The Rise of Capital Discipline*

Niels Joachim Gormsen Kilian Huber Jingoo Kwon

October 2026

Please click here for the latest version

Abstract

Modern economics assumes that firms have always used disciplined investment rules based on the cost of capital. However, these rules were only developed around 1960 and entered corporate practice slowly. We study their adoption using a text-based proxy from the annual reports of listed U.S. firms. The share of adopting firms went from 0% in 1963 to 25% in 1990. In the years after adoption, firms had higher market valuations (relative to book assets), higher returns on invested capital, and lower capital stocks, because previous accounting-based rules had led to value-destroying overinvestment. The findings may account for the failure of standard models to explain pre-1990 investment and highlight the importance of firms' investment rules.

*Gormsen: Copenhagen Business School and University of Chicago, njg.fi@cbs.dk. Huber: University of Chicago, kilianhuber@uchicago.edu. Kwon: Copenhagen Business School, jik.fi@cbs.dk.

1 Introduction

Modern firms are taught to allocate their capital with great discipline. Only investments whose expected profits exceed the cost of capital are supposed to get approved. Modern economists take such capital discipline for granted. Workhorse models in every field assume that firms have always followed investment rules based on the cost of capital. However, this assumption is difficult to reconcile with historical knowledge. The modern notion of capital discipline and the cost of capital framework were only developed around 1960. It was Modigliani and Miller (1958, 1963) that explained to economists how the cost of capital could be used to maximize value. When did modern investment rules enter corporate practice? How do they affect firm outcomes? And can their adoption explain the poor performance of academic models?

In this paper, we construct a text-based proxy for when firms adopted an investment rule based on the cost of capital. We find that firm market values increased (relative to book assets), returns on invested capital increased, and capital fell after firms adopted the rule. The findings suggest that modern capital discipline replaced accounting-based rules that had generated value-destroying overinvestment. We also document a role for executive education. The cost of capital was taught relatively intensively at the Chicago School of Business during the 1960s and 1970s and firms with Chicago-educated executives adopted the modern rules sooner. Finally, we argue that the slow adoption of modern rules can explain why textbook models fail to explain investment before 1990.

We begin by describing the history of corporate investment decisions. Surveys suggest that the vast majority of large firms used rules based on book profits before 1970, for instance, requiring that investments generate a minimum accounting profit. Some firms also discounted cash flows using interest rates and used other shortcuts. Using a simple model, we illustrate that such investment rules generate overinvestment in the eyes of the firm's investors because they understate the opportunity cost of capital. To properly account for the cost of capital, firms need to consider the cost of their equity financing. Modigliani and Miller clarified that investment decisions maximize economic value only when they consider the full cost of capital (WACC) in financial markets. In theory, firms using rules based on the cost of capital generate more value for their investors and act more in line with the standard assumptions used by economists.

To identify when firms adopted investment rules based on the cost of capital, we collect annual reports of listed U.S. firms from 1963 to 1991. We search for three keywords that are associated with modern capital discipline: "cost of capital," "discount rate," and "hurdle rate." We analyze only keyword mentions that are related to a firm discussing its approach

to investment. For 772 listed U.S. firms, the dataset records whether and when a keyword appears. We use the year of first mention as a proxy for when the firm adopted an investment rule based on the cost of capital.

The share of firms that have adopted rose from 0% in 1963 to 10% in 1980 and 25% in 1990. The data thus suggest a relatively slow diffusion even among large listed firms.[1] The finding of slow diffusion is consistent with managers sticking to established rules because of information barriers and weak incentives to innovate, as in Giroud and Mueller (2010) and Bloom, Eifert, Mahajan, McKenzie, and Roberts (2013).

We examine how firms' outcomes developed after adoption. Inspired by the theory, we first analyze the market-to-book ratio of firms, measured by Tobin's q. Before adoption, q follows a similar path for adopters and comparison firms. This suggests that adoption was not driven by unusual events at the adopting firms. After adoption, q increases slowly over time and remains elevated several years out. The q of adopters is roughly 0.1 points higher relative to non-adopters after adoption. This estimate is economically meaningful as the standard deviation of q in the sample is around 1. The findings are similar controlling for shocks to different industries and to firms with different initial size and valuation. We also obtain similar findings when comparing only firms that eventually adopt to each other, so that identification comes from differences in adoption timing.

Firms' capital stock and investment returns also changed after adoption. In the model, a cost of capital rule makes firms more selective: they undertake fewer projects, but the selected projects earn higher average returns. The data support these predictions, as adoption is associated with a lower capital stock, but higher returns on invested capital. Taken together, the results are consistent with the view that investment rules based on the cost of capital are associated with greater capital discipline, valuations, and returns. More generally, they underscore the importance of management practices in shaping firm outcomes (Bloom and Van Reenen 2007).

We consider whether executive education facilitated the transmission of the academic cost of capital framework to corporate practice. We focus on the Chicago Graduate School of Business, where the cost of capital was taught relatively soon and intensively during the 1960s, for example, by Fama and Miller (1972). Using historical yearbooks, we identify firms that employed graduates of Chicago's executive education program before 1980. We find that firms with Chicago-educated executives were approximately 16.5 percentage points more likely to be an "early adopter," which means they mentioned a cost of capital keyword in their annual report between 1963 and 1979. Early adopters also saw an increase in their

[1] We end the sample in 1991 because subsequent changes in reporting practices and accounting rules made financial terminology in reports less informative about internal investment rules.

q between 1963 and 1979.

A natural question is whether adopters would have experienced higher q even if they had not adopted the cost of capital rule. The flat pre-trends in the previously described analysis suggest there were no obvious pre-adoption shocks, so adoption might have been partly linked to incidental exposure to the modern concepts. One source of exposure may be whether the firm was located near a business school teaching the modern concepts. Geographic proximity was an important determinant of whether executives attended a school. For instance, firms in the six states within 200 miles of Chicago were more likely to employ a Chicago-educated executive.

Based on this finding, we conduct an instrumental variables analysis. The outcome is the change in q between 1963 and 1979, the endogenous regressor is early adoption between 1963 and 1979, and the instruments are indicators for the six states near Chicago. The instrumented coefficient on early adoption is significantly positive, consistent with the view that variation in adoption related to geographic location led to a higher q between 1963 and 1979. These results need to be interpreted with some caution, since they are based on a sample of 349 firms in an early period.

Taken together, the results in the paper support three points. First, the adoption of modern capital discipline was a slow process over several decades, rejecting the modern modeling assumption that firms have always used rules based on the cost of capital. Second, firms had higher valuations, earned higher returns, and employed less capital after adoption. These findings suggest that the previous investment rules based on book profits were value-destroying. Third, business school education is associated with firm decisions (Shue 2013).

A large literature shows that standard investment theory fails to explain aggregate U.S. investment before the 1990s, in particular, aggregate q is not closely linked to investment in the data (Philippon 2009; Peters and Taylor 2017; Andrei, Mann, and Moyen 2019). The results in this paper offer a natural explanation. Our model shows that the predictive power of q relies on firms following an investment rule based on the cost of capital. However, with only a quarter of large firms following such a rule in 1990, it is not surprising that textbook theory fails pre-1990. In addition, the model suggests that firms using other rules tend to overinvest, which is consistent with the high investment levels and low returns observed in the pre-1990 period.

A related literature shows that, by the early 2000s, the vast majority of large firms had adopted investment rules based on the firm-level cost of capital (Graham and Harvey 2001; Graham 2022), affecting their investment and market value (Krüger, Landier, and Thesmar 2015; Gormsen, Huber, and Jensen 2026). Since 2000, firms may have become "too disciplined" in their capital budgeting, for instance, using sticky discount rates and underin-

vesting when the cost of capital falls (Gormsen and Huber 2025). These developments may be a long-run overreaction, driven by tighter corporate governance, to the lack of discipline and overinvestment experienced before 1990 (Fukui, Gormsen, and Huber 2026).

2 History of Investment Rules

Investment decisions are complicated: firms need to evaluate potential projects whose future returns are uncertain and whose funding comes partly through debt and partly through equity. Over time, firms have used different approaches, moving from simple rules based on book profits to those based on the cost of capital and value creation.

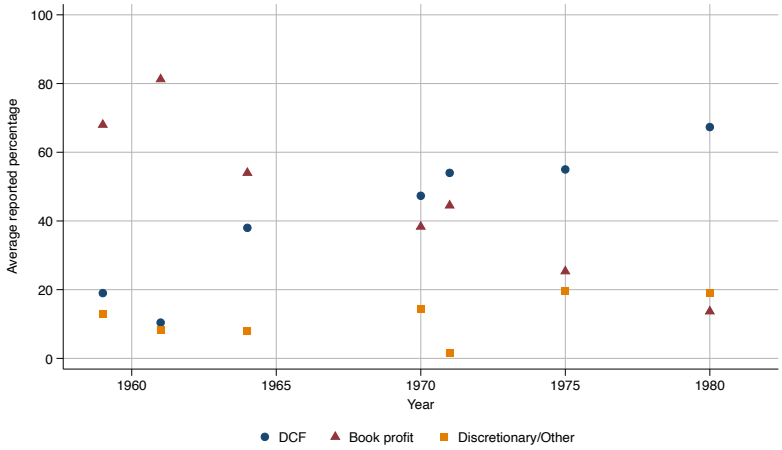
The dominant rules used in the 1950s and 1960s were based on book profits. Simply put, if a project's book profit, net of the initial investment, exceeds a certain threshold, the project gets accepted. Surveys conducted among large firms suggest that between 60% and 80% of firms used rules based on book profits to make investment decisions before 1970, as shown in Figure 1. Examples of book-based rules include profit thresholds, which require the net profit accumulated over time to exceed a minimum, or payback thresholds, which require the gross profit to pay back the initial outlay within a few years.

A smaller share of firms, between 20% and 40%, used rules based on discounted cash flows (DCF) before 1970. Fisher (1930) formalized the idea that future profits are less valuable and recommended a market interest rate as the discount rate. However, most investments are funded by internal cash and equity capital. Discounting using an interest rate does not maximize firm value because equity owners expect to earn higher returns. Dean (1951) recognized this issue and recommended a discount rate that additionally incorporated the cost of equity.

The modern approach to discounting was heavily influenced by Modigliani and Miller (1958, 1963). They recognized that the opportunity cost of capital—and therefore the proper discount rate in the eyes of the firm's investors—did not depend on the firm's immediate financing rates, but on alternative returns that investors could earn. Moreover, if the firm disciplined its investment using the appropriate cost of capital (known as WACC), it would maximize its equity market value, as we illustrate in Section 3.

Business schools soon incorporated these insights into their executive teaching. Quantitatively oriented programs, like the University of Chicago Graduate School of Business, taught investment rules based on the cost of capital already in the early 1960s (Fama and Miller 1972; Fourcade and Khurana 2013). By the early 1970s, these rules were taught at many schools, potentially impacting corporate practice (Binder and Chaput 1996). Surveys suggest that around 60% of firms applied DCF-based rules in 1980, although this figure in-

Figure 1: **Historical surveys on investment rules**



This figure summarizes surveys on investment rules used by firms. Circles denote discounted cash flow (DCF) methods, including net present value, internal rate of return, and profitability index. Triangles denote book profit-based methods, including accounting return, simple rate of return, payback, and MAPI. Squares denote methods based on discretion, necessity, urgency, or other approaches. Each point represents an average across different surveys within the same year. The surveys reported in the figure are Istvan (1961), Klammer (1972), Petty, Scott Jr, and Bird (1975), and Klammer and Walker (1984).

cludes many firms using interest rates and other shortcuts (Mao 1970). Book-based methods continued to be used by around 15% of firms in 1980 and other discretionary methods by around 20%.

The cost of capital gained further prominence in debates about corporate governance during the 1980s. Jensen and Meckling (1976) and Jensen (1986) pointed out that executives may not optimize value and instead build corporate empires. Many firms made shareholder value maximization their prime objective, which essentially meant they adopted investment rules based on the cost of capital. Private equity investors also expanded, acquiring firms with suboptimal rules and imposing discipline.

Taken together, the decades between the 1950s and the 1980s saw the rise of capital discipline. Many firms moved away from book-based and discretionary investment rules that had led to overinvestment and low returns. They adopted disciplined rules that emphasized the opportunity cost of capital and value creation.

3 Model of Investment Rules

We develop a model of a firm setting its investment rule. The model clarifies why the cost of capital is the opportunity cost of the firm's investment and allows the firm to discipline its decisions in line with the views of financial investors. In particular, firm market value and

Tobin's q (market relative to book value) are maximized when the firm uses an investment rule based on the cost of capital. In contrast, rules based on book profits or interest rates lead to overinvestment and lower q.

3.1 Firm investment decision

There are two periods, current ($t = 0$) and future ($t = 1$). Symbols and letters without a prime refer to current values at the beginning of period 0, while primes indicate future values at the end of period 1.

The firm chooses investment I in period 0. The expected value of the firm's future assets is $F(I)$ and depends on a linear production function:

$$F(I) = (1 + a)I, \tag{1}$$

where $a > 0$ captures expected cash flows, which may vary with technology or consumer demand. The firm incurs adjustment costs $\frac{b}{2B}I^2$ in period 0, where $B > 0$ is its beginning-of-period book value and $b > 0$. Each unit of investment has the same risk characteristics as the firm's existing assets and the firm's investment does not change risk pricing in financial markets.

The firm's objective is to maximize its net present value:

$$\text{NPV}(I) = \frac{F(I)}{1 + \delta} - I - \frac{b}{2B}I^2, \tag{2}$$

where δ is the firm's discount rate. For an interior solution, the first-order condition gives the firm's investment rule:

$$I^* = \frac{B(a - \delta)}{b(1 + \delta)}. \tag{3}$$

Equivalently, the investment rule can be represented as a "hurdle" threshold, where the firm invests until the marginal return per unit of expenditure equals the discount rate δ, such that $(1 + a)/(1 + bI^*/B) - 1 = \delta$.

3.2 Firm financing and market value

The current market value of equity is $E(I)$. Equity investors expect to earn a return $r_E(I)$, known as the cost of equity. Both $E(I)$ and $r_E(I)$ vary with investment I because investors base their valuations on their expectations of firm investment. There are no payouts to

equity holders, so investors expect future equity value to be:

$$E'(I) = (1 + r_E(I))E(I). \tag{4}$$

The firm finances its investment and adjustment costs with debt, so debt taken out in period 0 equals $D(I) = I + \frac{b}{2B}I^2$. The firm pays an interest rate r_D on its debt and never defaults. As a matter of accounting, the future equity value equals the firm's future assets minus debt repayments:

$$E'(I) = F(I) - (1 + r_D)D(I). \tag{5}$$

Combining the two value equations (4) and (5), investors determine the appropriate current value of the equity:

$$E(I) = \frac{F(I) - (1 + r_D)D(I)}{1 + r_E(I)}. \tag{6}$$

This condition determines the current equity market value as long as the law of one price holds (e.g., investors perceive no arbitrage opportunities).[2] The current equity value (6) can be rewritten as

$$E(I) = \frac{F(I)}{1 + k} - I - \frac{b}{2B}I^2, \tag{7}$$

where k is defined as the firm's cost of capital (WACC):

$$k = \frac{E(I)}{E(I) + D(I)}r_E(I) + \frac{D(I)}{E(I) + D(I)}r_D. \tag{8}$$

This cost of capital has an intuitive interpretation if the law of one price holds. It captures the return that the debt and equity investors could earn if they deployed their capital outside the firm. In that sense, it captures the proper opportunity cost in financial markets of the firm investing an additional unit.

We assume that $a > k$, which implies that at least the first unit of investment is optimal from the investors' point of view because it generates more value than the cost of capital. In addition, $k > r_D > 0$, which implies that the interest rate is lower than the cost of equity because debt has a primary claim over equity on the firm's future assets, a condition that usually holds in the real world.

[2]The condition does not require equity markets to be efficient or investors to be rational. See Gormsen and Huber (2026) for a model that includes irrational investors with biased expectations where this equation still holds.

3.3 Firm investment rules

The firm needs to formulate an investment rule by specifying a discount rate δ. We analyze three investment rules that have been used by firms historically: based on the cost of capital, the interest rate, and the book profit.

We show how the rules affect Tobin's q, defined as market value $E(I)$ over book value B (measured at the beginning of period 0 before debt financing and investment):

$$q(I) = \frac{E(I)}{B}. \tag{9}$$

Intuitively, q captures how efficiently the firm generates market value for its investors, given its book assets. A higher q implies that equity investors value the investment strategy of the firm more.

We also show how the rules influence the relation between q and the firm's investment rate $i = \frac{I}{B}$.

Cost of capital rule: $\delta = k$. A firm using the cost of capital k as its discount rate δ maximizes its equity market value and q. The firm's investment problem (2) is identical to the equation determining market value (7) as long as $\delta = k$. Hence, by setting $\delta = k$, a firm optimizing its NPV also optimizes market value and q. Intuitively, the cost of capital k is the discount rate used by financial investors and captures the investors' opportunity cost of providing capital to the firm. Setting $\delta = k$ implies that the firm internalizes this opportunity cost and gets rewarded with higher q.

Under the cost of capital rule, the investment rate equals $i_{\text{coc}} = (a - k)/[b(1 + k)]$. Substituting the investment rate into the equation determining market value (7) implies that $q = \frac{b}{2} i_{\text{coc}}^2$ and hence $i_{\text{coc}} = \sqrt{2q/b}$. This means q is a sufficient statistic for the investment rate i, keeping the adjustment cost parameter b constant. For instance, if technology or demand result in higher a, the firm's q and investment rate rise. Similarly, if k decreases, q and the investment rate rise. The cost of capital rule therefore generates the well-known textbook result that q always predicts investment.

Interest rate rule: $\delta = r_D$. Since the interest rate is usually lower than the cost of capital (i.e., $k > r_D$), the interest rate rule generates a lower discount rate and more investment, relative to the cost of capital rule. It therefore leads to lower equity market value and q.

Under the interest rate rule, the investment rate equals $i_D = (a - r_D)/[b(1 + r_D)]$, implying

$$q = \frac{b(1 + 2r_D - k)}{2(1 + k)} i_D^2 - \frac{k - r_D}{1 + k} i_D. \tag{10}$$

Now, the relation between q and the investment rate i is not always positive. If technology or demand result in higher a, q and i both increase. However, if k decreases and r_D is fixed, q rises and i is unchanged. And if k is fixed and r_D falls, q falls and i rises.

Book profit rule: $\delta = (a - bi_{\text{book}})/(1 + bi_{\text{book}})$**.** This rule expands investment as long as book profit before interest payments remains at least as large as $x > 0$. Book profit net of the investment and before interest payments is $B(ai - \frac{b}{2}i^2)$. The firm chooses the largest investment level consistent with this floor, so the investment rate is $i_{\text{book}} = [a + \sqrt{a^2 - 2bx/B}]/b$. The marginal return per unit of expenditure and the implied hurdle rate equal $\delta_{\text{book}} = (a - bi_{\text{book}})/(1 + bi_{\text{book}})$. Assuming that $2bx/B < a^2$ and $x > kD(Bi_{\text{book}})$, the profit target is attainable and equity market value remains positive. Under these assumptions, the implied hurdle rate of the book profit rule is lower than the cost of capital and the interest rate. The book profit rule thus yields overinvestment, a lower q, and lower equity market value relative to the cost of capital rule.

Under the book profit rule, total book profit equals x, implying $q = [x/B - k(i_{\text{book}} + \frac{b}{2}i_{\text{book}}^2)]/(1 + k)$ and hence

$$i_{\text{book}} = \frac{\sqrt{1 + \frac{2b}{k}\left[\frac{x}{B} - (1 + k)q\right]} - 1}{b}. \tag{11}$$

The relation between q and the investment rate i can be zero or negative. If technology or demand result in higher a, q falls and i rises. If k decreases, q rises and i is unchanged. The book profit rule therefore largely invalidates the classical textbook prediction of q predicting i.

4 Data

This section describes the data. We construct a proxy for when a firm began using an investment rule based on the cost of capital. The proxy captures when the firm first used keywords related to the cost of capital in its annual report. We combine the proxies with firm-level accounting data from Compustat and market values from CRSP. More details about the data collection are in Appendix B.

4.1 Text analysis of annual reports

We collect the annual reports of listed U.S. firms from historical archives and online databases. We search for the reports of all firms that were listed in the S&P 500 at least once between

1963 and 1991. We additionally include listed firms from the 1960 and 1985 Fortune 500 lists.

Our aim is to construct a proxy for when firms used an investment rule related to the cost of capital. In a first step, we read the annual reports using automated tools and record mentions of three keywords: "cost of capital," "discount rate," and "hurdle rate." In a second step, we remove mentions of the keywords that are not related to an investment rule. To do so, we manually review 500 appearances of each keyword and note terms that appear in conjunction with a keyword but imply that the use of the keyword is unrelated to an investment rule. For example, the terms "cost of capital acquisition" and "cost of capital facilities" indicate an unrelated use of a keyword. We remove all instances where a keyword appears within 10 words of such an unrelated term. We iterate this manual review procedure five times, at which point we encounter very few instances of the keyword in another context. More details on the unrelated terms are in Appendix B.

We define a firm's adoption year as the year in which one of the three keywords first appears in the firm's reports. Our baseline analysis uses an indicator for all years from the first observed adoption year, but we also conduct robustness checks where we use indicators for firms that regularly mention a keyword over time and use a fully balanced panel.

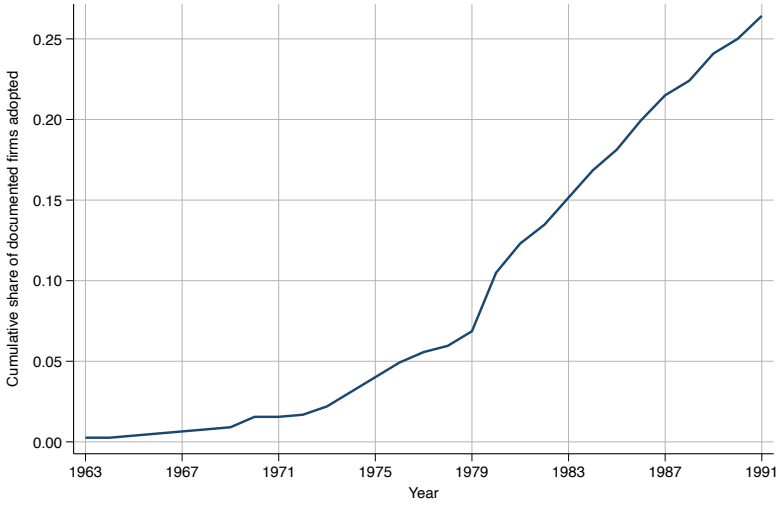
We begin our sample period in 1963 because it marks the publication year of Modigliani and Miller (1963), a seminal paper on the cost of capital. We end our sample period in 1991 because changes in reporting requirements and report structures after 1991 mean that tracking the content of reports is no longer informative about firms' internal decision-making after 1991. A key factor is that SEC reforms in 1992–95 and management consulting trends led many firms to discuss technical terms related to the cost of capital, shareholder returns, and financial markets in their reports, without really changing their investment rules.[3]

The mention of keywords is, of course, not a perfect measure of firm investment rules. However, it is plausibly correlated with how much firms rely on rules related to the cost of capital. We show in Figure 2 that the share of firms that had mentioned a keyword once increased from 0% in 1963 to around 25% in 1990, indicating a slow diffusion of the cost of capital framework over time.

A few examples illustrate how firms mention a relevant keyword. Genesco in 1973 states

[3]For instance, the SEC introduced a 1992 requirement to disclose shareholder return time series and the 1992 Caterpillar enforcement by the SEC led firms to disclose much more detailed forward-looking plans "through the eyes of management." In late 1992, the SFAS 106 and SFAS 109 accounting standards became effective, incentivizing firms to consider discount rates and future cash flows in detail. Finally, the 1995 Private Securities Litigation Reform Act made firms more likely to disclose forward-looking information and led to boilerplate disclosures (Johnson, Kasznik, and Nelson 2001). Another factor is that the concept of "economic value added" became popular as a management tool after 1991, leading some firms to discuss concepts related to value creation and the cost of capital without adopting them.

Figure 2: **Share of firms that have mentioned a cost of capital keyword**



This figure plots the share of firms that have mentioned a cost of capital keyword by the given year. The sample includes firms with at least one observed annual report between 1963 and 1991.

that "in view of the uncertain economic conditions and the high cost of capital, action was taken early in the year to substantially reduce planned expenditures." TRW in 1975 writes that "we must earn, at a minimum, a rate of return equal to the composite of our equity and our debt costs. As the cost of capital moves up, we must seek higher returns." Coca-Cola in 1984 reports that "management seeks solid investments that [...] offer long-term cash flow returns which exceed the company's overall cost of capital." Other firms in the dataset include Scott Paper (adopted in 1967), Tenneco (adopted in 1968), General Electric (1976), Exxon (1980), Target (1983), PepsiCo (1985), and Disney (1990).

4.2 Firm characteristics

We merge the keyword mentions to accounting data from Compustat and market values from CRSP. We manually match the firm names on the reports and ensure consistency over time, taking into account possible name changes. We measure Tobin's q as the firm's market value over book value. The market value is the sum of equity market value (from CRSP), current liabilities, long-term debt, preferred stock, deferred taxes, and investment tax credits (all from Compustat). The book value is total assets (from Compustat). In addition, we measure the firm's return on invested capital as operating income before depreciation over the lagged sum of long-term debt, current debt, and shareholders' equity minus cash (all from Compustat).

About 1,110 firms appeared in the S&P 500 at least once and we find at least one report

for 772 firms. The sample is broadly representative of industry composition in Compustat, as shown in Figure A1, and many firms' reports are observed regularly, as shown in Figure A2. We summarize the characteristics of the sample in Table 1.

Table 1: **Summary statistics**

	N	Mean	SD	P25	Median	P75
Tobin's Q	12,139	1.20	0.93	0.72	0.94	1.31
Capital (million 1991 USD)	13,822	1,966.71	3,815.34	222.50	581.80	1,712.44
ROIC (%)	12,639	29.31	30.22	17.12	25.27	36.30

This table reports summary statistics for the firm-year panel. The sample covers 1963–1991 and contains 13,935 firm-year observations for 772 firms. The table reports the mean, standard deviation, and the 25th, 50th, and 75th percentiles of each variable. Capital is measured using property, plant, and equipment and reported in millions of 1991 US dollars using the US CPI. The mean percentile of capital in the sample, relative to the Compustat universe, is 85. These percentiles are calculated using capital in 1991 US dollars. The return on invested capital (ROIC) divides the operating income before depreciation by the lagged value of [total debt plus stockholders' equity minus cash], multiplied by 100. The variables are winsorized at the 1st and 99th percentiles.

5 Adopting the Cost of Capital and Firm Outcomes

This section examines what happens to firm q, capital, and investment returns after adoption of the cost of capital rule.

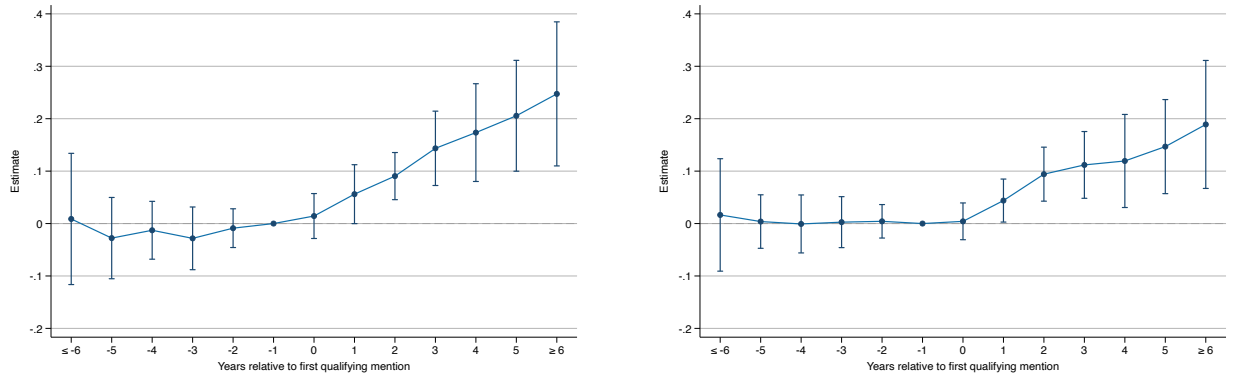
Firms relying on a cost of capital rule should achieve higher Tobin's q, as explained by the model of Section 3. We test this prediction by specifying

$$Q_{it} = \alpha_i + \alpha_t + \alpha_{s(i),t} + \alpha_{q^{BM}(i),t} + \alpha_{q^{Size}(i),t} + \sum_{k \neq -1} \beta_k \, \mathbf{1}[t - T_i = k] + \varepsilon_{it}, \tag{12}$$

where T_i is firm i's adoption year, α_i and α_t are firm and year fixed effects, and the year before adoption ($k = -1$) is the omitted year. The baseline specification includes firm and year fixed effects. The full specification adds year fixed effects interacted with: the firm's two-digit SIC industry, $\alpha_{s(i),t}$; indicators for the quintile of market value; and indicators for the quintile of book-to-market value. We measure the market value and book-to-market controls in 1959 (or the year with an observed value closest to 1959), to avoid any bias from including a potential outcome variable as a control.

For the years before adoption, the results in Figure 3 show coefficients that are close to zero and statistically not different from zero. This suggests that, on average, firms were not on different paths before their adoption year, compared to other firms, and that the adoption was not a response to unusual developments. The coefficients turn positive and statistically

Figure 3: **Tobin's q around adoption**



(a) Baseline (b) Full controls

These figures plot the coefficients β_k from equation (12) with Tobin's q as the outcome variable. The year before adoption is the omitted category. Panel (a) uses firm and year fixed effects. Panel (b) adds the year-industry fixed effects, where the industry is defined using two-digit SIC codes, the year-B/M fixed effects, constructed as the quintile of the firm's book-to-market from the year 1959 (or the closest observed year), and the year-size fixed effects, constructed as the quintile of the firm's market capitalization from the year 1959 (or the closest observed year). Standard errors are two-way clustered by firm and year. The bars show 90% confidence intervals.

significant from the first year after adoption. The buildup over time is consistent with the view that equity investors do not respond immediately to the adoption of a keyword but react over time as the firm changes its behavior. The pattern over time is similar in the specification with the baseline controls and the full controls.

To estimate an average effect, we use a specification where we include only one indicator for the years from the adoption year, $begin_{it}$, among the regressors:

$$Q_{i,t+2} = \alpha_i + \alpha_t + \alpha_{s(i),t} + \alpha_{q^{BM}(i),t} + \alpha_{q^{Size}(i),t} + \beta\, begin_{it} + \varepsilon_{it}. \tag{13}$$

The outcome is q two years ahead to account for the fact that q increases slowly over time. We find statistically significant coefficients between 0.12 and 0.17 in Table 2, depending on the controls. The magnitude is meaningful, given the standard deviation of q in the sample is about 1.

In Table A1, we find similar estimates using alternative samples and definitions of adoption. We restrict the sample to firms that adopt the cost of capital rule at some point between 1963 and 1991 in column 1, so the variation comes only from the timing of adoption. The sample in column 2 contains a fully balanced panel and in column 3 only firms with at least 20 observations. The sample in column 4 contains only firms where we observe both an annual report and q in the five-year period around the adoption year, ensuring that we measure the adoption year precisely. The adoption indicator in column 5 captures only firms that

Table 2: **Adoption and Tobin's q**

	(1)	(2)	(3)	(4)
Adoption	0.168***	0.135**	0.123**	0.129**
	(0.051)	(0.053)	(0.047)	(0.049)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Year × Industry FE		✓	✓	✓
Year × B/M FE			✓	✓
Year × Size FE				✓
Mean of Tobin's Q	1.176	1.176	1.176	1.176
Observations	11,931	11,931	11,931	11,931
R^2	0.664	0.733	0.777	0.780

This table reports the estimates of equation (13). The dependent variable is Tobin's q in year $t+2$, where t denotes the year of the adoption variable. Adoption is an indicator that equals 1 starting in the first year a cost of capital keyword is mentioned. All specifications include firm and year fixed effects. Column (1) reports the baseline specification. Column (2) adds the year-industry fixed effect, defined using two-digit SIC codes. Column (3) adds the year-B/M fixed effects, constructed as the quintile of the firm's book-to-market from the year 1959 (or the closest observed year). Column (4) adds the year-size fixed effects, constructed as the quintile of the firm's market capitalization from the year 1959 (or the closest observed year). The sample spans 1963–1991. Tobin's q is winsorized at the 1st and 99th percentiles. Standard errors are two-way clustered by firm and year and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

repeatedly mention a keyword, by setting the indicator to zero if a firm has not mentioned a keyword for five years. The adoption indicator in column 6 captures only mentions of the single keyword "cost of capital." In Table A2, we find similar estimates using alternative estimators that address the possibility of bias when adoption of a treatment is staggered (Goodman-Bacon 2021).

Firms typically used book profit-based rules before the cost of capital framework was developed, suggesting that they overinvested relative to equity investors' valuations, as explained in the model of Section 3. Firms following a cost of capital rule should be more selective: they pick fewer projects but, on average, these projects earn higher returns. Column 2 of Table 3 shows that firms' capital (measured using their property, plant, and equipment) decreased by approximately 0.08 log points after adoption. The estimate for capital is statistically significant at the 1 percent level with baseline controls and the 10 percent level with all controls. The return on invested capital (ROIC) of firms increased by 2.3 percentage points after adoption, as shown in column 4 of Table 3. The standard deviation of ROIC is 30 percentage points in the data, so the estimate is economically meaningful. The estimate for ROIC is statistically significant at the 5 percent level with baseline controls and the 10 percent level with all controls.

Table 3: **Adoption, Capital, and Investment Returns**

	(1) Log Capital	(2) Log Capital	(3) ROIC (%)	(4) ROIC (%)
Adoption	-0.180***	-0.078*	3.216**	2.314*
	(0.056)	(0.045)	(1.166)	(1.268)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Year × Industry FE		✓		✓
Year × B/M FE		✓		✓
Year × Size FE		✓		✓
Mean of LHS	5.952	5.952	29.170	29.170
Observations	11,876	11,876	11,434	11,434
R^2	0.940	0.961	0.500	0.621

This table reports the estimates of equation (13). The dependent variables are the log of capital (PPENT) and the percent return on invested capital (ROIC) in year $t+2$, where t denotes the year of the adoption variable. Adoption is an indicator that equals 1 starting in the first year a cost of capital keyword is mentioned. All specifications include firm and year fixed effects. Columns (1) and (2) report the results for log capital. Columns (3) and (4) report the results for ROIC. Columns (2) and (4) add year-industry, year-B/M, and year-size fixed effects. Industry is defined using two-digit SIC codes, B/M is constructed as the quintile of the firm's book-to-market from 1959 (or the closest observed year), and size is constructed as the quintile of the firm's market capitalization from the year 1959 (or the closest observed year). The dependent variables are winsorized at the 1st and 99th percentiles. Standard errors are two-way clustered by firm and year and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Executive Education and the Cost of Capital

In this section, we explore a specific channel through which the cost of capital framework reached firms: managerial education at the University of Chicago Graduate School of Business (henceforth GSB). The GSB taught the framework relatively intensively already during the 1960s. We find that firms geographically close to the GSB were more likely to have GSB-educated managers and to adopt the cost of capital rule before 1980, with potential consequences for their q.

6.1 Teaching the Cost of Capital at Chicago GSB

The cost of capital framework was formalized by Modigliani and Miller (1958, 1963). Miller moved from the Carnegie Institute of Technology to Chicago GSB in 1961. He was joined on the faculty by Eugene Fama in 1963. Both emphasized the cost of capital framework in their teaching throughout the 1960s. The Fama and Miller (1972) textbook is based on their 1960s course material and contains a detailed treatment of the cost of capital.

In comparison to Chicago GSB, other business schools seem to have been less focused on

teaching the cost of capital framework to executives. Details on the precise weights given to different topics are not easily available today, but we conducted several interviews with business school faculty active in the 1960s and 1970s. They suggested that Chicago GSB was teaching the cost of capital framework more intensively in the 1960s and 1970s, compared to other large business schools (e.g., Harvard, Stanford, MIT).[4]

This historical narrative raises the question whether firms with Chicago GSB-educated executives were more likely to adopt the cost of capital rule during the 1960s and 1970s. To address this question, we collect new data on executive education. Using historical yearbooks from the University of Chicago Library, we extract the names of firms that employed at least one executive who graduated from a GSB executive education program before 1980. The data reveal which firms had a GSB-educated executive at some point before 1980, but not the precise year in which a GSB-educated executive worked at a firm (because the yearbooks do not provide the full employment history of executives). In our analyses, we therefore study long differences from 1963 (start of our sample and publication year of Modigliani and Miller) to 1979.

6.2 Education, Adoption, and q

We regress the difference in q between 1979 and 1963 on an indicator for whether the firm was an "early adopter" (i.e., mentioned a cost of capital keyword in a report in the relevant period). Firms without an observed report in this period are set to missing. We control for fixed effects for industry, size quintiles, and book-to-market quintiles. We find that q increased by 0.3 points more when the firm was an early adopter in column 1 of Table 4. The coefficient is significant at the 1 percent level and consistent with the event study in Section 5.

In column 2 of Table 4, we report that having a GSB-educated executive increased the likelihood of being an early adopter by 17 percentage points. This finding is consistent with the view that a GSB education meant that firms encountered the cost of capital framework sooner and adopted it sooner.

Of course, firms adopting early may have experienced higher q even if their executives had not adopted the cost of capital. We therefore take a step back and study a less endogenous determinant of adopting: geographic proximity to Chicago. Six states are closer than 200 miles from Chicago: Illinois, Indiana, Michigan, Wisconsin, Iowa, Ohio. Firms headquartered in these states between 1963 and 1979 were more likely to have a GSB-educated executive,

[4]Modigliani moved to MIT in 1962, with joint appointments in the economics department and the business school, but the cost of capital framework seems to have been less prominent in teaching at MIT during the early 1960s.

Table 4: **Education, adoption, and the change in Tobin's q**

	(1) Δq	(2) Early adopter	(3) Chicago educated	(4) Δq
Early adopter	0.296***			0.574**
	(0.091)			(0.256)
Chicago educated		0.165***		
		(0.046)		
Near Chicago			0.133*	
			(0.076)	
Estimator	OLS	OLS	OLS	2SLS
Controls	✓	✓	✓	✓
Observations	349	349	349	349
R^2	0.658	0.311	0.185	0.653

This table reports regressions of outcome variables listed in the top row on the regressors listed in the left-hand column. The outcome in column 1 is the change in a firm's Tobin's q from 1963 to 1979. The main regressor in column 1 is an indicator for "early adopter" (i.e., firms who mentioned a cost of capital keyword in the 1963–1979 period). The outcome in column 2 is the early adopter indicator. The main regressor in column 2 is an indicator for firms with an executive who graduated from a GSB executive education program before 1980. The outcome in column 3 is the Chicago-educated indicator. The main regressor in column 3 is an indicator for firms whose headquarters were in a state within 200 miles of Chicago between 1963 and 1979. The outcome in column 4 is the change in a firm's Tobin's q from 1963 to 1979. The main regressor in column 4 is an indicator for "early adopter" instrumented using indicators for the six states near Chicago. All columns control for fixed effects for industry, size quintiles, and book-to-market quintiles. Standard errors are clustered by headquarters state and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

as shown in column 3 of Table 4, making it possible that their q also evolved differently.

We investigate whether the variation in geographic proximity to Chicago is associated with variation in q. To do so, we use indicators for the six states near Chicago as instruments for being an early adopter. The outcome is the change in q between 1963 and 1979. The two-stage least squares (2SLS) coefficient implies that being an early adopter is associated with a greater change in q, as shown in column 4 of Table 4. The coefficient is statistically different from zero at the 5 percent level.[5] The 2SLS point estimate is larger than the OLS estimate, which may be partly due to the keyword-based analysis introducing measurement error. However, the 2SLS estimate is not statistically different from the OLS estimate in column 1.

Taken together, the results are consistent with executive education influencing the adoption of the cost of capital rule and improving q. The sample of firms with available reports

[5]We conduct an Anderson–Rubin bootstrap test on the 2SLS coefficient to address the possibility of weak instruments, finding a p-value of 0.04. The first-stage regression has an F-statistic (for the six excluded instruments) of 59.6.

in this early period is small, so the results should be interpreted with caution.

7 Conclusion

We study how large U.S. firms adopted modern capital discipline between 1963 and 1991. Using annual reports, we construct a text-based proxy for when firms adopted investment rules based on the cost of capital. The framework spread slowly, rising from 0% of firms in 1963 to 25% in 1990. Before adoption, firms typically used approaches based on book profits, which led to excessive investment relative to the views of equity investors. After adoption, firms' market-to-book valuations were higher, their returns on invested capital were higher, and their capital stocks were lower.

Executive education appears to have facilitated adoption. Firms employing graduates of Chicago's executive education program were more likely to be early adopters between 1963 and 1979. Geographical proximity to Chicago predicted whether firms employed these executives. Instrumenting early adoption with headquarters location suggests a positive association between firm market value and adoption.

Standard models assume that firms have always evaluated investments against the cost of capital, leading to the prediction that investment depends on Tobin's q. Our evidence suggests that this assumption poorly describes large firms before the 1990s. We show in a simple model that rules based on book profits can generate a low or negative association between q and investment. This pattern is more consistent with the observed relation in aggregate U.S. data. The slow diffusion of capital discipline therefore offers a potential explanation for the weak empirical performance of textbook investment models.

References

Andrei, D., W. Mann, and N. Moyen (2019). Why did the q theory of investment start working? *Journal of Financial Economics 133*(2), 251–272.

Binder, J. J. and J. S. Chaput (1996). A positive analysis of corporate capital budgeting practices. *Review of Quantitative Finance and Accounting 6*(3), 245–257.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013). Does management matter? Evidence from India. *Quarterly Journal of Economics 128*(1), 1–51.

Bloom, N. and J. Van Reenen (2007). Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics 122*(4), 1351–1408.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies 91*(6), 3253–3285.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

de Chaisemartin, C. and X. D'Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

Dean, J. (1951). *Capital Budgeting: Top-Management Policy on Plant, Equipment, and Product Development*. New York: Columbia University Press.

Fama, E. F. and M. H. Miller (1972). *The Theory of Finance*. Hinsdale, Illinois: Dryden Press.

Fisher, I. (1930). *The theory of interest: As determined by impatience to spend income and opportunity to invest it*. Macmillan.

Fourcade, M. and R. Khurana (2013). From social control to financial economics: The linked ecologies of economics and business in twentieth century America. *Theory and Society 42*(2), 121–159.

Fukui, M., N. J. Gormsen, and K. Huber (2026). Sticky discount rates.

Giroud, X. and H. M. Mueller (2010). Does corporate governance matter in competitive industries? *Journal of Financial Economics 95*(3), 312–331.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Gormsen, N. J. and K. Huber (2025). Corporate discount rates. *American Economic Review 115*(6), 2001–2049.

Gormsen, N. J. and K. Huber (2026). Firms' perceived cost of capital.

Gormsen, N. J., K. Huber, and T. I. Jensen (2026). A market-based cost of capital.

Graham, J. R. (2022). Presidential address: Corporate finance and reality. *Journal of Finance 77*(4), 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*(2-3), 187–243.

Istvan, D. F. (1961). The economic evaluation of capital expenditures. *Journal of Business 34*(1), 45–51.

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review 76*(2), 323–329.

Jensen, M. C. and W. H. Meckling (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics 3*(4), 305–360.

Johnson, M. F., R. Kasznik, and K. K. Nelson (2001). The impact of securities litigation reform on the disclosure of forward-looking information by high technology firms. *Journal of Accounting Research 39*(2), 297–327.

Klammer, T. (1972). Empirical evidence of the adoption of sophisticated capital budgeting techniques. *Journal of Business 45*(3), 387–397.

Klammer, T. P. and M. C. Walker (1984). The continuing increase in the use of sophisticated capital budgeting techniques. *California Management Review 27*(1), 137–148.

Krüger, P., A. Landier, and D. Thesmar (2015). The WACC fallacy: The real effects of using a unique discount rate. *Journal of Finance 70*(3), 1253–1285.

Mao, J. C. (1970). Survey of capital budgeting: Theory and practice. *Journal of Finance 25*(2), 349–360.

Modigliani, F. and M. H. Miller (1958). The cost of capital, corporation finance and the theory of investment. *American Economic Review 48*(3), 261–297.

Modigliani, F. and M. H. Miller (1963). Corporate income taxes and the cost of capital: a correction. *American Economic Review 53*(3), 433–443.

Peters, R. H. and L. A. Taylor (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics 123*(2), 251–272.

Petty, J. W., D. F. Scott Jr, and M. M. Bird (1975). The capital expenditure decision-making process of large corporations. *The engineering economist 20*(3), 159–172.

Philippon, T. (2009). The bond market's q. *Quarterly Journal of Economics 124*(3), 1011–1056.

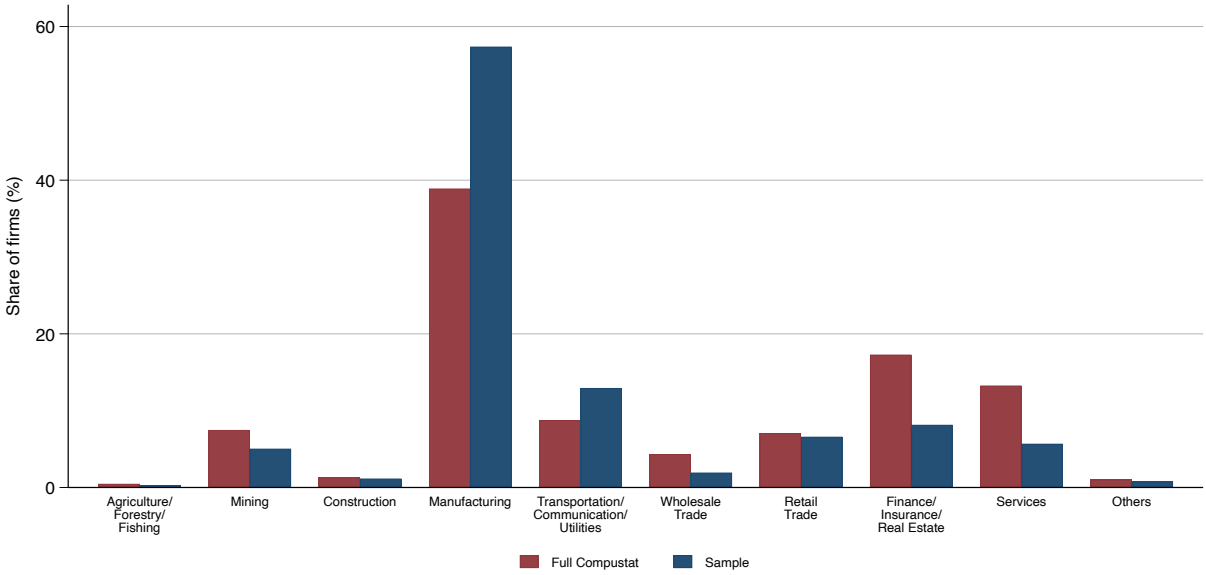
Rajan, R. G., P. Ramella, and L. Zingales (2023). What purpose do corporations purport? evidence from letters to shareholders.

Shue, K. (2013). Executive networks and firm policies: Evidence from the random assignment of MBA peers. *The Review of Financial Studies 26*(6), 1401–1442.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

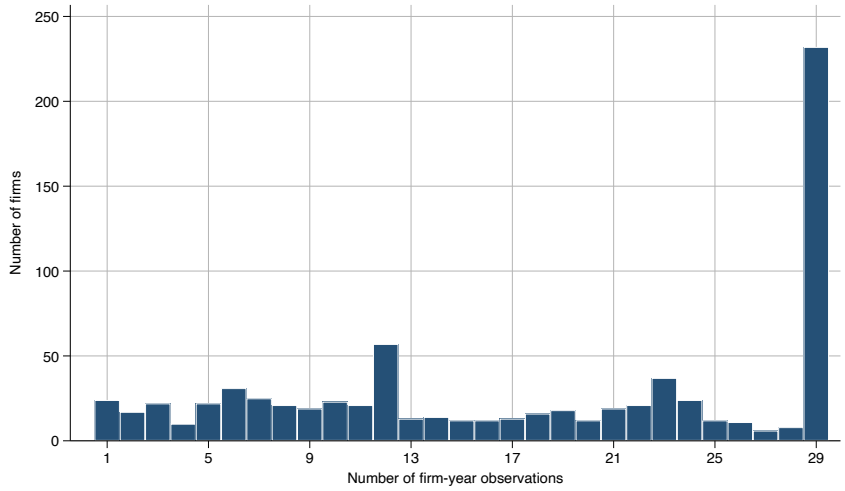
A Additional Figures and Tables

Figure A1: **Industry composition of sample compared to Compustat universe**



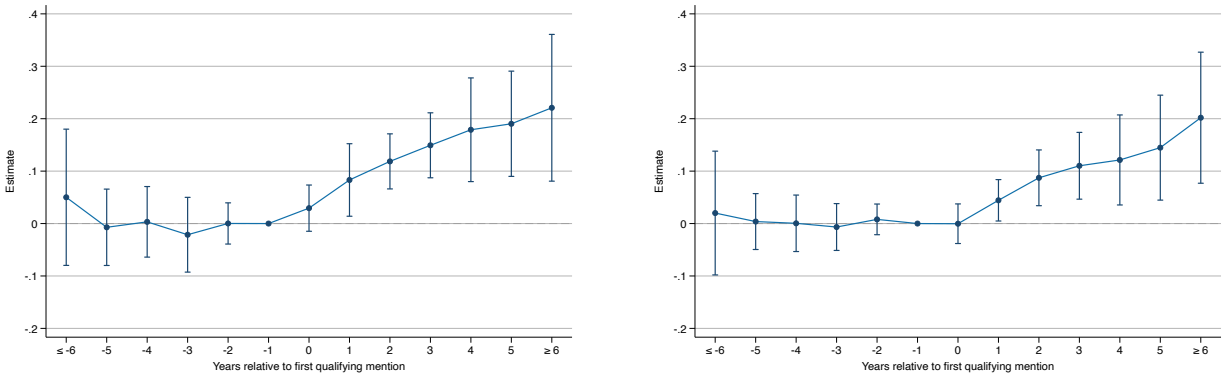
This figure shows the industry composition of firms in the sample compared to the full Compustat universe. We use the two-digit SIC codes to classify industries. The red bars report the share of firms in each industry in the full Compustat universe during our sample period. The blue bars report the corresponding shares in our sample.

Figure A2: **Histogram of observations per firm**



This figure shows the histogram of firm-year observations for firms in the sample. The sample period spans 1963–1991, and hence the max number of firm-year observations in the sample can be 29. Out of 772 firms covered by the data, 232 firms are covered in full during the sample period. The median number of firm-year observations in the sample is 19, and the average is 18.05.

Figure A3: **Tobin's q around adoption using only the "cost of capital" keyword**



(a) Baseline

(b) Full controls

The figure is identical to Figure 3, except the adoption indicator is based on whether firms mention the phrase "cost of capital," not on whether firms mention one of the three keywords.

Table A1: **Robustness Checks**

	(1) Adopters	(2) Balanced	(3) 20+ obs.	(4) ±2 years	(5) 5-year	(6) CoC-only
Adoption	0.131**	0.224***	0.183***	0.103*	0.059**	0.117**
	(0.048)	(0.078)	(0.062)	(0.050)	(0.028)	(0.051)
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Year × Industry FE	✓	✓	✓	✓	✓	✓
Year × B/M FE	✓	✓	✓	✓	✓	✓
Year × Size FE	✓	✓	✓	✓	✓	✓
Mean of Tobin's Q	1.122	1.224	1.216	1.106	1.177	1.177
Observations	3,803	4,031	8,143	3,075	10,704	11,957
R^2	0.787	0.793	0.786	0.792	0.788	0.780

This table reports additional results for Table 2. The dependent variable is Tobin's q in year $t+2$, where t denotes the observation year. Adoption is an indicator that equals 1 starting in the first year a cost of capital keyword is mentioned, except for columns (5) and (6). All specifications include firm, year, year-industry, year-B/M, and year-size fixed effects. Industries are defined using two-digit SIC codes. The year-B/M and year-size fixed effects are constructed by interacting year indicators with quintiles of firms' book-to-market and market capitalization in 1959 (or the closest observed year), respectively. Column (1) restricts the sample to firms that adopt by 1991. Column (2) requires a report and Tobin's q to be observed in every year from 1963 to 1991. Column (3) requires at least 20 firm-year observations during this period. Column (4) restricts the sample to firms with an observed report and Tobin's q from two years before through two years after adoption. Column (5) replaces the adoption indicator with an indicator that equals 1 only if a keyword is mentioned in the current year or any of the preceding four years. Column (6) replaces the adoption indicator with an indicator that equals 1 from the first year where the firm mentions "cost of capital." Tobin's q is winsorized at the 1st and 99th percentiles. Standard errors are two-way clustered by firm and year and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: **Robustness: alternative estimators**

Estimator	Estimate	Std. Error
Callaway and Sant'Anna (2021)	0.148***	(0.044)
Borusyak, Jaravel, and Spiess (2024)	0.157***	(0.046)
de Chaisemartin and D'Haultfœuille (2020)	0.094***	(0.025)
Sun and Abraham (2021)	0.118***	(0.031)

This table reports the coefficient on the adoption indicator estimated using four staggered adoption estimators (Callaway and Sant'Anna, 2021; Borusyak et al., 2024; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021). The specification contains the baseline controls. The sample is the 1963–1991 panel. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Data Appendix

This appendix documents the construction of the dataset.

B.1 Annual Report Collection

We identify firms listed on the S&P 500 in previous years using CRSP and firms appearing in the Fortune 500 list using
https://www.kaggle.com/datasets/lelandbartlett/fortune500-1955-2024.

Rajan, Ramella, and Zingales (2023) generously provided us with the annual reports of Fortune 500 firms. We additionally search for the annual reports of relevant firms from ProQuest, Mergent, and AnnualReports.com. We also check the Internet Archive (https://archive.org/search), the University of Washington Library (https://content.lib.washington.edu/reportsweb/index.html), and various corporate websites to collect reports not accessible elsewhere. In around 3% of cases, the annual report is mainly composed of a 10-K form. The data also include this type of report, although the impact on the findings is negligible.

B.2 Text Extraction

The text from the annual reports (PDF files) is extracted using a large language model (LLM).

Primary Pipeline. The primary pipeline for text extraction uses the `gemini-1.5-flash` LLM.

i) Divide the PDFs into smaller chunks of 8 pages or fewer.

ii) Submit the divided PDFs to the Gemini API.

iii) Prompt the LLM with the following prompt, extract the text of the output, and save it:

```
This file is [filename]. Give me all text in the file
literally, do not paraphrase, but correct any identifiable
typos or misspelled words. Omit any text related to the
copyright of the pdf file itself. Remove or add any
(un)necessary spaces, new lines, and punctuation where
appropriate. Return the result in plain text.
```

iv) Attempt extraction with the above prompt 3 times; if all attempts fail, assign the text file to the backup pipeline.

Backup Pipeline. The backup pipeline continues to use the `gemini-1.5-flash` LLM. For each file that fails to pass the primary pipeline, the procedure is as follows.

i) Further divide the PDF into even smaller chunks of 4 pages or fewer. Submit it to the LLM using the same prompt as above to ensure that the error is not caused by the constraint on output tokens.

ii) For files that fail again, extract their text using `pytesseract`.

iii) Submit the output of `pytesseract` to the LLM with the following prompt:

```
"Clean up the following OCR output produced by Pytesseract
from [filename]: [tesseract output]. Please correct any
spelling errors, typos, or grammatical mistakes. Do not
paraphrase or summarize."
```

iv) Store the output for files that are successfully extracted.

v) The output from Step iv) is combined with the output from the primary pipeline to constitute the final text dataset.

Machine-Readable Pipeline. Some files downloaded from AnnualReports.com are fully machine-readable. Accordingly, text for these files is extracted directly using the `pdfplumber` package from Python. The procedure is as follows:

i) Cut each file into chunks of 20 pages or fewer.

ii) Extract the text with `pdfplumber`. Save the output as a `txt` file.

B.3 Keyword Search

After extracting the text, we search for the three keywords: cost of capital, hurdle rate, and discount rate (Section 4). This section documents the search using cost of capital as the example. The same treatment of spacing and hyphenation applies to the other two terms. We search for the term "cost of capital" while allowing for (absence of) spaces or hyphens in between the words. For example, we search for

cost of capital, cost-of-capital, costofcapital, cost ofcapital, cost-of capital

and other variants.

We flag the `gvkey`-year pair with any matches. To ensure that we measure the "correct" cost of capital concept, we check terms within ten words and drop those with incorrect associations. For example, matches with the following keywords are dropped:

> *acquisition, asset, bond, consum, employed, equipment, expansion, expenditure, facilities, facility, good, impair, improvement, invested, investment, ized, lease, program, project, reimbursement, resource, share, stock*

This means that text containing "cost of capital acquisition" is not flagged.

After capturing the occurrences of "cost of capital" keywords in our text data, we construct an indicator for when "cost of capital" first appears in the report for each `gvkey`. The baseline adoption measure used in the paper is the union of the three terms (cost of capital, hurdle rate, and discount rate).

B.4 Accounting and Market Value Data

Compustat. We download the annual fundamentals from Compustat at the `gvkey`-fiscal year level and merge them with the text data using `gvkey` and year. Book equity is total assets minus total liabilities (`AT` $-$ `LT`), and the book-to-market ratio divides it by the CRSP market capitalization. Tobin's q is the sum of market capitalization, current liabilities (`DLC`), long-term debt (`DLTT`), preferred stock (`PSTK`), and deferred income taxes (`TXDITC`), divided by total assets (`AT`). The return on invested capital (ROIC) measure is constructed as operating income before depreciation (`OIBDP`) over the lagged sum of long-term debt (`DLTT`), debt in current liabilities (`DLC`), and stockholders' equity (`SEQ`) minus cash and short-term investments (`CHE`).

CRSP. We obtain the share price, shares outstanding, and market capitalization from CRSP, matched to each firm's fiscal year end (the closest CRSP monthly observation within 32 days of the Compustat `datadate`). Market capitalization enters the numerator of Tobin's q.

Winsorization. The Compustat levels and ratios and the CRSP variables are winsorized at the 1st and 99th percentiles over the full panel. The log transformations are applied to the winsorized values.