To stay agnostic about the sources of misspecification, the second simulation ab-

stracts from specific factor structures and instead directly calibrates parameters of the mean–variance frontier. Assuming a true annual risk-free rate of 3%, I simulate many inefficient portfolios corresponding to misspecified models. For each simulated portfolio, I compute both the implied zero-beta rate and the maximum Sharpe ratio of zero-investment, zero-beta portfolios. The results confirm a clear pattern: as inefficiency increases, the estimated zero-beta rate converges toward the mean return of the GMV portfolio. When misspecification is severe (Sharpe ratio around 1.1), the probability that the estimated zero-beta rate exceeds 8% is approximately 80%. Hence, the empirical magnitude of the zero-beta rate puzzle—several percentage points above Treasury yields—can naturally arise from plausible degrees of model misspecification.

These simulations confirm that greater model misspecification leads to higher estimated zero-beta rates, even when the true risk-free rate is low, hence transforming the zero-beta rate puzzle from an empirical anomaly into a measurable outcome of model misspecification.

**Equity Risk Premium Puzzle and Risk-Free Rate.** The preceding results suggest that high estimated zero-beta rates can emerge purely from model misspecification, even when the true risk-free rate is low. A natural complementary question is whether a genuinely high risk-free rate would be consistent with equilibrium asset pricing. Intuitively, a higher risk-free rate does not alleviate the challenge of explaining equity returns. As the risk-free rate rises, investors would require even higher compensation to hold risky assets in order to maintain equilibrium consistency, implying a larger risk premium on the true tangency portfolio. In this sense, a high risk-free rate amplifies the equity risk premium puzzle, rather than resolving it, thereby creating greater tension with structural models in macro-finance.

**Literature and Contributions.** This paper contributes directly to the empirical literature on zero-beta rate estimation. [Fama and MacBeth \(1973\)](#), [Gibbons \(1982\)](#), [Giglio and Xiu \(2021\)](#), among others, estimate the zero-beta rate as the intercept from cross-sectional regressions, while [Long \(1971\)](#), [Black \(1972\)](#), and [Black et al. \(1972\)](#) solve for the minimum-variance market-neutral portfolio weights and compute its mean return. [Di Tella et al. \(2026\)](#) extend this framework by modeling the zero-beta rate as a time-varying function of macroeconomic predictors. Despite methodological differences, all find robustly high estimated zero-beta rates across models—a phenomenon I reinterpret as evidence of misspecification rather than as an equilibrium property of financial markets.

The paper’s key contribution is to reconceptualize the zero-beta rate puzzle as an identification failure stemming from model misspecification. I provide theoretical, empirical, and simulation-based evidence linking model misspecification to zero-beta rate estimates. Specifically, I argue that the empirical robustness of zero-beta rate estimates arises jointly from pervasive model misspecification and the reliance on the minimum-variance zero-beta

portfolio estimator. The results caution against the use of factor-model-implied zero-beta rates to infer the magnitude of the risk premium or the convenience yield, as these estimates do not provide information about the true risk-free rate.

This paper also speaks to the huge factor model literature (Ross (1976); Huberman (1982); Chamberlain and Rothschild (1983); Ingersoll Jr (1984); Connor and Korajczyk (1986); Fama and French (1993); Carhart (1997); Hou et al. (2015); Stambaugh and Yuan (2017); Fama and French (2018); Kelly et al. (2019); Gu et al. (2021); etc.). The paper proposes a new economically grounded measure of model misspecification based on the maximum Sharpe ratio of zero-investment, zero-beta portfolios, which complements existing statistical metrics by directly connecting model fit to exploitable investment opportunities.

In addition, the paper is related to the literature on factor model misspecification. A large body of work has focused on identification and inference issues arising from misspecified factor models, particularly in the estimation of factor risk premia, as well as the construction of robust test statistics and confidence intervals (Kan and Zhang, 1999b,a; Kan et al., 2013; Gospodinov et al., 2014, 2017; Kleibergen and Zhan, 2020; Manresa et al., 2023; etc.). In contrast, this paper shifts attention to a different object—the zero-beta rate—and shows that misspecification can fundamentally undermine the estimation for it. In particular, the zero-beta rate becomes non-unique and depends on the particular zero-beta portfolio used under misspecification, leading to systematic ambiguity in the estimated level of the zero-beta rate.

**Paper Structure.** The remainder of the paper is organized as follows. Section 2 develops the theoretical link between model misspecification and zero-beta rate estimation. Section 3 presents the data, model specifications, and empirical analysis quantifying the misspecification channel. Section 4 explores the economic tension between the equity risk premium puzzle and the risk-free rate. Section 5 concludes.

## 2. Zero-Beta Rate and Model Misspecification

First and foremost, it is important to draw a clear conceptual distinction between the risk-free rate and zero-beta rate. The risk-free rate is the expected return on an asset that delivers a certain payoff in every future state of the world. It is a fundamental economic object and does not depend on any particular asset pricing model. In contrast, a zero-beta rate is defined within the context of a specific factor model or, equivalently, with respect to a given stochastic discount factor (SDF). It represents the expected return on any portfolio that is orthogonal to the SDF or, equivalently, carries zero exposure to the model’s factors. In empirical asset pricing, researchers often assume that a risk-free rate exists (proxied by the Treasury yield) and that it coincides with the model-implied zero-beta rate.

This distinction immediately raises a fundamental identification issue. Can the true, unobserved risk-free rate be recovered from risky returns through zero-beta rate estimation? The answer is no. When a risk-free asset is absent—that is, when a unit payoff is not traded—the mean of the SDF,  $\nu = E[M]$ , is not identified from gross returns. One can arbitrarily choose  $\nu$  and construct an SDF that prices all observed risky payoffs. More generally, in an incomplete market there exist infinitely many admissible SDFs, and therefore infinitely many factor models, that perfectly price the same set of risky returns. Each admissible SDF implies its own model-specific zero-beta rate, which can be viewed as an internally consistent “risk-free rate” within that model. Consequently, the true risk-free rate is fundamentally not identified from risky returns alone. This insight is formalized by [Hansen and Jagannathan \(1991\)](#) and also emphasized by [Cochrane \(2009\)](#).

Nevertheless, the empirical literature routinely estimates zero-beta rates in settings where a risk-free asset is absent. Remarkably, a consistent empirical pattern has emerged: estimated zero-beta rates remain persistently high across a wide range of risk factors and model specifications. While the fundamental identification issue above establishes that zero-beta rates cannot reveal the true risk-free rate, it is a theoretical statement under perfect pricing and it does not explain the empirical regularity that estimated zero-beta rates are consistently high in practice—a phenomenon I refer to as the *zero-beta rate puzzle*. Since empirical factor models are inevitably misspecified and fail to perfectly price the cross section of asset returns, the central question becomes how model misspecification affects zero-beta rate estimation. This paper studies this link and argues that model misspecification is the primary force behind the persistently high estimated zero-beta rates.

The remainder of this section develops this argument within a mean-variance framework and shows how model misspecification can systematically distort estimated zero-beta rates. [Section 2.1](#) revisits existing estimation approaches and shows that the zero-beta rate is typically computed as the expected return of the unit-investment, minimum-variance zero-beta portfolio. [Section 2.2](#) demonstrates that when a factor model is misspecified, the expected returns of zero-beta portfolios can take infinitely many values, implying that the zero-beta rate is not unique within the model itself. In such cases, the conventional estimator based on the minimum-variance zero-beta portfolio becomes arbitrary and lacks clear economic interpretation. [Section 2.3](#) further investigates the observed robustness of zero-beta rate estimates and shows that large model misspecification systematically pushes the estimated zero-beta rate toward the mean return of the global minimum-variance (GMV, hereafter) portfolio. Thus, the empirical robustness of zero-beta rate estimates may arise jointly from pervasive model misspecification and the reliance on the minimum-variance zero-beta portfolio estimator. Finally, [Section 2.4](#) introduces a general, investment-based measure of factor

model misspecification that will be used in empirical assessments.

### 2.1. Zero-Beta Rate Estimation

Suppose there are  $N$  assets in the market,  $\mathbf{R}_{t+1} \in \mathcal{R}^N$ , indexed by  $i = 1, 2, \dots, N$ . Let  $\boldsymbol{\mu}$  denote the expected return vector and  $\boldsymbol{\Sigma}$  the variance-covariance matrix. A factor model posits that an asset return  $R_{i,t+1}$  follows a  $K$ -factor structure and the expected return in excess of the zero-beta rate is determined by risk loadings and factor risk premia:

$$R_{i,t+1} = \alpha_{i,t} + \boldsymbol{\beta}'_{i,t} \mathbf{f}_{t+1} + \varepsilon_{i,t+1} \quad (1)$$

$$\mu_{i,t} - r_{z,t} = a_{i,t} + \boldsymbol{\beta}'_{i,t} \boldsymbol{\lambda}_t \quad (2)$$

Equation (1) is the time-series statistical model of realized asset gross returns where  $\alpha_{i,t}$  is the intercept term,  $\boldsymbol{\beta}_{i,t}$  is the  $K \times 1$  vector of risk loadings, and  $\mathbf{f}_{t+1}$  is the  $K \times 1$  vector of factors. In the cross-sectional beta pricing model of equation (2),  $a_{i,t}$  represents the cross-sectional pricing errors which should be zero if the factor model is perfect,  $\boldsymbol{\lambda}_t$  is the  $K \times 1$  vector of factor risk premia, and  $r_{z,t}$  denotes the zero-beta rate—expected return for not taking any (systematic) risk. Empirically, the pricing errors  $a_{i,t}$  may exist for many reasons such as omitted risk factors, limits to arbitrage, transaction costs, market inefficiency, and so on. This is a general factor model framework incorporating a wide range of model specifications. The conventional Fama-French type of models assume constant betas and pre-specify known risk factors. The Arbitrage Pricing Theory (APT) models retain the constant betas but rely on principal component analysis (PCA) to extract statistical factors. With the recent development of machine learning techniques, advanced conditional models are able to formulate betas as functions of asset characteristics.

This paper examines the zero-beta rate  $r_{z,t}$  implied by a factor model, which is often proxied in empirical studies using the U.S. Treasury bill yield, since Treasury bills are regarded as risk-free assets. However, because Treasury securities offer a convenience yield, their yields tend to be lower than the true frictionless risk-free rate. The literature has proposed several methods to estimate the zero-beta rate implied by a given factor model. In particular, I revisit two widely used approaches for estimating the zero-beta rate: the portfolio-based approach and the regression-based approach.<sup>1</sup>

---

<sup>1</sup>Appendix C.1 discusses another category of zero-beta rate estimation—the test-optimization approach—which determines the value of the zero-beta rate that makes a given factor model as close as possible to being correctly specified (see [Kandel, 1984, 1986](#); [Shanken, 1986](#); [Velu and Zhou, 1999](#); [Beaulieu et al., 2013, 2023, 2025](#); [Ferson et al., 2025](#)). This approach, however, is not the main focus of the paper.

### 2.1.1. Portfolio Approach

The portfolio approach estimates the zero-beta rate in a given factor model by identifying the unit-investment, minimum-variance portfolio that has zero exposure to factor risks, and then taking its expected return as the zero-beta rate. This approach originates from Long (1971), Black (1972), and Morgan (1975), who explicitly solve for the minimum-variance portfolio orthogonal to the market. Formally, they solve the optimization problem:  $\min_{\omega} \omega' \Sigma \omega$  subject to  $\omega' \iota = 1$  and  $\omega' \Sigma \omega_m = 0$  where  $\iota$  is a vector of ones and  $\omega_m$  denotes the market portfolio weights. The zero-covariance constraint  $\omega' \Sigma \omega_m = 0$  ensures that the portfolio is market-neutral. The analytical solution is  $\omega_{z,mv} = (1 - \kappa \iota' \omega_m) \omega_{gmv} + \kappa \omega_m$  where  $\omega_{gmv} = \Sigma^{-1} \iota / (\iota' \Sigma^{-1} \iota)$  is the global minimum-variance portfolio weights, and  $\kappa = (\omega_{gmv}' \Sigma \omega_m) / ((\iota' \omega_m) \omega_{gmv}' \Sigma \omega_m - \omega_m' \Sigma \omega_m)$ . The sample average return on this minimum-variance, market-neutral portfolio provides an estimate of the zero-beta rate,  $r_z = \mathbb{E}[\omega_z' \mathbf{R}_{t+1}]$ .

Extending this framework from the CAPM to multifactor models introduces additional challenges, as the multifactor-mimicking portfolio weights are often difficult to obtain, particularly when factors are latent. Di Tella et al. (2026) generalize the portfolio approach by directly imposing zero-beta constraints for all factors in the portfolio construction stage:  $\min_{\omega} \omega' \Sigma \omega$  subject to  $\omega' \iota = 1$  and  $\omega' \beta = \mathbf{0}_K$  where  $\beta$  denotes the  $N \times K$  matrix of estimated betas<sup>2</sup> and  $\mathbf{0}_K$  denotes a  $K \times 1$  vector of zeros. The analytical solution for the minimum-variance zero-beta portfolio weights is given by:

$$\omega_{z,mv} = \Sigma^{-1} \begin{bmatrix} \iota & \beta \end{bmatrix} \left( \begin{bmatrix} \iota' \\ \beta' \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \beta \end{bmatrix} \right)^{-1} \begin{bmatrix} 1 \\ \mathbf{0}_K \end{bmatrix} \quad (3)$$

Given the portfolio weights, the zero-beta rate is estimated as the sample average return of the minimum-variance zero-beta portfolio:  $r_z = \mathbb{E}[\omega_z' \mathbf{R}_{t+1}]$ . Note that whereas Di Tella et al. (2026) study the time-series dynamics of the zero-beta rate, the focus of this paper is on understanding its level<sup>3</sup>.

---

<sup>2</sup>Betas are estimated differently across models. In Fama–French-type models, betas are obtained from time-series regressions, whereas in machine-learning-based models, they are often estimated as linear or nonlinear functions of firm characteristics.

<sup>3</sup>Di Tella et al. (2026) model the zero-beta rate as a linear function of macroeconomic predictors,  $\mathbf{Y}_t$ . They project realized zero-beta portfolio returns onto these predictors to obtain the conditional expected return:  $\omega_{z,mv}' \mathbf{R}_{t+1} = \delta' \mathbf{Y}_t + \epsilon_{t+1}$ . The fitted value  $\delta' \mathbf{Y}_t$  thus provides a time-varying estimate of the zero-beta rate. A practical challenge arises because the zero-beta rate is needed to construct excess returns used in estimating betas via time-series regressions. Since it is not known ex ante but depends on the estimated betas through equation (3), estimation must proceed iteratively or within a GMM framework that jointly determines the betas and the time-varying zero-beta rate.

### 2.1.2. Regression Approach

The regression approach estimates the zero-beta rate as the intercept in the cross-sectional regression of expected returns on estimated betas:

$$\mathbb{E}[R_i] = \lambda_0 + \boldsymbol{\lambda}'\boldsymbol{\beta}_i + e_i \quad (4)$$

where  $\mathbb{E}[R_i]$  is the sample mean return of asset  $i$ ,  $\boldsymbol{\beta}_i$  is the estimated  $K \times 1$  vector of risk loadings,  $\boldsymbol{\lambda}$  is the  $K \times 1$  vector of factor risk premia, and  $\lambda_0$  is the intercept term, interpreted as the zero-beta rate. The regression can be estimated using either ordinary least squares (OLS) or generalized least squares (GLS). The empirical literature often advocates the use of GLS since it accounts for heteroskedasticity and cross-sectional dependence, and the GLS estimator is invariant to portfolio repackaging<sup>4</sup>. The regression approach is employed in classic studies such as Black et al. (1972), Fama and MacBeth (1973), and Gibbons (1982) via the Fama–MacBeth two-pass procedure, and in more recent work such as Giglio and Xiu (2021) through the three-pass procedure. When GLS is used, the estimated intercept takes the following form:

$$\hat{\lambda}_{0,GLS} = \begin{bmatrix} 1 & \mathbf{0}_K \end{bmatrix} \left( \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \right)^{-1} \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \mathbb{E}[\mathbf{R}] = \boldsymbol{\omega}'_z \mathbb{E}[\mathbf{R}] \quad (5)$$

where  $\boldsymbol{\omega}_z$  denotes the minimum-variance zero-beta portfolio weights defined in equation (3), and  $\mathbb{E}[\mathbf{R}]$  is the vector of sample mean returns. Thus, the GLS intercept is mathematically identical to the sample mean return on the minimum-variance zero-beta portfolio. Similarly, when OLS is used, the intercept corresponds to the sample mean return of a particular zero-beta portfolio. In this sense, the regression approach is equivalent to the portfolio approach.

### 2.2. Uniqueness of Zero-Beta Rate

The previous discussion of estimation approaches has already alluded to the issue that the zero-beta rate associated with a given factor model may not be uniquely identified. Among the infinitely many zero-beta portfolios, empirical studies typically select the minimum-variance one. This convention follows the classic Black-CAPM (Black, 1972; Black et al., 1972), in which the intercept of the security market line corresponds to the expected return on the minimum-variance market-neutral portfolio when either no risk-free asset exists or

---

<sup>4</sup>Being invariant to portfolio repackaging means that the estimated risk premia remain unchanged when existing assets are merely recombined or duplicated into new portfolios. This property is desirable because asset pricing results should depend on the underlying investment opportunities rather than on the particular choice of test assets.

borrowing at the risk-free rate is constrained. In this section, however, I re-examine this portfolio choice and show that it is arbitrary and lacks clear economic interpretation.

Roll (1980) discusses the zero-beta portfolios (orthogonal portfolios) under CAPM. It proves that the zero-beta rate can take all values if the market portfolio is not mean-variance efficient. This argument is generally correct for any multi-factor model. As an extension to Roll (1980), I emphasize the following Proposition about the uniqueness of the zero-beta rate within any factor model (The complete proof is provided in Appendix B.1).

**Proposition 1.** *There exists an infinite number of unit-investment, zero-beta portfolios obtained within a factor model.*

- (i) *If the factor model is correctly specified, all the zero-beta portfolios have the same expected returns and the zero-beta rate is unique within this model.*
- (ii) *If the factor model is misspecified, the zero-beta portfolios do not have equal expected returns and the zero-beta rate is indeterminate.*


The intuition can be developed from the vector form of equation (2),  $\mathbb{E}_t[\mathbf{R}_{t+1}] - r_{z,t}\boldsymbol{\iota} = \mathbf{a}_t + \boldsymbol{\beta}_t\lambda_t$  where  $\mathbf{R}_{t+1}$  is the  $N \times 1$  vector of returns,  $\boldsymbol{\iota}_N$  is a  $N \times 1$  vector of ones,  $\mathbf{a}_t$  is the  $N \times 1$  vector of pricing errors, and  $\boldsymbol{\beta}_t$  is the  $N \times K$  matrix of betas. Consider a unit-investment, zero-beta portfolio with a  $N \times 1$  vector of weights  $\boldsymbol{\omega}_z$ . By definition,  $\boldsymbol{\omega}'_z\boldsymbol{\beta}_t = \mathbf{0}_K$  and  $\boldsymbol{\omega}'_z\boldsymbol{\iota}_N = 1$ . Therefore, the zero-beta portfolio should have an expected return:  $\mathbb{E}_t[\mathbf{R}_{z,t+1}] = \boldsymbol{\omega}'_z\mathbb{E}_t[\mathbf{R}_{t+1}] = r_{z,t} \cdot \boldsymbol{\omega}'_z\boldsymbol{\iota}_N + \boldsymbol{\omega}'_z\mathbf{a}_t + \boldsymbol{\omega}'_z\boldsymbol{\beta}_t\lambda_t = r_{z,t} + \boldsymbol{\omega}'_z\mathbf{a}_t$ . A correctly specified factor model with zero pricing errors ( $a_{i,t} = 0$ ) should recover a unique zero-beta rate  $r_{z,t}$ . By contrast, a misspecified model with non-zero pricing errors fails to disentangle the zero-beta rate from pricing errors. Different combinations of pricing errors lead to different expected returns of zero-beta portfolios, and thus indeterminate zero-beta rates.

We can understand the above proposition within the standard textbook mean-variance framework, following Roll (1980). First of all, asset pricing theory suggests that a correctly specified factor model corresponds to a mean-variance efficient portfolio on the efficient frontier, denoted by  $p^*$ , whereas a misspecified factor model corresponds to an inefficient portfolio  $p$ . Figure 1 shows the unit-investment zero-beta portfolios (orthogonal portfolios) for an efficient portfolio  $p^*$  and an inefficient portfolio  $p$  in the mean-standard deviation space. The blue horizontal solid line represents the zero-beta portfolios with respect to  $p^*$ , with  $z_{p^*}$  denoting the corresponding minimum-variance portfolio. The gray shaded region depicts the zero-beta portfolios with respect to  $p$ , and the black dashed hyperbola shows their zero-beta frontier<sup>5</sup>, where  $z_p$  is the corresponding minimum-variance portfolio. For illustration, when

---

<sup>5</sup>The zero-beta frontier with respect to portfolio  $p$  is defined as the set of zero-beta portfolios for portfolio  $p$  that minimize variance for a given level of mean return.

Figure. 1. Unit-Investment Zero-Beta Portfolios



*Notes:* This figure shows the unit-investment zero-beta portfolios (orthogonal portfolios) for an efficient portfolio  $p^*$  and an inefficient portfolio  $p$  in the mean-standard deviation space. The black hyperbola represents the mean-variance frontier. The blue horizontal solid line represents the zero-beta portfolios with respect to  $p^*$ , with  $z_{p^*}$  denoting the corresponding minimum-variance portfolio. The gray shaded region depicts the zero-beta portfolios with respect to  $p$ , and the black dashed hyperbola shows their zero-beta frontier, where  $z_p$  is the corresponding minimum-variance portfolio. For illustration, when  $p$  has the same expected return as  $p^*$ , the zero-beta frontier with respect to  $p$  is tangent to the mean-variance frontier at  $z_{p^*}$ .

$p$  has the same mean as  $p^*$ , the zero-beta frontier with respect to  $p$  is tangent to the mean-variance frontier at  $z_{p^*}$ . Proposition 1 (i) shows that the zero-beta portfolios with respect to  $p^*$  share the same expected return—the zero-beta rate—represented by the blue horizontal solid line. In this case, estimating the zero-beta rate using any zero-beta portfolio yields the same result, and the expected return of these portfolios,  $r_{z^*}$ , recovers the unique zero-beta rate. By contrast, Proposition 1 (ii) shows that the zero-beta portfolios with respect to  $p$  are located inside the shaded area. Hence, the zero-beta rate becomes indeterminate, admitting infinitely many possible values. In this case, the minimum-variance zero-beta portfolio is no longer on the mean-variance frontier and using it to estimate the zero-beta rate therefore produces an arbitrary expected return,  $r_z$ , that does not have a clear economic interpretation.

Revisiting Hansen and Jagannathan (1991), Cochrane (2009), and related work, the argument that the risk-free rate cannot be identified from risky returns alone constitutes a first, economic layer of non-uniqueness in the zero-beta rate. In mean-variance language, any efficient portfolio  $p^*$  (other than the GMV portfolio) defines a factor model, stochastic discount factor (SDF), or beta-pricing relation that perfectly prices all risky assets. Each

such model implies a unique model-specific zero-beta rate,  $r_z^*$ , which need not coincide with the true risk-free rate. Thus, non-uniqueness arises because multiple factor models can perfectly price the same set of risky assets. This implies that attempts to infer the true risk-free rate from factor models are inherently limited, even when a model prices all risky assets exactly.

Proposition 1 highlights a second and distinct source of non-uniqueness in the zero-beta rate: model misspecification. When a factor model fails to perfectly price assets—which is the empirically more relevant case than perfect pricing—pricing errors enter the expected returns of zero-beta portfolios, making it impossible to disentangle the zero-beta rate from mispricing. Consequently, different zero-beta portfolios can have different expected returns, so the zero-beta rate becomes indeterminate, even within a given factor model. In the next section, I examine how this misspecification-induced non-uniqueness translates into estimation bias in the inferred zero-beta rate.

*Remark on the mean-variance framework.* I use mean-variance analysis as the primary conceptual framework for expositional convenience. The mean-variance frontier is equivalent to the stochastic discount factor (SDF) and beta-pricing representations of asset pricing, and all results can be translated across these frameworks. A key advantage of the mean-variance perspective is that it provides a tractable way to analyze varying degrees of model misspecification while remaining agnostic about its underlying source. In this framework, a correctly specified factor model corresponds to a portfolio on the mean-variance frontier, and increasing misspecification corresponds to increasing portfolio inefficiency. This permits a continuous characterization of model misspecification through the location of a portfolio relative to the frontier. By contrast, in beta-pricing representations, misspecification appears as a pricing-error vector, while in SDF representations it appears as a deviation from the set of admissible discount factors (Hansen-Jagannathan distance). Although these representations are equivalent, they do not naturally provide a simple way to vary the degree of misspecification independently of its specific source. The mean-variance framework therefore offers a convenient and economically interpretable device for studying how model misspecification affects zero-beta rate estimation.

*Remark on the sources of model misspecification.* I use the term model misspecification in a broad sense, remaining agnostic about its underlying sources. Any deviation of a factor model from perfectly pricing all risky assets indicates model misspecification, whether it arises from omitted factors, weak factors, genuine mispricing, measurement errors<sup>6</sup>, or other departures from perfect pricing. In the mean–standard-deviation space, model mis-

---

<sup>6</sup>The notion of misspecification also subsumes classical beta measurement errors, which can distort estimated factor exposures and thereby place the implied factor model inside the efficient frontier.

specification implies that the factors do not span or intersect the mean–variance frontier. Accordingly, I treat portfolio inefficiency and factor model misspecification as equivalent concepts throughout the analysis.

### 2.3. Zero-Beta Rate and Model Misspecification

When the factor model is misspecified, the zero-beta rate is no longer uniquely determined within the model. As a consequence, the standard estimator based on the minimum-variance zero-beta portfolio becomes essentially arbitrary, despite its computational convenience and desirable econometric properties. For example, GLS is often preferred to OLS from multiple econometric grounds. But essentially, GLS corresponds to the minimum-variance zero-beta portfolio, whereas OLS corresponds to an alternative but equally valid choice of zero-beta portfolio, and is therefore equally arbitrary in the estimation. Although the standard estimator is arbitrarily chosen, there show systematic patterns as model misspecification shapes its behavior in predictable ways, offering a potential explanation for the empirical robustness of zero-beta rate estimates across specifications. To formalize this idea, I characterize model misspecification through portfolio inefficiency in mean–standard deviation space. Proposition 2 establishes the relationship between this inefficiency and the behavior of zero-beta rate estimates. (model misspecification) and the estimated zero-beta rate (The complete proof is provided in Appendix B.2.)

**Proposition 2.** *Suppose a factor model corresponds to a portfolio  $p$  with mean and variance  $r_p$  and  $\sigma_p$ . The true unobserved Tangency portfolio is denoted as  $T^*$ , which corresponds to the true unobserved risk-free rate  $r_f$ . Denote the mean and standard deviation of the global minimum-variance portfolio (GMV) as  $r_{GMV}$  and  $\sigma_{GMV}$ . Using the unit-investment, minimum-variance zero-beta portfolio with respect to portfolio  $p$ , the zero-beta rate is:*

$$r_z = r_{GMV} - \sigma_{GMV}^2 \frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} \quad (6)$$

- (i) *If the inefficient portfolio  $p$  lies above the GMV portfolio return ( $r_p > r_{GMV}$ ), then the zero-beta rate increases with portfolio inefficiency—it rises with higher volatility (holding mean return constant) and with lower mean return (holding volatility constant):*

$$\frac{dr_z}{d\sigma_p} > 0, \quad \frac{dr_z}{dr_p} < 0 \quad (7)$$

- (ii) *The zero-beta rate depends on the location of portfolio  $p$ :*

$$- r_z < r_f \text{ if and only if } (r_p - r_z)/\sigma_p^2 > (r_{T^*} - r_f)/\sigma_{T^*}^2.$$

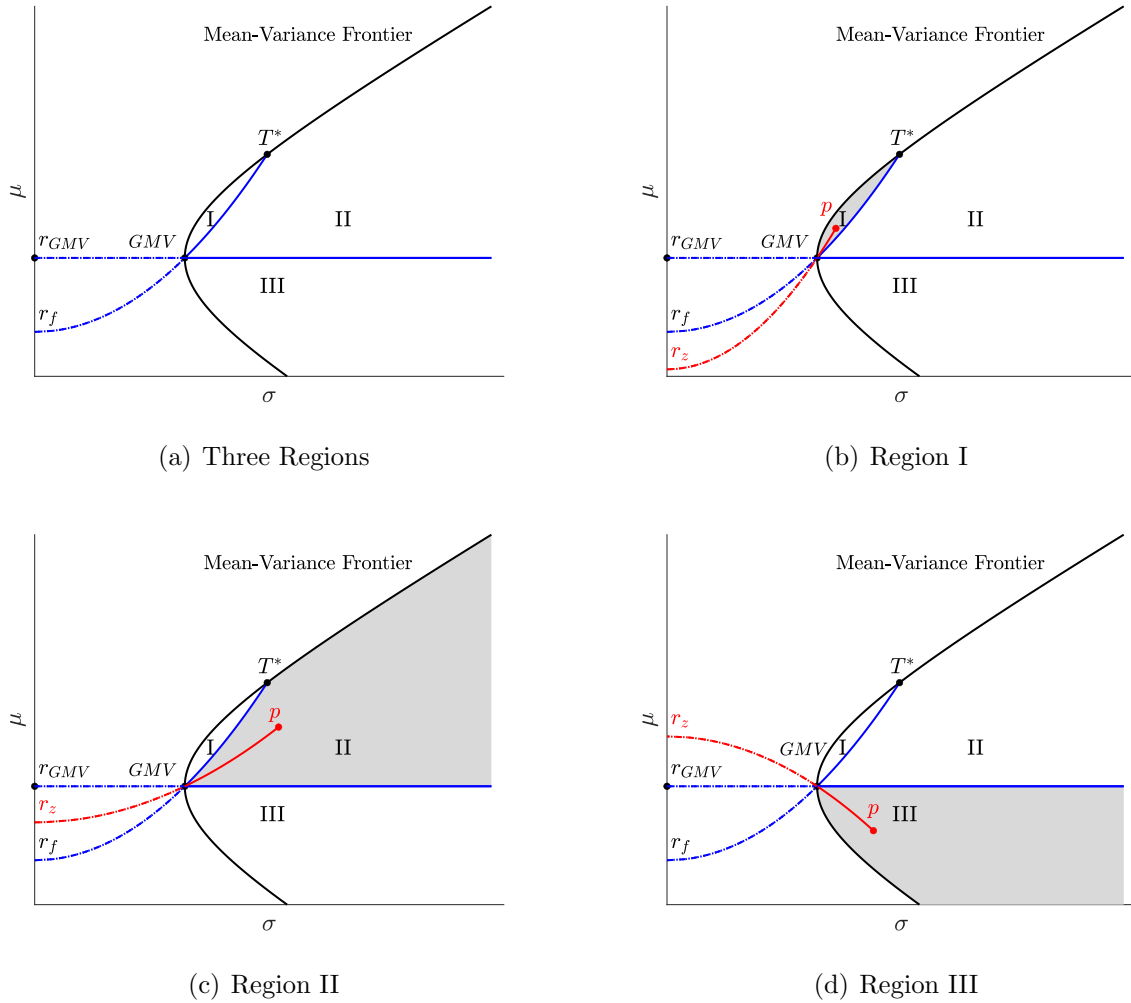
- $r_f \leq r_z \leq r_{GMV}$  if and only if  $(r_p - r_z)/\sigma_p^2 \leq (r_{T^*} - r_f)/\sigma_{T^*}^2$  and  $r_p \geq r_{GMV}$ .
- $r_z > r_{GMV}$  if and only if  $r_p < r_{GMV}$ .

(iii)  $r_z \rightarrow r_{GMV}$  as model misspecification increases.

Equation (6) provides the analytical formula for the expected return ( $r_z$ ) of a portfolio that minimizes variance subject to the constraint of having zero beta (zero covariance) with respect to a specific factor model portfolio  $p$ . This relationship implies that the zero-beta rate equals the return on the global minimum-variance (GMV) portfolio ( $r_{GMV}$ ) plus a “tilt” term—the precise adjustment required to satisfy the zero-beta constraint. Based on this formula, Proposition 2 (i) considers partial derivatives, showing that portfolio inefficiency—i.e., model misspecification—*inflates* the level of the estimated zero-beta rate. Assuming the inefficient portfolio  $p$  lies above the GMV portfolio, the estimated zero-beta rate increases monotonically with the degree of inefficiency along both risk and return dimensions. Holding the mean return constant, portfolios with higher volatility imply higher estimated zero-beta rates ( $dr_z/d\sigma_p > 0$ ). Conversely, holding volatility constant, portfolios with lower mean returns imply higher estimated zero-beta rates ( $dr_z/dr_p < 0$ ). Figure C.1 and C.2 visualize the relationship between portfolio inefficiency and the level of the estimated zero-beta rate in the mean–variance diagram.

Proposition 2 (ii) shows that the estimated zero-beta rate depends on the location of the associated inefficient portfolio  $p$  in mean–variance space. As illustrated in Figure 2 Panel (a), equation (6) divides the space into three regions using two boundaries. A first boundary is obtained by setting  $r_z = r_f$  in equation (6). The resulting curve passes through the unobserved tangency portfolio  $T^*$  and the GMV portfolio, and intersects the y-axis exactly at the risk-free rate  $r_f$ . Any portfolio on this boundary implies a zero-beta rate equal to the true (but unobserved) risk-free rate. A second boundary is derived by setting  $r_z = r_{GMV}$  in equation (6), yielding a horizontal line that passes through the GMV portfolio. Any portfolio on this boundary implies a zero-beta rate equal to  $r_{GMV}$ . If portfolio  $p$  lies in Region I, the estimated zero-beta rate is below the true risk-free rate ( $r_z < r_f$ ). In Panel (b), the red curve representing equation (6) intersects the y-axis below  $r_f$ . If  $p$  lies below the GMV portfolio in Region III, the zero-beta rate exceeds the GMV portfolio return ( $r_z > r_{GMV}$ ). In Panel (d), the red curve intersects the y-axis above  $r_{GMV}$ . If portfolio  $p$  lies in Region II, the zero-beta rate lies between the risk-free rate and the GMV return ( $r_f < r_z < r_{GMV}$ ). In Panel (c), the red curve intersects the y-axis between these two thresholds. This analysis shows that there is no obvious sign for the theoretical bias in zero-beta rate estimation with respect to the risk-free rate. Different forms of misspecification place the factor model  $p$  in different regions of the mean–variance space, potentially generating positive, negative, or zero biases.

Figure. 2. Zero-Beta Rate and Portfolio Inefficiency



*Notes:* This figure illustrates how zero-beta rate estimates vary with portfolio location in mean–variance space. Panel (a) depicts the three regions, while Panels (b), (c), and (d) show the corresponding zero-beta rate estimates within each region.

The central empirical question is therefore what pattern is observed in practice.

Empirically, regions II is the most relevant case to consider. Since the factor portfolio  $p$  typically exhibits a higher expected return than the GMV portfolio ( $r_p > r_{GMV}$ ), consistent with a positive risk–return trade-off<sup>7</sup>. This condition effectively rules out region III. Region I is also unlikely. As shown in Proposition 2 (ii), a factor model lies in Region I if and only if  $(r_p - r_z)/\sigma_p^2 > (r_{T^*} - r_f)/\sigma_{T^*}^2$ , meaning that the factor mean–variance portfolio must dominate the unobserved tangency portfolio in terms of excess return per unit variance (rather than Sharpe ratio). This condition is typically implausible, further ruling out Region I. Therefore, focusing on regions II, the estimated zero-beta rate  $r_z$  is highly likely to be upward biased relative to the true unobserved risk-free rate  $r_f$ .

Most importantly, predictable patterns of the estimated zero-beta rate can be inferred from portfolio inefficiency. Focusing on Figure 2 Panel (c), suppose there is another portfolio  $p'$  that is more inefficient than portfolio  $p$  (a more misspecified factor model) within region II.  $p'$  would locate further away from the mean-variance frontier, the red curve rotates clockwise, and the estimated zero-beta rate rises to the mean return of the GMV portfolio (see Figure C.3). This pattern suggests that as the inefficient portfolio  $p$  deviates further from the true tangency portfolio  $T^*$ , the estimated zero-beta rate converges toward  $r_{GMV}$ . In an extreme case with severe model misspecification ( $\frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} \rightarrow 0$ ) implies that the estimated zero-beta rate converges to  $r_{GMV}$  ( $r_z \rightarrow r_{GMV}$ ) according to equation (6). Intuitively, when the model is severely misspecified, the second term in equation (6) attributable to factor exposures embedded in the GMV portfolio vanishes. The estimated zero-beta rate converges to the mean return of the GMV portfolio. Importantly, this convergence pertains only to expected returns: the associated zero-beta portfolio need not collapse to the GMV portfolio itself, as its variance may remain distinct. To be rigorous, the ratio  $\frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2}$ , which can be interpreted as a GMV-anchored analogue of a Sharpe ratio, is closely related to the degree of model misspecification. This relationship will be made precise once the misspecification measure is introduced in Section 2.4.

In summary, Proposition 2 provides a comprehensive characterization of the relationship between model misspecification and the estimated zero-beta rate. The estimated rate depends on the location of the factor model portfolio in the mean–variance space. Without additional information, the zero-beta rate may coincide with, fall below, or exceed the true risk-free rate. In practice, however, model misspecification tends to inflate the estimated zero-beta rate, pushing it toward the mean return of the GMV portfolio. This convergence result is the key to understanding the zero-beta rate puzzle. Specifically, if all practical mod-

---

<sup>7</sup>If a factor portfolio has a lower expected return than the GMV portfolio, it is mean–variance inefficient, as there exist portfolios (e.g., the GMV) that deliver both higher expected return and lower risk.

els are strongly misspecified, all zero-beta rate estimates are close to a single benchmark, which explains the observed robustness in the estimation. This benchmark is the mean return of the GMV portfolio. An empirical evaluation of this implication will be presented in Section 3.

#### 2.4. Measure of Model Misspecification

The preceding analysis shows that model misspecification can account for the high estimates of zero-beta rates. An important but often overlooked step in the zero-beta-rate literature is to quantify the degree of model misspecification, as doing so is essential for evaluating the potential magnitude of the puzzle. When a factor model closely approximates an admissible stochastic discount factor (SDF), it should feature a unique zero-beta rate—though not necessarily the true risk-free rate. In contrast, when the model is substantially misspecified, its estimated zero-beta rate tends to coincide with the mean return of the global minimum-variance (GMV) portfolio. Thus, while earlier studies often report that zero-beta rate estimates appear robust across different factor models, this apparent robustness may simply reflect the fact that those models are misspecified to a similar degree.

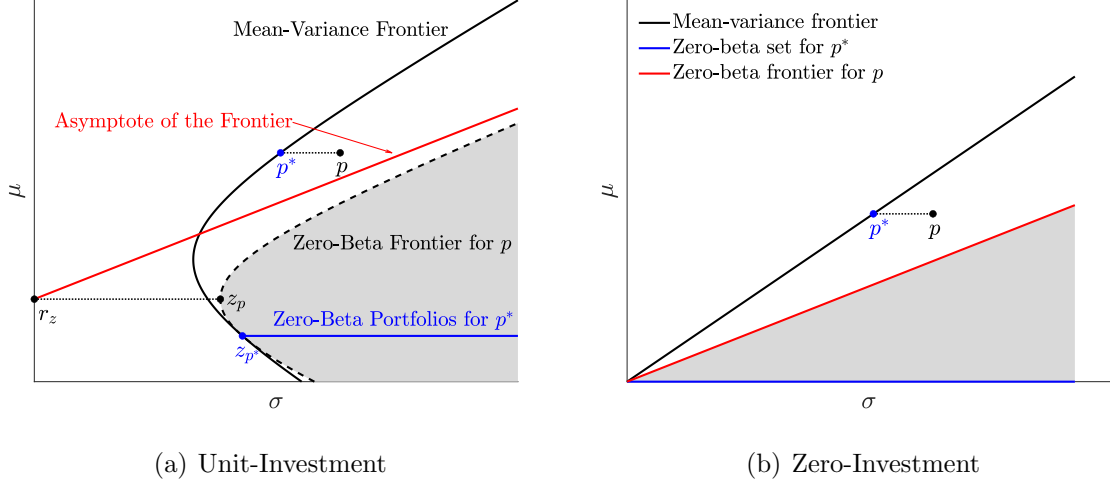
Proposition 3 introduces a new diagnostic tool of model misspecification grounded in the behavior of zero-beta portfolios. The idea builds on Proposition 1, which establishes that a correctly specified factor model should imply a unique zero-beta rate. Hence, one can evaluate model misspecification by examining how much the returns on zero-beta portfolios deviate from featuring a unique rate. (The complete proof is provided in Appendix B.3.)

**Proposition 3.** *The slope of the asymptote for the **unit-investment**, zero-beta frontier,  $S$ , equals the slope of the **zero-investment**, zero-beta frontier,  $S_z$ . These slopes correspond to the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios and thereby provide a measure of model misspecification.*

Proposition 3 establishes that the slope of the asymptote in the unit-investment setting and the slope of the zero-beta frontier in the zero-investment setting are equivalent measures of factor model misspecification. This is because any unit-investment, zero-beta portfolio can be orthogonally decomposed into a zero-investment, zero-beta portfolio and the unit-investment, minimum-variance zero-beta portfolio.

Figure 3 illustrates this misspecification measure using the mean–variance diagrams. The left panel depicts the zero-beta frontier in a unit-investment setting, while the right panel presents its counterpart in a zero-investment setting. In both panels, portfolio  $p^*$  denotes the true tangency portfolio, whereas an inefficient portfolio  $p$  represents one implied by a

Figure. 3. Measuring Model Misspecification



*Notes:* This figure shows the measure of factor model misspecification in both a unit-investment setting (left panel) and a zero-investment setting (right panel). Among unit-investment portfolios in Panel (a), the red solid line represents the asymptote of the zero-beta frontier (black dashed hyperbola) with respect to an inefficient portfolio  $p$ . Among zero-investment portfolios in Panel (b), the red solid line is the zero-beta frontier with respect to an inefficient portfolio  $p$ .

misspecified factor model. In Panel (a), the red solid line shows the asymptote of the zero-beta frontier (the black dashed hyperbola) corresponding to portfolio  $p$ . For a correctly specified model, the zero-beta set is a horizontal line (blue solid line) with a unique zero-beta rate. Thus, the slope of the asymptote measures how much zero-beta returns diverge from featuring a unique rate. According to Proposition 3, this asymptotic slope,  $S$ , equals to the slope of the linear zero-beta frontier corresponding to portfolio  $p$  in Panel (b),  $S_z$ , represented by the red solid line. For a correctly specified model, the zero-beta set coincides with the horizontal axis (blue solid line), implying zero expected returns. In contrast, for a misspecified model—represented by the inefficient portfolio  $p$ —the shaded area illustrates all possible zero-investment, zero-beta portfolios. The slope of the zero-beta frontier thus corresponds to the maximum Sharpe ratio attainable by such portfolios. This slope can be expressed as:

$$S^2 = S_z^2 = SR^2(p^*) - SR^2(p) \tag{8}$$

Equation (8) shows that the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios equals the difference in Sharpe ratios between the true tangency portfolio and the model-implied factor portfolio. This “Sharpe ratio spread”, akin to the GRS statistic<sup>8</sup>,

<sup>8</sup>The GRS statistics:  $J \propto \frac{SR^2(\mathbf{R}, \mathbf{F}) - SR^2(\mathbf{F})}{1 + SR^2(\mathbf{F})}$  where  $\mathbf{R}$  and  $\mathbf{F}$  denotes the test asset returns and factor returns. It represents the gain in Sharpe ratio from adding test assets to the factor set.

provides an economically interpretable measure of model misspecification.

As a supplement to Proposition 2 (iii), we can show that:

$$S_z^2 = S^2 = \frac{ac - b^2}{a} \cdot \left( 1 - \frac{\frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2}}{\frac{r_{p^*} - r_{GMV}}{\sigma_{p^*}^2 - \sigma_{GMV}^2}} \right) \in \left[ 0, \frac{ac - b^2}{a} \right] \quad (9)$$

where  $a = \boldsymbol{\nu}'\boldsymbol{\Sigma}\boldsymbol{\nu}$ ,  $b = \boldsymbol{\nu}'\boldsymbol{\Sigma}\boldsymbol{\mu}$ , and  $c = \boldsymbol{\mu}'\boldsymbol{\Sigma}\boldsymbol{\mu}$  are standard parameters of the mean-variance geometry. Portfolio  $p$  corresponds to the factor model, and portfolio  $p^*$  is the efficient portfolio with  $r_{p^*} = r_p$ . Equation (9) implies that:

$$\underbrace{S^2 \rightarrow \frac{ac - b^2}{a}}_{\text{model misspecification } \uparrow} \Leftrightarrow \frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} \rightarrow 0 \Rightarrow r_z \rightarrow r_{GMV} \quad (10)$$

Hence, shrinking GMV-anchored analogue of a Sharpe ratio  $\left(\frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2}\right)$  means rising model misspecification, which leads to the convergence of the estimated zero-beta rate towards the mean return of the GMV portfolio according to equation (6).

#### 2.4.1. Maximum Sharpe Ratio of Zero-Investment, Zero-Beta Portfolios

For empirical implementation, I focus on the zero-investment setting and construct the misspecification measure by estimating the *maximum Sharpe ratio of zero-investment, zero-beta portfolios*. Working in the zero-investment setting is more tractable and an additional advantage is that the risk-free rate cancels out in zero-investment portfolios, allowing model misspecification to be evaluated without knowing the true risk-free rate or estimating a zero-beta rate. Specifically, I construct the optimal zero-investment, zero-beta portfolio with no exposure to any systematic risk factors of the model and evaluate its investment performance. This is achieved by solving the following constrained mean-variance optimization problem:

$$\begin{aligned} \max_{\boldsymbol{\omega}} \quad & \boldsymbol{\omega}'\boldsymbol{\mu} - \frac{\gamma}{2}\boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega} \\ \text{s.t.} \quad & \boldsymbol{\omega}'\boldsymbol{\nu} = 0, \quad \boldsymbol{\omega}'\boldsymbol{\beta} = \mathbf{0}_K \end{aligned} \quad (11)$$

where  $\gamma$  is the risk aversion coefficient. The analytical solution for the optimal zero-investment, zero-beta portfolio weights is:

$$\boldsymbol{\omega}_z^* = \frac{1}{\gamma}\boldsymbol{\Sigma}^{-1} \left[ \mathbf{I} - \boldsymbol{\Pi} (\boldsymbol{\Pi}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\Pi})^{-1} \boldsymbol{\Pi}'\boldsymbol{\Sigma}^{-1} \right] \boldsymbol{\mu} \quad (12)$$

where  $\boldsymbol{\Pi} = [\boldsymbol{\nu}, \boldsymbol{\beta}]$ . Recall that the unconstrained optimal portfolio weights are given by

$\boldsymbol{\omega}^* = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$ . Comparing the two, equation (12) can be interpreted as the optimal portfolio based on the projected mean return  $\mathbf{P}_{\Pi} \boldsymbol{\mu}$ , where  $\mathbf{P}_{\Pi} \equiv \left[ \mathbf{I} - \boldsymbol{\Pi} (\boldsymbol{\Pi}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\Pi})^{-1} \boldsymbol{\Pi}' \boldsymbol{\Sigma}^{-1} \right]$  is a generalized projection matrix with weights  $\boldsymbol{\Sigma}^{-1}$ . This matrix projects the mean return vector onto the subspace orthogonal to  $\boldsymbol{\Pi}$ , so that  $\tilde{\boldsymbol{\alpha}} \equiv \mathbf{P}_{\Pi} \boldsymbol{\mu}$  can be interpreted as pricing errors (alphas) and  $\boldsymbol{\omega}_z^*$  represents the optimal portfolio of alphas that is beta neutral.

The maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios is:

$$S_z = \sqrt{\tilde{\boldsymbol{\alpha}}' \boldsymbol{\Sigma}^{-1} \tilde{\boldsymbol{\alpha}}} \quad (13)$$

Equation (13) indicates that  $S_z$  captures the “arbitrage”<sup>9</sup> investment opportunities that can be exploited under a given factor model without taking on factor risk. Intuitively, the higher the attainable Sharpe ratio from this zero-beta (factor-neutral) portfolio, the farther the model is from correctly pricing all assets—implying greater model misspecification.

#### 2.4.2. Comparison with Traditional Diagnostic Tools

This paper introduces the *maximum Sharpe ratio of zero-investment, zero-beta portfolios*,  $S_z$ , as a new diagnostic tool of evaluating model misspecification. While it is conceptually related to the GRS statistic via equations (8) and (13), its empirical construction and interpretation are fundamentally different. Instead of estimating alphas through time-series regressions to form a test statistic,  $S_z$  is built directly from zero-investment, zero-beta portfolios and evaluates model misspecification through the investment opportunities embedded in pricing errors. This new measure yields several advantages.

First,  $S_z$  does not require estimation of SDF parameters or factor risk premia, thereby avoiding the well-known estimation issues associated with misspecification and weak identification (Kan and Zhang, 1999b; Gospodinov et al., 2014; Kleibergen and Zhan, 2020; among others). By contrast, the Hansen-Jagannathan distance (HJD) is explicitly based on estimating SDF parameters and computing the minimum distance between a proposed SDF and the set of admissible SDFs that correctly price test assets, making it inherently parameter-dependent<sup>10</sup>.

Second,  $S_z$  provides an investment-based diagnostic in the spirit of GRS, but without requiring specification or estimation of a risk-free rate or zero-beta rate, both of which are central inputs in GRS-type procedures (Shanken, 1986; Gibbons et al., 1989; among others).

<sup>9</sup>Similar to Kim et al. (2021), the notion of “arbitrage” is that portfolios are constructed to hedge out the systematic risk.

<sup>10</sup> $HJD = \min_{\gamma} \mathbb{E}[M(\gamma) \mathbf{R} - 1]' \mathbf{W}^{-1} \mathbb{E}[M(\gamma) \mathbf{R} - 1]$  where  $M(\gamma)$  is the SDF depending on parameters  $\gamma$ ,  $\mathbf{R}$  is a matrix of asset returns, and  $\mathbf{W}$  is the weighting matrix.

Third, as will be shown in the empirical section, the optimal zero-investment, zero-beta portfolios underlying exhibit only modest portfolio turnover and limited short positions, making them more economically plausible and implementable. This stands in contrast to GRS-based implementations, where portfolios that achieve maximum Sharpe ratios in GRS-type constructions often involve highly unrealistic leverage and very large short positions (Fama and French, 2015).

Finally, this method is naturally suited to extensions such as out-of-sample evaluation and the incorporation of transaction costs, making it a practical and flexible tool for assessing economic model misspecification in realistic trading environments.

### 3. Empirical Assessment

#### 3.1. Data

I obtain monthly individual stock returns and characteristics from the [Global Factor Data](#) website organized by [Jensen, Kelly, and Pedersen \(2023\)](#) (JKP).<sup>11</sup> My sample spans January 1960 to December 2024, covering 780 months (65 years). In total, there are 3,658,843 stock-month observations for 28,828 unique stocks, averaging 4,691 stocks per month. For each characteristic, I fill missing values with the cross-sectional median by 2-digit SIC industry each month. After this step, I retain 136 characteristics with complete coverage across the full sample (see Appendix A.1 for the full list). All characteristics are lagged one month. JKP update characteristics using the most recent accounting data four months after the fiscal period ends, ensuring that lagged characteristics are in the public information set and avoiding look-ahead bias.

I use characteristic-sorted portfolios in model evaluation and zero-beta rate estimation. Following [Jensen et al. \(2023\)](#), I construct portfolios and factors for each characteristic and retain the two corner portfolios (top and bottom terciles), since much of the relevant information resides in the extremes ([Lettau and Pelger, 2020](#)). I also include the middle tercile portfolio sorted by size so that the market return is spanned by the testing portfolios. This yields a total of  $136 \times 2 + 1 = 273$  univariate-sorted portfolios. Each factor is constructed as the return spread between portfolios in the top and bottom terciles of a given characteristic. The factor's sign is adjusted, if necessary, to ensure that its average return over the sample period is positive. The timing of my portfolio formation differs a bit from standard practice: while Fama–French form portfolios annually in June and JKP form them monthly, I construct portfolios each December, aligned with the rolling out-of-sample periods in my

---

<sup>11</sup>I thank the authors for making the data easily accessible (a WRDS account with access to CRSP and Compustat is required).

following analysis. During this procedure, I store portfolio weights on individual stocks. My empirical results are robust to the portfolio formation method.

In conditional factor models (described in Section 3.2), lagged characteristics also serve as determinants of model parameters. Following Gu et al. (2020), Gu et al. (2021), and others, I cross-sectionally rank-normalize all characteristics into the  $(-1, 1)$  interval each month. Since stock characteristics often display high skewness and kurtosis, the rank transformation reduces sensitivity to outliers.

### 3.2. Candidate Factor Models

I study a wide range of factor models in this paper. For each model, I evaluate specifications with 1, 3, 6, and 9 factors. A general representation of a factor model is:

$$r_{i,t+1} = \alpha(\mathbf{z}_{i,t}) + \boldsymbol{\beta}(\mathbf{z}_{i,t})' \mathbf{f}_{t+1} + \varepsilon_{i,t+1} \quad (14)$$

where  $\mathbf{f}_{t+1}$  is a  $K$ -dimensional vector of factors,  $\alpha(\mathbf{z}_{i,t})$  and  $\boldsymbol{\beta}(\mathbf{z}_{i,t})$  denote the intercept and risk loadings, potentially functions of the 136 stock characteristics.

First, I consider *unconditional linear models with pre-specified factors*, the most widely used class of models. These models assume that a small set of observable, economically motivated factors explain stock returns, with constant intercepts and loadings:  $\alpha(\mathbf{z}_{i,t}) = \alpha_i$  and  $\boldsymbol{\beta}(\mathbf{z}_{i,t}) = \boldsymbol{\beta}$ . Prominent examples include Fama and French (1993), Carhart (1997), Hou et al. (2015), Stambaugh and Yuan (2017), and Fama and French (2018). In the 1-factor case, I include only the market factor. In the 3-factor case, I include market, size, and value factors. In the 6-factor case, I include market, size, value, profitability, investment, and momentum factors. In the 9-factor case, I add the betting-against-beta (BAB) factor (Frazzini and Pedersen, 2014) and two mispricing factors (Stambaugh and Yuan, 2017) to the 6-factor specification.<sup>12</sup> I refer to these models collectively as “FF”.

Second, I use *unconditional linear models with PCA factors*, rooted in the Arbitrage Pricing Theory (APT) (Ross, 1976; Huberman, 1982; Chamberlain and Rothschild, 1983; Ingersoll Jr, 1984; Connor and Korajczyk, 1986, among others). Unlike the CAPM or the Intertemporal CAPM (ICAPM), which derive from equilibrium models with explicit preferences and market assumptions, APT is a reduced-form framework.<sup>13</sup> It assumes a factor structure in which returns decompose into systematic and idiosyncratic components. With

---

<sup>12</sup>The market factor is the weighted average of all stocks. The size, value, profitability, investment, momentum, BAB, mispricing factors are constructed using characteristics “market\_equity”, “be\_me”, “ope\_me”, “at\_gr1”, “ret\_12\_1”, “betabab\_1260d”, “mispricing\_mgmt”, and “mispricing\_perf”.

<sup>13</sup>Extensions that embed APT in an equilibrium setting include Connor (1984) and Connor and Korajczyk (1988).

sufficiently many assets, idiosyncratic risk diversifies away, and the absence of arbitrage opportunities yields an approximate linear beta-pricing relation. As in the “FF” case, loadings are static:  $\alpha(\mathbf{z}_{i,t}) = \alpha_i$  and  $\beta(\mathbf{z}_{i,t}) = \beta$ . The APT naturally motivates the use of principal component analysis (PCA) to extract statistical factors. Cooper et al. (2021) demonstrate that such statistically constructed factors outperform most of the traditional “FF”-style multi-factor models, in both economic and statistical terms. Following this insight, I extract 1, 3, 6, and 9 factors using the standard PCA. I refer to these specifications collectively as “PCA”.

Third, I examine *conditional linear models with latent factors*, which allow risk loadings to vary with stock characteristics. Specifically, I implement the instrumented principal component analysis (IPCA) models of Kelly et al. (2019) with 1, 3, 6, and 9 factors. In this framework, the intercepts and loadings are modeled as linear functions of observable characteristics:  $\alpha(\mathbf{z}_{i,t}) = \Gamma'_\alpha \mathbf{z}_{i,t}$  and  $\beta(\mathbf{z}_{i,t}) = \Gamma'_\beta \mathbf{z}_{i,t}$ . Unlike FF models, which rely on pre-specified factors, or PCA models, which assume static loadings, IPCA jointly estimates latent factors and their time-varying exposures using an alternating least squares (ALS) algorithm. I collectively refer to these models as “IPCA”.<sup>14</sup>

Finally, I also study *conditional non-linear models with latent factors*, which leverage machine learning methods to capture richer relationships between characteristics and risk exposures. While IPCA models imposes linearity, neural networks can approximate complex non-linear mappings. I use the conditional autoencoder model of Gu et al. (2021). Autoencoders are neural networks designed for unsupervised dimension reduction, which can be viewed as nonlinear analogues of PCA. They aim to learn a compressed, low-dimensional representation of input data by training the network to reconstruct their own inputs as accurately as possible. A standard latent factor model can be interpreted as a simple autoencoder, while conditional autoencoders extend this by incorporating observable characteristics. The architecture consists of two networks: a multi-layer beta network capturing non-linear mappings from characteristics to loadings, and a single-layer factor network generating latent factors as linear combinations of portfolios. The two are then combined as in equation (14). My implementation follows Gu et al. (2021) but adds an intercept term in the beta network, allowing  $\alpha_{i,t}$  to vary flexibly with characteristics, and uses the 273 characteristic-sorted portfolios as the input layer to the factor network. Estimation relies on stochastic gradient descent (SGD), with learning rate tuning, LASSO ( $l_1$ ) penalization, and early stopping for

---

<sup>14</sup>Another strand of conditional linear factor models emphasizes time-varying risk premia in addition to time-varying loadings, pioneered by Ferson and Harvey (1991), who attribute much of cross-sectional return predictability to variations in risk premia than by variations in betas. Gagliardini et al. (2016) further develop econometric methods for large panels of individual stocks, modeling both risk premia and risk loadings as parametric functions of macro instruments and stock characteristics.

Table 1: Model Total  $R^2$  (%)

Models	In-Sample				Out-of-Sample			
	1-factor	3-factor	6-factor	9-factor	1-factor	3-factor	6-factor	9-factor
FF	87.0	93.3	95.0	95.9	86.6	92.1	93.9	94.5
PCA	92.7	97.3	98.4	98.9	92.3	96.2	97.4	98.1
IPCA	75.4	93.1	94.4	95.3	73.6	93.0	94.4	95.2
AE	83.1	90.8	94.2	95.0	80.5	88.8	92.8	92.1

*Notes:* This table reports the total  $R^2$  values for characteristic-sorted portfolios across four classes of factor models containing 1, 3, 6, and 9 factors.

regularization.<sup>15</sup> I refer to these models collectively as “AE”.<sup>16</sup>

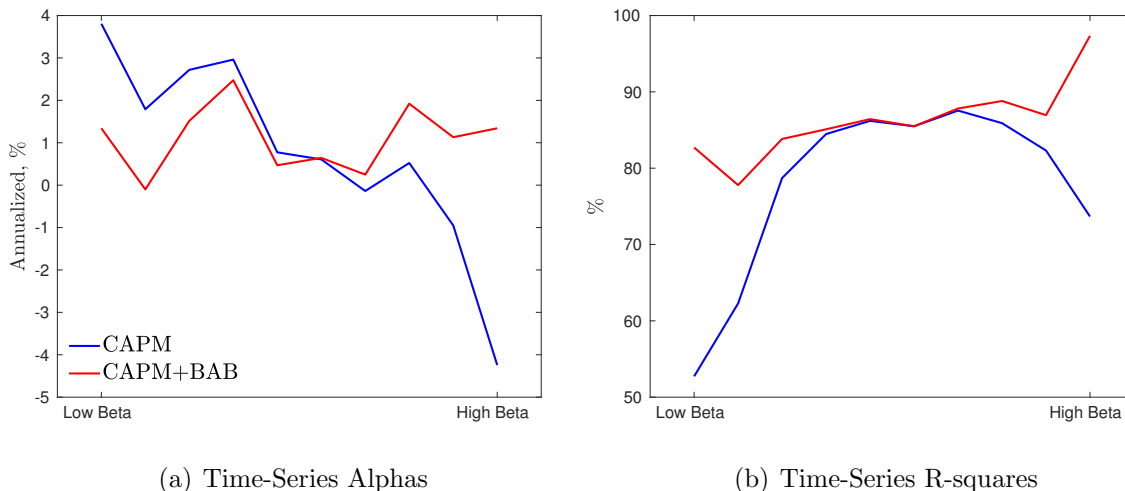
I evaluate factor models both in-sample and out-of-sample. For in-sample analysis, I run a one-time full-sample model estimation. For out-of-sample analysis, I estimate models using expanding windows and apply estimated model parameters in the out-of-sample period. For the conditional autoencoder models, in particular, I split the full sample into training, validation, and testing sets. The initial training period is 1960–1977 (18 years), the validation period is 1978–1989 (12 years), and the testing period is 1990–1991 (1 year). Following the literature (e.g., Gu et al., 2020), I refit the models annually. At each refit, the training sample expands by one year, while the validation sample is rolled forward with a fixed length, always including the most recent 12 years. This setup yields an out-of-sample period from 1990 to 2024, totaling 35 years. Since non-deep learning models typically do not require hyperparameter tuning, I combine the training and validation samples for estimation and use the same 1-year testing window for out-of-sample evaluation. To ensure comparability, both the in-sample and out-of-sample periods are set to span January 1990 through December 2024, totaling 420 months. When the sample period is extended to January 1965–December 2024 for the FF, PCA, and IPCA models, the results and conclusions remain unchanged.

I start with the statistical performance of various factor models. Kelly et al. (2019) and Gu et al. (2021) introduce total  $R^2$  to measure the model explanatory power of test assets using contemporaneous factor realizations:

<sup>15</sup>Other machine learning approaches include Feng et al. (2024), who use feed-forward networks to map characteristics into deep characteristics that generate latent deep factors, and Chen et al. (2024), who incorporate no-arbitrage directly into the loss function via a generative adversarial network (GAN) framework. Their architecture pairs an SDF network that constructs the pricing kernel with a conditional network that selects assets and moments yielding the largest mispricings, iterating until arbitrage opportunities are eliminated.

<sup>16</sup>I use two hidden layers in the beta network, with 32 and 16 neurons, respectively. The empirical results are robust to the choice of network depth.

Figure. 4. Ten Beta-Sort Portfolios: Time-Series Regression Results



*Notes:* This figure shows the time-series regressions results of pricing ten beta-sorted portfolios using two factor models: (1) the standard CAPM and (2) an extended two-factor model incorporating both the market factor (MKT) and the betting-against-beta factor (BAB). Panels A and B show the time-series alphas and R-squares for the 10 portfolios, respectively.

$$R_{\text{total}}^2 = 1 - \frac{\sum_{i,t} (r_{i,t} - \hat{\beta}'_i \hat{\mathbf{f}}_t)^2}{\sum_{i,t} r_{i,t}^2}. \quad (15)$$

Table 1 reports the total  $R^2$  values for characteristic-sorted portfolios across four classes of factor models containing 1, 3, 6, and 9 factors. All specifications display high time-series  $R^2$  values above 75%. These results indicate that existing factor returns account for nearly all of the common time-series variation in stock returns (Lopez-Lira and Roussanov, 2020). However, they may still perform poorly in explaining the cross-section of expected returns, leaving substantial pricing errors unaccounted for. Table D.1 and D.2 also report the statistical performance of these factor models on individual stocks.

### 3.3. A Simple Exercise: Ten Beta-Sorted Portfolios

Before analyzing the four classes of factor models on the 273 characteristic-sorted portfolios, I begin with a simple illustrative exercise using only ten test assets. The purpose of this exercise is to demonstrate a key conceptual point: even if one identifies a model that near-perfectly prices all risky assets, that model may correspond to an arbitrary stochastic discount factor (SDF) implying an arbitrary zero-beta rate—one that provides no information about the true risk-free rate.

Table 2: Ten Beta-Sort Portfolios: Cross-Sectional Regression Results

Models	$\lambda_0$ (ann., %)	$\lambda_{\text{MKT}}$ (ann., %)	$\lambda_{\text{BAB}}$ (ann., %)	$R^2$ (%)
CAPM	6.50 (6.41)	6.04 (5.31)		77.9
CAPM+BAB	1.90 (0.50)	10.67 (3.85)	-4.40 (-4.65)	86.3

*Notes:* This table reports the cross-sectional regression results of pricing ten beta-sorted portfolios using two factor models: (1) the standard CAPM and (2) an extended two-factor model incorporating both the market factor (MKT) and the betting-against-beta factor (BAB). The regression for the two-factor model is:  $\mathbb{E}[R_i] = \lambda_0 + \beta'_{i,\text{MKT}}\lambda_{\text{MKT}} + \beta'_{i,\text{BAB}}\lambda_{\text{BAB}} + e_i$ . Estimation uses generalized least squares (GLS). t-statistics are reported in parentheses.

I evaluate ten portfolios sorted by market beta as test assets and compare two factor models: (1) the standard CAPM and (2) an extended two-factor model incorporating both the market factor (MKT) and the betting-against-beta factor (BAB). Figure 4 reports the time-series regression results for pricing the ten beta-sorted portfolios under these two models. Panel A illustrates the well-known beta anomaly under the CAPM: low-beta portfolios exhibit positive alphas, while high-beta portfolios display negative alphas. Adding the BAB factor eliminates this pattern and improves the model’s time-series fit, as reflected in higher  $R^2$ ’s (Panel B). In the cross-section regressions estimated via generalized least squares (GLS) shown in Table 2, the CAPM produces a statistically significant annualized intercept  $\lambda_0$  of 6.50%, whereas the intercept becomes statistically insignificant under the two-factor model. The cross-sectional  $R^2$  also increases from 77.9% to 86.3%. These results indicate that the two-factor specification statistically passes the cross-sectional test: the MKT and BAB factors jointly capture the cross-section of ten beta-sorted portfolios, leaving an insignificant intercept.

Since the cross-sectional regressions are estimated using GLS, the intercept term  $\lambda_0$  is equivalent to the estimated zero-beta rate implied by the unit-investment, minimum-variance zero-beta portfolio (see Section 2.1). Accordingly, the zero-beta rate estimated from the ten beta-sorted portfolios under the two-factor model is 1.90% per year, though it is statistically insignificant. During the same sample period, the average one-month Treasury yield is 3.26%. This finding illustrates that the zero-beta rate can take on an arbitrary value even when the factor model achieves an almost perfect fit in pricing risky assets. In mean–variance terms, a near-perfect pricing model merely implies that the factor-model portfolio lies close to the mean–variance frontier, not necessarily close to the (unobserved) tangency portfolio

Table 3: Maximum Sharpe Ratios of Zero-Investment Zero-Beta Portfolios (Annualized)

Models	In-Sample				Out-of-Sample			
	1-factor	3-factor	6-factor	9-factor	1-factor	3-factor	6-factor	9-factor
FF	3.31	3.30	3.26	3.15	1.27	1.27	1.20	1.08
PCA	3.31	3.30	3.18	3.15	1.26	1.28	1.16	1.01
IPCA	3.30	3.29	3.26	3.21	1.28	1.27	1.12	0.99
AE	3.31	3.30	3.29	3.17	1.27	1.16	1.09	0.98

*Notes:* This table reports the in-sample and out-of-sample annualized maximum Sharpe ratios of zero-investment, zero-beta portfolios implied by the FF, PCA, IPCA, and AE models with 1, 3, 6, and 9 factors.

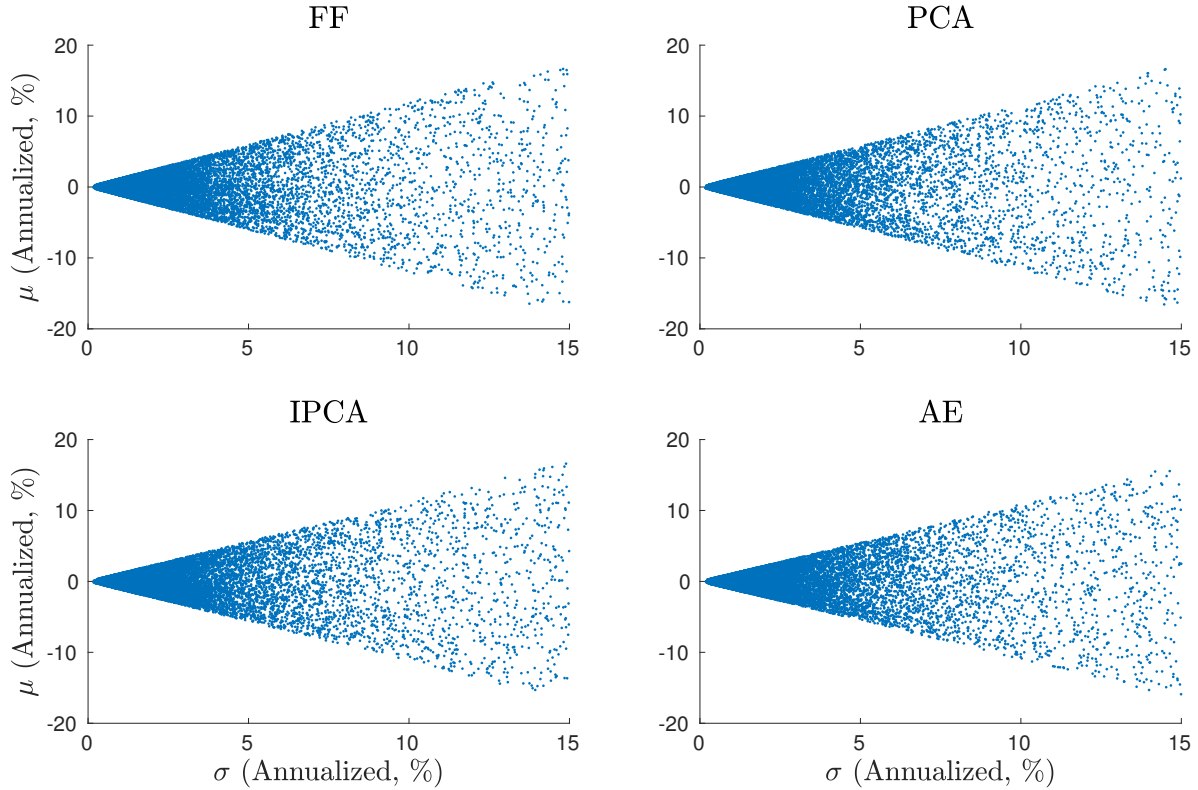
associated with the true risk-free rate.

### 3.4. Measuring Model Misspecification

Section 2.4 introduces an empirical measure of factor model misspecification. Specifically, I construct the optimal zero-investment, zero-beta portfolios implied by each factor model using the analytical solution in Equation (12). Portfolios are formed both in-sample and out-of-sample to avoid full-sample overfitting and look-ahead bias. The Sharpe ratios of these strategies serve as quantitative measures of model misspecification: the higher the attainable Sharpe ratio from a factor-neutral portfolio, the greater the potential profits from exploiting model-implied pricing errors, and thus, the greater the degree of misspecification.

Table 3 reports the annualized maximum Sharpe ratios of zero-investment, zero-beta portfolios implied by the FF, PCA, IPCA, and AE models with 1, 3, 6, and 9 factors. The first four columns show that in-sample Sharpe ratios are highly positive, ranging from 3.15 to 3.31 on an annualized basis. Since a correctly specified factor model should imply a zero Sharpe ratio for such portfolios, these high values indicate that all models considered in this paper are substantially misspecified in-sample. The last four columns of Table 3 examine the degree of misspecification out-of-sample. Out-of-sample Sharpe ratios are markedly lower than their in-sample counterparts—a typical manifestation of full-sample overfitting and estimation error—yet they remain economically large, ranging from 0.98 to 1.28. Overall, the evidence strongly suggests that prominent factor models, including those based on advanced machine learning methods (IPCA and AE), exhibit significant misspecification, as it remains highly profitable to invest in zero-beta portfolios with no exposure to systematic factor risks defined in given models. A comparison across models with varying numbers of factors

Figure. 5. Out-of-Sample Zero-Investment, Zero-Beta Portfolios (Six Factors)



*Notes:* This figure shows the out-of-sample, zero-investment zero-beta portfolios in the mean-standard deviation diagram for the FF, PCA, IPCA, and AE models with six factors.

suggests that increasing the number of factors from one to nine yields little improvement, as the corresponding Sharpe ratio reduction is small—even the nine-factor models imply an annualized Sharpe ratio for zero-investment, zero-beta portfolios around 1.

Figure 5 visualizes the misspecification measure by plotting the out-of-sample, zero-investment zero-beta portfolios in the mean-standard deviation diagram for the FF, PCA, IPCA, and AE models with six factors. I construct these portfolios for each factor model in the following steps. First, in each expanding estimation window, I compute the null space of estimated betas<sup>17</sup>. Second, I randomly simulate 50,000 linear combinations of basis vectors around the zero-vector in the beta null space. This makes sure that the simulated portfolios are not too far away from the origin. Next, I apply these beta-neutral weights to portfolio returns in out-of-sample periods. Finally, I evaluate the mean and standard deviation of these portfolios ex post. Consistent with our analytical results in Figure 3, zero-investment zero-beta portfolios fall into a triangular cone area with line zero-beta frontiers. Visible

<sup>17</sup>The null space of betas is characterized by a set of basis vectors. Any linear combination of the basis vectors will be zero-beta by definition.

deviations of these cones from the horizontal zero line indicate model misspecification, since correctly specified models should yield zero-beta portfolio mean returns equal to zero. The slope of these zero-beta frontiers correspond to the out-of-sample maximum Sharpe ratio values reported in Table 3.

In summary, the empirical assessment of factor model misspecification yields two key findings for the empirical asset pricing literature. First, prominent factor models exhibit economically large pricing errors that can be profitably exploited through investment strategies, indicating substantial model misspecification. Second, neither the adoption of advanced machine learning techniques nor the inclusion of additional factors (up to nine in this analysis) materially alleviates this misspecification.

### 3.5. *Inspecting the Robustness of Zero-Beta Rate Estimation*


Section 2.3 shows that factor model misspecification tends to introduce biases to the estimated zero-beta rate obtained from the unit-investment, minimum-variance zero-beta portfolio. Greater misspecification may push the estimated zero-beta rate toward the return on the global minimum-variance (GMV) portfolio,  $r_{GMV}$ . Building on this insight, I hypothesize that the observed robustness of zero-beta rate estimates across different factor models arises because these models exhibit similarly large degrees of misspecification, which causes the estimated zero-beta rate to appear close to the mean return of the GMV portfolio.

To inspect the underlying source of robustness, I examine two sets of characteristic-sorted portfolios that differ in their GMV portfolio returns. Specifically, I rank the 273 characteristic-sorted portfolios by their return variances and divide them into two groups of test assets. The first asset group contains the 136 high-variance characteristic-sorted portfolios, while the second group includes the 137 low-variance portfolios. I construct the GMV portfolios for both groups in-sample and out-of-sample. The analytical portfolio weights for the GMV portfolio are given by  $\Sigma^{-1}\boldsymbol{\iota}/\boldsymbol{\iota}'\Sigma^{-1}\boldsymbol{\iota}$ . For the in-sample construction, the full-sample estimate of  $\Sigma$  is used. For the out-of-sample construction, I estimate  $\Sigma$  using expanding windows and apply the resulting weights to the following month's returns. The resulting out-of-sample mean returns of the GMV portfolio are 12.3% and 10.1% for the two asset groups, respectively.<sup>18</sup> This partition enables an examination of whether zero-beta rate estimates differ systematically across asset universes characterized by distinct GMV portfolio returns.

---

<sup>18</sup>Although the literature typically reports full-sample zero-beta rate estimates, I primarily focus on the zero-beta rate implied by out-of-sample, unit-investment zero-beta portfolios. This approach provides a more realistic measure of the zero-beta return that investors could feasibly earn in practice. In-sample estimates are also reported in Appendix D.2.

Figure. 6. Out-of-Sample Zero-Beta Rate across Different Asset Universes



*Notes:* This figure shows the out-of-sample estimated zero-beta rates obtained from unit-investment, minimum-variance zero-beta portfolios across two asset universes that differ in their GMV portfolio returns. In the first asset group (black dashed line), which includes the full set of 273 characteristic-sorted portfolios, the mean GMV portfolio return is lower (10.1%). In the second asset group (blue dashed line), consisting of the 136 high-variance portfolios, the mean GMV portfolio return is higher (12.3%). Four classes of factor models—FF, PCA, IPCA, and AE—with 1, 3, 6, and 9 factors are analyzed. The estimated zero-beta rates are represented by circles, squares, diamonds, and triangles, respectively.

Figure 6 shows the out-of-sample estimated zero-beta rates obtained from unit-investment, minimum-variance zero-beta portfolios across two asset universes that differ in their GMV portfolio returns. In the first asset group (blue markers), which includes the 136 high-variance portfolios, the mean GMV portfolio return is 12.3% shown by the blue dashed line. In the second asset group (black markers), consisting of the 137 low-variance portfolios, the mean GMV portfolio return is 10.1% shown by the black dashed line. I emphasize the following observations. First, across both asset universes, the estimated zero-beta rates appear robust to the choice of factor model and are slightly increasing in the number of factors. The literature tends to interpret the estimation robustness as evidence that these estimates capture the true, unobserved risk-free rate. If that were the case, the zero-beta rates should be similar across different asset universes. However, the second observation is that the estimated zero-beta rates in the high-GMV universe are systematically higher than those in the low-GMV universe. This is inconsistent with the risk-free rate interpretation. Moreover,

the estimates are indeed closely aligned with the mean returns of GMV in their respective universes. This pattern suggests that the estimated zero-beta rates may primarily reflect the mean return of the GMV portfolio rather than the true risk-free rate, providing empirical support for my analytical hypothesis that large model misspecification biases zero-beta rate estimates towards a single high benchmark—the mean return of the GMV portfolio.

### 3.6. *Simulation Analysis*

I further analyze the relationship between the estimated zero-beta rate and factor model misspecification through simulation exercises. The purpose of these simulations is twofold. First, they demonstrate that when a factor model is correctly specified, it can recover the true, unobserved risk-free rate by estimating the expected return of the unit-investment, zero-beta portfolio. This result confirms that the estimation procedures themselves are statistically valid—there is nothing inherently wrong with the methods. Second, the simulations shed lights on quantifying the extent to which model misspecification can bias zero-beta rate estimates upward.

I conduct two distinct sets of simulations. The first simulates stock returns based on a return-generating process implied by a factor model, while the second directly calibrates and simulates the parameters of the mean–variance frontier (MVF). Both exercises are designed to evaluate how model misspecification distorts the estimated zero-beta rate. In the simulated-return framework, model misspecification is interpreted as omitted risk factors. Given any misspecified model, I estimate the zero-beta rate and the maximum Sharpe ratio of the zero-investment, zero-beta portfolio following the same empirical procedures described earlier. In contrast, in the MVF-based simulation, model misspecification is interpreted as portfolio inefficiency, which allows for analytical expressions of both the zero-beta rates and the maximum Sharpe ratios.

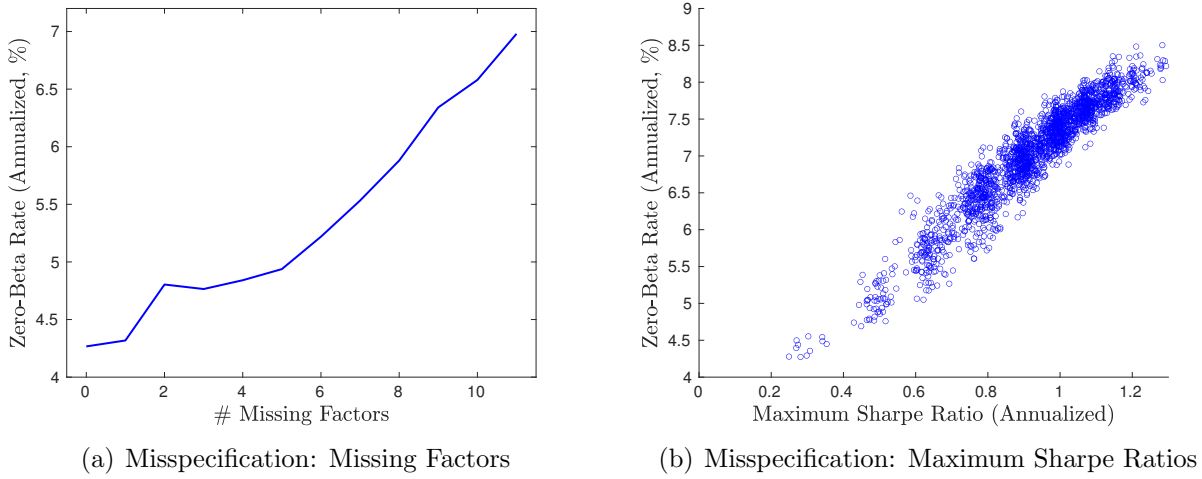
#### 3.6.1. *Simulating Stock Returns*

Following standard practice of simulation in the literature, I assume a return-generating process based on a factor model. Specifically, I simulate stock returns according to a “true” 12-factor model that includes the market, size, value, profitability, investment, momentum, betting-against-beta (BAB), management-based mispricing, performance-based mispricing, idiosyncratic volatility, liquidity and quality factors<sup>19</sup>. Factor construction follows the procedures described in Section 3.1. In this artificial economy, these 12 factors represent the

---

<sup>19</sup>The 12 factors are constructed using characteristics “market\_equity”, “be\_me”, “ope\_me”, “at\_gr1”, “ret\_12\_1”, “betabab\_1260d”, “mispricing\_mgmt”, “mispricing\_perf”, “ivol\_ff3\_21d”, “aliq\_at”, and “qmj”

Figure. 7. Simulating Stock Returns



*Notes:* This figure illustrates the relationship between zero-beta rate estimation and factor model misspecification in a simulation setting where stock returns are generated from a 12-factor model. Panel (a) quantifies misspecification by the number of omitted risk factors and plots the median estimated zero-beta rate against the number of missing factors. Panel (b) quantifies misspecification by the maximum Sharpe ratio attainable from zero-investment, zero-beta portfolios, showing the pairs of the estimated zero-beta rates and the corresponding Sharpe ratios.

complete set of priced risks. Using time-series regressions, I estimate factor loadings ( $\beta_i$ ) and residual volatilities ( $\sigma_i$ ) for each of 273 characteristic-sorted portfolio over the 780-month period from January 1960 to December 2024. I treat these estimated betas as the true risk exposures and resample factor return series using the empirical means and standard deviations of the 12 factors. Portfolio returns are then simulated using the true betas, resampled factor realizations, and calibrated residual volatilities according to the following process:  $R_{i,t} = r_f + \beta_i' F_t + \sigma_i \varepsilon_{i,t}$ , where  $\varepsilon_{i,t}$  is drawn from a standard normal distribution. The true risk-free rate,  $r_f$ , is assumed constant and set equal to the average 1-month Treasury bill yield of 4.27% (annualized) over the 780-month period. This procedure yields simulated returns for 273 portfolios over 780 months.

Suppose an econometrician is unaware of the true 12-factor return structure and attempts to estimate the zero-beta rate using an incomplete—and therefore misspecified—factor model. I start with measuring the degree of model misspecification by the number of omitted factors. Suppose the market factor is always included in the model. Among the remaining 11 factors, there are  $\binom{11}{1} = 11$  possible ways to omit one factor,  $\binom{11}{2} = 55$  possible ways to omit two factors, and so forth. For each group of incomplete models with the same number of missing factors, I estimate the zero-beta rates and record the median value (results are similar when using the mean). Panel (a) of Figure 7 plots the median estimated zero-beta rate against the

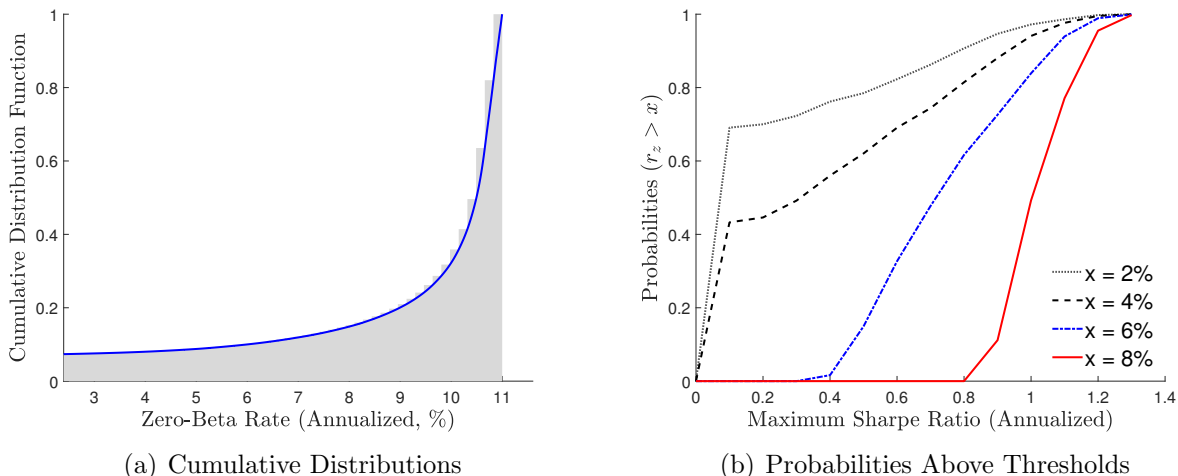
number of omitted factors. The upward-sloping curve provides clear evidence that greater model misspecification biases zero-beta rate estimates upward. When no factors are omitted (the correctly specified model), the estimated zero-beta rate successfully recovers the true, unobserved risk-free rate.

Connecting to earlier analysis, I also compute the measure of model misspecification proposed in this paper—the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios. This metric quantifies the magnitude of investment opportunities left unexplained by a misspecified model. For each misspecified specification, I obtain a pair consisting of the estimated zero-beta rate and the corresponding maximum Sharpe ratio. Panel (b) of Figure 7 displays the scatter plot of these pairs, revealing a strong positive relationship between model misspecification and zero-beta rate estimates. The tight, upward-sloping pattern indicates that as the degree of misspecification increases (i.e., as the beta-neutral Sharpe ratio rises), the estimated zero-beta rate becomes progressively higher, rising from its true value of 4.27% to over 8% when the maximum Sharpe ratio exceeds 1.2 (annualized). This result is consistent with the empirical evidence that most models exhibit an annualized maximum Sharpe ratio between 1.0 and 1.2 (Table 3), while the corresponding estimated zero-beta rates range from approximately 8% to 10% per year (Figure 6).

### 3.6.2. *Simulating Mean-Variance Parameters*

The second simulation abstracts from a specific factor structure to provide a more general analysis. Instead of simulating returns, I directly calibrate the geometric parameters of the mean-variance frontier (MVF) itself. This allows us to assess the statistical distribution of ZBR estimates conditional on the true shape of the investment opportunity set. The mean-variance frontier is determined by three parameters:  $a = \iota' \Sigma \iota$ ,  $b = \iota' \Sigma \mu$ ,  $c = \mu' \Sigma \mu$ . I calibrate  $a$  and  $b$  using the global minimum-variance (GMV) portfolios since  $r_{GMV} = b/a$  and  $\sigma_{GMV}^2 = 1/a$ . Assuming  $r_{GMV} = 11\%$  and  $\sigma_{GMV} = 6.5\%$  (annualized) solves  $a$  and  $b$ . I assume the true risk-free rate ( $r_f$ ) is 3% annually and the return of the true Tangency portfolio ( $r_{p^*}$ ) is 20% annually. Thus, parameter  $c$  is pinned down by  $c = r_f(b - ar_{p^*}) + br_{p^*}$  (proved in Appendix B.1 Equation B.9). In the mean-variance space, a correctly specified model corresponds to the true Tangency portfolio  $p^*$ , and any misspecified model is associated with an inefficient portfolio  $p$ . With the calibrated MVF, I randomly generate a large number of inefficient portfolios  $p$  with  $r_p$  and  $\sigma_p$ . This approach allows us to compute zero-beta rate

Figure. 8. Simulating Mean-Variance Parameters



*Notes:* This figure illustrates the relationship between zero-beta rate estimation and factor model misspecification in a simulation setting where mean-variance frontier parameters are calibrated. Panel (a) plots the Cumulative Distribution Function (CDF) of the estimated zero-beta rates from 100,000 inefficient portfolios  $p$  that are uniformly distributed inside the mean-variance frontier. Panel (b) shows the probability that the estimated zero-beta rate ( $r_z$ ) exceeds a given threshold ( $x$ ), as a function of the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios.

estimates as well as Sharpe ratios using analytical expressions:

$$r_z = r_{GMV} - \sigma_{GMV}^2 \frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} \quad (16)$$

$$S_z = \sqrt{\frac{ac - b^2}{a} \cdot \left(1 - \frac{\sigma_{p^*}^2 - \sigma_{GMV}^2}{\sigma_p^2 - \sigma_{GMV}^2}\right)} \quad (17)$$

where  $S_z$  denotes the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios given a factor model. Proposition 3 shows that this slope equals to the asymptote of the unit-investment, zero-beta frontier. (Appendix B.7 proves Equation 17).

Figure D.2 illustrates the procedure of randomly generating inefficient portfolios. Each inefficient portfolio (misspecified model) produces a pair of the zero-beta rate estimate and the maximum Sharpe ratios of zero-investment, zero-beta portfolios. Assuming that the inefficient portfolio is uniformly distributed inside the mean-variance frontier, Figure 8 Panel (a) plots the Cumulative Distribution Function (CDF) of the estimated zero-beta rates from 100,000 inefficient portfolios  $p$ . The distribution of zero-beta rates is clustered visibly in areas with high zero-beta rate estimates (close to the return of the GMV portfolio). The CDF's shape—rising slowly at first and then accelerating steeply at high values—indicates that

a large mass of the probability distribution is concentrated at high zero-beta rate values. Estimates of 10-11% ( $\approx r_{GMV}$ ) are common in calibrated true world with low risk-free rate. Panel (b) explicitly links this phenomenon to the maximum Sharpe ratio. It shows the probability that the estimated zero-beta rate ( $r_z$ ) exceeds a given threshold ( $x$ ), as a function of the maximum Sharpe ratio of the economy. The results are striking. When we empirically observe an annualized maximum Sharpe ratios of 1.2 (indicating a high degree of misspecification), the probability of the zero-beta rate estimate being greater than 8% (the red line) is over 80%. Conversely, in a low-misspecification economy with a Sharpe ratio of 0.4, the probability of the ZBR exceeding 8% is effectively zero.

Taken together, these two simulation exercises provide robust evidence that high estimates of the zero-beta rate are likely to be a direct consequence of factor model misspecification rather than a truly high risk-free rate.

#### 4. Equity Risk Premium Puzzle and Risk-Free Rate

Previous sections have shown, both theoretically and empirically, that even in a low risk-free rate environment, it is highly likely to obtain a high estimated zero-beta rate—often close to the mean return of the global minimum-variance (GMV) portfolio—when the underlying factor model is substantially misspecified. the zero-beta rate implied by a factor model provides little definitive information about the true risk-free rate, and hence offers limited insight into the magnitude of equity risk premium or convenience yield. In this section, I turn to a conceptually complementary question: could the true risk-free rate itself be high?

I argue that a high risk-free rate intensifies the equity risk premium puzzle, thereby creating greater tension with structural models in macro-finance. At first glance, it is tempting to assume that a high risk-free rate necessarily implies a low equity risk premium. This reasoning clearly applies to the market risk premium due to a negative mechanical relationship between the risk-free rate and the market excess return. The market portfolio itself, however, does not coincide with the true tangency portfolio—it may not even lie close to the tangency portfolio in the mean–variance space, particularly given the massive misspecification of the CAPM. When considering the risk premium on the true tangency portfolio, the relationship in fact reverses: a higher risk-free rate can imply a larger equilibrium equity risk premium, rather than a smaller one. To formalize this idea, I start from deriving an analytical relationship linking the risk-free rate  $r_f$ , the expected return on the GMV portfolio  $r_{GMV}$ , and the expected return on the (unobserved) tangency portfolio  $r_p^*$  (The complete proof is provided in Appendix B.4).

**Proposition 4.** *In a mean-variance framework, let  $a = \mathbf{1}'\Sigma\mathbf{1}$ ,  $b = \mathbf{1}'\Sigma\boldsymbol{\mu}$ ,  $c = \boldsymbol{\mu}'\Sigma\boldsymbol{\mu}$ . Denote the Tangency portfolio as  $p^*$ , risk-free rate as  $r_f$ , expected return and volatility of the GMV portfolio as  $r_{GMV}$  and  $\sigma_{GMV}$ , respectively. Assume that  $r_{GMV} > r_f$ . The term  $(ac - b^2)/a$  represents the squared slope of the asymptote of the mean-variance frontier and can be interpreted as the maximum Sharpe ratio attainable by all zero-investment portfolios. If there exists a lower bound such that  $(ac - b^2)/a \geq L^2$ , then*

$$r_{p^*} - r_f \geq r_{GMV} - r_f + \frac{\sigma_{GMV}^2 L^2}{r_{GMV} - r_f} \geq 2\sigma_{GMV}L \quad (18)$$

Equation (18) encapsulates some fundamental tensions in asset pricing. First, consider the function  $f(x) = x + C/x$ , where  $x = r_{GMV} - r_f$  and  $C = \sigma_{GMV}^2 L^2$ . This function is convex (U-shaped) and has its minimum at  $x = \sqrt{C} = \sigma_{GMV}L$ . This defines the minimum possible tangency portfolio risk premium ( $r_p^* - r_f \geq 2\sigma_{GMV}L$ ). Interestingly, the lower bound for tangency portfolio risk premium does not depend on the risk-free rate. If the investment opportunities implied by  $L$  are economically meaningful, then the tangency portfolio risk premium cannot be small. For instance, if  $L \approx 0.5$ ,  $\sigma_{GMV} \approx 7\%$ , then  $r_p^* - r_f \geq 7\%$ . Since  $L$  represents the maximum Sharpe ratio attainable by all zero-investment portfolios, it should be bounded below by the maximum Sharpe ratio of zero-investment, zero-beta portfolios. Empirical evidence in the literature indicates that such Sharpe ratio opportunities may exceed 0.5 after accounting for different types of implementation costs, implying even larger tangency portfolio risk premia. Such high risk premia imply high prices of risk that are difficult to reconcile within standard macro-finance structural models. Second, if we accept a high risk-free rate that is close to  $r_{GMV}$ , the GMV portfolio premium  $x = r_{GMV} - r_f$  becomes very small. As  $x$  shrinks, the term  $C/x$  increases, generating even larger tangency portfolio risk premia  $r_p^* - r_f$ .

In summary, higher risk-free rate implies larger equity risk premium for researchers to explain. In other words, a high risk-free rate magnifies the equity risk premium puzzle, rather than resolving it.

## 5. Conclusion

This paper revisits the long-standing zero-beta rate puzzle through the lens of factor model misspecification. I show that the pervasive finding of high estimated zero-beta rates across a broad range of models need not reflect a high unobserved risk-free rate, but may instead arise from systematic model inefficiencies. Theoretically, when the factor model is misspecified, the zero-beta rate is no longer uniquely determined within the model. As a

result, the common practice of focusing on the minimum-variance zero-beta portfolio induces a systematic bias: as the degree of misspecification increases, the estimated zero-beta rate converges to the mean return of the global minimum-variance portfolio. This convergence toward a single benchmark level helps explain the persistently elevated estimates observed in the data.

To quantify this mechanism, I introduce a new measure of model misspecification based on the maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios. This measure provides a direct, investment-based link between statistical misspecification and economic inefficiency. Empirical analysis using a comprehensive cross-section of characteristic-sorted portfolios shows that all major factor models—including machine-learning-based ones—remain substantially misspecified, with zero-beta portfolios delivering Sharpe ratios exceeding one even out of sample. Two simulation exercises confirm that such degrees of misspecification are sufficient to fully reproduce the empirically observed high zero-beta rates.

Overall, this study reframes the zero-beta rate puzzle from a mystery of financial equilibrium into a measurable outcome of factor model misspecification. What appears as a stable empirical fact—the persistently high zero-beta rate—is, in fact, the byproduct of a shared structural flaw in factor models. The results caution against using model-implied zero-beta rates to infer fundamental quantities such as the equity risk premium or convenience yield.

## References

- Acharya, V. V., Laarits, T., 2025. When do treasuries earn the convenience yield?: A hedging perspective. Tech. rep., Forthcoming, Journal of Finance.
- Bansal, R., Coleman, W. J., 1996. A monetary explanation of the equity premium, term premium, and risk-free rate puzzles. *Journal of political Economy* 104, 1135–1171.
- Beaulieu, M.-C., Dufour, J.-M., Khalaf, L., 2013. Identification-robust estimation and testing of the zero-beta capm. *Review of Economic Studies* 80, 892–924.
- Beaulieu, M.-C., Dufour, J.-M., Khalaf, L., 2025. Identification-robust and simultaneous inference in multifactor asset pricing models. *Journal of Econometrics* 248, 105915.
- Beaulieu, M.-C., Dufour, J.-M., Khalaf, L., Melin, O., 2023. Identification-robust beta pricing, spanning, mimicking portfolios, and the benchmark neutrality of catastrophe bonds. *Journal of Econometrics* 236, 105464.
- Black, F., 1972. Capital market equilibrium with restricted borrowing. *The Journal of business* 45, 444–455.
- Black, F., Jensen, M. C., Scholes, M., 1972. The capital asset pricing model: Some empirical tests. *Studies in the Theory of Capital Markets* .
- Carhart, M. M., 1997. On persistence in mutual fund performance. *The Journal of finance* 52, 57–82.
- Chamberlain, G., Rothschild, M., 1983. Arbitrage, factor structure, and mean-variance analysis on large asset markets.
- Chen, L., Pelger, M., Zhu, J., 2024. Deep learning in asset pricing. *Management Science* 70, 714–750.
- Cieslak, A., Li, W., Pflueger, C. E., 2025. Inflation and treasury convenience. USC Marshall School of Business Research Paper Sponsored by iORB .
- Cochrane, J. H., 2009. *Asset pricing: Revised edition*. Princeton university press.
- Connor, G., 1984. A unified beta pricing theory. *Journal of Economic Theory* 34, 13–31.
- Connor, G., Korajczyk, R. A., 1986. Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of financial economics* 15, 373–394.

- Connor, G., Korajczyk, R. A., 1988. Risk and return in an equilibrium apt: Application of a new test methodology. *Journal of financial economics* 21, 255–289.
- Cooper, I., Ma, L., Maio, P., Philip, D., 2021. Multifactor models and their consistency with the apt. *The Review of Asset Pricing Studies* 11, 402–444.
- Di Tella, S., Hébert, B., Kurlat, P., Wang, Q., 2026. The zero-beta interest rate. *Journal of political Economy* 134, 000–000.
- Fama, E. F., French, K. R., 1993. Common risk factors in the returns on stocks and bonds. *Journal of financial economics* 33, 3–56.
- Fama, E. F., French, K. R., 2015. Incremental variables and the investment opportunity set. *Journal of Financial Economics* 117, 470–488.
- Fama, E. F., French, K. R., 2018. Choosing factors. *Journal of financial economics* 128, 234–252.
- Fama, E. F., MacBeth, J. D., 1973. Risk, return, and equilibrium: Empirical tests. *Journal of political economy* 81, 607–636.
- Feng, G., He, J., Polson, N. G., Xu, J., 2024. Deep learning in characteristics-sorted factor models. *Journal of Financial and Quantitative Analysis* 59, 3001–3036.
- Ferson, W. E., Harvey, C. R., 1991. The variation of economic risk premiums. *Journal of political economy* 99, 385–415.
- Ferson, W. E., Siegel, A. F., Wang, J. L., 2025. Factor model comparisons with conditioning information. *Journal of Financial and Quantitative Analysis* 60, 1401–1426.
- Frazzini, A., Pedersen, L. H., 2014. Betting against beta. *Journal of financial economics* 111, 1–25.
- Gagliardini, P., Ossola, E., Scaillet, O., 2016. Time-varying risk premium in large cross-sectional equity data sets. *Econometrica* 84, 985–1046.
- Gibbons, M. R., 1982. Multivariate tests of financial models: A new approach. *Journal of financial economics* 10, 3–27.
- Gibbons, M. R., Ross, S. A., Shanken, J., 1989. A test of the efficiency of a given portfolio. *Econometrica: Journal of the Econometric Society* pp. 1121–1152.

- Giglio, S., Xiu, D., 2021. Asset pricing with omitted factors. *Journal of Political Economy* 129, 1947–1990.
- Gospodinov, N., Kan, R., Robotti, C., 2014. Misspecification-robust inference in linear asset-pricing models with irrelevant risk factors. *The Review of Financial Studies* 27, 2139–2170.
- Gospodinov, N., Kan, R., Robotti, C., 2017. Spurious inference in reduced-rank asset-pricing models. *Econometrica* 85, 1613–1628.
- Gu, S., Kelly, B., Xiu, D., 2020. Empirical asset pricing via machine learning. *The Review of Financial Studies* 33, 2223–2273.
- Gu, S., Kelly, B., Xiu, D., 2021. Autoencoder asset pricing models. *Journal of Econometrics* 222, 429–450.
- Hansen, L. P., Jagannathan, R., 1991. Implications of security market data for models of dynamic economies. *Journal of political economy* 99, 225–262.
- Hou, K., Xue, C., Zhang, L., 2015. Digesting anomalies: An investment approach. *The Review of Financial Studies* 28, 650–705.
- Huberman, G., 1982. A simple approach to arbitrage pricing theory. *Journal of Economic Theory* 28, 183–191.
- Ingersoll Jr, J. E., 1984. Some results in the theory of arbitrage pricing. *The Journal of Finance* 39, 1021–1039.
- Jensen, T. I., Kelly, B., Pedersen, L. H., 2023. Is there a replication crisis in finance? *The Journal of Finance* 78, 2465–2518.
- Kan, R., Robotti, C., Shanken, J., 2013. Pricing model performance and the two-pass cross-sectional regression methodology. *The Journal of Finance* 68, 2617–2649.
- Kan, R., Zhang, C., 1999a. Gmm tests of stochastic discount factor models with useless factors. *Journal of Financial Economics* 54, 103–127.
- Kan, R., Zhang, C., 1999b. Two-pass tests of asset pricing models with useless factors. *the Journal of Finance* 54, 203–235.
- Kandel, S., 1984. The likelihood ratio test statistic of mean-variance efficiency without a riskless asset. *Journal of Financial Economics* 13, 575–592.

- Kandel, S., 1986. The geometry of the maximum likelihood estimator of the zero-beta return. *The Journal of Finance* 41, 339–346.
- Kelly, B. T., Pruitt, S., Su, Y., 2019. Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134, 501–524.
- Kim, S., Korajczyk, R. A., Neuhierl, A., 2021. Arbitrage portfolios. *The Review of Financial Studies* 34, 2813–2856.
- Kleibergen, F., Zhan, Z., 2020. Robust inference for consumption-based asset pricing. *The Journal of Finance* 75, 507–550.
- Krishnamurthy, A., Vissing-Jorgensen, A., 2012. The aggregate demand for treasury debt. *Journal of Political Economy* 120, 233–267.
- Lettau, M., Pelger, M., 2020. Factors that fit the time series and cross-section of stock returns. *The Review of Financial Studies* 33, 2274–2325.
- Long, J. B., 1971. Notes on the black valuation model for risky securities. Unpublished manuscript .
- Lopez-Lira, A., Roussanov, N. L., 2020. Do common factors really explain the cross-section of stock returns? Jacobs Levy Equity Management Center for Quantitative Financial Research Paper .
- Manresa, E., Peñaranda, F., Sentana, E., 2023. Empirical evaluation of overspecified asset pricing models. *Journal of Financial Economics* 147, 338–351.
- Morgan, I. G., 1975. Prediction of return with the minimum variance zero-beta portfolio. *Journal of Financial Economics* 2, 361–376.
- Nagel, S., 2016. The liquidity premium of near-money assets. *The Quarterly Journal of Economics* 131, 1927–1971.
- Roll, R., 1980. Orthogonal portfolios. *Journal of Financial and Quantitative analysis* 15, 1005–1023.
- Ross, S. A., 1976. The arbitrage theory of capital asset pricing. In: *Journal of Economic Theory*, Elsevier, pp. 341–360.
- Shanken, J., 1986. Testing portfolio efficiency when the zero-beta rate is unknown: a note. *The Journal of Finance* 41, 269–276.

Stambaugh, R. F., Yuan, Y., 2017. Mispricing factors. *The review of financial studies* 30, 1270–1315.

Velu, R., Zhou, G., 1999. Testing multi-beta asset pricing models. *Journal of Empirical Finance* 6, 219–241.

## Appendix A. Data

### A.1. Stock Characteristics

Table A.1: Stock Characteristics

Name	Description	Paper
age	Firm age (since listing), measured in months.	Jensen, Kelly and Pedersen (2023).
aliq_at	Ortiz-Molina liquidity measure scaled by assets ( $ALIQ/AT$ ).	Ortiz-Molina and Phillips (2014).
aliq_mat	Ortiz-Molina liquidity measure scaled by market assets ( $ALIQ/MAT$ ).	Ortiz-Molina and Phillips (2014).
ami_126d	Amihud illiquidity (average $ R /VOL$ ) over 126 days.	Amihud (2002).
at_be	Assets-to-book equity ( $AT/BE$ ).	Fama and French lineage (shown in JKP).
at_gr1	1-year growth in total assets ( $AT_t/AT_{t-12} - 1$ ).	JKP construction.
at_me	Assets-to-market equity ( $AT/ME$ ).	JKP construction.
at_turnover	Asset turnover ( $SALE/AT$ ).	JKP construction.
be_gr1a	1-year change in book equity scaled by assets ( $(BE_t - BE_{t-12})/AT_t$ ).	JKP construction.
be_me	Book-to-market equity ( $BE/ME$ ).	Rosenberg, Reid and Lanstein (1985).
beta_60m	CAPM beta estimated over 60 months.	Fama-MacBeth / CAPM estimates (JKP).
beta_dimson_21d	Dimson-style beta (21-day window with lead/lag market adjustments).	Dimson (1979) style (JKP).
betabab_1260d	Betting-against-beta metric (long low-beta, short high-beta) over 1260 days.	Frazzini and Pedersen (2014) family.
betadown_252d	Downside beta estimated over 252 days (restricted to days market return negative).	Ang, Chen and Xing (2006).
bev_mev	Book enterprise value to market enterprise value ( $BEV/MEV$ ).	Penman, Richardson and Tuna (2007).
bidaskhl_21d	Bid-ask high-low spread estimator over 21 days (Corwin and Schultz method).	Corwin and Schultz (2012).
capex_abn	Abnormal capital expenditures (deviation from expected CAPX).	Titman, Wei and Xie (2004).
capx_gr1	1-year growth in capital expenditures ( $CAPX_t/CAPX_{t-12} - 1$ ).	JKP construction.
capx_gr2	2-year growth in capital expenditures ( $CAPX_t/CAPX_{t-24} - 1$ ).	JKP construction.

Continued on next page

Table A.1 – continued from previous page

Column 1	Column 2	Column 1
capx_gr3	3-year growth in capital expenditures ( $CAPX_t/CAPX_{t-36} - 1$ ).	JKP construction.
cash_at	Cash and short-term investments scaled by assets ( $CHE/AT$ ).	Palazzo (2012) and JKP.
chcsho_12m	Net stock issues / change in shares over 12 months (CHC-SHO 12m).	Pontiff and Woodgate (2008).
coa_gr1a	Change in current operating assets 1-year scaled by assets.	JKP construction (current operating assets family).
col_gr1a	Change in current operating liabilities 1-year scaled by assets.	JKP construction.
cop_at	Cash from operations scaled by assets ( $COP/AT$ ).	JKP construction / cash-flow measures.
cop_atl1	Lagged cash-from-operations scaled by lagged assets ( $COP/AT_{t-12}$ ).	JKP construction.
corr_1260d	Correlation of stock excess returns with market over 1260 days.	JKP construction.
coskew_21d	Co-skewness with market over 21 days (co-skew measure).	JKP construction (skewness family).
cowc_gr1a	Change in current operating working capital 1-year.	Richardson, Sloan, Soliman and Tuna (2005).
dbnetis_at	Net debt issuance scaled by assets ( $DBNETIS/AT$ ).	JKP construction (issuance family).
debt_gr3	Growth in book debt over 3 years ( $DLTT$ change over 3 years).	Lyandres, Sun and Zhang (2008).
debt_me	Debt scaled by market equity ( $DEBT/ME$ ).	Bhandari / JKP family.
dgp_dsale	Change in gross profit minus change in sales ( $\Delta GP - \Delta SALE$ ).	Abarbanell and Bushee (1998) lineage.
div12m_me	Dividend yield over 12 months ( $DIV_{12m}/ME$ ).	Litzenberger and Ramaswamy (1979).
dolvol_126d	Average dollar trading volume over 126 days ( $\overline{DOLVOL}_{126d}$ ).	Chordia, Subrahmanyam and Anshuman (2001).
dolvol_var_126d	Variability (std) of dollar volume over 126 days.	Chordia et al. (2001).
dsale_dinv	Change in sales minus change in inventory ( $\Delta SALE - \Delta INV$ ).	Abarbanell and Bushee (1998).
dsale_drec	Change in sales minus change in receivables ( $\Delta SALE - \Delta REC$ ).	Abarbanell and Bushee (1998).
dsale_dsga	Change in sales minus change in SG&A ( $\Delta SALE - \Delta XSGA$ ).	Abarbanell and Bushee (1998).

Continued on next page

Table A.1 – continued from previous page

Column 1	Column 2	Column 1
earnings_variability	Variability (volatility) of earnings (NI) across periods (e.g., std of NI).	Earnings volatility literature (JKP cites relevant sources).
ebit_bev	EBIT scaled by book enterprise value ( $EBIT/BEV$ ) — operating profitability measure.	Soliman (2008) / JKP.
ebit_sale	EBIT scaled by sales ( $EBIT/SALE$ ) — profit margin.	Soliman (2008).
ebitda_mev	EBITDA scaled by market enterprise value ( $EBITDA/MEV$ ).	Loughran and Wellman (2011) (profitability families).
emp_gr1	Employment (employees) 1-year growth ( $EMP_t/EMP_{t-12} - 1$ ).	JKP construction (labor efficiency family).
eq_dur	Equity duration (duration-like measure of equity cash flows).	Dechow, Sloan and Soliman (2004).
eqnpo_12m	Net equity payout over 12 months (EQNPO 12m).	Daniel and Titman (2006); Boudoukh et al. (2007).
f_score	Piotroski F-score (composite score 0–9 from fundamentals).	Piotroski (2000).
fcf_me	Free cash flow scaled by market equity ( $FCF/ME$ ).	Lakonishok, Shleifer and Vishny (1994).
fnl_gr1a	Change in financial liabilities 1-year scaled by assets.	JKP construction (financial-liabilities family).
gp_at	Gross profit scaled by assets ( $GP/AT$ ) — gross profitability.	Novy-Marx (2013).
gp_at1l	Lagged gross profit scaled by lagged assets ( $GP_{t-1}/AT_{t-12}$ ).	JKP construction.
inv_gr1	Inventory 1-year growth ( $INV_t/INV_{t-12} - 1$ ).	JKP construction.
inv_gr1a	Change in investment/inventory 1-year scaled by assets.	JKP construction.
iskew_capm_21d	Idiosyncratic skewness from CAPM residuals over 21 days.	Bali, Engle and Murray (2016).
iskew_ff3_21d	Idiosyncratic skewness from FF3 residuals over 21 days.	Bali, Engle and Murray (2016).
ival_me	Intrinsic value scaled by market equity ( $IVAL/ME$ ) — intrinsic value measure.	Frankel and Lee (1998); JKP notes on scaling.
ivol_capm_21d	Idiosyncratic volatility from CAPM residuals (21 days).	Ang, Hodrick, Xing and Zhang (2006).
ivol_capm_252d	Idiosyncratic volatility from CAPM residuals (252 days).	Ang et al. (2006).
ivol_ff3_21d	Idiosyncratic volatility from FF3 residuals (21 days).	Ang et al. (2006).
kz_index	Kaplan–Zingales (KZ) index of financing constraints (composite).	Lamont, Polk and Saa-Requejo (2001) lineage.

Continued on next page

Table A.1 – continued from previous page

Column 1	Column 2	Column 1
lnoa_gr1a	Change in long-term net operating assets 1-year scaled by assets.	JKP construction.
lti_gr1a	Change in long-term investments 1-year scaled by assets.	JKP construction.
market_equity	Market equity (ME) — price times shares outstanding ( $PRC \times SHARES$ ).	CRSP/Compustat standard.
mispricing_mgmt	Mispricing composite (management-based signals; multi-component).	Stambaugh and Yuan (2017) and JKP.
mispricing_perf	Mispricing composite (performance-based signals).	Stambaugh and Yuan (2017) and JKP.
ncoa_gr1a	Change in non-current operating assets 1-year scaled by assets.	JKP construction.
ncol_gr1a	Change in non-current operating liabilities 1-year scaled by assets.	Richardson et al. (2005) family.
netdebt_me	Net debt scaled by market equity ( $NETDEBT/ME$ ).	JKP construction / Penman et al. (2007) family.
nfna_gr1a	Change in net financial assets 1-year scaled by assets.	JKP construction.
ni_ar1	First-order autocorrelation of net income (earnings persistence AR(1)).	Earnings persistence literature (JKP).
ni_be	Net income scaled by book equity ( $NI/BE$ ) — ROE family.	Haugen and Baker (1996) lineage.
ni_ivol	Idiosyncratic volatility of net income (earnings volatility).	Francis et al. (2004) style measures.
ni_me	Net income scaled by market equity ( $NI/ME$ ) — earnings-to-price family.	Basu (1983) lineage.
nncoa_gr1a	Change in net non-current operating assets 1-year scaled by assets.	JKP construction.
noa_at	Net operating assets scaled by assets ( $NOA/AT$ ).	JKP construction.
noa_gr1a	Change in net operating assets 1-year ( $NOA_t/NOA_{t-12} - 1$ ).	JKP construction.
o_score	Ohlson O-score (bankruptcy/distress probability measure).	Ohlson-style distress measures (JKP references).
oaccruals_at	Operating accruals scaled by assets ( $OACCRUALS/AT$ ).	Sloan (1996) / Richardson et al. (2005) family.
oaccruals_ni	Percent operating accruals (operating accruals scaled by net income).	Hafzalla, Lundholm and Van Winkle (2011).
ocf_at	Operating cash flow scaled by assets ( $OCF/AT$ ).	Bouchaud, Krueger, Landier and Thesmar (2019) cited in JKP.
ocf_at_chg1	Change in operating cash flow to assets over 1 year ( $OCF_t/AT_t - OCF_{t-12}/AT_{t-12}$ ).	JKP construction.

Continued on next page

Table A.1 – continued from previous page

Column 1	Column 2	Column 1
ocf_me	Operating cash flow scaled by market equity ( $OCF/ME$ ).	Desai, Rajgopal and Venkatachalam (2004) family.
op_at	Operating profit scaled by assets ( $OP/AT$ ).	Ball et al. (2015/2016) operating-profit family.
op_atl1	Lagged operating profit scaled by lagged assets ( $OP_{t-1}/AT_{t-12}$ ).	JKP construction.
ope_be	Operating profit scaled by book equity ( $OP/BE$ ).	Fama and French / Ball et al. lineage.
ope_bell	Operating profit scaled by lagged book equity.	JKP construction.
opex_at	Operating expenses scaled by assets ( $OPEX/AT$ ).	Novy-Marx (2011) / JKP.
pi_nix	Earnings before tax and extraordinary items scaled by net income including extraordinary items (PI/NIX).	JKP construction.
ppeinv_gr1a	Change in PPE plus inventory 1-year scaled by assets ( $(PPEINV_t - PPEINV_{t-12})/AT_{t-12}$ ).	JKP construction / investment-change literature.
prc	Stock price (PRC), typically adjusted close.	CRSP/Compustat standard.
prc_highprc_252d	Price relative to 252-day high ( $PRC/\max(PRC_{252d})$ ).	George and Hwang (2004) style measure.
qmj	Quality Minus Junk composite (aggregate of quality signals).	Asness, Frazzini and Pedersen (2019).
qmj_growth	QMJ growth subcomponent (growth-related z-scores).	Asness et al. (2019).
qmj_prof	QMJ profitability subcomponent (profitability z-scores).	Asness et al. (2019).
qmj_safety	QMJ safety subcomponent (safety/z-score measures).	Asness et al. (2019).
rd5_at	R&D scaled to assets (5-year aggregated/averaged) ( $(\overline{R\&D}_5)/AT$ ).	Chan, Lakonishok and Sougiannis (2001) family.
rd_me	R&D scaled by market equity ( $R\&D/ME$ ).	Chan et al. (2001).
rd_sale	R&D scaled by sales ( $R\&D/SALE$ ).	Chan et al. (2001).
resff3_12_1	Residual momentum: residuals from FF3, 12-month horizon, scaled by residual std (JKP variant).	Blitz, Huij and Mertens (2011) adjustments noted in JKP.
resff3_6_1	Residual momentum: residuals from FF3, 6-month horizon.	Blitz, Huij and Mertens (2011).
ret_12_1	Price momentum: cumulative return $t - 12$ to $t - 1$ (12 months).	Jegadeesh and Titman (1993) momentum family.
ret_12_7	Price momentum: cumulative return $t - 12$ to $t - 7$ .	JKP / momentum literature.
ret_1_0	Most recent monthly return ( $R_{t-1 \rightarrow t}$ ).	JKP / return family.

Continued on next page

Table A.1 – continued from previous page

Column 1	Column 2	Column 1
ret_3_1	Cumulative return $t - 3$ to $t - 1$ (3 months).	Momentum / JKP construction.
ret_60_12	Long-horizon momentum/reversal: cumulative return $t - 60$ to $t - 12$ .	Novy-Marx / De Bondt and Thaler lines (JKP).
ret_6_1	Cumulative return $t - 6$ to $t - 1$ (6 months).	Jegadeesh and Titman (1993) family.
ret_9_1	Cumulative return $t - 9$ to $t - 1$ (9 months).	JKP / momentum family.
rmax1_21d	Maximum 1-day return within a 21-day window ( $\max_{1d}$ ).	JKP (Asness et al. style).
rmax5_21d	Maximum 5-day return within a 21-day window ( $\max_{5d}$ ).	Asness et al. (2020) style.
rmax5_rvol_21d	Highest 5-day return scaled by return volatility ( $RMAX5/RVOL$ ) over 21 days.	Asness et al. (2020) and JKP.
rskew_21d	Total return skewness over 21 days.	Bali et al. (2016) family.
rvol_21d	Return volatility (std) over 21 days.	Ang, Engle, and colleagues (JKP references).
sale_bev	Sales scaled by book enterprise value ( $SALE/BEV$ ).	JKP construction.
sale_emp_gr1	Sales per employee growth (1-year) ( $SALE/EMP$ growth).	Abarbanell and Bushee (1998) lineage.
sale_gr1	Sales 1-year growth ( $SALE_t/SALE_{t-12} - 1$ ).	JKP construction.
sale_gr3	Sales 3-year growth ( $SALE_t/SALE_{t-36} - 1$ ).	JKP construction.
sale_me	Sales scaled by market equity ( $SALE/ME$ ).	JKP construction.
seas_11_15an	Annual seasonality: average returns in months $t - 11$ to $t - 15$ (annual lags).	Heston and Sadka (2008).
seas_11_15na	Non-annual seasonality: months $t - 11$ to $t - 15$ non-annual lags.	Heston and Sadka (2008).
seas_16_20an	Annual seasonality: months $t - 16$ to $t - 20$ (annual).	Heston and Sadka (2008).
seas_16_20na	Non-annual seasonality: months $t - 16$ to $t - 20$ .	Heston and Sadka (2008).
seas_1_1an	One-year lagged return (annual seasonality, month $t - 12$ ).	Heston and Sadka (2008).
seas_1_1na	One-year non-annual seasonality (non-annual lag).	Heston and Sadka (2008).
seas_2_5an	Annual seasonality: months $t - 2$ to $t - 5$ .	Heston and Sadka (2008).
seas_2_5na	Non-annual seasonality: months $t - 2$ to $t - 5$ .	Heston and Sadka (2008).
seas_6_10an	Annual seasonality: months $t - 6$ to $t - 10$ .	Heston and Sadka (2008).
seas_6_10na	Non-annual seasonality: months $t - 6$ to $t - 10$ .	Heston and Sadka (2008).
taccruals_at	Total accruals scaled by assets ( $TACCRUALS/AT$ ).	Richardson, Sloan, Soliman and Tuna (2005).
taccruals_ni	Percent total accruals (total accruals scaled by net income).	Hafzalla et al. (2011) and JKP.
tangibility	Asset tangibility measure (PPE and tangible asset share; JKP formula).	Tangibility literature; JKP construction.
tax_gr1a	Tax expense change 1-year scaled by assets (tax surprise).	Thomas and Zhang (2011).
turnover_126d	Share turnover averaged over 126 days ( $\overline{TURN}_{126d}$ ).	Liu (2006).

Continued on next page

Table A.1 – continued from previous page

Column 1	Column 2	Column 1
turnover_var_126d	Volatility of turnover over 126 days (std).	Chordia et al. (2001).
zero_trades_126d	Fraction of zero-trade days over 126 days.	Lesmond et al. (1999).
zero_trades_21d	Fraction of zero-trade days over 21 days.	Lesmond et al. (1999).
zero_trades_252d	Fraction of zero-trade days over 252 days.	Lesmond et al. (1999).

## Appendix B. Proofs

### B.1. Proof of Proposition 1

*Proof.*

Step 1: Mean-variance frontier.

This proof works with the unit-investment frontier. Deriving the analytical expression for the mean–variance frontier repeats the same procedure in Chapter 5 of [Cochrane \(2009\)](#). For a given target return  $r_{p^*}$ , the variance is minimized by solving the following problem:

$$\min_{\omega} \omega' \Sigma \omega \quad s.t. \quad \omega' \iota = 1, \quad \omega' \mu = r_{p^*} \quad (\text{B.1})$$

where  $\mu$  is the asset mean returns,  $\Sigma$  is the variance-covariance matrix, and  $\iota$  is a vector of ones. Set up the Lagrangian:

$$\mathcal{L} = \frac{1}{2} \omega' \Sigma \omega - \lambda_1 \omega' \iota - \lambda_2 (\omega' \mu - r_{p^*}) \quad (\text{B.2})$$

The first-order condition is given by:

$$\Sigma \omega = \lambda_1 \iota + \lambda_2 \mu = \begin{bmatrix} \iota & \mu \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \implies \omega = \Sigma^{-1} \begin{bmatrix} \iota & \mu \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad (\text{B.3})$$

Premultiply equation (B.3) by  $\begin{bmatrix} \iota' \\ \mu' \end{bmatrix}$  we have:

$$\begin{bmatrix} \iota' \\ \mu' \end{bmatrix} \omega = \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} = \begin{bmatrix} \iota' \\ \mu' \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \mu \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad (\text{B.4})$$

Denote

$$\mathbf{A} \equiv \begin{bmatrix} \iota' \\ \mu' \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \mu \end{bmatrix} = \begin{bmatrix} \iota \Sigma^{-1} \iota & \iota \Sigma^{-1} \mu \\ \mu \Sigma^{-1} \iota & \mu \Sigma^{-1} \mu \end{bmatrix} \equiv \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad (\text{B.5})$$

where  $a = \boldsymbol{\iota}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\iota}$ ,  $b = \boldsymbol{\iota}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$ , and  $c = \boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$ .

Thus

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\omega}_{p^*} = \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} \quad (\text{B.6})$$

The variance of this efficient portfolio  $p^*$  is:

$$\begin{aligned} \sigma_{p^*}^2 &= \boldsymbol{\omega}_{p^*}' \boldsymbol{\Sigma} \boldsymbol{\omega}_{p^*} \\ &= \begin{bmatrix} 1 & r_{p^*} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\iota} \\ \boldsymbol{\mu} \end{bmatrix} \boldsymbol{\Sigma}^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} \\ &= \begin{bmatrix} 1 & r_{p^*} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} \\ &= \frac{a}{ac - b^2} (ar_{p^*}^2 - 2br_{p^*} + c) \end{aligned} \quad (\text{B.7})$$

Therefore, the mean-variance frontier corresponds to a parabola.

Step 2: Unit-investment, zero-beta portfolios with respect to an efficient portfolio  $p^*$ .

The covariance between an arbitrary portfolio  $j$  and an efficient portfolio  $p^*$  is given by:

$$\sigma_{j,p^*} = \boldsymbol{\omega}_j' \boldsymbol{\Sigma} \boldsymbol{\omega}_{p^*} = \boldsymbol{\omega}_j' \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} = \begin{bmatrix} 1 & r_j \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_{p^*} \end{bmatrix} = c - br_{p^*} + (ar_{p^*} - b)r_z \quad (\text{B.8})$$

Here, the third equality follows from  $\boldsymbol{\omega}_j' \boldsymbol{\iota} = 1$ , since portfolio  $j$  is assumed to be a unit-investment portfolio. According to equation (B.8), if portfolio  $j$  is zero-beta (i.e., has zero covariance) with respect to the efficient portfolio  $p^*$ , then  $\sigma_{j,p^*} = 0$  implies

$$r_j = r_f = \frac{c - br_{p^*}}{b - ar_{p^*}} \quad (\text{B.9})$$

given that  $r_{p^*} \neq b/a$  where  $b/a$  is the return of the GMV portfolio.<sup>20</sup> Since the efficient Tangent portfolio  $p^*$  is unique, the corresponding zero-beta rate is also unique. Specifically, under a correctly specified factor model, there exists a factor portfolio  $p^*$  on the mean-variance frontier, and all zero-beta portfolios with respect to  $p^*$  must have the same expected return. Thus, the zero-beta rate is uniquely identified and corresponds to the frictionless risk-free

---

<sup>20</sup>Throughout the proof, I maintain the assumption that the efficient portfolio  $p^*$  lies above the GMV portfolio.

rate in the economy. This establishes Proposition 1 (i).

Step 3: Unit-investment, zero-beta portfolios with respect to an arbitrary portfolio  $p$ .

For unit-investment, zero-beta portfolios with respect to an arbitrary portfolio  $p$  targeting an expected return  $r_z$ , we solve the following problem:

$$\min_{\omega} \omega' \Sigma \omega \quad s.t. \quad \omega' \iota = 1, \quad \omega' \mu = r_z, \quad \omega' \Sigma \omega_p = 0 \quad (\text{B.10})$$

Set up the Lagrangian:

$$\mathcal{L} = \frac{1}{2} \omega' \Sigma \omega - \lambda_1 \omega' \iota - \lambda_2 (\omega' \mu - r_z) - \lambda_3 \omega' \Sigma \omega_p \quad (\text{B.11})$$

The first-order condition is given by:

$$\begin{aligned} \Sigma \omega &= \lambda_1 \iota + \lambda_2 \mu + \lambda_3 \Sigma \omega_p = \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} \\ \implies \omega &= \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} \end{aligned} \quad (\text{B.12})$$

Premultiply equation (B.12) by  $\begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix}$  we have:

$$\begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \omega = \begin{bmatrix} 1 \\ r_z \\ 0 \end{bmatrix} = \begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} \quad (\text{B.13})$$

Denote

$$\mathbf{H} \equiv \begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} = \begin{bmatrix} a & b & 1 \\ b & c & r_p \\ 1 & r_p & \sigma_p^2 \end{bmatrix} = \left[ \begin{array}{c|c} A & 1 \\ \hline 1 & r_p \\ \hline 1 & r_p & \sigma_p^2 \end{array} \right] \quad (\text{B.14})$$

Thus

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} = \mathbf{H}^{-1} \begin{bmatrix} 1 \\ r_p \\ 0 \end{bmatrix} \quad \text{and} \quad \omega_z = \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 1 \\ r_z \\ 0 \end{bmatrix} \quad (\text{B.15})$$

where I denote  $\boldsymbol{\omega}_z$  as the weights of a portfolio  $z$  that is orthogonal to  $p$ . The variance of the unit-investment, zero-beta portfolio  $z$  is given by:

$$\begin{aligned}
\sigma_z^2 &= \boldsymbol{\omega}'_z \boldsymbol{\Sigma} \boldsymbol{\omega}_z \\
&= \begin{bmatrix} 1 & r_z & 0 \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\mu}' \\ \boldsymbol{\omega}'_p \boldsymbol{\Sigma} \end{bmatrix} \boldsymbol{\Sigma}^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} & \boldsymbol{\Sigma} \boldsymbol{\omega}_p \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 1 \\ r_z \\ 0 \end{bmatrix} \\
&= \begin{bmatrix} 1 & r_z & 0 \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 1 \\ r_z \\ 0 \end{bmatrix}
\end{aligned} \tag{B.16}$$

To understand the relationship between  $\sigma_z$  and  $r_z$ ,  $\mathbf{H}^{-1}$  needs to be calculated. I start by expressing the Schur complement of  $\mathbf{A}$ :

$$S = \sigma_p^2 - \begin{bmatrix} 1 & r_p \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 1 \\ r_p \end{bmatrix} = \sigma_p^2 - \frac{a}{ac - b^2} (ar_p^2 - 2br_p + c) \tag{B.17}$$

Recall from equation (B.7) that the variance of an efficient portfolio  $p^*$  is  $\sigma_{p^*}^2 = \frac{a}{ac - b^2} (ar_{p^*}^2 - 2br_{p^*} + c)$ . Hence, if we assume the arbitrary portfolio  $p$  has the same expected return as the efficient portfolio  $p^*$ :  $r_p = r_{p^*}$ . Then,  $\frac{a}{ac - b^2} (ar_p^2 - 2br_p + c) = \sigma_{p^*}^2$ , and the Schur complement  $S = \sigma_p^2 - \sigma_{p^*}^2$ , which measures the horizontal distance between portfolio  $p$  and  $p^*$  with the same expected return.

Note that:

$$\begin{aligned}
|\mathbf{H}| &= ac\sigma_p^2 + 2br_p - c - ar_p^2 - b^2\sigma_p^2 \\
&= \frac{1}{\sigma_p^2} \left[ \underbrace{(a\sigma_p^2 - 1)}_{\equiv a_z} \underbrace{(c\sigma_p^2 - r_p^2)}_{\equiv c_z} - \underbrace{(b\sigma_p^2 - r_p)}_{\equiv b_z} \right] \\
&= \frac{1}{\sigma_p^2} (a_z c_z - b_z^2)
\end{aligned} \tag{B.18}$$

Hence, we can compute  $\mathbf{H}^{-1}$ :

$$\mathbf{H}^{-1} = \frac{\sigma_p^2}{a_z c_z - b_z^2} \begin{bmatrix} c_z & -b_z & br_p - c \\ -b_z & a_z & b - ar_p \\ br_p - c & b^2 - ac & ac - b^2 \end{bmatrix} \tag{B.19}$$

Substitute  $\mathbf{H}^{-1}$  into equation (B.16) we obtain:

$$\sigma_z^2 = \begin{bmatrix} 1 & r_z & 0 \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 1 \\ r_z \\ 0 \end{bmatrix} = \frac{\sigma_p^2}{a_z c_z - b_z^2} (a_z r_z^2 - 2b_z r_z + c_z) \quad (\text{B.20})$$

Hence, the unit-investment, zero-beta frontier is characterized by a parabola. The unit-investment, zero-beta portfolios can attain infinitely many expected returns, rendering the zero-beta rate indeterminate and unidentified. This establishes Proposition 1 (ii).  $\square$

## B.2. Proof of Proposition 2

*Proof.*

This proposition explores the relation between portfolio inefficiency and the level of the estimated zero-beta rate from the unit-investment, minimum-variance zero-beta portfolio. Recall from equation (B.20) that the zero-beta frontier with respect to an inefficient portfolio  $p$  is  $\sigma_z^2 = \frac{\sigma_p^2}{a_z c_z - b_z^2} (a_z r_z^2 - 2b_z r_z + c)$ . The variance is minimized at

$$r_z = \frac{b_z}{a_z} = \frac{b\sigma_p^2 - r_p}{a\sigma_p^2 - 1} \quad (\text{B.21})$$

$$= \frac{b/a \cdot \sigma_p^2 - 1/a \cdot r_p}{\sigma_p^2 - 1/a} \quad (\text{B.22})$$

$$= \frac{r_{GMV} \cdot \sigma_p^2 - \sigma_{GMV}^2 \cdot r_p}{\sigma_p^2 - \sigma_{GMV}^2} \quad (\text{B.23})$$

$$= r_{GMV} - \sigma_{GMV}^2 \frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} \quad (\text{B.24})$$

where  $r_z$  is the estimated zero-beta rate,  $r_{GMV} = b/a$ , and  $\sigma_{GMV}^2 = 1/a$ .

For the tangency portfolio  $p^*$ , recall from Appendix B.1 Step 2 that  $r_f = \frac{c - br_{p^*}}{b - ar_{p^*}}$ . Combined with the MVF formula,  $\sigma_{p^*}^2 = \frac{a}{ac - b^2} (ar_{p^*}^2 - 2br_{p^*} + c)$ , it can be shown that

$$r_f = \frac{b_z}{a_z} = \frac{b\sigma_{p^*}^2 - r_{p^*}}{a\sigma_{p^*}^2 - 1} = r_{GMV} - \sigma_{GMV}^2 \frac{r_{p^*} - r_{GMV}}{\sigma_{p^*}^2 - \sigma_{GMV}^2} \quad (\text{B.25})$$

Relationships between  $r_{GMV}$ ,  $r_f$ , and  $r_z$

First, let  $r_{GMV} - \sigma_{GMV}^2 \frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} = r_f$ , we have:

$$\frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} = \frac{r_{GMV} - r_f}{\sigma_{GMV}^2} \quad (\text{B.26})$$

Rearrange equation (B.24) and (B.25) we know:

$$\frac{r_p - r_{GMV}}{\sigma_p^2 - \sigma_{GMV}^2} = \frac{r_p - r_z}{\sigma_p^2}, \quad \frac{r_{GMV} - r_f}{\sigma_{GMV}^2} = \frac{r_{p^*} - r_f}{\sigma_{p^*}^2} \quad (\text{B.27})$$

Graphically, this means that in the mean-variance diagram ( $\mu$ - $\sigma^2$ ), a line crossing  $GMV$  and  $p^*$  intersects with the vertical axis at  $r_f$ , and a line crossing  $GMV$  and  $p$  intersects with the vertical axis at  $r_z$ . Hence,  $r_z \geq r_f$  if and only if  $(r_p - r_z)/\sigma_p^2 \leq (r_{p^*} - r_f)/\sigma_{p^*}^2$ .

Lastly, equation (B.24) implies that  $r_z > r_{GMV}$  if and only if  $r_p < r_{GMV}$ .

### Zero-Beta Rate and Portfolio Inefficiency

I separately consider two types of portfolio inefficiency.

(1) *Risk inefficiency, holding the mean return fixed* ( $r_p = r_{p^*}$ ,  $\sigma_p^2 > \sigma_{p^*}^2$ ).

Consider the derivative:  $\frac{dr_z}{d\sigma_p^2} = \frac{ar_p - b}{(a\sigma_p^2 - 1)^2}$ . Assuming that portfolio  $p$  lies above the  $GMV$  portfolio,  $r_p > r_{gmv} = b/a$ , we know  $\frac{dr_z}{d\sigma_p^2} > 0$ . Hence,  $r_z$  increases with the  $\sigma_p^2$ , holding the mean return fixed. Since  $\sigma_p^2 > \sigma_{p^*}^2$ , we have  $r_z > r_z^*$ .

(2) *Return inefficiency, holding the volatility fixed* ( $\sigma_p = \sigma_{p^*}$ ,  $r_p < r_{p^*}$ ).

Consider the derivative:  $\frac{dr_z}{dr_p} = -\frac{1}{a\sigma_p^2 - 1}$ . In this case,  $\frac{dr_z}{dr_p} < 0$  holds unambiguously since  $\sigma_p > 1/a = \sigma_{gmv}$ . Hence,  $r_z$  increases as  $r_p$  falls, holding the volatility fixed. Finally, compute  $r_z - r_z^* = \frac{r_{p^*} - r_p}{a\sigma_p^2 - 1} > 0$ .

□

### B.3. Proof of Proposition 3

*Proof.*

Step 1: Mean-variance frontier.

This proof works with the zero-investment frontier, as opposed to unit-investment frontier. Deriving the analytical expression for the mean-variance frontier follows a procedure similar to that in Chapter 5 of [Cochrane \(2009\)](#). For a given target return  $r_{p^*}$ , the variance is minimized by solving the following problem:

$$\min_{\omega} \omega' \Sigma \omega \quad s.t. \quad \omega' \iota = 0, \quad \omega' \mu = r_{p^*} \quad (\text{B.28})$$

where  $\boldsymbol{\mu}$  is the asset mean returns,  $\boldsymbol{\Sigma}$  is the variance-covariance matrix, and  $\boldsymbol{\iota}$  is a vector of ones. Set up the Lagrangian:

$$\mathcal{L} = \frac{1}{2}\boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega} - \lambda_1\boldsymbol{\omega}'\boldsymbol{\iota} - \lambda_2(\boldsymbol{\omega}'\boldsymbol{\mu} - r_{p^*}) \quad (\text{B.29})$$

The first-order condition is given by:

$$\boldsymbol{\Sigma}\boldsymbol{\omega} = \lambda_1\boldsymbol{\iota} + \lambda_2\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \implies \boldsymbol{\omega} = \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad (\text{B.30})$$

Premultiply equation (B.30) by  $\begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\mu}' \end{bmatrix}$  we have:

$$\begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\mu}' \end{bmatrix} \boldsymbol{\omega} = \begin{bmatrix} 0 \\ r_{p^*} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\mu}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad (\text{B.31})$$

Denote

$$\mathbf{A} \equiv \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\mu}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\iota}\boldsymbol{\Sigma}^{-1}\boldsymbol{\iota} & \boldsymbol{\iota}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} \\ \boldsymbol{\mu}\boldsymbol{\Sigma}^{-1}\boldsymbol{\iota} & \boldsymbol{\mu}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} \end{bmatrix} \equiv \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad (\text{B.32})$$

where  $a = \boldsymbol{\iota}\boldsymbol{\Sigma}^{-1}\boldsymbol{\iota}$ ,  $b = \boldsymbol{\iota}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$ , and  $c = \boldsymbol{\mu}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$ .

Thus

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_{p^*} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\omega}_{p^*} = \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_{p^*} \end{bmatrix} \quad (\text{B.33})$$

The variance of this efficient portfolio  $p^*$  is:

$$\begin{aligned} \sigma_{p^*}^2 &= \boldsymbol{\omega}_{p^*}'\boldsymbol{\Sigma}\boldsymbol{\omega}_{p^*} \\ &= \begin{bmatrix} 0 & r_{p^*} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\iota} \\ \boldsymbol{\mu} \end{bmatrix} \boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_{p^*} \end{bmatrix} \\ &= \begin{bmatrix} 0 & r_{p^*} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_{p^*} \end{bmatrix} \\ &= \frac{a}{ac - b^2} r_{p^*}^2 \end{aligned} \quad (\text{B.34})$$

Therefore, the mean-variance frontier can be expressed as  $r_{p^*} = \sqrt{\frac{|ac - b^2|}{a}} \sigma_{p^*}$ , which cor-

responds to a straight line emanating from the origin.

Step 2: Zero-beta portfolios with respect to an efficient portfolio  $p^*$ .

The covariance between an arbitrary portfolio  $j$  and an efficient portfolio  $p^*$  is given by:

$$\sigma_{j,p^*} = \boldsymbol{\omega}'_j \boldsymbol{\Sigma} \boldsymbol{\omega}_{p^*} = \boldsymbol{\omega}'_j \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_{p^*}^* \end{bmatrix} = \begin{bmatrix} 0 & r_j \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_{p^*}^* \end{bmatrix} = \frac{a}{ac - b^2} r_j r_{p^*} \quad (\text{B.35})$$

Here, the third equality follows from  $\boldsymbol{\omega}'_j \boldsymbol{\iota} = 0$ , since portfolio  $j$  is assumed to be a zero-investment portfolio. According to equation (B.35), if portfolio  $j$  is zero-beta (i.e., has zero covariance) with respect to the efficient portfolio  $p^*$ , then  $\sigma_{j,p^*} = 0$  implies  $r_j = 0$ , given that  $a \neq 0$  and  $r_{p^*} \neq 0$ . This establishes Proposition ?? (i). Specifically, if the factor model is correctly specified, there exists a factor portfolio  $p^*$  on the mean–variance frontier, and all zero-beta portfolios with respect to  $p^*$  must have zero expected returns.

Step 3: Zero-beta portfolios with respect to an arbitrary portfolio  $p$ .

For zero-investment, zero-beta portfolios with respect to an arbitrary portfolio  $p$  targeting an expected return  $r_z$ , we solve the following problem:

$$\min_{\boldsymbol{\omega}} \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega} \quad \text{s.t.} \quad \boldsymbol{\omega}' \boldsymbol{\iota} = 0, \quad \boldsymbol{\omega}' \boldsymbol{\mu} = r_z, \quad \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega}_p = 0 \quad (\text{B.36})$$

Set up the Lagrangian:

$$\mathcal{L} = \frac{1}{2} \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega} - \lambda_1 \boldsymbol{\omega}' \boldsymbol{\iota} - \lambda_2 (\boldsymbol{\omega}' \boldsymbol{\mu} - r_z) - \lambda_3 \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega}_p \quad (\text{B.37})$$

The first-order condition is given by:

$$\begin{aligned} \boldsymbol{\Sigma} \boldsymbol{\omega} &= \lambda_1 \boldsymbol{\iota} + \lambda_2 \boldsymbol{\mu} + \lambda_3 \boldsymbol{\Sigma} \boldsymbol{\omega}_p = \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} & \boldsymbol{\Sigma} \boldsymbol{\omega}_p \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} \\ \implies \boldsymbol{\omega} &= \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\mu} & \boldsymbol{\Sigma} \boldsymbol{\omega}_p \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} \end{aligned} \quad (\text{B.38})$$

Premultiply equation (B.38) by  $\begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\mu}' \\ \boldsymbol{\omega}'_p \boldsymbol{\Sigma} \end{bmatrix}$  we have:

$$\begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \omega = \begin{bmatrix} 0 \\ r_z \\ 0 \end{bmatrix} = \begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} \quad (\text{B.39})$$

Denote

$$\mathbf{H} \equiv \begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} = \begin{bmatrix} a & b & 0 \\ b & c & r_p \\ 0 & r_p & \sigma_p^2 \end{bmatrix} = \left[ \begin{array}{c|c} A & 0 \\ \hline 0 & r_p \\ \hline 0 & \sigma_p^2 \end{array} \right] \quad (\text{B.40})$$

Thus

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} = \mathbf{H}^{-1} \begin{bmatrix} 0 \\ r_p \\ 0 \end{bmatrix} \quad \text{and} \quad \omega_z = \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 0 \\ r_z \\ 0 \end{bmatrix} \quad (\text{B.41})$$

where I denote  $\omega_z$  as the weights of a portfolio  $z$  that is orthogonal to  $p$ .

The variance of the zero-beta portfolio  $z$  is given by:

$$\begin{aligned} \sigma_z^2 &= \omega_z' \Sigma \omega_z \\ &= \begin{bmatrix} 0 & r_z & 0 \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} \iota' \\ \mu' \\ \omega_p' \Sigma \end{bmatrix} \Sigma^{-1} \Sigma \Sigma^{-1} \begin{bmatrix} \iota & \mu & \Sigma \omega_p \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 0 \\ r_z \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & r_z & 0 \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 0 \\ r_z \\ 0 \end{bmatrix} \end{aligned} \quad (\text{B.42})$$

To understand the relationship between  $\sigma_z$  and  $r_z$ ,  $\mathbf{H}^{-1}$  needs to be calculated. I start by expressing the Schur complement of  $\mathbf{A}$ :

$$S = \sigma_p^2 - \begin{bmatrix} 0 & r_p \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} 0 \\ r_p \end{bmatrix} = \sigma_p^2 - \frac{a}{ac - b^2} r_p^2 \quad (\text{B.43})$$

Recall from equation (B.34) that the variance of an efficient portfolio  $p^*$  is  $\sigma_{p^*}^2 = \frac{a}{ac - b^2} r_{p^*}^2$ . Hence, if we assume the arbitrary portfolio  $p$  has the same expected return as the efficient portfolio  $p^*$ :  $r_p = r_{p^*}$ . Then,  $\frac{a}{ac - b^2} r_p^2 = \frac{a}{ac - b^2} r_{p^*}^2 = \sigma_{p^*}^2$ , and the Schur complement  $S = \sigma_p^2 - \sigma_{p^*}^2$ , which measures the horizontal distance between portfolio  $p$  and  $p^*$  with the same expected return.

Note that:

$$\begin{aligned}
|\mathbf{H}| &= ac\sigma_p^2 - ar_p^2 - b^2\sigma_p^2 \\
&= (ac - b^2)\sigma_p^2 - ar_p^2 \\
&= (ac - b^2)\sigma_p^2 - (ac - b^2)\sigma_{p^*}^2 \\
&= (ac - b^2)(\sigma_p^2 - \sigma_{p^*}^2)
\end{aligned} \tag{B.44}$$

Hence, we can compute  $\mathbf{H}^{-1}$ :

$$\mathbf{H}^{-1} = \frac{1}{(ac - b^2)(\sigma_p^2 - \sigma_{p^*}^2)} \begin{bmatrix} c\sigma_p^2 - r_p^2 & -b\sigma_p^2 & br_p \\ -b\sigma_p^2 & a\sigma_p^2 & ar_p \\ br_p & ar_p & ac - b^2 \end{bmatrix} \tag{B.45}$$

Substitute  $\mathbf{H}^{-1}$  into equation (B.42) we obtain:

$$\sigma_z^2 = \begin{bmatrix} 0 & r_z & 0 \end{bmatrix} \mathbf{H}^{-1} \begin{bmatrix} 0 \\ r_z \\ 0 \end{bmatrix} = \frac{a\sigma_p^2}{(ac - b^2)(\sigma_p^2 - \sigma_{p^*}^2)} r_z^2 \tag{B.46}$$

Define the zero-beta frontier with respect to an arbitrary portfolio  $p$  as the set of zero-beta portfolios for  $p$  that minimize variance for a given level of mean return. Then, the zero-beta frontier can be expressed as:

$$r_z = \sqrt{\frac{(ac - b^2)(\sigma_p^2 - \sigma_{p^*}^2)}{a\sigma_p^2}} \sigma_z = \sqrt{1 - \frac{\sigma_{p^*}^2}{\sigma_p^2}} \sqrt{\frac{ac - b^2}{a}} \sigma_z \tag{B.47}$$

From Equation (B.34) we know that the maximum Sharpe ratio of all assets is  $SR^2(p^*) = \left(\frac{r_{p^*}}{\sigma_{p^*}}\right)^2 = \frac{ac - b^2}{a}$ . Thus, we compute the slope of the zero-investment, zero-beta frontier:

$$S_z = \frac{r_z}{\sigma_z} = \sqrt{1 - \frac{\sigma_{p^*}^2}{\sigma_p^2}} \sqrt{\frac{ac - b^2}{a}} \tag{B.48}$$

$$= \sqrt{1 - \frac{r_{p^*}^2/\sigma_{p^*}^2}{r_p^2/\sigma_p^2}} \sqrt{\frac{ac - b^2}{a}} \tag{B.49}$$

$$= \sqrt{1 - \left(\frac{SR^2(p^*)}{SR^2(p)}\right)^2} SR^2(p^*) \tag{B.50}$$

$$= \sqrt{SR^2(p^*) - SR^2(p)} \tag{B.51}$$

where  $SR^2(p^*)$  denotes the Sharpe ratio of the efficient portfolio  $p^*$  and  $SR^2(p)$  denotes the

Sharpe ratio of the inefficient portfolio  $p$ . Therefore, the slope of the zero-investment, zero-beta frontier—maximum Sharpe ratio attainable by zero-investment, zero-beta portfolios—quantifies how much a portfolio’s Sharpe ratio falls short of the optimal. Therefore, it provides a measure of model misspecification.

To complete Proposition 3, it remains to show that the slope of the asymptote for the unit-investment, zero-beta frontier equals the slope of the zero-investment, zero-beta frontier.

### Equivalence of Slopes

First of all, I conjecture that any unit-investment, zero-beta portfolio weights  $w_{z,u}$  can be decomposed into two orthogonal components:

$$w_{z,u} = w_{z,mv} + w_z \quad (\text{B.52})$$

where  $w_{z,mv}$  denotes the unit-investment, minimum-variance zero-beta portfolio weights, and  $w_z$  denotes the zero-investment, zero-beta portfolio weights. Let me check the constraints and confirm this conjecture.

Investment constraints hold:  $\iota'w_{z,u} = 1 = \iota'w_{z,mv} + \iota'w_z = 1 + 0$ . Zero-beta constraints hold:  $\beta'w_{z,u} = \mathbf{0}_K = \beta'w_{z,mv} + \beta'w_z = \mathbf{0}_K + \mathbf{0}_K$ . Hence, any zero-investment, zero-beta portfolio corresponds to a unit-investment, zero-beta portfolio. The weights shift is the weights of the unit-investment, minimum-variance zero-beta portfolio.

Similar to the property that the GMV portfolio is orthogonal to any zero-investment portfolio ( $\omega_{gmv}'\Sigma\omega_z \propto \iota'\Sigma^{-1}\Sigma\omega_z = \iota'\omega_z = 0$ ), the unit-investment, minimum-variance zero-beta portfolio is orthogonal to any zero-investment, zero-beta portfolio. This is proved using Equation (3):

$$w'_{z,mv}\Sigma w_z = \underbrace{\begin{bmatrix} 1 & \mathbf{0}'_K \end{bmatrix} \left( \begin{bmatrix} \iota' \\ \beta' \end{bmatrix} \Sigma^{-1} \begin{bmatrix} \iota & \beta \end{bmatrix} \right)^{-1}}_C \cdot \begin{bmatrix} \iota' \\ \beta' \end{bmatrix} \Sigma^{-1} \Sigma \omega_z \quad (\text{B.53})$$

$$= C \cdot \begin{bmatrix} \iota' \\ \beta' \end{bmatrix} \omega_z \quad (\text{B.54})$$

$$= C \cdot \mathbf{0}_{K+1} \quad (\text{B.55})$$

$$= 0 \quad (\text{B.56})$$

Due to the orthogonal decomposition, we know returns  $r_{z,u} = r_{z,mv} + r_z$  and variances:  $\sigma_{z,u}^2 = \sigma_{z,mv}^2 + \sigma_z^2$ .

From the previous procedure, we know that the zero-investment, zero-beta frontier is a straight line with slope  $S_z$ . The formula for this frontier can be written as  $r_z = S_z \cdot \sigma_z$ . Since

this frontier corresponds to the unit-investment, zero-beta hyperbola frontier, we can simply write the formula for the hyperbola as  $(r_{z,u} - r_{z,mv})^2 = S_z^2 \cdot (\sigma_{z,u}^2 - \sigma_{z,mv}^2)$ .

According to the properties of a hyperbola, the slope of the asymptote is exactly  $S_z$ .

Now we complete the proofs of Proposition 3.

□

#### B.4. Proof of Proposition 4

*Proof.*

We start with the formula for the tangency portfolio return,  $r_p^*$ :

$$r_p^* = \frac{c - br_f}{b - ar_f} \quad (\text{B.57})$$

We can rewrite this using  $r_{gmv} = b/a$ ,  $a = 1/\sigma_{gmv}^2$ , and  $b = r_{gmv}/\sigma_{gmv}^2$ .

$$r_p^* = \frac{c/a - (b/a)r_f}{b/a - r_f} = \frac{c/a - r_{gmv}r_f}{r_{gmv} - r_f} \quad (\text{B.58})$$

Rearranging this gives:

$$r_p^*(r_{gmv} - r_f) = c/a - r_{gmv}r_f \quad (\text{B.59})$$

We add and subtract  $r_{gmv}^2$  to the right-hand side:

$$r_p^*(r_{gmv} - r_f) = (c/a - r_{gmv}^2) + r_{gmv}^2 - r_{gmv}r_f \quad (\text{B.60})$$

$$r_p^*(r_{gmv} - r_f) = (c/a - r_{gmv}^2) + r_{gmv}(r_{gmv} - r_f) \quad (\text{B.61})$$

Isolating the first term on the right:

$$(r_p^* - r_{gmv})(r_{gmv} - r_f) = c/a - r_{gmv}^2 \quad (\text{B.62})$$

Now we show that  $c/a - r_{gmv}^2 = S^2\sigma_{gmv}^2$ :

$$c/a - r_{gmv}^2 = \frac{c}{a} - \left(\frac{b}{a}\right)^2 = \frac{ac - b^2}{a^2} \quad (\text{B.63})$$

Define  $S^2 \equiv \frac{ac - b^2}{a}$  and note  $\sigma_{gmv}^2 = \frac{1}{a}$ , we have:

$$\frac{ac - b^2}{a^2} = \left(\frac{ac - b^2}{a}\right) \left(\frac{1}{a}\right) = S^2\sigma_{gmv}^2 \quad (\text{B.64})$$

This gives us the central identity:

$$(r_p^* - r_{gmv})(r_{gmv} - r_f) = S^2 \sigma_{gmv}^2 \quad (\text{B.65})$$

We apply our empirical assumption  $S^2 \geq L^2$  to the identity:

$$(r_p^* - r_{gmv})(r_{gmv} - r_f) \geq L^2 \sigma_{gmv}^2 \quad (\text{B.66})$$

Assuming  $r_{gmv} > r_f$ , we can divide by  $(r_{gmv} - r_f)$ :

$$r_p^* - r_{gmv} \geq \frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f} \quad (\text{B.67})$$

Add  $(r_{gmv} - r_f)$  to both sides to get the total tangency premium:

$$(r_p^* - r_{gmv}) + (r_{gmv} - r_f) \geq (r_{gmv} - r_f) + \frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f} \quad (\text{B.68})$$

$$r_p^* - r_f \geq r_{gmv} - r_f + \frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f} \quad (\text{B.69})$$

This proves the first part of the proposition.

This part is a direct application of the Arithmetic Mean-Geometric Mean (AM-GM) inequality, which states that for any non-negative  $A$  and  $B$ ,  $A + B \geq 2\sqrt{AB}$ .

Let  $A = (r_{gmv} - r_f)$  and  $B = \frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f}$ .

$$r_{gmv} - r_f + \frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f} \geq 2\sqrt{(r_{gmv} - r_f) \cdot \left(\frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f}\right)} \quad (\text{B.70})$$

The  $(r_{gmv} - r_f)$  terms inside the square root cancel out:

$$r_{gmv} - r_f + \frac{L^2 \sigma_{gmv}^2}{r_{gmv} - r_f} \geq 2\sqrt{L^2 \sigma_{gmv}^2} = 2L\sigma_{gmv} \quad (\text{B.71})$$

This proves the second part of the proposition.

Conclusion: We have shown that by combining the standard geometry of the mean-variance frontier with an empirical lower bound  $L$  on the maximum zero-beta Sharpe ratio, we arrive at the full inequality:

$$r_{p^*} - r_f \geq r_{gmv} - r_f + \frac{\sigma_{gmv}^2 L^2}{r_{gmv} - r_f} \geq 2\sigma_{gmv} L \quad (\text{B.72})$$

□

### B.5. Minimum-Variance Market-Neutral Portfolio Weights

*Proof.*

Solve the following variance minimization problem:

$$\min_{\omega} \omega' \Sigma \omega \quad s.t. \quad \omega' \iota = 1, \quad \omega' \Sigma \omega_{p^*} = 0 \quad (\text{B.73})$$

Set up the Lagrangian:

$$\mathcal{L} = \frac{1}{2} \omega' \Sigma \omega - \lambda_1 \omega' \iota - \lambda_2 \omega' \Sigma \omega_{p^*} \quad (\text{B.74})$$

The first-order condition is given by:

$$\Sigma \omega = \lambda_1 \iota + \lambda_2 \Sigma \omega_{p^*} \quad \implies \quad \omega = \lambda_1 \Sigma^{-1} \iota + \lambda_2 \omega_{p^*} \quad (\text{B.75})$$

From  $\omega' \iota = 1$ , we have:

$$\lambda_1 \iota' \Sigma^{-1} \iota + \lambda_2 \iota' \omega_{p^*} = 1 \quad \implies \quad \lambda_1 = \frac{1 - \lambda_2 \iota' \omega_{p^*}}{\iota' \Sigma^{-1} \iota} \quad (\text{B.76})$$

Note that the weights for global minimum variance (GMV) portfolio is:

$$\omega_{gmv} = \frac{\Sigma^{-1} \iota}{\iota' \Sigma^{-1} \iota} \quad (\text{B.77})$$

Thus,

$$\omega = (1 - \lambda_2 \iota' \omega_{p^*}) \omega_{gmv} + \lambda_2 \omega_{p^*} \quad (\text{B.78})$$

Now impose the orthogonality condition:

$$\begin{aligned} \omega' \Sigma \omega_{p^*} = 0 &= [(1 - \lambda_2 \iota' \omega_{p^*}) \omega_{gmv} + \lambda_2 \omega_{p^*}]' \Sigma \omega_{p^*} \\ \implies \lambda_2 &= \frac{\omega'_{gmv} \Sigma \omega_{p^*}}{(\iota' \omega_{p^*}) \omega'_{gmv} \Sigma \omega_{p^*} - \omega'_{p^*} \Sigma \omega_{p^*}} \end{aligned}$$

In summary,

$$\omega = (1 - \kappa \iota' \omega_{p^*}) \omega_{gmv} + \kappa \omega_{p^*}, \quad \text{where } \kappa = \frac{\omega'_{gmv} \Sigma \omega_{p^*}}{(\iota' \omega_{p^*}) \omega'_{gmv} \Sigma \omega_{p^*} - \omega'_{p^*} \Sigma \omega_{p^*}} \quad (\text{B.79})$$

□

B.6. Proof of Equation (3)

*Proof.*

Solve the following variance minimization problem:

$$\min_{\boldsymbol{\omega}} \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega} \quad \text{s.t.} \quad \boldsymbol{\omega}' \boldsymbol{\iota} = 1, \quad \boldsymbol{\omega}' \boldsymbol{\beta} = \mathbf{0}_K \quad (\text{B.80})$$

where  $\boldsymbol{\Sigma}$  is the variance-covariance matrix,  $\boldsymbol{\iota}$  is a vector of ones,  $\boldsymbol{\beta}$  is a  $K \times 1$  vector of betas, and  $\mathbf{0}_K$  is a  $K \times 1$  vector of zeros. Set up the Lagrangian:

$$\mathcal{L} = \frac{1}{2} \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega} - \lambda_1 \boldsymbol{\omega}' \boldsymbol{\iota} - \lambda_2 \boldsymbol{\omega}' \boldsymbol{\beta} \quad (\text{B.81})$$

The first-order condition is given by:

$$\boldsymbol{\Sigma} \boldsymbol{\omega} = \lambda_1 \boldsymbol{\iota} + \lambda_2 \boldsymbol{\beta} = \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \implies \boldsymbol{\omega} = \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad (\text{B.82})$$

Premultiply equation (B.38) by  $\begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix}$  we have:

$$\begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix} \boldsymbol{\omega} = \begin{bmatrix} 1 \\ \mathbf{0}_K \end{bmatrix} = \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad (\text{B.83})$$

$$\implies \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \left( \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \right)^{-1} \begin{bmatrix} 1 \\ \mathbf{0}_K \end{bmatrix} \quad (\text{B.84})$$

Thus,

$$\boldsymbol{\omega}_z^* = \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \left( \begin{bmatrix} \boldsymbol{\iota}' \\ \boldsymbol{\beta}' \end{bmatrix} \boldsymbol{\Sigma}^{-1} \begin{bmatrix} \boldsymbol{\iota} & \boldsymbol{\beta} \end{bmatrix} \right)^{-1} \begin{bmatrix} 1 \\ \mathbf{0}_K \end{bmatrix} \quad (\text{B.85})$$

□

B.7. Proof of Equation (17)

*Proof.* Recall from Equation (B.20) that the unit-investment, zero-beta frontier with respect to portfolio  $p$  is

$$\sigma_z^2 = \frac{\sigma_p^2}{a_z c_z - b_z^2} (a_z r_z^2 - 2b_z r_z + c_z) \quad (\text{B.86})$$

The slope of the asymptote is

$$S_z^2 = \frac{a_z c_z - b_z^2}{a_z \sigma_p^2} \quad (\text{B.87})$$

Since  $a_z = a\sigma_p^2 - 1$ ,  $b_z = b\sigma_p^2 - r_p$ , and  $c_z = c\sigma_p^2 - r_p^2$ , let's rearrange the denominator:

$$a_z c_z - b_z^2 = (a\sigma_p^2 - 1)(c\sigma_p^2 - r_p^2) - (b\sigma_p^2 - r_p)^2 \quad (\text{B.88})$$

$$= ac\sigma_p^3 + 2br_p - c - ar_p^2 - b^2\sigma_p^2 \quad (\text{B.89})$$

$$= \sigma_p^2(\sigma_{p^*}^2 - \sigma_p^2)(ac - b^2) \quad (\text{B.90})$$

Hence, the slope of the asymptote can be expressed as (using  $\sigma_{GMV}^2 = 1/a$ ):

$$S_z^2 = \frac{a_z c_z - b_z^2}{a_z \sigma_p^2} \quad (\text{B.91})$$

$$= \frac{\sigma_p^2(\sigma_{p^*}^2 - \sigma_p^2)(ac - b^2)}{\sigma_p^2(a\sigma_p^2 - 1)} \quad (\text{B.92})$$

$$= \frac{ac - b^2}{a} \frac{\sigma_{p^*}^2 - \sigma_p^2}{\sigma_p^2 - \sigma_{GMV}^2} \quad (\text{B.93})$$

$$= \frac{ac - b^2}{a} \left[ 1 - \frac{\sigma_{p^*}^2 - \sigma_{GMV}^2}{\sigma_p^2 - \sigma_{GMV}^2} \right] \quad (\text{B.94})$$

This completes the proof of Equation (17). □

## Appendix C. Additional Analysis on Zero-Beta Rate Estimation

### C.1. Test-Optimization Approach of Zero-Beta Rate Estimation

I summarize the existing zero-beta rate estimation methods into three categories: the *regression approach*, the *test-optimization approach*, and the *zero-beta portfolio approach*. The basic idea of the test-optimization approach is to optimally solve for a zero-beta rate that makes the given factor model perform as well as possible when subjected to formal tests. The idea originates with [Kandel \(1984\)](#), [Kandel \(1986\)](#), and [Shanken \(1986\)](#). These papers construct the likelihood function of the data subject to the model restrictions and then solve for the zero-beta rate that maximizes the constrained likelihood. Equivalently, the estimate can be seen as minimizing the relevant test statistic—variants of the likelihood ratio test (LRT) used across the papers. Intuitively, the zero-beta rate is “tilted” just enough to make the model look as good as it possibly can in sample. [Velu and Zhou \(1999\)](#) uses GMM to

estimate the zero-beta rate that makes the model’s pricing errors as small as possible. More recently, [Ferson et al. \(2025\)](#) shows that the same logic implies choosing the zero-beta rate that minimizes the gap between the maximum squared Sharpe ratio attainable with the test assets and that implied by the model’s factors:

$$\hat{\lambda}_0^* = \max_{\lambda_0} L(\lambda_0) = \min_{\lambda_0} LRT(\lambda_0) = \min_{\lambda_0} \varepsilon'W\varepsilon = \min_{\lambda_0} (SR^2(r, f) - SR^2(f)) \quad (\text{C.1})$$

The first equality corresponds to the constrained likelihood maximization in [Kandel \(1984, 1986\)](#), the second to the test-statistic minimization in [Shanken \(1986\)](#), the third to the pricing errors ( $\varepsilon$ ) minimization with a GMM framework in [Velu and Zhou \(1999\)](#), and the fourth to the Sharpe ratio criterion in [Ferson et al. \(2025\)](#). Put differently, the zero-beta rate estimate is chosen to make the model’s factors lie as close as possible to the mean–variance frontier of returns.


Instead of looking for a single value of the zero-beta rate and then calculating the a standard error around it, [Beaulieu et al. \(2013, 2023, 2025\)](#) ask a different question: “For which possible values of the zero-beta rate would a hypothesis test fail to reject the given factor model at a given significance level?” The set of all such “non-rejected” values forms the confidence interval of the zero-beta rate.

In summary, the test-optimization approach does not so much validate the model as reveal the most “forgiving” estimate of the zero-beta rate consistent with the data. However, forcing the factor model to be accepted—or to fit the data as perfectly as possible—may impose an overly strong assumption.

### *C.2. Zero-Beta Rate and Portfolio Inefficiency*

[Proposition 2](#) (ii) explores how portfolio inefficiency impacts the associated zero-beta rate estimation. Here, I illustrate the results in a hypothetical mean-variance space by investigating two distinct types of portfolio inefficiency. In [Figure C.1](#) (risk inefficiency), I fix the expected return for inefficient portfolios ( $r_{p^*} = r_{p_1} = r_{p_2}$ ) and progressively increase their volatilities ( $\sigma_{p^*}^2 < \sigma_{p_1}^2 < \sigma_{p_2}^2$ ). Here,  $p^*$  represents the efficient portfolio on the mean-variance frontier (MVF), while  $p_1$  and  $p_2$  have the same return but excess, uncompensated risk. This panel shows how the zero-beta frontier and its minimum variance zero-beta portfolio shift as the reference portfolio moves horizontally away from the efficient frontier. Holding the mean return fixed, the estimated zero-beta rate (expected return of the minimum-variance zero-beta portfolio) rises with the volatility of the inefficient portfolio  $p$ . In [Figure C.2](#) (return inefficiency), I fix the volatility for inefficient portfolios ( $\sigma_{p^*}^2 = \sigma_{p_1}^2 = \sigma_{p_2}^2$ ) and

Figure. C.1. Zero-Beta Rate and Portfolio Inefficiency (Constant Mean)




*Notes:* This figure shows the relations between portfolio inefficiency and the level of estimated zero-beta rate in the mean-standard deviation diagram. Holding the mean return constant, the estimated zero-beta rate rises with the volatility of the inefficient portfolio  $p$ . Panel (b) holds the volatility constant, the estimated zero-beta rate rises as the mean return of  $p$  falls.

progressively decrease their expected return ( $r_{p^*} > r_{p_1} > r_{p_2}$ ). Here,  $p^*$  again represents the efficient increase, while  $p_1$  and  $p_2$  take on the same amount of risk for a lower reward. This panel shows how the zero-beta frontier and its minimum variance zero-beta portfolio shift as the reference portfolio moves vertically downward from the efficient frontier. Holding the volatility fixed, the estimated zero-beta rate rises as the mean return of  $p$  falls.

### C.3. Zero-Beta Rate in the Mean-Variance Diagram

Equation (6) provides the formula for the expected return ( $r_z$ ) of the minimum-variance zero-beta portfolio with respect to the factor model benchmark portfolio  $p$ . Figure C.3 considers two inefficient portfolios in Region II. When portfolio  $p'$  is more inefficient than  $p$  (i.e., farther from the mean-variance frontier), the red curve representing equation (6) rotates clockwise, and the estimated zero-beta rate shifts upward toward the mean return of the GMV portfolio. In the mean-variance diagram, the zero-beta rate contour curves derived from (6) are depicted in Figure C.4. The space of inefficient portfolios can be divided into three regions, as illustrated in Figure C.4. If portfolio  $p$  lies in region I, then the estimated zero-beta rate is downward biased relative to the true unobserved risk-free rate ( $r_z < r_f$ ).

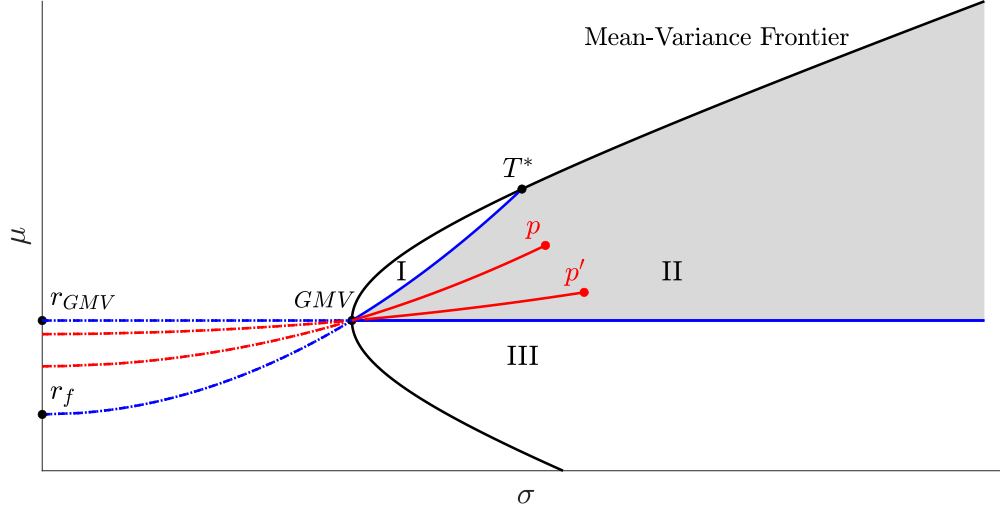
Figure. C.2. Zero-Beta Rate and Portfolio Inefficiency (Constant Volatility)



*Notes:* This figure shows the relations between portfolio inefficiency and the level of estimated zero-beta rate in the mean-standard deviation diagram. Holding the volatility constant, the estimated zero-beta rate rises as the mean return of  $p$  falls.

If portfolio  $p$  lies below the GMV portfolio in region III, then the zero-beta rate is higher than the GMV portfolio return ( $r_z > r_{GMV}$ ). If portfolio  $p$  lies in region II, then the zero-beta rate is upward biased relative to the true risk-free rate and it is lower than the GMV portfolio return ( $r_f < r_z < r_{GMV}$ ). Therefore,  $r_z$  could be downward biased or upward biased relative to  $r_f$  depending on the location of portfolio  $p$ .  $r_z$  may also be exactly equal to the unobserved  $r_f$  if the factor model happens to be lying on the boundary of region I and II, where  $(r_p - r_f)/\sigma_p^2 = (r_{p^*} - r_f)/\sigma_{p^*}^2$ . This boundary is shown as the blue solid contour line in Figure C.4. This line extends leftward and intersects the vertical axis at  $r_f$ . In addition, C.4 plots other zero-beta rate contour lines, where portfolios lying on the same blue dashed line imply an identical zero-beta rate, corresponding to the intercept on the vertical axis if the line were extended leftward (not shown).

Figure. C.3. Zero-Beta Rate and Portfolio Inefficiency in Region II



*Notes:* This figure considers two inefficient portfolios in Region II. When portfolio  $p'$  is more inefficient than  $p$  (i.e., farther from the mean–variance frontier), the red curve representing equation (6) rotates clockwise, and the estimated zero-beta rate shifts upward toward the mean return of the GMV portfolio.

#### C.4. Instrumented PCA (IPCA) Estimation

I estimate IPCA models following Kelly et al. (2019):

$$r_{i,t+1} = \alpha(\mathbf{z}_{i,t}) + \boldsymbol{\beta}(\mathbf{z}_{i,t})' \mathbf{f}_{t+1} + \varepsilon_{i,t+1} \quad (\text{C.2})$$

$$\alpha(\mathbf{z}_{i,t}) = \mathbf{z}_{i,t}' \Gamma_{\alpha} + \nu_{\alpha,i,t} \quad (\text{C.3})$$


$$\boldsymbol{\beta}(\mathbf{z}_{i,t}) = \mathbf{z}_{i,t}' \Gamma_{\beta} + \nu_{\beta,i,t} \quad (\text{C.4})$$

where  $r_{i,t+1}$  is the total return of stock  $i$ ,  $\mathbf{f}_{t+1}$  is a  $K$ -dimensional vector of factors,  $\alpha(\mathbf{z}_{i,t})$  and  $\boldsymbol{\beta}(\mathbf{z}_{i,t})$  denote the intercept and risk loadings, modeled as linear functions of the 136 stock characteristics,  $\mathbf{z}_{i,t}$ .  $\nu_{\alpha,i,t}$  and  $\nu_{\beta,i,t}$  represents components in  $\alpha$  and  $\boldsymbol{\beta}$  that are orthogonal to characteristics. Substitute Equations (C.3) and (C.4) into (C.2), we have:

$$r_{i,t+1} = \mathbf{z}_{i,t}' (\Gamma_{\alpha} + \Gamma_{\beta} \mathbf{f}_{t+1}) + \tilde{\varepsilon}_{i,t+1} \quad (\text{C.5})$$

where  $\tilde{\varepsilon}_{i,t+1} = \varepsilon_{i,t+1} + \nu_{\alpha,i,t} + \nu_{\beta,i,t} \mathbf{f}_{t+1}$  is a composite error term. Denote the  $N \times 1$  total return vector as  $\mathbf{r}_{t+1}$ , the  $N \times 1$  composite error vector as  $\tilde{\boldsymbol{\varepsilon}}_{t+1}$ , and the  $N \times C$  ( $C = 136$ )

Figure. C.4. Zero-Beta Rate Contour Lines and Portfolio Inefficiency



*Notes:* This figure illustrates the estimated zero-beta rate contour lines in mean–variance space. The black hyperbola represents the mean–variance frontier. Portfolios located on the blue solid contour line imply a zero-beta rate equal to the true risk-free rate. This line extends leftward and intersects the vertical axis at  $r_f$ . Portfolios lying on the same blue dashed contour line imply an identical zero-beta rate, corresponding to the intercept on the vertical axis if the line were extended leftward (not shown). The space of inefficient portfolios can be divided into three regions. If portfolio  $p$  lies in region I, then  $r_z < r_f$ ; if it lies in region II, then  $r_f < r_z < r_{GMV}$ ; and if it lies in region III, then  $r_z > r_{GMV}$ .

stock characteristics matrix as  $\mathbf{Z}_{t+1}$ . Then, the vector form of an IPCA model is:

$$\mathbf{r}_{t+1} = \mathbf{Z}_{t+1} (\Gamma_\alpha + \Gamma_\beta \mathbf{f}_{t+1}) + \tilde{\boldsymbol{\varepsilon}}_{t+1} \quad (\text{C.6})$$

The optimization objective is to minimize the sum of squared model composite errors:

$$\min_{\Gamma_\alpha, \Gamma_\beta, \mathbf{f}_{t+1}} \sum_{t=1}^T \left( \mathbf{r}_{t+1} - \mathbf{Z}_{t+1} \Gamma_\alpha - \mathbf{Z}_{t+1} \Gamma_\beta \mathbf{f}_{t+1} \right)' \left( \mathbf{r}_{t+1} - \mathbf{Z}_{t+1} \Gamma_\alpha - \mathbf{Z}_{t+1} \Gamma_\beta \mathbf{f}_{t+1} \right) \quad (\text{C.7})$$

Table C.1: IPCA Model Performance and Alphas

Panel A: Zero-Investment Managed Portfolios				
	1	3	6	9
Total $R^2$ (%)	66.5	91.3	95.2	96.8
$W_\alpha$ p-value (%)	0.00	0.00	25.2	32.6
$\mathbb{E}[\alpha_{p,t}]$ (Annualized, %)	0.16	0.32	0.40	0.53
Panel B: Unit-Investment Characteristic-Sorted Portfolios				
	1	3	6	9
Total $R^2$ (%)	75.4	93.1	94.4	95.3
$\mathbb{E}[\alpha_{p,t}]$ (Annualized, %)	7.07	11.13	9.74	4.33

*Notes:* Panel A reports total  $R^2$ 's, Wald test p-values for  $\mathcal{H}_0 : \Gamma_\alpha = 0$ , average alphas across portfolios and time,  $\mathbb{E}[\alpha_{p,t}]$ , where  $\alpha_{p,t}$  is an element of the portfolio alpha vector  $\alpha_p = \mathbf{Z}'_t \Gamma_\alpha / N_t$  using zero-investment managed portfolios,  $\mathbf{X}_{t+1} \equiv \mathbf{Z}_{t+1} \mathbf{r}_{t+1} / N_{t+1}$ . Panel B reports total  $R^2$ 's and average alphas across portfolios and time,  $\mathbb{E}[\alpha_{p,t}]$ , where  $\alpha_{p,t}$  is an element of the portfolio alpha vector  $\alpha_p = \boldsymbol{\omega}'_t \mathbf{Z}_t \Gamma_\alpha$  and  $\boldsymbol{\omega}_t$  denotes the unit-investment portfolio weights.

The first-order conditions (F.O.C.) for this problem are:

$$\hat{\mathbf{f}}_{t+1} = \left( \hat{\Gamma}'_\beta \mathbf{Z}'_{t+1} \mathbf{Z}_{t+1} \hat{\Gamma}_\beta \right)^{-1} \hat{\Gamma}'_\beta \mathbf{Z}'_{t+1} \left( \mathbf{r}_{t+1} - \mathbf{Z}_{t+1} \hat{\Gamma}_\alpha \right) \quad (\text{C.8})$$

$$\text{vec}(\hat{\Gamma}') = \left( \sum_{t=1}^T (\mathbf{Z}'_{t+1} \mathbf{Z}_{t+1}) \otimes (\tilde{\mathbf{f}}_{t+1} \tilde{\mathbf{f}}'_{t+1}) \right)^{-1} \left( \sum_{t=1}^T (\mathbf{Z}'_{t+1} \otimes \tilde{\mathbf{f}}_{t+1}) \mathbf{r}_{t+1} \right) \quad (\text{C.9})$$

where I denote  $\Gamma \equiv [\Gamma_\alpha \ \Gamma_\beta]$  and  $\tilde{\mathbf{f}}_{t+1} \equiv [1 \ \mathbf{f}'_{t+1}]'$ . These system of equations are solved numerically using the Alternative Least Squares (ALS) algorithm.

Kelly et al. (2019) presents a managed-portfolio interpretation of IPCA. The estimated IPCA factors and parameters can also be viewed as minimizing the pricing errors of managed portfolio returns, which are constructed as weighted averages of individual stock returns interacted with instruments:  $\mathbf{X}_{t+1} \equiv \mathbf{Z}_{t+1} \mathbf{r}_{t+1} / N_{t+1}$ , where  $N_{t+1}$  denotes the number of non-missing stock returns at time  $t + 1$ . After estimating the model parameters, Kelly et al. (2019) proceeds to test whether alphas arise as a function of characteristics:  $\mathcal{H}_0 : \Gamma_\alpha = 0$ . I follow the same bootstrap procedure to compute the p-values of this test.

Table C.1 Panel A reports total  $R^2$  values and Wald test p-values for  $\mathcal{H}_0 : \Gamma_\alpha = 0$  using

the zero-investment managed portfolios  $X$  with 1, 3, 6, and 9 IPCA factors. Consistent with Kelly et al. (2019), the IPCA models explain portfolio return variation well, with total  $R^2$  values exceeding 66%. With more than six factors, the Wald test p-values rise above 1%, leading to a failure to reject  $\Gamma_\alpha = 0$  at the 1% level. For the zero-investment managed portfolios, I compute the average alphas across portfolios and time,  $\mathbb{E}[\alpha_{p,t}]$ , where  $\alpha_{p,t}$  is an element of the portfolio alpha vector  $\boldsymbol{\alpha}_p = \mathbf{Z}'_t \mathbf{Z}_t \Gamma_\alpha / N_t$ . These average alphas range from 0.16% to 0.53% on an annualized basis, suggesting that characteristic-related mispricings are small in such zero-investment managed portfolios.

Note that each stock characteristic vector  $\mathbf{z}_{i,t+1}$  is rank-normalized to the  $(-1, 1)$  interval, with elements summing to zero. This normalization ensures that the managed portfolios are zero-investment portfolios, where the risk-free (zero-beta) rate cancels out. As an alternative investigation, I construct 273 unit-investment characteristic-sorted portfolios based on the portfolio weights (see Section 3.1) and use the estimated  $\Gamma_\alpha$  and  $\Gamma_\beta$  to assess IPCA model performance. Table C.1 Panel B reports total  $R^2$  values for these portfolios, which are similar to those in Panel A. More importantly, for unit-investment portfolios, Panel B also presents the average alphas across portfolios and time,  $\mathbb{E}[\alpha_{p,t}]$ , where  $\alpha_{p,t}$  is an element of the portfolio alpha vector  $\boldsymbol{\alpha}_p = \boldsymbol{\omega}'_t \mathbf{Z}_t \Gamma_\alpha$  and  $\boldsymbol{\omega}_t$  denotes the unit-investment portfolio weights. With  $\Gamma_\alpha$  statistically zero, the average alphas should also be zero; however, Panel B shows values ranging from 4.33% to 11.13% on an annualized basis. This evidence suggests that the zero-alpha conclusion should be interpreted with caution.

### C.5. Revisiting Black (1972) and Black, Jensen, and Scholes (1972)

Using the conceptual framework regarding the uniqueness of the zero-beta rate discussed in Section 2, we can revisit the classic theoretical contribution of Black (1972) and its empirical counterpart, Black, Jensen, and Scholes (1972). This perspective clarifies both the power and the limitations of the zero-beta CAPM and its empirical implementation.

Relative to the standard Sharpe–Lintner CAPM, Black (1972) relaxes only one assumption: *free riskless borrowing and lending*. Instead of assuming that investors can borrow or lend unlimited amounts at a common risk-free rate, Black (1972) considers an economy without a riskless asset. In this setting, expected returns satisfy:

$$\mathbb{E}[R_i] = \mathbb{E}[R_z] + \beta (\mathbb{E}[R_m] - \mathbb{E}[R_z]) \tag{C.10}$$

where  $R_z$  is the return on the zero-beta portfolio relative to the market portfolio. The zero-beta portfolio replaces the risk-free asset in the pricing relation.

Black et al. (1972) provide the empirical counterpart by estimating a two-factor repre-

sentation:

$$\mathbb{E}[R_i] = \underbrace{(1 - \beta)\mathbb{E}[R_z]}_{\text{explains the flat SML}} + \beta\mathbb{E}[R_m] \quad (\text{C.11})$$

This formulation highlights the role of the intercept term,  $(1 - \beta_i)\mathbb{E}[R_z]$ . It provides a natural explanation for the empirical "beta anomaly" or flat security market line. If  $\beta_i > 1$ , then  $(1 - \beta_i) < 0$ , generating negative alphas. If  $\beta_i < 1$ , then  $(1 - \beta_i) > 0$ , generating positive alphas.

Thus, the presence of the zero-beta intercept can rationalize why high-beta assets underperform and low-beta assets outperform relative to the standard CAPM prediction.

Although Black and coauthors remove the assumption of free borrowing in the riskless asset, the model still assumes unrestricted short-selling and borrowing in risky assets. In practice, constraints in risky-asset markets may be just as severe as constraints in riskless borrowing. Therefore, the model replaces one strong assumption with another that may be equally unrealistic. This is a traditional caveat emphasized in the literature.

From the perspective of the uniqueness of zero-beta rate, a deeper issue arises. Black's framework assumes that the market portfolio lies on the mean-variance frontier as illustrated in . Under this assumption, the zero-beta portfolio is uniquely defined, and  $\mathbb{E}[R_z]$  is well identified. However, if the market portfolio is not mean-variance efficient—as empirical evidence often suggests—then the zero-beta portfolio is not unique. In that case,  $\mathbb{E}[R_z]$  is not uniquely identified, the intercept in the two-factor model is not well defined, and the empirical two-factor loses its structural foundation.

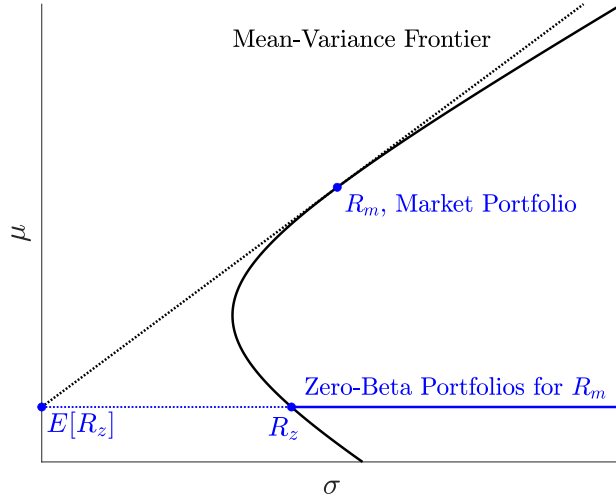
In summary, revisiting [Black \(1972\)](#) and [Black et al. \(1972\)](#) through the lens of zero-beta rate uniqueness clarifies both their contribution and their limitations. While relaxing free riskless borrowing generates an elegant alternative CAPM and helps explain the beta anomaly, the theory depends crucially on market efficiency and on the uniqueness of the zero-beta portfolio—assumptions that may not hold empirically.

## Appendix D. Additional Empirical Results

### D.1. Factor Model Statistical Performance

In addition to total  $R^2$ , predictive  $R^2$  measures the model explanatory power of test assets using the factor risk premia, calculated as the prevailing sample average of factors up to the last month:

Figure. C.5. Revisiting Black (1972) and Black, Jensen, and Scholes (1972)



*Notes:* This figure shows the unit-investment zero-beta portfolios (orthogonal portfolios) for the market portfolio  $R_m$ , assuming that the market is on the mean-variance frontier (black hyperbola). The blue horizontal solid line represents the zero-beta portfolios with respect to  $R_m$ , with  $R_z$  denoting the corresponding minimum-variance portfolio and  $E[R_z]$  the zero-beta rate with respect to  $R_m$ .

$$R_{\text{pred}}^2 = 1 - \frac{\sum_{i,t} (r_{i,t} - \hat{\beta}'_i \hat{\lambda}_{t-1})^2}{\sum_{i,t} r_{i,t}^2}. \quad (\text{D.1})$$

For comparison, total  $R^2$  measures how well the realized factor returns explain realized asset returns, whereas predictive  $R^2$  evaluates how well a model's conditional expected returns explain realized asset returns. Table D.1 and D.2 present both performance metrics for individual stocks and characteristic-sorted portfolios across the FF, PCA, IPCA, and AE models, constructed in-sample and out-of-sample, respectively. Consistent with Gu et al. (2021), I find that IPCA and AE models outperform the standard FF and PCA benchmarks in terms of both total and predictive  $R^2$ , for both individual stocks and characteristic-sorted portfolios. As expected, predictive  $R^2$  values for individual stocks are negative under the FF and PCA models. Comparing IPCA and AE, I find that IPCA achieves slightly higher total  $R^2$ , whereas AE delivers higher predictive  $R^2$ , particularly for individual stocks.

Table D.2 replicates Tables 1 and 2 of Gu et al. (2021), but the magnitudes of predictive  $R^2$  for characteristic-sorted portfolios are considerably larger in my results. For IPCA and AE, this difference likely stems from portfolio construction. I form 273 extreme-tercile characteristic-sorted portfolios, whereas Gu et al. (2021) uses managed portfolios constructed

Table D.1: In-Sample Model  $R^2$ 

Models	Test Assets	Metrics	# Factors			
			1	3	6	9
FF	Individual Stocks	Total $R^2$	10.9	17.0	18.7	19.5
		Pred $R^2$	0.93	0.90	0.89	0.87
	Portfolios	Total $R^2$	89.4	94.1	95.2	95.4
		Pred $R^2$	3.42	3.42	3.42	3.42
PCA	Individual Stocks	Total $R^2$	4.5	10.3	13.1	14.3
		Pred $R^2$	0.01	0.28	0.34	0.34
	Portfolios	Total $R^2$	94.1	97.4	98.1	98.3
		Pred $R^2$	3.42	3.42	3.42	3.42
IPCA	Individual Stocks	Total $R^2$	12.8	15.1	15.8	16.0
		Pred $R^2$	0.81	0.79	0.78	0.77
	Portfolios	Total $R^2$	78.7	94.0	95.0	95.5
		Pred $R^2$	3.01	3.10	3.18	3.17
AE	Individual Stocks	Total $R^2$	12.4	13.3	13.5	13.5
		Pred $R^2$	1.12	1.02	1.09	1.16
	Portfolios	Total $R^2$	84.1	92.9	92.9	93.3
		Pred $R^2$	3.52	3.50	3.51	3.37

*Notes:* This table reports the in-sample total  $R^2$  and predictive  $R^2$  in percentages (%) for FF, PCA, ICA, and AE models with 1, 3, 6, and 9 factors.

from the characteristics matrix,  $x_t = (Z'_{t-1}Z_{t-1})Z_{t-1}r_t$ . For FF and PCA, the discrepancy primarily reflects my inclusion of an intercept in model estimation. Since both FF and PCA are estimated via OLS, including an intercept implies that predicted returns reduce to historical averages. Consequently, the predictive  $R^2$  for FF and PCA essentially capture the predictive  $R^2$  associated with expanding mean returns.

#### D.2. In-Sample Zero-Beta Rate Estimates

In Section 3.5, I examine two sets of characteristic-sorted portfolios that differ in their GMV portfolio returns to test this hypothesis of spurious robustness. Specifically, I rank the characteristic-sorted portfolios by their return variances and select the 130 portfolios with the highest variances as an alternative universe of test assets. The first asset group thus contains the full set of 273 characteristic-sorted portfolios, while the second group includes only the

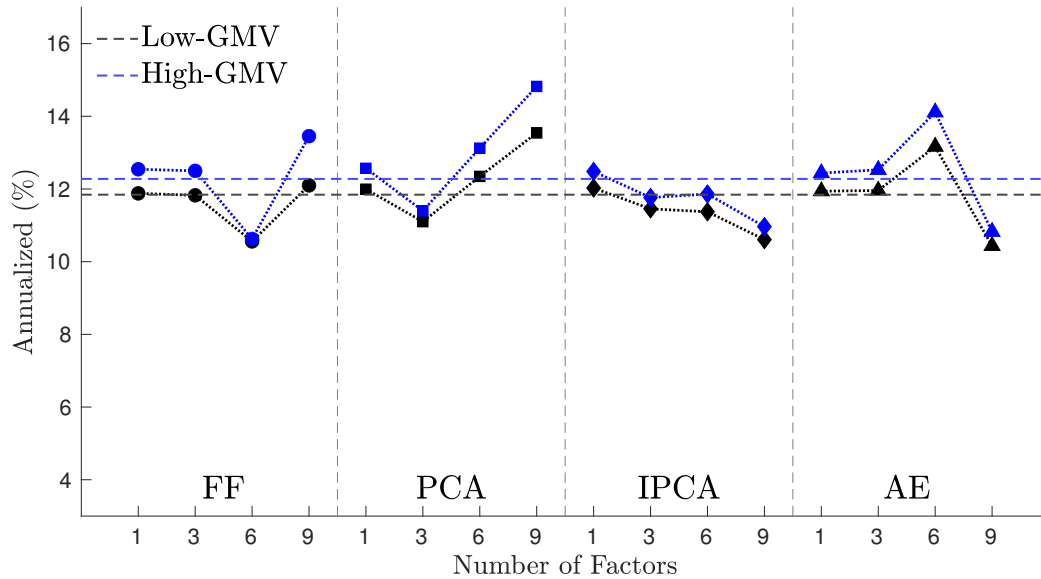
Table D.2: Out-of-Sample Model  $R^2$ 

Models	Test Assets	Metrics	# Factors			
			1	3	6	9
FF	Individual Stocks	Total $R^2$	6.5	7.4	2.7	0.3
		Pred $R^2$	-0.23	-0.21	-0.24	-0.25
	Portfolios	Total $R^2$	86.6	92.1	93.6	93.9
		Pred $R^2$	2.89	2.88	2.89	2.89
PCA	Individual Stocks	Total $R^2$	7.4	6.9	7.1	7.2
		Pred $R^2$	-1.09	-1.06	-1.07	-1.07
	Portfolios	Total $R^2$	92.3	96.2	97.1	97.5
		Pred $R^2$	2.88	2.88	2.88	2.88
IPCA	Individual Stocks	Total $R^2$	11.0	13.3	13.9	14.1
		Pred $R^2$	0.54	0.55	0.54	0.52
	Portfolios	Total $R^2$	73.6	93.0	93.8	94.4
		Pred $R^2$	3.15	3.22	3.23	3.27
AE	Individual Stocks	Total $R^2$	10.2	11.3	11.3	11.3
		Pred $R^2$	0.66	0.78	0.71	0.81
	Portfolios	Total $R^2$	79.5	91.7	93.3	92.5
		Pred $R^2$	3.23	3.02	3.27	3.13

*Notes:* This table reports the out-of-sample total  $R^2$  and predictive  $R^2$  in percentages (%) for FF, PCA, ICA, and AE models with 1, 3, 6, and 9 factors.

136 high-variance portfolios. The analytical portfolio weights for the GMV portfolio are given by  $\Sigma^{-1}\mathbf{1}/\mathbf{1}'\Sigma^{-1}\mathbf{1}$ . The resulting in-sample mean returns of the GMV portfolio are 11.8% and 12.3% for the two asset groups, respectively. This partition enables an examination of whether zero-beta rate estimates differ systematically across asset universes characterized by distinct GMV portfolio returns.

Figure. D.1. In-Sample Zero-Beta Rate across Different Asset Universes



*Notes:* This figure shows the in-sample estimated zero-beta rates obtained from unit-investment, minimum-variance zero-beta portfolios across two asset universes that differ in their GMV portfolio returns. In the first asset group (black dashed line), which includes the full set of 273 characteristic-sorted portfolios, the mean GMV portfolio return is lower (11.8%). In the second asset group (blue dashed line), consisting of the 136 high-variance portfolios, the mean GMV portfolio return is higher (12.3%). Four classes of factor models—FF, PCA, IPCA, and AE—with 1, 3, 6, and 9 factors are analyzed. The estimated zero-beta rates are represented by circles, squares, diamonds, and triangles, respectively.


Figure D.1 shows the in-sample estimated zero-beta rates obtained from unit-investment, minimum-variance zero-beta portfolios across two asset universes that differ in their GMV portfolio returns. In the first asset group (black dashed line), which includes the full set of 273 characteristic-sorted portfolios, the mean GMV portfolio return is lower (11.8%). In the second asset group (blue dashed line), consisting of the 130 high-variance portfolios, the mean GMV portfolio return is higher (12.3%). Across both asset universes, the estimated zero-beta rates appear robust to the choice of factor model and to the number of factors. The literature tends to interpret this spurious robustness as evidence that these estimates capture the true, unobserved risk-free rate. If that were the case, the zero-beta rates should be similar across different asset universes. However, the results show that the zero-beta rates are systematically higher in the universe with the higher GMV portfolio return, although the difference between zero-beta rates is small due to the small difference between GMV

mean returns (11.8% vs 12.3%). Moreover, the average estimated rates lie close to the mean GMV portfolio returns within their respective asset groups. This pattern suggests that the estimated zero-beta rates may primarily reflect the mean return of the GMV portfolio rather than the true risk-free rate, providing empirical support for my analytical conjecture that substantial model misspecification biases zero-beta rate estimates upward.


### *D.3. Simulating Mean-Variance Parameters*

Recall from Section 3.6.2, I calibrate the mean-variance frontier such that the true risk-free rate ( $r_f$ ) is 3% annually, the return of the true Tangency portfolio ( $r_{p^*}$ ) is 20% annually, the mean and standard deviation of the global minimum-variance portfolio is  $r_{GMV} = 11\%$  and  $\sigma_{GMV} = 6.5\%$  (annualized), respectively. Then, inefficient portfolios, corresponding to misspecified factor models, are generated in the following procedure illustrated in Figure D.2. The black curves represent the mean-variance frontier. For a randomly generated inefficient portfolio  $p_1$ , Panel (a) plots the zero-beta frontier and the minimum-variance zero-beta portfolio,  $z_{p_1}$ . Applying equations (6) and (17), I compute the analytical zero-beta rate and the maximum Sharpe ratios of zero-investment, zero-beta portfolios associated with portfolio  $p_1$ . In Panel (b), I plot  $p_1$  in the rate-misspecification space. Similarly in Panel (c) and (d), another inefficient portfolio  $p_2$  is generated and represented by a pair of values of zero-beta rate and the misspecification measure. Repeating this portfolio generating process for 100,000 times, Panel (e) shows the uniform distribution of inefficient portfolios created in the mean-variance space. I restrict that the mean of the inefficient portfolio is higher than  $r_{GMV}$  and not 20% higher than the mean return of the tangency portfolio. Standard deviations of the inefficient portfolios are lower than 40%. Results and conclusions from this simulation exercise is not affected by these boundary choices. All inefficient portfolios are plotted in Panel (f) in the rate-misspecification space. Based on this scatter plot, Figure 8 compute the Cumulative Distribution Function (CDF) of the estimated zero-beta rates from 100,000 inefficient portfolios  $p$  and the probability that the estimated zero-beta rate ( $r_z$ ) exceeds a given threshold ( $x$ ), as a function of the maximum Sharpe ratio of the economy.

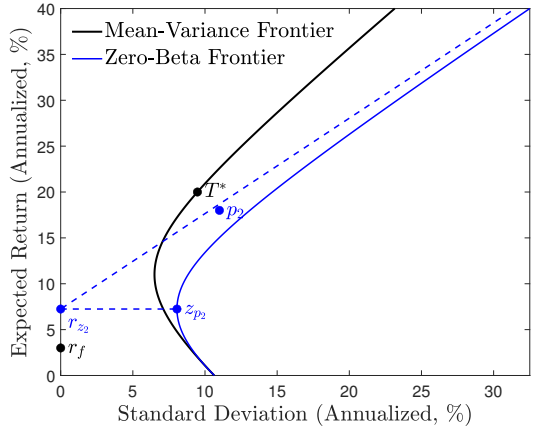
Figure. D.2. Simulated Zero-Beta Rate Estimates



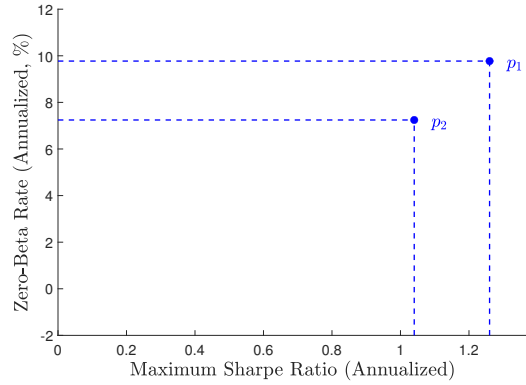
(a) Mean-Variance Space



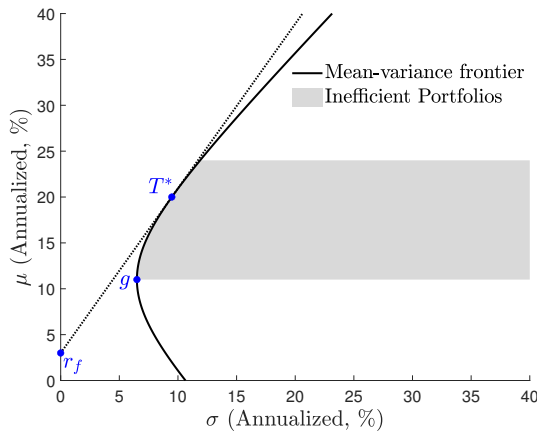
(b) Rate-Misspecification Space



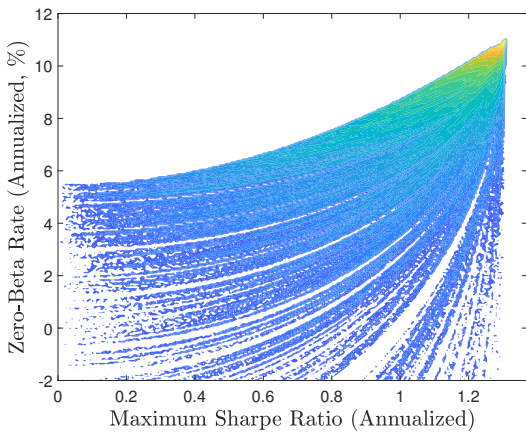
(c) Mean-Variance Space



(d) Rate-Misspecification Space



(e) Mean-Variance Space



(f) Rate-Misspecification Space

*Notes:* This figure illustrates the procedure of generating inefficient portfolios (misspecified model) and plot pairs of zero-beta rate estimates and maximum Sharpe ratios of zero-investment, zero-beta portfolios.