The Cost of Capital and Investment *

Niels Joachim Gormsen and Kilian Huber

November 2026

Abstract

A central decision facing every firm is how much to invest. Firms nowadays undertake investment projects if the expected returns exceed a discount rate set by the firm. Economists usually assume that this discount rate equals the opportunity cost of capital. However, this assumption is inspired by theory and the actual rules used by firms are not easily observed. A recent literature opens the black box of how firms make investment decisions using textual analysis and surveys. In practice, firms' discount rates deviate from the cost of capital assumption typically made by economists: firms' discount rates are sticky over time, excessively dispersed across firms, and substantially elevated. In the long run, the elevation and cross-firm differences lead to capital misallocation and reductions in firm value. In the shorter run, the stickiness mutes the sensitivity of investment with respect to interest rates. Firms' discount rates also generate a distinct form of monetary non-neutrality and help explain the behavior of investment over the business cycle.

*Gormsen is at Copenhagen Business School and The University of Chicago, njg.fi@cbs.dk. Huber is at The University of Chicago, kilianhuber@uchicago.edu.

Few decisions are as consequential for firms as how much to invest. Firms often make investment decisions using a simple rule: invest in projects that yield returns above a certain threshold, called the "discount rate."

Economists usually assume that firms use the "cost of capital" as their discount rate. This cost of capital captures the opportunity cost of investing, namely, the return that firms could earn by deploying their capital elsewhere. A threshold based on an opportunity cost is intuitive and in line with basic economic principles, so it has become the default assumption across all fields of economics. But the assumption is deduced from theory and economists have historically not been able to observe the internal investment rules used by firms in practice.

A recent line of research has opened the black box of firm decisions using surveys and textual analysis of corporate statements. What emerges are discount rates that deviate fundamentally from the cost of capital threshold assumed by economists. Real-world discount rates are sticky over time, excessively dispersed across firms, and substantially elevated. These deviations are partly driven by agency issues, behavioral biases, and heterogeneous frictions.

Firms' investment rules have widespread implications across many fields of economics. In the long run, they generate misallocation of capital and reductions in firms' market value. Over shorter horizons, they mute the effect of interest rates and financial markets on investment. In addition, they generate a distinct form of monetary non-neutrality and help explain the behavior of investment over the business cycle.

Investment Rules in Theory and in Practice Firms' investment rules can be summarized by a discount rate, δ_t at time t. The discount rate consists of two elements:

$$\delta_t = r_t + \kappa_t. \tag{1}$$

The baseline, r_t, is the firm's perceived cost of capital. The cost of capital is defined as an opportunity cost: the return firms forgo by using capital for an investment. For risk-free investments, the cost of capital is simply the interest rate and thus easily observed. But for risky investments, the cost of capital is the return investors could expect to earn on investments with similar risk in financial markets. While clearly defined, the cost of capital for any risky project is unobserved and difficult to estimate. Firms typically attempt to estimate the cost of capital in line with the standard definition, but their internal estimates

deviate from the values usually assumed by economists. In fact, firms' perceptions of their cost of capital are almost unrelated to the values used by economists. This deviation is the first major difference between the investment rules applied by firms and those assumed by economists.

The second major difference is that firms do not base their investment decisions directly on their perceived cost of capital. Instead, firms deliberately use a discount rate that deviates from their perceived cost of capital. We capture this deviation by the time-varying discount rate wedge κ_t. The wedge depends on frictions and biases that make discount rates sticky and elevated relative to the perceived cost of capital.

Firms began using investment rules based on discount rates and the cost of capital in the 1960s. In the 1980s and 1990s, the rules became widely adopted as consultancies and private equity investors pushed firms toward more disciplined capital budgeting. After adopting a rule based on the cost of capital, firms became more disciplined, their investment rates fell, and markets rewarded them with higher valuation ratios. By the 2000s, rules based on discount rates were used by the vast majority of large corporations.

Already in the 1970s, surveys of managers suggested that not all firms set their discount rates based only on the cost of capital. However, the survey data did not allow researchers to fully understand how discount rates were used and to test the benchmark assumptions, for instance, the relation between discount rates and the cost of capital or the impact of discount rates on real investment.

Recent data based on firm statements allow researchers to better understand firms' investment rules. The data suggest that the cost of capital and discount rates matter. An increase in a firm's discount rate is associated with a lower investment rate and higher realized returns on its investment projects. Moreover, the data establish new facts about how firms determine their perceived cost of capital and discount rates.

Perceived Cost of Capital Most firms regularly calculate a perceived cost of capital. In part, firms incorporate the standard opportunity cost definition used by economists, for instance, interest rates and certain risk premia. However, the perceptions also incorporate much variation that is not in line with the standard definition. As a result, most firms use a perceived cost of capital that deviates substantially from their opportunity cost: at most 20 percent of the variation in the perceived cost of capital can be justified by variation in the opportunity cost.

Why do firms' perceptions deviate so fundamentally from the benchmark? One reason is

that obtaining an unbiased cost of capital is simply too difficult. It is not possible to observe the opportunity cost of capital ex ante, because it depends on the unobserved distribution of future stock returns. Instead, firms are forced to estimate it and may make mistakes doing so, in line with a large literature on biased expectations.

Discount Rate Wedges Firms' discount rates are sticky relative to the perceived cost of capital. Over horizons shorter than two years, the average discount rate is essentially insensitive to changes in the perceived cost of capital. The extent of incorporation increases with the horizon. Only over horizons longer than ten years do discount rates move one-to-one with the perceived cost of capital. The slow incorporation reflects that discount rates change infrequently: less than 20 percent of discount rates observed within two years differ from the previous observation. When discount rates do change, they move almost one-to-one with the perceived cost of capital.

The discount rates directly reported by firms are often higher than the perceived cost of capital. Early surveys found very high average discount rates, twice as high as the average cost of capital. These very high levels led researchers to question the validity of the survey data, in particular, because such high discount rates are inconsistent with realized returns. However, the conference call data reveal that part of the high level is due to cost accounting. Many firms exclude overhead costs from their project return forecasts, so firms add an adjustment to the discount rate to capture the overhead costs. After correcting for the overhead cost adjustment, the true average discount rates are substantially lower and consistent with realized returns. The true average discount rates are still a few percentage points above the average cost of capital. This average wedge implies a persistent deviation from the standard assumption that the two objects are always equal.

The stickiness of discount rates and their high levels may not reflect mistakes by firms. Instead, they may reflect conscious responses by firms to organizational and behavioral frictions. For instance, manager statements suggest that high and stable discount rates prevent empire building, give divisions reliable targets, and alleviate complexity. Financial constraints and real option considerations may also lead to wedges.

Implications Firms' practical investment rules have large consequences for investment and the economy. We summarize the implications below, separately for the perceived cost of capital and discount rate wedges.

Implications of the perceived cost of capital. In the long run, the perceived cost of capital

strongly influences firms' discount rates and thereby their real investment decisions. One implication is that firms do not maximize their market value. Firms whose perceived cost of capital is too low effectively invest in projects with a lower return than what investors can earn in financial markets. This investment behavior represents a destruction of value in the eyes of investors, so market values are lower. Similarly, too high a perceived cost of capital implies that firms are forgoing value-creating investments, which also lowers market values.

The persistent deviations between the perceived cost of capital and the opportunity cost of capital also distort the allocation of capital across firms. Through the lens of a stylized model, the observed deviations can reduce aggregate total factor productivity (TFP) by 3 to 5 percent.

Moreover, firms' perceived cost of capital is, on average, slightly elevated relative to standard measures. This finding implies that firms tend to invest too little and, if they do invest, are biased toward short-term projects.

Implications of discount rate wedges. Stickiness in discount rates mutes the response to interest rates and financial markets in general. For example, firms noticed the decline in interest rates between 2002 and 2019 and reduced their perceived cost of capital, but they did not adjust their discount rates by nearly as much. By 2019, firms had let the wedge between discount rates and the perceived cost of capital grow by around 2.5 percentage points. This wedge accounts for much of the well-documented "missing investment" puzzle in the United States. An adjusted Q-model that incorporates realistic discount rates eliminates most of the apparent investment shortfall.

Moreover, because discount rates are sticky in nominal terms, increases in expected inflation directly lower real discount rates and raise real investment. This direct link from expected inflation to investment has implications for the transmission of demand shocks, which generate expected inflation and therefore crowd in rather than crowd out investment; for monetary policy, where changes in expected inflation affect real output even when all other prices are flexible; and for the comovement of investment and consumption over the business cycle.

1 Investment Rules in Theory and in Practice

1.1 The Investment Rule

Firms today invest in projects whose expected return exceeds a given threshold δ_t, known as a "discount rate," "hurdle," or "required return:"

$$\text{expected return of project} > \delta_t. \tag{2}$$

The discount rate is thus the minimum return that the firm is willing to accept on a new investment project. It is also the return on the last unit of capital invested by the firm, so it determines the firm's marginal product of capital. In addition to the simple rule (2), many firms require the net present value (NPV) of a project to be positive, using the discount rate δ_t to calculate the NPV.[1]

Such a rule has not always been standard practice. Business schools began teaching rigorous concepts such as the cost of capital and discount rates in the 1960s, building on Modigliani and Miller (1958, 1963). Sophisticated investors and corporate boards subsequently began to demand that investment decisions be justified relative to the cost of capital. Mentions of the cost of capital in annual reports rose steadily in the 1970s and 1980s, as consulting, private equity, and venture capital firms pushed more companies toward disciplined capital budgeting. During this period, firms reduced their investment rates after adopting a rule based on the cost of capital and former "empire builders" saw their market-to-book ratios rise, suggesting that equity investors valued the disciplined allocation of capital (Gormsen et al. 2026b). By the early 2000s, investment rules based on discount rates had become standard practice.

Today, almost all large firms report in surveys that they employ methods based on a discount rate to make investment decisions, either in the form of the threshold rule (2) or an NPV rule (e.g., Trahan and Gitman 1995, Poterba and Summers 1995, Graham 2022). In practice, the managers of divisions forecast the expected cash flows of projects (e.g., sales and costs), which determine the expected returns of projects. In contrast, the discount rate is typically set by the chief executives of the firm and is independent of the cash flows entered

[1]The rules based on (2) and the NPV always lead to the same investment decision for conventional projects, where the NPV of the project declines smoothly in the discount rate, which is the case in standard models (Brealey et al. 2011, pages 109–113). Investment rules based on the stochastic discount factor can also be represented using a discount rate in the case of standard projects. The NPV rule is conceptually preferable in the case of non-standard projects.

by the division managers (Jagannathan et al. 2016, Graham 2022). In line with corporate practice, firms make decisions based on discount rates in essentially all academic models and firms are taught to use discount rates in business schools.

The effect of the discount rate is large, as already recognized in the foundational theory by Jorgenson (1963). Take a textbook Q-model of firm investment. Under a standard calibration (e.g., Philippon 2009, Koby and Wolf 2020), an increase in the discount rate by 1 percentage point reduces the firm's investment rate (investment divided by the existing capital stock) by 1 to 3 percentage points. This magnitude is comparable to the drop in aggregate U.S. investment during the Great Financial Crisis of 2007–09. Micro data confirm the importance of discount rates for firm outcomes. Higher discount rates are associated with lower investment, both at the level of divisions (Krüger et al. 2015) and entire firms (Gormsen and Huber 2025). In addition, firms using higher long-run discount rates earn higher realized returns, both on individual projects and for the whole firm (Gormsen and Huber 2025, 2026, Barry et al. 2026).[2]

1.2 The Benchmark: Discount Rates Equal the Cost of Capital

Across fields of economics, the standard assumption is that firms' discount rates equal the opportunity cost of capital. The assumption is made in micro- and macroeconomics, industrial organization, labor, and finance; in models with irrational and rational agents alike; and in models used by central banks and treasuries.

The intuition is most easily understood in an imaginary risk-free world. In this world, profit-maximizing investors shift their capital across projects until every project earns the same marginal return, the risk-free rate. As a result, the risk-free rate represents the next-best alternative use of the investors' capital, in short, the opportunity "cost of capital." The

[2]Changes in interest rates are also associated with investment (e.g., Romer and Romer 2004, Christiano et al. 2005, Gilchrist and Zakrajšek 2007, Ottonello and Winberry 2020, Cloyne et al. 2023). Interest rates affect both firms' discount rates and liquidity constraints, so it is difficult to isolate the role of discount rates in the overall effect of interest rates. Corporate tax reforms have also affected investment, in part because the reforms allowed accelerated expensing and relaxed liquidity constraints. These channels are distinct from the effect of discount rates and the traditional cost of capital channel in the spirit of Modigliani and Miller (1958). In this review, we use "cost of capital" to mean investors' expected returns accounting for the debt tax shield (known as WACC), following Modigliani and Miller (1963). In comparison, the tax literature typically analyzes the "tax-adjusted user cost of capital" (Chodorow-Reich 2025), which additionally incorporates the effects of accelerated expensing and bonus depreciation in a separate "tax term" (e.g., Cummins et al. 1994, House and Shapiro 2008, Zwick and Mahon 2017, Chodorow-Reich et al. 2026). Changes in the debt tax shield have little short-run effect on discount rates and investment (Goodman et al. 2025), consistent with the view that discount rates are sticky with respect to the cost of capital.

firm therefore acts optimally in the eyes of its investors if it invests only in projects whose return exceeds the risk-free rate.

The real world is more complicated. Projects vary in their risk and other characteristics. If the firm encounters a project with an expected return p, investors consider the next-best alternative: how much they can earn on another project outside the firm that has similar characteristics. The expected return on that similar project is the proper definition of the opportunity cost of capital in a risky world, which we denote r^*. Investors only want the firm to invest in the project if p exceeds r^*, so r^* is the firm's optimal discount rate in the eyes of its investors. If the firm follows this rule, investors want to provide capital to the firm and the firm's market value is maximized. In terms of equation (1), the benchmark is that firms perceive the cost of capital correctly, $r = r^*$, and add no wedge, $\kappa = 0$.

The benchmark implies the traditional "cost of capital channel:" when expected returns on similar projects in financial markets change, firms immediately adjust their discount rates and investment.

1.3 Why Firms Deviate from the Benchmark

Equation (1) separates two ways in which firms may deviate from the benchmark. The first is that firms may misperceive the opportunity cost of capital, so the perceived cost of capital is not equal to the true opportunity cost ($r \neq r^*$). The second is that firms may consciously choose a discount rate that differs from their perceived cost of capital, $\kappa \neq 0$.

Perceiving the cost of capital. The first deviation arises from a fundamental measurement problem. The opportunity cost of capital is the expected return of another project that has similar characteristics to the project under consideration by the firm. Expected returns are unknown because they depend on uncertain future outcomes. Firms therefore need to estimate a "perceived cost of capital." However, there is no agreed-upon, off-the-shelf estimation method, even though researchers have tackled this challenge for decades (Fama and French 1997).

Firms are commonly taught to calculate the "WACC" (weighted average cost of capital) and use it as the perceived cost of capital.[3] But the WACC, in turn, depends on the

[3]The logic behind the WACC method is as follows. The average investor who has provided capital to the firm will earn the average return on all the projects the firm has undertaken. The WACC measures this average return (averaging over all debt and equity holders). Assuming that investments with identical characteristics earn identical returns (i.e., assuming the law of one price), the WACC should also capture how much an investor could expect to earn on an alternative investment that is unrelated to the firm but has

expected return on the firm's debt and equity in financial markets (not just interest rates). It is not clear how the components of this expected return should be estimated, leaving much room for idiosyncratic choices (Frank and Shen 2016, Jagannathan et al. 2017). Firms could employ simple models, such as the Capital Asset Pricing Model (CAPM), which may measure fundamental risk but do not actually capture investors' expected returns (Stein 1996). More recent return models developed by finance researchers often still remain far from unbiased. A further challenge is that the WACC logic provides the cost of capital of the firm's average project, but not of individual projects or divisions. Given these challenges, it is not clear whether firms' perceived cost of capital is correlated with true expected returns, and whether firms could implement the benchmark even if they wanted to.

Setting the discount rate. Assume that the firm has estimated a perceived cost of capital. According to the benchmark, the firm should now use this cost of capital as its discount rate in order to maximize its value. However, firms face challenges not captured by the benchmark that may incentivize them to set discount rates differently.

One class of potential challenges is organizational. High and stable discount rates may stop senior executives from building corporate "empires" (Jensen 1986) and motivate the individual divisions to develop high-return projects and agreements (Barry et al. 2026). Relatedly, senior executives may use sticky discount rates as a management tool that gives divisions reliable targets for potential projects, without divisions fearing that future discount rate changes render their development efforts redundant. Firms may additionally set a sticky discount rate to avoid the costs of renegotiating a new rate with divisions. They may also use one common discount rate for all divisions to avoid rent-seeking struggles between divisions (Scharfstein and Stein 2000, Rajan et al. 2000). Moreover, firms may be financially constrained, implying they do not have sufficient liquid funds to accept all the projects whose expected returns lie above the cost of capital.

Another class of challenges is behavioral. Firms may find it too complex to recalculate their discount rate over time or adjust it for individual divisions (Mankiw and Reis 2002, Gabaix 2025, Puri 2025), instead opting for the simplicity of one stable, common discount rate.

Taken together, the two deviations imply that it is not at all clear whether discount rates equal the opportunity cost of capital and how they move over time. The remainder of this

identical characteristics to the firm's average project. Following that logic, the WACC captures the average investor's expected return on the alternative investment, which is by definition the firm's cost of capital.

review documents what the data say about each.

2 Measuring Firms' Discount Rates

A key measurement challenge is that discount rates are set internally by firms and not usually reported in annual reports, analyst statements, or other easily accessible data sources. Researchers initially addressed the challenge using surveys, starting in the 1970s. The surveys provide insights about the methods used by firms in the real world.

More recently, researchers have relied on statements by firm executives on conference calls—quarterly conversations between executives and analysts—to measure discount rates. Many firms report discount rates at different points in time on conference calls, allowing researchers to construct a panel and conduct dynamic analyses of changes over time and shock transmission.

One interesting feature of the data is that the discount rates reported on surveys far exceed the realized returns that firms earn on their projects. In principle, discount rates should be the expected return on the lowest-return project. So, discount rates above average realized returns puzzled many observers.

A resolution to this puzzle is that firms report different types of discount rates on surveys. Some firms purposely report higher rates to account for the fact that overhead costs are not fully incorporated in their internal calculation of project returns. Hence, the level of reported discount rates needs to be corrected for cost accounting before being compared to realized returns and the cost of capital. Once corrected, the discount rates reported on conference calls are fully consistent with realized returns, but still greater than the cost of capital.

2.1 Survey-Based Measurement and Different Discount Rate Types

Surveys provide insights on the methods used by firms. In the 1970s, around half of the surveyed firms—in some way—considered their cost of capital when setting their discount rates (e.g., Brigham 1975, Gitman and Forrester Jr. 1977, Schall et al. 1978). A 1991 survey reveals that the planning horizons of firms respond to the cost of capital (Poterba and Summers 1995). The Duke CFO Surveys show that some firms adjust their discount rates in response to changes in the cost of capital (Graham and Harvey 2001), but also that many firms prefer simple and stable rules in their decision-making (Graham 2022).

Several surveys record numerical values of discount rates, although the type of discount

rate varies. For example, a 1985 survey asks about discount rates for risk-free tax allowances (Summers 1986), which leads some firms to report lower discount rates than those used for their typical risky investments. As a result, the discount rates reported on that survey are not directly comparable to firm-level returns and the firm-level cost of capital.

Similarly, a 1991 survey asks 228 firms about the rate used to discount earnings "in the firm's largest division" (Poterba and Summers 1995). Such division-level discount rates are used by firms to evaluate division-level returns, which encompass the inflows and outflows directly generated by a project, but not the overhead costs due to the project. Overhead costs include, for example, the planning, administrative, and marketing expenses that accrue to the headquarters.[4] The division-level return of a project is greater than the overall firm-level return because of the exclusion of overhead costs. Hence, the discount rate applied to evaluate division-level returns is purposely inflated relative to the discount rate that would be applied to evaluate overall firm-level returns.

In the 1991 survey, firms report an average division-level discount rate of 17 percent (in nominal terms), which is above standard estimates of the cost of capital in 1991. Because the discount rates are for division-level earnings and thus purposely inflated, it is not immediately clear how far the firm-level discount rate diverges from the cost of capital.

A 2003 survey asks 86 firms about "the hurdle your company has used for a typical project" (Jagannathan et al. 2016), finding a reported average of 15 percent. The Duke CFO surveys in 2007, 2011, 2012, 2017, and 2019 ask about "your firm's hurdle rate (the rate of return that an investment must beat in order to be adopted)" (Graham 2022), finding reported averages around 14 to 15 percent. These surveys do not specify the type of discount rate that is to be reported, so some discount rates in these samples may be inflated to account for overhead costs.

2.2 Measurement Using Conference Calls

Listed firms organize quarterly conference calls with financial analysts. On these calls, some firms report their discount rates and perceived cost of capital. Gormsen and Huber (2025, 2026) manually read through thousands of conference call transcripts and construct a panel dataset. The data start in 2002 and are regularly updated.[5] Observations in the data are based only on non-hypothetical, unambiguous statements by firm managers.

[4] Total overhead (measured as SG&A) is 31 percent of invested capital of Compustat firms (Gormsen and Huber 2025).

[5] See costofcapital.org.

Many firms report one firm-level discount rate that is applied to all of the firm's projects or is representative of the firm's typical project. A few firms also report additional, division-specific discount rates, which are recorded and analyzed separately. For the cost of capital, firms usually report their firm-level perceived cost of capital, which captures the funding cost of a typical project representative of the overall firm.[6]

Firms with at least one observed discount rate or perceived cost of capital account for roughly 50 percent of the total assets of listed firms in Compustat in our sample period. Large firms are more likely to hold regular conference calls, so large firms are more likely to be in the sample, relative to other firms in Compustat. In fact, of the 100 largest U.S. firms, 40 mention a discount rate at least once. Apart from size, other firm characteristics, such as firm growth, leverage, market valuation, and profitability, are not associated with the propensity to report a discount rate. On the calls, firms can discuss their investment strategies using various metrics, including details about individual projects, expected future returns, overall investment rates, and discount rates. Many firms do not report discount rates or the perceived cost of capital because analysts are content with the other types of information.

Firms' statements on conference calls are subject to scrutiny from analysts who often compare previous statements to realized decisions. Moreover, the calls appear as evidence in securities lawsuits (Rogers et al. 2011). As a result, firms tend to report discount rates accurately, as confirmed by the association between discount rates and real outcomes (see Section 1.1).

3 Facts on the Perceived Cost of Capital

Economists' benchmark assumption is that firms know their opportunity cost of capital and use it as their perceived cost of capital. This section summarizes how firms set their perceived cost of capital in practice and how far it deviates from the opportunity cost.

[6]The firm-level cost of capital is typical in the sense that it is for a project with the same characteristics as the firm's average project. The firm-level discount rate and perceived cost of capital are comparable to each other, as they both relate to the firm's typical project.

3.1 The Perceived Cost of Capital and the Opportunity Cost of Capital

Most firms regularly calculate a perceived cost of capital. To implement the benchmark, a firm would need to observe investors' expected returns on its debt and equity and adopt these expected returns as its perceived cost of capital. As discussed in Section 1.3, expected returns are not observed, so firms need to estimate them, and there is no agreed-upon method for doing so.

Surveys show that firms mix and match methods. Around 80 percent of large firms use the CAPM, 70 percent use multi-factor models, and 40 percent use historical returns in the Duke CFO Survey (Graham 2022), with other surveys finding similar results (Brounen et al. 2004, Jacobs and Shivdasani 2012, Mukhlynina and Nyborg 2016, Jagannathan et al. 2016).

In part, firms' perceptions incorporate the standard drivers of the opportunity cost of capital. In the conference call data, the average perceived cost of capital moves almost one-to-one with the earnings yield over time, and the average perceived cost of debt moves with long-term Treasury rates. In the cross section, the perceived cost of capital increases in market beta and decreases in leverage and size (Gormsen and Huber 2026).

Importantly, the perceptions incorporate much variation that is not in line with the opportunity cost. Gormsen and Huber (2026) find that at most 20 percent of the variation in the perceived cost of capital reflects variation in expected returns. The remaining 80 percent is "excess dispersion." This analysis can be conducted without knowing the "true" expected returns ex ante. By comparing firms' perceptions of their cost of capital to the future realized returns on firms' debt and equity, we can assess the extent to which the perceived cost of capital lines up with ex ante expected returns. The large dispersion reflects that firms' perceptions of their cost of capital line up poorly with future realized returns and hence ex ante expected returns.

The large excess dispersion suggests that firms do not calculate their cost of capital in line with the benchmark assumptions made by economists. By extension, the finding implies that firms do not invest in line with the assumptions of standard models.

3.2 Other Drivers of the Perceived Cost of Capital

While firms' perceived cost of capital is largely unrelated to the opportunity cost of capital, it is more closely related to the cost of capital perceived by financial analysts. This finding suggests that firms may use the same cognitive process to form their expectations as analysts.

Similarly to the perceived cost of capital, analysts' forecasts do not align well with expected returns in the cross section of firms (Bouchaud et al. 2019, Engelberg et al. 2020, Jensen 2026, Bordalo et al. 2025), but incorporate time variation in expected returns to a larger extent (Thesmar and Verner 2025, Bastianello 2025). Other aspects of firms' decisions, for instance, cash flow forecasts, are also associated with information provided by analysts (e.g., Gennaioli et al. 2016, Dessaint et al. 2026a). Why analysts and firms are aligned remains open: firms may influence analysts or vice versa, and the relation may vary across firms and time. Recent work sheds light on the subjective models and beliefs used by analysts (Nagel and Xu 2022, Bastianello et al. 2025, Décaire and Graham 2026, Ke 2026).

Another salient driver of the perceived cost of capital is firms' climate characteristics. Since around 2016, firms with high carbon emissions have perceived a higher cost of capital than firms with low emissions (Gormsen et al. 2026c). This climate wedge has grown with the climate concerns of financial market participants. Because the wedge passes through to discount rates and investment, it operates as an implicit carbon tax on high-emission firms (Chitarro et al. 2025, Pedersen 2026).

3.3 The Market-Based Cost of Capital

One possible explanation for the deviations is that the opportunity cost of capital cannot be directly observed. Expected returns are unobserved, so there is no readily available estimate of the cost of capital, and long-run expected returns in particular are hard to estimate (Keloharju et al. 2021, van Binsbergen et al. 2023, Cho and Polk 2024). Common methods for calculating the expected returns underlying the cost of capital include the CAPM, the implied cost of capital (ICC) based on analyst forecasts, and factor models aimed at predicting short-run returns. These methods deliver biased estimates of long-run expected returns, and hence of the cost of capital. If no unbiased estimate exists, firms cannot implement the benchmark even if they wanted to.

Recent work has begun to approach this challenge using new statistical tools (Nagel 2021, Kelly and Xiu 2023, Bianchi et al. 2026). Hommel et al. (2025) find that traditional methods using the CAPM and other factor models perform poorly, whereas a lasso procedure performs better. Their procedure estimates a cost of capital without requiring a current market price as an input, which is useful for projects and firms where a price is not observed.

Gormsen et al. (2026a) provide unbiased out-of-sample estimates of long-run expected stock returns that can be used to calculate the cost of capital. Rather than predicting returns

directly, they forecast firms' cash flows with machine learning and back out expected returns from observed market prices and the present value identity. The resulting "market-based cost of capital" (MCC) is an unbiased predictor of ten-year realized returns out of sample, with a slope coefficient close to one, and it has greater predictive power than existing alternatives. The MCC may serve as an unbiased measure of firms' opportunity cost of capital and thereby influence market values (see Section 5.1).

As of now, firms' perceptions are far from the MCC. At most 20 percent of the variation in the perceived cost of capital can be justified by variation in the MCC, the same share as for expected returns measured using realized returns (Section 3.1). The MCC is unrelated to market beta and volatility and to the return expectations of financial analysts, in contrast to firms' perceived cost of capital (Sections 3.1 and 3.2).

4 Facts on Firms' Discount Rates

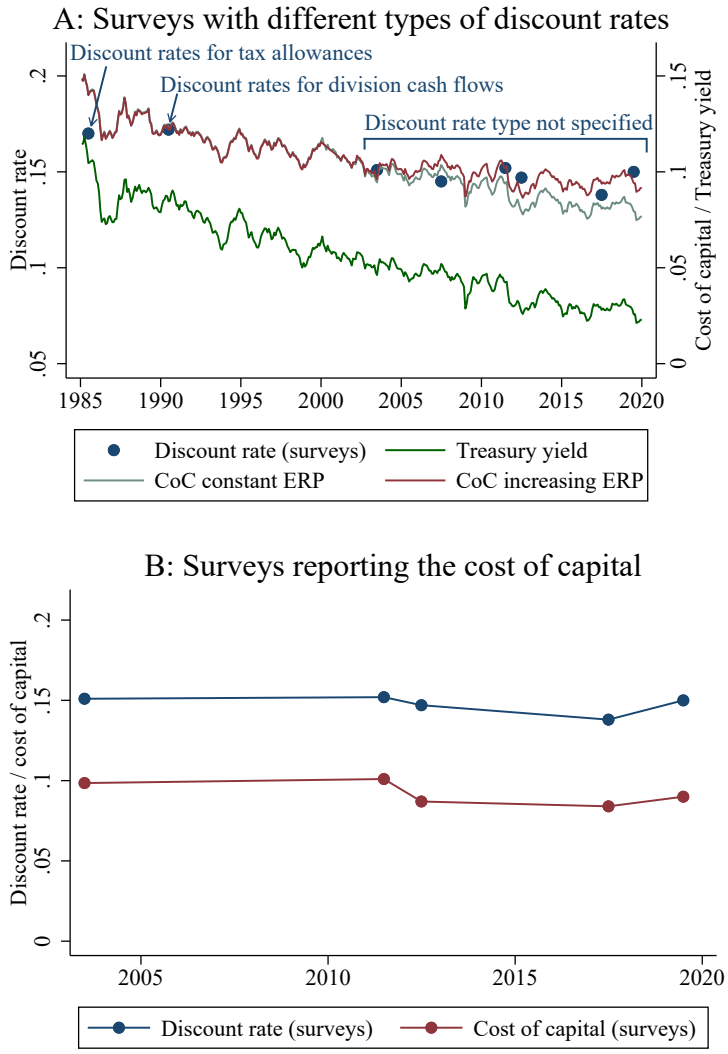
The key fact about the relation between discount rates and the perceived cost of capital is horizon dependency. Over short time horizons, the average firm rarely moves its discount rate when its perceived cost of capital changes. Over long horizons, the average discount rate strongly incorporates changes in the perceived cost of capital.

In this section, we describe the history of thought about discount rate dynamics; the benefits of within-firm panel data; and the insights about horizon dependency based on conference call panel data. We also summarize evidence about whether firms use division-specific discount rates or one common discount rate.

4.1 History of Thought on Dynamics

The default assumption in all fields of economics is to take the benchmark view for granted: any shock to the cost of capital immediately affects firm discount rates and thus real outcomes (e.g., Jorgenson 1963, Tobin 1969, Cochrane 1991). Under this view, firms update their discount rates whenever the cost of capital changes. The key question is whether the benchmark view is a reasonable and useful approximation for firm behavior (Poterba and Summers 1995, Rognlie 2019).

Surveys have asked firms how frequently they adjust their discount rates. Taken together, the survey evidence is intriguing, but ambiguous. Around 40 percent to 50 percent of firms do not change their discount rate every year, according to Brigham (1975) and Meier and



In Panel A, the navy markers report the average discount rate from different surveys. The 1985 average is for 95 firms asked about their discount rate for risk-free depreciation tax allowances, which should be lower than discount rates for investments (Summers 1986). The 1991 average is for 228 firms asked about their division-level discount rate (Poterba and Summers 1995), which does not account for overhead costs. The 2003 to 2019 surveys do not specify which type of discount rate the firm should report, so some discount rates in these samples may account for overhead costs and others may not (Jagannathan et al. 2016, Graham 2022). We plot two measures of the average cost of capital in the U.S., defined as $\omega_t \times (1-\tau) \times i_t^{\text{debt}} + (1-\omega_t) \times i_t^{\text{equity}}$, where ω_t is the aggregate market leverage, τ is the corporate tax rate, i_t^{debt} is the Treasury yield plus a 2 percentage point (ppt.) debt risk premium, and i_t^{equity} is the market cost of equity. "CoC constant ERP" assumes that the equity risk premium was constant at 6 ppt. throughout the sample. "CoC increasing ERP" assumes that the equity risk premium was constant from 1985 to 2000 and increased by 2 ppt. from 2000 to 2019, which is consistent with the recent literature (Campbell and Thompson 2008). "Treasury yield" is the long-run interest rate on Treasury bonds. Panel B includes only surveys that report both an average discount rate and an average cost of capital for firms in the sample. The figure is from Fukui et al. (2026).

Figure 1: Average Discount Rates from Different Surveys

Tarhan (2007). However, 89 percent of firms use the cost of capital as their discount rate, according to Bruner et al. (1998), and higher borrowing costs affect the investment plans of 63 percent of firms, according to Sharpe and Suarez (2021). In Graham (2022), 41 percent of firms report not changing their discount rate "by 1 percent or more over the past 10 years," suggesting that many firms do not make large changes to their discount rate in one instance. The result could be consistent with the benchmark view, since the cost of capital does not tend to experience large changes and firms may thus update their discount rates by small incremental amounts at regular intervals, but also with infrequent adjustment.

It is challenging to draw lessons about discount rate dynamics from the average discount rates reported on different surveys over time. In general, one cannot infer the flexibility of a firm-level object from a time series of averages. For instance, the macro-finance literature does not use average inflation over time to infer the degree of price stickiness (Bils and Klenow 2004, Nakamura and Steinsson 2008). Similarly, average discount rates over time are not informative about the degree of discount rate stickiness. A further challenge is that different surveys measure different types of discount rates, as detailed in Section 2.1.[7] Finally, each survey is based on a different sample of firms, so varying sample composition could account for any pattern.

In line with the conceptual challenges, averages over time from different surveys offer no clear conclusion, as shown in Figure 1.A. The average U.S. cost of capital changed by more than the average surveyed discount rate over some horizons and not others. A few surveys also report a specific cost of capital for the firms in the sample, plotted in Figure 1.B. Here, it looks like the discount rate and cost of capital comove, but it is not clear whether this is driven by varying sample composition or a true economic relation. The average surveyed discount rate fell by less than interest rates (here, the Treasury yield, see also Sharpe and Suarez 2021, Graham 2022) in recent decades, but the relevant benchmark is the cost of capital, not an interest rate.

Taken together, traditional models and surveys raise intriguing questions, but leave open whether the benchmark view is a useful approximation and whether it holds for the large firms dominating aggregate investment.

[7]For instance, firms reported the rates used to discount risk-free tax allowances in 1985; the rates used to discount division-level earnings in 1991; and potentially a mix of unspecified types in the later years.

4.2 Panel Data Reveal Dynamic Effects

Testing firm-level stickiness requires repeated observations for the same firm in the form of panel data. Such panel data overcome the conceptual challenges associated with using repeated cross-sections, different discount rate types on surveys, and varying firm samples.

Panel data based on conference calls show that the incorporation of changes in the perceived cost of capital into discount rates is horizon-dependent. Figure 2 plots the relation over different time horizons. The leftmost coefficient implies that a 1 percentage point increase in a firm's perceived cost of capital over horizons shorter than 2 years is, on average, not associated with a significant change in the discount rate over the same period. The coefficient is around 0.25 percentage points for horizons of 3–4 years (and statistically different from 0, 1, and the 1–2 year coefficient). These estimates imply a strong degree of stickiness in discount rates in the short and medium run.

Many firms leave their discount rates entirely unchanged for many quarters, as shown in Figure 3. In contrast, firms update their perceived cost of capital much more frequently. Firms therefore consciously choose sticky discount rates, even though they are aware that their cost of capital has changed.

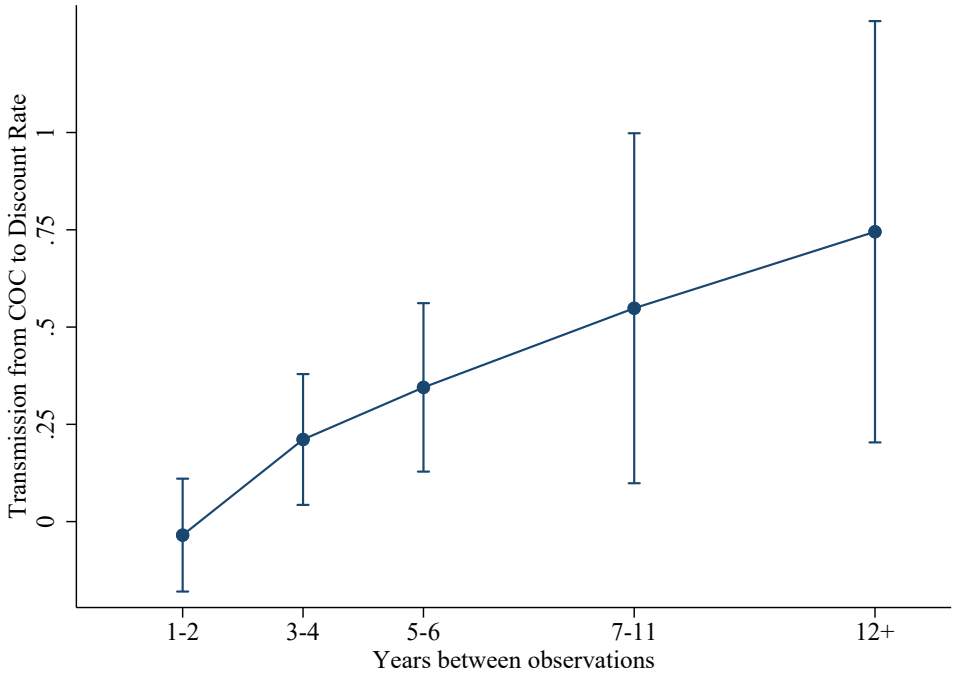
Over longer horizons, the relation between discount rates and the cost of capital grows closer to 1. In the long run, discount rates depend strongly on the cost of capital. When different firms are compared with one another at the same point in time, firms whose cost of capital is 1 percentage point higher have, on average, a 1 percentage point higher discount rate (Gormsen and Huber 2025).

Taken together, discount rates move with the perceived cost of capital over long horizons, but not in the short run. This pattern has strong implications for the real economy and generates new conceptual mechanisms for how investment responds to financial and policy shocks.

4.3 Elevated Discount Rates

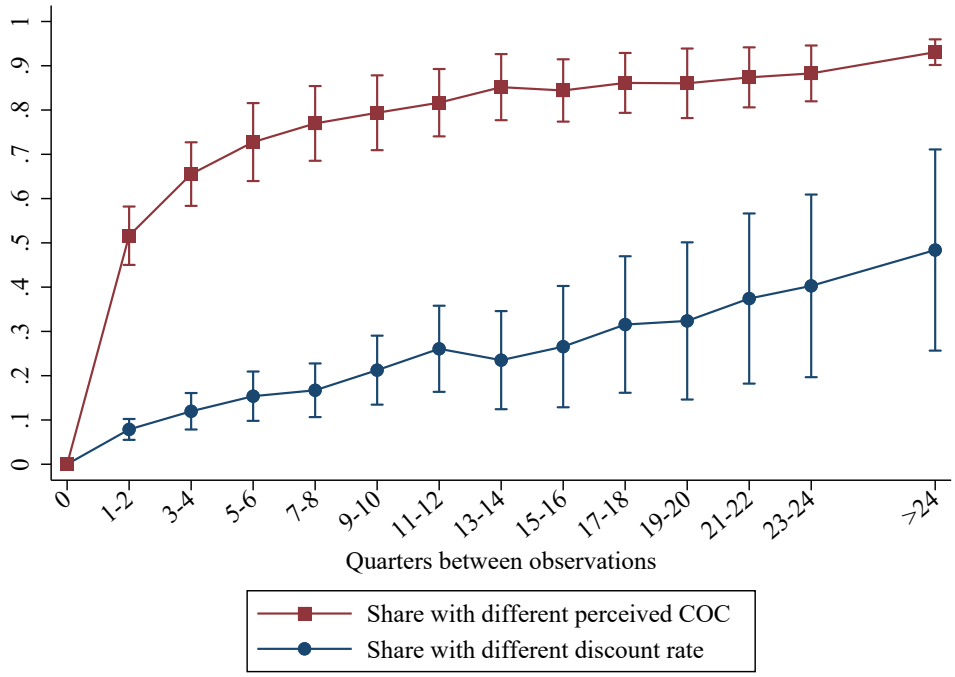
The average discount rate reported in the surveys (Section 2.1) is usually around 15 percent, which is well above standard measures of the cost of capital and above the return to capital earned by firms. Expected stock returns have historically been around 10 percent and the average return on invested capital (ROIC) for U.S. listed firms is around 13 percent in recent years (Mauboussin and Callahan 2022, Ayyagari et al. 2024, Gormsen and Huber 2025).

The relatively high level of discount rates was a long-standing puzzle in the literature.



The figure is based on conference call data. The coefficients are estimated by regressing the average yearly within-firm change in the discount rate on the average yearly within-firm change in the perceived cost of capital over the same period. The change in the cost of capital is interacted with indicators for the number of years between the observations. The figure plots the estimated coefficients on the interactions. The regression sample contains all within-firm changes in discount rates observed in the conference call panel data (i.e., all differences between different discount rate observations of the same firm). The vertical lines measure 90 percent confidence intervals. Standard errors are clustered by firm and quarter of the later observation. The figure is from Gormsen and Huber (2025).

Figure 2: Slow Incorporation of the Cost of Capital Into Discount Rates



The figure is based on conference call data. The coefficients with circle markers are estimated by regressing an indicator for whether a firm has a different discount rate between two observations on indicators for the number of years between the two observations. The first coefficient implies that less than 10 percent of discount rates are different for observations that are at most two quarters apart. The regression sample contains all within-firm changes in discount rates observed in our dataset (i.e., all differences between different discount rate observations of the same firm). The coefficients with square markers use an indicator for whether a firm has a different perceived cost of capital between two observations. The vertical lines measure 90 percent confidence intervals. Standard errors are clustered by firm and quarter of the later observation. The figure is from Gormsen and Huber (2025).

Figure 3: Infrequent Discount Rate Adjustment

The puzzle can be resolved once we recognize that discount rates reported by firms may not reflect the total returns earned by firms. In particular, since firms often report purposely inflated discount rates (Section 2.1), the level of these inflated discount rates does not reflect the overall level of returns that firms expect to earn. That also means, however, that the level of reported discount rates cannot always be directly compared to the firm-level cost of capital, since the firm-level cost of capital is meant to be compared to a discount rate capturing the total return earned by firms.

On conference calls, firms often describe in detail how they calculate their discount rates. This context helps researchers to identify which of the reported discount rates capture overall firm-level returns and which are purposely inflated. Among the firms that report discount rates capturing firm-level returns (i.e., including all overhead costs), the average discount rate is 11.5 percent. This average discount rate is below the average realized ROIC of roughly 13 percent and thus consistent with firms' realized returns.

The average discount rate of 11.5 percent is greater than the average cost of capital (as reported by firms themselves on conference calls), which is around 8 percent. This highlights that discount rates are elevated, even if they are elevated less than the initial surveys may have implied.

4.4 Project-Specific Versus Firm-Wide Discount Rates

According to the benchmark view, firms apply different discount rates to individual projects depending on project risk and characteristics. Implementing this project-specific approach is challenging because many methods estimate the perceived cost of capital for the whole firm (see Section 1.3). Firms would need to adjust this firm-level cost of capital for individual projects or find comparable firms with similar characteristics (Welch 2022, Jagannathan et al. 2017). The share of firms using multiple division-specific rates varies across surveys. It is around 30 percent in Bierman (1993); 50 percent in Graham and Harvey (2001), Block (2003), and Meier and Tarhan (2007); and includes "most firms" in Poterba and Summers (1995).

Whether firms use multiple discount rates or not has real implications. Compared to a world where all firms use division-specific discount rates, firms seem to overinvest in divisions whose cost of capital is higher than that of the firm's typical project and underinvest in divisions whose cost of capital is relatively lower (Krüger et al. 2015, Dessaint et al. 2021). This finding is consistent with at least some firms using just one discount rate based on the

CAPM.

The use of multiple discount rates substantially shapes whether financial markets implement social preferences and help to overcome social challenges (Chitarro et al. 2025, Pedersen 2026). For instance, Gormsen et al. (2026c) analyze how a lower cost of capital for "green" (environmentally friendly) projects in financial markets affected firms' investment rules. In response, firms using division-specific discount rates reduced the rates applied to greener divisions, relative to brown divisions, thereby shifting capital toward the greener divisions.

5 Consequences of Firms' Investment Rules

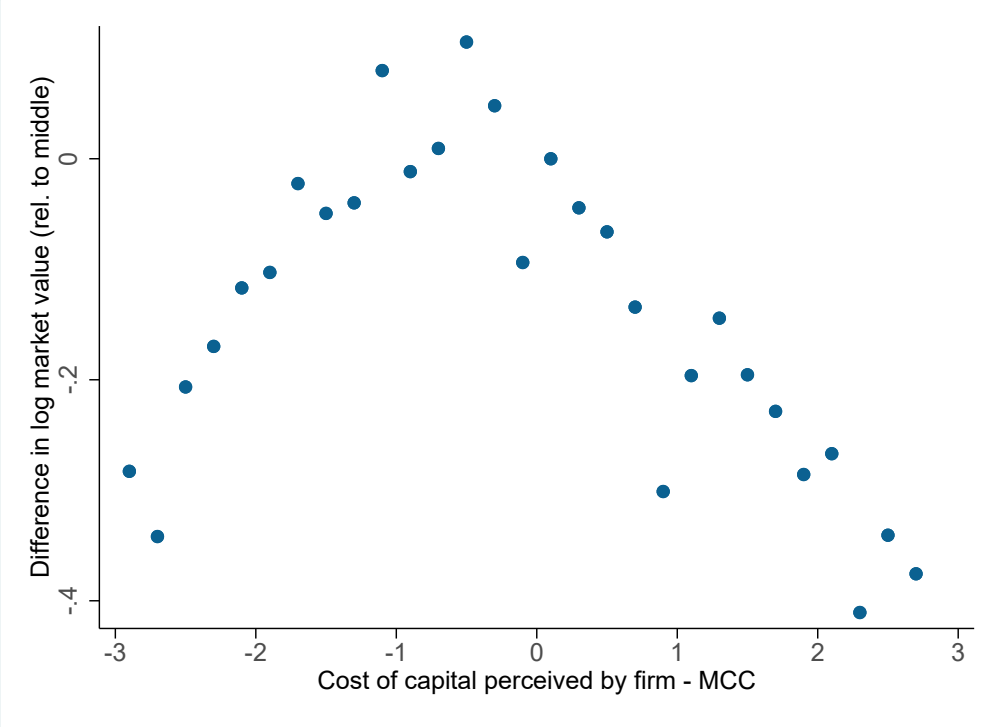
Firms' practical investment rules have large consequences for investment and the economy. We summarize the consequences separately for the two elements of equation (1). The perceived cost of capital is incorporated into discount rates in the long run, so its deviations from the opportunity cost shape the long-run allocation of capital. The discount rate is elevated and sticky, so it shapes the level of investment and how investment responds to shocks over shorter horizons.

5.1 Consequences of the Perceived Cost of Capital

In the long run, discount rates move one-to-one with the perceived cost of capital (Section 4.2). The persistent deviations between the perceived cost of capital and the opportunity cost of capital (Section 3) therefore feed into firms' investment decisions and influence the allocation of capital.

One consequence is that firms do not maximize their market value. Firms whose perceived cost of capital is too low effectively invest in projects with a lower return than what investors can earn in financial markets. This investment behavior represents a destruction of value in the eyes of investors, so market values are lower. Similarly, too high a perceived cost of capital implies that firms are forgoing value-creating investments, which also leads to lower market values. Market value should therefore be an inverted-U function of the deviation between the perceived cost of capital and the opportunity cost of capital, with the peak at a deviation of zero.

The MCC makes it possible to test this prediction, since it provides an unbiased estimate of the opportunity cost of capital for each firm (Section 3.3). Figure 4 plots the average log market value of firms with similar book equity against the deviation between their perceived



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 of Gormsen et al. (2026a). Bins for the cost of capital deviation are 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. The figure is from Gormsen et al. (2026a).

Figure 4: Inverted U: Firm Market Value and Deviations from the MCC

cost of capital and their MCC. The relation is an inverted U: market value is highest when the deviation is close to zero and declines in both directions. The decline is roughly symmetric (Gormsen et al. 2026a). Consistent with the inverted U, equity markets respond more positively to investment announcements by firms whose perceived cost of capital is above their MCC, that is, firms that invest too little relative to the market's view.

A second consequence of how firms set their perceived cost of capital is the misallocation of capital. Under the benchmark, riskier firms with a higher opportunity cost of capital receive less capital, which guarantees that capital is allocated efficiently (Hsieh and Klenow

2009, David et al. 2020). In the data, the perceived cost of capital is close to unrelated to the opportunity cost, so capital is allocated based on perceptions that have little to do with underlying risk. Through the lens of a stylized model, the excess dispersion in the perceived cost of capital reduces aggregate TFP by 5 percent relative to the benchmark allocation, or by 3 percent if the social optimum is for firms to maximize a long-run fundamental value consistent with the CAPM (Gormsen and Huber 2026). These findings suggest that the large TFP losses documented in the capital misallocation literature are partly driven by how firms perceive their cost of capital.

5.2 Consequences of Elevated Discount Rates

Firms' discount rates are elevated for two reasons. First, the perceived cost of capital is itself slightly elevated relative to standard measures of expected returns (Section 3.1). Second, the discount rate exceeds the perceived cost of capital by a few percentage points, even after correcting for cost accounting (Section 4.3). Both elevations enter the discount rate in equation (1) and therefore have the same effect on investment: firms forgo projects whose expected returns lie above the opportunity cost of capital. Investment is lower than under the benchmark, and the accepted projects have shorter horizons, because a higher discount rate penalizes cash flows that arrive further in the future.

5.3 Consequences of Sticky Discount Rates

Stickiness in discount rates mutes the response of investment to interest rates and financial markets in general. The effect is particularly strong at short horizons, but it also matters for horizons spanning multiple years.

Muted response to the cost of capital. The secular decline in interest rates between 2002 and 2019 is a case in point. Firms noticed the decline and reduced their perceived cost of capital accordingly, but they did not adjust their discount rates by nearly as much. By 2019, firms had let the wedge between discount rates and the perceived cost of capital grow by around 2.5 percentage points. This is a sizable change compared to other financial shocks: the Fed's quantitative easing programs in 2008 and 2010 reduced the cost of capital by less than 0.5 percentage points (Krishnamurthy and Vissing-Jørgensen 2011).

This wedge accounts for much of the well-documented "missing investment" puzzle in the United States. Under the benchmark, the decline in the cost of capital should have raised

investment. Instead, aggregate investment has lagged far behind what standard models imply, even accounting for intangible capital (Gutiérrez and Philippon 2017, Crouzet et al. 2022). The cumulative investment shortfall according to a standard Q-model, which takes the benchmark as given, exceeded 20 percent of the capital stock by 2019 (Figure 5). An adjusted Q-model that incorporates the observed discount rates, rather than the benchmark, eliminates most of the apparent shortfall (Gormsen and Huber 2025).

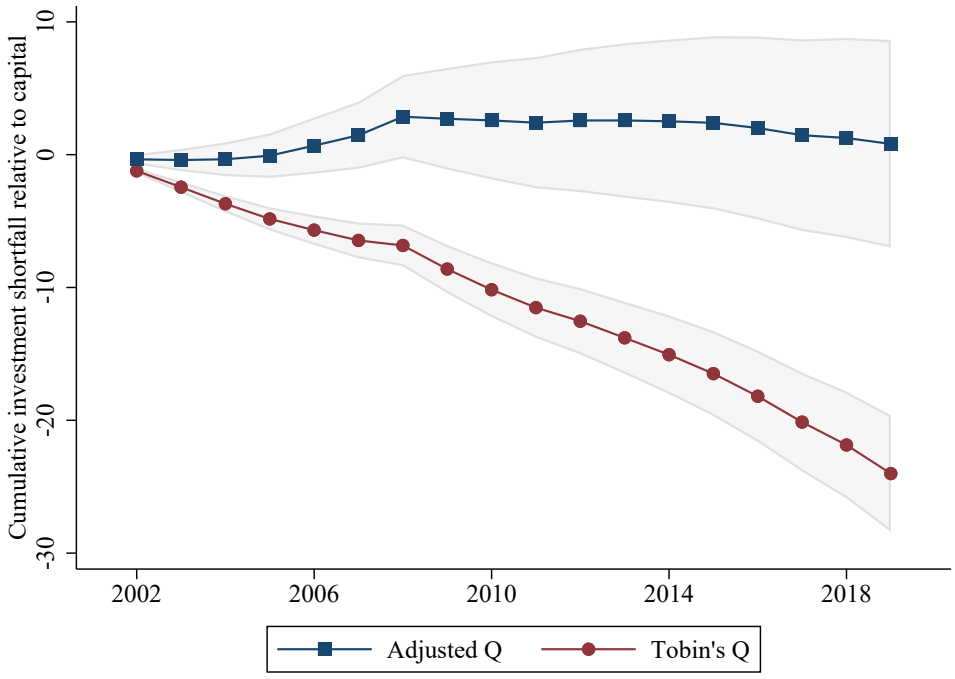
The same stickiness implies that the central bank's interest rate is less effective at stimulating investment in the short run than in textbook models. The investment sensitivity in models with sticky discount rates is consistent with the empirical evidence, in contrast to the large sensitivity implied by textbook models (Koby and Wolf 2020), and firms with stickier discount rates react less strongly to interest rate shocks (Drechsel et al. 2026).

Expected inflation and investment. Because discount rates are sticky in nominal terms, increases in expected inflation lower real discount rates and raise real investment. In benchmark models, including neoclassical and New Keynesian ones, real investment responds only to the real cost of capital and real investment opportunities, not to expected inflation on its own. In contrast, under nominally sticky discount rates, any increase in expected inflation implies that firms' discount rates fall in real terms, so firms accept more projects (Fukui et al. 2026). This direct link from expected inflation to investment has three implications.

First, demand shocks crowd in rather than crowd out investment. In standard models, shocks to household demand and government spending cause investment and consumption to move in opposite directions, which is at odds with U.S. data and gives rise to the "co-movement puzzle." Under sticky discount rates, these shocks generate expected inflation and thereby raise investment alongside consumption.

Second, monetary policy affects real outcomes even when all other prices and wages are flexible. Monetary policy shocks that change inflation expectations, such as changes in the inflation target and forward guidance, directly change real discount rates and thus investment. Sticky discount rates are therefore a distinct source of monetary non-neutrality.

Third, sticky discount rates imply very different investment dynamics from standard models: shocks to expected project returns and cash flows have larger effects, while changes in the cost of capital have weaker effects.



The figure plots the cumulative investment shortfall in the U.S. in percent of the capital stock. We separately estimate the investment shortfall implied by Tobin's Q and by adjusted Q. We calculate Tobin's Q using data on market value from the Flow of Funds (in the numerator) and data on tangible plus intangible capital from the BEA (in the denominator). Adjusted Q is calculated by correcting Tobin's Q for the wedge between discount rates and the perceived cost of capital. We estimate the relation between investment and Q using data for 1990-2002, separately for each type of Q. For the years after 2002, we then calculate cumulative residuals between observed investment and the values predicted by the 1990-2002 data. The pre-2002 adjusted Q is based on backwards extrapolated discount rate wedges. 95 percent confidence intervals (shaded region) are calculated using Newey-West standard errors adjusted for 5 lags. Investment is aggregate U.S. investment from the BEA, which includes intangibles. The figure is from Gormsen and Huber (2025).

Figure 5: Accounting for Missing Investment

5.4 The Relation Between Sticky Discount Rates and Other Mechanisms

5.4.1 Sticky Discount Rates Compared to Adjustment Costs and Inattention

To fix the implausibly high investment sensitivity to interest rate shocks in textbook models, researchers typically rely on high adjustment costs. However, these adjustment costs are often not grounded in data on the actual frictions faced by firms (Philippon 2009). In contrast, sticky discount rates generate the observed investment elasticities based on firms' investment rules measured in the data.

In general, sticky discount rates are not isomorphic to adjustment costs. Under sticky discount rates, investment responds more strongly to expected inflation and less strongly to the cost of capital, relative to standard models. Under high adjustment costs or inattention, investment responds weakly in general. The predictions of sticky discount rates also differ from models with sticky wages or prices. Hence, sticky discount rates offer distinct predictions compared to standard models.

Modigliani and Cohn (1979) explain stock prices in times of high inflation by assuming that a subset of stock owners makes a discounting mistake. In comparison, sticky discount rates explain firm investment in times of high inflation, using a very different friction in how firms discount.[8]

5.4.2 Complementary Channels: Cash Flows and Financial Constraints

The traditional cost of capital channel is weak in the short run because discount rates are sticky. However, this fact does not imply that monetary shocks have no effect on investment in general equilibrium. Two alternative channels can still lead to strong investment responses: (1) demand and expected cash flows and (2) financial constraints.

Demand and Expected Cash Flows First, monetary shocks change not only firms' cost of capital, but also household demand, for example, by moving asset prices and borrowing rates. From the point of view of firms, greater household demand raises the expected cash flows of firms' projects and thus, in the framework of equation (2), the expected returns of

[8]The two models deliver different predictions because expected inflation lowers firms' real discount rates under sticky discount rates, whereas it raises investors' real expected return in Modigliani and Cohn. However, Modigliani and Cohn assume that only some stock owners make this particular mistake and do not discuss firm behavior.

projects. Demand shocks can thereby increase investment even if discount rates are sticky. The response of consumer spending often builds up slowly, so this channel may generate the stylized hump-shaped response of investment to monetary shocks (Auclert et al. 2020).

The literature suggests that changes in firms' cash flow expectations affect investment (e.g., Malmendier and Tate 2005, Greenwood and Hanson 2015, Gennaioli et al. 2016, Coibion et al. 2020, Selgrad and Siani 2026). In fact, sticky discount rates amplify the general equilibrium effects of changes in firms' cash flows even more (due to the effects of expected inflation, see Section 5.3 on comovement), so combining models of sticky discount rates with models of cash flow expectations (e.g., Angeletos and Lian 2018, Barrero 2022, Bordalo et al. 2026) could generate additional insights.[9]

Financial Constraints Some firms do not have sufficient liquid funds to undertake all projects whose expected returns lie above their discount rates. Monetary shocks can make more liquid funds available to such constrained firms and thereby raise investment (e.g., Bernanke et al. 1999, Auclert et al. 2025).

Conceptually, the traditional cost of capital channel implicit in the benchmark view differs from the financial constraints channel. The cost of capital channel operates if the expected returns of alternative projects in financial markets change and firms accordingly adjust their discount rates (see Section 1.2). In contrast, the financial constraints channel operates if the liquid funds available to firms change. As a result, the financial constraints channel can affect investment even if firms' cost of capital and discount rates do not adjust. For instance, when liquid funds become scarce, constrained firms may maintain the same discount rate (i.e., required long-run return), but additionally introduce the extra condition that accepted projects generate liquid funds already in the first years of operation. In this case, investment falls because of the financial constraints channel independently of the traditional cost of capital channel and discount rates.

Sticky discount rates and financial constraints are complementary features of firm investment decisions. Relative to models with financial constraints, the predictions of sticky discount rate models (e.g., with respect to demand shocks and inflation non-neutrality, see Section 5.3) are distinct. Constrained firms are often young, small, and dependent on bank fi-

[9] In some classical models, firms' expectations of cash flows and their discount rates are "coherent," which means they jointly produce firms' current market values. In these models, biases in cash flows and discount rates cancel out and firms make value-maximizing decisions even if they have biased cash flow expectations. However, such coherence relies on strong assumptions (Ben-David and Chinco 2026a) and is rejected by data on firm decisions (Gormsen and Huber 2026, Dessaint et al. 2026b).

nancing (Kashyap et al. 1993, Gertler and Gilchrist 1994, Hadlock and Pierce 2010), whereas all types of firms maintain sticky discount rates.

6 Firms' Objectives and Corporate Governance

The behavior of firms' discount rates raises fundamental questions about what firms do. According to the standard models generating the benchmark view, firms care only about their current market value. However, these models may misspecify firms' objectives and therefore fail to predict firms' discount rate setting. The literature has explored several alternative hypotheses for firms' objectives that may be informative about firms' decision rules.

One hypothesis is that firms target fundamental values because contemporaneous financial markets are mispriced (Stein 1996, Gormsen and Huber 2026). The question for researchers is then: what is this fundamental or future value and what behavior does it predict? The CAPM is one candidate as it captures fundamental risk in simple models, accounts for some variation in the perceived cost of capital, and is frequently taught in business schools. However, the CAPM is not the full answer because it captures relatively little variation in the actual discount rates used by firms and in expected stock returns even over future horizons.

Another hypothesis is that firms' stakeholders care about non-standard metrics. For instance, firms may target financial ratios like earnings per share (EPS) or return on equity (ROE) to please investors (Terry 2023, Pennacchi and Santos 2021, Ben-David and Chinco 2026b). Surveys suggest that firms consider not just investors, but also additional stakeholders in their decisions, such as employees and customers (Brounen et al. 2004, Graham 2022). However, market values remain the most commonly cited objectives in these surveys.

Senior managers and division-level managers may also have non-standard preferences. Managers may care about the size of the overall firm and their division because a corporate empire brings prestige (Baumol 1959, Jensen 1986, Stein 2003) or care about cultural factors (Grennan and Li 2023). Managers may also suffer from inertia (DellaVigna and Gentzkow 2019), prefer the "quiet life" (Hicks 1935, Bertrand and Mullainathan 2003, Giroud and Mueller 2010), and want full discretion over their decisions (Pugsley and Hurst 2011), even if it comes at the expense of market values.

Sticky discount rates (and, in general, inflexible investment rules) could be a corporate governance response to the non-standard preferences of investors and managers, in particular

because it is hard to agree on the true cost of capital. For example, sticky discount rates stop division-level managers from lowering discount rates and overinvesting. They could also serve as a disciplining device to force managers to develop and accept only high-quality projects, in particular in firms facing little competition (Giroud and Mueller 2010). Indeed, the data show that firms facing higher organizational frictions and lower competition maintain stickier discount rates (Fukui et al. 2026). In addition, the high average level of discount rates could mean firms need to raise less equity, which keeps EPS and ROE relatively high in some cases.

Firms do not hide, but openly discuss their discount rates on conference calls. For instance, 59 percent of firms with an elevated or sticky discount rate argue that their behavior is good for firm performance and stability, whereas 33 percent argue it is the right way to account for risk. As discussed in Section 1.3, it is possible that both organizational frictions and behavioral drivers play a role.

One may wonder whether the market for corporate control should discipline firms to maximize market values. Corporate arbitrageurs could, in a frictionless world, take over firms, impose market value-maximizing behavior, and sell at a profit (Jensen and Ruback 1983, Jensen 1986). The difficulty with this logic is the corporate version of the limits to arbitrage argument (Shleifer and Vishny 1997). Arbitrageurs need to accumulate a large capital base, expose themselves to idiosyncratic risk, and buy at a premium. In addition, they may need to convince other investors to influence managerial decisions, which may be difficult given the overlap between managers' and analysts' views.

7 Conclusion

Economists have long assumed that firms base their investment decisions on the opportunity cost of capital. The assumption was deduced from theory, because the investment rules firms actually use could not be observed.

Surveys and textual analysis of corporate statements have allowed researchers to look inside the black box. Firms use investment rules based on discount rates, but these discount rates differ fundamentally from the cost of capital. The first component of discount rates is firms' perceived cost of capital, which is almost unrelated to the opportunity cost of capital required by financial investors. The second component is a discount rate wedge, which firms deliberately introduce to give their discount rates a life of their own. These wedges imply that discount rates are higher and stickier than the perceived cost of capital. Discount rates move one-to-one with the perceived cost of capital only over horizons longer than ten years.

The real-world behavior of firms' discount rates has wide consequences. Firms whose perceptions are further from the opportunity cost have lower market values, and the dispersion in perceptions across firms reduces aggregate TFP by 3 to 5 percent. Because discount rates are sticky, investment responds weakly to interest rates and to financial markets: firms let the wedge between their discount rates and their perceived cost of capital grow by 2.5 percentage points as interest rates fell after 2002, which accounts for much of the missing investment in the United States. And because discount rates are sticky in nominal terms, expected inflation lowers real discount rates and raises investment, so that demand shocks crowd in investment, monetary policy has real effects even when all other prices are flexible, and investment comoves with consumption over the business cycle.

Taken together, the research suggests that the investment rule at the center of essentially all models of firm behavior is not the one firms use. Replacing it with the rule observed in practice changes the analysis of capital allocation, of the transmission of financial and monetary shocks, and of business cycles.

The recent findings raise as many questions as they answer. For one, why do firms use discount rates the way they do? The early findings suggest that the answer may be a combination of forces. For instance, firms may not know their cost of capital and make mistakes when estimating it. They may also incorporate frictions and organizational constraints into their discount rates. Understanding how these forces influence firms' discount rates is an important next step for the literature.

A separate question is how to better incorporate the real-world behavior of discount rates into models. The research summarized in this piece shows that accounting for real-world discount rates can substantially alter the predictions of standard models—and help explain realized outcomes. In order to fully incorporate real-world discount rates into models, we need tractable expressions of the rule, which could take the shape of reduced-form representations. But in certain instances, we will need microfoundations of discount rates to fully explain economic phenomena.

Finally, the research has a practical role. Firms adopted investment rules based on the cost of capital in the first place because business schools began teaching them in the late 1960s, and firms that adopted the rules were rewarded by equity investors. The findings reviewed here suggest a second round of the same process. To the extent that firms misperceive their cost of capital, better teaching and research aimed at practitioners can narrow the gap: an unbiased estimate of the opportunity cost of capital now exists, and firms whose perceived cost of capital is closer to it have higher market values.

References

Angeletos, G.-M. and C. Lian (2018): "Forward Guidance without Common Knowledge," *American Economic Review*, 108, 2477–2512.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," .

Auclert, A., M. Rognlie, L. Straub, and T. Winberry (2025): "New Keynesian Economics with Household and Firm Heterogeneity," .

Ayyagari, M., A. Demirgüç-Kunt, and V. Maksimovic (2024): "The Rise of Star Firms: Intangible Capital and Competition," *Review of Financial Studies*, 37, 882–949.

Barrero, J. M. (2022): "The Micro and Macro of Managerial Beliefs," *Journal of Financial Economics*, 143, 640–667.

Barry, J. W., B. Carlin, A. Crane, and J. R. Graham (2026): "Project Development with Delegated Bargaining: The Role of Elevated Hurdle Rates," *Journal of Financial Economics*, 178, 10424.

Bastianello, F. (2025): "Time-series and Cross-Section of Risk Premia Expectations: Evidence from Financial Analysts," .

Bastianello, F., P. Décaire, and M. Guenzel (2025): "Mental Models and Financial Forecasts," .

Baumol, W. J. (1959): *Business Behavior, Value and Growth*, Harcourt, Brace & World New York.

Ben-David, I. and A. Chinco (2026a): "Crimes Against Campbell-Shiller," .

——— (2026b): "The Max EPS Paradigm for Corporate Finance," .

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of macroeconomics*, 1, 1341–1393.

Bertrand, M. and S. Mullainathan (2003): "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences," *Journal of Political Economy*, 111, 1043–1075.

Bianchi, F., D. Lee, S. C. Ludvigson, and S. Ma (2026): "The Prestakes of Stock Market Investing," .

Bierman, H. (1993): "Capital Budgeting in 1992: A Survey," *Financial Management*, 22, 24–24.

Bils, M. and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985.

Block, S. (2003): "Divisional Cost of Capital: a Study of its Use by Major US Firms," *The Engineering Economist*, 48, 345–362.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2025): "Finance Without Exotic Risk," *Journal of Financial Economics*, 173, 104145.

Bordalo, P., N. Gennaioli, A. Shleifer, and S. J. Terry (2026): "Real Credit Cycles," *American Economic Review*, 116, 1274–1308.

Bouchaud, J.-P., P. Krueger, A. Landier, and D. Thesmar (2019): "Sticky Expectations and the Profitability Anomaly," *Journal of Finance*, 74, 639–674.

Brealey, R. A., S. C. Myers, and F. Allen (2011): "Principles of Corporate Finance,

10th Edition," .

Brigham, E. F. (1975): "Hurdle rates for screening capital expenditure proposals," *Financial Management*, 17–26.

Brounen, D., A. De Jong, and K. Koedijk (2004): "Corporate Finance in Europe: Confronting Theory with Practice," *Financial mMnagement*, 71–101.

Bruner, R. F., K. M. Eades, R. S. Harris, and R. C. Higgins (1998): "Best Practices in Estimating the Cost of Capital: Survey and Synthesis," *Financial Practice and Education*, 8, 13–28.

Campbell, J. Y. and S. B. Thompson (2008): "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21, 1509–1531.

Chitarro, L., M. Piazzesi, M. Schneider, and M. Sena (2025): "Asset Returns as Carbon Taxes," .

Cho, T. and C. Polk (2024): "Putting the Price in Asset Pricing," *Journal of Finance*, 79, 3943–3984.

Chodorow-Reich, G. (2025): "The Neoclassical Theory of Firm Investment and Taxes: A Reassessment," Tech. rep.

Chodorow-Reich, G., M. Smith, O. M. Zidar, and E. Zwick (2026): "Tax Policy and Investment in a Gobal Economy," *American Economic Review*, forthcoming.

Christiano, L. J., M. S. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

Cloyne, J., C. Ferreira, M. Froemel, and P. Surico (2023): "Monetary Policy, Corporate Finance, and Investment," *Journal of the European Economic Association*, 21, 2586–2634.

Cochrane, J. H. (1991): "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations," *Journal of Finance*, 46, 209–237.

Coibion, O., Y. Gorodnichenko, and T. Ropele (2020): "Inflation Expectations and Firm Decisions: New Causal Evidence," *Quarterly Journal of Economics*, 135, 165–219.

Crouzet, N., J. C. Eberly, A. L. Eisfeldt, and D. Papanikolaou (2022): "The Economics of Intangible Capital," *Journal of Economic Perspectives*, 36, 29–52.

Cummins, J. G., K. A. Hassett, and R. G. Hubbard (1994): "A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments," *Brookings Papers on Economic Activity*, 1994, 1–74.

David, J. M., L. Schmid, and D. Zeke (2020): "Risk-Adjusted Capital Allocation and Misallocation," FRB of Chicago Working Paper No. 2020-34.

Décaire, P. and J. R. Graham (2026): "Valuation Fundamentals," .

DellaVigna, S. and M. Gentzkow (2019): "Uniform Pricing in U.S. Retail Chains," *Quarterly Journal of Economics*, 134, 2011–2084.

Dessaint, O., T. Foucault, and L. Frésard (2026a): "The Horizon of Investors' Information and Corporate Investment," .

Dessaint, O., N. Gondhi, and J. Peress (2026b): "Learning about Discount Rates," .

Dessaint, O., J. Olivier, C. A. Otto, and D. Thesmar (2021): "CAPM-Based Company (Mis)valuations," *Review of Financial Studies*, 34, 1–66.

Drechsel, T., D. Lewis, D. Melcangi, and L. Pilossoph (2026): "The Investment Channel of Monetary Policy: Disentangling Firm Heterogeneity," .

Engelberg, J., R. D. McLean, and J. Pontiff (2020): "Analysts and Anomalies," *Journal of Accounting and Economics*, 69, 101249.

Fama, E. F. and K. R. French (1997): "Industry Costs of Equity," *Journal of Financial Economics*, 43, 153–193.

Frank, M. Z. and T. Shen (2016): "Investment and the Weighted Average Cost of Capital," *Journal of Financial Economics*, 119, 300–315.

Fukui, M., N. J. Gormsen, and K. Huber (2026): "Sticky Discount Rates," .

Gabaix, X. (2025): "A Theory of Complexity Aversion," .

Gennaioli, N., Y. Ma, and A. Shleifer (2016): "Expectations and Investment," *NBER Macroeconomics Annual*, 30, 379–431.

Gertler, M. and S. Gilchrist (1994): "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics*, 109, 309–340.

Gilchrist, S. and E. Zakrajšek (2007): "Investment and the Cost of Capital: New Evidence from the Corporate Bond Market," NBER Working Paper 13174.

Giroud, X. and H. Mueller (2010): "Does Corporate Governance Matter in Competitive Industries?" *Journal of Financial Economics*, 95, 312–331.

Gitman, L. J. and J. R. Forrester Jr. (1977): "A Survey of Capital Budgeting Techniques Used by Major U.S. Firms," *Financial Management*, 66–71.

Goodman, L., A. Isen, and J. Richmond (2025): "Tax Policy, Investment, and Firm Financing: Evidence from the US Interest Limitation," .

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

——— (2026): "Firms' Perceived Cost of Capital," *Quarterly Journal of Economics*, forthcoming.

Gormsen, N. J., K. Huber, and T. I. Jensen (2026a): "A Market-Based Cost of Capital," .

Gormsen, N. J., K. Huber, and J. Kwon (2026b): "The Rise of Capital Discipline," Working paper, Copenhagen Business School and University of Chicago.

Gormsen, N. J., K. Huber, and S. Oh (2026c): "Climate Capitalists," *Journal of Financial Economics*, forthcoming.

Graham, J. R. (2022): "Presidential Address: Corporate Finance and Reality," *Journal of Finance*, 77, 1975–2049.

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.

Greenwood, R. and S. G. Hanson (2015): "Waves in Ship Prices and Investment," *Quarterly Journal of Economics*, 130, 55–109.

Grennan, J. and K. Li (2023): "Corporate Culture: A Review and Directions for Future Research," *Handbook of Financial Decision Making*, 112–132.

Gutiérrez, G. and T. Philippon (2017): "Investment-Less Growth: An Empirical Investigation," NBER Working Paper 22897.

Hadlock, C. J. and J. R. Pierce (2010): "New evidence on measuring financial constraints: Moving beyond the KZ index," *Review of Financial Studies*, 23, 1909–1940.

Hicks, J. R. (1935): "Annual Survey of Economic Theory: The Theory of Mnopoly," *Econometrica: Journal of the Econometric Society*, 1–20.

Hommel, N., A. Landier, and D. Thesmar (2025): "Corporate Valuation: An Empirical Comparison of Discounting Methods," .

House, C. L. and M. D. Shapiro (2008): "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," *American Economic Review*, 98, 737–68.

Hsieh, C.-T. and P. J. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.

Jacobs, M. T. and A. Shivdasani (2012): "Do You Know Your Cost of Capital?" *Harvard Business Review*, 90, 118–124.

Jagannathan, R., J. Liberti, B. Liu, and I. Meier (2017): "A Firm's Cost of Capital," *Annual Review of Financial Economics*, 9, 259–282.

Jagannathan, R., D. A. Matsa, I. Meier, and V. Tarhan (2016): "Why Do Firms Use High Discount Rates?" *Journal of Financial Economics*, 120, 445–463.

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

Jensen, M. C. and R. S. Ruback (1983): "The Market for Corporate Control: The Scientific Evidence," *Journal of Financial Economics*, 11, 5–50.

Jensen, T. I. (2026): "Subjective Risk and Return," .

Jorgenson, D. (1963): "Capital Theory and Investment Behavior," *American Economic Review*, 53, 247–259.

Kashyap, A. K., J. C. Stein, and D. W. Wilcox (1993): "Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance," *American Economic Review*, 83, 78–98.

Ke, S. B. (2026): "Dissecting Misreactions in Earnings Expectations: Evidence from Analyst Reports," .

Kelly, B. and D. Xiu (2023): "Financial Machine Learning," *Foundations and Trends in Finance*, 13, 205–363.

Keloharju, M., J. T. Linnainmaa, and P. Nyberg (2021): "Long-Term Discount Rates Do Not Vary Across Firms," *Journal of Financial Economics*, 141, 946–967.

Koby, Y. and C. Wolf (2020): "Aggregation in Heterogeneous-Firm Models: Theory and Measurement," .

Krishnamurthy, A. and A. Vissing-Jørgensen (2011): "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity*, 2, 215–265.

Krüger, P., A. Landier, and D. Thesmar (2015): "The WACC Fallacy: The Real Effects of Using a Unique Discount Rate," *Journal of Finance*, 70, 1253–1285.

Malmendier, U. and G. Tate (2005): "CEO Overconfidence and Corporate Investment,"

Journal of Finance, 60, 2661–2700.

Mankiw, N. G. and R. Reis (2002): "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117, 1295–1328.

Mauboussin, M. J. and D. Callahan (2022): "Return on Invested Capital: How to Calculate ROIC and Handle Common Issues," *Consilient Observer: Counterpoint Global Insights*.

Meier, I. and V. Tarhan (2007): "Corporate Investment Decision Practices and the Hurdle Rate Premium Puzzle," .

Modigliani, F. and R. A. Cohn (1979): "Inflation, Rational Valuation and the Market," *Financial Analysts Journal*, 24–44.

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

——— (1963): "Corporate Income Taxes and the Cost of Capital: A Correction," *American Economic Review*, 53, 433–443.

Mukhlynina, L. and K. G. Nyborg (2016): "The Choice of Valuation Techniques in Practice: Education Versus Profession," .

Nagel, S. (2021): "Machine Learning in Asset Pricing," .

Nagel, S. and Z. Xu (2022): "Asset Pricing with Ffading Memory," *Review of Financial Studies*, 35, 2190–2245.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

Ottonello, P. and T. Winberry (2020): "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 88, 2473–2502.

Pedersen, L. H. (2026): "Can Sustainable Finance Save the Planet?" *Journal of Finance: Insights and Perspectives*.

Pennacchi, G. G. and J. A. Santos (2021): "Why Do Banks Target ROE?" *Journal of Financial Stability*, 54, 100856.

Philippon, T. (2009): "The Bond Market's Q," *Quarterly Journal of Economics*, 124, 1011–1056.

Poterba, J. M. and L. H. Summers (1995): "A CEO Survey of US Companies' Time Horizons and Hurdle Rates," *MIT Sloan Management Review*, 37, 43–53.

Pugsley, B. W. and E. Hurst (2011): "What Do Small Businesses Do?" *Brookings Papers on Economic Activity*, 43, 73–142.

Puri, I. (2025): "Simplicity and isk," *Journal of Finance*, 80, 1029–1080.

Rajan, R., H. Servaes, and L. Zingales (2000): "The Cost of Diversity: The Diversification Discount and Inefficient Investment," *Journal of Finance*, 55, 35–80.

Rogers, J. L., A. Van Buskirk, and S. L. Zechman (2011): "Disclosure Tone and Shareholder Litigation," *The Accounting Review*, 86, 2155–2183.

Rognlie, M. (2019): "Comment on "Accounting for Factorless Income"," in *NBER Macro Annual 2018, volume 33*, University of Chicago Press, 235–248.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Deriva-

tion and Implications," *American Economic Review*, 94, 1055–1084.

Schall, L. D., G. L. Sundem, and W. R. Geijsbeek (1978): "Survey and Analysis of Capital Budgeting Methods," *Journal of Finance*, 33, 281–287.

Scharfstein, D. S. and J. C. Stein (2000): "The Dark Side of Internal Capital Markets: Divisional Rent-Seeking and Inefficient Investment," *Journal of Finance*, 55, 2537–2564.

Selgrad, J. and K. Siani (2026): "Monetary Policy and Investment Plans," .

Sharpe, S. A. and G. A. Suarez (2021): "Why Isn't Business Investment More Sensitive to Interest Rates? Evidence from Surveys," *Management Science*, 67, 720–741.

Shleifer, A. and R. W. Vishny (1997): "The Limits of Arbitrage," *Journal of Finance*, 52, 35–55.

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

——— (2003): "Agency, Information and Corporate Investment," *Handbook of the Economics of Finance*, 1, 111–165.

Summers, L. H. (1986): "Investment Incentives and the Discounting of Depreciation Allowances," NBER Working Paper 1941.

Terry, S. J. (2023): "The Macro Impact of Short-Termism," *Econometrica*, 91, 1881–1912.

Thesmar, D. and E. Verner (2025): "Beliefs and Stock Market Fluctuations: New Evidence from the Past Seven Decades," .

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

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., M. Boons, C. C. Opp, and A. Tamoni (2023): "Dynamic Asset (Mis)pricing: Build-Up Versus Resolution Anomalies," *Journal of Financial Economics*, 147, 406–431.

Welch, I. (2022): "Corporate Finance, 5th Edition," .

Zwick, E. and J. Mahon (2017): "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review*, 107, 217–48.