A Market-Based Cost of Capital*

Niels Joachim Gormsen   Kilian Huber   Theis Ingerslev Jensen

March 2026

Please click here for the latest version

Abstract

How should firms set their cost of capital and does it matter for their market value? A classical intuition suggests that firms should base their cost of capital on the opportunity cost in financial markets if they want to maximize their value. This approach is taught to firm managers and used as a bedrock assumption in modern models of firm behavior. But the cost of capital is unobserved and there exists no general method that can estimate the cost of capital without bias; no dataset of an unbiased cost of capital that can be used by practitioners and academics; and no evidence validating the classical intuition that the cost of capital shapes market values. We introduce a new method based on machine learning forecasts of cash flows to estimate a firm-level "market-based cost of capital" (MCC). The resulting MCC is a strong and unbiased predictor of future firm-level equity returns at the 10-year horizon. In practice, most firms' perceptions of the cost of capital differ substantially from the MCC. In the data, firms using a perceived cost of capital closer to the MCC achieve higher market values.

*We are grateful for helpful comments from Paul Fontanier, Stefano Giglio, Bryan Kelly, Lasse Pedersen and Alp Simsek; and conference and seminar participants at 2025 HBS Juniors Conference, UCLA Anderson, Yale (accounting), Yale (finance). Jensen gratefully acknowledges support from the Center for Big Data in Finance (grant no. DNRF167). Gormsen is at The University of Chicago and Copenhagen Business School; www.voices.uchicago.edu/gormsen. Huber is at The University of Chicago; www.kilianhuber.org, and Jensen is at Yale University; www.tijensen.com.

1 Introduction

A fundamental principle in economics is that firms should determine the opportunity cost of their capital before investing. The idea goes back to at least Adam Smith who described individuals "exerting [...] to find out the most advantageous employment for whatever capital." The intuition is simple. The opportunity cost of an investment—known as the "cost of capital"—is the expected return on an alternative, external investment that has similar risk to the firm's investment. If that cost of capital is high, the firm should forgo its investment and instead reallocate the capital toward the external investment. By following this rule, the firm maximizes its economic profits and thereby its market value.

Modern scholars and practitioners take this principle as given. Essentially all academic models assume that firms make investment decisions based on the cost of capital, including neoclassical frameworks like Tobin's Q and Modigliani-Miller, models with financial frictions, and behavioral models with irrational expectations. In the real world, firm executives attempt to incorporate the cost of capital in their decision-making, inspired by decades of business school teaching and consultants.

Despite the importance of the cost of capital in academia and practice, fundamental issues remain. First, the cost of capital is not directly observed, since it refers to the hypothetical "most advantageous" return of a similar external investment. There is no general method that estimates a firm's cost of capital without bias (i.e., without predictable forecast errors). Second, there is no dataset measuring an unbiased cost of capital at the firm level, so it is unclear which measure should inform the real-world decisions of firms and academic analyses. Third, there exists no evidence validating the classical intuition that firms improve their market value when they use an appropriate cost of capital.

We contribute by overcoming all three issues. First, we develop a new method that uses machine learning to forecast expected cash flows and then solves for the cost of capital. The method estimates a firm's "market-based cost of capital" (MCC), which has two additional, equivalent interpretations: it is the expected return on an external investment with similar risk to the firm's typical investment and the rate used by the stock market to discount the firm's cash flows.

Second, we implement the method and present firm-level data on the estimated MCC. The MCC is unbiased, in the sense that it predicts the returns on the firm's assets in financial markets without predictable errors. The MCC differs substantially from existing biased proxies for the cost of capital, including those based on the Capital Asset Pricing

Model (CAPM), the implied cost of capital (ICC), and interest rates.

Third, we evaluate the classical intuition that the cost of capital determines the value of firms. We compare the MCC to the perceived cost of capital actually used by firms in their investment decisions. We find that firm market value is greater when the perceived cost of capital is closer to the MCC. The finding is consistent with the classical intuition that firms maximize their market value when they use an unbiased estimate of the cost of capital. It suggests that adopting the MCC in lieu of their biased perceptions may allow firms to improve their market value.

We begin the paper by defining the object of interest and clarifying its economic role. The cost of capital relates to expected returns in financial markets, firm value, and investment decisions, but measuring it at the firm level is difficult in practice. Existing approaches, like the CAPM and the implied cost of capital, do not provide a general unbiased estimate of the cost of capital because they are biased with respect to realized returns in financial markets.

2 Conceptual Overview

In this section, we define the cost of capital that we are interested in measuring, before describing our method for doing so.

2.1 The Cost of Capital and Expected Returns

Financial investors provide capital to the firm in the form of debt and equity. The total debt of the firm is D and total equity is E. The firm invests its capital in assets with an expected long-run rate of return r_t^{firm}. In addition, the firm deducts its payments to debt holders from its taxes, so it also earns $\text{D} \times \tau \times r_t^{\text{debt}}$, where τ is the tax rate and r_t^{debt} the debt rate. In total, the firms' expected return on capital is:

$$\text{ROR}_t = \frac{(\text{D+E}) \times r_t^{\text{firm}} + \text{D} \times \tau \times r_t^{\text{debt}}}{\text{D+E}} \tag{1}$$

$$= r_t^{\text{firm}} + \frac{\text{D}}{\text{D+E}} \tau r_t^{\text{debt}}. \tag{2}$$

Instead of investing in the firm, debt holders could purchase an alternative, external debt security that has the same risk and long-run term as the debt of the firm. For simplicity, assume that the external debt generates an expected return equal to the firm's debt rate of r_t^{debt}. Similarly, equity holders could invest in an external equity security that has the same risk and term as the firm's equity and generates an expected return of r_t^{equity}. If all the firm's investors reallocated their capital to the external securities, their expected return would be:

$$\text{OCC}_t = \frac{\text{D} \times r_t^{\text{debt}} + \text{E} \times r_t^{\text{equity}}}{\text{D+E}}, \tag{3}$$

which is the investors' opportunity cost of the capital (OCC) invested in the firm.

Assuming the law of one price holds, investments with the same risk and term yield the same expected return. Hence, the expected return on the firm's debt equals r_t^{debt}, which is also known as the firm's "cost of debt." Similarly, the expected return on the firm's equity is r_t^{equity}, called the firm's "cost of equity."

The law of one price also implies that the return on investing in the firm (ROR) is identical to the opportunity cost of capital (OCC), as both have the same risk and term. By equalizing (2) and (3), we can solve for the expected equilibrium return on the firm's assets:

$$r_t^{\text{firm}} = \frac{\text{D}}{\text{D+E}}(1-\tau)\, r_t^{\text{debt}} + \frac{\text{E}}{\text{D+E}}\, r_t^{\text{equity}} = \text{WACC}. \tag{4}$$

This formulation shows that the firm can use the expected returns on debt and equity to determine the return that financial investors expect the firm's assets to generate. This expected return r_t^{firm} is known as the firm-level after-tax cost of capital or the weighted average cost of capital (WACC). Intuitively, it captures the funding cost of a potential investment that is similar to the firm's existing assets. The measurement in our paper estimates the WACC at the firm level. In the remainder of the paper, we use the terms "cost of capital" and WACC interchangeably.

2.2 The Cost of Capital, Firm Value, and Real Investment

The cost of capital plays an important role in academia and corporate practice. First, the cost of capital determines firm value. The net present value v_t of a firm equals the sum of expected free cash flows discounted by the matching expected return on the firm's assets. Assuming that the same expected return applies to all future cash flows, the cost of capital is the appropriate discount rate:

$$v_t = \sum_{h=1}^{\infty} \frac{\mathbb{E}_t[x_{t+h}]}{(1+r_t)^h}, \tag{5}$$

where x_t is the expected free cash flow to debt and equity holders and r_t the cost of capital. This formulation assumes that firms are infinitely-lived, following the literature and practical valuations (Koller et al. 2020; Rosenbaum and Pearl 2021).

The cost of capital also influences firm investment decisions. Imagine a firm considering a long-run investment that has expected return r_t^{project} and is otherwise similar to the existing assets. If $r_t^{\text{project}} > WACC$, the expected return of the investment exceeds that of similar external securities. If the firm invests, it can expect to pay the debt rate to its debt holders and to have a greater return left for its equity holders than comparable securities. As a result, the equity market value of the firm increases. In contrast, if $r_t^{\text{project}} < WACC$ and the firm invests, the market value of the firm falls. This logic implies that firms should base their investment decision on the cost of capital (WACC) if they want to maximize market value. Standard academic models assume that firms apply this rule, including neoclassical frameworks, models with frictions, and behavioral models with irrational agents.[1]

Executives are widely taught to use the cost of capital, based on expected returns as in (4), in their investment decisions. Surveys suggest that the vast majority of listed firms indeed know and regularly update a "perceived cost of capital," their internal estimate of the WACC (e.g., Trahan and Gitman 1995; Graham and Harvey 2001).

In practice, firms use a required return on investment, called discount rate or hurdle, as the minimum expected return that potential investments need to meet. Firms typically set discount rates using the perceived cost of capital as the basis and then adding additional terms, for example, to account for the idiosyncratic risk of potential projects, real option values, inflated return expectations provided by division managers, or financial constraints. Although discount rates often do not exactly equal the perceived cost of capital, existing work shows that the perceived cost of capital is a key input that determines the allocation of capital and returns across firms. For instance, the perceived cost of capital is strongly correlated with discount rates, real investment rates, and realized returns on invested capital across firms (Gormsen and Huber 2025b). Moreover, in the long run, changes in the perceived cost of capital strongly affect discount rates (although the short-run incorporation is weak, Gormsen and Huber 2025a).

[1] If the firm's objective is not to maximize current market value, it may prefer an alternative return threshold, for instance, a future expected return (Stein 1996) or a fundamental risk benchmark (Gormsen and Huber 2025b).

2.3 The Measurement Challenge

Taken together, the cost of capital is a key input into academic models and the real decision-making of firms. However, expected returns on securities are fundamentally unobserved. They depend on the risk perceptions and preferences of financial investors. As a result, firms do not know their cost of capital. The CAPM is often used in practice to approximate the cost of capital, but it is well known that the CAPM is biased with respect to expected returns.

In general, there exists no method that estimates an unbiased cost of capital and no readily available firm-level dataset. This absence is especially glaring given that, according to standard teaching, firms need to know their cost of capital to maximize their market value. In the remainder of this paper, we develop an estimation method of the cost of capital, present firm-level estimates, and finally evaluate its economic relation to market value.

3 Method to Estimate a Market-Based Cost of Capital

3.1 Overview of the Method

Our aim is to estimate the firm-level cost of capital, given by (4). There exist established estimates for firm leverage, tax rates, and the cost of debt, which we adopt in our implementation in Section 4. Our method focuses on the remaining term: the cost of equity. On average, listed firms are financed with two-thirds equity, so the cost of equity is a key component of the firm-level cost of capital.

Analogous to the firm's overall value in (5), the firm's equity market value p_t is:

$$p_t = \sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{fcfe}_{t+h}]}{(1 + r_t^{\text{equity}})^h}, \tag{6}$$

where $\mathbb{E}_t[\text{fcfe}_{t+h}]$ is a forecast for the free cash flow accruing to equity holders in period $t+h$. Firm market value p_t is directly observed, but we do not easily observe cash flow forecasts that capture substantial variation. The key challenge in estimating the cost of equity r_t^{equity} is therefore determining cash flow forecasts.

Our method focuses on the estimation of cash flow forecasts with strong predictive power, which can then be used to derive the cost of equity.

3.2 Forecasting Cash Flows and Net Present Value

Since we cannot estimate infinite-horizon cash flows, we define an explicit forecasting horizon of H periods and a terminal value period. We impose the terminal value condition that all firms' cash flows grow at the same constant rate g_t:

$$p_t = \sum_{h=1}^{10} \frac{\mathbb{E}_t[\text{fcfe}_{t+h}]}{(1+r_t^{\text{equity}})^h} + \frac{\mathbb{E}_t[\text{fcfe}_{t+10}](1+g_t)}{r_t^{\text{equity}} - g_t} \frac{1}{(1+r_t^{\text{equity}})^{10}}. \tag{7}$$

In our estimation, we set $H = 10$. The terminal value condition can be motivated as an economic prior that competition eliminates above-average growth in the long run. Consistent with this view, the terminal growth rates used in valuation practice do not vary much across firms (Decaire and Guenzel 2023).

We need to forecast cash flows in every period t for all horizons $h \leq 10$. We employ a machine-learning technique based on decision tree learning. This technique essentially splits firms into groups based on potential predictors of cash flows, and then estimates which groups are associated with higher future cash flows. However, traditional tree learning is time-intensive, and would be infeasible for a large number of firms, periods, and horizons. We therefore build on the LightGBM model of Ke et al. (2017), which allows running many simultaneous decision trees efficiently and fast.[2] Since we need forecasts made at time t without look-ahead bias, we only use data available up to four months before t for each forecast. We avoid in-sample overfitting by relying on out-of-sample predictions, as detailed in our implementation Section 4.

3.3 Cash Flows Forecasts and the MCC

With the cash flow forecasts in hand, we can solve (7) for the "market-based cost of equity" (MCE), our estimate of r_t^{equity}. This MCE equates the discounted sum of expected cash flows with the observed market value. It is therefore the expected return on the firm's equity in financial markets.

We insert the MCE into the definition of the cost of capital (4). Combined with conventional measures of firm leverage, tax rates, and the cost of debt, we calculate our estimate of the firm-level cost of capital. We term it the "market-based cost of capital" (MCC) since it

[2] The model is a gradient-boosted decision tree ensemble (GBDT). Such GDBTs perform well on tabular data, outperforming deep learning methods across a wide variety of tabular benchmarks (Shwartz-Ziv and Armon 2022).

is consistent with observed market values. Estimating the MCC requires cash flow forecasts with strong predictive power. If we instead used weak forecasts, the resulting cost of capital would be severely biased.

Following standard asset pricing techniques, a tempting alternative approach might be to directly predict the cost of equity, for example, using a factor model. However, this approach is tailored to estimating short-run expected returns. The cost of equity depends on long-run returns that match the horizon of firms' investments, usually at least a decade long. Existing return data do not cover many decades, so this approach is infeasible.

Our method is inspired by seminal work calculating an "implied cost of capital" (e.g., Gebhardt et al. 2001; Claus and Thomas 2001; Easton 2004; Ohlson and Juettner-Nauroth 2005), which relies on a net present value formulation and analyst forecasts to derive the cost of equity. Since analyst forecasts are biased, the implied cost of capital is also biased. This has made it challenging to use it to assess the importance of the cost of capital for market value.

4 Estimating a Market-Based Cost of Capital

4.1 Data

We construct a database for around 16,000 firms listed in 93 countries. We start the samples in 1950 for U.S. stocks and 1985 for non-U.S. stocks to ensure that accounting data are available. Data on market values originally come from CRSP for U.S.-listed firms and from Compustat for non-U.S.-listed firms. The accounting data are from Compustat. We rely on the database provided by Jensen et al. (2023) through WRDS and GitHub. Finally, we add analyst forecasts from I/B/E/S, and we always use the median forecasts from the consensus file. The data set is global, covering stocks listed in 93 countries.

We exclude financial firms (i.e., groups 45–48 in the Fama and French 1997 classification) because free cash flows for financial firms are conceptually different from free cash flows from non-financial firms.[3] We also require firms to have a non-missing and positive accounting value of assets and shares outstanding.

[3]We aim to estimate firms' unlevered free cash flows by separating cash flows due to operations from cash flows due to financing. However, this separation is difficult for financial firms because their business model is all about financing.

4.2 Estimating the cost of equity (MCE)

Our method requires cash flow forecasts to determine the cost of equity. We forecast real cash flows in U.S. dollars. The forecasts are made in every month t for up to ten years in the future. The terminal growth rate for cash flows is the IMF's 10-year real GDP growth forecast in the country.

We employ two measures of cash flows to equity holders: free cash flows to equity and residual income. Details on the data construction are in Appendix B.

When using the free cash flow to equity, we directly rely on (6). The free cash flow to equity is:

$$\text{fcfe}_t = e_t - \Delta\text{oa} + \Delta\text{ol} + \Delta\text{debt}, \tag{8}$$

where e_t is company earnings (net income), Δoa_t is the change in operating assets, Δol_t is the change in operating liabilities, and Δdebt_t is net debt issuance. We forecast free cash flows to equity in $t+1, \ldots, t+10$, and solve for r_t^{equity} numerically.

When using the residual income measure (Edwards and Bell 1961; Ohlson 1995), we rely on a slightly modified net present value formulation:

$$p_t = b_t + \sum_{h=1}^{10} \frac{\mathbb{E}_t[\text{ri}_{t+h}]}{(1+r_t^{\text{equity}})^h} + \frac{\mathbb{E}_t[\text{ri}_{t+10}](1+g_t)}{r_t^{\text{equity}} - g_t} \frac{1}{(1+r_t^{\text{equity}})^{10}}, \tag{9}$$

where b_t is the firm's book equity and ri_t is residual income:

$$\text{ri}_t = e_t - r_t^{\text{equity}} b_{t-1}. \tag{10}$$

To estimate the cost of equity from this model, we forecast net income in $t+1, \ldots, t+10$, and infer the future book equity as:

$$b_t = b_{t-1} + e_t - d_t, \tag{11}$$

where d_t is the firm's dividends. We estimate dividends by multiplying our earnings forecast with the firm's most recent payout ratio (d_t/e_t).[4] The method for inferring the future book equity is inspired by Gebhardt et al. (2001). We then solve for r_t^{equity} numerically.

We fit a separate prediction model for each of the two cash flow measures and each of the

[4]Our method for inferring the future book equity can result in negative book equity, in which case the residual income from (10) is not well defined. We set the residual income to zero in cases where the book equity is negative.

ten horizons. Each model pools information across all countries. Once we have estimated the two models separately, we

The explanatory variables in the prediction models are 153 characteristics from Jensen et al. (2023); 75 characteristics we derive from their data set related to current cash flow levels and changes; 12 characteristics based on analyst forecasts from IBES related to the expected cash flows for the firm over the next three years and the analyst long-term growth rate forecast, as well as the corresponding median within the firm's industry and country, respectively; 45 industry dummies based on the classification from Fama and French (1997), expanded to also include the computer software group, but without firms from the four financial groups; and 93 country dummies. In total, the models are fit on 240 continuous and 138 dummy characteristics. Some of the explanatory variables—especially those that require analyst data—are missing for a large fraction of our sample. Removing observations with missing explanatory variables would thus drastically reduce our sample and tilt it towards relatively large stocks with analyst coverage. Fortunately, LightGBM handles missing values natively by learning how missing values map to the dependent variable. As a result, we keep observations with missing values and delegate their handling to the learning algorithm.

We avoid in-sample overfitting. The model is trained on an expanding training sample, with the last 10 years used as the validation sample. We select hyperparameters by evaluating 20 candidate sets, chosen to maximize entropy over a predefined search space spanning five specific hyperparameters, and select the ones with the lowest mean squared error in the validation sample. When the hyperparameters have been found, we re-estimate the model on the combined training and validation data. The first model is trained on a data set ending on December 31st, 1983, and we re-estimate the model parameters and hyperparameters every year. All forecasts are made out-of-sample (OOS) on a monthly basis, with the first forecast made on January 31st, 1984, and the last on December 31st, 2023.

4.3 Estimating the cost of debt

We observe traded bond yields for U.S. firms, so we can calculate the cost of debt as:

$$\mathbb{E}_t[r_{t+1}^{\text{debt}}] = \text{yield}_t - [\text{prob. of default}]_t \times (1 - [\text{recovery rate}]_t). \tag{12}$$

We compute the yield as the weighted average over all the firm's outstanding bonds using bond market value as weights. Data on bond yields are from Dick-Nielsen et al. (2023), based on dealer quotes over 1973-2001 and traded prices over 2002-2023. The probability of default is the average realized default rate of bonds with the same rating over the past 3 years, following Campello et al. (2008). The recovery rate is again rating specific and from Altman and Kishore (1998). To adjust for differences in bond maturity, we subtract the duration-matched risk-free rate and add the ten-year U.S. treasury yield:

$$r^{\text{debt}} = \mathbb{E}_t[r_{t+1}] - r^f_{\text{duration-matched}} + r^f_{10}. \tag{13}$$

This serves as our cost of debt estimate for firms with traded bonds.

We project the cost of debt for firms with traded bonds on their distance-to-default DtD_t, estimated as in Bharath and Shumway (2008):

$$r^{\text{debt}}_t = a_t + b_t \times \text{DtD}_t + \varepsilon. \tag{14}$$

We then use the estimates from this regression to predict the cost of debt for firms without traded bonds.

4.4 Estimating the firm-level cost of capital (MCC)

We calculate the MCC, our market-based estimate of the firm-level cost of capital, using the WACC formula (4). We directly observe the market value of equity and proxy for the market value of debt using the book value of debt.[5] We estimate the tax rate as an equal-weighted average of the firm's effective tax rate and the median effective tax rate of all firms in the same year and country. The firm's effective tax rate is Compustat item `txt` divided by `pi`, capped at zero and one, to control for outliers (e.g., due to carry-forward tax losses).

5 Properties of Cash Flow Forecasts and the MCC

We show that the new cash flow forecasts have relatively strong predictive power for realized cash flows. Accordingly, the predicted cost of equity is an unbiased estimator of the firm-level cost of capital.

[5]In some textbooks, the leverage measure is the firm's long-run target leverage (Ross et al. 2022). Under that view, one can view our market-based current leverage measure as a proxy for the target leverage.

Figure 1
Cash flow forecasts

This figure shows the out-of-sample R^2 of the cash flow prediction models. We consider two different cash flow variables, as indicated by the panel header, and ten forecast horizons, as indicated by the x-axis. The results are based on global stocks. The sample is balanced, meaning that stocks must have non-missing realized and predicted cash flow for both variables and all ten horizons. Figure A2 shows the results from the unbalanced sample.

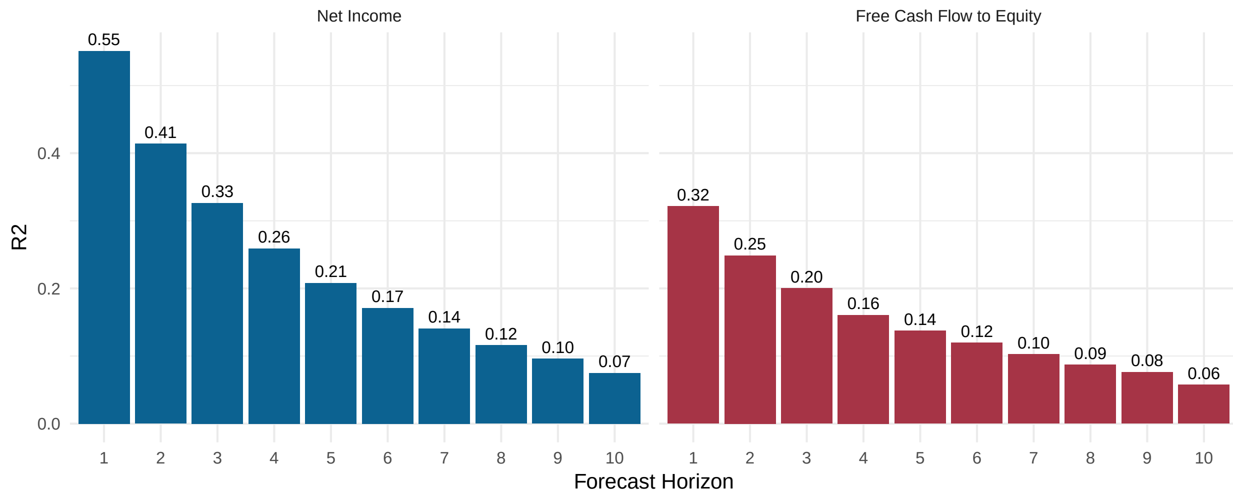

5.1 Cash flow predictability

Figure 1 shows the out-of-sample R^2 of the cash flow prediction models for each of the two cash flow targets and ten forecast horizons. The R^2 is positive across all the outcome variables used in our prediction model. A positive R^2 is not guaranteed when conducting out-of-sample predictions. We predict short-horizon cash flows more accurately than long-horizon cash flows. Furthermore, net income is more predictable than free cash flows to equity because accounting rules smooth underlying cash flows. For example, a large investment strongly reduces free cash flows to equity in one year, but has weaker effects on net income over several years due to depreciation allowances. Figure A3 shows that the R^2 is relatively stable over time, with a slight upward trend.

We compare our cash flow forecasts to existing benchmarks from the literature. Few papers predict long-run cash flows, so we mainly consider earnings forecasts.[6] The first two

[6] Analysts typically produce earnings forecasts one quarter or one year ahead. For papers trying to forecast short-run cash flows with machine learning, see Hillenbrand and McCarthy (2023), Van Binsbergen et al. (2023), Zhang et al. (2024), and De Silva and Thesmar (2024). Among these, the longest horizon considered is four years ahead, by De Silva and Thesmar (2024).

benchmark methods we consider come from Claus and Thomas (2001, CT) and Gebhardt et al. (2001, GLS), which are part of the analyst-based ICC literature. Both methods take analyst forecasts over the next two years as input, but differ in how they extrapolate cash flows beyond year two. CT assumes that cash flows grow at analysts' long-term growth rate forecasts, while GLS assumes that each firm's return on equity converges linearly to its corresponding industry average over a 12-year period. The third model is a vector autoregression (VAR) method similar in spirit to Vuolteenaho (2002), Campbell et al. (2010), and Cho et al. (2024). Finally, we consider the random walk model, which assumes that all future cash flows are equal to the most recent realization. The random walk (RW) model of earnings is supported by early research in accounting (Ball and Watts 1972), and earnings surprises are often defined with respect to the RW model (see, e.g., Foster et al., 1984 and Martineau, 2022).

Figure 2 shows that the out-of-sample R^2 of the MCC cash flow forecasts exceeds that of all the benchmarks. At a one-year horizon, the difference relative to CT and GLS (which at this horizon are effectively analyst forecasts) is relatively small.[7] However, by year three, both have a negative out-of-sample R^2. The VAR model performs slightly worse than the ICC methods at the one- and two-year horizons, but maintains a positive R^2 at longer horizons. The RW model performs relatively poorly at short horizons but is the best benchmark at the ten-year horizon. Overall, the results suggest that the MCC cash flow forecasts outperform existing methods in the literature—especially at the long horizons that are key for the cost of capital.[8]

5.2 Return predictability

The market-based cost of equity (MCE) is meant to predict returns over a very long-run horizon. We estimate the association between the MCE and future ten-year realized equity

[7] The dependent variable in Figure 2 is for earnings realized over a constant yearly horizon, but the analyst forecasts used by CT and GLS are for the next fiscal year, which is typically closer than one year. Figure A7 shows that the ML model also beats analysts when we map their forecasts to a constant one-year horizon. It also shows that the analysts' out-of-sample R^2 improves when the forecast is demeaned to remove bias arising from well-known analyst conflicts of interest (Richardson et al. 2004). However, even after demeaning the forecasts and outcome, the ML model remains superior.

[8] Figure A1 show a binned scatter plots of predictions relative to realizations for each of the models. The CT, GLS, and RW models have systematic biases for various parts of the distributions, whereas both our method and VAR appear approximately unbiased. Our method outperforms VAR because it remains unbiased while making more dispersed predictions. For example, at the ten-year horizon, our method predicts earnings to assets of -0.2 for the bottom percentile of firms and 0.3 for the top percentile. By contrast, the comparable spread in VAR predictions is only 0.01 to 0.12.

Figure 2
Cash flow predictability

The figure shows the out-of-sample R^2 for net income forecasts at horizons 1 through 10 years for different methods. The machine learning model is LightGBM trained on firm characteristics; the benchmarks are implied cash flow forecasts from the Claus and Thomas (2001), Gebhardt et al. (2001), and a VAR model similar to the one in Vuolteenaho (2002), as well as a random walk (last observed net income). All forecasts are scaled by total assets and winsorized at the 1st and 99th percentiles. The sample is balanced, requiring all methods to have non-missing predictions for each firm-month-horizon observation.

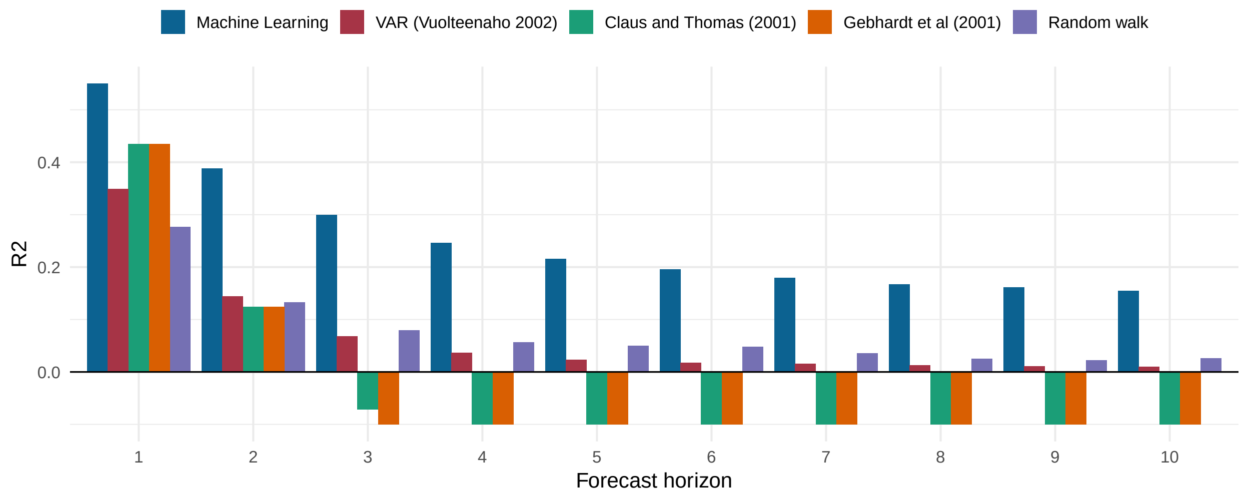

returns:

$$r^{\text{realized},i}_{t\to t+120} = \alpha_{c,t} + b \times \left[\left(1 + \text{MCE}^i_t\right)^{10} - 1\right] + \varepsilon^i_{t+h}, \tag{15}$$

where $r^i_{t\to t+120}$ is the buy-and-hold realized return from t to $t+120$ months (ten years), $\alpha_{c,t}$ is a country-by-time fixed effect, and MCE^i_t is the market-based cost of equity converted to a ten-year horizon. Standard errors are clustered by industry and time.

One challenge with analyzing ten-year returns is that some firms delist, potentially introducing survivorship bias in the sample. Bessembinder (2018), for example, shows that the median time a stock is listed in the CRSP database is 7.5 years and that firms often delist for performance-related reasons. To address this issue, we include all delisted firms in the sample, even if they delist during a ten-year period. We impute returns after delisting using the median return for firms in the same country and month, and continue compounding this imputed return. The imputation approach has no impact on a firm that goes bankrupt and has a delisting return of -100%, but ensures that firms that, say, get acquired continue to accumulate returns.[9]

[9]In unreported results, we find that our results are weaker if we remove stocks that delist from the

Figure 3

The MCE, realized equity returns, and alternative benchmarks

This figure shows the relationship between future 10-year realized returns and the current 10-year cost of equity for different estimation methods. MCE is the market-based cost of equity from this paper. CAPM computes the cost of equity as the risk-free rate plus beta times an expanding-window equity risk premium. MLR predicts 10-year cumulative excess returns directly using LightGBM, and then adds the current risk-free rate. VAR uses a firm-level vector autoregression similar to the one used in Vuolteenaho (2002). CT, GLS, PEG, and OJ, are analyst-based ICC estimates following Claus and Thomas (2001), Gebhardt et al. (2001), Easton (2004), Ohlson and Juettner-Nauroth (2005), respectively. Each month, stocks are sorted into 100 portfolios by their estimated cost of equity. For each portfolio, we compute the average cost of equity and the average realized buy-and-hold return over the next 10 years, and then average across months. Both axes are annualized. The solid red line shows the best linear fit, and the dashed blue line shows the 45-degree line. The sample includes all global stocks, 1983–2023.

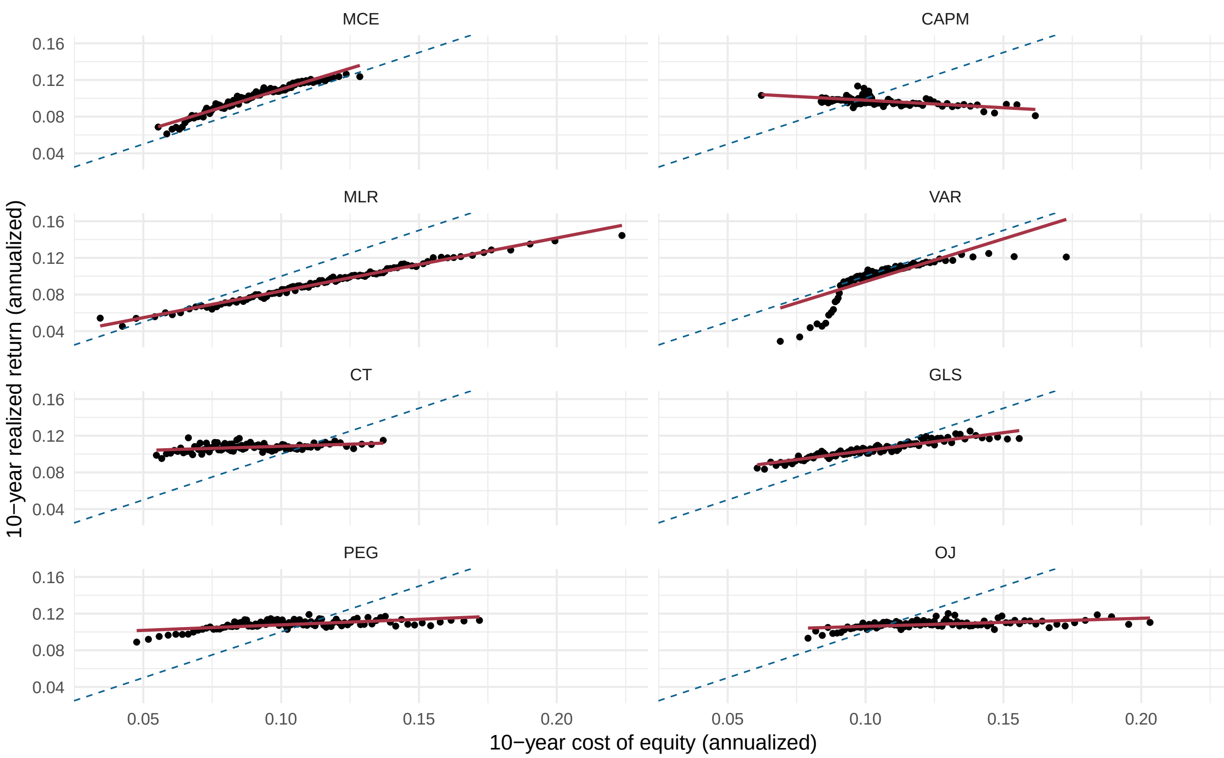

We plot a binned scatter plot showing the association between the MCE and future realized ten-year equity returns in the top-left panel of Figure 3. The relation is strong, as all points are close to the 45-degree line. The spread in the MCE is 8% and the spread in realized ten-year returns is 10%, suggesting that the MLE is relatively well calibrated.

We also plot the relation between alternative benchmarks for the cost of equity and realized returns. We first consider a method that we call ML-direct (MLR), which uses our prediction model to directly estimate ten-year predicted equity returns, rather than first predicting cash flows and then deriving the MLE. The MLR predicts returns positively, but with a lower slope than the MLE. It also has a spread of 25%, far too large relative to the realized return spread of 10%.

The other benchmark methods perform more poorly than the MLE and MLR. The GLS-based implied cost of capital has the relatively most positive slope and the slope of the CAPM is negative.

Table 1, Panel A, reports regressions of ten-year realized returns on the MCE, controlling for month fixed effects. The point estimates in the U.S. and the global sample are close to 1 and statistically indistinguishable from 1. The estimates are also significantly different from 0. The MCE is an unbiased estimate of expected returns. Panel B reports the average error, that is, the ten-year realized minus predicted return. In the U.S. sample, the average error is a significant 25% (approximately 2% annualized). This result suggests that realized returns in the out-of-sample period exceeded investors' ex ante expectations in the U.S., consistent with Binsbergen et al. (2025). In the global and non-U.S. samples, however, the average errors are closer to zero and statistically insignificant, indicating no systematic errors overall.

Table 1, Panel B, also shows the out-of-sample R^2 relative to two benchmark models, defined as:

$$R^2_{\text{bench}} = 1 - \frac{\sum(r_i - \hat{r}_i)^2}{\sum(r_i - \hat{r}_{i,\text{bench}})^2}, \tag{16}$$

where r_i is the realized return, $\hat{r}_i$ is the out-of-sample MCE prediction, and $\hat{r}_{i,\text{bench}}$ is the out-of-sample benchmark prediction. The more standard in-sample R^2 uses the in-sample mean instead of the benchmark prediction, but the in-sample mean is not available in real-time and is therefore an inappropriate benchmark for our out-of-sample predictions. Instead, we consider two alternative cost of equity estimates that could have been used in real time. MKT is the ten-year risk-free rate at the estimation date plus a rolling estimate of the

sample, because firms that go bankrupt are relatively more likely to have a low MCC. Further, the results are relatively insensitive to whether we impute using the cross-sectional median or mean.

	10-year realized return		
	U.S.	Global	Global ex-U.S.
	(i)	(ii)	(iii)
Panel A: Regression estimates			
MCE	0.93***	0.96***	1.01***
	(0.08)	(0.18)	(0.23)
Observations	1,106,464	2,610,489	1,504,025
Panel B: Predictive accuracy ($\alpha = 0, \beta = 1$)			
Avg. error	0.25	0.14	0.06
	(0.12)	(0.10)	(0.15)
OOS R^2_{CAPM} (%)	2.1	3.8	5.3
OOS R^2_{MKT} (%)	0.4	2.1	4.1

The table reports regressions of 10-year realized buy-and-hold returns on the current 10-year cost of equity, estimated separately for U.S., global, and global ex-U.S. samples. All regressions include month fixed effects. The average pricing error is the mean difference between realized returns and the predicted cost of equity. OOS R^2_{CAPM} and R^2_{MKT} report out-of-sample predictive R^2, computed as $1 - \sum(r_i - \hat{r}_i)^2 / \sum(r_i - \hat{r}_{i,\text{bench}})^2$, where the benchmark prediction is either the CAPM-implied cost of equity or the country-level expanding-window average market return. Both benchmarks use only information available at the time of prediction. We report Driscoll-Kraay standard errors using a 120 month window in parentheses. Statistical significance is denoted by *** p<0.01, ** p<0.05, * p<0.1.

Table 1: Return predictability

historical ten-year local market excess return based on data through this date. CAPM is the 10-year risk-free rate plus the same rolling estimate of the local market risk premium, multiplied by the stock's market beta, using the last 60 months of data. The out-of-sample R^2s are all positive, suggesting that an investor would achieve higher accuracy using the MCE than either the historical market return or the CAPM.

Table 2 reports point estimates for the other methods. The point estimate for the MLE is the only one close to 1. The only other benchmarks with a significantly positive point estimate are MLR, VAR, and GLS. Further, Table A1 shows that the mean-squared error of the MLE is lower than that of the other benchmarks.

	10-year realized return							
	MCE	CAPM	MLR	VAR	CT	GLS	PEG	OJ
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
Panel A: Regression estimates								
Cost of equity	0.96***	-0.13	0.50***	0.66***	0.12	0.43**	0.10	0.07
	(0.18)	(0.13)	(0.16)	(0.21)	(0.24)	(0.20)	(0.10)	(0.10)
Observations	2,610,489	3,724,501	4,872,824	3,463,929	974,231	1,852,071	1,000,734	985,775
Panel B: Predictive accuracy ($\alpha = 0, \beta = 1$)								
Avg. error	0.14	0.21	-0.19	-0.04	0.34	-0.10	-0.05	-0.56
	(0.10)	(0.14)	(0.28)	(0.09)	(0.09)	(0.09)	(0.10)	(0.08)
OOS R^2_{CAPM} (%)	3.6	0.0	0.9	2.8	1.5	3.0	1.3	-1.3
OOS R^2_{MKT} (%)	2.1	-1.1	-0.5	1.6	0.0	1.2	-0.5	-2.6
Panel C: MSE relative to MCE								
ΔMSE	—	0.74	0.94	0.14	0.20	0.07	0.34	0.88
$p^{\text{bs}}_{\text{ind}}$		0.000	0.000	0.000	0.000	0.063	0.000	0.000
$p^{\text{bs}}_{\text{time}}$		0.001	0.004	0.083	0.011	0.243	0.048	0.010

The table reports regressions of 10-year realized buy-and-hold returns on the current 10-year cost of equity for eight estimation methods. MCE is the market-based cost of equity from this paper. CAPM computes the cost of equity as the risk-free rate plus beta times an expanding-window country-specific equity risk premium. MLR predicts 10-year cumulative excess returns directly using LightGBM, and then adds the current risk-free rate. VAR uses a firm-level vector autoregression similar to the one used in Vuolteenaho (2002). CT, GLS, PEG, and OJ, are analyst-based ICC estimates following Claus and Thomas (2001), Gebhardt et al. (2001), Easton (2004), Ohlson and Juettner-Nauroth (2005), respectively. All regressions include month fixed effects. The average pricing error is the mean difference between realized returns and the predicted cost of equity. OOS R^2_{CAPM} and R^2_{MKT} report out-of-sample predictive R^2, computed as $1 - \sum(r_i - \hat{r}_i)^2 / \sum(r_i - \hat{r}_{i,\text{bench}})^2$, where the benchmark prediction is either the CAPM-implied cost of equity or the country-level expanding-window average market return. Both benchmarks use only information available at the time of prediction. Panel C reports $\Delta\text{MSE} = \text{MSE}_{\text{method}} - \text{MSE}_{\text{MCE}}$, computed on the id-month pairs where both MCC and the comparison method have non-missing predictions. Positive values indicate that MCE has lower MSE. $p^{\text{bs}}_{\text{ind}}$ and $p^{\text{bs}}_{\text{time}}$ are one-sided bootstrap p-values for the null that $\Delta\text{MSE} \leq 0$, obtained by resampling Fama-French 49 industries and 10 equal-sized time blocks, respectively, with replacement (10,000 iterations). We report Driscoll-Kraay standard errors using a 120 month window in parentheses. Statistical significance is denoted by *** p<0.01, ** p<0.05, * p<0.1.

Table 2: Return benchmarks – unbalanced – fixed effects

6 The Importance of the Cost of Capital for Firm Value

Using a simple model, we formalize the classical intuition that firms achieve a higher market value when they use an unbiased cost of capital, like the MCC, in their investment decisions. We confirm that there is an inverted U relationship in the data: firms have greater market value when their perceived cost of capital is closer to the MCC.

6.1 Model of inverted U

We sketch a simple model to formalize how firms' perceptions about their cost of capital affect firms' capital allocation and equity market value.

We consider an infinite-horizon discrete-time model with periods $t = 0, 1, 2, \ldots$ and a cross-section of firms indexed by i. Each firm i owns capital stock $K_{i,t}$ which depreciates at rate $\gamma \in (0, 1)$ per period. The law of motion for capital is:

$$K_{i,t+1} = (1 - \gamma)K_{i,t} + I_{i,t}, \tag{17}$$

where $I_{i,t}$ is gross investment at date t.

In each period, capital generates stochastic output according to the production function:

$$Y_{i,t}(K_{i,t}) = (\alpha_i + \eta_{i,t})K_{i,t} - \frac{\beta_i}{2}K_{i,t}^2, \tag{18}$$

where $\alpha_i, \beta_i > 0$ are firm-specific productivity parameters and $\eta_{i,t}$ is an i.i.d. productivity shock with $E[\eta_{i,t}] = 0$.

Firm problem. Firm i maximizes the perceived present value of future cash flows using the perceived cost of capital $r_{i,t}^{\text{perc.}}$. The perceived cost of capital may differ from the firm-level cost of capital, r_i^{firm}, that financial investors use to discount firm i's cash flows. We model the deviation as:

$$r_{i,t}^{\text{perc.}} = r_i^{\text{firm}} + \delta_{i,t}. \tag{19}$$

For simplicity, we assume that r_i^{firm} is constant and that the cost of capital deviation, $\delta_{i,t}$, is orthogonal to r_i^{firm} and the fundamentals α_i and β_i.

Firm i chooses capital to maximize its perceived present value of cash flows:

$$\max_{\{K_{i,t}\}_{t=0}^{\infty}} \mathbb{E}\left[\sum_{t=0}^{\infty} \frac{1}{(1+r_i^{\text{perc.}})^t}\left[Y_{i,t}(K_{i,t}) - I_{i,t}\right]\right]. \tag{20}$$

Taking expectations over the productivity shocks and noting that firms choose a constant capital stock in the absence of adjustment costs, the firm's first-order condition implies that the optimal choice of capital, K_i^*, is given by:

$$K_i^* = \bar{K}_i - \frac{\delta_i}{\beta_i}, \tag{21}$$

where $\bar{K}_i \equiv (\alpha_i - r_i^{\text{firm}} - \gamma)/\beta_i$ is the capital the firm would choose if it used the correct discount rate ($r_i^{\text{perc.}} = r_i^{\text{firm}}$). In equilibrium, investment equals depreciation: $I_i^* = \gamma K_i^*$.

Market value. Investors price assets using a stochastic discount factor M_t. For firm i, the equity price satisfies:

$$P_i = \mathbb{E}\left[\sum_{t=1}^{\infty} M_t \cdot (Y_{i,t}(K_i^*) - \gamma K_i^*)\right]. \tag{22}$$

We assume that each firm is too small to influence the pricing of risk in the economy, such that the stochastic discount factor is not influenced by the choice of capital of the individual firm. We assume the stochastic discount factor satisfies conditions under which the multi-period discount factor can be represented as $(1+r_i^{\text{firm}})^{-t}$ for a constant firm-specific discount rate r_i^{firm} that reflects both time preference and systematic risk. These assumptions yield the value function,

$$P_i(K) = \sum_{t=1}^{\infty} \frac{1}{(1+r_i^{\text{firm}})^t}[Y_i(K) - \gamma K] = \frac{Y_i(K) - \gamma K}{r_i^{\text{firm}}}. \tag{23}$$

Result 1 (An Inverted U Between Market Value and Cost of Capital Wedges)**.** *The market value of firm i is given by*

$$P_i(K_i^*) = P_i(\bar{K}_i) - \frac{\delta_i^2}{2\beta_i r_i^{\textit{firm}}}, \tag{24}$$

where $P_i(\bar{K}_i)$ is the market value of the firm when the cost of capital deviation is zero, $\delta_i = r_i^{\textit{perc.}} - r_i^{\textit{firm}} = 0$.

Result 1 implies an inverted U relation: firm market value is maximized when firms use the same cost of capital as financial markets, r_i^{firm}, in their investment decisions. If firms use $r_i^{\text{perc.}}$ below or above r_i^{firm}, market value decreases. The inverted U is illustrated in Figure 4.

Figure 4

Model: Inverted U between firm market value and deviations from MCC

This figure shows the relation between market value and cost of capital deviations in our model. The model is calibrated such that market value is $1 when the cost of capital deviation is 0.

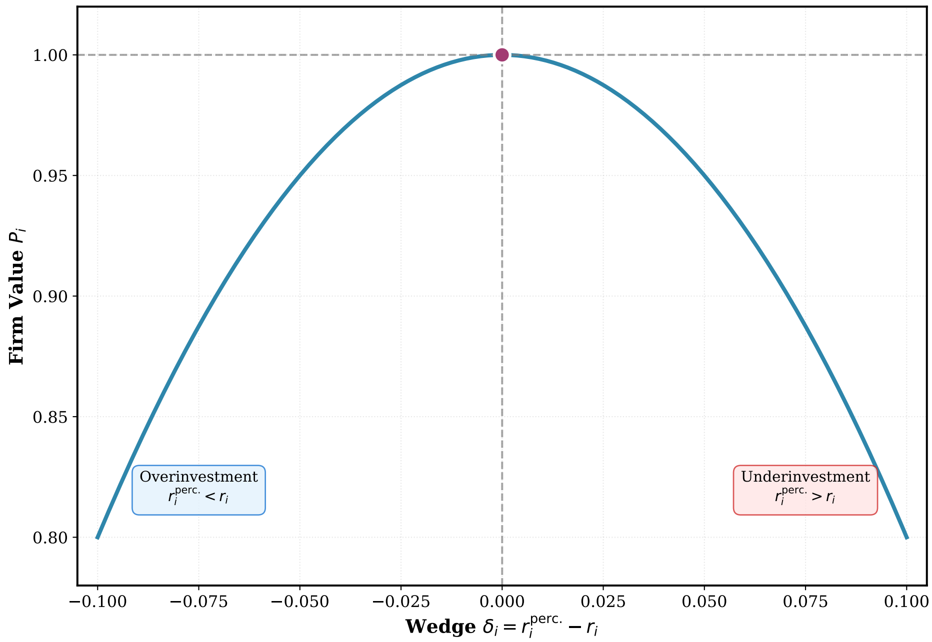

Our second result characterizes the relationship between investment and firm value in the cross-section of firms. From equation (6), we have $I_i^* = \gamma K_i^* = \gamma \bar{K}_i - \gamma \delta_i / \beta_i$. Since investment is proportional to capital, the correlation between investment and market value depends on the sign of the deviation.

Result 2 (Investment-value correlation)**.** *In the cross-section of firms, the correlation between investment and value depends on the sign of firms' cost of capital deviation*

$$\mathit{Corr}(I_i, P_i) > 0 \textit{ for firms for which } \delta_i > 0$$
$$\mathit{Corr}(I_i, P_i) < 0 \textit{ for firms for which } \delta_i < 0$$

Result 2 shows that the comovement between market values and investment depends on the sign of the cost of capital deviation. Firms with a positive deviation underinvest, and higher investment is therefore met with a positive response from investors. Firms with a negative deviation overinvest, and higher investment is therefore met with a negative response from investors.

6.2 Empirical results on inverted U

We investigate the relation between firm market value and the cost of capital deviation in the data. To measure the deviation, we calculate the difference between a firm's perceived cost of capital and its MCC. The "perceived cost of capital" is the cost of capital that influences the firm's real investment decisions. We use data collected by Gormsen and Huber (2025b) to measure the perceived cost of capital. The data are hand-collected from conference calls between firm managers, financial investors, and analysts. They measure explicit statements by managers on the calls about their perceptions of their firm-level weighted average cost of capital. In the data, firms with a higher perceived cost of capital have higher returns on invested capital and lower investment rates, consistent with the view that the perceived cost of capital influences capital allocation in the long run.

We are interested in whether firms with similar book equity are valued differently in equity markets depending on their cost of capital deviation. Figure 5 illustrates the average log market value for firms in different bins of the cost of capital deviation. The figure plots point estimates from a regression of log market value on the perceived cost of capital. The regression additionally controls for fixed effects for book equity deciles (relative to other firms in the same quarter-year) interacted with country fixed effects, which ensures that the point estimates capture differences in how effectively firms in the same country use a given level of book equity to generate market value. In addition, we control for quarter-by-year-by-country fixed effects, which ensures that macroeconomic shocks do not influence the estimates.

Figure 5 reveals an inverted U relation. Firm market value is lower when the cost of capital perceived by the firm (perc. CoC) deviates more from the MCC. This finding suggests that the cost of capital used by firms is strongly associated with market value, consistent with the classical intuition about the opportunity cost of capital and modern models of firm behavior.

Table 3 suggests that the association between the cost of capital deviation and market value is statistically and economically significant. In columns 1 to 2, the regressor is the absolute value of the deviation. The point estimates are all significant at the 1 percent level. They are of similar magnitude in the U.S. and the full sample.

In columns 3 to 5, we use the simple subtraction as the regressor. A one percentage

Figure 5

Inverted U: firm market value and deviations from the MCC

The figure shows the relation between firm market value and the cost of capital deviation (i.e., the difference between perc. CoC and MCC). The perc. CoC is reported by the firm on conference calls and influences the firm's investment decisions. MCC is the market-based cost of capital from this paper. We construct bins for the cost of capital deviation in percentage points. The bins are 0.2 percentage points wide, except for the leftmost bin, which includes all observations below -2.8 (roughly the 1st percentile), and the rightmost bin, which pools all observations above 2.8. The figure plots point estimates from a regression of firm log market value on indicators for the bins. The point estimates give the difference in log market value relative to the middle bin (difference just above 0). The regression controls for fixed effects for book equity deciles-by-country and quarter-by-year-by-country. Book equity deciles are based on the firm's median decile in the distribution of book equity relative to other firms in the same quarter-year. Table 3 reports related regressions.

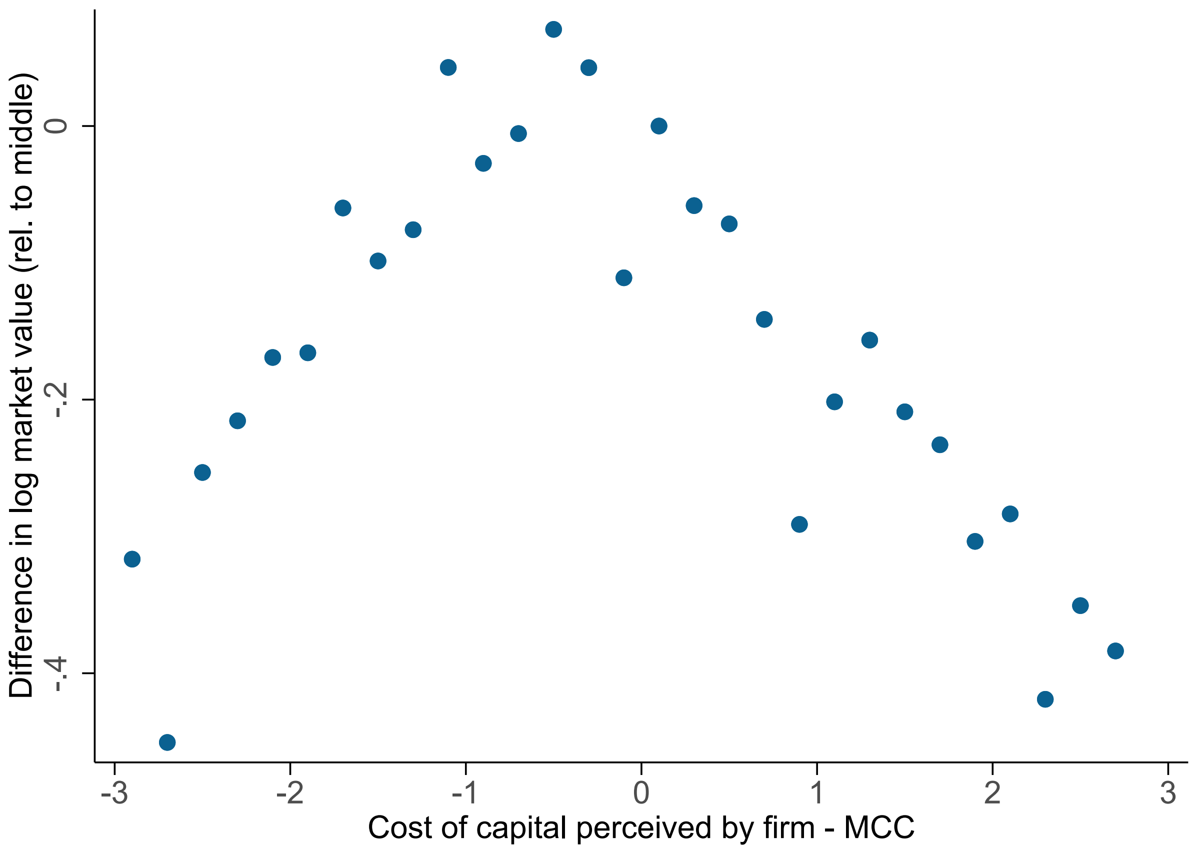

	(1) U.S.	(2) Global	(3) $\leq$ -0.3	(4) $\geq$ 0.3	(5) > -0.3, < 0.3
|Perc. CoC − MCC|	-0.083***	-0.097***			
	(0.019)	(0.018)			
Perc. CoC − MCC			0.096*	-0.068***	-0.139
			(0.050)	(0.025)	(0.196)
Observations	3,120	5,226	1,517	3,066	643
Within R^2	0.0200	0.0278	0.0112	0.0144	0.00155

The table reports regressions of firm log market value on the deviation between the cost of capital perceived by the firm and the MCC. Perc. CoC is the cost of capital reported by the firm on conference calls that influences the firm's investment decisions. MCC is the market-based cost of capital from this paper. The regressor in columns 1 to 2 is the absolute difference between Perc. CoC and MCC. The regressor in columns 3 to 5 is the simple subtraction. The sample includes only U.S. firms in column 1; all firms in column 2; only firms where the simple subtraction is at most 0.3 percentage points in column 3 (i.e., firms where the perc. CoC is low relative to the MCC); only firms where the simple subtraction is at least 0.3 in column 4; and firms where the simple subtraction is close to 0 in column 5. All specifications control for fixed effects for book equity deciles-by-country and quarter-by-year-by-country. Book equity deciles are based on the firm's decile in the distribution of book equity relative to other firms in the same quarter-year. Standard errors are clustered at the firm level. Statistical significance is denoted by *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Inverted U: firm market value and deviations from the MCC

point increase in the deviation is associated with a 9.6 log point value increase in the range of relatively large negative deviations (below -0.3 percentage points) analyzed in column 3. In contrast, a one percentage point increase in the deviation is associated with a 6.8 log point value decrease in the range of relatively large positive deviations (above 0.3 percentage points) analyzed in column 4. The magnitude is thus similar on both sides. The point estimate is statistically insignificant in the middle range close to zero in column 5, albeit statistically noisy. This pattern is consistent with the model that suggested an inverted U shape in Figure 4.

Taken together, the empirical results suggest that firms using an unbiased perceived cost of capital, like the MCC, are valued more positively in equity markets.

7 Conclusion

We develop a new method to estimate the firm-level cost of capital using market prices and machine-learning forecasts of future cash flows. The key challenge is that the textbook object firms should use in investment decisions—the expected return required by financial markets on debt and equity with comparable risk—is not directly observed. Existing empirical proxies, including the CAPM and analyst-based implied cost of capital measures, are biased. Our approach combines forecasts of future cash flows with a valuation equation that backs out the discount rate consistent with observed market values. The result is an unbiased market-based cost of capital (MCC).

Empirically, we show that the method performs well along two dimensions. First, the cash-flow forecasts exhibit meaningful out-of-sample predictive power, especially relative to widely used benchmarks. Second, the implied market-based cost of equity strongly predicts future long-run realized returns and does so more accurately than standard alternatives. These findings suggest that the MCC captures variation in expected returns well.

We then address a core question in finance and economics: how does the cost of capital used by firms matter for market value? Using hand-collected data on firms' perceived cost of capital, we find that market value is highest when firms' perceived cost of capital is closest to the MCC, consistent with the classical view that firms create value by aligning investment decisions with the opportunity cost of capital faced by investors.

References

Altman, E. I. and V. M. Kishore (1998): "Default and returns on high yield bonds: analysis through 1997," *NYU Salomon Center Report, January1995*.

Ball, R. and R. Watts (1972): "Some time series properties of accounting income," *The Journal of Finance*, 27, 663–681.

Bessembinder, H. (2018): "Do stocks outperform treasury bills?" *Journal of financial economics*, 129, 440–457.

Bharath, S. T. and T. Shumway (2008): "Forecasting default with the Merton distance to default model," *The Review of Financial Studies*, 21, 1339–1369.

Binsbergen, J. v., S. Hua, J. Peeters, and J. Wachter (2025): "Is the United States a lucky survivor? A hierarchical Bayesian approach," *The Journal of Finance*, 80, 2355–2388.

Campbell, J. Y., C. Polk, and T. Vuolteenaho (2010): "Growth or glamour? Fundamentals and systematic risk in stock returns," *The Review of Financial Studies*, 23, 305–344.

Campello, M., L. Chen, and L. Zhang (2008): "Expected returns, yield spreads, and asset pricing tests," *The Review of Financial Studies*, 21, 1297–1338.

Cho, T., L. Kremens, D. Lee, and C. Polk (2024): "Scale or yield? A present-value identity," *The Review of Financial Studies*, 37, 950–988.

Claus, J. and J. Thomas (2001): "Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets," *Journal of Finance*, 56, 1629–1666.

De Silva, T. and D. Thesmar (2024): "Noise in expectations: Evidence from analyst forecasts," *The Review of Financial Studies*, 37, 1494–1537.

Decaire, P. and M. Guenzel (2023): "What Drives Very Long-Run Cash Flow Growth Expectations?" *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

Dick-Nielsen, J., P. Feldhütter, L. H. Pedersen, and C. Stolborg (2023): "Corporate bond factors: Replication failures and a new framework," *Available at SSRN 4586652*.

Easton, P. D. (2004): "PE Ratios, PEG Ratios, and Estimating the Implied Expected Rate of Return on Equity Capital," *Accounting Review*, 79, 73–95.

Edwards, E. O. and P. W. Bell (1961): *The theory and measurement of business income*, Univ of California Press.

Fama, E. F. and K. R. French (1997): "Industry costs of equity," *Journal of financial economics*, 43, 153–193.

Foster, G., C. Olsen, and T. Shevlin (1984): "Earnings releases, anomalies, and the behavior of security returns," *Accounting Review*, 574–603.

Gebhardt, W. R., C. M. C. Lee, and B. Swaminathan (2001): "Toward an Implied Cost of Capital," *Journal of Accounting Research*, 39, 135–176.

Gormsen, N. J. and K. Huber (2025a): "Corporate Discount Rates," *American Economic Review*, 115, 2001–2049.

——— (2025b): "Firms' Perceived Cost of Capital," .

Graham, J. R. and C. R. Harvey (2001): "The Theory and Practice of Corporate Finance: Evidence from the Field," *Journal of Financial Economics*, 60, 187–243.

Hillenbrand, S. and O. McCarthy (2023): "The Optimal Stock Valuation Ratio," *Available at SSRN 4288780*.

Jensen, T. I., B. Kelly, and L. H. Pedersen (2023): "Is there a replication crisis in finance?" *The Journal of Finance*, 78, 2465–2518.

Ke, G., Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu (2017): "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, 30.

Koller, T., M. Goedhart, and D. Wessels (2020): *Valuation: Measuring and Managing the Value of Companies*, John Wiley & Sons, 7 ed.

Martineau, C. (2022): "Rest in peace post-earnings announcement drift," *Critical Finance Review*, 11, 613–646.

Ohlson, J. A. (1995): "Earnings, book values, and dividends in equity valuation," *Contemporary accounting research*, 11, 661–687.

Ohlson, J. A. and B. E. Juettner-Nauroth (2005): "Expected EPS and EPS growth as determinantsof value," *Review of accounting studies*, 10, 349–365.

Richardson, S., S. H. Teoh, and P. D. Wysocki (2004): "The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives," *Contempo-*

rary accounting research, 21, 885–924.

Rosenbaum, J. and J. Pearl (2021): *Investment banking: valuation, LBOs, M&A, and IPOs*, John Wiley & Sons.

Ross, S., R. Westerfield, J. Jaffe, B. Jordan, and K. Shue (2022): *Corporate Finance*, McGraw-Hill Education, 13th ed.

Shwartz-Ziv, R. and A. Armon (2022): "Tabular data: Deep learning is not all you need," *Information Fusion*, 81, 84–90.

Stein, J. C. (1996): "Rational Capital Budgeting in an Irrational World," *Journal of Business*, 69, 429–455.

Trahan, E. A. and L. J. Gitman (1995): "Bridging the Theory-Practice Gap in Corporate Finance: A Survey of Chief Financial Officers," *Quarterly Review of Economics and Finance*, 35, 73–87.

Van Binsbergen, J. H., X. Han, and A. Lopez-Lira (2023): "Man versus machine learning: The term structure of earnings expectations and conditional biases," *The Review of financial studies*, 36, 2361–2396.

Vuolteenaho, T. (2002): "What drives firm-level stock returns?" *The Journal of Finance*, 57, 233–264.

Zhang, Y., Y. Zhu, and J. T. Linnainmaa (2024): "Man versus Machine Learning Revisited," *Available at SSRN 4899584*.

Supplemental Appendix

Appendix A Figures and Tables

Method	MSE	Bias	Bias^2	Var($\hat{r}$)	Var(r)	2 Cov($\hat{r}, r$)	Cor	N
MCE	19.20	0.15	0.023	0.16	19.29	0.27	0.0764	2,105,618
VAR	19.33	0.00	0.000	0.23	19.29	0.19	0.0455	2,105,618
MLR	19.88	-0.18	0.031	1.32	19.29	0.76	0.0751	2,105,618
CAPM	19.94	0.28	0.076	0.63	19.29	0.06	0.0085	2,105,618

The table decomposes the mean squared error (MSE) of ten-year return predictions for the four methods with the broadest coverage. $\text{MSE} = \text{Bias}^2 + \text{Var}(\hat{r}) + \text{Var}(r) - 2, \text{Cov}(\hat{r}, r)$. MSE is the average squared prediction error. Bias is the average realized minus predicted return. $\text{Var}(\hat{r})$ and $\text{Var}(r)$ are the variances of predicted and realized returns, respectively. $\text{Cov}(\hat{r}, r)$ and $\text{Cor}(\hat{r}, r)$ are the covariance and correlation between predicted and realized returns. N is the number of observations. The sample is balanced across methods, so $\text{Var}(r)$ is identical across rows.

Table A1: Return predictability – MSE decomposition

Figure A1
Cash flow predictability

This figure shows binned scatter plots of realized earnings divided by assets at t on cash flow predictions at t, for forecast horizons $h = 1, 5, 10$ years.

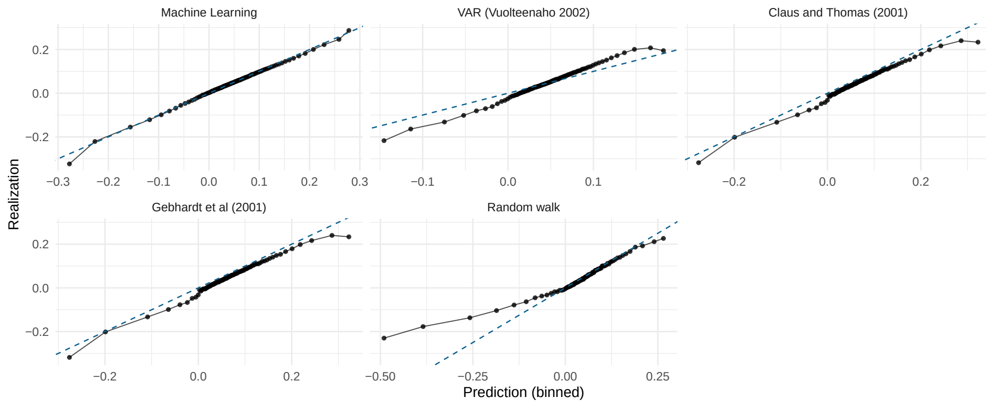

Panel A. $h = 1$

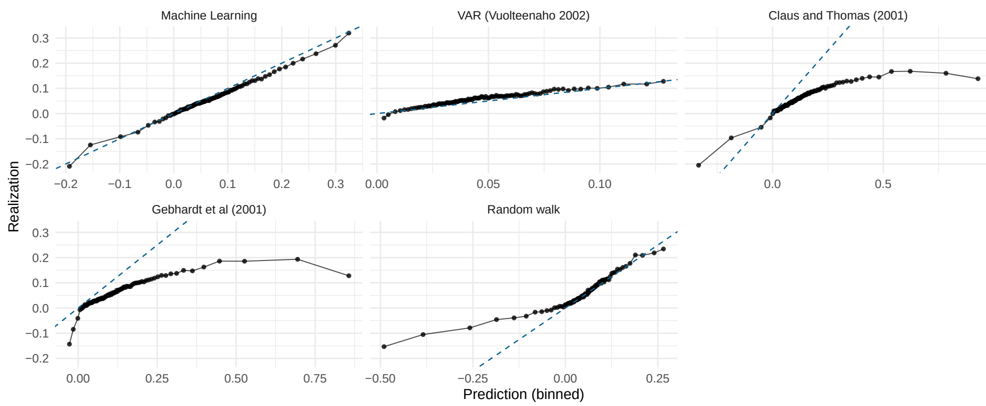

Panel B. $h = 5$

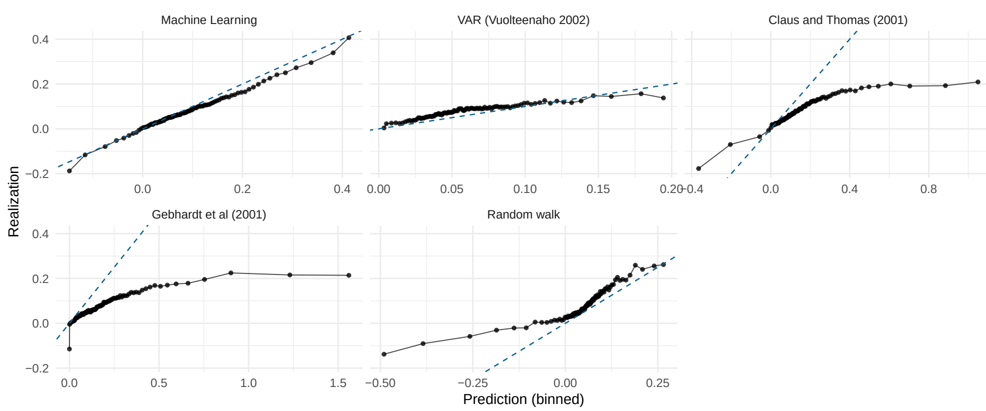

Panel C. $h = 10$

Figure A2
Cash flow forecasts (unbalanced)

This figure shows the out-of-sample R^2 of the cash flow prediction models. We consider three different cash flow variables, as indicated by the panel header, and ten forecast horizons, as indicated by the x-axis. The sample is unbalanced, meaning that we use all data available and that the samples differ across variables and horizons. Figure 1 shows the results from the balanced sample.

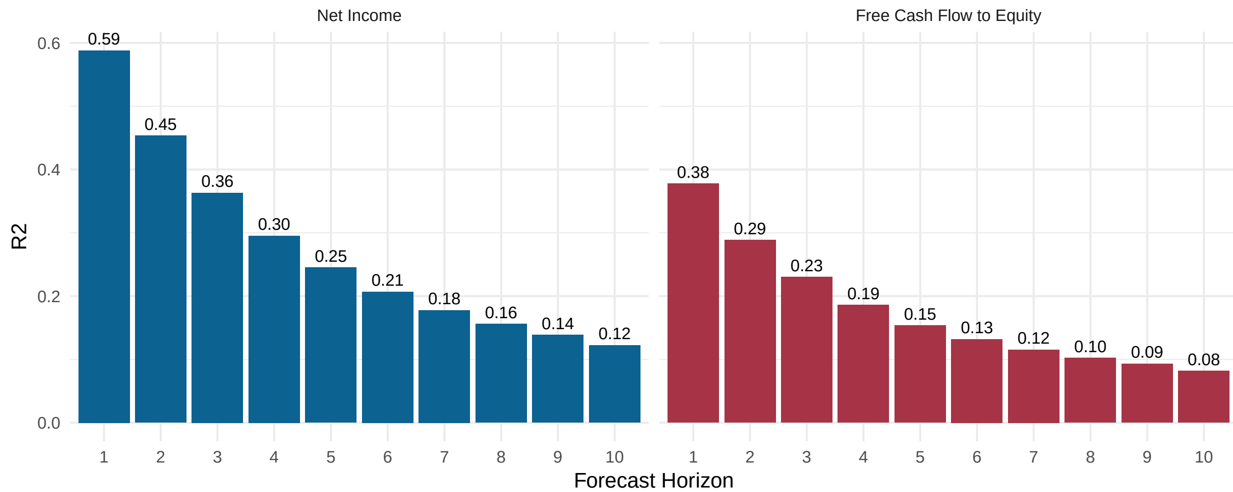

Figure A3
Cash flow forecasts over time (balanced)

This figure shows the out-of-sample R^2 of the cash flow prediction models every month. The panel header shows the cash flow variable, and the x-axis is the cash flow realization date. The sample is balanced. Figure A4 shows the results from the unbalanced sample.

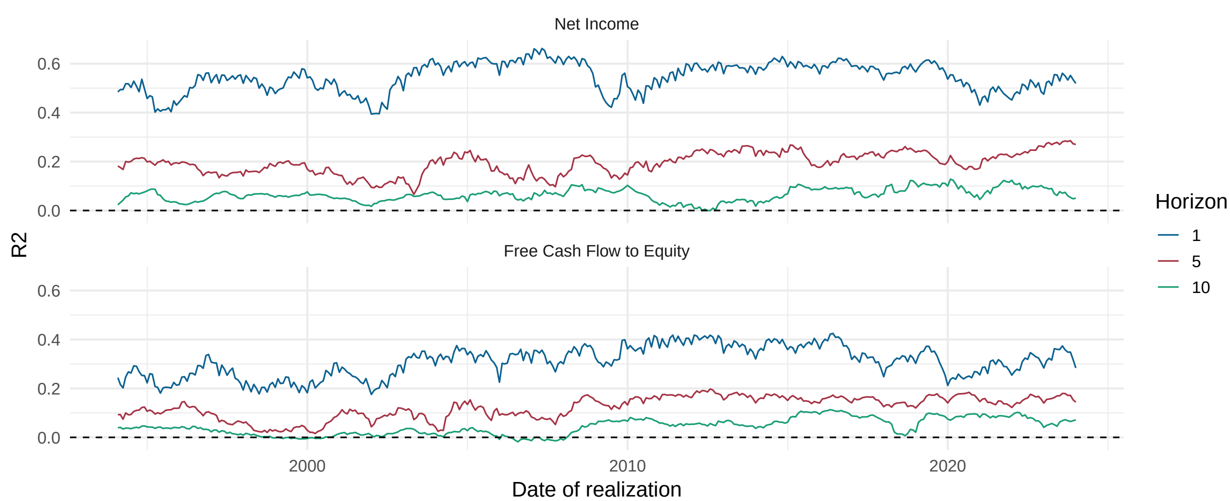

Figure A4
Cash flow forecasts over time (unbalanced)

This figure shows the out-of-sample R^2 of the cash flow prediction models every month. The panel header shows the cash flow variable, and the x-axis is the cash flow realization date. The sample is unbalanced. Figure A3 shows the results from the balanced sample.

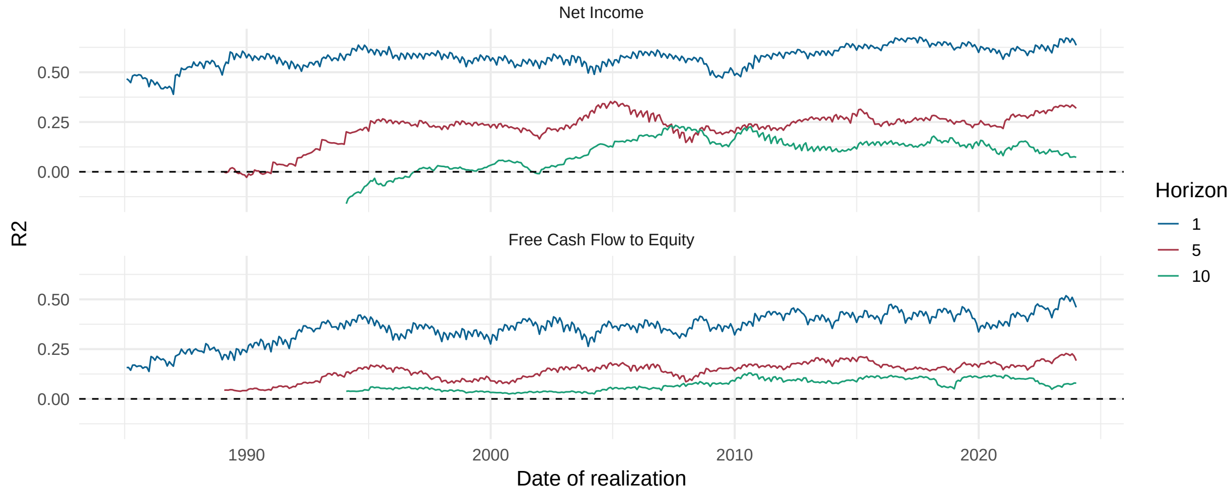

Figure A5
Long-run returns and the cost of equity

This figure shows the relationship between future 10-year realized returns and the current 10-year cost of equity. Specifically, each month we sort stocks into 100 portfolios by their cost of equity, and then compute the average 10-year cost of equity and the average realized buy-and-hold return over the next 10 years. We then average these two numbers for each portfolio over the full sample, 1983–2023. For ease of interpretation, the cost of equity and realized returns have been annualized. The solid red line shows the best linear fit, and the dotted blue line shows the 45-degree line. The sample is **U.S. stocks.**

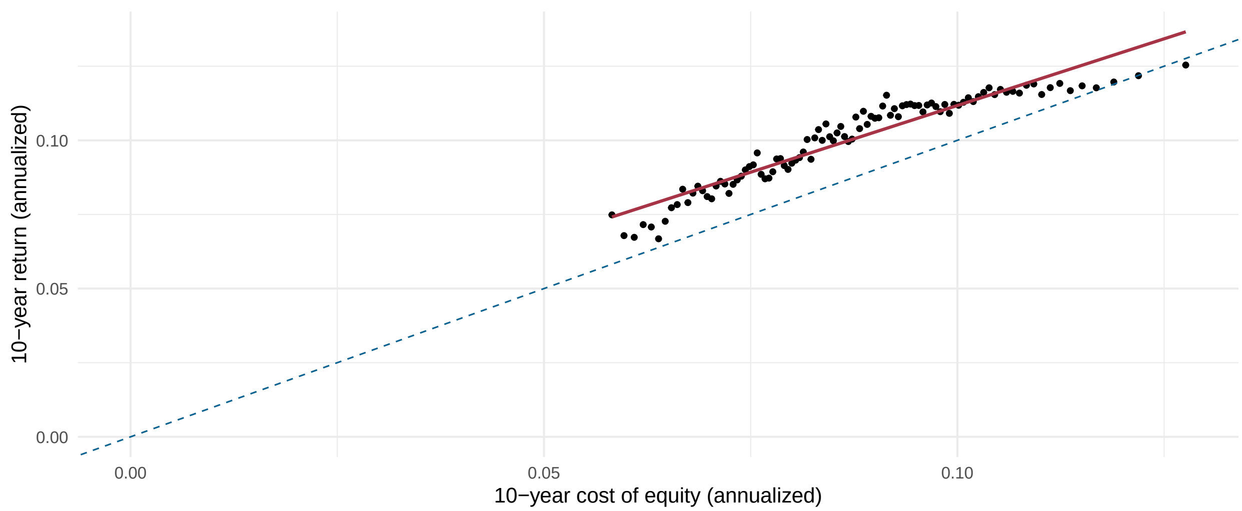

Figure A6

Long-run returns and the cost of equity

This figure shows the relationship between future 10-year realized returns and the current 10-year cost of equity. Specifically, each month we sort stocks into 100 portfolios by their cost of equity, and then compute the average 10-year cost of equity and the average realized buy-and-hold return over the next 10 years. We then average these two numbers for each portfolio over the full sample, 1983–2023. For ease of interpretation, the cost of equity and realized returns have been annualized. The solid red line shows the best linear fit, and the dotted blue line shows the 45-degree line. The sample consists of **non-U.S. stocks** listed in developed or emerging markets, as classified by MSCI.

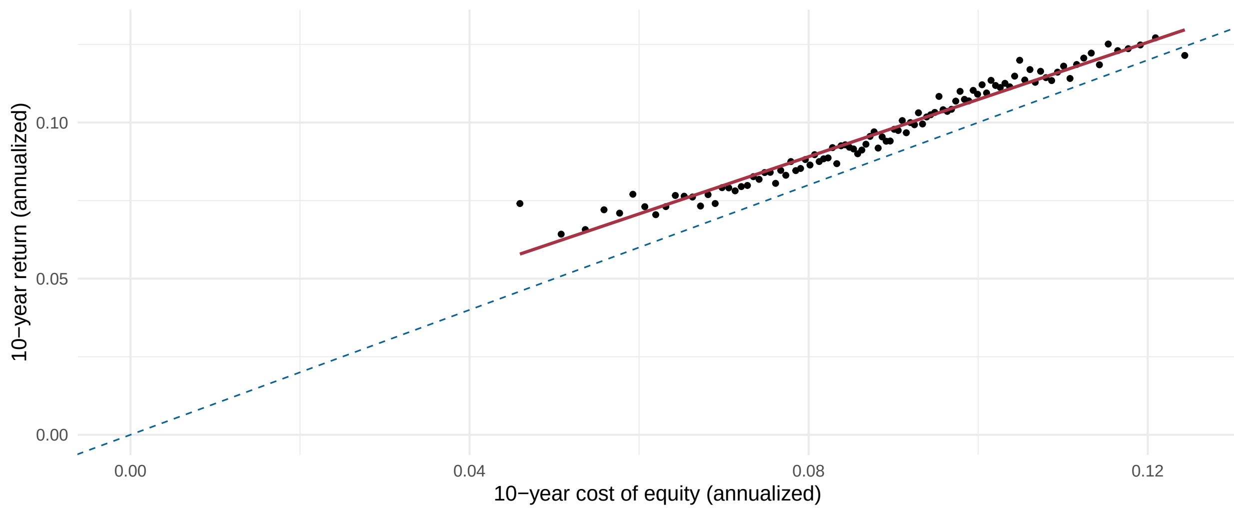

Figure A7
ML forecasts versus analysts

This figure compares the out-of-sample R^2 of one-year-ahead net income forecasts from the machine learning model, IBES analyst consensus forecasts, and a random walk (last observed net income). To ensure a comparable evaluation, the sample is restricted to analyst forecasts where the fiscal year end is exactly one year from the forecasting time, matching the constant-horizon property of the ML predictions. All forecasts and realizations are scaled by total assets and winsorized at the 1st and 99th percentiles. Firm-months where IBES and Compustat actual EPS disagree are excluded, and the sample requires non-missing predictions from all three methods. "Raw forecasts" computes R^2 as $1 - \text{MSE}/\text{Var}(y)$. "De-meaned forecasts" first subtracts the cross-sectional mean from both predictions and realizations before computing R^2, in an attempt to remove the well-known analyst level bias.

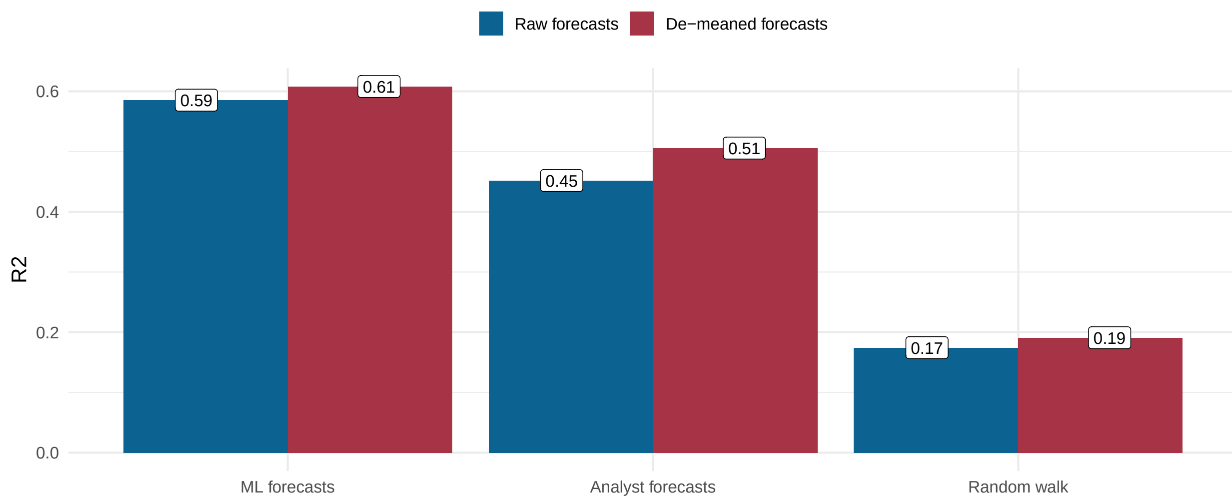

Appendix B Data Construction

We describe in detail how we construct the two cash flow measures.

- The valuation equations assume that the first cash flow occurs exactly one year from now, and that cash flows thereafter occur in one-year increments, but fiscal year ends do not necessarily allow for this one-year constant horizon. We handle this by defining a constant-horizon cash flow over 1 year that interpolates between the adjacent accounting reports. For example, consider a company with a fiscal year-end of December. To determine the dependent variable for this company in, say, 2020/06, we take an equal-weighted average of its annual earnings for the fiscal period of 2020/12 and 2021/12.

- Fiscal year ends differ across companies, which can create a mismatch when comparing annual earnings across companies. Our constant-horizon approach, described in the previous bullet point, handles this issue.

- The model draws on data from 93 countries, so to address currency mismatches, we convert all cash flows to U.S. dollars using the exchange rate at fiscal period-end. We get exchange rates from Compustat.

- The inflation rate varies widely across our sample. Therefore, we adjust the dependent variable for the realized inflation between the forecast date and the cash flow realization, using the `CPIAUCNS` series from FRED. This means our model effectively forecasts real cash flows.

- When forecasting long-run cash flow, we face the problem that not all firms available at the forecast date have realized cash flows over all forecasting horizons. For example, the company may have delisted or been taken private. One approach would be to fit the model only to observations from firms with non-missing realized cash flows, but this would introduce a survivorship bias.[A1] Instead, we keep all firms that were alive on the forecast date, and impute missing cash flows by carrying forward the last non-missing cash flow with the firm's stock return, if available, and zero otherwise. By using the stock return, we also incorporate delisting returns, which is a flexible way of handling both good and bad delisting reasons.[A2]

[A1] For example, if firms typically drop out of the sample for performance-related reasons, conditioning on firms that survive would introduce an upward bias in our cash flow forecasts.

[A2] For example, a firm that goes bankrupt will typically see its cash flows approach zero, whereas a firm that gets acquired will not.

- Equity issuances after the forecast date used, say, to acquire another company, could create a mechanical boost in future cash flow (the dependent variable) that would not necessarily accrue to the firm's current owners (because their ownership share would be diluted due to the equity issuance). Therefore, we divide the realized cash flows at all points by the number of shares outstanding, such that the dependent variable is the realized cash flow from owning one share at the forecast date.

- Working with a per-share variable, we need to adjust for stock splits that mechanically change the per share value without any change in the fundamentals (e.g., a two-for-one split would cut the cash flow per share in half, but since owners get two shares for each share they owned, they are fundamentally unaffected). We make this adjustment using the `CFACSHR` variable from CRSP for U.S. stocks, and the `AJEXDI` variable from Compustat for non-U.S. stocks.

- Finally, per share variables are somewhat arbitrary in the sense that the company can choose its number of shares outstanding, and this choice will lead to a difference in scale across companies. To handle this issue, we divide all realized cash flows by the most recent asset per share at the forecast date.