Firm Planning Horizon and Monetary Policy Transmission*

Marco Olivari[†] *Boston University*

Markus Schwedeler[‡] *NL Analytics*

Job Market Paper

October 29, 2026

Most recent version available here

Abstract

How do firms' planning horizons—defined as how far into the future they explicitly reason—shape investment and the transmission of monetary policy? We construct a firm-year measure of revealed planning horizons from corporate earnings calls by developing an algorithm that detects explicit future time references. The measure reveals large and persistent cross-firm heterogeneity, only partly explained by firm fundamentals and also associated with corporate governance. Using high-frequency monetary shocks at different points on the yield curve, we show that firms with longer planning horizons cut investment more in response to forward-guidance tightening but less to current-rate tightening. To rationalize these facts, we develop a New Keynesian model of investment with heterogeneous finite planning horizons and calibrate it to our measured distribution. The calibrated model successfully reproduces the observed differences in investment responses without targeting them, giving empirical content to theories of limited forward reasoning. Unlike benchmark models, where more distant forward guidance becomes steadily stronger or weaker, ours implies that forward guidance is most powerful when rate changes are one to two years ahead.

*We thank Masao Fukui, Adam Guren, and Tarek Hassan for their invaluable support and helpful comments. We also thank Marie Alder (discussant), Marta Bilghese, Roberto Corrao, Maryam Farboodi, Santiago Franco, Stefania Garetto, Laurence Kotlikoff, David Lagakos, Juan Cruz Lopez Del Valle, Fabian Seyrich (discussant), Taylor Watson, and seminar participants for helpful comments and discussions. Part of the research was conducted while Marco was visiting the European Central Bank. The views expressed here do not necessarily reflect those of the European Central Bank.

[†]olivari@bu.edu

[‡]markus@nlanalytics.tech

1 Introduction

How do changes in expected future economic conditions affect firms' investment? In standard macroeconomic models, firms value capital using the entire future stream of investment returns, incorporating expected future shocks to demand, costs, and interest rates into investment decisions today. It is possible, however, that firms evaluate these conditions over a much shorter horizon (Modigliani and Cohen, 1961; Graham, 2022). For example, corporate finance textbooks used in MBA courses teach managers to evaluate returns over a finite planning horizon and summarize later returns in a continuation value (Brealey et al., 2011). The planning horizon sets the reach of this explicit forward reasoning and therefore shapes how anticipated future shocks translate into investment today.

Firms' planning horizons are particularly relevant for the investment channel of monetary policy, as central banks have increasingly sought to stimulate economic activity by influencing expected future interest rates (Bernanke, 2020). Firms with different planning horizons may respond differently to policy news depending on when the announced rate changes occur. Indeed, theories of limited forward reasoning, including finite planning, limit the power of forward guidance (Woodford, 2019; Gabaix, 2020). Assessing the quantitative relevance of this mechanism requires evidence on how far ahead firm managers plan, yet we lack a systematic measure of firms' planning horizons and evidence on their role in investment responses to monetary policy.

In this paper, we construct a new firm-year measure of *revealed planning horizons* from dates and durations mentioned in U.S. corporate earnings-call transcripts to study how firms' planning horizons shape the investment channel of monetary policy. The measure seeks to capture the reach of managers' explicit forward reasoning through the future economic conditions they discuss when explaining their plans to investors. The measure reveals substantial and persistent cross-firm heterogeneity. Firm fundamentals, such as asset and liability structures, explain only a limited share of this variation, while management, board, and ownership characteristics together explain at least as much.

We document two main empirical findings. First, investment by firms with longer revealed planning horizons responds significantly more to forward-guidance shocks. Second, investment by firms with shorter revealed planning horizons responds more to current-rate shocks. To rationalize both findings, we develop a New Keynesian model of investment in which firms have different finite planning horizons. Disciplined by the empirical distribution of revealed planning horizons, the model successfully reproduces these untargeted investment responses and delivers three aggregate implications. First, our measured horizons support the quantitative relevance of the limited forward reasoning emphasized by Woodford (2019) and Gabaix (2020): for tightening five years ahead, the investment contraction at announcement is less than one-fifth as large as under infinite planning. Second, the measured distribution implies that forward guidance has its strongest announcement-date effect on aggregate investment when the announced rate changes are one to two years ahead. This contrasts both with the increasing potency of more distant news under infinite planning and with its declining potency in benchmark models of limited forward reasoning. Third,

the observed lengthening of firms' planning horizons over the past two decades has strengthened the Federal Reserve's ability to influence current investment through announcements of more distant rate changes.

In the first part of the paper, we develop a transparent, rule-based algorithm to construct our measure of revealed planning horizons from earnings calls. We analyze over 200,000 earnings-call transcripts from approximately 5,500 listed U.S. firms over more than two decades, following a growing literature that uses these discussions to recover firm-level information (Hassan et al., 2019, 2025). The algorithm identifies temporal references—such as "over the next five years" or "by fiscal 2028"—and translates them into temporal distances from the call date in years. To capture the reach of managers' planning as reflected in these discussions, we retrieve the upper tail of the distribution of these temporal references within each call.

Revealed planning horizons capture a distinct dimension of firm heterogeneity. Firms' average horizons vary widely, from 1.7 years at the 25th percentile to 3.8 years at the 75th, with some exceeding a decade. Planning horizons tend to be longer in industries characterized by long-lived investments, such as utilities and real estate, but large differences remain within narrowly defined industries. These differences persist over time, with firm fixed effects accounting for a substantial share of the variation. Consistent with economic intuition, firms with longer horizons tend to be larger and older and to have longer-lived assets and a larger share of long-term debt. These firm fundamentals explain only a limited share of the differences in planning horizons. Corporate governance accounts for at least part of the residual variation: longer horizons are associated with greater ownership by long-term institutional investors, greater long-term CEO compensation shares, larger boards, and higher shares of female directors.

In the second part of the paper, we estimate differential investment responses of firms with short versus long measured horizons to high-frequency monetary shocks affecting different parts of the yield curve. As a baseline, we use Jarociński's (2024) current-rate and forward-guidance surprises. Current-rate surprises mainly move current and near-term interest rates, with smaller effects at longer maturities. Forward-guidance surprises move expected future interest rates while leaving near-term rates largely unchanged. We estimate local projections of cumulative capital growth, comparing firms in the same sector and quarter that face the same announcement but differ in their average horizons.

Firms with longer revealed planning horizons contract investment more after contractionary forward-guidance surprises but less after contractionary current-rate surprises. One year after the announcement, with each surprise scaled to a 25-basis-point rate increase, a firm one standard deviation higher in our horizon measure experiences approximately 0.66 percentage points more cumulative capital decline after forward guidance and 0.25 points less after current-rate news. These differences are economically meaningful, equivalent to approximately 4.8 and 1.8 percent of annual investment, respectively. The differences in investment responses across firms with different revealed planning horizons are robust to allowing size, age, liquidity, leverage, asset maturity, and debt maturity to affect responses to both shocks. These findings suggest that planning horizons capture

an additional dimension of heterogeneity in investment responses to monetary policy, complementing Drechsel et al. (2026), who highlight heterogeneity beyond standard firm characteristics.

In the final part of the paper, we rationalize these investment patterns and quantify their aggregate implications in a New Keynesian model by adapting Woodford's (2019) finite-planning framework for households to firms' investment decisions. In our model, firms differ only in their planning horizons. At each date, a firm explicitly evaluates returns within its planning window and values capital beyond it using a steady-state continuation value, mirroring the structure of capital budgeting taught in MBA textbooks. We derive a finite-horizon Q-theory linking current investment to the shadow value of capital. Interest rates affect investment directly through discounting and indirectly through their effects on returns of capital. Planning horizons determine how far ahead firms anticipate each effect.

Disciplined by revealed planning horizons and estimated interest rate paths, the model successfully reproduces the contrasting investment responses despite being untargeted. Following a forward-guidance shock, longer-horizon firms incorporate future rate increases that lie beyond shorter-horizon firms' planning windows. The direct effect arises because these rate increases reduce the discounted value of investment returns, while the indirect effect reflects the anticipated contraction in aggregate demand, which further depresses returns of capital. Longer-horizon firms therefore contract investment more. Following a current-rate shock, the current increase enters both short- and long-horizon firms' planning windows. The weaker contraction of longer-horizon firms comes from the indirect effect: returns of capital initially fall as aggregate demand contracts, then recover as demand rebounds while capital remains scarce. Longer-horizon firms anticipate more of this recovery, which supports their perceived value of capital and moderates their investment contraction. Quantitatively, the predicted differences in investment responses are similar in magnitude to those observed empirically for both shocks.

The model delivers three aggregate implications for monetary transmission. First, with the horizon distribution disciplined by our data, it reproduces the theoretical prediction of Woodford (2019) and Gabaix (2020): limited forward reasoning substantially weakens the effects of distant forward guidance. For example, tightening five years ahead produces an investment contraction at announcement that is 83 percent smaller than under infinite-horizon planning. This substantial attenuation arises with planning horizons measured from earnings calls, giving the key theoretical prediction of this literature quantitative support from firm-level evidence.

Second, forward guidance has its strongest effect on aggregate investment at announcement when the rate increase is one to two years ahead. A more distant rate increase means a longer period of low returns of capital, so firms that anticipate it cut investment more today. At the same time, a more distant increase falls outside more firms' planning horizons, so fewer firms respond to the announcement. For near-term increases, larger cuts outweigh the decline in the number of responding firms; for more distant increases, the reverse holds. The resulting hump-shaped pattern is consistent with evidence that near-term policy news can have larger real effects than immediate surprises, while more distant guidance need not be more powerful (Goodhead and Kolb, 2025). In contrast, infinite

planning and benchmark models of limited forward reasoning predict monotonic effects in opposite directions. Under infinite planning, all firms incorporate even distant rate changes, so moving the increase further into the future strengthens the investment contraction without reducing the share of firms responding. In benchmark models of limited forward reasoning, the attenuation of distant news dominates throughout, making more distant announcements progressively less powerful.

Third, we find that measured planning horizons have lengthened over time: the median rises by more than one-third between 2002 and 2025. In this sense, firms in our sample have become more forward-looking. We use the model to compare monetary transmission under the 2002 and 2025 horizon distributions. This lengthening makes the investment contraction at announcement 73 percent larger for tightening five years ahead, but 6 percent smaller for tightening one year ahead. These results suggest that the increasing effectiveness of forward guidance emphasized by Bernanke (2020) and Goodhead (2024) can also reflect the lengthening of firms' planning horizons. Longer planning horizons complement the communication mechanism studied by Campbell et al. (2019) by making distant policy news relevant to current investment decisions as communication makes that news more informative.

Related Literature. This paper contributes to four strands of literature. First, we complement survey evidence on firms' reported planning and forecast horizons (Poterba and Summers, 1995; Graham and Harvey, 2001; Altig et al., 2021; Graham, 2022) by measuring the horizons managers discuss with investors in earnings calls. Our approach measures how far into the future managers discuss their firms' plans, allowing us to connect this dimension of corporate decision-making to observed investment behavior and theories of finite planning. By extracting the time dimension of corporate language, we also contribute to the growing literature that uses corporate text to recover firm-level information (Gentzkow et al., 2019; Hassan et al., 2019, 2024, 2025; Caldara et al., 2020; Caldara and Iacoviello, 2022; Jamilov et al., 2023; Sautner et al., 2023). The algorithm can be combined with other word searches to measure the horizons associated with particular topics. Unlike Appel and Herpfer (2026), who independently and contemporaneously construct a related measure using machine learning and large language models, we rely on simple, deterministic rules. These rules make each extracted horizon traceable to its source phrase and conversion and ensure reproducible results. Our measure captures the upper reach of firms' discussions of the future, rather than the duration of disclosed business plans (Dessaint et al., 2023) or the relative emphasis on short- and long-term language (Brochet et al., 2015).[1]

Second, we document how investment responses to current-rate and forward-guidance surprises vary with firms' planning horizons. This complements evidence on heterogeneity associated with firm and financial characteristics (Gertler and Gilchrist, 1994; Ottonello and Winberry, 2020; Darmouni

[1] Related measures capture implied discount rates (Sampson and Shi, 2023), reported discount rates and perceived costs of capital (Gormsen and Huber, 2025, 2026), capital durability or useful lives (Souder and Bromiley, 2012; Souder et al., 2016), production duration (Antràs and Tubdenov, 2025), realized time to build (Fernandes and Rigato, 2024), or executive temporal focus and depth (Nadkarni and Chen, 2014; Nadkarni et al., 2016; Weis and Klarner, 2022). These objects may correlate with planning reach but can also reflect discounting, technology, project maturity, or the direction of managerial attention.

et al., 2022; Deng and Fang, 2022; Gürkaynak et al., 2022; Cloyne et al., 2023; Altavilla et al., 2024; Anderson and Cesa-Bianchi, 2024; Jungherr et al., 2024; Alder et al., 2025; Caglio et al., 2026; Jeenas, 2026), uncertainty (Kroner, 2021), and attention (Song and Stern, 2025). Yet, as Drechsel et al. (2026) show, much of the variation in investment responses to monetary policy remains unexplained by observed firm characteristics. Our findings also complement Selgrad and Siani (2026), who focus primarily on how monetary policy affects firms' short- and long-term investment plans, interpreting their evidence through implementation frictions, including time to build and partial irreversibility. We study realized investment responses to monetary policy across firms with different horizons of explicit forward reasoning. We also complement work linking policy-induced market valuation changes to investment (Jeenas and Lagos, 2024; de Fraisse et al., 2026). Finite planning makes investment depend on a horizon-specific shadow value that need not coincide with market-based Q, making the mapping from market valuations to investment potentially sensitive to the timing of policy news.

Third, we contribute to the literature on forward guidance by showing that its effects on investment depend jointly on the timing of the announced rate changes and firms' planning horizons. We complement empirical work on the maturity (Gurkaynak et al., 2005; Adams and Barrett, 2025) and information content (Nakamura and Steinsson, 2018; Bauer and Swanson, 2023a; Jarociński, 2024) of monetary surprises by relating firms' investment responses to their planning horizons. Our model connects this firm-level evidence to the literature on the aggregate effects of forward guidance (Del Negro et al., 2023; McKay et al., 2016). Whereas the investment response grows with the distance of policy news under infinite planning, the measured distribution of firms' horizons generates the strongest announcement-date response at intermediate maturities.[2]

Finally, we extend the finite-planning framework of Woodford (2019) to firms' investment decisions. We complement its estimation by Gust et al. (2022) and its applications (Xie, 2020; Woodford and Xie, 2022; Dupraz et al., 2024; Gust and Lopez-Salido, 2025; Lustenhouwer and Mavromatis, 2025; Na and Xie, 2025) by disciplining firms' planning windows with measured horizons. The relevant heterogeneity concerns how far ahead firms reason, rather than impatience (Ameriks et al., 2003) or economic lifetimes (Lorenzoni, 2019). Our finite-horizon Q-theory shows how planning windows govern both the direct effect through discounting and the indirect effect through anticipated changes in returns of capital, rationalizing the contrasting investment responses to current-rate news and forward guidance.

Importantly, capital accumulation also helps distinguish mechanisms proposed to address the forward-guidance puzzle (Del Negro et al., 2023). Limited common knowledge (Angeletos and Lian, 2018), reflective equilibrium (García-Schmidt and Woodford, 2019), and bounded depth of reasoning (Farhi and Werning, 2019) attenuate higher-order feedback, whereas cognitive discounting (Gabaix, 2020) and finite planning attenuate across calendar time. See Gabaix (2019) for a review of the foundations of behavioral inattention. Although these mechanisms can imply similar attenuation

[2]Our focus on firm investment complements evidence on financial-market and household responses to monetary communication (Lewis et al., 2020; Swanson, 2021; Ahn et al., 2022; D'Acunto et al., 2022; Coibion et al., 2023; Sutherland, 2023).

of general equilibrium feedback (Angeletos and Lian, 2023), they have different implications for investment in our model. Returns to capital initially fall as demand contracts, then recover as capital becomes scarce. Attenuating higher-order feedback weakens both phases, whereas discounting or truncating distant dates places less weight on the recovery and can therefore amplify the investment contraction. Our measured horizons discipline this truncation, linking aggregate monetary transmission to differences in investment responses across firms.

Roadmap. The paper is organized as follows. Section 2 constructs the planning-horizon measure. Section 3 documents its properties. Section 4 presents the firm-level evidence on monetary transmission. Section 5 develops the model. Section 6 presents its quantitative implications for monetary transmission. Section 7 concludes.

2 Measuring Firm Planning Horizons

This section develops and validates our measure of planning horizons. We first present the conceptual framework linking managers' forward-looking discussions to firms' planning horizons. We then describe the earnings-call data and the algorithm used to construct the measure. Finally, we assess its validity by relating measured horizons to other horizon proxies used in the literature and comparing the algorithm's classifications with those of a large language model.

2.1 Conceptual Framework

A firm's *planning horizon* marks the furthest future date for which managers explicitly assess economic conditions when evaluating investment. Beyond this horizon, managers summarize subsequent returns through a continuation value rather than explicitly tracing the effects of shocks at each later date. Corporate finance textbooks organize investment valuation around this endpoint (Brealey et al., 2011). Writing H for the firm's planning horizon, measured in periods from today, t, the investment's value is

$$V_t = \sum_{h=1}^{H} \frac{CF_{t+h}}{(1+r)^h} + \frac{TV_{t+H}}{(1+r)^H}, \tag{1}$$

where CF_{t+h} denotes forecast investment returns, r is a constant discount rate,[3] and TV_{t+H} is the continuation value at $t+H$ of all subsequent returns. If the continuation value correctly accounts for all state-contingent returns beyond H, this finite forecast gives the same valuation as the infinite-horizon formulation of standard Q-theory (Hayashi, 1982).

Managers may, however, use simpler assumptions to value returns beyond their explicit forecast horizon. As Graham (2022) documents, many firms use simple rules to evaluate investments and

[3] Gormsen and Huber (2025) and Gormsen and Huber (2026) use earnings-call language to measure firms' discount rates and perceived costs of capital. Fukui, Gormsen, and Huber (2025) build on these data to study the implications of sticky nominal discount rates for monetary transmission. These studies concern discounting, represented by r in equation (1), while we focus on the planning horizon H.

regard their forecasts as reliable only over short horizons. For example, a widely used approach to estimating the continuation value is the Gordon–Shapiro model, which projects cash flows beyond the forecast horizon at a constant growth rate (Brealey et al., 2011; Damodaran, 2012). More generally, simplified projections of returns beyond the explicit forecast horizon can omit temporary changes in those returns, making investment responses depend on how far ahead managers plan.

The planning horizon therefore determines how far ahead firms explicitly incorporate a shock's effects into their investment valuation. Consider two otherwise identical firms with three- and ten-year planning horizons that use constant-growth continuation values. Both learn of a temporary increase in year-five cash flow. The gain enters the ten-year firm's forecast but leaves the three-year firm's continuation value unchanged because its year-four cash flow, long-run growth rate, and discount rate do not change. The same news can therefore lead to different investment decisions.

In earnings calls, managers sometimes explicitly describe distant developments as too far away to incorporate into current forecasts. Asked why anticipated AI-driven consumption growth had not raised its 2027–28 forecasts, Unisys replied in February 2026 that "it's too far out for us to adjust, multiple year out consumption estimates." Similarly, Calfrac described its rolling forecasting window before explaining that its planning could not indicate outcomes farther ahead because "it is too far away from us from today's point of view." These statements corroborate the idea that distance can limit firms' responses to anticipated developments. Appendix Table A.4 presents these passages and further examples.

The planning horizon is distinct from time to build, asset life, and depreciation rates. For example, consider a ski resort investing in a new chairlift.[4] Construction may take only a few months, while managers explicitly forecast returns over the next few years. Yet the lift's lifetime can span decades. This long service life does involve maintenance costs and does not by itself imply a low depreciation rate. This example shows why these characteristics do not directly measure the planning horizon, even if they are correlated with it.

Our empirical objective is therefore to measure the planning horizon H in equation (1). In doing so, we follow Graham's (2022) call to ground corporate finance models in evidence on how firms make decisions.

2.2 Measurement

Planning horizons are difficult to observe directly. Yet managers' discussions of future plans and expectations provide evidence that we can use to infer their underlying planning horizons.

Measurement Approach. We construct a firm's *revealed planning horizon* from explicit future references its managers mention in earnings calls. Earnings calls are used by managers to explain their investment decisions in terms of expected demand, costs, and financing conditions. A manager

[4] A public example is the 2006 master plan for the Erciyes ski-resort development in Türkiye. Its appraisal combines an implied initial depreciation rate of 10%, an 8% discount rate, asset lives of up to 30 years, and a 15-year forecast horizon. The projections rely largely on assumed growth rates and consider alternative scenarios, such as weaker demand revealed in year 4. Erciyes master plan, pp. 103–109.

may explain that the firm needs more capacity to meet the demand it expects over the next five years. We take the dates and durations in such explanations as indications of how far ahead managers explicitly reason. Our approach builds on research that uses earnings-call discussions to study firms' decisions (Hassan et al., 2025; Gormsen and Huber, 2025; Gormsen and Huber, 2026). In our baseline, the firm-year revealed planning horizon is the 95th percentile of the number of years ahead referred to in each sentence, pooled across the firm's calls during that year. This statistic summarizes the more distant part of managers' discussions, which we use to infer their planning horizon H. Table 1 provides illustrative examples of sentences in our sample containing explicit future dates and durations.

Repeated scrutiny gives managers the incentive to communicate credibly. Managers can choose strategically which plans to disclose, so the horizon which is implied by their statements may differ from the one they use internally. However, analysts and investors can compare earlier statements with subsequent investment and performance, limiting managers' incentives to misrepresent their plans (Hassan et al., 2025; Gormsen and Huber, 2025).

We interpret a longer revealed horizon as evidence that a firm plans farther ahead. A single mention of a distant date might reflect a comment unrelated to the firm's investment plans rather than being informative about the underlying planning. For this reason, we use the 95th percentile rather than the most distant reference captured by our algorithm. We also exclude the question-and-answer session because the horizons discussed there may reflect analysts' interests rather than managers' plans. Differences in communication style could still lead some managers to discuss more distant dates without planning farther ahead. We therefore compare revealed horizons with independently observed firm characteristics and managerial incentives in Section 2.3, looking for evidence consistent with the planning interpretation.

Data from Earnings Calls. We obtain earnings-call transcripts from Refinitiv Eikon. Our baseline sample comprises 212,132 earnings calls from 5,583 U.S.-headquartered firms over 2002–2025, with available Refinitiv economic-sector classifications. The full transcripts contain 80.7 million sentences. For the baseline measure, we focus only on the management presentation section before the Q&A with analysts.

Text-Extraction Algorithm. We develop a transparent, rule-based algorithm to extract future dates and durations from corporate earnings calls. These references can be classified into three main groups. First, *calendar-anchored references* name a date or calendar period, such as "October 20, 2027," "Q4 2027," or "early 2028." Second, *speech-relative references* express timing relative to the call date, as in "coming quarter" or "the next 18 months." Third, *bare durations* state a length of time without naming a date, as in "a 3-year investment program." The rules cover days, weeks, months, quarters, years, and decades, as well as calendar qualifiers such as "early," "mid," and "late." Expressions such as "in the future" and "long term" provide neither a specific date nor a measurable duration, so we exclude them. Appendix A describes extraction and conversion in more detail, and the Online Methodological Note provides the technical documentation.

Table 1: Examples of Explicit Future References in Earnings Calls

Firm and call date	Statement
Farmland Partners 5 Aug 2021	It's hard to predict out beyond **2 or 3 years**, but I think we will have that land appreciation trend with us for quite some time.
Bowlero 16 Nov 2022	We buy back opportunistically, and we acquire every center or sign every lease that meets our hurdle rate. Hurdle rate being about 25% compounded over **5 years** with a terminal value factored in.
EQT 24 Jul 2014	Along these lines, neither OVC nor MVP is explicitly assumed to occur though we did assume a 4% terminal growth rate in distributable cash flow **after 2019**.
Heidrick & Struggles 28 Jul 2015	You know, we don't forecast beyond the **next quarter** revenue because the visibility in this business is about as good as a cloudy day at Willis Tower here in Chicago.
Information Services Group 15 Mar 2019	We made this investment in anticipation of the stronger client demand for digital that we expect to see over the **next 24 months**.
Jabil 17 Jun 2015	We made this decision based on optimizing our construction costs and modeling our needs for additional capacity, capacity based on our growth projections for the **next two to three years**.
Monolithic Power Systems 25 Jul 2018	In anticipation of higher revenue in the **next 2 to 3 years**, we are stepping up investment in our foundry capabilities.
Teleflex 1 Nov 2018	We are also increasing investment in R&D to support our expected acceleration in growth from new products over the **2019 to 2021** time period.
Coterra Energy 23 Feb 2023	Based on the current service cost environment, we estimate that if the company invests $2.0 billion to $2.1 billion per year over the **next 3 years**, it will generate a compound annual growth rate of 0% to 5% for both BOE and natural gas and closer to 5% for oil.
GenOn Energy 29 Feb 2012	We expect that industry retirements, including those already announced, will result in market prices rising sufficiently to justify investments of between approximately $586 million and $726 million over the **next 10 years**.
Pool Corporation 16 Oct 2014	Over **a 5- to 7-year time horizon**, we expect that our base business expense growth will approximate one-half our gross profits growth, with year-to-year variations based on our being more focused on return on invested capital than on the relationship to gross profits in a particular year.
Backblaze 15 Feb 2023	While our computer backup business remains the larger of our 2 cloud service offerings at the moment, our strategy and increasing investments center around capitalizing on the approximately $100 billion total **2025** market opportunity for B2 cloud storage based on projections from IDC and company analysis.
Inovio Pharmaceuticals 11 May 2015	Overall, based on management's projections and analysis, the Company believes that cash, cash equivalents and short-term investments are sufficient to meet its planned working capital requirements, including the cost of its planned Phase 3 clinical trial of VGX-3100 through the **end of 2018**.
Vonage 6 Aug 2020	Based on the performance and predictability of the consumer segment, we project it will produce in excess of $600 million of after-tax equity free cash flow over the **next 5 years** with significant terminal value after that.

Notes: Statements are from Refinitiv Eikon earnings-call transcripts, with explicit future references highlighted in bold.

We convert the extracted dates and durations into years. For a date or calendar period, we measure the time from the call date to that date or period; for a fixed duration, we use its stated length. In the baseline specification, we assign calendar periods and duration ranges their midpoints. We exclude references classified as backward-looking and count overlapping matches for the same expression only once. We then average the values assigned to the retained future references within each sentence, so each sentence containing at least one such reference contributes one observation.

We use explicit rules rather than a black box machine learning approach to make the mea-

surement transparent. In fact, each reference can be traced to the rule that identifies it and the calculation that converts it into years. The procedure is also modular, allowing us to combine its sentence-level observations with separate measures of topics, risks, or expectations and study how far ahead managers discuss these subjects.

Our baseline measure is defined at the firm-year level. We pool sentence-level observations across all calls in a year and take their 95th percentile. We average these annual horizons across all available years to obtain the firm-level measure. This percentile is also less sensitive than the maximum to an isolated distant reference. We calculate horizons at alternative percentiles to assess how the choice of percentile affects our results.

2.3 Validation of the Measure

We validate the economic interpretation of the measured revealed planning horizons using independently measured firm fundamentals and corporate governance characteristics commonly associated with longer-term orientation. We also compare the algorithm's sentence classifications with those of a large language model to assess agreement on whether sentences contain future references.

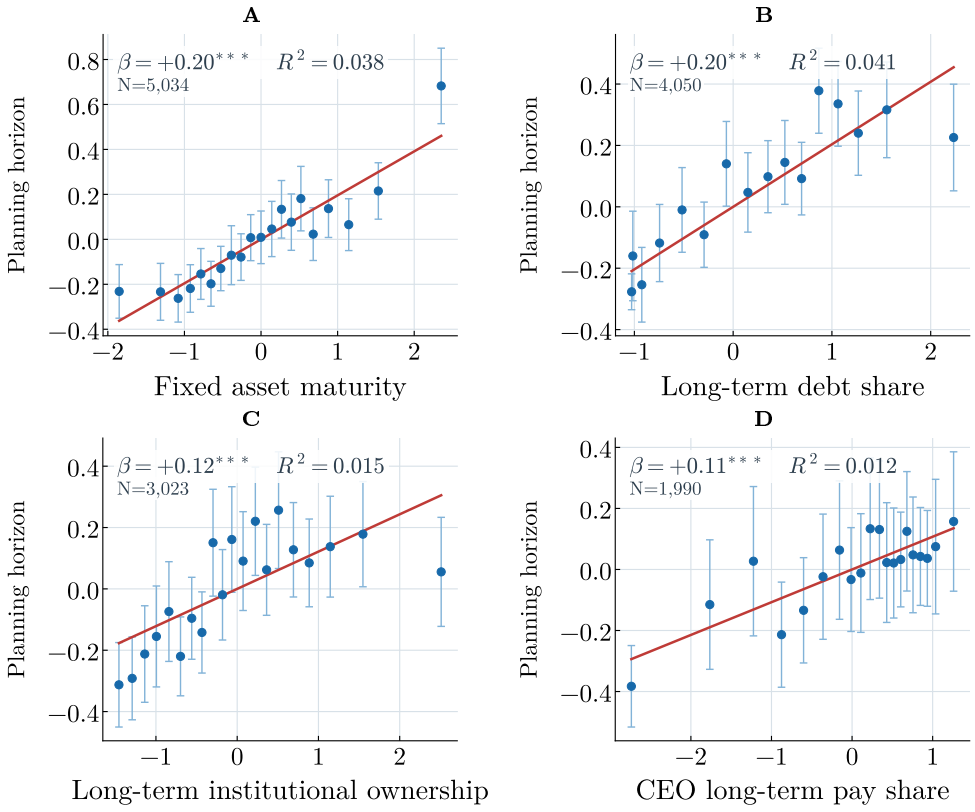
Validation against external benchmarks. Previous research has used investment, ownership, compensation, and financing characteristics to study firms' orientation toward the future. We compare our revealed planning horizons with these characteristics to validate the economic interpretation of our measure. For example, Gaspar et al. (2005) and Derrien et al. (2013) use portfolio turnover to measure shareholders' investment horizons, while Souder and Bromiley (2012) and Souder et al. (2016) use asset durability to measure firms' investment horizons. Research links the structure of executive compensation to managerial short-termism and investment horizons (Bizjak et al., 1993; Bolton et al., 2006; Souder and Shaver, 2010; Edmans et al., 2017). Models of long-term finance link financing commitment and maturity to firms' investment horizons (von Thadden, 1995; Milbradt and Oehmke, 2015).

Figure 1 shows that our revealed planning horizons are positively and significantly associated with all four characteristics: fixed asset maturity, the long-term debt share, long-term institutional ownership, and the share of CEO compensation outside salary and bonus.[5] These patterns support interpreting our earnings-call measure as informative about how far ahead firms plan. Yet the regressions have low R^2: each characteristic explains only a small share of the variation in revealed horizons. Most of the variation in our measure is therefore unaccounted for by these individual proxies.

Comparison with a large language model. We compare the algorithm's sentence-level classifications with those of Gemini 2.5 Flash-Lite for a sample of 900 sentences. Among sentences flagged by the algorithm, 98.9 percent are also flagged by Gemini; the algorithm identifies 56.3 percent of the sentences flagged by Gemini. Inspecting the disagreements shows that Gemini sometimes

[5]We thank Amir Amel-Zadeh, Fiona Kasperk, and Martin Schmalz for sharing their corporate ownership data.

Figure 1: Validation against proxies for firm horizons



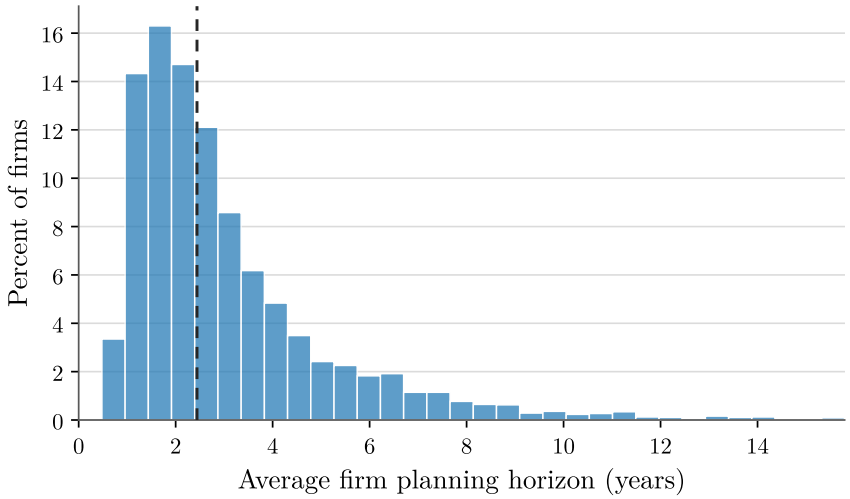
Notes: The panels plot average firm planning horizons against four proxies for firm horizons. Dots show mean horizons and characteristics in up to 20 quantile bins of the characteristic; bars show 95% confidence intervals for mean horizons. Red lines show pooled ordinary least squares fits using firm-level observations. Panels report the slope, R^2, and number of firms. Each panel uses the largest available sample of firms with data on revealed planning horizons and the plotted characteristic. Planning horizon is the average of call-level horizons over 2002–2025. Characteristics are averaged across available annual observations over 2002–2025; institutional ownership covers 2004–2020. Panel A uses fixed asset maturity, net property, plant, and equipment divided by depreciation excluding amortization, from Compustat following Hubert de Fraisse (2024). Panel B uses the Compustat share of total debt due after five years. Panel C uses the share of outstanding equity held by institutions with three-year portfolio turnover, averaged over four quarters, of at most 35%, following Derrien, Kecskés, and Thesmar (2013), using Corporate Ownership Data from Amel-Zadeh, Kasperk, and Schmalz (2022). Panel D uses CEO compensation outside salary and bonus divided by total compensation, from ExecuComp. Fixed asset maturity enters in logs; the three shares enter in levels. Within each panel, horizons and characteristics are winsorized at the 1st and 99th percentiles and standardized to mean zero and unit standard deviation. Significance stars use HC1 heteroskedasticity-robust standard errors: $*$ $p < 0.10$, $**$ $p < 0.05$, and $***$ $p < 0.01$. Appendix C.2 provides further variable definitions and construction details.

assigns horizons to broader forward-looking language, such as "going forward" or "the near future," when the sentence specifies no date or duration. We prefer the rule-based algorithm because it allows us to control which expressions enter the measure and how they are converted into distances. Appendix A.3 provides the prompt and detailed results.

3 Heterogeneity in Revealed Planning Horizons

We use our measure to study differences in firms' planning horizons. We first describe the distribution across firms and compare horizons across and within industries. We then show that persistent

Figure 2: Distribution of firm planning horizons



Notes: The histogram shows the distribution of average firm planning horizons, measured in years. Bars show the percentage of firms in each horizon bin, with each firm weighted equally; the dashed line marks the sample median. Planning horizon is the average of call-level horizons over 2002–2025.

firm differences account for a substantial share of the variation. Finally, we examine how these differences relate to firm fundamentals, managerial compensation, board composition, and ownership.

3.1 Differences across Firms and Sectors

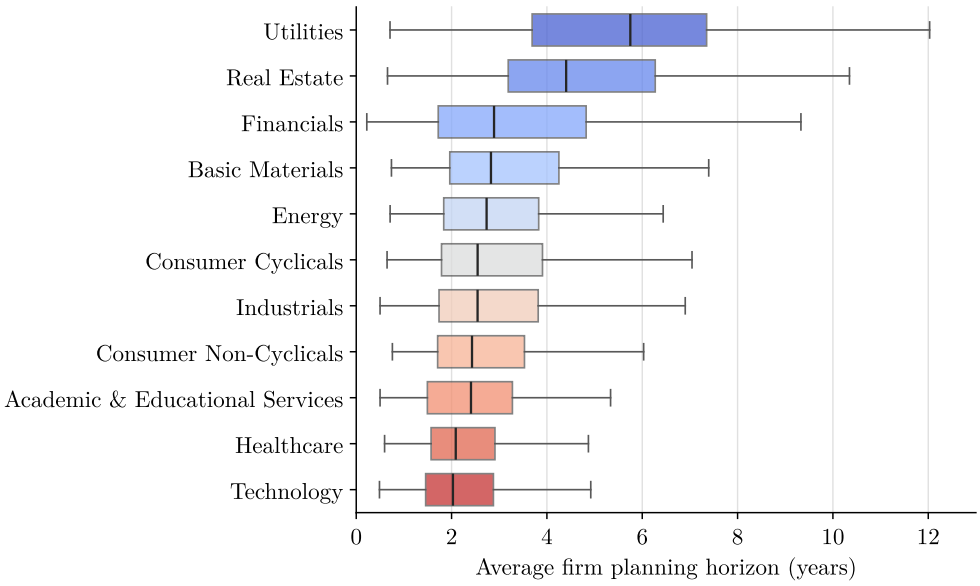
Investment decisions have consequences for years ahead, making differences in how far ahead firms plan potentially important for their investment choices. Figure 2 shows the distribution of firms' average revealed planning horizons, H_i. The median is 2.4 years, and the middle half of firms have horizons between 1.7 and 3.8 years. The horizon at the upper end of this range is more than twice as long as that at the lower end. Some firms plan much further ahead, giving the distribution a long right tail and raising the mean horizon to 3.3 years. Survey evidence also documents multiyear horizons and substantial variation across firms. The survey underlying Graham (2022) reports a mean reliable planning horizon of 2.3 years. Shrader et al. (2004) find strategic planning horizons averaging 2.7 years and ranging from 1 to 11 years among manufacturing firms with formal plans.[6]

Figure 3 shows substantial differences across sectors. Utilities and real estate have median horizons of 5.7 and 4.4 years, respectively, compared with 2.0 years in technology and 2.1 years in healthcare. Selgrad and Siani (2026) also find the longest investment-plan horizons in utilities and real estate. The longer horizons in utilities and real estate are consistent with their long-lived infrastructure and property investments and their multiyear regulatory and contractual commitments.

Large differences also appear among firms in the same four-digit SIC industry. Marriott and

[6] Additional survey evidence points to multiyear horizons in corporate strategy. The National Association of Corporate Directors (2025) reports that 70 percent of responding directors' boards have strategic planning horizons of three to five years. PwC (2015) reports that 42 percent of surveyed public-company directors evaluate strategy over one to three years, while 58 percent use horizons of five years or longer.

Figure 3: Distribution of firm planning horizons across sectors



Notes: The box plots show the distribution of average revealed planning horizons, measured in years, across sectors. Boxes span the 25th to 75th percentiles of firm horizons within each sector; the line inside each box marks the median. The horizontal lines extending from each box reach the most extreme observations within 1.5 interquartile ranges of the box edges; outlier points are not shown. Planning horizon is the average of call-level horizons over 2002–2025. Firms are assigned to their most frequently observed sector over this period. The 11 sectors each contain at least 30 firms and are ordered by median horizon.

Hilton, both hotel operators, have average revealed planning horizons of 2.5 and 8.4 years, respectively. Among telecommunications companies, the corresponding horizons are 2.6 years for Verizon and 6.2 years for AT&T. In home-improvement retail, Lowe's and Home Depot have horizons of 5.6 and 7.8 years, respectively.[7] Together, the sector patterns and company comparisons suggest that sector affiliation matters, but substantial differences in planning horizons also exist among firms within the same industry. These differences among firms with similar activities call for a closer examination. We next investigate how persistent they are and how much firm fundamentals and corporate governance can explain.

3.2 Heterogeneity and Persistence in Planning Horizons

Persistent differences in planning horizons could help explain systematic differences in firms' investment behavior. We therefore ask whether firms that plan further ahead tend to continue doing so over time. We use the firm-year panel to follow the same firms over time. A state-space model and Kalman filtering separate annual variation in horizons into persistent sector and firm differences, changes within firms, and measurement error. We find substantial persistent differences among firms in the same sector.

[7] Appendix Table B.2 reports additional illustrative comparisons.

Table 2: Sources of variation in firm planning horizons

			Between firms		Within firm	%
Specification	$\widehat{\rho}$	Time	Sector	Firm	variation	explained
Baseline	0.66	1.4	14.3	48.8	35.5	47.3
Time-varying sector effects	0.66	1.4	14.1	49.1	35.3	47.3
Business sectors	0.64	1.3	14.8	48.5	35.3	49.9

Notes: The table reports variance decompositions of annual firm planning horizons on the same sample of 3,759 firms and 36,315 firm-years over 2003–2025. The baseline uses 11 TRBC economic sectors, year and firm effects, and a persistent within-firm component. The second specification adds sector-year states using the same economic-sector classification; their contribution is included in the sector column. The third specification uses 30 TRBC business sectors. Persistence varies across sectors; $\widehat{\rho}$ is the firm-weighted posterior mean of the sector-specific persistence parameters. Time, sector, firm, and within-firm entries are percentages of variation net of measurement error. Sector and firm components include the contribution of finite-sample averages of the time-varying state. The final column reports the percentage of observed annual variation attributable to the underlying horizon after separating measurement error. The sample requires at least four calls per firm-year and three usable years per firm. Missing sector classifications enter an Other group, excluded from the sector counts shown.

We model the annual horizon $H_{i,t}$, the average of a firm's call-level horizons in year t, as

$$H_{i,t} = \lambda_t + \gamma_{s(i)} + \alpha_i + h_{i,t} + \varepsilon_{i,t}, \qquad h_{i,t} = \rho_{s(i)} h_{i,t-1} + \eta_{i,t}.$$

Here, λ_t captures common year effects, $\gamma_{s(i)}$ persistent sector differences, and α_i persistent firm differences within sectors. The state $h_{i,t}$ captures persistent movements in a firm's planning horizon, while $\varepsilon_{i,t}$ represents measurement error in the annual average. We allow the persistence parameter $\rho_{s(i)}$ to vary across sectors. The model treats calls within a firm-year as repeated measurements of a common annual horizon. Under a standard independent-error specification, call-level errors are independent across calls and years and of the serially uncorrelated state innovations. Differences across calls therefore help quantify measurement noise in the annual average, while persistence across years helps identify these persistent movements. We estimate the model using Bayesian methods; Appendix B.1 describes the estimation and variance decomposition.

Table 2 reports the resulting variance decomposition. The model attributes approximately half of observed variation to measurement error. This implies a signal-to-noise ratio close to one, similar to that reported by Hassan et al. (2019) for political risk. In the baseline specification, the firm component accounts for approximately half of variation net of measurement error, capturing differences among firms in the same sector over the observed period. Sector differences account for a further approximately 15 percent. Together, these components account for more than 60 percent of variation net of measurement error, with most cross-firm heterogeneity arising among firms in the same sector. Changes within firms also persist: the firm-weighted mean of sector-specific persistence is 0.66, so horizon shifts tend to last beyond the year in which they occur, although their persistence differs across sectors.

The finding that differences among firms within sectors account for the largest share of variation

net of measurement error is robust. The last two rows of Table 2 allow sector conditions to change over time or group firms into 30 business sectors rather than 11 economic sectors. Both specifications give sector effects greater scope to explain the observed heterogeneity, yet the firm component is still the largest. Appendix Table B.1 further shows how much observed firm characteristics can explain and whether finer industry classifications alter this conclusion. Adding firm characteristics explains only a small share of variation, leaving a large residual firm component. The firm component also exceeds the sector component under every SIC classification considered, including four-digit industries. The appendix separately reports diagnostic estimates of CEO effects as a possible source of firm differences. The characteristics and sector comparisons therefore show that substantial differences in firms' planning horizons extend beyond sector affiliation and observable economic characteristics.

3.3 Average Planning Horizons across Firms

Having established persistent differences in planning horizons across firms, we ask how firms with longer average horizons differ in their economic characteristics and corporate governance. We now focus on the cross-section, using one observation per firm and averaging horizons and characteristics over time. The regressions compare firms within the same sector. Figure 4 reports a separate regression for each characteristic.

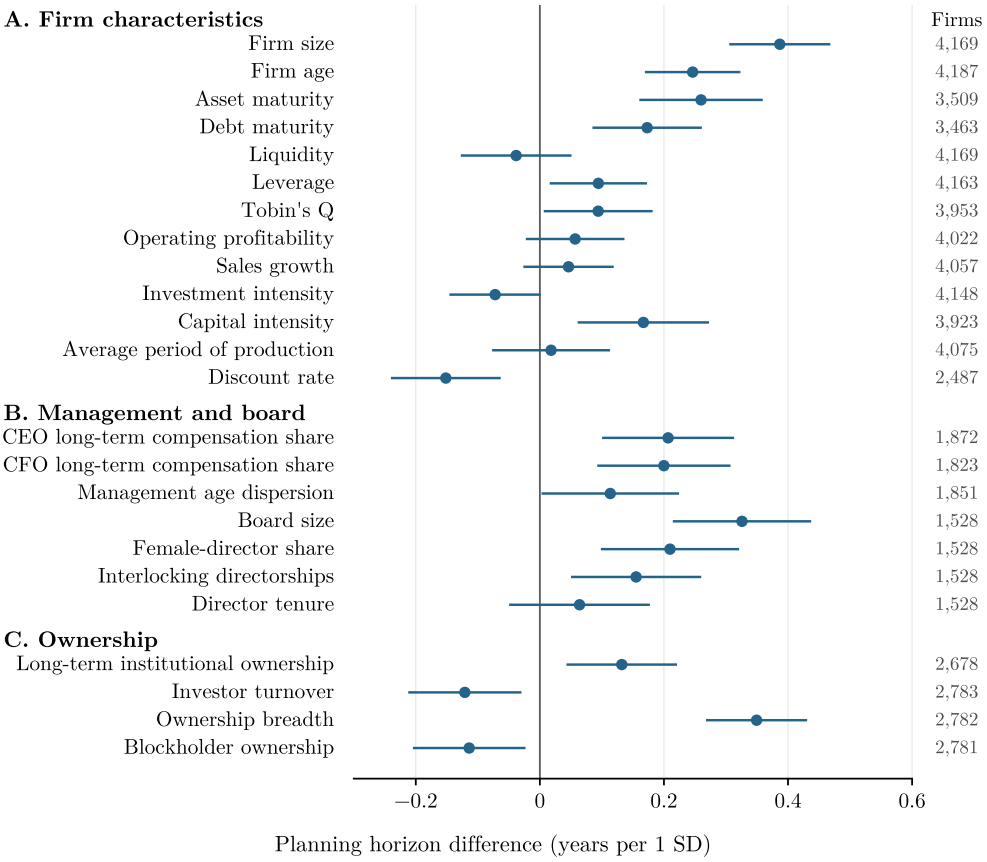
The relationships with firm characteristics are broadly consistent with economic intuition. Longer horizons are associated with larger, older firms and with longer-lived assets and longer-term debt. The positive associations with capital intensity, leverage, and Tobin's Q are also intuitively plausible. Several less precisely estimated relationships are compatible with this broad picture: profitability, sales growth, and the average period of production have positive point estimates, although their confidence intervals include zero. Higher predicted discount rates from Gormsen and Huber (2025)[8] are associated with shorter planning horizons, although the two play distinct roles in valuation (Section 2.1).

Despite their intuitive individual associations, standard firm characteristics have limited collective explanatory power for planning horizons. In Appendix Table B.3, column (1), we regress firms' average planning horizons on a large number of standard firm characteristics jointly, including sector fixed effects. These characteristics explain approximately 8 percent of the variation in average planning horizons among firms within the same sector.

Corporate governance helps explain part of this remaining variation. In Appendix Table B.3, column (2), we include management, board, and ownership characteristics in the preceding regression and explain approximately 17 percent of within-sector variation in average planning horizons. We allocate this explanatory power across the four groups using a Shapley value decomposition, which averages each group's additional contribution over different orders of inclusion. In this decomposition, standard firm characteristics account for approximately 5 percent of within-sector variation, while management, board, and ownership characteristics together account for approximately 12 per-

[8]We thank Niels Joachim Gormsen and Kilian Huber for sharing their discount-rate data.

Figure 4: Average planning horizon and firm characteristics



Notes: The figure reports associations between average revealed planning horizons and firm characteristics (Panel A), management and board characteristics (Panel B), and ownership characteristics (Panel C). Each point is the coefficient from a separate ordinary least squares regression with one characteristic and Refinitiv business-sector fixed effects, using that characteristic's available sample. Characteristics are standardized; coefficients measure differences in planning horizons in years per one-standard-deviation increase in the characteristic. Bars show 95% confidence intervals based on HC1 heteroskedasticity-robust standard errors. Each regression uses firms with revealed planning horizons and characteristic data over 2004–2020; firm counts are reported by row. Planning horizon is the average of call-level horizons over 2004–2020, measured in years. Characteristics are averaged across available annual observations in years with observed horizons over the same period, which matches the ownership-data coverage. Data come from Compustat, ExecuComp, ISS/RiskMetrics Directors, Amel-Zadeh, Kasperk, and Schmalz (2022), and Gormsen and Huber (2025). Horizons and characteristics are winsorized at the firm-level 1st and 99th percentiles. Appendix C.2 provides further variable definitions and construction details.

cent. The small common sample warrants caution about the precise shares, but the results suggest that how firms are managed and owned provides information about planning horizons beyond their economic characteristics.

Management and board characteristics also connect firms' planning horizons to how managers are paid and overseen. Longer measured horizons are associated with greater long-term compensation shares for both CEOs and CFOs. Using ISS/RiskMetrics Directors data, we also find that longer horizons are associated with larger boards, higher female-director shares, and more interlocking directorships with other public companies.

Firms' planning horizons are also related to their shareholders' investment horizons and ownership structure. Longer horizons are associated with greater long-term institutional ownership and lower investor turnover, suggesting that firms' planning horizons align with their shareholders' holding periods. Ownership breadth is positively associated with planning horizons, whereas blockholder ownership is negatively associated with them. These links between ownership and firms' attention to the future are the focus of a separate paper, in which we study these relationships in greater detail (Olivari and Schwedeler, 2026).

4 Planning Horizon, Investment and Monetary Transmission

Section 3 shows that firms differ systematically in their planning horizons. We now show that the relationship between planning horizons and monetary transmission depends on the type of tightening. Relative to firms with shorter horizons, firms with longer horizons cut investment more after forward-guidance tightening but less after current-rate tightening.

A monetary policy announcement can change current interest rates or expectations of rates further ahead. Changes in either can affect investment through expected financing costs and returns, but firms may differ in how far ahead they explicitly assess these changes. News about future rates may fall beyond a short planning horizon while lying within a longer one. This distinction is particularly relevant for forward guidance, through which central banks seek to influence current decisions by changing expectations of future interest rates (Bernanke, 2020). We therefore compare the investment responses of firms with long and short planning horizons to current-rate and forward-guidance news.

4.1 Data

We combine U.S. firms' revealed planning horizons with quarterly financial data from Compustat and high-frequency monetary policy shocks. We study shocks dated from 1991Q1 to 2019Q1 and follow firms' investment responses over subsequent quarters. At the one-year response horizon, the baseline estimation sample contains 4,160 firms and 196,997 firm-quarter observations. Appendix C reports the sample restrictions and descriptive statistics.

Our baseline monetary shocks come from Jarociński (2024), who decomposes changes in fed funds futures, Treasury yields, and stock prices within 30-minute windows around FOMC announcements.[9] We focus on the current-rate and forward-guidance components. The current-rate shock mainly moves current and near-term rates, with smaller effects at longer maturities. The forward-guidance shock changes expected future rates while leaving near-term rates largely unchanged, capturing Odyssean guidance, or news about commitments to future policy rates (Campbell et al., 2012).[10] The decomposition purges these policy components of the Fed information effect by identifying

[9] The decomposition assumes a sufficient degree of independence across shocks and exploits their tendency to generate occasional large market movements, which help reveal their distinct effects across assets.

[10] Delphic guidance, by contrast, conveys forecasts of the economy and the associated policy path.

information about the economic outlook as a separate shock. We sum the announcement surprises within each quarter.

The two shocks let us compare investment responses to news about current and future policy rates. Changes in the current rate fall within both short and long planning horizons, whereas changes in future rates may lie beyond a shorter horizon. Firms with longer horizons can therefore assess more of the future changes explicitly when deciding how much to invest today. We test whether the relationship between revealed planning horizons and investment responses differs according to when the announced rate changes take effect.

We average each firm's annual revealed planning horizons, constructed from earnings calls as described in Section 2, to capture persistent differences in how far ahead firms plan. Section 3 supports this choice by showing that persistent differences between firms in the same sector account for the largest share of variation after accounting for measurement error. Averaging over time also reduces the influence of noise in individual horizon measurements.

We measure investment responses through cumulative growth in firms' real capital stocks. The outcome is the log change in capital from the quarter before the shock to each response quarter, multiplied by 100. This captures the accumulated effect of investment after allowing for depreciation. We construct capital using the perpetual-inventory method of Stein and Stone (2013), adding quarterly capital expenditure to the existing stock after depreciation and adjusting for changes in capital-goods prices. Appendix C.5 provides the construction details.

4.2 Empirical Strategy

Our empirical approach builds on Ottonello and Winberry (2020), combining quarterly Compustat data with high-frequency monetary policy shocks. We closely follow their choice of fixed effects and controls. We compare the investment responses of firms in the same sector and quarter that face the same monetary policy news but differ in their planning horizons. Using panel local projections (Jordà, 2005), we estimate a separate regression for each response quarter $h = 0, \ldots, 12$. Our general empirical specification is

$$\begin{aligned} \Delta_h \log K_{it} = \alpha_i^{(h)} + \gamma_{j(i),t}^{(h)} + \beta_{FG}^{(h)}(\widetilde{H}_i FG_t) + \beta_{CR}^{(h)}(\widetilde{H}_i CR_t) \\ + \eta_{FG}^{(h)\prime}(Z_{it} FG_t) + \eta_{CR}^{(h)\prime}(Z_{it} CR_t) \\ + \rho^{(h)}(\widetilde{H}_i g_{t-1}) + \theta^{(h)\prime} X_{it} + \varepsilon_{it}^{(h)}. \end{aligned} \tag{2}$$

Here, $\Delta_h \log K_{it}$ denotes firm i's cumulative real-capital growth from quarter $t-1$ to quarter $t+h$, expressed in log percentage points. The variable $\widetilde{H}_i$ is the logarithm of firm i's revealed planning horizon, winsorized at 0.5 percent in each tail and then standardized. The variables CR_t and FG_t denote the current-rate and forward-guidance shocks, respectively. Both horizon–shock interactions enter the baseline regression, so the coefficient on each measures how investment responses vary with planning horizons, conditional on the interaction with the other shock.

Firm fixed effects, $\alpha_i^{(h)}$, account for persistent differences in capital growth. Sector-by-quarter

fixed effects, $\gamma^{(h)}_{j(i),t}$, absorb conditions shared by firms in the same sector and quarter, including the common response to monetary policy news.[11] The baseline controls X_{it} include lagged log real assets, lagged real-sales growth, the lagged share of current assets in total assets, and an indicator for the final quarter of the firm's fiscal year. We also interact planning horizons with lagged GDP growth, g_{t-1}, to account for differences in firms' investment sensitivity to the business cycle.

One concern is that revealed planning horizons proxy for other firm characteristics that shape investment responses to monetary policy. Interacting these characteristics, collected in Z_{it}, with the monetary shocks controls for differences in responses associated with them. Given the broad range of potentially relevant characteristics, we use a baseline whose fixed effects and controls closely follow Ottonello and Winberry (2020), excluding these additional shock interactions. We examine the additional interactions in the robustness checks in Section 4.4, adding one characteristic at a time and estimating a separate regression for each monetary shock.

Positive values of both shock series indicate tightening. We scale the forward-guidance shock to raise the quarterly-average two-year Treasury yield by 25 basis points on impact, and the current-rate shock to raise the quarterly-average federal funds rate by the same amount. Each coefficient, $\beta^{(h)}_{FG}$ or $\beta^{(h)}_{CR}$, measures how the cumulative capital-growth response to the corresponding tightening differs with a one-standard-deviation increase in log planning horizon. A negative coefficient indicates a lower response among longer-horizon firms, while a positive coefficient indicates a higher response. Standard errors are clustered by firm and quarter.

4.3 Investment Responses by Planning Horizon

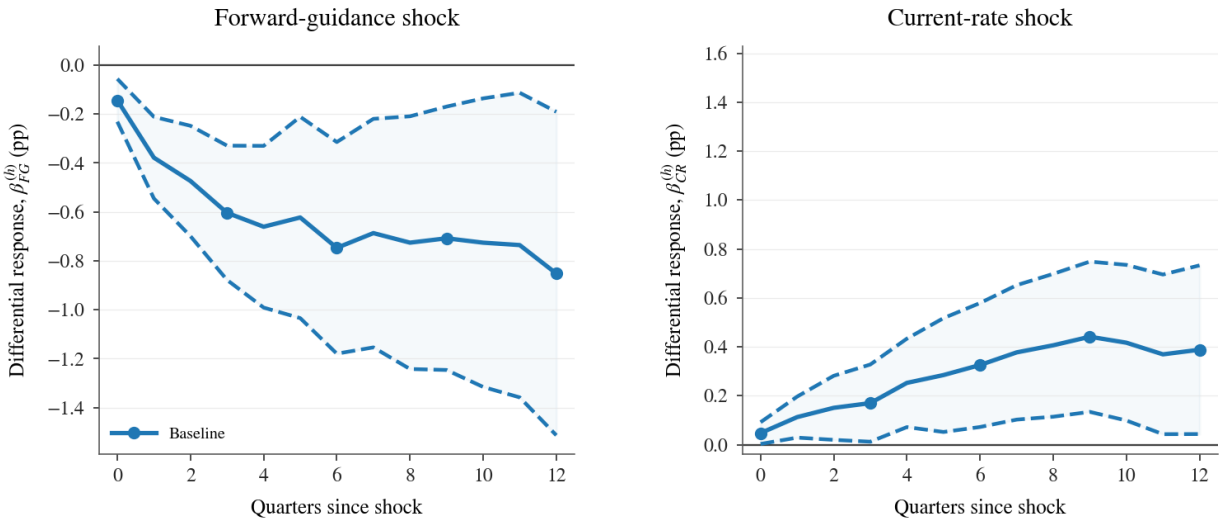
Figure 5 shows that longer planning horizons are associated with lower capital-growth responses to forward-guidance tightening and higher responses to current-rate tightening. These differences are already present in the announcement quarter and persist throughout the three-year response window. Most of the forward-guidance gap develops during the first year, whereas the current-rate gap grows more gradually and is largest around two years after the announcement. Estimates become less precise at later response quarters, but the pointwise 90 percent confidence intervals exclude zero throughout for both shocks.

These differences are sizeable relative to firms' annual investment. One year after the announcement, a one-standard-deviation higher log planning horizon is associated with a cumulative capital-growth response that is 0.66 percentage points lower following forward-guidance tightening and 0.25 percentage points higher following current-rate tightening. Each shock is scaled to raise its corresponding interest rate by 25 basis points on impact. Using the sample median annualized investment-to-capital ratio, these differences correspond to approximately 4.8 and 1.8 percent of annual investment, respectively.[12]

[11] Following Ottonello and Winberry (2020), our baseline specification excludes financial firms and utilities and uses eight broad SIC sectors: agriculture, forestry, and fishing; mining; construction; manufacturing; transportation, communications, electric, gas, and sanitary services; wholesale trade; retail trade; and services. Appendix D assesses robustness to alternative sector classifications.

[12] We express the capital-growth differences relative to the sample median annualized investment-to-capital ratio of 13.9 percent.

Figure 5: Investment responses by planning horizon



Notes: The panels plot $\beta_{FG}^{(h)}$ (left) and $\beta_{CR}^{(h)}$ (right), estimated jointly from equation (2), at each response quarter. The estimates show how much the cumulative real-capital-growth response differs, in percentage points, for firms with a one-standard-deviation higher log planning horizon. For each shock, the impact increase in its corresponding interest rate is normalized to 25 basis points. The shading marks pointwise 90 percent confidence intervals calculated using standard errors clustered by firm and quarter.

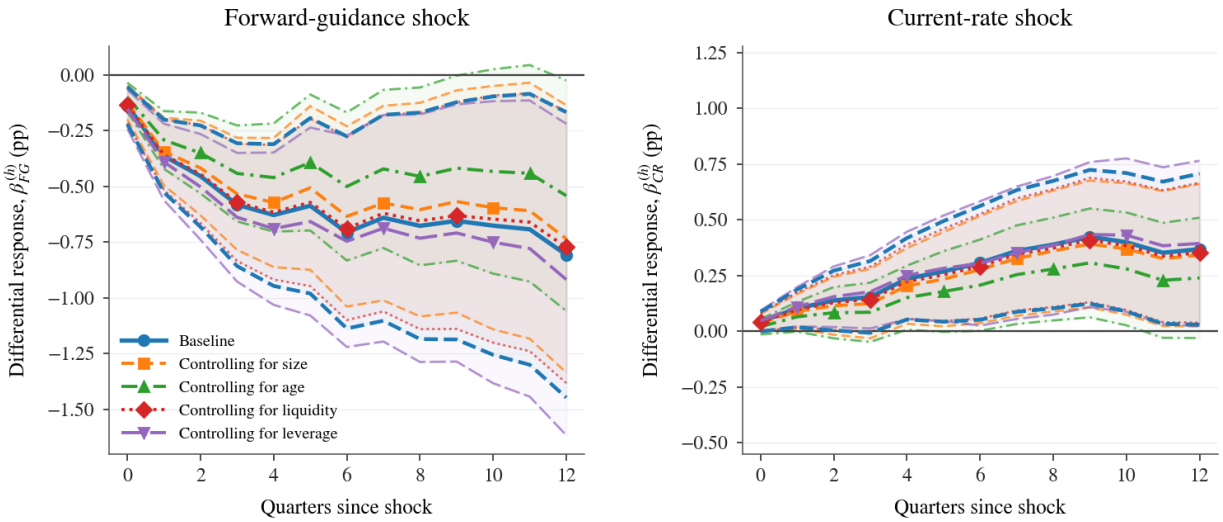
The forward-guidance finding aligns with the intuition that firms looking further ahead consider more distant changes in interest rates when making investment decisions. This intuition alone cannot explain the current-rate finding, since current rates fall within both short and long planning horizons. Section 5 develops a model in which longer-horizon firms anticipate more of the recovery in returns to capital following an investment contraction, moderating their response to current-rate tightening. The association with revealed horizons therefore depends on whether policy news concerns current or future rates, rather than reflecting uniformly greater sensitivity to monetary tightening.

4.4 Robustness

Other firm characteristics may account for the response differences associated with planning horizons. We therefore allow investment responses to vary with these characteristics and check whether the horizon-related differences persist. We also repeat the analysis with alternative horizon measures and test whether the contrast extends to yield surprises at different maturities. Appendix D describes additional robustness checks.

Firm demographics and financial position. The response differences could arise because firms with longer planning horizons also differ in characteristics already associated with monetary transmission: size (Gertler and Gilchrist, 1994), age (Cloyne et al., 2023), liquidity (Jeenas, 2026), and leverage (Ottonello and Winberry, 2020). We therefore allow investment responses to vary with

Figure 6: Investment responses controlling for interactions with size, age, liquidity, and leverage



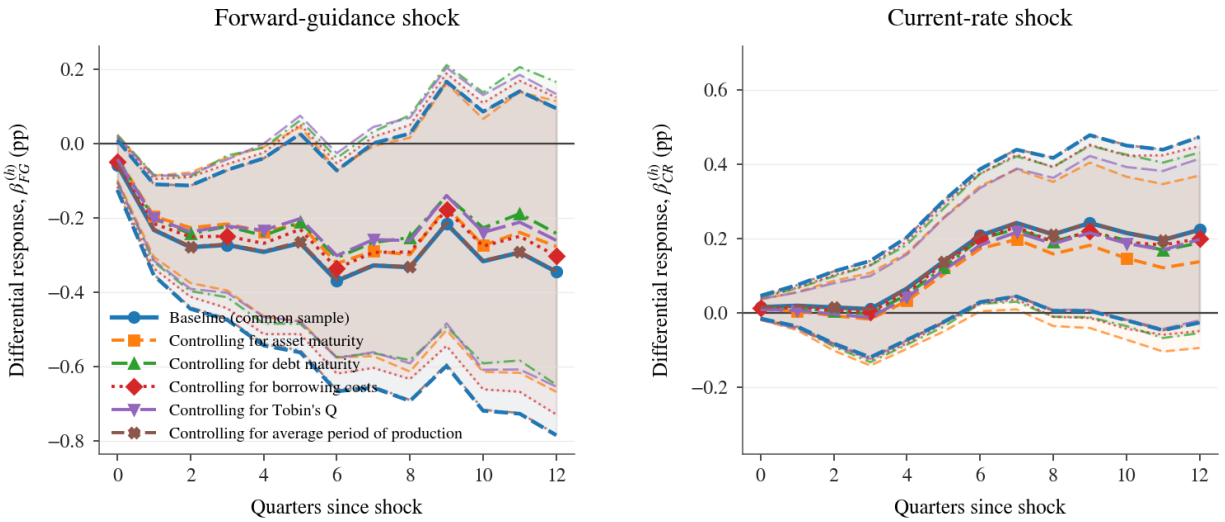
Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Size, age, liquidity, or leverage is added in levels and interacted with the shock and lagged GDP growth. Samples vary with control availability. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

each characteristic separately, estimating each monetary shock in a separate regression. Figure 6 shows that the main pattern holds after accounting for each characteristic: longer-horizon firms cut investment more after forward-guidance tightening and less after current-rate tightening. Both one-year response differences remain statistically significant after accounting for each characteristic. The horizon-related response differences persist after allowing investment responses to vary with size, age, liquidity, or leverage, considered one at a time.

Investment and financing characteristics. Firms may plan further ahead because their assets last longer, their debt matures later, or their production takes more time. These characteristics could also shape their investment responses to monetary policy. Figure 7 therefore considers asset maturity, debt maturity, borrowing costs, Tobin's Q, and production period, adding each separately on a common sample. Debt maturity is the share of debt due after one year.[13] The one-year point estimates reproduce the baseline pattern under every control. Sample coverage accounts for a substantial part of the smaller estimates: before adding any of these controls, restricting the sample reduces the forward-guidance difference from -0.63 to -0.29 percentage points and the current-rate difference from 0.24 to 0.07. With the additional controls, the forward-guidance estimates range from -0.29 to -0.23 and are all statistically significant at one year. The current-rate estimates range from 0.03 to 0.07 and are imprecise, including before adding the controls. The signs are

[13] This measure provides broader coverage than measures requiring debt repayments at longer maturities. Appendix Figure D.2 reports robustness checks using debt-maturity cutoffs from two to five years.

Figure 7: Investment responses controlling for interactions with asset and debt maturity, borrowing costs, Tobin's Q, and production period

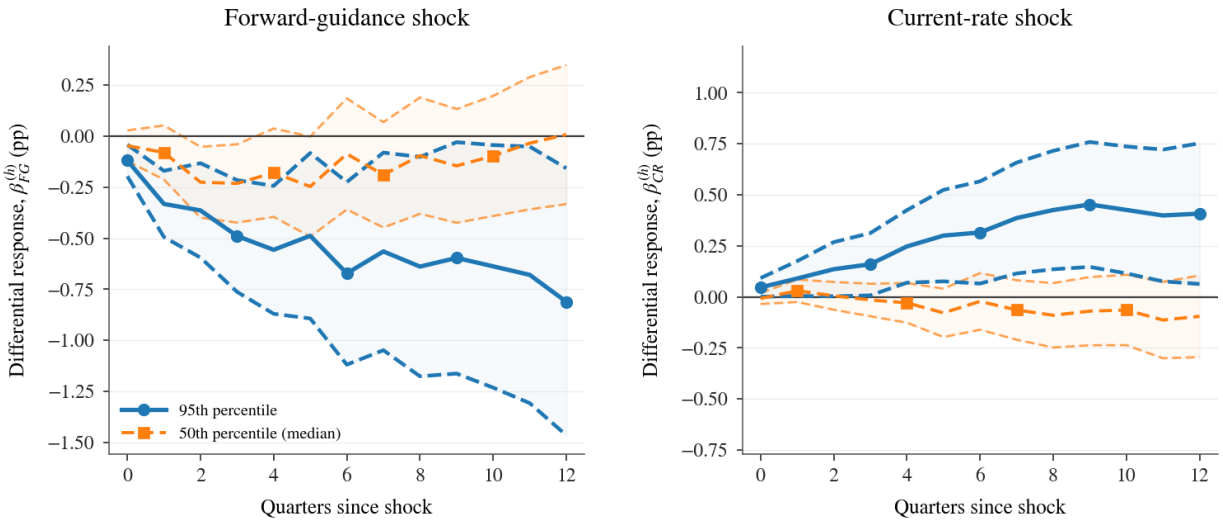


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Asset maturity, debt maturity, borrowing costs, Tobin's Q, and production period enter one at a time in levels and interacted with the shock and lagged GDP growth. Debt maturity is the share of debt due after one year. All lines, including the baseline, use the same observations within each panel and response quarter. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

preserved, but restricting the sample changes the estimates more than adding these controls. The current-rate difference is also imprecisely estimated on this common sample. Appendix D provides further asset and liability comparisons, including debt-maturity cutoffs from two to five years.

Upper-tail versus median horizons. The 95th percentile measures the more distant part of a firm's forward-looking discussions, while the median captures the typical time span. We include both measures jointly in Figure 8 to ask whether more distant references provide information about investment responses beyond the typical time span of a firm's discussions. We estimate separate regressions for each shock, using standardized logs of both horizon measures and interacting each with lagged GDP growth. The remaining controls and fixed effects follow the baseline. Holding the median fixed, the 95th-percentile horizon predicts a larger investment contraction after forward-guidance tightening and a smaller contraction after current-rate tightening, with both relationships statistically significant throughout the response window. The median predicts some additional forward-guidance responses early on, but this relationship subsequently fades; its current-rate estimates are imprecise throughout. These results support our choice of the 95th percentile, which captures information about investment responses beyond that provided by the median. Appendix D compares alternative upper-tail percentiles.

Figure 8: Investment responses by upper-tail and median planning horizons

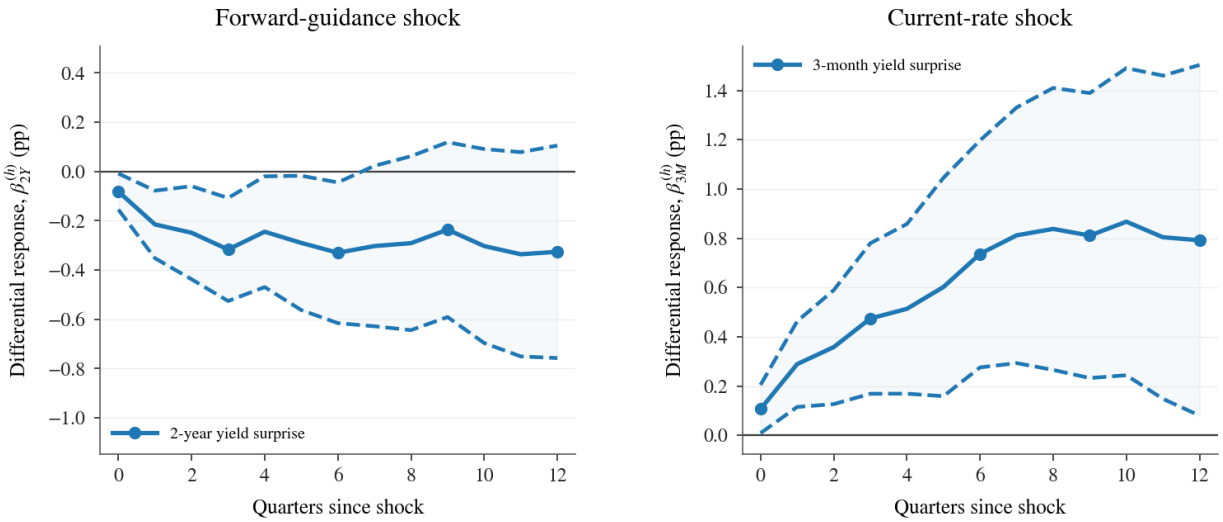


Notes: The panels plot interaction coefficients for the 95th-percentile and median planning horizons with forward-guidance (left) and current-rate (right) shocks. Each panel estimates $\Delta_h \log K_{it} = \alpha_i^{(h)} + \gamma_{j(i),t}^{(h)} + \beta_{s,95}^{(h)} \widetilde{H}_i^{95} s_t + \beta_{s,50}^{(h)} \widetilde{H}_i^{50} s_t + \rho_{95}^{(h)} \widetilde{H}_i^{95} g_{t-1} + \rho_{50}^{(h)} \widetilde{H}_i^{50} g_{t-1} + \theta^{(h)\prime} X_{it} + \varepsilon_{it}^{(h)}$, separately for $s = FG$ and $s = CR$. The variables $\widetilde{H}_i^{95}$ and $\widetilde{H}_i^{50}$ are the standardized logs of the firm's 95th-percentile and median horizon measures. The controls X_{it} and lagged GDP growth g_{t-1} are defined as in equation (2). Each coefficient gives the percentage-point difference in the cumulative real-capital-growth response associated with a one-standard-deviation higher log horizon, holding the other horizon measure fixed. For each shock, the impact increase in its corresponding interest rate is normalized to 25 basis points. The shading marks pointwise 90 percent confidence intervals calculated using standard errors clustered by firm and quarter.

Maturity of policy news. We also use yield surprises to assess whether the relationship between planning horizons and investment responses varies with maturity. This comparison uses three-month and two-year yields directly, without the baseline decomposition into current-rate and forward-guidance shocks. In Figure 9, longer-horizon firms have lower capital-growth responses to an increase in the two-year yield and higher responses to an increase in the three-month yield, matching the signs of the forward-guidance and current-rate estimates, respectively. The two-year difference is statistically significant through six quarters after the announcement and less precisely estimated thereafter, while the three-month difference is significant throughout the response window. This pattern is consistent with a role for the timing of rate changes. However, the yield surprises contain different combinations of news and do not isolate the effect of moving the same policy change to a later date.

Additional evidence. Appendix D reports further checks of whether the response differences depend on the measures used, the firms compared, or the regression specification. The shock comparisons use alternative current-rate measures from Swanson (2021), Miranda-Agrippino and Ricco (2021), Bauer and Swanson (2023b), Jarociński and Karadi (2020), and Lewis (2025), and an alternative forward-guidance measure from Swanson (2021) (Appendix Figure D.3). The response differences are also present when we compare firms within narrower sectors (Appendix Figure D.4),

Figure 9: Investment responses to two-year and three-month yield surprises



Notes: The panels plot coefficients on planning-horizon interactions with two-year (left) and three-month (right) yield surprises. Both interactions enter jointly with the baseline controls and fixed effects. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after an increase of 25 basis points on impact in the quarterly-average yield at the corresponding maturity. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

making differences in activities within broad sectors a less likely explanation. Further checks use firm characteristics averaged over time (Appendix Figures D.5 and D.6). We assess sensitivity to the measurement of capital and investment using industry-specific depreciation (Appendix Figure D.7) and alternative investment-rate outcomes (Appendix Figure D.8). We also use an alternative specification with aggregate controls (Appendix Figure D.9). Additional checks consider lagged investment dynamics and inference following Almuzara and Sancibrián (2024), and smooth local projections following Barnichon and Brownlees (2019), using the panel extension of Hartley and Mejia (2025) (Appendix Figure D.10). Finally, alternative upper-tail percentiles produce the same ordering (Appendix Figure D.11), as do comparisons between horizon groups that do not impose a linear relationship between horizons and investment responses (Appendix Figure D.12). The discount-rate point estimates do not reproduce the baseline contrast between investment responses to current-rate and forward-guidance tightening (Appendix Figure D.13). Across these checks, longer planning horizons are generally associated with lower investment responses to forward-guidance tightening and higher responses to current-rate tightening, although the size and precision of the differences vary.

The estimates show how investment responses vary with planning horizons, but leave open the mechanism behind this relationship. A model is needed to understand how differences in planning could generate the contrasting responses to current-rate and forward-guidance tightening.

5 A Model of Investment with Heterogeneous Finite Planning Horizons

The empirical results show that investment at firms with longer revealed planning horizons falls less after current-rate tightening and more after forward-guidance tightening. This contrast generally holds across a broad range of controls and checks addressing mechanisms already discussed in the literature. These findings suggest that planning horizons capture an additional dimension of heterogeneity in investment responses to monetary policy.

The structure of capital budgeting discussed in Section 2 provides a natural starting point. Managers combine the present value of returns forecast within a finite planning horizon with a discounted continuation value for later returns. Woodford's (2019) finite-planning framework formalizes this structure for households' consumption choices, combining explicit finite plans with an approximate continuation value. We adapt this framework to firms' investment decisions.

We develop a New Keynesian model with investment in which capital producers have different finite planning horizons. They choose current investment by explicitly evaluating returns within their planning windows and valuing capital beyond those windows using a steady-state continuation value, while households and price setters have rational expectations and solve their usual infinite-horizon problems. A longer horizon exposes a firm to more of the future policy rate path, which strengthens its response to forward guidance. A longer horizon also lets the firm anticipate more of the recovery in returns of capital that follows the aggregate investment contraction, dampening its response to a current-rate shock. Appendix E provides the derivations and proofs.

5.1 Capital Producers with Finite Planning Horizons

Capital producers are the only agents with finite planning horizons. They form a unit mass indexed by i and are partitioned into finitely many types. Producer i enters date t with capital $K_{i,t}$, rents it to the wholesale sector at the competitive rental rate ϱ_t^K, and chooses investment $I_{i,t}$. The rental rate is the flow return on installed capital. Capital and dividends evolve according to

$$K_{i,t+1} = (1-\delta)K_{i,t} + I_{i,t}, \qquad D_{i,t} = \varrho_t^K K_{i,t} - I_{i,t} - \Psi(I_{i,t}, K_{i,t}), \tag{3}$$

where $\delta \in (0,1)$ is the depreciation rate, $D_{i,t}$ is the dividend paid to the household, and Ψ is a convex adjustment cost, homogeneous of degree one and normalized so that $\Psi(\delta,1) = \Psi_I(\delta,1) = 0$. Producer i's type is its exogenously given, fixed planning horizon $H_i \in \mathcal{H}$, where $\mathcal{H}$ is a finite set of positive integers. The horizon is the number of periods ahead through which the producer reasons when choosing current investment. Let μ_H denote the mass of producers with horizon H, with $\sum_{H\in\mathcal{H}} \mu_H = 1$. Producers otherwise are identical, so $K_{H,t}$ and $I_{H,t}$ denote the common capital and investment of type-H producers.

An optimal finite-horizon plan. The producer's problem adapts Woodford's (2019) finite-horizon planning framework to capital accumulation. At date t, producer i plans investment through

$t+H_i$, taking as given the announced path of gross real interest rates R_t and the returns of capital it expects over this window. Let $\mathcal{R}^{-1}_{t,t+j} \equiv \prod_{\ell=0}^{j-1} R^{-1}_{t+\ell}$, with $\mathcal{R}^{-1}_{t,t} = 1$, denote the cumulative discount factor. The producer solves

$$\max_{\{I_{i,t+j}\}_{j=0}^{H_i}} \mathbb{E}_t^{H_i} \left[\sum_{j=0}^{H_i} \mathcal{R}^{-1}_{t,t+j} \, D_{i,t+j} + \mathcal{R}^{-1}_{t,t+H_i+1} \, v(K_{i,t+H_i+1}) \right], \tag{4}$$

subject to the capital law of motion (3) throughout the planning window. The operator $\mathbb{E}_t^{H_i}$ denotes the expectations of producer i over H_i periods. In the deterministic experiments studied here, all agents know the announced interest rate path, but producers' expectations of endogenous variables can differ from realized outcomes. $v(K)$ is the continuation value function for capital carried beyond the planning window.

The finite plan departs from the fully rational infinite-horizon benchmark in two respects: how the producer forms expectations within the planning window and how continuation values beyond the window are specified.

The first departure concerns expectations within the planning window. These expectations are not model-consistent because, when forming its date-t plan, the producer treats the terminal date as fixed at $t + H_i$ and does not anticipate its future replanning. At $t + 1$, however, it forms a new horizon-H_i plan, so only the current investment $I_{i,t}$ from its date-t plan is implemented. Assumption 5.1 formalizes the expectations used in that plan.

Assumption 5.1 (Finite-Horizon Expectations)**.** In its date-t plan, producer i treats every endogenous decision at date $t + j$ as determined by a problem with $H_i - j$ periods remaining, for $0 \leq j \leq H_i$.

This assumption implies two related restrictions on the producer's forecasts. First, its remaining planning horizon shrinks at successive future dates. Anticipating a fresh H_i-period plan at a later date would require reasoning beyond today's endpoint, $t + H_i$. Second, the producer applies the same shrinking horizon when forecasting the household's, price setters', and other capital producers' decisions.[14] This prevents forecasts of their behavior from requiring explicit reasoning beyond that endpoint.

The second departure concerns how the producer values capital beyond its planning window. The producer uses the continuation valuation associated with the deterministic steady state, under which installed capital is worth its replacement cost at the margin. Assumption 5.2 formalizes this continuation valuation.

Assumption 5.2 (Terminal Values)**.** Every decision problem truncated at the end of the planning window uses the corresponding deterministic-steady-state continuation function to value the state variables carried beyond the window. In particular, the continuation function in the producer's

[14] Assigning the same remaining horizon to all agents provides a tractable specification of the producer's forecasts. Identical horizons are one way to implement the restriction that explicit forward reasoning does not extend beyond $t + H_i$.

problem (4) is

$$v(K) = \bar{R}\,K, \qquad \bar{R} = 1/\beta, \tag{5}$$

where $\bar{R}$ is the steady-state gross real interest rate and β is the household's discount factor.

Following Woodford (2019), these continuation functions can be interpreted as having been learned from past experience rather than deduced from the structural model.[15] The steady-state continuation function is not the only terminal rule consistent with finite planning. The essential restriction is that the terminal valuation not incorporate policy beyond $t + H_i$; otherwise, later announcements would in general affect current investment through the continuation value. The assumption fixes continuation functions, not terminal states, which remain endogenous.

Definition 5.1 (Horizon-H perceived path)**.** Given the date-t predetermined variables and the announced path through $t + H$, the horizon-H perceived path formed at t, $\{Z^{(H)}_{t+j|t}\}_{j=0}^{H}$, is the equilibrium of the economy obtained by applying Assumptions 5.1 and 5.2 to every decision problem inside the producer's plan. At date $t + j$ every such decision, including the producer's own, is determined by a problem with $H - j$ periods remaining.

For any variable Z outside the producer's control, a producer with horizon H forms expectations according to

$$\mathbb{E}^H_t[Z_{t+j}] = Z^{(H)}_{t+j|t}, \qquad 0 \le j \le H. \tag{6}$$

The common truncation applies only inside the perceived path; the actual household and price setters solve their infinite-horizon problems. The origin label t is fixed throughout the plan, and the announced path is common across horizons and therefore carries no horizon label. Because capital producers inside a horizon-H perceived path face the same perceived prices and Ψ is homogeneous of degree one, they choose a common investment rate, so the index i can be dropped. Their first-order conditions are

$$1 + \Psi_I\big(I^{(H)}_{t+j|t}, K^{(H)}_{t+j|t}\big) = Q^{(H)}_{t+j|t}, \qquad 0 \le j \le H, \tag{7}$$

$$Q^{(H)}_{t+j|t} = R^{-1}_{t+j}\left(\varrho^{K,(H)}_{t+j+1|t} - \Psi_K\big(I^{(H)}_{t+j+1|t}, K^{(H)}_{t+j+1|t}\big) + (1-\delta)\,Q^{(H)}_{t+j+1|t}\right), \qquad 0 \le j \le H-1, \tag{8}$$

$$Q^{(H)}_{t+H|t} = R^{-1}_{t+H}\,\bar{R}. \tag{9}$$

Here $Q^{(H)}_{t+j|t}$ is the multiplier on the date-$(t + j)$ capital law and hence the shadow value, within the date-t plan, of capital installed at $t + j$. Equation (9) applies Assumption 5.2 and closes the recursion at the truncation date. Appendix E derives equations (7)–(9) and the corresponding terminal conditions for the perceived household and price setters.

We characterize equilibrium responses with a first-order approximation around the deterministic zero-inflation steady state. For any positive variable X_t, lowercase letters denote log deviations from steady state, $x_t \equiv \log(X_t/\bar{X})$.

[15] Woodford (2019) motivates this interpretation using the learned value functions of chess-playing programs. Section IV of the same paper extends the framework so that value functions are gradually re-learned when policy or fundamentals change persistently. We do not pursue that extension.

Linearizing equations (8) and (9) around the deterministic steady state gives the perceived shadow-value recursion

$$q^{(H)}_{t+j|t} = \frac{1}{\bar{R}}\bar{\varrho}^K r^{K,(H)}_{t+j+1|t} - r_{t+j} + \frac{1}{\bar{R}} q^{(H)}_{t+j+1|t}, \qquad 0 \le j \le H-1, \tag{10}$$

with $q^{(H)}_{t+H|t} = -r_{t+H}$ at the truncation date, where $\bar{\varrho}^K$ is the steady-state rental rate.

Finite-horizon Q theory. Linearizing equation (7) and evaluating it at $j = 0$ gives the rule implemented by a type-H producer,

$$\iota_{H,t} = \frac{q^{(H)}_{t|t}}{\phi\delta}, \tag{11}$$

where $\iota_{H,t} \equiv i_{H,t} - k_{H,t}$ is the log deviation of the investment rate and $\phi \equiv \Psi_{II}(\delta, 1) > 0$ measures the local curvature of adjustment costs. The implemented choice therefore retains the standard Q-theory form: the shadow value of capital is the sufficient statistic for investment. However, finite planning makes the shadow value horizon-specific.

The length of the producer's plan determines both how much of the announced rate path enters its valuation and how far into the future it incorporates the announcement's general equilibrium effects. Solving (10) backward from its terminal condition gives the shadow value governing current investment,

$$q^{(H)}_{t|t} = -\sum_{\ell=0}^{H} \frac{1}{\bar{R}^\ell} r_{t+\ell} + \bar{\varrho}^K \sum_{\ell=1}^{H} \frac{1}{\bar{R}^\ell} r^{K,(H)}_{t+\ell|t}. \tag{12}$$

The two terms in equation (12) correspond to the direct and indirect effects on the shadow value of capital. The first is the direct effect through discounting and depends on the announced real-rate path, including the current rate. The announced rate path is exogenous and common to every perceived path, so it carries no horizon label; producers differ in how much of it their planning windows contain.

The second term captures the indirect effect through perceived future returns of capital, expressed as the discounted value of their deviations from steady state. It begins at $\ell = 1$ because current investment augments next period's capital, whereas current returns of capital are earned on predetermined capital. Along the perceived path, lower aggregate investment leaves less future capital, raising the marginal product of capital and thereby contributing to a recovery in returns of capital. A longer window incorporates more of this perceived response. The complete path of these returns is determined jointly with the rest of the horizon-H perceived path in Section 5.3.

Holding the current state fixed, the horizon imposes a sharp cutoff: a rate change dated $t + H$ affects current investment, whereas one dated $t + H + 1$ does not. A longer horizon therefore exposes the producer to more of the announced rate path, strengthening its direct response to forward guidance, but also to more of the perceived recovery in returns of capital, dampening its response to a current-rate shock.

5.2 The Rest of the Economy

Household. The representative household consumes c_t, supplies labor n_t, and trades a one-period real bond in zero net supply. It owns the retailers and capital producers. Let r_t denote the log deviation of the gross real return on the bond, and let w_t denote the log deviation of the real wage. The Euler and labor-supply conditions are for $t \geq 0$

$$\sigma c_t = \sigma c_{t+1} - r_t, \qquad w_t = \sigma c_t + \varphi n_t. \tag{13}$$

Here, $\sigma > 0$ is the inverse intertemporal elasticity of substitution and $\varphi \geq 0$ is the inverse Frisch elasticity of labor supply.

Production and factor prices. A competitive wholesale sector produces the wholesale good y_t using aggregate capital k_t and labor n_t with a Cobb-Douglas technology. Productivity is a_t, and $\alpha \in (0,1)$ is the capital share. Let mc_t denote real marginal cost, equivalently the wholesale-good price relative to the final-good price under perfect competition. Capital is rented from capital producers at rental rate deviation r_t^K, while labor is paid the real wage w_t. Production and cost minimization imply for $t \geq 0$

$$y_t = a_t + \alpha k_t + (1-\alpha) n_t, \tag{14}$$

$$r_t^K = mc_t + y_t - k_t, \qquad w_t = mc_t + y_t - n_t. \tag{15}$$

Price setting. Monopolistically competitive retailers buy the wholesale good and transform it one for one into differentiated retail varieties. A final-good aggregator combines these varieties into final output. Final and wholesale output coincide at first order. Retailers face Calvo price rigidity with non-reset probability $\theta_p \in (0,1)$ and receive a subsidy that removes the steady-state markup distortion. The resulting Phillips curve is for $t \geq 0$

$$\pi_t = \beta \pi_{t+1} + \kappa_p \, mc_t, \qquad \kappa_p \equiv \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p}, \tag{16}$$

where π_t is inflation, β is the household's discount factor, and κ_p is the slope of the Phillips curve.

Monetary policy. The monetary authority announces an exogenous path $\{r_t\}_{t \geq 0}$ for the real interest rate deviation. All agents know the announced path at date 0. A current-rate shock raises r_t at near-term dates, whereas a forward-guidance shock raises r_t only at future dates.

Aggregate capital. Aggregate capital and the investment rate are capital-weighted averages across producer types,

$$k_t = \sum_H \omega_H k_{H,t}, \qquad \iota_t = \sum_H \omega_H \iota_{H,t}, \qquad \omega_H \equiv \frac{\mu_H \bar{K}_H}{\bar{K}}, \tag{17}$$

where $\iota_t \equiv i_t - k_t$ is the aggregate investment rate. The steady state pins down only aggregate capital, so the shares ω_H are free parameters. Capital accumulates at the type and aggregate levels,

$$k_{H,t+1} = k_{H,t} + \delta \iota_{H,t}, \tag{18}$$

$$k_{t+1} = k_t + \delta \iota_t. \tag{19}$$

Resource constraint. Final output is used for consumption and investment. Let $s_C \equiv \bar{C}/\bar{Y}$ and $s_I \equiv \bar{I}/\bar{Y}$ denote the steady-state consumption and investment shares, with $s_C + s_I = 1$. For $t \geq 0$, the linearized resource constraint is

$$y_t = s_C c_t + s_I\big(k_t + \iota_t\big). \tag{20}$$

5.3 Equilibrium

Equilibrium links two kinds of paths. The horizon-specific perceived paths of Definition 5.1 determine capital producers' current investment choices, and the single realized path is generated when those choices are implemented. The realized equilibrium combines those choices with the behavior of the actual household and price setters, which solve their infinite-horizon problems and anticipate both the choices and the producers' subsequent replanning.

Horizon-H perceived system. At every date, equilibrium contains one perceived path for each $H \in \mathcal{H}$. Because producers within a perceived path choose a common investment rate, their capital stocks enter that perceived economy only through aggregate capital. The realized path nevertheless contains the type-specific stocks $\{k_{H,t}\}_{H\in\mathcal{H}}$.

To first order, the conditions below determine the quantities $\{c^{(H)}_{t+j|t}, n^{(H)}_{t+j|t}, y^{(H)}_{t+j|t}, \iota^{(H)}_{t+j|t}, k^{(H)}_{t+j+1|t}\}_{j=0}^{H}$, prices $\{w^{(H)}_{t+j|t}, r^{K,(H)}_{t+j|t}, mc^{(H)}_{t+j|t}, \pi^{(H)}_{t+j|t}\}_{j=0}^{H}$, and marginal capital values $\{q^{(H)}_{t+j|t}\}_{j=0}^{H}$ along the horizon-H perceived path. Each holds for $0 \leq j \leq H$. Labor and capital markets clear, and perceived bonds are in zero net supply.

- The perceived capital producers value capital and choose investment,

$$q^{(H)}_{t+j|t} = -\sum_{\ell=0}^{H-j} \frac{1}{\bar{R}^{\ell}} r_{t+j+\ell} + \bar{\varrho}^{K} \sum_{\ell=1}^{H-j} \frac{1}{\bar{R}^{\ell}} r^{K,(H)}_{t+j+\ell|t}, \tag{10$'$}$$

$$\iota^{(H)}_{t+j|t} = \frac{q^{(H)}_{t+j|t}}{\phi\delta}; \tag{11$'$}$$

- Capital is initialized at the realized aggregate capital deviation, $k^{(H)}_{t|t} = k_t$, and accumulates according to

$$k^{(H)}_{t+j+1|t} = k^{(H)}_{t+j|t} + \delta\, \iota^{(H)}_{t+j|t}; \tag{19$'$}$$

- The perceived household optimizes,

$$\sigma c^{(H)}_{t+j|t} = -\sum_{\ell=0}^{H-j} r_{t+j+\ell}, \qquad w^{(H)}_{t+j|t} = \sigma c^{(H)}_{t+j|t} + \varphi\, n^{(H)}_{t+j|t}; \tag{13$'$}$$

- The perceived price setters optimize,

$$\pi^{(H)}_{t+j|t} = \kappa_p \sum_{\ell=0}^{H-j} \beta^{\ell} mc^{(H)}_{t+j+\ell|t}; \tag{16$'$}$$

- The perceived goods market clears,

$$y^{(H)}_{t+j|t} = s_C\, c^{(H)}_{t+j|t} + s_I\big(k^{(H)}_{t+j|t} + \iota^{(H)}_{t+j|t}\big); \tag{20$'$}$$

- Output follows the production function,

$$y^{(H)}_{t+j|t} = a_{t+j} + \alpha k^{(H)}_{t+j|t} + (1-\alpha)\, n^{(H)}_{t+j|t}; \tag{14$'$}$$

- Prices are consistent with technology,

$$w^{(H)}_{t+j|t} = mc^{(H)}_{t+j|t} + y^{(H)}_{t+j|t} - n^{(H)}_{t+j|t}, \qquad r^{K,(H)}_{t+j|t} = mc^{(H)}_{t+j|t} + y^{(H)}_{t+j|t} - k^{(H)}_{t+j|t}. \tag{15$'$}$$

The mark ($'$) after an equation number identifies a condition derived from the earlier equation with the same number. The shadow-value equation, the Euler equation, and the Phillips curve are finite forward solutions closed by the corresponding terminal values in Assumption 5.2; their derivations appear in Appendices E.2 and E.3.

Only the continuation valuations are fixed at their steady-state values. Terminal capital, quantities, and prices are determined jointly with the rest of each perceived path.

Definition 5.2 (Equilibrium)**.** Given initial capital deviations $\{k_{H,0}\}_{H\in\mathcal{H}}$ and announced paths $\{a_t, r_t\}_{t\geq 0}$, an equilibrium is a sequence of quantities $\{c_t, n_t, y_t, \{\iota_{H,t}, k_{H,t+1}\}_{H\in\mathcal{H}}\}_{t\geq 0}$, prices $\{w_t, r^K_t, mc_t, \pi_t\}_{t\geq 0}$, and a horizon-$H$ perceived path formed at every date for each $H \in \mathcal{H}$, such that at every date (i) the household optimizes (13), (ii) price setters optimize (16), (iii) each capital producer implements the date-t action of its own perceived path, so its investment satisfies (11), its capital accumulates (18), and type-specific capital and investment aggregate according to (17), (iv) the goods market clears according to (20), labor and capital markets clear, and bonds are in zero net supply, (v) output follows the production function (14) and factor prices equal marginal products (15), and (vi) the realized path satisfies the terminal conditions $\lim_{j\to\infty} c_{t+j} = 0$ and $\lim_{j\to\infty} \pi_{t+j} = 0$. Every perceived path satisfies the conditions of the horizon-H perceived system stated above.

We consider announced paths that are at their deterministic steady-state values from some finite date onward and parameter values for which aggregate-capital dynamics are locally stable.

The terminal conditions above then select the perfect-foresight path of the actual household and price setters.[16]

5.4 Horizon Heterogeneity: Analytical Results

Equations (11) and (12) express the investment rate of a single producer. We now compare producers facing the same announced real-rate path from the same realized state. Consider a long horizon H and a short horizon H', with $H > H'$. Subtracting their investment rates gives the pairwise decomposition

$$\iota_{H,t} - \iota_{H',t} = -\frac{1}{\phi\delta} \sum_{j=H'+1}^{H} \frac{1}{\bar{R}^j} r_{t+j} + \underbrace{\frac{\bar{\varrho}^K}{\phi\delta} \left[\sum_{j=1}^{H} \frac{1}{\bar{R}^j} r^{K,(H)}_{t+j|t} - \sum_{j=1}^{H'} \frac{1}{\bar{R}^j} r^{K,(H')}_{t+j|t} \right]}_{\Delta\Gamma_{H,H',t}}. \tag{21}$$

The investment gap reflects differences in direct effects through discounting and indirect effects through perceived returns of capital. The first term is the difference in direct effects, summed over dates beyond the shorter horizon and within the longer one. At dates shared by both horizons, the direct effects are identical across producers and therefore cancel from the investment gap. The second term, $\Delta\Gamma_{H,H',t}$, is the difference in indirect effects: it compares the discounted valuations of the two perceived paths of returns of capital. The longer-horizon producer includes returns at additional future dates. At dates shared by both horizons, the two producers can also perceive different returns of capital, because each solves a planning problem with a different horizon.

With productivity at steady state, combining household labor supply, goods-market clearing, production, and factor pricing gives

$$r^{K,(H)}_{t+j|t} = \left[\sigma + \frac{(1+\varphi)s_C}{1-\alpha} \right] c^{(H)}_{t+j|t} + \frac{1+\varphi}{1-\alpha} \left[s_I - \frac{1+\varphi\alpha}{1+\varphi} \right] k^{(H)}_{t+j|t} + \frac{(1+\varphi)s_I}{1-\alpha} \iota^{(H)}_{t+j|t}. \tag{22}$$

Equation (22) separates the demand and capital-scarcity effects on returns of capital. The positive coefficients on consumption and the investment rate capture the demand effect: with capital predetermined on impact, lower demand reduces employment and the marginal product of capital, while lower wages reduce marginal cost. Both forces depress returns. The subsequent capital shortfall creates an opposing scarcity effect, raising the marginal product of capital at a given level of employment. After accounting for adjustments in employment and marginal cost, the coefficient on capital is negative by the model's steady-state identities. Appendix E provides the detailed proof. A capital shortfall therefore raises returns holding consumption and the investment rate fixed. The total response depends on the balance between these opposing contributions.

[16] Woodford (2019) shows that finite-horizon planning by all private decision makers can select an outcome under an interest rate peg despite rational-expectations indeterminacy. Here finite planning is confined to capital producers, so the infinitely lived household and price setters still require the boundary conditions above.

Proposition 5.1 (Common-window real-rate path)**.** *Suppose the economy begins at its deterministic steady state, $k_t = 0$, there are no productivity shocks, and the announced real-rate path satisfies $r_{t+j} \geq 0$ for $0 \leq j \leq H'$, with at least one strict inequality, and $r_{t+j} = 0$ for $j > H'$. Then for the horizons $H > H'$,*

$$\iota_{H,t} - \iota_{H',t} = \Delta\Gamma_{H,H',t} > 0. \tag{23}$$

Proposition 5.1 shows that when announced rate changes lie within the planning horizons of both the long- and short-horizon producers, the shorter-horizon producer cuts investment more on impact. Both producers incorporate every announced rate change, so their direct effects through discounting are identical. The difference therefore comes entirely from the indirect effect: the longer-horizon producer anticipates more of the recovery in returns of capital as the aggregate investment contraction makes capital scarce, leading it to cut investment less on impact. Appendix E.4 provides the proof.

Figure 10 illustrates Proposition 5.1 for producers with planning horizons $H_S = 4$ and $H_L = 10$ quarters. For illustration, we set $\beta = 0.99$, $\delta = 0.025$, $\alpha = 0.33$, $\varphi = 1$, $\sigma = 1.25$, and $\phi = 7$. Panel (a) shows a 50-basis-point increase in the annualized real rate, announced at $j = 0$ and implemented for one quarter at $j = 4$, so the increase lies within both producer types' planning horizons.

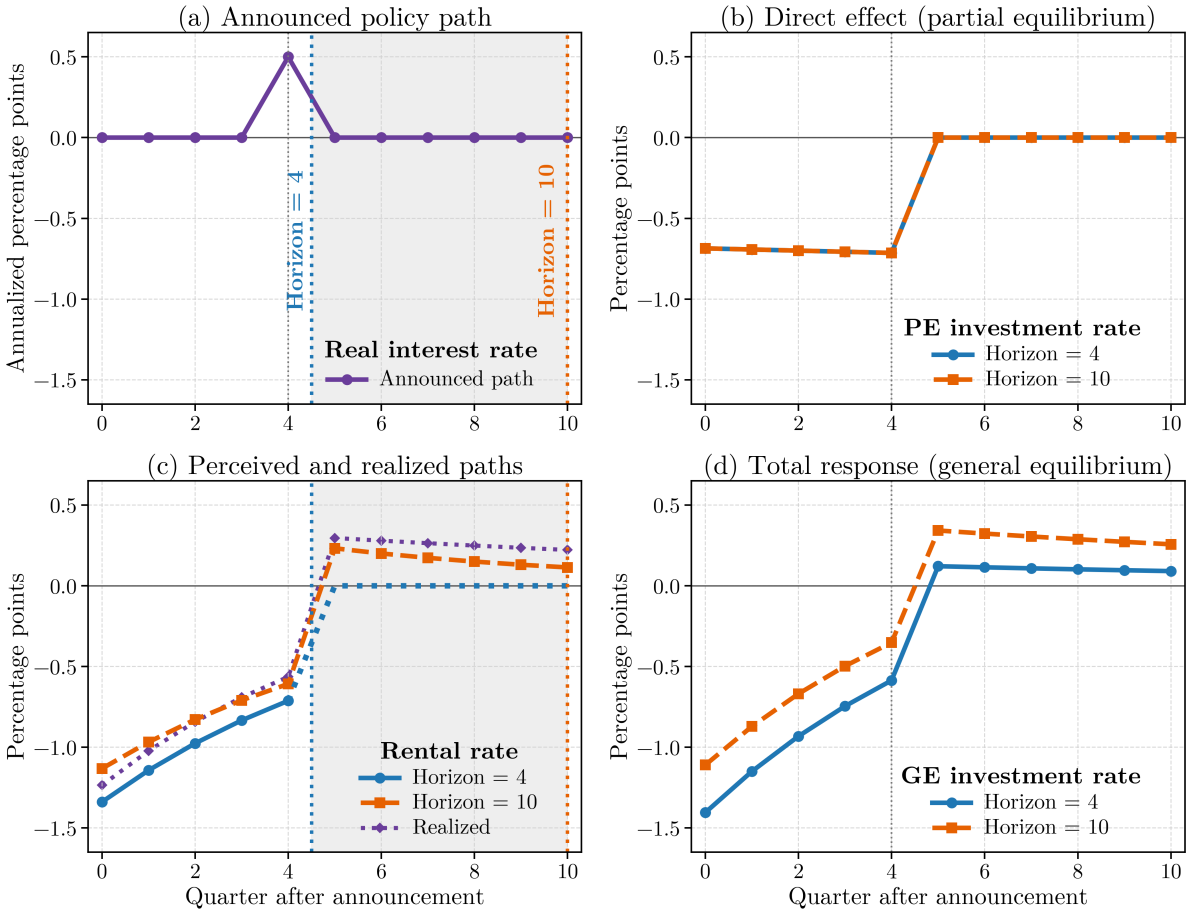
Panel (b) isolates the direct effect through discounting by holding perceived rental-rate deviations at zero. The higher real rate lowers the shadow value of capital and investment rates from announcement through date 4. Since the entire rate increase lies within both planning horizons, both producers incorporate the same change in discounting. Their partial equilibrium investment-rate responses are therefore identical.

Panel (c) shows the perceived rental-rate paths through which general equilibrium adjustments generate the indirect effect. Through implementation, the negative contributions of consumption and the investment rate in Equation (22) outweigh the positive contribution of capital scarcity. Once the rate increase has passed, consumption returns to steady state while capital remains depressed. Capital scarcity and the rebound in investment rates then jointly raise returns of capital.

Both types perceive the initial decline in returns of capital, but their perceived paths differ. The short type's explicit path ends at $H_S = 4$; its dotted continuation at zero represents the steady-state terminal continuation rather than an explicit forecast. The long type incorporates the subsequent recovery in returns of capital. Anticipating that recovery also reduces the perceived shortfall in returns of capital at dates shared by both horizons. The resulting difference between the discounted paths of returns implies $\Delta\Gamma_{H_L,H_S,0} > 0$. The purple line reports the realized rental rate generated as producers implement their current choices and replan, so it need not coincide with either date-0 perceived path.

Panel (d) shows realized general equilibrium investment-rate responses as producers replan each period. The indirect effect deepens both types' initial investment contractions relative to partial equilibrium. This effect is stronger for the shorter-horizon producer, whose plan excludes the subsequent recovery in returns of capital.

Figure 10: Investment responses to a rate increase within both planning horizons



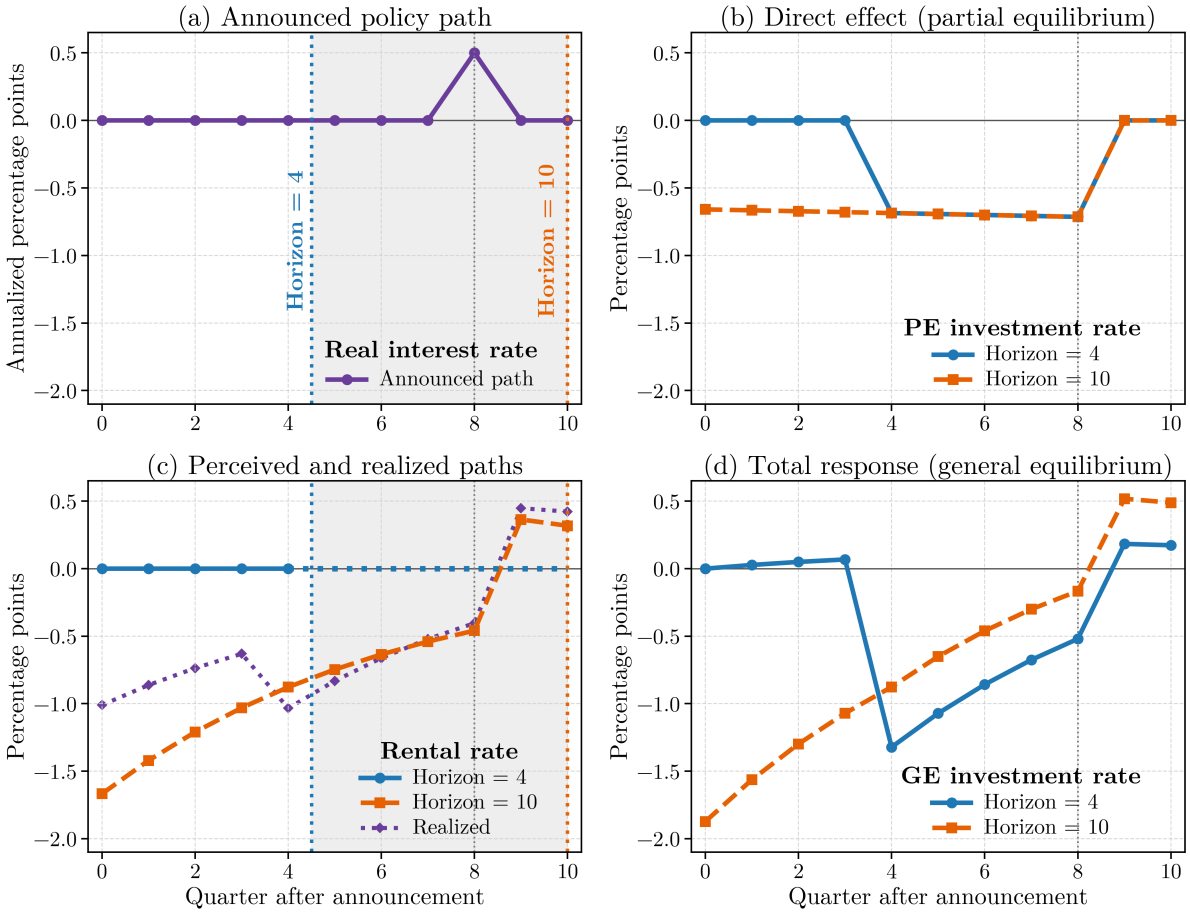
Notes: A one-quarter, 50-basis-point annualized real-rate increase is announced at date 0 and implemented in quarter 4. The two producer types begin at steady state with equal capital weights and have planning horizons of four and ten quarters, identified by the colored vertical lines. Panel (a) shows the announced real-rate path. Thin gray vertical lines mark implementation; shading in panels (a) and (c) marks dates beyond the shorter planning horizon. Panel (b) shows the partial equilibrium investment-rate responses, holding perceived rental-rate deviations at zero. Panel (c) shows rental-rate paths perceived at announcement together with the realized path. Panel (d) shows realized general equilibrium investment-rate responses as producers replan each period.

Proposition 5.2 (Delayed forward-guidance path)**.** *Suppose the economy begins at its deterministic steady state, $k_t = 0$, there are no productivity shocks, and the announced real-rate path satisfies $r_{t+j} \geq 0$ for $H' < j \leq H$, with at least one strict inequality, and $r_{t+j} = 0$ otherwise. Then, for horizons $H > H'$,*

$$\iota_{H',t} = 0, \qquad \iota_{H,t} - \iota_{H',t} = \iota_{H,t} < 0. \tag{24}$$

Proposition 5.2 shows that a rate increase outside the shorter planning horizon but inside the longer one induces only the longer-horizon producer to cut investment on announcement. The shorter-horizon producer perceives a steady-state path, while higher future real rates lower the longer-horizon producer's shadow value of capital through discounting and weaker consumption demand. The difference in indirect effects, $\Delta\Gamma_{H,H',t}$, can be negative or positive, depending on the balance between the demand contraction and the subsequent recovery in returns of capital.

Figure 11: Investment responses to a rate increase within only the longer planning horizon



Notes: A one-quarter, 50-basis-point annualized real-rate increase is announced at date 0 and implemented in quarter 8. The two producer types begin at steady state with equal capital weights and have planning horizons of four and ten quarters, identified by the colored vertical lines. Panel (a) shows the announced real-rate path. Thin gray vertical lines mark implementation; shading in panels (a) and (c) marks dates beyond the shorter planning horizon. Panel (b) shows the partial equilibrium investment-rate responses, holding perceived rental-rate deviations at zero. Panel (c) shows rental-rate paths perceived at announcement together with the realized path. Panel (d) shows realized general equilibrium investment-rate responses as producers replan each period.

Appendix E.4 shows that, after accounting for capital scarcity, each rate increase enters the longer-horizon producer's investment response on announcement with a negative coefficient.

Figure 11 illustrates Proposition 5.2 by moving the one-quarter, 50-basis-point rate increase in Figure 10 to $j = 8$, outside the short planning horizon $H_S = 4$ but inside the long horizon $H_L = 10$, as panel (a) shows.

Panel (b) isolates the direct effect through discounting by holding perceived rental-rate deviations at zero. On announcement, the short type does not respond because the date-8 increase lies outside its planning horizon. The long type incorporates the increase and immediately lowers its investment rate. As producers replan, the increase first enters the short type's planning horizon at $j = 4$. From that date through implementation, both types incorporate the same rate increase, so their partial equilibrium investment-rate responses coincide.

Panel (c) compares rental-rate paths perceived at announcement with the realized path. The short type perceives a steady-state path because the rate increase lies beyond its horizon. The long type expects returns of capital below steady state through implementation, followed by a recovery once consumption returns to steady state while capital remains depressed. In this illustration, the discounted initial decline dominates the subsequent recovery, so the indirect effect reinforces the long type's initial investment contraction. The purple realized path reflects rational households' consumption response and producers' implemented investment decisions. It falls immediately and drops further at $j = 4$, when the increase enters the short type's planning horizon.

Panel (d) shows realized general equilibrium investment-rate responses as producers replan each period. The shorter-horizon producer does not respond on announcement but subsequently raises its investment rate slightly as the capital shortfall increases perceived returns. At $j = 4$, the rate increase enters its planning horizon and its investment rate falls sharply. From this date through implementation, its investment-rate contraction is larger than the longer-horizon producer's despite their identical direct effects. The longer-horizon producer's contraction gradually moderates as replanning incorporates more of the recovery in returns of capital.

Propositions 5.1 and 5.2 show how the model can reproduce the qualitative empirical pattern, as illustrated by the one-quarter rate increases above. A persistent tightening can include rate changes both within and beyond the shorter planning horizon, so the investment gap depends on the balance between direct and indirect effects. For example, under the calibration above, these contributions cancel for a 50-basis-point annualized real-rate increase at announcement that decays geometrically with quarterly persistence $\rho \simeq 0.57$. Producers with four- and ten-quarter horizons then cut investment equally on announcement; at higher persistence, the longer-horizon producer cuts more. We next assess the model's quantitative predictions using the estimated real-rate paths and the measured distribution of planning horizons.

6 Quantitative Implications for Monetary Transmission

We now calibrate the model to our measured planning horizon distribution to assess whether it can reproduce the magnitudes of the empirical investment differentials without targeting them.

6.1 Calibration

Table 3 reports the baseline quarterly calibration. Following Christiano et al. (2005) and Galí and Gertler (2007), we set $\beta = 0.99$, $\delta = 0.025$, $\alpha = 0.33$, $\varphi = 1$, and $\theta_p = 0.75$. The steady-state conditions imply an investment share of $s_I = \delta\alpha/(\beta^{-1} - 1 + \delta) = 0.235$. Following Winberry et al. (2025), we set $\sigma = 1.25$. In our baseline, we set $\phi = 7$, in line with the investment responses to

Table 3: Calibration inputs for the quantitative model.

Parameter	Value	Source / calibration basis
Discount factor, β	0.99	Galí and Gertler (2007)
Depreciation rate, δ	0.025	Christiano et al. (2005)
Capital share, α	0.33	Galí and Gertler (2007)
Inverse Frisch elasticity, φ	1	Galí and Gertler (2007)
Price non-reset probability, θ_p	0.75	Galí and Gertler (2007)
Inverse intertemporal elasticity, σ	1.25	Winberry et al. (2025)
Adjustment-cost curvature, ϕ	7	Cummins et al. (1996) Zwick and Mahon (2017)

tax changes estimated by Cummins et al. (1996) and Zwick and Mahon (2017)[17] and we assess sensitivity to alternative values below.

We discretize the empirical distribution of planning horizons on a quarterly grid and winsorize its upper tail at forty quarters (ten years). We aggregate across types using the steady-state capital weights ω_H defined in equation (17). Appendix Table E.1 reports every type and its weight. We hold all parameters other than H common across types to isolate planning-horizon heterogeneity.

We specify the interest-rate paths for the model's current-rate and forward-guidance experiments using estimated responses of the quarterly average three-month US Treasury bill rate. In the baseline, these nominal-rate responses serve as proxies for the corresponding ex ante real-rate responses. We use the three-month maturity because it matches the model's quarterly period. Appendix Figure E.2 shows the resulting paths. Appendix E.5 describes their estimation and extension beyond the estimated quarters, using geometric decay for the forward-guidance path with persistence broadly in line with the estimated interest-rate response to a monetary policy shock. It also reports robustness checks using ex post real rates and lower persistence.

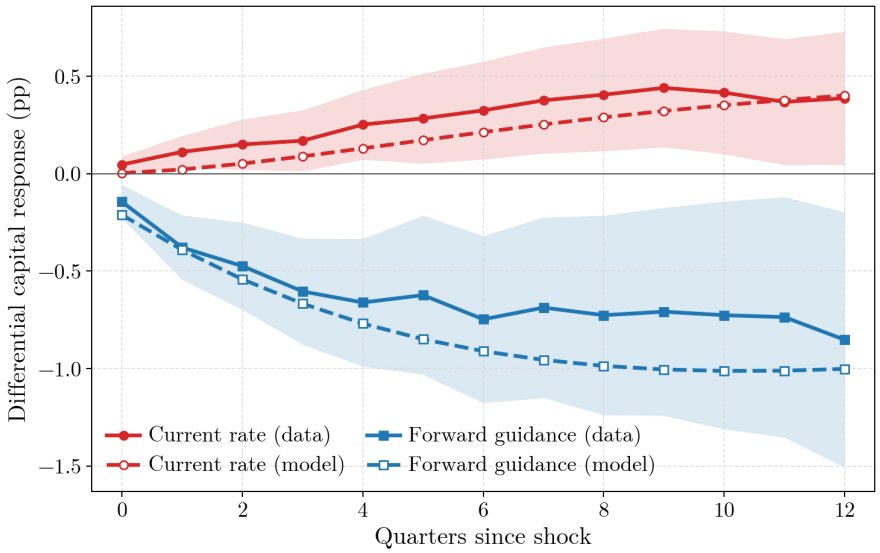
6.2 Reproducing Untargeted Empirical Investment Differentials

We feed the empirical monetary policy paths into the calibrated model and compare the resulting planning-horizon differentials with their empirical counterparts in Figure 12. Each differential measures the additional cumulative capital-growth response associated with a one-standard-deviation increase in log planning horizon. A negative differential means a larger contraction for longer-horizon firms; a positive differential means a smaller contraction. Under the estimated policy paths, the model reproduces the empirical contrast: firms with longer planning horizons experience a larger cumulative capital contraction after forward-guidance tightening and a smaller contraction after current-rate tightening.

Following forward-guidance tightening, both differentials are negative from impact and track each other closely during the first year. The model subsequently overstates the magnitude of the

[17] Using the annual-to-quarterly conversion, $\phi = 7$ corresponds to an annual investment–Q coefficient of 0.57, compared with the US estimate of 0.65 (95% CI: $[0.50, 0.80]$) in Cummins et al. (1996). It implies an investment–q sensitivity of 5.71, compared with the tax-price elasticity of 7.2 (95% CI: $[5.17, 9.22]$) in Zwick and Mahon (2017). These confidence intervals are calculated from the reported standard errors.

Figure 12: Planning-horizon differentials after current-rate and forward-guidance shocks



Notes: The figure compares empirical (solid lines) and model-implied (dashed lines) differences in cumulative capital growth following current-rate and forward-guidance tightening. Differentials are measured in percentage points for a one-standard-deviation increase in log planning horizon. Model predictions use the calibration in Table 3. Shaded bands show 90 percent confidence intervals for the empirical estimates.

differential, reaching −1.00 percentage points by quarter 12 compared with −0.85 in the data, but stays within the empirical 90 percent confidence band throughout.

Following current-rate tightening, the model differential is initially close to zero and rises to 0.40 percentage points by quarter 12, compared with 0.39 in the data. The model understates the empirical differential during the first year but lies within its 90 percent confidence band from quarter 2 onward. The small initial differential reflects the persistence of the estimated real-rate path: longer-horizon firms incorporate later rate increases that fall beyond shorter-horizon firms' planning horizons. The additional direct effect of these later increases partly offsets the positive contribution of indirect effects to the initial differential.

The model gives empirical content to theories of limited forward reasoning by reproducing untargeted investment differentials using the measured distribution of planning horizons.

The qualitative contrast between current-rate and forward-guidance responses holds across several robustness checks: Appendix E shows that the model reproduces this pattern across a range of alternative investment adjustment-cost values, when effective federal funds rates or one-year Treasury yields supply the nominal-rate inputs, and when ex post real-rate paths replace nominal-rate paths as proxies for ex ante real-rate paths.

Alternative explanations Differences in measured planning horizons may reflect differences in firms' discount rates. Such an interpretation must jointly explain the weaker response of longer-horizon firms to current-rate tightening and their stronger response to forward-guidance tightening. In an infinite-planning benchmark with common investment technology and fixed returns of capital,

firms with lower discount rates place more weight on future rate increases and cut investment more after both types of tightening. This partial equilibrium prediction does not reproduce the empirical pattern. In general equilibrium, however, more patient firms also place greater weight on the subsequent recovery in returns of capital. The empirical pattern can arise if the indirect contribution to the investment differential dominates the direct contribution after current-rate tightening but not after forward-guidance tightening. Appendix Figure D.13 shows larger estimated contractions for firms with lower predicted hurdle rates after both shocks, consistent with indirect effects not overturning the direct ranking.

Capital durability could also explain differences in measured planning horizons. In an infinite-planning benchmark with our installation technology and common adjustment-cost curvature, fixed adjustments to rental payments equalize returns net of depreciation across firms. Their cumulative capital-growth responses are then identical: the lower surviving value of faster-depreciating capital is offset by the greater installation costs saved by holding an additional unit. This cancellation holds even when returns of capital adjust in general equilibrium, so depreciation differences alone do not generate the empirical contrast in this benchmark (Appendix E.9). Empirically, the one-year planning-horizon differentials have opposite signs after current-rate and forward-guidance tightening when investment responses are allowed to vary with asset maturity or tangibility (Appendix Figure D.1). The estimated differentials are also nearly unchanged when capital is constructed using industry-specific depreciation rates (Appendix Figure D.7).

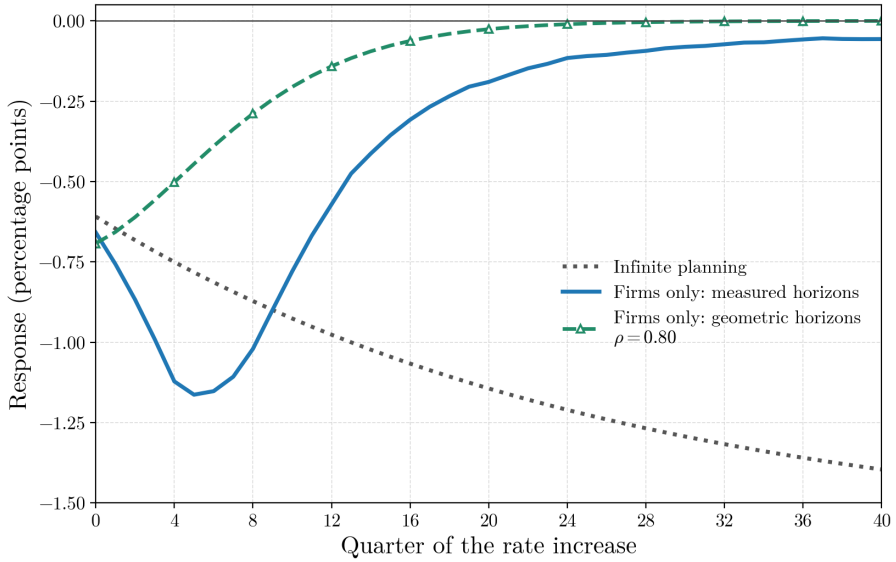
6.3 Implications for the Aggregate Investment Response

We now assess how the measured distribution of planning horizons shapes aggregate investment responses to monetary policy. We study a one-quarter, 50-basis-point increase in the annualized real rate, announced at date zero, and vary the date at which it takes effect. An immediate increase represents current-rate tightening; an increase at a future date represents forward guidance. Figure 13 plots the aggregate investment response at announcement as a function of the quarter in which the rate increase occurs. We compare the response under the measured distribution of finite planning horizons (solid line) with two otherwise identical benchmark economies. In the infinite-planning benchmark (dotted line), capital producers plan over an infinite horizon. In the geometric benchmark (dashed line), their planning horizons follow the geometric distribution considered by Woodford (2019), with the degree of attenuation calibrated in line with the literature.[18]

Under infinite planning, more distant forward guidance produces a larger investment contraction at announcement. All firms anticipate the rate increase, and delaying its implementation prolongs the anticipated period of weak demand and low returns of capital, strengthening the incentive to

[18] Planning horizons satisfy $\Pr(H = h) = (1 - \rho)\rho^h$, for $h = 0, 1, \ldots$. The share of firms whose plans include a rate increase T quarters ahead is therefore ρ^T. Averaging the direct effect through discounting across these horizons gives the same geometric attenuation as a cognitive-discounting rule with parameter ρ (Gabaix, 2020), although indirect effects through perceived returns of capital can differ. We set $\rho = 0.80$, matching the value of the cognitive-discounting coefficient in firms' pricing equation in Gabaix (2020). This exceeds the posterior mean estimates of ρ in Gust et al. (2022), which range from 0.46 to 0.69 across specifications. Appendix E.7 provides a detailed comparison of finite planning and cognitive discounting.

Figure 13: Announcement-date investment response by implementation date



Notes: Each experiment announces at date zero a one-quarter, 50-basis-point annualized real-rate increase, with rates at steady state at all other dates. The horizontal axis gives the implementation quarter; the vertical axis gives the investment response at date zero in percentage points. The solid line uses the measured distribution of producer planning horizons, weighted by firm counts; the dashed line uses geometric weights $\Pr(H = h) = (1 - \rho)\rho^h$ for $h = 0, 1, \ldots$, with $\rho = 0.80$; the dotted line shows infinite planning. Households have rational expectations in all three economies, and all other parameters are unchanged.

cut investment today. This increasing potency of distant policy news is the investment counterpart of the forward-guidance puzzle (Del Negro et al., 2023, McKay et al., 2016).[19]

Both the measured distribution of finite planning horizons and the geometric benchmark overturn the infinite-planning prediction for distant forward guidance. As the right-hand side of Figure 13 shows, with measured horizons, a rate increase twenty quarters ahead—the five-year horizon considered by McKay et al. (2016)—reduces investment by 0.19 percent at announcement, compared with 1.14 percent under infinite planning, a reduction of 83 percent. The substantial attenuation is qualitatively the same as that obtained from an assumed geometric distribution with a calibrated parameter, but here it arises from the distribution of revealed planning horizons measured from earnings calls. The result therefore connects the attenuation emphasized theoretically by Woodford (2019) and Gabaix (2020) to firm-level evidence on how far ahead firms plan.

At shorter horizons, however, the measured distribution and the geometric benchmark yield different predictions. Under the geometric benchmark, the investment contraction weakens steadily as implementation is delayed. With the measured distribution of finite planning horizons, the contraction first strengthens and then weakens, producing a hump: forward guidance has its strongest effect on investment at announcement when the rate increase is one to two years ahead. The contraction is largest at five quarters, when investment falls by 1.16 percent, compared with 0.78 percent under infinite planning—a 49 percent larger contraction.

[19] Our exercise concerns the investment channel of the forward-guidance puzzle. Household consumption satisfies the rational Euler equation in all three economies; finite planning restricts capital producers' investment decisions.

The hump reflects two opposing effects of delaying the rate increase: stronger investment cuts among responding firms and a declining share of firms that respond. For near-term increases, delaying implementation prolongs the anticipated period of low returns of capital and leaves less of the subsequent recovery within responding firms' plans, strengthening their investment cuts. Limited anticipation of this recovery also makes near-term cuts larger than under infinite planning. At the same time, a more distant increase falls outside more firms' planning horizons, reducing the responding share. With the measured distribution, stronger cuts initially dominate, while the declining share dominates at more distant dates. At the standard calibration of $\rho = 0.80$ used in the geometric benchmark, the responding share falls by 20 percent with each additional quarter of delay. At this calibration, any increase in cuts among responding firms is too small to offset that decline, so the aggregate contraction weakens monotonically.

The hump-shaped pattern in investment is consistent with evidence that forward guidance can be most powerful at intermediate horizons. Goodhead and Kolb (2025) find evidence that, in the United States, communication about rates two years ahead can have stronger contractionary effects than communication one or three years ahead. Similarly, Brubakk et al. (2022) find that the effects of forward guidance in Norway and Sweden weaken as the guidance horizon lengthens.

Capital and limited forward reasoning The hump-shaped investment response in our model arises from the interaction between capital and finite planning. We assess how its strength depends on the measured distribution of planning horizons and on whether households or capital producers have finite horizons, before comparing these responses with those generated by other limits on forward reasoning.

The measured distribution of planning horizons matters for the strength of the investment response beyond the average horizon. For a geometric distribution to match the mean of the measured distribution, its parameter must be $\rho \simeq 0.92$, well above the values typically used in the literature.[20] This geometric distribution also produces a hump (Appendix Figure E.9). However, its maximum investment contraction is about 36 percent smaller than under the measured distribution. Matching the average planning horizon is therefore not enough to reproduce the strength of the investment response.

Whether households or capital producers have finite planning horizons also matters for the strength of monetary transmission. Applying the measured horizon distribution to a finite-horizon household consumption rule, while allowing capital producers to plan over infinite horizons, also produces a hump in output at announcement (Appendix Figure E.8). However, at every implementation date shown, output falls by the same amount or less than under infinite planning. When capital producers have finite horizons and households are rational, near-term rate increases produce

[20] For geometric planning, Gust et al. (2022) estimate ρ between 0.46 and 0.69, while Dupraz et al. (2024) use 0.50 and 0.80. For cognitive discounting, Gabaix (2020) uses 0.85 for households and 0.80 for firms and interprets Fuhrer and Rudebusch's (2004) estimates as implying 0.65. Eichenbaum et al. (2026) estimate 0.302, while Ilabaca et al. (2020) estimate 0.706 for households and 0.412 for firms. Other values include 0.68 in Andrade et al. (2019), 0.46 in Afsar et al. (2024), 0.68 in Kolasa et al. (2025), 0.88 for the United States in Benchimol et al. (2025), and 0.79 for the 1982–2007 sample in Meggiorini and Milani (2026). For higher-order feedback, Angeletos and Lian (2018) use 0.75.

larger output contractions than under infinite planning. Thus, both cases generate a hump, but only finite planning by capital producers amplifies the near-term output response in this experiment.

We also assess whether other departures from full-information rational expectations generate the same investment responses. It is important to distinguish attenuation of higher-order feedback from attenuation across calendar time. The first weakens the feedback operating through agents' expectations of other agents' responses, as in dispersed information, level-k thinking, and reflective equilibrium (Angeletos and Lian, 2018; Farhi and Werning, 2019; García-Schmidt and Woodford, 2019). The second reduces the weight agents place on distant outcomes, as in cognitive discounting (Gabaix, 2020) and finite planning (Woodford, 2019). As emphasized by La'O (2019) and Angeletos and Lian (2023), these different microfoundations yield similar aggregate implications in a standard New Keynesian model without capital: they dampen the influence of future expectations on current activity and weaken distant forward guidance.

The two groups yield different predictions for forward guidance in a model with capital. Appendix Figure E.10 illustrates this distinction by calibrating alternative expectation restrictions and applying them to both households and capital producers. Under this calibration, attenuation of higher-order feedback allows more distant forward guidance to deepen the investment contraction at announcement. These restrictions leave the direct effect through discounting unchanged and weaken the indirect effect through perceived returns of capital. Under the cognitive-discounting and geometric finite-planning benchmarks,[21] the investment contraction at announcement becomes progressively weaker as the direct effect fades with delay and the indirect effect does not offset that decline. Weakening feedback from other agents therefore does not by itself reproduce the declining effectiveness of forward guidance in these comparisons.

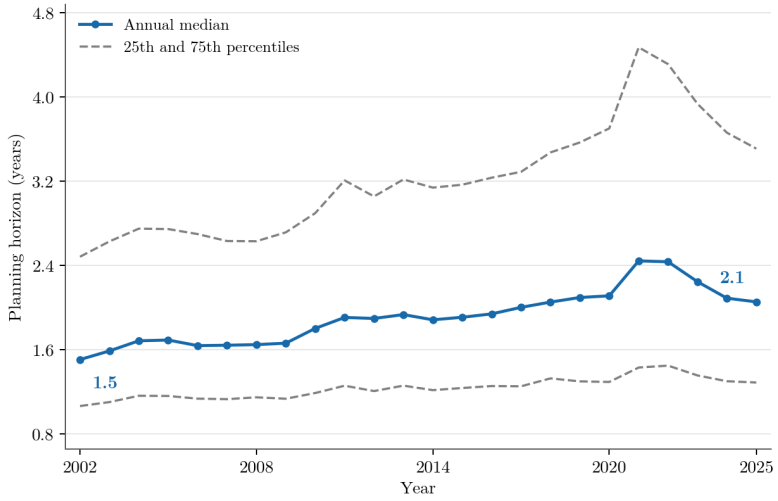
Cumulative investment response A weaker response at announcement need not imply a smaller cumulative investment contraction. As firms replan, rate increases initially beyond their horizons enter their investment decisions, shifting adjustment toward implementation. Appendix Figure E.1 shows the cumulative investment responses over the full transition as a function of the quarter in which the rate increase occurs. Under the measured horizon distribution, the cumulative contraction is more than twice as large as under infinite planning for rate increases implemented up to three years ahead. For an increase five years ahead, it is still 41 percent larger despite the much weaker announcement response. For sufficiently distant increases, the cumulative contraction becomes smaller than under infinite planning. Measured planning horizons therefore change both when investment adjusts and how much it contracts cumulatively.

6.4 The Effects of Lengthening Planning Horizons

We study how long-term changes in firms' planning horizons affect the aggregate investment response to monetary policy. Figure 14 reports the annual median and the 25th and 75th percentiles of the

[21] The geometric comparison averages separately solved economies in which households and capital producers share a common planning horizon. It differs from the economy with heterogeneous producer horizons and rational households in Appendix Figure E.9.

Figure 14: Annual evolution of firm planning horizons



Notes: The solid line shows the annual median planning horizon across firms in all countries; dashed lines show the 25th and 75th percentiles. Each firm's annual horizon is the average of its call-level horizons in that year.

distribution of firms' horizons. Planning horizons change substantially over the period considered: the median grows by more than one-third (36 percent), from 1.50 years in 2002 to 2.05 years in 2025. The 25th and 75th percentiles also rise, and the gap between them widens. The increase is gradual over most of the period, with a sharp rise between 2020 and 2021 that persists through 2022 before horizons decline toward their pre-pandemic level.[22]

We compare economies with the fitted 2002 and 2025 planning-horizon distributions, holding all other model parameters fixed, to isolate the implications of a rise in firms' planning horizons for monetary transmission.[23] In each economy, we announce at date zero a one-quarter, 50-basis-point increase in the annualized real rate and vary the date at which it occurs. Table 4 shows investment responses at announcement and cumulatively over the full transition for rate increases five years ahead, one year ahead, and immediately.

The rise in planning horizons strengthens the announcement response to distant forward guidance while leaving the response to near-term rate increases largely unchanged. For a rate increase five years ahead, investment falls at announcement by 0.25 percent under the 2025 distribution, compared with 0.15 percent under the 2002 distribution—a 73 percent larger contraction. A larger share of firms has planning horizons that include the announced increase, allowing more firms to adjust investment when the news arrives. By comparison, announcement contractions are only 6 percent smaller for an increase one year ahead and 2 percent smaller for an immediate increase.

The cumulative investment contraction following a rate increase five years ahead is essentially unchanged. For near-term tightening, however, the cumulative contraction is smaller under the 2025 distribution: 23 percent smaller for an increase one year ahead and 21 percent smaller for an

[22] The pandemic increase may reflect firms reconsidering their long-term strategies in response to lasting changes in demand and business operations; see Maqui and Morris (2020).

[23] We fit generalized-gamma distributions with parameters that vary by year using observations from 2002 through 2025. Appendix E.5 describes the estimation, and Appendix Figure E.6 shows the fitted and empirical distributions.

immediate increase. Thus, the shift in horizons strengthens the initial response to distant guidance but weakens the cumulative response to near-term tightening.

The rise in planning horizons following the Covid pandemic makes the model's announcement response to distant forward guidance stronger still, while further reducing the cumulative contraction after near-term tightening. Relative to the 2002 distribution, the fitted 2021 distribution produces a 140 percent larger announcement contraction for a rate increase five years ahead and a 33 percent smaller cumulative contraction for an increase one year ahead (Appendix Figure E.7 and Appendix Table E.2). The subsequent decline in planning horizons moderates both model-implied effects by 2025.

Forward guidance can become more powerful through changes both in central banks' ability to shape policy expectations and in firms' responses to those expectations. Bernanke (2020) argues that better execution and greater public understanding can increase its effectiveness, while Campbell et al. (2019) show how imperfect communication limits central banks' influence on expectations about distant rates. Planning horizons impose a further limit: even a clearly communicated distant rate change can lie beyond the period over which a firm explicitly evaluates investment returns. Our counterfactual shows that longer horizons strengthen the announcement response of investment to a given distant rate increase, providing a complementary explanation for the growing influence of forward guidance on current investment.

Central bank communication may also influence how far ahead firms choose to plan. Information about future policy and economic conditions could reduce the cost of gathering information for longer-term investment assessments or make those assessments more useful. If firms respond by extending their planning horizons, improvements in communication could contribute to the observed lengthening and thereby strengthen the investment response to distant guidance. Our model takes planning horizons as given and does not establish this explanation. A question for future research is whether central bank communication encourages firms to plan further ahead by changing the costs and benefits of evaluating distant investment returns.

7 Conclusion

How do changes in expected future economic conditions affect firms' investment? Our answer is that their effects depend on how far ahead firms plan. Our measure of revealed planning horizons, constructed from corporate earnings calls, documents large and persistent differences in how far ahead firms plan. We study monetary policy announcements as they provide a clean setting in which news arrives today but concerns changes to future economic activity. This allows us to examine how investment responses to such news vary with firms' planning horizons. Among firms in the same sector facing the same announcement, those with longer revealed horizons cut investment more after forward-guidance tightening and less after current-rate tightening.

Our New Keynesian model with heterogeneous planning horizons gives empirical content to theories of limited forward reasoning (Woodford, 2019; Gabaix, 2020) by quantifying their impli-

Table 4: Investment responses under the 2002 and 2025 horizon distributions

Rate increase	2002	2025	Change in contraction (%)
Panel A: Investment response at announcement (%)			
Five years ahead	-0.146	-0.254	+73.5
One year ahead	-1.114	-1.048	-6.0
Immediate	-0.671	-0.660	-1.6
Panel B: Cumulative investment response (pp-quarters)			
Five years ahead	-4.434	-4.452	+0.4
One year ahead	-1.860	-1.427	-23.3
Immediate	-0.275	-0.218	-20.8

Notes: Each experiment announces at date zero a one-quarter, 50-basis-point annualized real-rate increase. The columns use the fitted 2002 and 2025 planning-horizon distributions. Panel A reports investment at announcement; Panel B sums quarterly investment responses over the full transition. Percentage changes compare absolute contractions with those under the 2002 distribution and use unrounded responses. A positive change denotes a larger contraction. Firms have equal steady-state capital, households have rational expectations, and all other parameters are fixed.

cations at the horizons observed in the data. At these horizons, finite planning can more than double the cumulative investment contraction relative to infinite planning. The model also changes these theories' qualitative predictions: forward guidance has its strongest effect on investment at announcement when rate increases are one to two years ahead.

The investment implications of planning horizons may extend beyond monetary policy. Our measure makes it possible to examine whether firms that plan further ahead devote more resources to research and development and other long-term investments, and how these choices relate to their subsequent profitability. Linking these investments to the technologies firms develop and adopt could clarify whether differences in planning horizons contribute to differences in productivity growth.

Understanding what makes firms plan further ahead and how policy can shape their planning horizons is an important direction for future research. Our evidence on long-term CEO compensation and institutional ownership provides a starting point for examining whether changes in managerial and shareholder incentives lead firms to plan further ahead. A related question is whether lower policy uncertainty or clearer central bank communication influences the horizon over which firms evaluate investments. Answering these questions would clarify whether incentives and policy influence investment by changing the horizon over which firms assess future returns.

References

Adams, Jonathan J. and Philip Barrett (2025). *What Are Empirical Monetary Policy Shocks? Estimating the Term Structure of Policy News*. IMF Working Paper WP/25/128. International Monetary Fund.

Afsar, Atahan et al. (2024). *A Behavioral Hybrid New Keynesian Model: Quantifying the Importance of Belief Formation Frictions*.

Ahn, Hie Joo, Choongryul Yang, and Shihan Xie (Oct. 2022). "Effects of Monetary Policy on Household Expectations: The Role of Homeownership". In: *Finance and Economics Discussion Series* 2022-065, pp. 1–50.

Alder, Marie, Nuno Coimbra, and Urszula Szczerbowicz (2025). "Corporate Debt Structure and Heterogeneous Monetary Policy Transmission". In: *SSRN Electronic Journal*.

Almuzara, Martín and Víctor Sancibrián (2024). *Micro Responses to Macro Shocks*. Staff Reports 1090. Revised August 2026. Federal Reserve Bank of New York.

Altavilla, Carlo, Refet S. Gürkaynak, and Rogier Quaedvlieg (Oct. 2024). "Macro and Micro of External Finance Premium and Monetary Policy Transmission". In: *Journal of Monetary Economics* 147, p. 103634.

Altig, David et al. (2021). "Firm-Level Planning Frequencies and Forecast Horizons". In: *SSRN Electronic Journal*.

Amel-Zadeh, Amir, Fiona Kasperk, and Martin C. Schmalz (2022). *Mavericks, Universal, and Common Owners—The Largest Shareholders of U.S. Public Firms*. Working Paper 9926. CESifo.

Ameriks, John, Andrew Caplin, and John Leahy (Aug. 2003). "Wealth Accumulation and the Propensity to Plan". In: *The Quarterly Journal of Economics* 118.3, pp. 1007–1047.

Anderson, Gareth and Ambrogio Cesa-Bianchi (July 2024). "Crossing the Credit Channel: Credit Spreads and Firm Heterogeneity". In: *American Economic Journal: Macroeconomics* 16.3, pp. 417–446.

Andrade, Joaquim, Pedro Cordeiro, and Guilherme Lambais (2019). *Estimating a Behavioral New Keynesian Model*.

Angeletos, George-Marios and Chen Lian (Sept. 2018). "Forward Guidance without Common Knowledge". In: *American Economic Review* 108.9, pp. 2477–2512.

— (2023). "Dampening General Equilibrium: Incomplete Information and Bounded Rationality". In: *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Academic Press. Chap. 20, pp. 613–645.

Angeletos, George-Marios and Karthik A Sastry (Nov. 2021). "Managing Expectations: Instruments Versus Targets". In: *The Quarterly Journal of Economics* 136.4, pp. 2467–2532.

Antràs, Pol and Vitalii Tubdenov (May 2025). "Measuring the Average Period of Production". In: *AEA Papers and Proceedings* 115, pp. 624–630.

Appel, Ian R and Christoph Herpfer (2026). "Corporate Horizons". Working paper.

Barnichon, Regis and Christian Brownlees (2019). "Impulse Response Estimation by Smooth Local Projections". In: *The Review of Economics and Statistics* 101.3, pp. 522–530.

Bauer, Michael D. and Eric T. Swanson (May 2023a). "A Reassessment of Monetary Policy Surprises and High-Frequency Identification". In: *NBER Macroeconomics Annual* 37, pp. 87–155.

— (Mar. 2023b). "An Alternative Explanation for the "Fed Information Effect"". In: *American Economic Review* 113.3, pp. 664–700.

Benchimol, Jonathan, Lahcen Bounader, and Mario Dotta (2025). *Estimating Behavioral Inattention.*

Bernanke, Ben S. (Apr. 2020). "The New Tools of Monetary Policy". In: *American Economic Review* 110.4, pp. 943–983.

Bizjak, John M., James A. Brickley, and Jeffrey L. Coles (1993). "Stock-based incentive compensation and investment behavior". In: *Journal of Accounting and Economics* 16.1–3, pp. 349–372.

Bolton, Patrick, José Scheinkman, and Wei Xiong (2006). "Executive Compensation and Short-Termist Behaviour in Speculative Markets". In: *Review of Economic Studies* 73.3, pp. 577–610.

Brealey, Richard A., Stewart C. Myers, and Franklin Allen (2011). *Principles of Corporate Finance.* 10th ed. New York: McGraw-Hill/Irwin.

Brochet, Francois, Maria Loumioti, and George Serafeim (Sept. 2015). "Speaking of the Short-Term: Disclosure Horizon and Managerial Myopia". In: *Review of Accounting Studies* 20.3, pp. 1122–1163.

Brubakk, Leif et al. (2022). "The Macroeconomic Effects of Forward Communication". In: *Journal of International Money and Finance* 120, 102536.

Caglio, Cecilia R., R. Matthew Darst, and Şebnem Kalemli-Özcan (Aug. 2026). *Collateral Heterogeneity and Monetary Policy Transmission: Evidence from Loans to SMEs and Large Firms.* Working Paper 28685. Revised August 2026. National Bureau of Economic Research.

Caldara, Dario and Matteo Iacoviello (2022). "Measuring Geopolitical Risk". In: *American Economic Review* 112.4, pp. 1194–1225.

Caldara, Dario et al. (2020). "The Economic Effects of Trade Policy Uncertainty". In: *Journal of Monetary Economics* 109, pp. 38–59.

Campbell, Jeffrey R. et al. (2012). "Macroeconomic Effects of Federal Reserve Forward Guidance". In: *Brookings Papers on Economic Activity* 2012.1, pp. 1–80.

Campbell, Jeffrey R. et al. (2019). "The Limits of Forward Guidance". In: *Journal of Monetary Economics* 108, pp. 118–134.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy". In: *Journal of Political Economy* 113.1, pp. 1–45.

Cloyne, James et al. (Dec. 2023). "Monetary Policy, Corporate Finance, and Investment". In: *Journal of the European Economic Association* 21.6, pp. 2586–2634.

Coibion, Olivier et al. (Oct. 2023). "Forward Guidance and Household Expectations". In: *Journal of the European Economic Association* 21.5, pp. 2131–2171.

Cummins, Jason G., Kevin A. Hassett, and R. Glenn Hubbard (1996). "Tax Reforms and Investment: A Cross-Country Comparison". In: *Journal of Public Economics* 62.1–2, pp. 237–273.

D'Acunto, Francesco, Daniel Hoang, and Michael Weber (Mar. 2022). "Managing Households' Expectations with Unconventional Policies". In: *The Review of Financial Studies* 35.4, pp. 1597–1642.

Damodaran, Aswath (2012). *Investment Valuation: Tools and Techniques for Determining the Value of Any Asset*. 3rd ed. John Wiley & Sons.

Darmouni, Olivier, Oliver Giesecke, and Alexander Rodnyansky (Apr. 2022). *The Bond Lending Channel of Monetary Policy*. SSRN Scholarly Paper. Rochester, NY.

de Fraisse, Antoine Hubert et al. (2026). "Pre-Trends Ahead of Monetary Policy Surprises: Recovering the Response of Corporate Investment". Working paper.

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson (Mar. 2023). "The Forward Guidance Puzzle". In: *Journal of Political Economy Macroeconomics* 1.1, pp. 43–79.

Deng, Minjie and Min Fang (May 2022). "Debt Maturity Heterogeneity and Investment Responses to Monetary Policy". In: *European Economic Review* 144, p. 104095.

Derrien, François, Ambrus Kecskés, and David Thesmar (2013). "Investor Horizons and Corporate Policies". In: *Journal of Financial and Quantitative Analysis* 48.6, pp. 1755–1780.

Dessaint, Olivier, Thierry Foucault, and Laurent Frésard (June 2023). *The Horizon of Investors' Information and Corporate Investment*. SSRN Scholarly Paper. Rochester, NY.

Drechsel, Thomas et al. (Sept. 2026). *The Investment Channel of Monetary Policy: Disentangling Firm Heterogeneity*. Discussion Paper 22000. Centre for Economic Policy Research.

Dupraz, Stéphane, Hervé Le Bihan, and Julien Matheron (Apr. 2024). "Make-up Strategies with Finite Planning Horizons but Infinitely Forward-Looking Asset Prices". In: *Journal of Monetary Economics* 143, p. 103542.

Edmans, Alex, Vivian W. Fang, and Katharina A. Lewellen (2017). "Equity Vesting and Investment". In: *The Review of Financial Studies* 30.7, pp. 2229–2271.

Eichenbaum, Martin, João Guerreiro, and Jana Obradović (2026). *Ricardian Non-Equivalence*.

Farhi, Emmanuel and Iván Werning (Nov. 2019). "Monetary Policy, Bounded Rationality, and Incomplete Markets". In: *American Economic Review* 109.11, pp. 3887–3928.

Fernandes, Adriano and Rodolfo Rigato (July 2024). "K Wasn't Built in a Day: Investment with Endogenous Time to Build". Working paper.

Fukui, Masao, Niels Joachim Gormsen, and Kilian Huber (July 2025). "Sticky Discount Rates". Working paper.

Gabaix, Xavier (2019). "Behavioral Inattention". In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Ed. by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson. Vol. 2. Elsevier. Chap. 4, pp. 261–343.

— (Aug. 2020). "A Behavioral New Keynesian Model". In: *American Economic Review* 110.8, pp. 2271–2327.

Galí, Jordi and Mark Gertler (2007). "Macroeconomic Modeling for Monetary Policy Evaluation". In: *Journal of Economic Perspectives* 21.4, pp. 25–46.

García-Schmidt, Mariana and Michael Woodford (Jan. 2019). "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis". In: *American Economic Review* 109.1, pp. 86–120.

Gaspar, José-Miguel, Massimo Massa, and Pedro Matos (2005). "Shareholder Investment Horizons and the Market for Corporate Control". In: *Journal of Financial Economics* 76.1, pp. 135–165.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (Sept. 2019). "Text as Data". In: *Journal of Economic Literature* 57.3, pp. 535–574.

Gertler, Mark and Simon Gilchrist (1994). "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms". In: *The Quarterly Journal of Economics* 109.2, pp. 309–340.

Goodhead, Robert (2024). "The Economic Impact of Yield Curve Compression: Evidence from Euro Area Forward Guidance and Unconventional Monetary Policy". In: *European Economic Review* 164, 104716.

Goodhead, Robert and Benedikt Kolb (2025). "Monetary Policy Communication Shocks and the Macroeconomy". In: *Economica* 92.365, pp. 173–198.

Gormsen, Niels Joachim and Kilian Huber (June 2025). "Corporate Discount Rates". In: *American Economic Review* 115.6, pp. 2001–2049.

— (2026). "Firms' Perceived Cost of Capital". In: *The Quarterly Journal of Economics*. Forthcoming.

Graham, John R and Campbell R Harvey (May 2001). "The Theory and Practice of Corporate Finance: Evidence from the Field". In: *Journal of Financial Economics* 60.2, pp. 187–243.

Graham, John R. (2022). "Presidential Address: Corporate Finance and Reality". In: *The Journal of Finance* 77.4, pp. 1975–2049.

Gurkaynak, Refet S, Brian Sack, and Eric T Swanson (May 2005). "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements". In: *International Journal of Central Banking* 1(1), pp. 55–93.

Gürkaynak, Refet, Hatice Gökçe Karasoy-Can, and Sang Seok Lee (2022). "Stock Market's Assessment of Monetary Policy Transmission: The Cash Flow Effect". In: *The Journal of Finance* 77.4, pp. 2375–2421.

Gust, Christopher, Edward Herbst, and David López-Salido (Oct. 2022). "Short-Term Planning, Monetary Policy, and Macroeconomic Persistence". In: *American Economic Journal: Macroeconomics* 14.4, pp. 174–209.

Gust, Christopher and David Lopez-Salido (Aug. 2025). "Optimal Monetary Policy with Uncertain Private Sector Foresight". In: *Finance and Economics Discussion Series* 2024-059, pp. 1–32.

Hartley, Jonathan S. and Jackson Mejia (2025). *Smooth Panel Local Projections*. Working Paper. January 8.

Hassan, Tarek A et al. (Nov. 2019). "Firm-Level Political Risk: Measurement and Effects". In: *The Quarterly Journal of Economics* 134.4, pp. 2135–2202.

Hassan, Tarek A et al. (July 2024). "Sources and Transmission of Country Risk". In: *Review of Economic Studies* 91.4, pp. 2307–2346.

Hassan, Tarek A. et al. (Aug. 2025). "Text as Data in Economic Analysis". In: *Journal of Economic Perspectives* 39.3, pp. 193–220.

Hayashi, Fumio (1982). "Tobin's Marginal q and Average q: A Neoclassical Interpretation". In: *Econometrica* 50.1, pp. 213–224.

Ilabaca, Francisco, Greta Meggiorini, and Fabio Milani (Jan. 2020). "Bounded Rationality, Monetary Policy, and Macroeconomic Stability". In: *Economics Letters* 186, p. 108522.

Jamilov, Rustam, Hélène Rey, and Ahmed Tahoun (May 2023). *The Anatomy of Cyber Risk*. Working Paper 206. Institute for New Economic Thinking.

Jarociński, Marek (May 2024). "Estimating the Fed's Unconventional Policy Shocks". In: *Journal of Monetary Economics* 144, p. 103548.

Jarociński, Marek and Peter Karadi (Apr. 2020). "Deconstructing Monetary Policy Surprises— The Role of Information Shocks". In: *American Economic Journal: Macroeconomics* 12.2, pp. 1–43.

Jeenas, Priit (2026). "Firm Balance Sheet Liquidity, Monetary Policy Shocks, and Investment Dynamics". In: *Journal of Political Economy*. Forthcoming.

Jeenas, Priit and Ricardo Lagos (Mar. 2024). "Q -Monetary Transmission". In: *Journal of Political Economy* 132.3, pp. 971–1012.

Jordà, Òscar (Feb. 2005). "Estimation and Inference of Impulse Responses by Local Projections". In: *American Economic Review* 95.1, pp. 161–182.

Jungherr, Joachim et al. (2024). *Corporate Debt Maturity Matters For Monetary Policy*. Working Paper 2024-30. Federal Reserve Bank of San Francisco.

Kolasa, Marcin, Sahil Ravgotra, and Pawel Zabczyk (2025). *Monetary Policy and Exchange Rate Dynamics in a Behavioral Open Economy Model*.

Kroner, T. Niklas (2021). "Firm-Level Uncertainty and the Transmission of Forward Guidance to Investment". In: *SSRN Electronic Journal*.

La'O, Jennifer (2019). *Comment on Monetary Policy Analysis When Planning Horizons Are Finite*.

Lewis, Daniel J, Christos Makridis, and Karel Mertens (2020). *Do Monetary Policy Announcements Shift Household Expectations?* Discussion Paper 14360. Centre for Economic Policy Research.

Lewis, Daniel J. (2025). "Announcement-Specific Decompositions of Unconventional Monetary Policy Shocks and Their Effects". In: *The Review of Economics and Statistics* 107.4, pp. 1086–1103.

Lorenzoni, Guido (2019). "Comment on Monetary Policy Analysis When Planning Horizons Are Finite". In: *NBER Macroeconomics Annual* 33, pp. 67–74.

Lustenhouwer, Joep and Kostas Mavromatis (2025). "The Effects of Fiscal Policy When Planning Horizons Are Finite". In: *Journal of Money, Credit and Banking* 57.2-3, pp. 549–582.

Maqui, Eduardo and Richard Morris (2020). "The Long-Term Effects of the Pandemic: Insights from a Survey of Leading Companies". In: *ECB Economic Bulletin* 8.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (Oct. 2016). "The Power of Forward Guidance Revisited". In: *American Economic Review* 106.10, pp. 3133–3158.

Meggiorini, Greta and Fabio Milani (2026). *Behavioral New Keynesian Models: Learning versus Cognitive Discounting*.

Milbradt, Konstantin and Martin Oehmke (2015). "Maturity Rationing and Collective Short-Termism". In: *Journal of Financial Economics* 118.3, pp. 553–570.

Miranda-Agrippino, Silvia and Giovanni Ricco (July 2021). "The Transmission of Monetary Policy Shocks". In: *American Economic Journal: Macroeconomics* 13.3, pp. 74–107.

Modigliani, Franco and Kalman J. Cohen (1961). *The Role of Anticipations and Plans in Economic Behavior and Their Use in Economic Analysis and Forecasting*. Illinois: University of Illinois.

Na, Seunghoon and Yinxi Xie (Dec. 2025). "Expectations and the UIP Puzzles When Foresight Is Limited". In: *Journal of International Economics* 158, p. 104183.

Nadkarni, Sucheta and Jianhong Chen (2014). "Bridging Yesterday, Today, and Tomorrow: CEO Temporal Focus, Environmental Dynamism, and Rate of New Product Introduction". In: *The Academy of Management Journal* 57.6, pp. 1810–1833.

Nadkarni, Sucheta, Tianxu Chen, and Jianhong Chen (2016). "The Clock Is Ticking! Executive Temporal Depth, Industry Velocity, and Competitive Aggressiveness". In: *Strategic Management Journal* 37.6, pp. 1132–1153.

Nakamura, Emi and Jón Steinsson (Aug. 2018). "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect". In: *The Quarterly Journal of Economics* 133.3, pp. 1283–1330.

National Association of Corporate Directors (2025). *2025 Governance Outlook*. National Association of Corporate Directors.

Olivari, Marco and Markus Schwedeler (2026). "Institutional Ownership, Blockholders, and Attention to the Future". Working paper.

Ottonello, Pablo and Thomas Winberry (2020). "Financial Heterogeneity and the Investment Channel of Monetary Policy". In: *Econometrica* 88.6, pp. 2473–2502.

Poterba, James M. and Lawrence H. Summers (Oct. 1995). "A CEO Survey of U.S. Companies' Time Horizons and Hurdle Rates". In: *Sloan Management Review* 37.1, pp. 43–53.

PwC (2015). *Governing for the Long Term: Looking Down the Road with an Eye on the Rear-View Mirror*. Annual Corporate Directors Survey. PwC.

Sampson, Rachelle C. and Yuan Shi (Jan. 2023). "Are U.S Firms Becoming More Short-term Oriented? Evidence of Shifting Firm Time Horizons from Implied Discount Rates, 1980–2013". In: *Strategic Management Journal* 44.1, pp. 231–263.

Sautner, Zacharias et al. (2023). "Firm-Level Climate Change Exposure". In: *The Journal of Finance* 78.3, pp. 1449–1498.

Selgrad, Julia and Kerry Siani (Sept. 2026). "Monetary Policy and Investment Plans". Working paper.

Shrader, Charles B. et al. (2004). "Planning and Firm Performance: Effects of Multiple Planning Activities and Technology Policy". In: *International Journal of Management and Decision Making* 5.2/3, pp. 171–195.

Song, Wenting and Samuel Stern (Oct. 2025). "Firm Inattention and the Efficacy of Monetary Policy: A Text-Based Approach". In: *Review of Economic Studies* 92.5, pp. 3438–3469.

Souder, David and Philip Bromiley (2012). "Explaining Temporal Orientation: Evidence from the Durability of Firms' Capital Investments". In: *Strategic Management Journal* 33.5, pp. 550–569.

Souder, David and J. Myles Shaver (2010). "Constraints and incentives for making long horizon corporate investments". In: *Strategic Management Journal* 31.12, pp. 1316–1336.

Souder, David et al. (2016). "A Behavioral Understanding of Investment Horizon and Firm Performance". In: *Organization Science* 27.5, pp. 1202–1218.

Stein, Luke C. D. and Elizabeth C. Stone (2013). *The Effect of Uncertainty on Investment, Hiring, and R&D: Causal Evidence from Equity Options*. Working Paper 1649108. SSRN.

Sutherland, Christopher S. (2023). "Forward Guidance and Expectation Formation: A Narrative Approach". In: *Journal of Applied Econometrics* 38.2, pp. 222–241.

Swanson, Eric T. (Mar. 2021). "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets". In: *Journal of Monetary Economics* 118, pp. 32–53.

Von Thadden, Ernst-Ludwig (1995). "Long-Term Contracts, Short-Term Investment and Monitoring". In: *The Review of Economic Studies* 62.4, pp. 557–575.

Weis, Maximilian and Patricia Klarner (Dec. 2022). "A CEO's Future Temporal Depth and Organizational Resilience". In: *Schmalenbach Journal of Business Research* 74.4, pp. 659–693.

Winberry, Thomas et al. (Dec. 2025). "New Keynesian Economics with Household and Firm Heterogeneity". Working paper.

Woodford, Michael (2019). "Monetary Policy Analysis When Planning Horizons Are Finite". In: *NBER Macroeconomics Annual* 33, pp. 1–50.

Woodford, Michael and Yinxi Xie (Jan. 2022). "Fiscal and Monetary Stabilization Policy at the Zero Lower Bound: Consequences of Limited Foresight". In: *Journal of Monetary Economics* 125, pp. 18–35.

Xie, Yinxi (2020). *Fiscal and Monetary Policy Interaction under Limited Foresight*. Working paper. Columbia University.

Zwick, Eric and James Mahon (2017). "Tax Policy and Heterogeneous Investment Behavior". In: *American Economic Review* 107.1, pp. 217–248.

A Detailed Construction of the Planning-Horizon Measure

This appendix describes extraction and conversion in detail, reports the source sentences for the conversion examples, and defines the aggregation into revealed planning horizons. It also records the prompt used in the sentence-level comparison with a large language model. Further details on the extraction rules and conversion procedures are provided in the Online Methodological Note.

A.1 Extraction and Conversion

For every sentence in our corpus, we identify each expression that refers to time and convert it into a *distance*: the number of years ahead of the date on which the sentence was spoken. The measurement is produced by a deterministic, rule-based pipeline that (i) detects candidate time expressions in the raw text, (ii) normalizes each candidate into a canonical interpretation, (iii) resolves overlapping matches in favor of the most specific reading, and (iv) converts the result into a distance in years from the speech date. Because the rules are explicit, every extracted distance can be traced back to the exact phrase and rule that produced it, and the procedure is fully replicable.

Although surface forms vary widely, candidate expressions fall into four families, distinguished by how their distances are computed; Table A.1 reports example expressions from each family together with the distances the pipeline assigns. Expressions excluded from our baseline are labelled "not in baseline." Table A.2 further shows the pipeline's detections in sentences drawn from the transcripts.

- *Calendar-anchored references* contain an explicit calendar anchor: a year ("2027", "full-year 2027"), a dated period ("Q4 2027", "April 2026", "second half of 2027"), a qualified year portion ("early 2027", "mid-2028"), or an exact date ("April 15, 2027"). Conversion measures the distance from the speech date to the referenced period; a calendar window is represented by its midpoint, so "Q4 2027", spoken in early August 2026, spans months 14 to 17 ahead and yields 15.5 months.

- *Speech-relative references* carry a directional or remainder cue that anchors them to the moment of speech: "next quarter", "this year", "the coming quarters", "the next 18 months", "the remainder of the year", "by the end of the decade". Their distance depends on when the sentence was spoken: "the rest of the year" resolves to 8.5 months in mid-April but to 1.5 months in mid-November.

- *Bare durations* name a length of time without anchoring it to any date: "three years" *(not in baseline)*, "a 10-month period", "a 24-month forecast", "a four-year investment program" *(not in baseline)*. These are invariant to the speech date: "four-year program" maps to 48 months whenever it is said. Ranges such as "two to three years" *(not in baseline)* or "an 18- to 24-month transition" target the midpoint of their bounds (30 and 21 months), with the bounds retained for robustness checks.

- *Standalone context-dependent calendar references (not in baseline)* name a calendar location without a year: "January", "March 31", "H2", "the first half of the year". Resolving them

requires both the speech date and the orientation of the sentence. The reference receives its earliest upcoming reading, so that "January" spoken in August 2026 resolves to January 2027, 5.5 months ahead, and it is retained only when the tense of the sentence supports a present- or future-oriented reading: a sentence stating that margins *improved* in January is classified as backward-looking by a subsequent tense analysis and excluded, and when tense cannot be resolved the reference is conservatively treated as not forward-looking. Standalone context-dependent calendar references are excluded from the v1 baseline; the description here documents the broader candidate taxonomy.

A set of rules handles expressions that contain temporal vocabulary but do not specify future dates or durations. Rate units ("shipments per week", "7% per annum") are flagged as non-temporal and excluded from the planning-horizon measure.

The detector further embeds a set of safeguards against false positives. Four-digit numbers are accepted as calendar years only within the range 1900–2099, and numeric tokens are rejected as year candidates when the surrounding context marks them as quantities, for instance when they are followed by a percentage sign ("the mid 90% range") or by a magnitude word such as *million* or *basis points*. Lexically ambiguous tokens are resolved by surface form: lowercase "may" is always read as the modal verb, and only capitalized "May" as the month. Expressions carrying past cues such as "ago" are flagged as backward-looking and excluded from the forward-distance measure. Finally, duration phrases whose head noun signals age or history rather than a plan ("100-year history", "65-year-old patient", "20-year experience") are subsequently rejected by a vocabulary filter, to reduce matches that describe age or history rather than a future duration.

Four conventions apply throughout. First, all distances are computed in months and reported in years by dividing by twelve; sub-monthly references (days, weeks, exact dates) are converted at 30 days per month, so an exact date 257 days ahead yields a distance of 8.6 months, or 0.71 years. Second, whenever a reference denotes an interval rather than a point (a calendar period such as "Q4 2027", a directional anchor such as "next year", or a duration range such as "two to three years"), the distance targets the midpoint of the interval. Applying the midpoint uniformly keeps the interpretation consistent when qualifiers narrow the window: spoken in mid-April 2026, "next year" resolves to 14.5 months, "early next year" and "late next year" to the midpoints of the early and late portions of that year (11.5 and 17.5 months), and "late next year" coincides with the explicitly anchored "late 2027". The exact windows assigned to the "early", "mid", and "late" qualifiers, together with the complete specification of the detection rules, fiscal-calendar alignment, and the disambiguation of two-digit years, are part of the detailed conversion specification. Third, references to periods that have fully elapsed by the speech date, like backward-looking references, are excluded from the measure, since our object of interest is forward orientation. Fourth, when two matched expressions overlap in the text, the pipeline keeps the more specific interpretation: "the rest of 2026" is resolved as a single remainder expression, not as a bare mention of "2026".

Table A.1: Example Time Expressions and Their Converted Distances

Expression	**Family**	**Call date**	**Months**	**Distance (years)**
"Q4 2022"	Calendar-anchored	Nov 11, 2021	12.1	1.01
"full year 2025"	Calendar-anchored	Jan 30, 2025	5.5	0.46
"early 2007"	Calendar-anchored	Feb 1, 2006	14.0	1.17
"late 2017"	Calendar-anchored	Feb 2, 2017	8.2	0.68
"second half of 2013"	Calendar-anchored	Aug 8, 2012	13.7	1.14
"December 31st, 2033"	Calendar-anchored	Mar 14, 2024	119.3	9.94
"next quarter"	Speech-relative	Apr 16, 2002	4.0	0.33
"this year"	Speech-relative	Jan 17, 2022	5.7	0.48
"rest of the year"	Speech-relative	Feb 4, 2021	10.9	0.91
"remainder of the year"	Speech-relative	Jan 30, 2025	11.0	0.92
"coming quarters"	Speech-relative	Oct 29, 2008	9.0	0.75
"next five years"	Speech-relative	Mar 6, 2012	60.0	5.00
"next few years"	Speech-relative	Nov 7, 2013	36.0	3.00
"by the end of the decade"	Speech-relative	Aug 8, 2019	4.7	0.39
"5 years" *(not in baseline)*	Bare duration	Oct 29, 2021	60.0	5.00
"1 or 2 years" *(not in baseline)*	Bare duration	Jan 30, 2018	18.0	1.50
"12-month period"	Bare duration	Apr 27, 2012	12.0	1.00
"7-year contract"	Bare duration	May 8, 2024	84.0	7.00
"12- to 18-month period"	Bare duration	Oct 14, 2008	15.0	1.25
"September" *(not in baseline)*	Context-dependent	Jun 30, 2004	2.5	0.21
"June 30" *(not in baseline)*	Context-dependent	May 6, 2003	1.8	0.15
"H2" *(not in baseline)*	Context-dependent	May 30, 2023	4.0	0.33
"second half of the year" *(not in baseline)*	Context-dependent	Jan 25, 2007	8.2	0.68

Notes: The entries illustrate a broader detector configuration and are pending verification against the v1 baseline; corresponding source sentences appear in Table A.2.

A.1.1 Source sentences for the conversion examples

Table A.2 reports the source sentences for Table A.1. Bold marks the illustrated expressions; superscript NB denotes *(not in baseline)*.

Table A.2: Source Sentences for the Illustrative Conversion Examples

Sentence	Company	Call date
"Yes, just to elaborate on that because you also spoke about sort of the return-to-office dynamic being a positive that you expected to sort of benefit you from, say, **Q4 2022** or so."	Pexip Holding	Nov 11, 2021
"We aim to have a continued high space in cost reductions and have an ambition to target for savings in the range of SEK3.5 billion to SEK4 billion in the **full year 2025**."	Electrolux	Jan 30, 2025
"Additionally, we have budgeted 15 million for our POS rollout in **early 2007**."	Tractor Supply	Feb 1, 2006
"And just one the final question, do you have any plans on returning to capital markets and issuing some Eurobonds **this year**, because I understand you have some $500 million due sometime **late 2017**, so are you planning to tap markets again or –?"	Turkiye Is Bankasi	Feb 2, 2017
"We continue to expect that we will complete the profit line migration process during the **second half of 2013**."	Tangoe	Aug 8, 2012
"However, the amendment extended the maturity date to **December 31st, 2033**, in terms of the amendment are otherwise substantially similar to the original loan agreement."	Montauk Renewables	Mar 14, 2024
"So I believe that that might have more of an effect **next quarter** then it did **this quarter**."	ManpowerGroup	Apr 16, 2002
"And that's why this is a very important reason why we are talking about profit and loss 7% to 9% **this year**, whereas it should be 8% to 10% if Turkey number was normal."	Qatar National Bank	Jan 17, 2022
"But we've got a lot on tap for the **rest of** the **year** to further improve retention and churn well."	GoPro	Feb 4, 2021
"As you have probably seen already, we have yet again raised our cash guidance substantially for the **remainder of** the **year**."	ITM Power	Jan 30, 2025
"Some points to consider as you think about our financial performance over the **coming quarters**."	Viavi Solutions	Oct 29, 2008
"We think the combination of growth and very, very strong cash flows through the **next five years** offers shareholders real value."	Inmarsat	Mar 6, 2012
"Again, I'm very optimistic that we'll have great results with the federal government over the **next few years**."	Gartner	Nov 7, 2013
"I don't think until semiconductor really becomes of a trillion-dollar industry, maybe **by** the **end of** the **decade** or the next **decade** that you'll see a type of stability that let's say there is no cycles."	VAT Group	Aug 8, 2019
"And we show that at $60, even prices well below where we're at now, that we can generate $25 billion of excess cash over **5 years**^{NB}."	Chevron	Oct 29, 2021
"But we intend and the direction, as we clearly stated for all our shareholders, to be in line or to be aware of the direction which we will be getting over the **next few years** or maybe **1 or 2 years**^{NB}."	IDFC	Jan 30, 2018
"That will stretch over at least a **12-month period**, so that's maybe CAD12 million per quarter."	Canam Group	Apr 27, 2012
"The first product is for existing power plants, and that's a **7-year contract**."	New Fortress Energy	May 8, 2024
"So it's a phased approach over a **12- to 18-month period** on an individual facility."	Zep	Oct 14, 2008
"And so I am trying to remember, I think we actually start to produce samples and get the treater qualified in, our current schedule anyway in the **September**^{NB} timeframe."	Park Aerospace	Jun 30, 2004
"We expect it will be resolved by **June 30**^{NB}."	Crown Castle	May 6, 2003
"In Rice, as I mentioned, we will implement a revised pricing strategy as we look to reset our margins and do expect some margin recovery in **H2**^{NB}."	Tiger Brands	May 30, 2023
"We're all seeing great interest in our next generation solutions that should start to contribute to revenues and share gains in the **second half of** the **year**^{NB}."	Microsemi	Jan 25, 2007

A.2 Aggregation

For sentence s in call c by firm i in year t, let m_{itcsr} denote the converted distance in months associated with reference r. After baseline selection, overlap resolution, and the exclusion of backward-looking references, let

$$R^{+}_{itcs} = \{r : m_{itcsr} > 0\}$$

be the retained set. When this set is nonempty, the sentence-level distance in years is

$$d_{itcs} = \frac{1}{12|R^{+}_{itcs}|} \sum_{r \in R^{+}_{itcs}} m_{itcsr}.$$

Each sentence therefore supplies at most one observation, regardless of how many future references it contains. Sentences with no retained positive distance contribute no observation.

Let $\mathcal{S}^{Pr}_{itc}$ denote the prepared-presentation sentences in call c with a nonmissing distance, and let $\mathcal{C}_{it}$ denote firm i's calls in year t. We pool sentence-level observations across these calls and define the firm-year revealed planning horizon as

$$H_{it} = Q_{0.95}\big(\{d_{itcs} : c \in \mathcal{C}_{it},\ s \in \mathcal{S}^{Pr}_{itc}\}\big),$$

where $Q_{0.95}$ is the empirical 95th percentile, with each qualifying sentence contributing one observation. Let $\mathcal{T}_i$ denote the set of years for which firm i has a measured annual horizon. The firm-level measure is

$$H_i = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} H_{it}.$$

Each observed annual horizon receives equal weight. A year with no retained observations has no measured annual horizon and is excluded from $\mathcal{T}_i$; a firm with no measured annual horizons has no firm-level measure.

A call-level summary applies the same percentile to observations within a single call. Quarterly summaries pool sentences across a firm's calls within a quarter. Full-transcript variants also include sentences from the question-and-answer session, and alternative-percentile measures replace the 95th percentile with another percentile.

A.3 Comparison with a Large Language Model

We compare the algorithm's sentence-level classifications with those of Gemini 2.5 Flash-Lite for a sample of 900 sentences from the 2025 validation data. Gemini classifies a sentence as forward-looking when its estimated horizon is positive; the algorithm does so when it retains a future temporal reference. Table A.3 reports the complete comparison for these 900 sentences.

Table A.3: External Comparison Against Gemini 2.5 Flash-Lite

	Gemini 2.5 Flash-Lite	
Algorithm	**False**	**True**
False	6.67% (60)	40.56% (365)
True	0.56% (5)	52.22% (470)

Notes: The comparison contains 900 sentences. Rows are algorithm classifications and columns are Gemini classifications; entries are percentages of the comparison sample, with counts in parentheses. Percentages may not sum to 100 because of rounding. Gemini is positive when its estimated forward horizon is greater than zero. The algorithm is positive when it retains a future temporal reference under the v1 baseline. Treating Gemini as the reference gives precision of 98.9%, recall of 56.3%, accuracy of 58.9%, and F1 of 0.718. These statistics describe agreement within the comparison sample, not accuracy against human ground truth.

Prompt. The sentence-level comparison with Gemini 2.5 Flash-Lite uses the following prompt template for estimating the distance to a referenced period:

```
You are an expert financial analyst. Analyze the temporal reference in this
earnings call sentence.

Context:
- Speech Date: {speech_date}
- Fiscal Quarter: Q{fiscal_period}

Sentence:
"{sentence}"

Task:
Calculate the months from the speech date to the end of the mentioned period.
Return ONLY a JSON object: {"months_estimation": <float_or_null>}
```

Table A.4: Earnings-Call Discussions of Distant Outcomes

Firm and call date	Selected passage
Unisys 25 Feb 2026	*After raising its near-term revenue forecast, Unisys is asked about further adjustments to consumption estimates for 2027–28:* Analyst: "[. . .] but for [2027–28], you just haven't really adjusted that yet [. . .] I mean consumption levels should continue to increase. Is that fair?" Management: "Yeah, look, I would say **it's too far out for us to adjust, multiple year out consumption estimates** [. . .]" *Management adds that some consumption growth is already incorporated in these estimates.*
Xylem 30 Jul 2024	*On future demand from PFAS regulation:* Management: "And it's only going to potentially push out what's already a long phase and timeline of three years to test and five years to mitigate PFAS. [. . .] So while it may stumble a little bit out of the gates, and there's going to be some challenges, we're already seeing challenges in litigation, that's why we didn't bake it into our long-term framework. **We just think it's too far out.**"
AutoNation 22 Jul 2015	*On the long-run effect of car sharing on vehicle sales:* Analyst: "[. . .] are you open to that possibly being a threat that could bring the SAAR well below 14 million or do you think that is still just way too far out to even think about at this point?" Management: "**I think that is way too far away.** [. . .] So that is my point of view and we will see. So far my batting average on these things is pretty good. So we will see."
Calfrac Well Services 29 Jul 2015	*On visibility into future work:* Management: "[. . .] that the way that Calfrac operates is forecast, is in a rolling mode, we normally analyze the next quarter in front of us and the three consecutive quarters. And we're in a cycle that's basically taking us into the 12 to 18 months. So we are very cautious and careful. **And that's why if you try to think that our planning is basically going to give you ideas on what's going to happen in 2016, it is too far away from us from today's point of view.**"
Dream Office REIT 18 Feb 2022	*On financing a proposed development:* Management: "The trickier one is 212 King, but we're – because that's just a huge building, you've got to build it all at once. **We have a partner there already, but it's too far away to actually plan the financials.**"
Godawari Power & Ispat 30 May 2022	*On when higher coal prices would affect production costs:* Analyst: "But then your cargo would have hedged up to a certain point, right? [. . .] So I wanted to know, this higher cost, when it is going to hit us, this higher cost of production." Management: "We are – at the moment, we are very well covered until next 3 months, which is June, July, August. **So September, we still are uncovered, but I think September is too far away to think right now.**"
TJX Companies 17 Aug 2022	*On revising sales forecasts for the second half of the year:* Management: "[. . .] when we went to this more conservative sales for the back half, I would tell you, we did that knowing that we want to be conservative. [. . .] **it's too far out for us to change the numbers right now.** We have – Q4 is where we feel like there's a lot of opportunity for us."

Notes: The table presents illustrative excerpts from firms' earnings calls, including question-and-answer exchanges. The baseline planning-horizon measure uses management presentations and excludes question-and-answer sessions. Italicized context, bold emphasis, and omission marks [. . .] are added by the authors. Firm names link to the source transcripts.

B Additional Facts about Planning Horizons

B.1 State-Space Models of Planning Horizons

Baseline specification. Let $H_{i,t,c}$ be the planning horizon measured from call c of firm i in calendar year t, and let $n_{i,t}$ denote the number of calls observed in that firm-year. The annual outcome is

$$H_{i,t} = \frac{1}{n_{i,t}} \sum_{c=1}^{n_{i,t}} H_{i,t,c}.$$

The estimation sample keeps firm-years with at least four calls and firms with at least three such years. Requiring the 20 firm characteristics listed in Panel A of Table B.1 yields a common sample of 36,315 firm-year observations for 3,759 firms from 2003 through 2025. We estimate the observation equation

$$H_{i,t} \mid \Theta \sim \mathcal{N}\left(\lambda_t + \gamma_{s(i)} + \alpha_i + h_{i,t}, \frac{\sigma_\varepsilon^2}{n_{i,t}}\right).$$

The components λ_t, $\gamma_{s(i)}$, and α_i represent calendar-year variation, persistent sector differences, and persistent firm differences within sectors, where $s(i)$ denotes the sector containing firm i. The value of $\gamma_{s(i)}$ is common to firms belonging to the same sector. The observation variance is inversely proportional to $n_{i,t}$, as implied by averaging call-level errors with variance σ_ε^2. The likelihood also incorporates the within-firm-year dispersion of the call-level horizons, which provides direct information about this variance.

The latent within-firm state follows

$$h_{i,1} \sim \mathcal{N}(0, \sigma_h^2), \qquad h_{i,t} \mid h_{i,t-1} \sim \mathcal{N}\left(\rho_{s(i)} h_{i,t-1}, \sigma_h^2(1-\rho_{s(i)}^2)\right), \qquad 0 < \rho_{s(i)} < 1.$$

This parameterization makes σ_h^2 the stationary variance of the latent state. The sector-specific innovation variance is therefore

$$\sigma_{\eta,s}^2 = \sigma_h^2(1-\rho_s^2).$$

For observations separated by Δt years, the transition uses $\rho_{s(i)}^{\Delta t}$ and the corresponding stationary transition variance.

Priors and posterior estimation. The year coefficients receive Gaussian priors with standard deviation 20. The sector, firm, and state standard deviations receive half-t_3 priors with scale 5, while the call-level measurement-error standard deviation receives a half-t_3 prior with scale 10. Sector-specific persistence follows the hierarchy

$$\operatorname{logit}(\rho_s) = \mu_\rho + \tau_\rho u_s, \qquad u_s \sim \mathcal{N}(0,1),$$

$$\mu_\rho \sim \mathcal{N}\left(\operatorname{logit}(0.77), 1\right), \qquad \tau_\rho \sim \text{half-t}_3(0,1).$$

The sector and firm effects are centered so that $\gamma_{s(i)}$ captures differences across sectors and α_i captures differences among firms within sectors.

We sample the posterior using a blocked Metropolis-within-Gibbs algorithm. The year, sector, and firm effects have conditional Gaussian updates. Complete paths of $h_{i,t}$ are drawn by forward filtering and backward sampling. Persistence is updated from the state-collapsed Kalman likelihood, and the scale parameters are sampled using the inverse-gamma representation of the half-t priors. The baseline estimates use four chains, each with 1,200 warm-up and 2,400 retained draws. Variance components and variance shares are calculated for every posterior draw before their posterior summaries are computed.

Variance decomposition. For posterior draw m, annual measurement noise is

$$V_\varepsilon^{(m)} = \frac{1}{M} \sum_{(i,t)\in\mathcal{O}} \frac{\sigma_\varepsilon^{2(m)}}{n_{i,t}},$$

where $\mathcal{O}$ is the set of observed firm-years and $M = |\mathcal{O}|$. Because the panel is finite and unbalanced, part of the persistent state $h_{i,t}$ survives when a firm's observations are averaged over time. Define

$$w_{i,t} = \frac{n_{i,t}}{\sum_{r\in T_i} n_{i,r}}$$

and

$$D_i^{(m)} = \sigma_h^{2(m)} \sum_{t\in T_i} \sum_{r\in T_i} w_{i,t} w_{i,r} \left(\rho_{s(i)}^{(m)}\right)^{|t-r|}, \qquad D^{(m)} = \frac{1}{N} \sum_i D_i^{(m)}.$$

The genuine within-firm component is

$$V_{\text{within}}^{(m)} = \sigma_h^{2(m)} - D^{(m)}.$$

The persistent amount $D^{(m)}$ is allocated between the sector and firm components using measurement-error-corrected between-sector and within-sector ANOVA shares. The components then satisfy

$$V_{\text{total}}^{(m)} = V_\lambda^{(m)} + V_\gamma^{(m)} + V_\alpha^{(m)} + V_{\text{within}}^{(m)} + V_\varepsilon^{(m)}.$$

The four variance components in Table 2 divide genuine variation, $V_{\text{total}} - V_\varepsilon$, into time, sector, between-firm, and within-firm components. The final column reports the ratio of genuine variation to total observed variation.

Alternative state processes. The baseline and time-varying-sector specifications in Table 2 are nested in the general model

$$H_{i,t} = \lambda_t + \gamma_{s(i),t} + \alpha_i + h_{i,t} + \varepsilon_{i,t},$$

$$\gamma_{s(i),t} = \gamma_{s(i)} + q_{s(i),t}, \qquad h_{i,t} = \rho_{s(i)} h_{i,t-1} + \eta_{i,t}, \qquad q_{s(i),t} = \phi q_{s(i),t-1} + \zeta_{s(i),t},$$

where $s(i)$ denotes the sector containing firm i. The permanent component $\gamma_{s(i)}$ is common to firms in the same sector, while $q_{s(i),t}$ captures movements shared by those firms during year t.

The baseline imposes $q_{s,t} = 0$ and allows ρ_s to vary across sectors. The specification with time-varying sector effects introduces $q_{s,t}$, with sector-specific persistence in the firm-specific state. The sector state follows a stationary AR(1) process, with

$$\text{logit}(\phi) \sim \mathcal{N}\left(\text{logit}(0.5), 1.5^2\right), \qquad \sigma_q \sim \text{half-t}_3(0, 5).$$

At each posterior draw, $q_{s,t}$ is orthogonalized with respect to the calendar-year and permanent sector components under the firm-year sample weights. Its reported variance therefore captures sector-specific movements over time rather than common aggregate movements or permanent differences across sectors.

Observed characteristics. Panel A of Table B.1 examines how much genuine variation can be associated with observed firm characteristics. Let $x_{i,t}$ contain the 20 firm characteristics listed in Panel A of Table B.1. For a firm in sector s, each characteristic is decomposed as

$$x_{i,t} = \bar{x}_s + (\bar{x}_i - \bar{x}_s) + (x_{i,t} - \bar{x}_i).$$

These terms represent differences across sectors, persistent differences among firms within sectors, and changes within firms. They enter the corresponding levels of the model:

$$\gamma_{s(i)} = \bar{x}'_{s(i)} \beta_S + g_{s(i)}, \qquad \alpha_i = \left(\bar{x}_i - \bar{x}_{s(i)}\right)' \beta_B + a_i, \tag{25}$$

$$h_{i,t} = \rho_s h_{i,t-1} + (x_{i,t-1} - \bar{x}_i)' \beta_W + \eta_{i,t}.$$

Estimation follows the same blocked Bayesian procedure as the baseline model, augmented with Gaussian updates for β_S, β_B, and β_W. The characteristic coefficients have independent $\mathcal{N}(0, 2.5^2)$ priors. Within-firm characteristics enter with a one-year lag and affect the transition of the latent state. Because the three fitted pathways can be correlated, their variance contributions are allocated using the Shapley decomposition of the variance of their sum. The specification with characteristics in Panel A uses the same firm-year observations as the baseline in Table 2.

Table B.1: Firm characteristics, CEO effects, and alternative sector definitions

Specification	$\widehat{\rho}$	Time	Between firms: Sector	Between firms: Firm	CEO	Within firm variation	% explained
Panel A. Observed firm characteristics							
With firm characteristics (20 variables)							
Total	0.66	1.5	13.2	51.1	–	34.2	48.1
Characteristics	–	–	5.5	2.9	–	0.3	–
Residual	–	–	7.7	48.2	–	34.0	–
Panel B. CEO effects and characteristics: same sample							
Baseline (CEO sample)	0.76	2.6	23.3	35.2	–	39.0	47.5
CEO effects	0.75	2.6	18.2	39.4	3.3	36.4	47.4
With firm and CEO characteristics and CEO effects							
Total	0.74	2.6	15.9	43.3	3.3	34.9	49.1
Characteristics	–	–	7.4	2.8	0.1	0.7	–
Residual	–	–	8.5	40.5	3.2	34.2	–
Panel C. Alternative sector definitions							
67 2-digit SIC sectors	0.84	0.9	26.5	51.8	–	20.7	59.0
251 3-digit SIC sectors	0.85	0.6	33.5	48.2	–	17.6	66.0
388 4-digit SIC sectors	0.86	0.5	38.4	44.8	–	16.2	68.8

Notes: Panel A uses the Table 2 sample of 3,759 firms and 36,315 firm-years over 2003–2025. The diagnostic CEO estimates in Panel B use the same 1,918 firms, 21,107 firm-years, and 4,079 CEOs over 2003–2024. Panels A and B use the baseline sector classification. Panel C uses the Table 2 sample with the SIC classifications shown. Total combines characteristic and residual contributions. Panel B additionally uses CEO age, tenure, log total compensation, long-term compensation share, chair status, and sex, including annual and CEO-profile contributions; CEO effects are centered within firms. The smaller CEO sample and the fact that most CEOs are observed at only one firm limit our ability to separate persistent managerial influences from firm characteristics. Persistence varies across sectors; $\widehat{\rho}$ is the firm-weighted posterior mean of the sector-specific persistence parameters. Time, sector, firm, CEO (where included), and within-firm entries are percentages of variation net of measurement error. Sector and firm components include the contribution of finite-sample averages of the time-varying state. The final column reports the percentage of observed annual variation attributable to the underlying horizon after separating measurement error. The sample requires at least four calls per firm-year and three usable years per firm. Missing sector classifications enter an Other group, excluded from the sector counts shown.

Firm characteristics: Firm age (log); Firm size (log assets); Cash / assets; Frequency of accounting losses; Dividend-paying frequency; Book leverage; Asset growth; Gross profitability; Operating cash flow / assets; Capital expenditure / assets; Intangibles / assets; Net debt issuance / lagged assets; Inventory days (log); Sales growth; Acquisitions / assets; Tangible-asset share; Operating profitability; Tobin's Q; Repurchases / lagged assets; Employment growth. Full variable definitions and data sources are provided in Appendix C.2.

B.2 Within-Industry Comparisons

Table B.2: Within-industry dispersion in firm planning horizons, 2002 to 2025

	Shorter-horizon firm		Longer-horizon firm		Gap
Industry or comparison	Firm	Years	Firm	Years	Years
Panel A. Pairs with the same reported four-digit SIC code					
Heavy trucks	Navistar	2.43	PACCAR	12.19	9.75
Equity REITs	Equity Residential	5.63	Realty Income	13.50	7.86
Hotel operators	Marriott	2.51	Hilton	8.44	5.93
Automakers	Tesla	2.68	Ford	8.05	5.37
Water utilities	American States Water	8.30	California Water Service	12.95	4.65
Copper mining	Southern Copper	4.16	Freeport-McMoRan	8.46	4.30
Quick-service restaurants	Chipotle	2.78	Domino's	6.63	3.85
General retailers	Dollar General	1.99	Target	5.78	3.79
Full-service restaurants	Cheesecake Factory	1.69	Darden Restaurants	5.38	3.69
Pharmaceutical preparations	Zoetis	1.68	Pfizer	5.34	3.66
Telecommunications	Verizon	2.62	AT&T	6.24	3.61
Electric utilities	American Electric Power	4.87	Southern Co.	8.48	3.61
Nonalcoholic beverages	Monster Beverage	1.04	Coca-Cola	4.63	3.59
Theme parks	SeaWorld	1.94	Six Flags	5.50	3.56
Alternative asset managers	Ares	4.72	Blackstone	8.17	3.46
Supermarkets	Kroger	1.76	Albertsons	4.75	2.99
Apparel	Under Armour	3.37	Levi Strauss	6.35	2.98
Cosmetics	Revlon	1.72	Estée Lauder	4.66	2.94
Footwear	Crocs	1.81	Nike	4.60	2.79
Education providers	Duolingo	3.14	Coursera	5.85	2.71
Integrated oil and refining	Valero	1.85	Chevron	4.52	2.68
Home-improvement retail	Lowe's	5.59	Home Depot	7.77	2.17
Airlines	Spirit Airlines	2.32	American Airlines	4.01	1.70
Business software	Oracle	2.57	Salesforce	3.51	0.94
Ride-hailing	Uber	1.45	Lyft	2.16	0.71
Panel B. Other economically related pairs					
Household products[N]	Colgate-Palmolive	1.62	Procter & Gamble	6.96	5.34
Aerospace and defense[E]	General Dynamics	2.29	Lockheed Martin	7.50	5.21
Regional banks[N]	Fifth Third	2.97	Huntington	6.87	3.90
Parcel delivery[P]	FedEx	4.13	UPS	6.56	2.43
Large banks[N]	JPMorgan	2.24	Wells Fargo	4.35	2.10
Payment networks[N]	Mastercard	2.57	Visa	3.79	1.23
Retail pharmacies[P]	CVS Health	2.14	Walgreens Boots Alliance	3.27	1.14

Notes: Entries are averages of call-level planning horizons within each firm over 2002 to 2025, measured in years. The gap is the longer-horizon value minus the shorter-horizon value. Rows are ordered within panels by the unrounded gap. Panel A pairs share the same reported four-digit SIC code. In Panel B, [N] indicates the same detailed NAICS industry, [E] the same sector, and [P] manually selected business-model peers across formal classifications. The comparisons are descriptive, not causal. Source: authors' calculations.

B.3 Explanatory Power of Observed Characteristics

Table B.3: Explaining differences in firms' planning horizons

	Common sample: 804 firms		Largest available sample	
	Each group separately	All groups jointly	Each group separately	
	(1)	(2)	(3)	
Characteristics	R^2 (%)	R^2 (%)	R^2 (%)	Firms
Firm characteristics	7.9	5.0	6.5	2,037
Management	7.3	4.4	5.3	1,699
Board	6.1	3.4	3.6	1,436
Ownership	6.6	4.3	3.8	2,677
All characteristics	–	17.1	–	–

Notes: Entries report the percentage of within-sector variation in average planning horizons explained by the characteristics. Columns (1) and (2) use the same 804 firms. Column (2) allocates the joint regression's R^2 across groups using a Shapley decomposition; these contributions sum to the total in the last row, subject to rounding. "Each group separately" in columns (1) and (3) includes only the indicated characteristic group. Column (3) uses each group's largest available sample, with firm counts shown alongside. All regressions include TRBC business-sector fixed effects. The outcome and characteristics are firm averages over 2004–2020. The outcome and continuous characteristics are winsorized at the first and ninety-ninth percentiles.

Firm characteristics (31): Firm size (log assets); Firm age (log); Asset maturity; Share of debt due after two years; Cash-to-assets ratio; Book leverage; Tobin's Q; Operating profitability; Sales growth; Intangible assets / total assets; Capital expenditure / lagged assets; Acquisition expenditure / lagged assets; Tangible-asset share; Capital per employee (log); Inventory days (log); Sales-growth volatility; Dividend-paying frequency; Operating cash flow / lagged assets; Current assets / current liabilities; Interest coverage; Share repurchases / lagged assets; Net equity issuance / lagged assets; Net debt issuance / lagged assets; Employment growth; Asset growth; Gross profitability; Profitability volatility; Cash-flow volatility; Frequency of accounting losses; Operating noncash working capital / assets; Discount rate.

Management (31): CEO long-term pay share; CEO total compensation (log); CEO stock awards / total pay; CEO option awards / total pay; CEO equity ownership; CEO age; CEO tenure; Female CEO; CEO board membership; CEO is board chair; CEO share of total reported executive pay; CEO unvested equity / annual pay; CFO long-term pay share; CFO total compensation (log); CFO stock awards / total pay; CFO option awards / total pay; CFO equity ownership; CFO age; Female CFO; CFO board membership; CFO unvested equity / annual pay; Other executives' average age; Dispersion of other executives' ages; Other executives' average pay (log); Other executives' long-term pay share; Other executives' stock awards / total pay; Other executives' option awards / total pay; Other executives' average equity ownership; Female share of other executives; Share of other executives serving on the board; Dispersion of other executives' log pay. Long-term pay denotes compensation outside salary and bonus. Other executives exclude CEOs.

Board (14): Number of directors; Average director tenure; Independent-director share; Outside public-company board seats per director; Average director age; Dispersion of director ages; Female-director share; Share of directors recorded as financial experts; Share of directors on the audit committee; Share of directors on the compensation committee; Share of directors on the nomination committee; Share of directors on the governance committee; Presence of a non-executive chair; Presence of a lead director.

Ownership (16): Equity share held by long-term institutional investors; Long-term investors' share of institutional holdings; Institutional investors' portfolio turnover; Number of institutional investors (log); Equity held by 13D/G-reporting owners; Total institutional equity ownership; Dual-class indicator; Fraction of institutional investors classified as long-term; Largest owner's equity share; Equity in ownership blocks above 5%; Equity in ownership blocks above 10%; Equity share attributed to Forms 3/4/5 filings; Number of all recorded owners (log); Largest owner's voting share; Votes in voting blocks above 5%; Votes in voting blocks above 10%.

C Compustat Variables and Sample Construction

This appendix documents the financial variables used in Sections 3 and 4, the institutional-ownership measure in Figure 1, and the construction and cleaning of the firm-quarter panel.

C.1 Compustat Inputs

The quarterly panel is indexed by the standard Compustat firm identifier (`gvkey`) and the calendar-quarter date derived from `datacqtr`. Industry classification comes from SIC and NAICS codes, and the sample is restricted to U.S.-incorporated firms using `fic`.

The capital and investment measures are built from net property, plant, and equipment and capital expenditures (`ppentq` and `capxy`). The baseline firm controls use total assets, current assets, sales, and the investment deflator (`atq`, `actq`, `saleq`, and `ipd`). The broader balance-sheet characteristics used in robustness exercises additionally draw on cash holdings, current and long-term debt, current and total liabilities, and operating income (`cheq`, `dlcq`, `dlttq`, `lctq`, `ltq`, and `oiadpq`).

Sales growth and size classifications rely on quarterly sales (`saleq`). Additional robustness characteristics use interest expense, cost of goods sold, SG&A expenses, preferred dividends, common equity, receivables, payables, share price, shares outstanding, deferred taxes, dividends, fiscal-quarter information, and acquisition expenditures (`xintq`, `cogsq`, `xsgaq`, `dd1q`, `ceqq`, `rectq`, `apq`, `prccq`, `cshoq`, `txditcq`, `dvpq`, `fqtr`, and `aqcy`). The panel used in Section 4 also carries quarterly monetary shocks and macro controls merged by quarter. Transcript-based horizon measures are merged separately by firm-quarter and by firm average; the average horizon measures used there are therefore constant within firm.

C.2 Data Sources and Variable Definitions

This appendix describes the data used in Figures 1 and 4, as well as Tables 2, B.1, and B.3. A firm's revealed planning horizon is the average of its call-level horizons, each measured as the 95th percentile of future time distances in presentation sentences. Figure 1 covers 2002–2025, with institutional ownership available for 2004–2020. Figure 4 uses 2004–2020 for all characteristics.

Firm characteristics. Compustat provides annual balance-sheet, income-statement, and employment data. Firm size is log total assets, and firm age is $\log(1 + \text{age})$, counting the first Compustat year as age one. Liquidity, leverage, and asset intangibility are cash and short-term investments, total debt, and intangible assets, respectively, divided by total assets. Tobin's Q is total assets plus market equity minus book equity, divided by total assets. Operating profitability, R&D intensity, and investment intensity are operating income before depreciation, R&D expenditure, and capital expenditure, respectively, divided by lagged assets. Sales growth is the annual percentage change in sales, and capital intensity is net property, plant, and equipment (PPE) per employee.

Asset and debt maturity, production, and discount rates. Following Hubert de Fraisse (2024), fixed asset maturity is net property, plant, and equipment divided by depreciation excluding amortization, $\texttt{ppent}/(\texttt{dp} - \texttt{am})$. Missing amortization is set to zero, and annual observations with nonpositive denominators do not contribute to the firm average. Figure 1 uses averages over available years in 2002–2025. Weighted asset maturity combines fixed asset maturity with a one-year maturity for current assets, using their respective asset-value shares as weights. Debt maturity is the share of total debt due after two years. Figure 1 uses the share of total debt due after five years. The maturity measures use Compustat. The average period of production is 365 times inventories divided by cost of goods sold, following Antràs and Tubdenov (2025). Table B.3 uses the predicted investment hurdle rate from Gormsen and Huber (2025), averaged across available years. Figure 4 uses the annual predicted rate divided by total assets before averaging across years.

Management. ExecuComp provides executive compensation, equity ownership, and ages. CEO and CFO long-term compensation shares measure the fraction of total pay outside salary and bonus. CEO equity ownership is the fraction of outstanding shares held by the CEO, excluding options. Management age dispersion is the standard deviation of ages among reported executives other than the CEO.

Boards. ISS/RiskMetrics Directors provides annual board rosters. Board size is the number of directors, and the female-director share is the fraction who are female. Interlocking directorships measure directors' average number of other public-company board positions. Director tenure is their average years of service on the board.

Ownership. We use Corporate Ownership Data from Amel-Zadeh et al. (2022), which combine SEC Forms 13F, 13D/G, and 3/4/5. Long-term institutional ownership is the share of outstanding equity held by institutions with three-year portfolio turnover, averaged over four quarters, of at most 35%, following Derrien, Kecskés, and Thesmar (2013). Investor turnover is the holdings-weighted average of institutions' four-quarter average portfolio churn rates, following Gaspar, Massa, and Matos (2005). Ownership breadth is the number of institutional investors. Blockholder ownership is the fraction of equity held in individual stakes above 10%.

Annual panel. Tables 2 and B.1 use a firm-year panel. Each annual horizon is the average of the firm's call-level horizons in that year. In specifications with firm characteristics, the annual characteristics are decomposed into sector means, differences in firm means within sectors, and annual deviations from firm means. The annual deviations enter the transition of the underlying horizon with a one-year lag, as described in Appendix B.1. The panel therefore preserves both differences across firms and changes within firms.

Cross-sectional samples. Figure 4 and Table B.3 use a cross-section with one observation per firm over 2004–2020. Each firm's horizon is the average of its call-level horizons over that period.

Table C.1: Summary statistics for firm-level regression variables

Variable	Mean	Median	St. dev.	Obs.	Corr. with H_i	Corr. with log assets
Quarterly log capital growth	0.023	0.008	0.055	207,832	-0.017	-0.029
Log real assets	5.984	5.990	2.111	207,836	0.141	–
Sales growth	0.026	0.019	0.199	207,836	-0.000	0.001
Share current assets	0.522	0.525	0.240	203,118	-0.076	-0.428
Leverage	0.241	0.181	0.357	200,550	0.024	0.046
Liquidity	0.195	0.103	0.216	207,768	-0.044	-0.334
Firm age	16.688	13.000	11.292	207,836	0.065	0.420
Tobin's Q	2.432	1.669	3.435	203,060	-0.010	-0.230
Large-firm indicator	0.879	1.000	0.327	207,836	0.043	0.523

Notes: The sample imposes the baseline Stein-Stone regression window used in the heterogeneity code, including the shock-availability requirement and the 2019Q1 endpoint. The reported mean, median, and standard-deviation columns are medians of the corresponding quarterly cross-sectional moments. Correlations are medians of quarterly cross-sectional correlations. The large-firm indicator is the Gertler-Gilchrist size classification used in the size-heterogeneity robustness. Observation counts vary with variable availability.

Characteristics use equal-year averages across available annual observations in years with observed horizons; quarterly ownership measures are first averaged within years. These firm averages enter regressions with sector fixed effects. In Figure 1, characteristics are averaged across available annual observations within that figure's sample period.

Transformations. Asset maturity, capital intensity, production period, and ownership breadth enter as $\log(1 + \text{firm mean})$. Figure 4 winsorizes horizons and characteristics at the firm-level 1st and 99th percentiles. In Figure 1, firm-average fixed asset maturity is bounded below at zero and transformed as $\log(1+\text{fixed asset maturity})$ before winsorization. Both horizons and characteristics are winsorized at the firm-level 1st and 99th percentiles and standardized to mean zero and unit standard deviation. In Figure 4, characteristics are standardized and horizons are expressed in years.

C.3 Descriptive Statistics

Tables C.1 and C.2 summarize the firm-level objects and shock series used in the heterogeneity regressions in Section 4 under the current baseline code window. The firm-level table reports medians over quarterly cross-sectional moments rather than moments from a single pooled cross-section. The shock table isolates the quarterly moments of the two Jarociński (2024) factors that enter the baseline multi-shock specification.

C.4 Cleaning and Panel Construction

We begin from quarterly Compustat fundamentals and convert the raw date field `datacqtr` into a calendar-quarter index `dateq`. We verify uniqueness first at the `gvkey-datadate` level and then at the `gvkey-dateq` level. When duplicate firm-quarter observations remain, the final update (`updq==3`) is retained; if a duplicate still remains after that step, observations with missing sales are dropped.

Table C.2: Summary statistics for the normalized quarterly Jarociński shock factors

Shock	Description	Mean	Median	St. dev.	Min	Max	Quarters
MP	$U1$ current-rate and near-term policy-news factor.	-0.22	0.00	0.89	-4.15	2.35	113
FG	$U2$ Odyssean forward-guidance factor.	-0.02	-0.04	0.37	-0.91	1.11	113

Notes: The table reports quarterly sums of the rate-impact-normalized Jarociński (2024) announcement-day $U1$ and $U2$ factors over the 1991Q1–2019Q1 window used by the current baseline local-projection code. Entries are normalized shock units: one unit of *MP* is scaled to a 25 basis-point impact response in quarterly-average annualized FEDFUNDS, and one unit of *FG* is scaled to a 25 basis-point impact response in quarterly-average annualized GS2. In the original high-frequency Jarociński units, these normalized shocks correspond to about 11.1 basis points of $U1$ and 13.8 basis points of $U2$, respectively. Positive realizations are contractionary. In the regressions, these shocks are kept in the same rate-impact units before interacting with the firm-level horizon measure.

We then exclude financial firms (SIC 6000–6799), utilities (SIC 4900–4999), and firms with `fic` different from `"USA"`.

After sorting the data by firm and quarter, we declare the Compustat panel and use `tsfill` to insert missing quarters within firm histories. This places every firm on a regular quarterly grid, so leads, lags, and cumulative outcomes are defined on a consistent quarterly timeline. We then merge quarterly monetary shocks and quarterly macro controls. The resulting panel contains both quarter-specific transcript-based horizon measures and firm-level average horizon measures; the latter are the fixed-within-firm objects used in Section 4.

C.5 Constructed Variables Used in the Heterogeneity Analysis

The outcome in the baseline local projections is 100 times cumulative log growth in the firm's capital stock from quarter $t-1$ to quarter $t+h$, that is,

$$\Delta \log K_{i,t+h} = 100\,(\log K_{i,t+h} - \log K_{i,t-1})\,.$$

These cumulative outcomes are constructed horizon by horizon for the Stein-Stone capital measure.

Baseline capital measure: Stein-Stone. Our baseline measure follows Stein and Stone (2013). Nominal capital evolves according to

$$K_{it}^{SS} = (1-\delta)K_{i,t-1}^{SS}\frac{P_t}{P_{t-1}} + I_{it}^{SS}, \qquad K_{it}^{R,SS} = \frac{K_{it}^{SS}}{P_t}.$$

The construction uses the following components.

- K_{i0}^{SS} is the initial value taken from Compustat net property, plant, and equipment, `ppentq`. Before the recursion is applied, short interior gaps in `ppentq` are linearly interpolated when the missing run is no longer than four quarters and is bounded on both sides by observed values.
- I_{it}^{SS} is quarterized capital expenditure from Compustat `capxy`. In quarter 1, the quarterly flow equals reported `capxy`. In quarters 2–4, the quarterly flow is `capxy_t - capxy_t-1` whenever

the two observations are consecutive quarters.

- $\delta = 1 - (1 - 0.10)^{1/4}$ is the quarterly depreciation rate implied by a 10 percent annual depreciation assumption.
- P_t is the producer price index for capital goods, specifically the Producer Price Index by Commodity: Final Demand: Private Capital Equipment.
- The ratio P_t/P_{t-1} converts lagged nominal capital into current prices before current nominal investment is added. Real capital is then obtained by dividing nominal capital by P_t.

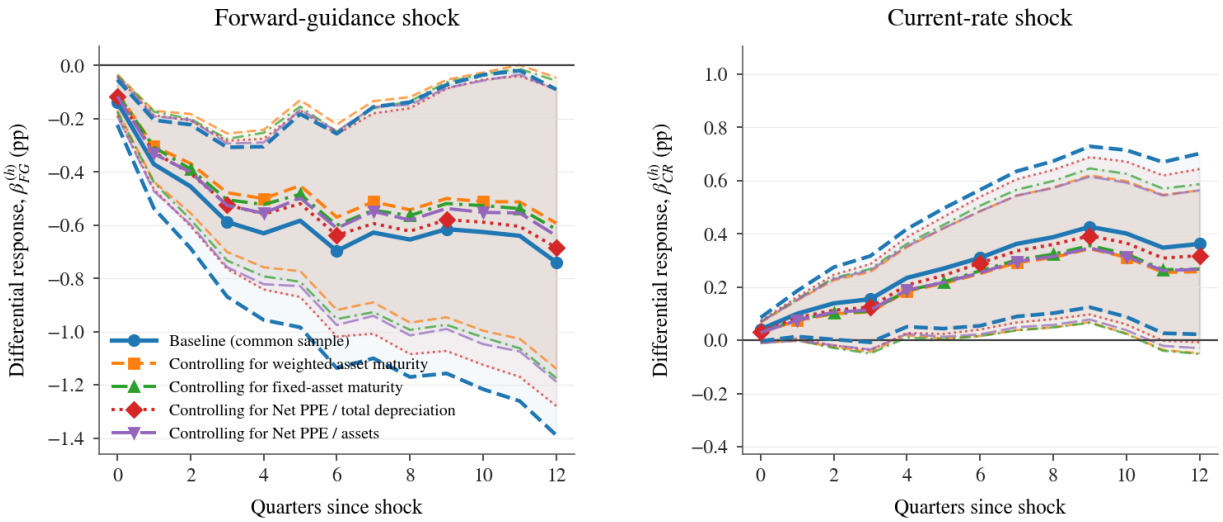
The baseline controls in Section 4 are lagged log real assets, lagged sales growth, and the lagged share of current assets in total assets. Real assets and real sales are deflated by `ipd`, and sales growth is measured as $\log(\text{sales}_{i,t-1}) - \log(\text{sales}_{i,t-2})$. The broader robustness set includes leverage, liquidity, firm age, and the Gertler-Gilchrist size classification based on the previous 10 years of real sales. In these specifications, leverage is lagged debt over lagged assets, liquidity is lagged cash holdings over lagged assets, and age is measured as years since the firm's first appearance in the quarterly Compustat panel. Average period of production follows Antràs and Tubdenov (2025) and is measured as inventories divided by cost of goods sold.

C.6 Final Sample Restrictions

We first drop observations with a missing baseline average horizon measure. We then apply the following filters. First, firm-quarters with acquisition expenditures above 5 percent of assets are removed, where quarterly acquisitions are constructed from year-to-date `aqcy`. Second, observations with non-positive capital or non-positive assets are excluded. Third, investment rates are trimmed at the 0.5th and 99.5th percentiles of the baseline distribution. Fourth, observations with sales growth outside $[-1, 1]$, negative sales, or negative cash holdings are dropped. Because lagged controls enter the regressions, the quarter immediately following a flagged acquisition, sales-growth outlier, negative-sales observation, or negative-cash observation is also removed.

We additionally require an investment spell of at least 10 quarters. Although the underlying Stein-Stone panel extends through 2025Q3 and the merged Swanson series extends through 2019Q2, the local-projection code for Section 4 uses observations only through 2019Q1. After these restrictions are imposed, the continuous horizon and balance-sheet variables entering interaction terms are winsorized at the 0.5 percent tails and standardized; some balance-sheet variables are also demeaned within firm before they are interacted with shocks.

Figure D.1: Investment responses controlling for interactions with asset maturity and composition



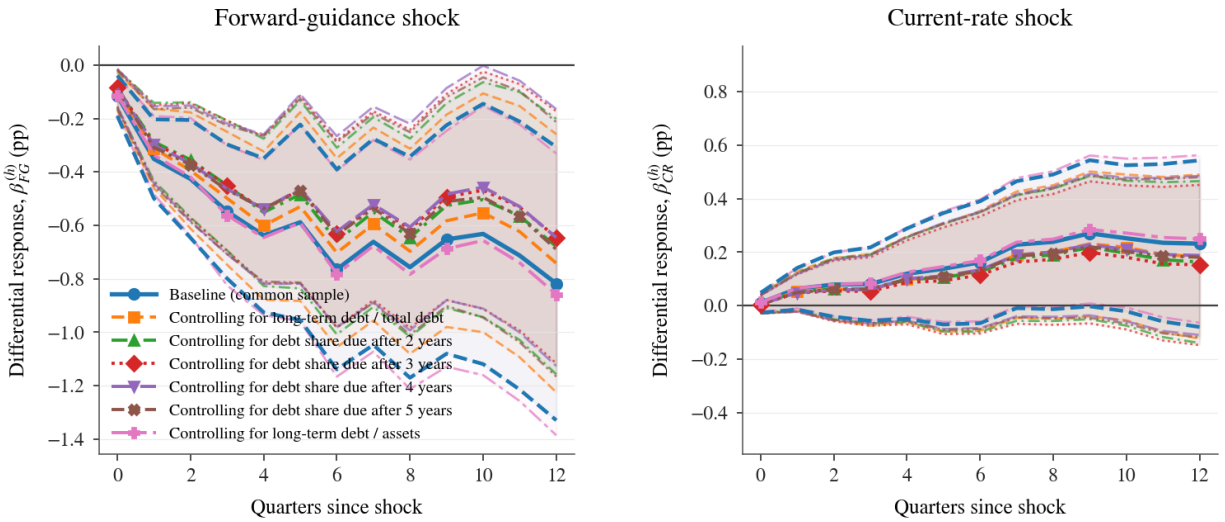
Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Each asset measure enters separately in levels and interacted with the shock and lagged GDP growth. Annual controls are matched from the preceding fiscal year. All lines, including the baseline, use the same observations within each panel and response quarter, requiring all four asset measures to be available. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

D Additional Robustness

Asset maturity and composition. We examine whether asset maturity and composition account for the investment-response differences associated with planning horizons. We add four measures separately: weighted asset maturity, fixed-asset maturity, an alternative fixed-asset maturity measure, and asset tangibility. Fixed-asset maturity is net property, plant, and equipment divided by depreciation excluding amortization; the alternative uses total depreciation and amortization. Weighted maturity combines the fixed-asset measure with a one-year maturity for current assets, weighted by their shares of the combined value of these assets. Tangibility is net property, plant, and equipment divided by total assets. The three maturity measures enter in logs.

Each control enters in levels and interacted with the monetary shock and lagged GDP growth. We estimate responses to current-rate and forward-guidance shocks separately and use the same sample across controls within each shock and response quarter. Annual controls are matched from the preceding fiscal year, winsorized at the 0.5 percent tails, and standardized. Figure D.1 shows that the controls reduce the estimated differences but preserve their signs.

Figure D.2: Investment responses controlling for interactions with debt maturity and financing structure

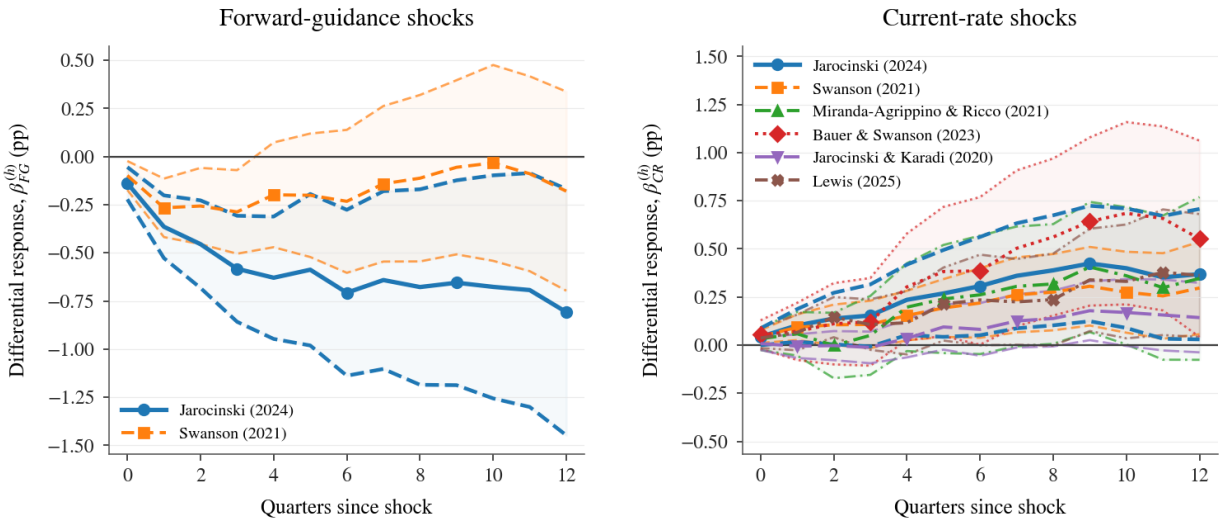


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Each debt measure enters separately in levels and interacted with the shock and lagged GDP growth. Annual controls are matched from the preceding fiscal year. All lines, including the baseline, use the same observations within each panel and response quarter, requiring all six debt measures to be available. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Debt maturity and financing structure. We examine whether debt maturity or reliance on long-term debt accounts for the relationship between planning horizons and investment responses. The maturity measures are the shares of total debt due after one, two, three, four, and five years. The share due after one year equals long-term debt divided by total debt. A separate measure divides long-term debt by total assets and therefore captures financing structure. Each control enters separately in levels and interacted with the monetary shock and lagged GDP growth, following the procedure used for the asset controls in Figure D.1.

Figure D.2 shows positive current-rate and negative forward-guidance point estimates under each control. The current-rate estimates are imprecise on this common sample, including in the regression without an additional debt control. The comparison therefore distinguishes the consequences of using a smaller sample from those of controlling for debt maturity.

Figure D.3: Investment responses to alternative monetary policy shocks

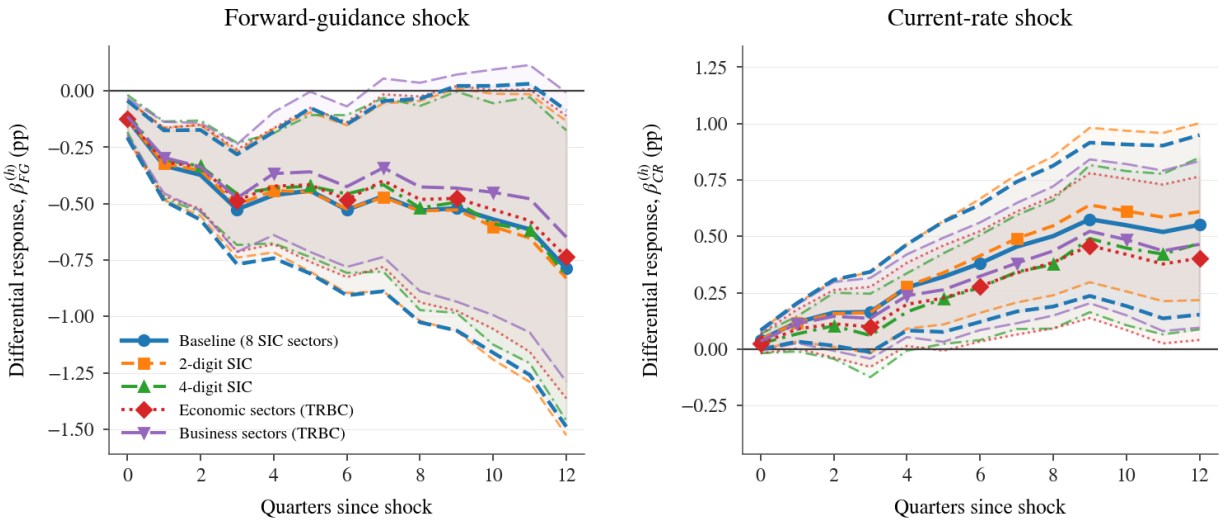


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Each line uses the shock series indicated in the legend; samples vary with shock availability. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Alternative monetary shocks. We compare the monetary surprises of Jarociński (2024) with alternative measures of current-rate and forward-guidance news. The current-rate comparisons use Swanson (2021), Miranda-Agrippino and Ricco (2021), Bauer and Swanson (2023b), Jarociński and Karadi (2020), and Lewis (2025). The forward-guidance comparison uses Swanson (2021). Each shock enters a separate regression, including the Jarociński series used as the comparator. The regressions use the baseline controls and fixed effects, with sample coverage determined by the availability of each shock.

Figure D.3 shows that the alternative current-rate estimates are generally positive, although their magnitudes and precision differ. Both forward-guidance measures produce negative point estimates. Swanson's estimates are smaller in magnitude and imprecise over most of the response window.

Figure D.4: Investment responses under alternative sector classifications

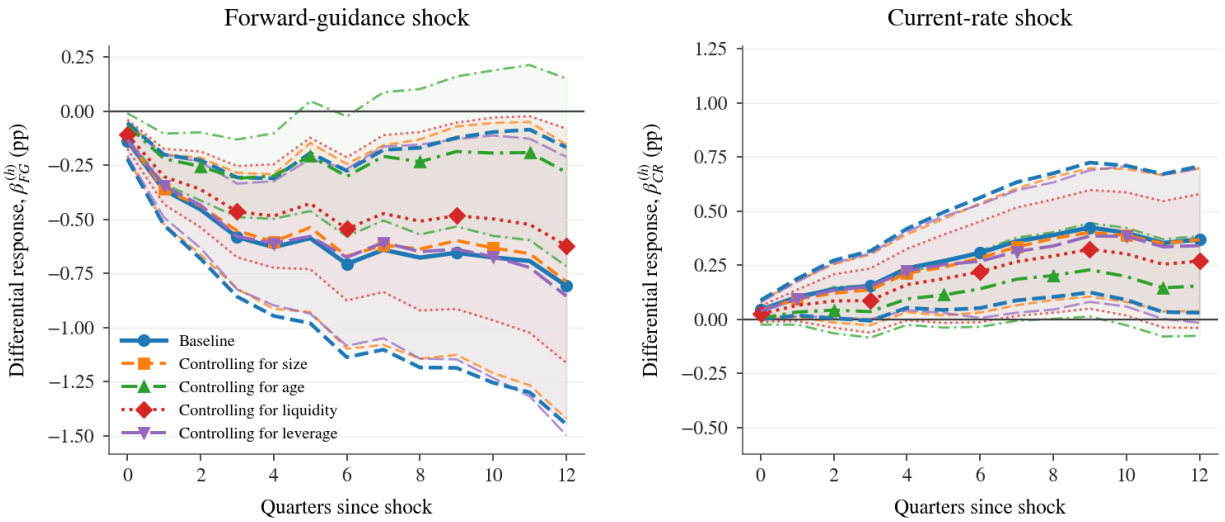


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Both horizon–shock interactions enter jointly with the baseline controls and firm fixed effects. Sector-by-quarter fixed effects use the classification indicated in the legend. The baseline uses eight broad SIC sectors, following Ottonello and Winberry (2020). All lines use the same observations at each response quarter. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Sector definitions. We assess whether the results depend on comparing firms within broad sectors. We estimate the specification with both shocks using the baseline eight SIC sectors, two-digit and four-digit SIC industries, and the economic and business sectors of the Refinitiv Business Classification supplied with the earnings-call data. At each response quarter, all five specifications use the same observations after requiring every classification and removing singleton observations under each set of fixed effects.

Figure D.4 shows negative forward-guidance and positive current-rate point estimates across these classifications. The magnitudes and precision vary. For example, the one-year current-rate coefficient is 0.27 percentage points under the broad-sector classification and 0.16 under four-digit SIC; the latter's 90 percent confidence interval includes zero.

Figure D.5: Investment responses controlling for interactions with firm averages of size, age, liquidity, and leverage



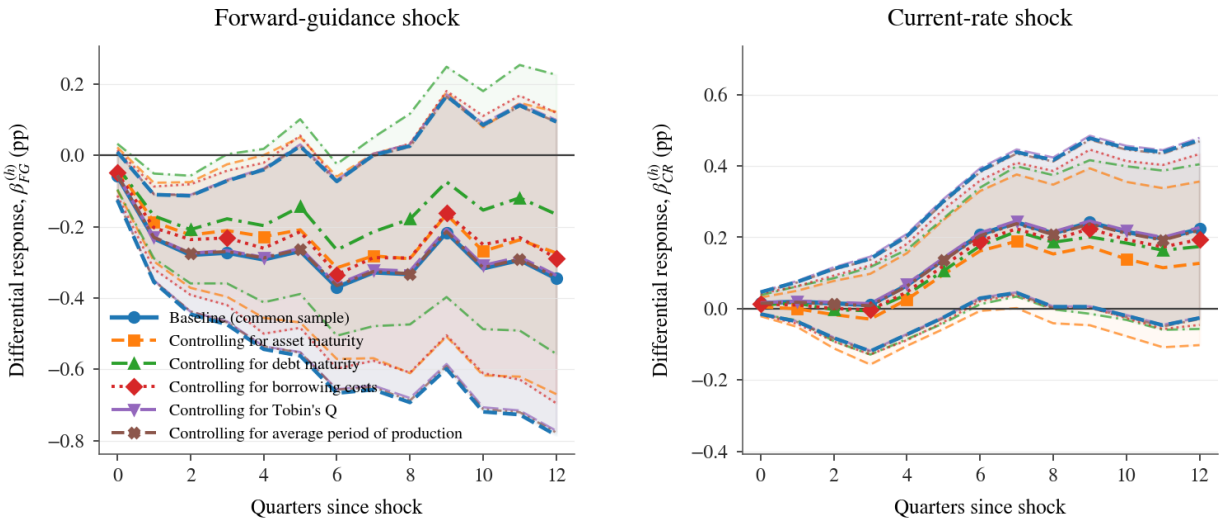
Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Each firm-average characteristic is interacted with the shock and lagged GDP growth. Averages use available observations over 1991Q1–2019Q1 and are computed before selecting the estimation sample for each response quarter. Size and age use averages of the category indicators in the specifications of Figure 6; liquidity and leverage use averages of their winsorized levels before within-firm demeaning. Samples match the corresponding time-varying specifications at each response quarter and vary with control availability. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Fixed firm characteristics. We examine whether the investment-response differences depend on allowing other firm characteristics to vary over time. We construct firm averages from available observations over 1991Q1–2019Q1, before imposing restrictions specific to each response quarter. Each average is interacted with the monetary shock and lagged GDP growth; its level is absorbed by firm fixed effects. Shocks are estimated separately with the baseline controls and fixed effects.

For Figure D.5, size and age are averages of the category indicators used in Figure 6. Size distinguishes firms below the 30th percentile of ten-year average sales from other firms; age distinguishes firms younger than 15 years, aged 15–50, and older than 50. Liquidity and leverage use averages of their winsorized levels before within-firm demeaning. Each regression uses the same observations as its corresponding time-varying specification at each response quarter, so samples can differ across controls.

The firm-average controls generally preserve the response pattern. Average age produces the largest reduction in both differences; its forward-guidance estimate is statistically significant at one year, while its current-rate estimate is not.

Figure D.6: Investment responses controlling for interactions with firm averages of asset and debt maturity, borrowing costs, Tobin's Q, and production period

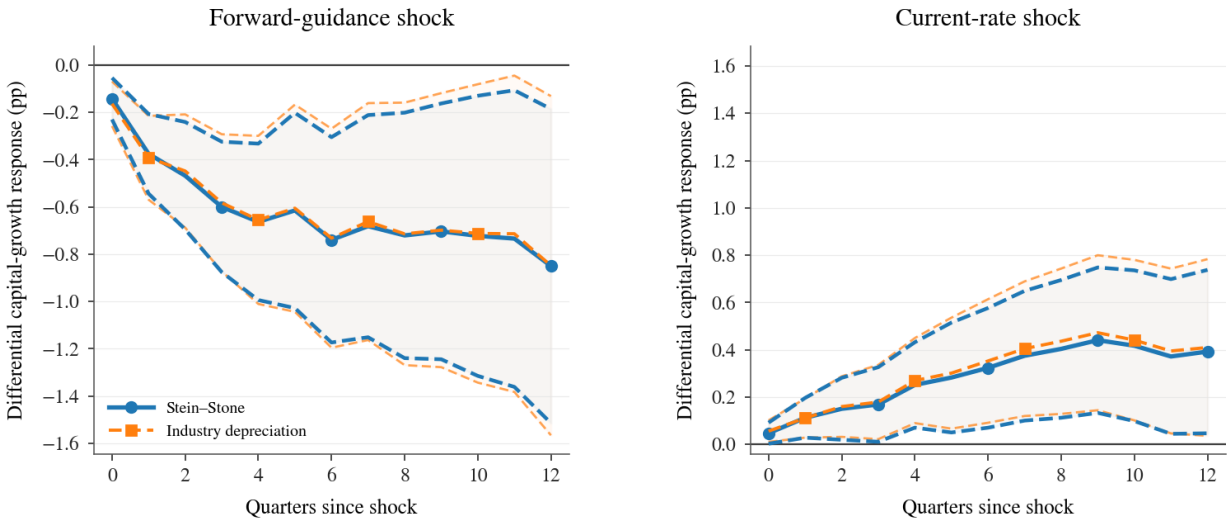


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each shock is estimated separately using the baseline controls and fixed effects. Each firm-average characteristic is interacted with the shock and lagged GDP growth. Averages use available observations over 1991Q1–2019Q1 and are computed before selecting the estimation sample for each response quarter. All controls use the transformations in Figure 7; debt maturity is the share of debt due after one year. All lines use the same observations within each panel and response quarter, matching the corresponding time-varying specifications. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

For Figure D.6, we use firm averages of asset maturity, debt maturity, borrowing costs, Tobin's Q, and production period. Debt maturity is the share of debt due after one year. All controls follow the transformations in the time-varying specifications of Figure 7, with winsorization and standardization applied before averaging. All lines use the common sample of that comparison at each response quarter.

The firm-average controls preserve the response pattern at one year, although the forward-guidance evidence becomes less precise when controlling for asset or debt maturity. These controls reduce the forward-guidance difference from -0.29 to -0.23 and -0.20 percentage points, respectively, and their 90 percent confidence intervals include zero. The forward-guidance estimates under the other controls are statistically significant. None of the current-rate estimates is statistically distinguishable from zero at one year.

Figure D.7: Investment responses with industry-specific depreciation

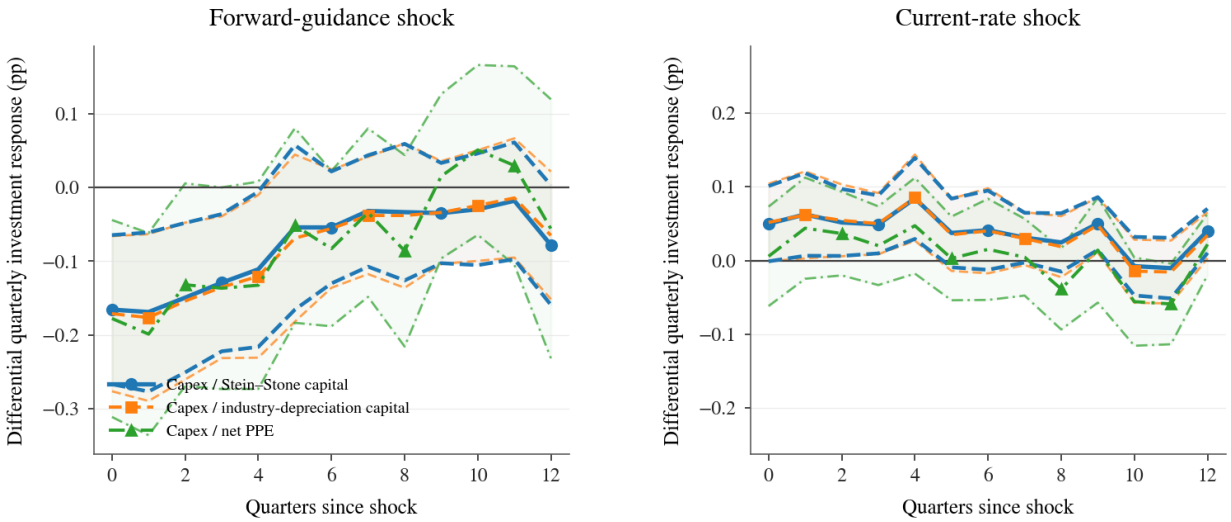


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Both horizon–shock interactions enter jointly with the baseline controls and fixed effects. Capital is constructed by perpetual inventory using either a fixed 10 percent annual depreciation rate (Stein–Stone) or industry-quarter depreciation rates. Both lines use the same observations at each response quarter. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Alternative capital construction. We examine whether the planning-horizon response differences depend on the depreciation assumption used to construct capital. We compare the baseline Stein–Stone series, which assumes 10 percent annual depreciation, with a series using industry-quarter depreciation rates. Both use the same capital expenditures and capital-goods price index. The regressions include both horizon–shock interactions with the baseline controls and fixed effects, using identical observations at each response quarter.

Figure D.7 shows nearly overlapping response profiles. At one year, the forward-guidance coefficients are -0.66 and -0.65, while the current-rate coefficients are 0.25 and 0.27, for Stein–Stone and industry-specific depreciation, respectively. All four pointwise 90 percent confidence intervals exclude zero. The estimated response differences are therefore similar when depreciation varies across industries and quarters.

Figure D.8: Investment responses using alternative capital-expenditure rates

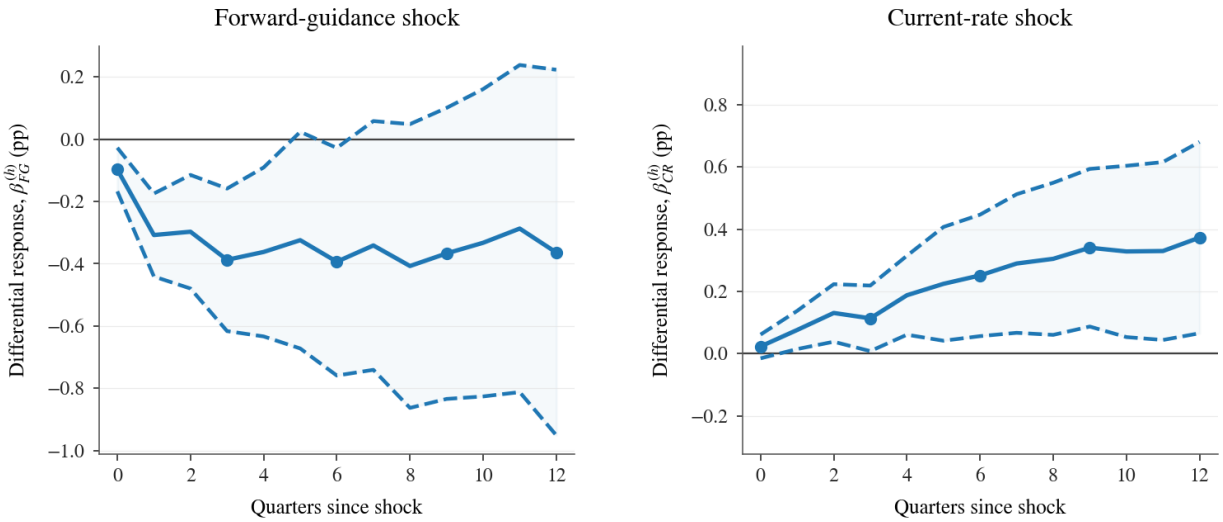


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. The outcome is $100 I_{i,t+h}/K_{i,t+h-1}$, using the denominator indicated in the legend. The Stein–Stone and industry-depreciation ratios use real capex and capital deflated by the capital-goods producer price index; the net-PPE ratio uses nominal capex and book net property, plant, and equipment. Each rate is trimmed at 0.1 percent in each tail using fixed cutoffs. Both horizon–shock interactions enter jointly with the baseline controls and fixed effects. All lines use the same observations at each response quarter. Coefficients measure percentage-point differences in quarterly investment rates per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Alternative investment rates. We also examine quarterly capital expenditure divided by the previous quarter's capital stock, using Stein–Stone capital, capital constructed with industry-quarter depreciation, and book net property, plant, and equipment. These outcomes measure investment rates in the response quarter rather than cumulative capital growth. The two constructed-capital ratios use real values; the net-PPE ratio uses nominal values. We trim 0.1 percent from each tail using fixed cutoffs for each measure and estimate the baseline specification with both shocks on identical observations.

Figure D.8 shows similar responses for the two constructed-capital ratios: at one year, forward-guidance coefficients are approximately -0.11 to -0.12 and current-rate coefficients are about 0.085, with pointwise 90 percent intervals excluding zero. These quarterly response differences generally fade beyond one year, which can coexist with persistent cumulative capital differences. The net-PPE results are less consistent: both one-year intervals include zero, and some later estimates change sign. At eleven quarters, its current-rate coefficient is negative and its pointwise 90 percent interval excludes zero.

Figure D.9: Investment responses with aggregate controls

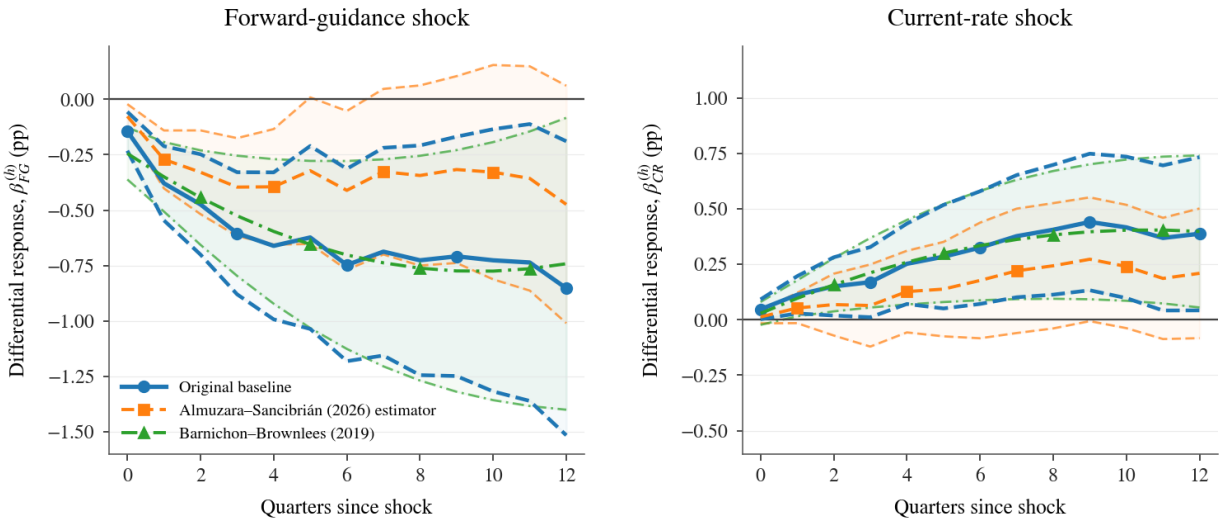


Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Both horizon–shock interactions enter jointly with the baseline firm controls, firm fixed effects, and sector-specific seasonal effects. Aggregate controls include four lags each of real GDP growth, CPI inflation, and unemployment, together with indicators for 2007Q4–2009Q2 and 2019Q1. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Aggregate controls. We assess sensitivity to how the regressions account for aggregate conditions. This specification includes four lags each of real GDP growth, CPI inflation, and the unemployment rate, together with indicators for 2007Q4–2009Q2 and 2019Q1. Firm fixed effects account for persistent differences in capital growth, while sector-specific indicators for the four calendar quarters account for seasonal patterns. The regression also includes both monetary shocks directly, both horizon–shock interactions, the horizon–lagged-GDP interaction, the baseline firm controls, and the fiscal-year-end indicator. Thus, common economic conditions are captured by observed aggregate variables and seasonal effects rather than unrestricted sector-by-quarter effects.

Figure D.9 shows positive current-rate and negative forward-guidance point estimates. The forward-guidance difference is smaller than in the baseline, and its confidence interval includes zero at several longer response horizons.

Figure D.10: Investment responses under alternative estimation and inference methods



Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. The figure compares the baseline, the one-lag Almuzara–Sancibrián specification, and smooth local projections. Both horizon–shock interactions enter jointly. The spline uses the baseline sample; the one-lag specification requires lagged observations. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals. Baseline and spline intervals use standard errors clustered by firm and quarter; spline intervals condition on the selected penalty and exclude smoothing-selection uncertainty and regularization bias. One-lag intervals use quarter clustering, HC2 adjustment, and Imbens–Kolesár degrees of freedom.

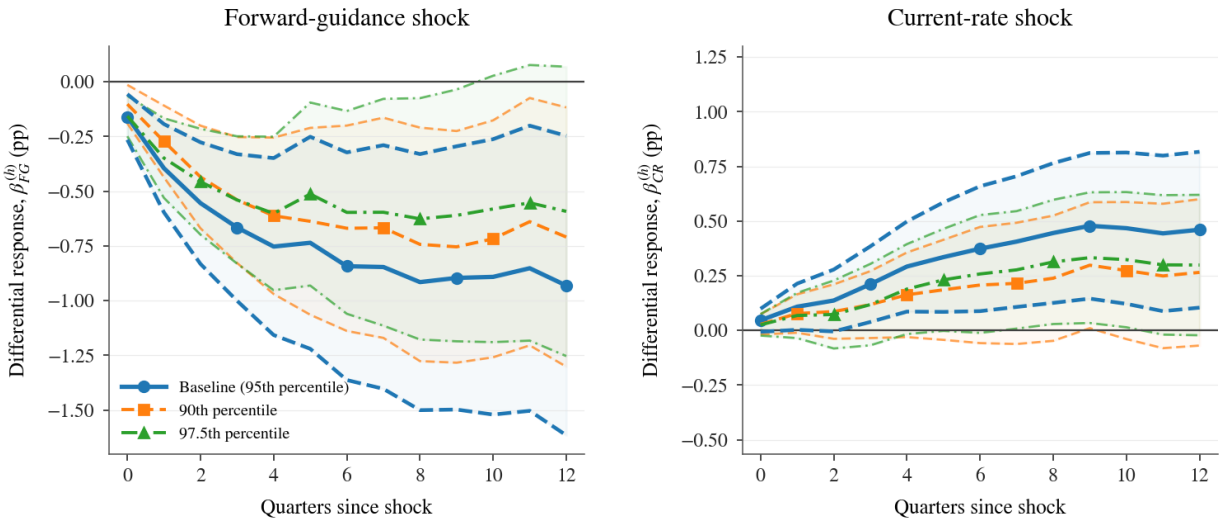
Dynamic specification and inference. We apply the procedure of Almuzara and Sancibrián (2024), using their August 2026 revision, to the specification with both shocks. The specification adds one lag of capital growth and one lag of each horizon–shock interaction. The regression includes the baseline controls and fixed effects. Standard errors are clustered by quarter.[24]

We also estimate smooth local projections following Barnichon and Brownlees (2019), using the panel extension of Hartley and Mejia (2025). This approach smooths the response coefficients across quarters using the baseline controls, fixed effects, and estimation samples. Its confidence intervals use firm and quarter clustering and condition on the selected smoothing penalty; they exclude uncertainty from penalty selection and regularization bias.

Figure D.10 shows smaller estimated differences under the lag-augmented specification. At one year, the current-rate coefficient is 0.13 percentage points, with a 90 percent interval of $[-0.06, 0.31]$; the forward-guidance coefficient is -0.39, with an interval of $[-0.66, -0.13]$. The current-rate intervals include zero throughout the response window. The smooth local projections closely track the baseline responses.

[24]The implementation extends the paper's balanced-panel refinement to our unbalanced panel with sector-by-quarter fixed effects.

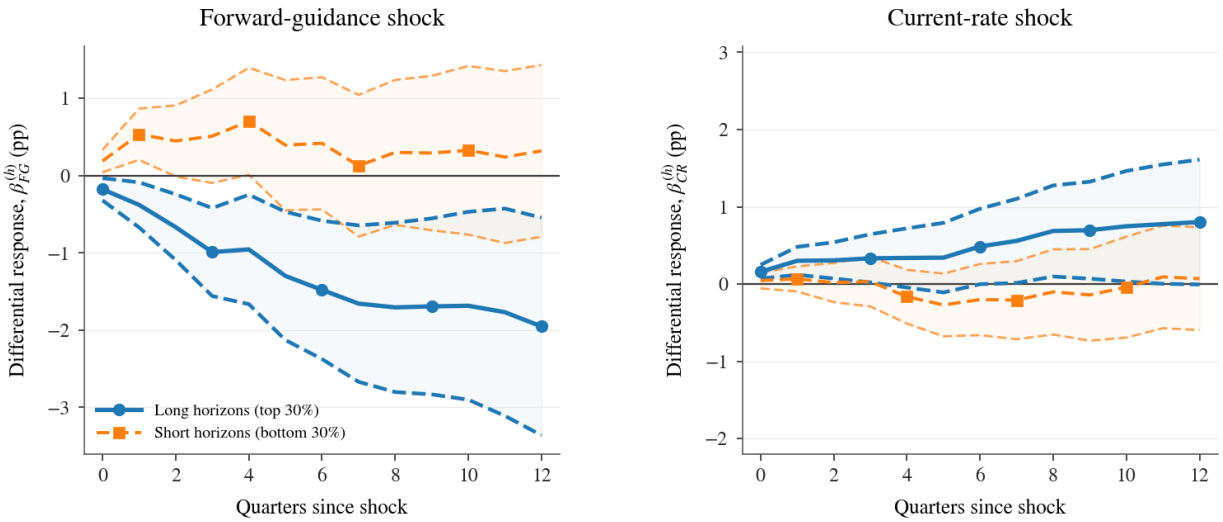
Figure D.11: Investment responses using alternative planning-horizon percentiles



Notes: The panels plot planning-horizon interaction coefficients for forward-guidance (left) and current-rate (right) shocks. Each line uses the horizon percentile indicated in the legend. Both horizon–shock interactions enter jointly with the baseline controls and fixed effects. All lines use the same observations at each response quarter among firms with at least 15 earnings calls containing retained future references and 200 such references. Coefficients measure percentage-point differences in cumulative real-capital growth per one-standard-deviation difference in log planning horizon after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Alternative horizon percentiles. We compare planning horizons measured by the 90th, 95th, and 97.5th percentiles of retained future distances. Each percentile enters a separate baseline regression with both horizon–shock interactions. To ensure comparable coverage, we require firms to have at least 15 earnings calls containing retained future references and at least 200 such references. All three specifications use the same observations at each response quarter. Figure D.11 shows negative forward-guidance and positive current-rate estimates for all three percentiles. The direction of the response differences therefore does not depend on choosing the 95th percentile, although their precision varies.

Figure D.12: Investment responses by planning-horizon group

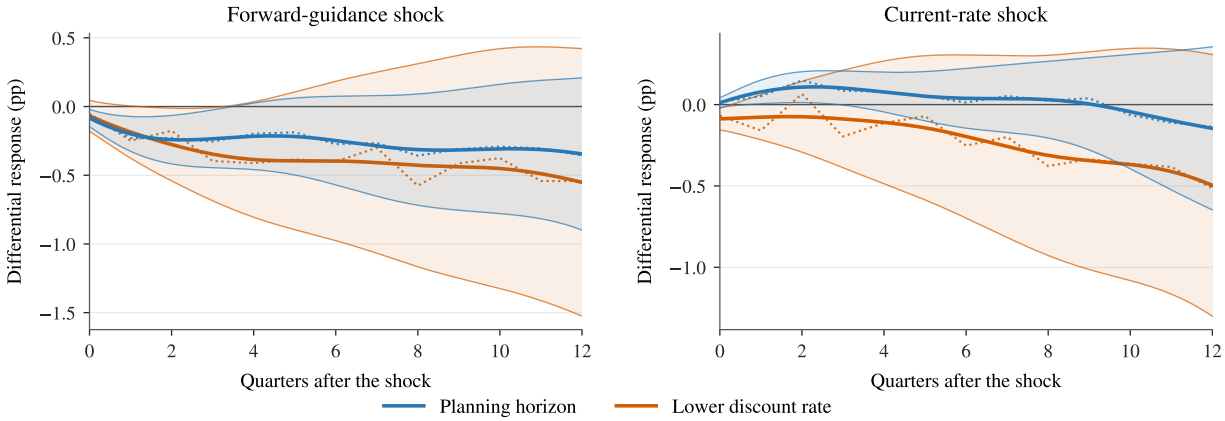


Notes: The panels plot planning-horizon group interaction coefficients for forward-guidance (left) and current-rate (right) shocks. The regression is $\Delta_h \log K_{it} = \alpha_i^{(h)} + \gamma_{j(i),t}^{(h)} + \beta_{FG,T}^{(h)} D_i^T FG_t + \beta_{FG,B}^{(h)} D_i^B FG_t + \beta_{CR,T}^{(h)} D_i^T CR_t + \beta_{CR,B}^{(h)} D_i^B CR_t + \mathcal{C}_{it}^{(h)} + \varepsilon_{it}^{(h)}$. The indicators D_i^T and D_i^B identify firms at or above the 70th percentile and below the 30th percentile of the pooled firm-quarter horizon distribution, respectively. The middle group is the reference category and is included in the sample. All four shock interactions enter jointly. The term $\mathcal{C}_{it}^{(h)}$ contains the baseline firm controls, the fiscal-year-end indicator, and interactions of both group indicators with lagged GDP growth. Each coefficient measures a percentage-point difference in cumulative real-capital growth relative to the middle group after a 25-basis-point tightening in the shock's target rate. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter.

Groups of firms with different horizons. We replace the continuous horizon interactions with indicators for short- and long-horizon firms to assess sensitivity to the linear specification. The cutoffs are the 30th and 70th percentiles of planning horizons in the firm-quarter estimation panel. Firms below the lower cutoff form the short-horizon group, and firms at or above the upper cutoff form the long-horizon group. The middle group is the reference category and is included in the estimation sample.

Both group indicators are interacted with the current-rate shock, the forward-guidance shock, and lagged GDP growth. The specification includes the baseline firm controls, firm and sector-by-quarter fixed effects, and fiscal-year-end indicator. Figure D.12 shows that long-horizon firms cut investment more after forward-guidance tightening and less after current-rate tightening than firms in the middle group. Short-horizon firms cut investment less after forward-guidance tightening than firms in the middle group. These comparisons support the direction of the continuous-horizon results without imposing a linear relationship.

Figure D.13: Planning horizons, discount rates, and monetary transmission



Notes: Differences in cumulative capital-growth responses after a 25-basis-point tightening, associated with a one-standard-deviation higher log planning horizon (blue) or lower firm-average predicted hurdle rate (orange). For ease of comparison, the hurdle rate is multiplied by −1, so that higher values indicate lower discount rates. Solid lines are smooth panel LP estimates following Barnichon and Brownlees (2019) and the panel extension of Hartley and Mejia (2025), using cubic splines and a quadratic penalty target. The penalty is chosen to allow six effective response parameters per path. Dotted lines are the corresponding ordinary LP estimates. Shaded bands show pointwise 90 percent confidence intervals, with standard errors clustered by firm and quarter. Both series use available observations over the common period 2002Q3–2019Q1, with the rate series restricted to rate-covered observations. At one year, the samples contain 135,314 observations from 4,079 firms and 85,797 observations from 2,812 firms, respectively. Source: Gormsen and Huber (2025) for predicted hurdle rates.

Discount rates. The discount-rate point estimates do not reproduce the baseline contrast between investment responses to current-rate and forward-guidance tightening. In the smaller sample covered by the predicted hurdle rates of Gormsen and Huber (2025), the specification including the rate control yields statistically significant coefficients with the baseline signs at one response quarter for each shock, although their magnitudes are smaller. Most attenuation, however, occurs under the sample restriction: adding the rate control produces little further change relative to the baseline estimated on exactly the same observations.

A more direct comparison asks whether discount rates themselves reproduce the investment-response pattern associated with planning horizons. Figure D.13 reports separate local projections using either firms' average planning horizons or average predicted discount rates over the common period 2002–2019. Longer horizons and lower discount rates are both associated with stronger initial contractions after forward-guidance tightening. However, their initial current-rate point estimates have opposite signs: longer horizons indicate weaker contractions, whereas lower discount rates indicate stronger contractions. The discount-rate point estimates are consistent with the direct discounting channel: holding returns on capital fixed, lower discount rates strengthen investment responses to both types of tightening.

The confidence intervals for the smooth estimates in Figure D.13 condition on the chosen smoothing penalty and do not account for penalty selection or smoothing bias.

E Model Appendix: Derivations

This appendix collects the derivations and proofs behind Section 5. Notation follows the main text: bare lowercase letters denote log deviations of realized objects from the zero-inflation steady state, and perceived objects carry the horizon superscript (H) and the origin label t. Where a derivation works inside a perceived path with a fixed origin, we abbreviate

$$x_j^{(H)} \equiv x_{t+j|t}^{(H)}, \qquad 0 \le j \le H, \tag{26}$$

so that j measures distance from the origin and $H - j$ is the remaining horizon at date $t + j$. Announced paths are exogenous and common to all perceived paths and carry no perceived labels.

E.1 Agents' Problems

This appendix derives the optimality conditions of the model's agents, in the order in which Section 5.2 introduces them and ending with the capital producers of Section 5.1. The producer's plan is linearized in Appendix E.2. Throughout, MC_t denotes the level of real marginal cost and mc_t its log deviation.

E.1.1 Household

The household takes the paths of the real wage W_t, the gross real return R_t, and the transfers Ω_t as given, holds rational expectations, and solves

$$\max_{\{C_t, N_t, B_t\}_{t\ge 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} \right] \tag{27}$$

subject to the date-by-date budget constraint

$$C_t + \frac{B_t}{R_t} = W_t N_t + B_{t-1} + \Omega_t, \qquad t \ge 0, \tag{28}$$

with B_{-1} given, and to the no-Ponzi condition

$$\lim_{T\to\infty} \mathbb{E}_0 \left[\left(\prod_{m=0}^{T} \frac{1}{R_m} \right) B_T \right] \ge 0. \tag{29}$$

With $\beta^t \lambda_t$ the multiplier on (28), so that λ_t is measured in units of date-t utility, the conditions for consumption, labor, and bonds are

$$\lambda_t = C_t^{-\sigma}, \qquad \chi N_t^{\varphi} = \lambda_t W_t, \qquad \frac{\lambda_t}{R_t} = \beta\, \mathbb{E}_t \lambda_{t+1}. \tag{30}$$

Eliminating λ_t between the first two gives the labor-supply condition $\chi N_t^{\varphi} = C_t^{-\sigma} W_t$ of Section 5.2.

Dividing the bond condition in (30) by λ_t gives the Euler equation

$$C_t^{-\sigma} = \beta R_t \mathbb{E}_t\left[C_{t+1}^{-\sigma}\right], \tag{31}$$

where R_t is known at date t.

At an optimum the household also satisfies the transversality condition

$$\lim_{T\to\infty} \mathbb{E}_0\left[\beta^{T+1}\lambda_{T+1}B_T\right] = 0. \tag{32}$$

Conditions (30) and (32) are jointly sufficient.

The producer's valuation criterion is specified separately below using reciprocal cumulative gross returns. Log-linearizing the household conditions around the steady state, where $\beta\bar{R} = 1$, gives

$$\sigma c_t = \sigma \mathbb{E}_t c_{t+1} - r_t, \qquad w_t = \sigma c_t + \varphi n_t. \tag{33}$$

Within a perceived path the labor-supply condition holds at every date of the plan and the Euler equation holds at the dates $j < H$; the terminal date is closed by the household anchor derived in Appendix E.3.

E.1.2 Wholesale Firm

The competitive wholesale firm produces $X_t = A_t K_t^{\alpha} N_t^{1-\alpha}$, renting capital from capital producers at the rate ϱ_t^K and hiring labor at the wage W_t, and sells its output at the relative price MC_t. Taking all three prices as given, it solves, date by date,

$$\max_{K_t,\, N_t} \; MC_t\, A_t K_t^{\alpha} N_t^{1-\alpha} - \varrho_t^K K_t - W_t N_t. \tag{34}$$

Its conditions equate each factor price to the value of the corresponding marginal product,

$$\begin{aligned} \varrho_t^K &= MC_t\, \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} = MC_t\, \alpha \frac{X_t}{K_t}, \\ W_t &= MC_t\, (1-\alpha) A_t K_t^{\alpha} N_t^{-\alpha} = MC_t\, (1-\alpha) \frac{X_t}{N_t}, \end{aligned} \tag{35}$$

which are the factor-price conditions of Section 5.2. Adding the two factor bills, $\varrho_t^K K_t + W_t N_t = [\alpha + (1-\alpha)] MC_t X_t = MC_t X_t$, so the firm earns no profit.

Price dispersion separates final from wholesale output only at second order, as shown below, so $Y_t = X_t$ to first order. Taking logs of the production function and of the marginal-product forms in (35),

$$y_t = a_t + \alpha k_t + (1-\alpha) n_t, \qquad r_t^K = mc_t + y_t - k_t, \qquad w_t = mc_t + y_t - n_t, \tag{36}$$

where the two factor-price relations follow from the marginal-product forms on the right of (35).

E.1.3 Price Setting

Retailers are the source of the price rigidity in Section 5.2. This block derives the Phillips curve stated there and the dispersion result just used.

The final-good firm's problem. A competitive final-good firm bundles retail varieties into Y_t with a CES aggregator of elasticity $\varepsilon > 1$. Buying each variety at the price $P_t(\ell)$, it chooses the bundle that delivers a given Y_t at least cost,

$$\min_{\{Y_t(\ell)\}_{\ell\in[0,1]}} \int_0^1 P_t(\ell) Y_t(\ell)\, d\ell \qquad \text{s.t.} \qquad \left(\int_0^1 Y_t(\ell)^{(\varepsilon-1)/\varepsilon} d\ell\right)^{\varepsilon/(\varepsilon-1)} \geq Y_t. \tag{37}$$

With ζ_t the multiplier on the constraint, the condition for variety ℓ is $Y_t(\ell) = (P_t(\ell)/\zeta_t)^{-\varepsilon} Y_t$; substituting it back into the aggregator identifies ζ_t as the minimized cost of one unit of the bundle. Writing P_t for it, the demand schedule and price index are

$$Y_t(\ell) = \left(\frac{P_t(\ell)}{P_t}\right)^{-\varepsilon} Y_t, \qquad P_t = \left(\int_0^1 P_t(\ell)^{1-\varepsilon}\, d\ell\right)^{1/(1-\varepsilon)}, \tag{38}$$

and total expenditure satisfies $\int_0^1 P_t(\ell) Y_t(\ell)\, d\ell = P_t Y_t$, so the final-good firm earns no profit.

The retailer's problem. Retailers buy the wholesale good at nominal price P_t^w and transform it one for one, so $MC_t = P_t^w / P_t$. Future nominal profits are weighted by the household's marginal utility of the corresponding real payout. A retailer resetting at date t keeps its price with probability θ_p each period and, while unable to reset, faces demand $Y_{t+m|t} = (P_t^*/P_{t+m})^{-\varepsilon} Y_{t+m}$. With the sales subsidy τ_p, the reset problem is

$$\max_{P_t^*} \; \mathbb{E}_t \sum_{m=0}^{\infty} (\beta\theta_p)^m \left(\frac{C_{t+m}}{C_t}\right)^{-\sigma} \frac{P_t}{P_{t+m}} \left[(1+\tau_p) P_t^* - P_{t+m} MC_{t+m}\right] Y_{t+m|t} \tag{39}$$

Differentiating with respect to P_t^* and using $\partial Y_{t+m|t}/\partial P_t^* = -\varepsilon Y_{t+m|t}/P_t^*$ gives the first-order condition

$$\mathbb{E}_t \sum_{m=0}^{\infty} (\beta\theta_p)^m \left(\frac{C_{t+m}}{C_t}\right)^{-\sigma} \frac{P_t}{P_{t+m}} Y_{t+m|t} \left[(1+\tau_p)(1-\varepsilon) + \varepsilon \frac{P_{t+m} MC_{t+m}}{P_t^*}\right] = 0 \tag{40}$$

The subsidy $\tau_p = 1/(\varepsilon - 1)$ removes the steady-state markup, so the deterministic zero-inflation steady state has $\overline{MC} = 1$ and $(1+\tau_p)(1-\varepsilon) = -\varepsilon$.

Price index and dispersion. A mass $1-\theta_p$ of retailers charges P_t^* and a mass θ_p keeps last period's price, so

$$P_t^{1-\varepsilon} = (1-\theta_p)(P_t^*)^{1-\varepsilon} + \theta_p P_{t-1}^{1-\varepsilon}, \tag{41}$$

and price dispersion $\Delta_t \equiv \int_0^1 (P_t(\ell)/P_t)^{-\varepsilon} d\ell$ evolves as

$$\Delta_t = (1-\theta_p)\left(\frac{P_t^*}{P_t}\right)^{-\varepsilon} + \theta_p\, \Pi_t^{\varepsilon}\, \Delta_{t-1}, \qquad X_t = \Delta_t Y_t. \tag{42}$$

Around zero inflation, the first-order expansion of (42) is $\log \Delta_t = \theta_p \log \Delta_{t-1}$ up to second-order terms, so dispersion initialized at zero remains zero to first order and $Y_t = X_t$, as claimed in Section 5.2.

Log-linear pricing block. Write $p_t^* \equiv \log(P_t^*/P_t)$. Expanding (40) and (41) to first order around the zero-inflation steady state gives the reset recursion and the index relation

$$p_t^* = (1-\beta\theta_p)\, mc_t + \beta\theta_p\, \mathbb{E}_t\left[p_{t+1}^* + \pi_{t+1}\right], \qquad (1-\theta_p)\, p_t^* = \theta_p\, \pi_t, \tag{43}$$

and combining the two yields the New Keynesian Phillips curve

$$\pi_t = \beta\, \mathbb{E}_t \pi_{t+1} + \kappa_p\, mc_t, \qquad \kappa_p \equiv \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p}. \tag{44}$$

The same relations hold within a perceived path at the dates $j < H$; the terminal date is closed by the pricing anchor derived in Appendix E.3.

E.1.4 Capital Producers

Producer i plans at date t with horizon H_i, taking as given the announced path of gross real returns and the competitive rental rates it expects over the planning window. It solves

$$\max_{\{I_{i,t+j}\}_{j=0}^{H_i}} \; \mathbb{E}_t^{H_i}\left[\sum_{j=0}^{H_i} \mathcal{R}_{t,t+j}^{-1} D_{i,t+j} + \mathcal{R}_{t,t+H_i+1}^{-1}\, v\big(K_{i,t+H_i+1}\big)\right] \tag{45}$$

subject to, for $0 \le j \le H_i$,

$$\begin{aligned} K_{i,t+j+1} &= (1-\delta)K_{i,t+j} + I_{i,t+j}, \\ D_{i,t+j} &= \varrho_{t+j}^K K_{i,t+j} - I_{i,t+j} - \Psi\big(I_{i,t+j}, K_{i,t+j}\big), \end{aligned} \tag{46}$$

with $K_{i,t}$ given. Attaching to the date-$(t+j)$ capital law the multiplier $\mathcal{R}_{t,t+j}^{-1} Q_{i,t+j}$, the Lagrangian is

$$\mathbb{E}_t^{H_i}\left[\sum_{j=0}^{H_i} \mathcal{R}_{t,t+j}^{-1}\Big\{D_{i,t+j} + Q_{i,t+j}\big((1-\delta)K_{i,t+j} + I_{i,t+j} - K_{i,t+j+1}\big)\Big\} + \mathcal{R}_{t,t+H_i+1}^{-1} v\big(K_{i,t+H_i+1}\big)\right],$$

with $D_{i,t+j} = \varrho_{t+j}^K K_{i,t+j} - I_{i,t+j} - \Psi(I_{i,t+j}, K_{i,t+j})$. Scaling the multiplier by $\mathcal{R}_{t,t+j}^{-1}$ expresses $Q_{i,t+j}$ in date-$(t+j)$ goods.

Investment at $t+j$ appears at that date only, through $-I_{i,t+j} - \Psi(I_{i,t+j}, K_{i,t+j})$ in the dividend and through $+I_{i,t+j}$ in the capital law. Its condition is therefore static. Dividing by $\mathcal{R}^{-1}_{t,t+j}$,

$$\frac{\partial}{\partial I_{i,t+j}} : \; -1 - \Psi_I(I_{i,t+j}, K_{i,t+j}) + Q_{i,t+j} = 0.$$

Capital carried into $t+j+1$ appears at three places. It is subtracted in the date-$(t+j)$ capital law, weighted by $\mathcal{R}^{-1}_{t,t+j} Q_{i,t+j}$; it raises the date-$(t+j+1)$ dividend by $\varrho^K_{t+j+1} - \Psi_K(I_{i,t+j+1}, K_{i,t+j+1})$; and it enters the date-$(t+j+1)$ capital law through its undepreciated part $1-\delta$, weighted by $\mathcal{R}^{-1}_{t,t+j+1} Q_{i,t+j+1}$. For $j < H_i$, setting $\partial/\partial K_{i,t+j+1} = 0$,

$$\mathcal{R}^{-1}_{t,t+j} Q_{i,t+j} = \mathcal{R}^{-1}_{t,t+j+1} \left(\varrho^K_{t+j+1} - \Psi_K(I_{i,t+j+1}, K_{i,t+j+1}) + (1-\delta)\, Q_{i,t+j+1} \right).$$

Dividing by $\mathcal{R}^{-1}_{t,t+j}$ and using

$$\frac{\mathcal{R}^{-1}_{t,t+j+1}}{\mathcal{R}^{-1}_{t,t+j}} = R^{-1}_{t+j}$$

gives the shadow-value recursion.

At $j = H_i$ the capital carried past the window appears at two places only, the date-$(t+H_i)$ capital law and the continuation valuation $\mathcal{R}^{-1}_{t,t+H_i+1} v(K_{i,t+H_i+1})$. The same step gives

$$Q_{i,t+H_i} = R^{-1}_{t+H_i}\, v'(K_{i,t+H_i+1}),$$

the condition at the truncation date.

Producers sharing a horizon face the same perceived prices and Ψ is homogeneous of degree one, so they choose a common investment rate and the index i drops, leaving (7) and (8). Assumption 5.2 fixes v, and Appendix E.3 evaluates v' to give (9).

E.2 Linearization and Aggregation

This appendix linearizes the producer's plan, derives the solved investment equation and the linearized resource constraint, and proves that aggregate capital is a sufficient first-order state.

Adjustment-cost identities. The adjustment cost is homogeneous of degree one, so $\Psi(I,K) = K\,\psi(I/K)$ with $\psi(x) \equiv \Psi(x,1)$, and the normalization sets $\psi(\delta) = \psi'(\delta) = 0$ with $\phi \equiv \psi''(\delta) > 0$; the quadratic example is $\psi(x) = \frac{\phi}{2}(x-\delta)^2$. The partial derivatives are $\Psi_I = \psi'(x)$ and $\Psi_K = \psi(x) - x\psi'(x)$, so $\Psi_K(\delta,1) = 0$, and to first order in $x-\delta$,

$$\Psi_I = \phi\,(x-\delta), \qquad \Psi_K = -\delta\phi\,(x-\delta). \tag{47}$$

Writing q for the log deviation of the shadow value and using $\bar{Q} = 1$, so that $Q = 1 + q$ to first order, the first-order condition $1 + \Psi_I = Q$ gives $x - \delta = q/\phi$, and substituting into (47),

$$\Psi_K = -\delta\, q \tag{48}$$

at first order. Since $I/K = \delta e^{\iota}$ with $\iota \equiv i - k$, the investment-rate deviation satisfies $x - \delta = \delta\iota$ and hence

$$\iota = \frac{q}{\phi\delta}. \tag{49}$$

The shadow-value recursion. Along a perceived path the producer's recursion from Section 5.1 reads, in levels and with the one-period reciprocal gross return R_{t+j}^{-1},

$$Q_j^{(H)} = \frac{1}{R_{t+j}} \left(\varrho_{j+1}^{K,(H)} - \Psi_K\big(I_{j+1}^{(H)}, K_{j+1}^{(H)}\big) + (1-\delta)\, Q_{j+1}^{(H)} \right).$$

Using $Q_j^{(H)} = 1 + q_j^{(H)}$ to first order, $\bar{R} = \bar{\varrho}^K + 1 - \delta$, $\beta = 1/\bar{R}$, the expansions $1/R_{t+j} = \frac{1}{\bar{R}}(1 - r_{t+j})$ and $\varrho_{j+1}^{K,(H)} = \bar{\varrho}^K(1 + r_{j+1}^{K,(H)})$, and the envelope identity (48), the first-order terms give

$$q_j^{(H)} = \frac{1}{\bar{R}} \bar{\varrho}^K r_{j+1}^{K,(H)} - r_{t+j} + \frac{1}{\bar{R}} q_{j+1}^{(H)}, \qquad 0 \le j \le H - 1. \tag{50}$$

The coefficient on $q_{j+1}^{(H)}$ is $\frac{1}{\bar{R}}$ rather than $\frac{1}{\bar{R}}(1-\delta)$: the term $-\Psi_K$ contributes $\delta q_{j+1}^{(H)}$ by (48), which combines with the undepreciated continuation $(1-\delta)q_{j+1}^{(H)}$ to a full $q_{j+1}^{(H)}$. At the truncation date, the level condition $Q_H^{(H)} = R_{t+H}^{-1}\bar{R} = \bar{R}/R_{t+H}$ derived in Appendix E.3 linearizes to

$$q_H^{(H)} = -r_{t+H}\,: \tag{51}$$

the steady-state marginal value $\bar{R}$ discounted by the last announced rate in the window.

The solved investment equation. Solving (50) backward from (51) gives, by induction on j,

$$q_j^{(H)} = -\sum_{m=j}^{H} \frac{1}{\bar{R}^{m-j}}\, r_{t+m} + \bar{\varrho}^K \sum_{m=j+1}^{H} \frac{1}{\bar{R}^{m-j}}\, r_m^{K,(H)}, \qquad 0 \le j \le H\,: \tag{52}$$

the expression holds at $j = H$ by (51), and substituting the date-$(j+1)$ expression into (50) extends it to date j. Evaluated at $j = 0$, (52) is the finite-horizon shadow value $q_{t|t}^{(H)}$ displayed in Section 5.1; with (49) and the implemented choice $\iota_{H,t} = \iota_0^{(H)}$, it yields the solved investment equation

$$\iota_{H,t} = \frac{1}{\phi\delta} \left[-\sum_{j=0}^{H} \frac{1}{\bar{R}^j}\, r_{t+j} + \bar{\varrho}^K \sum_{j=1}^{H} \frac{1}{\bar{R}^j}\, r_j^{K,(H)} \right]. \tag{53}$$

The resource constraint. For the resource constraint $Y_t = C_t + I_t + \Xi_t$, note that $\Psi(I, K) = K\psi(I/K)$ with $\psi(\delta) = \psi'(\delta) = 0$ implies $\Psi(I_{H,t}, K_{H,t}) = O\big((\iota_{H,t})^2\big)K_{H,t}$, so aggregate adjustment costs $\Xi_t = \sum_H \mu_H \Psi(I_{H,t}, K_{H,t})$ are second order and drop from the constraint at first order. Hence $\bar{Y} y_t = \bar{C} c_t + \sum_H \mu_H \bar{I}_H \, i_{H,t}$, which with the aggregation below is

$$y_t = s_C \, c_t + s_I \big(k_t + \iota_t\big). \tag{54}$$

The same steps apply within a perceived path.

Aggregation. Steady-state investment satisfies $\bar{I}_H = \delta \bar{K}_H$ for every type, so with $\omega_H \equiv \mu_H \bar{K}_H / \bar{K}$ the investment weights equal the capital weights: $\mu_H \bar{I}_H / \bar{I} = \omega_H$. Linearizing $K_t = \sum_H \mu_H K_{H,t}$ and $I_t = \sum_H \mu_H I_{H,t}$ gives $k_t = \sum_H \omega_H k_{H,t}$ and $i_t = \sum_H \omega_H i_{H,t}$; since $i_{H,t} = k_{H,t} + \iota_{H,t}$,

$$i_t = k_t + \iota_t, \qquad \iota_t \equiv \sum_{H \in \mathcal{H}} \omega_H \, \iota_{H,t}, \tag{55}$$

which is the form used in (17). The type capital law $K_{H,t+1} = (1-\delta)K_{H,t} + I_{H,t}$ linearizes to $k_{H,t+1} = k_{H,t} + \delta \iota_{H,t}$. The steady state pins down only aggregate capital. Every perceived path is at the steady state and every type invests at rate δ, so any collection of positive weights $\{\omega_H\}_{H \in \mathcal{H}}$ summing to one is consistent with the steady-state conditions.

Sufficiency of aggregate capital. At first order, given initial type capital deviations $\{k_{H,0}\}_{H \in \mathcal{H}}$ and announced paths, the realized paths of all aggregate variables depend on the initial deviations only through $k_0 = \sum_H \omega_H k_{H,0}$, and aggregate capital evolves as $k_{t+1} = k_t + \delta \iota_t$. Within a horizon-$H$ perceived path, all producers face the same perceived prices and plan to the same terminal date. The homogeneity identities above then imply a common investment rate. The path therefore carries a single capital stock, initialized at the realized aggregate k_t, and its solution depends on initial conditions only through k_t and the announced paths. Each implemented rate $\iota_{H,t} = \iota_0^{(H)}$ is then a function of k_t and those paths. Aggregating the linearized type capital laws with the fixed weights ω_H gives $k_{t+1} = k_t + \delta \iota_t$, so the realized capital block closes in k_t; the remaining realized variables are determined by the conditions of Definition 5.2 given $\{k_t, \iota_t\}$ and the announced paths. The fixed weights suffice because, in the exact aggregate law $K_{t+1} = \sum_H \mu_H [(1-\delta)K_{H,t} + I_{H,t}]$, deviations of the capital weights from ω_H multiply cross-type differences in investment rates, and both factors are first order, so the error is second order. Away from first order this fails: types implement different investment rates, the capital distribution moves, and realized dynamics depend on the joint distribution, as noted in Section 5.3.

E.3 Terminal Anchors

This appendix derives the continuation functions of Assumption 5.2 and the terminal conditions of the horizon-H perceived path of Section 5.3. Each anchor is the valuation that is correct in

the stationary environment in which it was learned, applied to the state actually carried past the truncation date.

The producer anchor. Consider a producer in an economy held at its deterministic steady state, with constant rental rate $\bar{\varrho}^K$ and gross return $\bar{R}$. Its stationary valuation problem is

$$\bar{V}(K) = \max_I \left\{ \bar{\varrho}^K K - I - \Psi(I,K) + \frac{1}{\bar{R}} \bar{V}\big((1-\delta)K + I\big) \right\}. \tag{56}$$

Because Ψ is homogeneous of degree one and the capital law is linear, the problem is homogeneous: $\bar{V}(\lambda K) = \lambda \bar{V}(K)$ for $\lambda > 0$, so $\bar{V}(K) = \bar{v}\,K$ for a constant $\bar{v}$. The first-order condition of (56) is $1 + \Psi_I = \bar{v}/\bar{R}$; at the steady-state investment rate, $\Psi_I = 0$, so $\bar{v} = \bar{R}$. The envelope condition $\bar{v} = \bar{\varrho}^K - \Psi_K + (1-\delta)\bar{v}/\bar{R}$ then reproduces the steady-state identity $\bar{R} = \bar{\varrho}^K + 1 - \delta$. The learned valuation is therefore

$$v(K) = \bar{R}\,K, \qquad v'(K) = \bar{R}, \tag{57}$$

which is the producer part of Assumption 5.2.

The household anchor. The perceived household at the truncation date $t+H$ has remaining horizon zero: it values the real bond position carried past $t+H$ with the continuation value of the stationary saving problem. Let $u(C) \equiv (C^{1-\sigma}-1)/(1-\sigma)$ and let $\bar{y}$ denote steady-state non-financial income, so that $\bar{C} = \bar{y}$ when bonds are in zero net supply. The continuation function solves the stationary Bellman equation $\mathcal{W}(B) = \max_C\{u(C) + \beta\,\mathcal{W}(B')\}$ subject to $C + \beta B' = \bar{y} + B$, whose solution is

$$\mathcal{W}(B) = \frac{1}{1-\beta}\,u\big(\bar{C} + (1-\beta)B\big). \tag{58}$$

It is strictly concave: the annuity value of extra wealth raises permanent consumption and lowers its marginal utility. What the anchor fixes is the whole function, and the two properties used below are its level and curvature at the symmetric position,

$$\mathcal{W}'(B) = \big(\bar{C} + (1-\beta)B\big)^{-\sigma}, \qquad \mathcal{W}'(0) = \bar{C}^{-\sigma}, \qquad \mathcal{W}''(0) = -\sigma(1-\beta)\,\bar{C}^{-\sigma-1} < 0. \tag{59}$$

The zero-depth household solves

$$\max_{C,N,B} \frac{C^{1-\sigma}-1}{1-\sigma} - \chi\frac{N^{1+\varphi}}{1+\varphi} + \beta\,\mathcal{W}(B) \qquad \text{s.t.} \qquad C + \frac{B}{R_{t+H}} = WN + B^- + \Omega,$$

where B^- is its inherited position. The bond and labor conditions are

$$C^{-\sigma} = \beta R_{t+H}\,\mathcal{W}'(B), \qquad \chi N^{\varphi} = C^{-\sigma} W. \tag{60}$$

The bond condition retains the household's own terminal position, because $\mathcal{W}'$ is not constant. Writing $b \equiv B/\bar{C}$ and using $\beta\bar{R} = 1$, it linearizes to

$$\sigma c_H^{(H)} = -r_{t+H} + \eta_B \, b_{H+1}^{(H)}, \qquad \eta_B \equiv \sigma(1-\beta) > 0, \tag{61}$$

where η_B is the curvature term: it is $-\mathcal{W}''(0)\bar{C}/\mathcal{W}'(0)$, and it would be zero under an affine continuation function. The perceived households inside a given perceived path are identical and perceived bonds are in zero net supply, so the symmetric terminal position is $b_{H+1}^{(H)} = 0$. Imposing it gives the terminal Euler condition of the perceived path, which is (13′) at $j = H$,

$$\sigma c_H^{(H)} = -r_{t+H}, \tag{62}$$

which is therefore stated there with perceived bond-market clearing already substituted in. The labor condition is unchanged from the interior dates. The producer discounts the terminal valuation by the final reciprocal gross return R_{t+H}^{-1}. Substituting (57) into the producer's truncation condition gives $Q_H^{(H)} = \bar{R}/R_{t+H}$ and therefore the terminal condition $q_H^{(H)} = -r_{t+H}$ used in Appendix E.2.

The pricing anchor. A perceived retailer that resets at the truncation date $t+H$ values the relative price it carries forward with the discounted stream of profits that price would earn while it remains in force, evaluated at the zero-inflation steady state and in steady-state marginal-utility units: with ϱ denoting the reset relative price, the continuation function is

$$\mathcal{G}(\varrho) = \sum_{m=1}^{\infty} (\beta\theta_p)^m \big[(1+\tau_p)\varrho - 1\big] \varrho^{-\varepsilon} \, \bar{Y} = \frac{\beta\theta_p}{1-\beta\theta_p} \big[(1+\tau_p)\varrho - 1\big] \varrho^{-\varepsilon} \, \bar{Y}, \tag{63}$$

since at that steady state the price level is constant, marginal cost is one, demand is $\bar{Y}\varrho^{-\varepsilon}$, survival contributes θ_p^m, and steady-state marginal-utility discounting contributes β^m. The zero-depth reset problem, with current profits weighted by the perceived marginal utility at the truncation date relative to its steady-state value, $\lambda_H^{(H)} \equiv \big(C_H^{(H)}\big)^{-\sigma}$ and $\bar{\lambda} \equiv \bar{C}^{-\sigma}$, is

$$\max_{\varrho} \; \frac{\lambda_H^{(H)}}{\bar{\lambda}} \big[(1+\tau_p)\varrho - MC_H^{(H)}\big] \varrho^{-\varepsilon} Y_H^{(H)} + \mathcal{G}(\varrho). \tag{64}$$

Using $(1+\tau_p)(1-\varepsilon) = -\varepsilon$, its first-order condition reduces to

$$\frac{\lambda_H^{(H)}}{\bar{\lambda}} \, Y_H^{(H)} \big(MC_H^{(H)} - \varrho\big) + \frac{\beta\theta_p}{1-\beta\theta_p} \, \bar{Y} \, (1-\varrho) = 0.$$

The weight multiplies a term that vanishes at the steady state, so it drops at first order, and the first-order approximation is

$$p_H^{*,(H)} = (1-\beta\theta_p) \, mc_H^{(H)}. \tag{65}$$

Combining (65) with the price-index relation in (43), which holds within the perceived path as a structural relation, gives the terminal pricing condition of the perceived path, which is (16′) at $j = H$,

$$\pi_H^{(H)} = \kappa_p \, mc_H^{(H)}. \tag{66}$$

At the interior dates $j < H$, a perceived retailer resetting at $t+j$ plans explicitly through $t+H$ and values the dates beyond $t+H$ at which its price may still be in force with the same steady-state continuation. Its linearized reset condition is the truncated sum

$$\frac{p_j^{*,(H)}}{1-\beta\theta_p} = \sum_{m=0}^{H-j} (\beta\theta_p)^m \Big[mc_{j+m}^{(H)} + \sum_{s=j+1}^{j+m} \pi_s^{(H)} \Big] + \frac{(\beta\theta_p)^{H-j+1}}{1-\beta\theta_p} \sum_{s=j+1}^{H} \pi_s^{(H)}, \tag{67}$$

where the tail term carries only cumulated within-window inflation because prices and marginal costs beyond the window are held at their steady-state values. A direct computation on (67) verifies the recursion

$$p_j^{*,(H)} = (1-\beta\theta_p)\, mc_j^{(H)} + \beta\theta_p \big(p_{j+1}^{*,(H)} + \pi_{j+1}^{(H)} \big), \qquad j < H, \tag{68}$$

with terminal condition (65); combined with the index relation, this delivers exactly the perceived pricing condition (16′), namely the Phillips curve at the dates $j < H$ together with (66) at the truncation date.

The primed conditions. The three forward-looking primed conditions in Section 5.3 impose their source relations at each date before the truncation date and are closed by the terminal values in Assumption 5.2. For the capital value this gives (52). For the household, the perceived interior Euler equation, which is (13) at the dates $j < H$, and the terminal condition (62) give

$$\sigma c_j^{(H)} = -\sum_{\ell=0}^{H-j} r_{t+j+\ell}, \qquad 0 \le j \le H. \tag{69}$$

The perceived Phillips curve and (66) solve to

$$\pi_j^{(H)} = \kappa_p \sum_{\ell=0}^{H-j} \beta^\ell mc_{j+\ell}^{(H)}, \qquad 0 \le j \le H. \tag{70}$$

The investment condition, the production function, the factor-price conditions, the labor-supply condition, and the resource constraint are static, so they coincide with their sources. The capital condition adds the initialization $k_{t|t}^{(H)} = k_t$, which states that each producer forecasts from realized aggregate capital. Appendix E.2 shows that aggregate capital is a sufficient first-order state.

E.4 Proofs for Section 5.4

The pairwise decomposition. The solved investment equation (53) holds for every horizon, with the announced rate path common to all perceived paths because the central bank is not

assigned a planning horizon. Subtracting the expression for H' from the expression for $H > H'$, the contributions to the direct effect at the dates $j \leq H'$ cancel, leaving

$$\iota_{H,t} - \iota_{H',t} = -\frac{1}{\phi\delta} \sum_{j=H'+1}^{H} \frac{1}{\bar{R}^j} r_{t+j} + \frac{\bar{\varrho}^K}{\phi\delta} \left[\sum_{j=1}^{H} \frac{1}{\bar{R}^j} r_j^{K,(H)} - \sum_{j=1}^{H'} \frac{1}{\bar{R}^j} r_j^{K,(H')} \right],$$

which is the pairwise decomposition displayed in Section 5.4.

Steady-state sign of the capital coefficient. The coefficient on capital in Equation (22) is negative when

$$s_I < \frac{1+\varphi\alpha}{1+\varphi}. \tag{71}$$

The steady-state capital law, factor-pricing condition, and producer optimality condition imply, respectively,

$$\bar{I} = \delta\bar{K}, \qquad \bar{\varrho}^K \bar{K} = \alpha\bar{Y}, \qquad \bar{\varrho}^K = \bar{R} - 1 + \delta, \tag{72}$$

where the steady-state subsidy normalizes marginal cost to one. Therefore

$$s_I \equiv \frac{\bar{I}}{\bar{Y}} = \frac{\delta\alpha}{\bar{R} - 1 + \delta}. \tag{73}$$

Substituting this identity into the right-hand side of (71) gives

$$\frac{1+\varphi\alpha}{1+\varphi} - s_I = \frac{(1+\varphi\alpha)(\bar{R}-1) + \delta(1-\alpha)}{(1+\varphi)(\bar{R}-1+\delta)} > 0. \tag{74}$$

The strict inequality follows from $\bar{R} = \beta^{-1} > 1$, $\delta > 0$, $0 < \alpha < 1$, and $\varphi \geq 0$.

The full perceived-path block. Suppress the origin label t and write $x_j \equiv r_{t+j}$. Define the coefficients in Equation (22) as

$$A_c \equiv \sigma + \frac{(1+\varphi)s_C}{1-\alpha} > 0, \qquad A_k \equiv \frac{1+\varphi}{1-\alpha}\left[s_I - \frac{1+\varphi\alpha}{1+\varphi}\right], \qquad A_\iota \equiv \frac{(1+\varphi)s_I}{1-\alpha} > 0. \tag{75}$$

Equations (73) and (74) imply $A_k < 0$. For a horizon-H perceived path, the household Euler equation and the investment first-order condition imply

$$c_j^{(H)} = -\frac{1}{\sigma} \sum_{m=j}^{H} x_m, \qquad \iota_j^{(H)} = \frac{q_j^{(H)}}{\phi\delta}. \tag{76}$$

Substituting (76) and $r_{j+1}^{K,(H)} = A_c c_{j+1}^{(H)} + A_k k_{j+1}^{(H)} + A_\iota \iota_{j+1}^{(H)}$ into the shadow-value recursion gives

$$q_j^{(H)} = \frac{1}{\bar{R}} \bar{\varrho}^K A_k k_{j+1}^{(H)} - x_j - \frac{1}{\bar{R}} \frac{\bar{\varrho}^K A_c}{\sigma} \sum_{m=j+1}^{H} x_m + \frac{1}{\bar{R}} \left(1 + \frac{\bar{\varrho}^K A_\iota}{\phi \delta} \right) q_{j+1}^{(H)}, \qquad 0 \le j < H. \tag{77}$$

Define

$$\lambda_F \equiv \frac{1}{\bar{R}} \bar{\varrho}^K A_k < 0, \qquad \widetilde{\beta} \equiv \frac{1}{\bar{R}} \left(1 + \frac{\bar{\varrho}^K A_\iota}{\phi \delta} \right) > 0, \qquad z_j^{(H)} \equiv x_j + \frac{1}{\bar{R}} \frac{\bar{\varrho}^K A_c}{\sigma} \sum_{m=j+1}^{H} x_m. \tag{78}$$

Because $z_H^{(H)} = x_H$, the terminal condition is $q_H^{(H)} = -z_H^{(H)}$. The full system relevant for the proposition can therefore be written as

$$q_j^{(H)} = \lambda_F k_{j+1}^{(H)} - z_j^{(H)} + \widetilde{\beta} q_{j+1}^{(H)} \quad (j < H), \qquad q_H^{(H)} = -z_H^{(H)}, \qquad k_0^{(H)} = k_t, \qquad k_{j+1}^{(H)} = k_j^{(H)} + \frac{q_j^{(H)}}{\phi}. \tag{79}$$

The argument below requires only $\widetilde{\beta} > 0$, not $\widetilde{\beta} < 1$, because every planning problem has a finite horizon.

At a node with h remaining steps, current capital k, and effective disturbances $z_0, \ldots, z_h$, write the shadow value as $q_h(k; z_0, \ldots, z_h)$. Equation (79) implies

$$q_0(k; z_0) = -z_0, \qquad q_h = \lambda_F \left(k + \frac{q_h}{\phi} \right) - z_0 + \widetilde{\beta} q_{h-1} \left(k + \frac{q_h}{\phi}; z_1, \ldots, z_h \right), \quad h \ge 1. \tag{80}$$

Define

$$a_0 \equiv 0, \qquad b_0 \equiv 1, \qquad b_h \equiv \frac{1}{1 - (\lambda_F + \widetilde{\beta} a_{h-1})/\phi}, \qquad a_h \equiv b_h (\lambda_F + \widetilde{\beta} a_{h-1}), \quad h \ge 1, \tag{81}$$

and

$$C_{0,0}^F \equiv 1, \qquad C_{h,0}^F \equiv b_h, \qquad C_{h,m}^F \equiv b_h C_{h-1,m-1}^F \quad (1 \le m \le h), \qquad C_{h,m}^F = \prod_{\ell=0}^{m} b_{h-\ell}. \tag{82}$$

Lemma E.1 (Impact response in the full perceived-path system). *For every $h \ge 0$,*

$$q_h(k; z_0, \ldots, z_h) = a_h k - \sum_{m=0}^{h} \widetilde{\beta}^m C_{h,m}^F z_m. \tag{83}$$

Moreover, $-\phi < a_h < 0$ for $h \ge 1$, $a_1 < a_0 = 0$, and $a_h < a_{h-1}$ for $h \ge 2$. In addition,

$1 = b_0 > b_1 > b_2 > \cdots > 0$. *Consequently,*

$$0 < C^F_{H,m} < C^F_{H',m} \qquad \textit{whenever } H > H' \geq m. \tag{84}$$

Proof. Equation (83) holds at $h = 0$ because $q_0(k; z_0) = -z_0$. Suppose it holds at $h-1$. Substitution into (80) gives

$$q_h = (\lambda_F + \widetilde{\beta} a_{h-1}) \left(k + \frac{q_h}{\phi} \right) - z_0 - \sum_{m=1}^{h} \widetilde{\beta}^m C^F_{h-1,m-1} z_m.$$

Solving for q_h yields (83) with the coefficients in (81)–(82).

For the signs, define

$$F(u) \equiv \frac{\lambda_F + \widetilde{\beta} u}{1 - (\lambda_F + \widetilde{\beta} u)/\phi}, \qquad u \leq 0.$$

Since $\lambda_F < 0$ and $\widetilde{\beta} > 0$, $F(u) \in (-\phi, 0)$ and F is strictly increasing. Because $a_1 = F(0) < 0 = a_0$, induction gives $a_h = F(a_{h-1}) < F(a_{h-2}) = a_{h-1}$ for $h \geq 2$. Thus $-\phi < a_h < a_{h-1} < 0$. Equation (81) then implies $b_h \in (0, 1)$, and the strict decline of a_{h-1} implies $b_h < b_{h-1}$, including $b_1 < b_0 = 1$. Finally, $C^F_{H,m}$ is the product of $b_H, \ldots, b_{H-m}$. If $H > H' \geq m$, each factor in $C^F_{H,m}$ is strictly smaller than the corresponding factor in $C^F_{H',m}$, proving (84). □

At the root of the horizon-H perceived path,

$$q_0^{(H)} = a_H k_t - \sum_{m=0}^{H} \widetilde{\beta}^m C^F_{H,m} z_m^{(H)}. \tag{85}$$

Proof of Proposition 5.1. The announced path satisfies $x_m \geq 0$ for $0 \leq m \leq H'$ and $x_m = 0$ for $m > H'$. Hence, at every date in the common window,

$$z_j^{(H)} = x_j + \frac{1}{\bar{R}} \frac{\bar{\varrho}^K A_c}{\sigma} \sum_{m=j+1}^{H'} x_m = z_j^{(H')} \equiv z_j \geq 0, \qquad 0 \leq j \leq H', \tag{86}$$

whereas $z_j^{(H)} = 0$ for $j > H'$. At least one z_j is strictly positive. Since $k_t = 0$, Equation (85) gives

$$q_0^{(H)} = -\sum_{m=0}^{H'} \widetilde{\beta}^m C^F_{H,m} z_m, \qquad q_0^{(H')} = -\sum_{m=0}^{H'} \widetilde{\beta}^m C^F_{H',m} z_m.$$

It follows from (84) that

$$\begin{aligned} \iota_{H,t} - \iota_{H',t} &= \frac{q_0^{(H)} - q_0^{(H')}}{\phi\delta} \\ &= \frac{1}{\phi\delta} \sum_{m=0}^{H'} \widetilde{\beta}^m \left(C^F_{H',m} - C^F_{H,m} \right) z_m > 0. \end{aligned} \tag{87}$$

The difference in direct effects in Equation (21) is zero because $x_m = 0$ for every $m > H'$. Therefore $\iota_{H,t} - \iota_{H',t} = \Delta\Gamma_{H,H',t} > 0$, as claimed. □

Proof of Proposition 5.2. Use the notation and coefficients of Lemma E.1, and write $x_m \equiv r_{t+m}$. Define

$$\chi \equiv \frac{1}{\bar{R}} \frac{\bar{\varrho}^K A_c}{\sigma} > 0, \qquad \mathcal{W}_{H,j} \equiv \widetilde{\beta}^j C^F_{H,j} > 0. \tag{88}$$

The shorter-horizon producer perceives no rate movement, so $z_j^{(H')} = 0$ for every $0 \le j \le H'$. Because $k_t = 0$, Equation (85) gives

$$q_0^{(H')} = 0, \qquad \iota_{H',t} = 0. \tag{89}$$

For the longer-horizon producer,

$$z_j^{(H)} = x_j + \chi \sum_{m=j+1}^{H} x_m.$$

Applying Equation (85) at $k_t = 0$ and reversing the order of summation gives

$$\begin{aligned} q_0^{(H)} &= -\sum_{j=0}^{H} \mathcal{W}_{H,j} \left(x_j + \chi \sum_{m=j+1}^{H} x_m \right) \\ &= -\sum_{m=H'+1}^{H} \left[\mathcal{W}_{H,m} + \chi \sum_{j=0}^{m-1} \mathcal{W}_{H,j} \right] x_m. \end{aligned} \tag{90}$$

Define the total loading of date-m guidance on the root shadow value by

$$\mathcal{G}_{H,m} \equiv \mathcal{W}_{H,m} + \chi \sum_{j=0}^{m-1} \mathcal{W}_{H,j} > 0. \tag{91}$$

The announced path is nonnegative and contains at least one strictly positive rate movement. Hence

$$\iota_{H,t} = -\frac{1}{\phi\delta} \sum_{m=H'+1}^{H} \mathcal{G}_{H,m} x_m < 0 = \iota_{H',t}, \tag{92}$$

which proves the investment ranking.

For completeness, combining (92) with the pairwise decomposition yields

$$\Delta\Gamma_{H,H',t} = \frac{1}{\phi\delta} \sum_{m=H'+1}^{H} \left(\frac{1}{\bar{R}^m} - \mathcal{G}_{H,m} \right) x_m. \tag{93}$$

Thus the difference in indirect effects can be positive or negative, depending on the discounted equilibrium loadings $\mathcal{G}_{H,m}$, even though the total investment difference is strictly negative. □

E.5 Empirical Rate Paths and Standardized-Shock Counterfactuals

Table E.1 reports the planning-horizon distribution used in the calibration. Figure E.2 reports the interest-rate inputs used for the model-data comparison in Section 6.2. Figure E.1 reports, for the experiments of Figure 13, the cumulative investment response over the whole transition by implementation date. We also describe the annual horizon distributions used in Section 6.4 and report the corresponding investment responses.

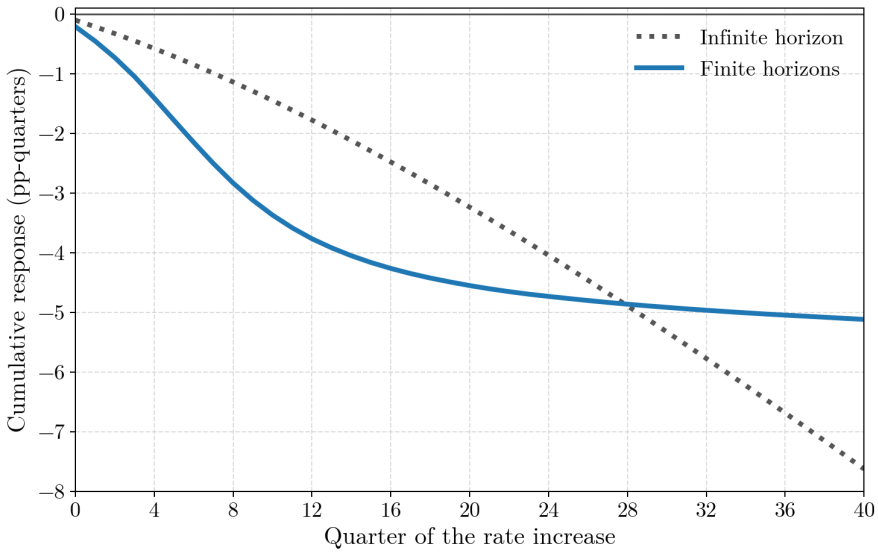
Figure E.2 reports the interest-rate paths used in the quantitative analysis. We estimate the responses of the quarterly average three-month US Treasury bill rate to current-rate and forward-guidance shocks using local projections over 1991Q1–2019Q4. Each regression includes a constant, one monetary policy shock, and one lag of the nominal rate. We use the resulting nominal-rate responses as proxies for the model's real-rate paths. The current-rate response declines after an early peak and is set to zero from quarter 13 onward. The forward-guidance response is more persistent, so we extend it beyond quarter 12 by geometric decay with quarterly persistence 0.95. The solid lines show the estimated responses, and the dashed lines show these extensions.

Figure E.3 assesses the sensitivity of the investment differentials to the interest-rate measure and the persistence of the forward-guidance extension. We estimate responses of the ex post real rate, defined as the nominal rate less next-quarter annualized GDP-deflator inflation. For both nominal and ex post real rates, we consider quarterly persistence of 0.95 and 0.90. All four specifications produce a positive current-rate differential and a negative forward-guidance differential. With persistence 0.95, the ex post real-rate inputs yield responses close to the baseline. Lower persistence substantially reduces the magnitude of the forward-guidance differential at later dates. Thus, the qualitative pattern holds across these specifications, while the quantitative fit depends on the assumed persistence of tightening.

Figure E.4 shows that the qualitative contrast between current-rate and forward-guidance differentials also arises when effective federal funds rates or one-year Treasury yields supply the model's interest-rate inputs. Both alternatives produce a positive current-rate differential and a negative forward-guidance differential. The one-year Treasury inputs generate a forward-guidance differential closer to the empirical estimates, whereas the federal funds inputs produce a substantially larger negative differential at later horizons. Thus, the contrast across shocks holds across these interest-rate measures, while the quantitative fit depends on the rate path used.

Figure E.5 shows that the contrast between current-rate and forward-guidance differentials holds across alternative values of investment adjustment costs. We vary the adjustment-cost parameter ϕ from 5 to 9 around the baseline value of 7, holding all other parameters, the planning-horizon distribution, and the interest-rate paths fixed. After the announcement quarter, every specification produces a positive current-rate differential and a negative forward-guidance differential. Lower adjustment costs increase the current-rate differential, while the forward-guidance differentials vary less across specifications. The model responses generally lie within the empirical 90 percent confidence intervals.

Figure E.1: Cumulative investment response by implementation date



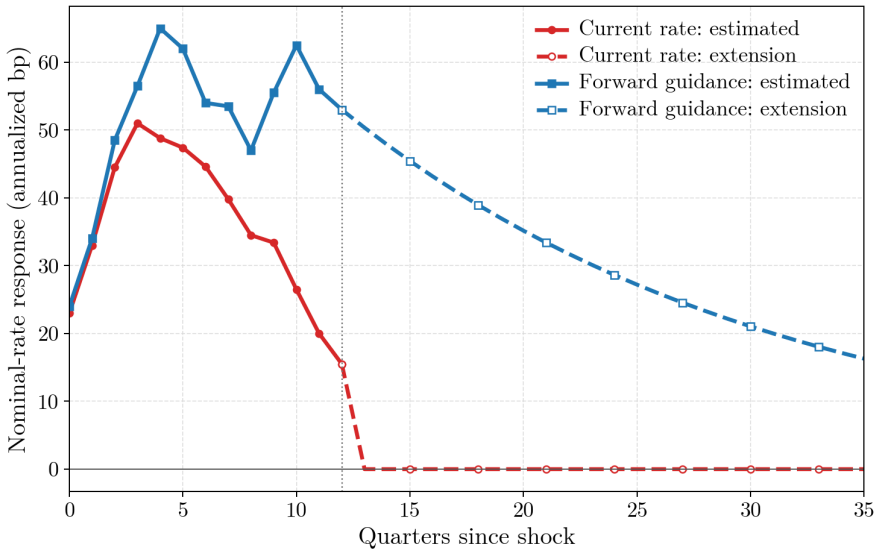
Notes: Each experiment announces at date zero a one-quarter, 50-basis-point annualized real-rate increase, implemented in the quarter on the horizontal axis. Rates are at steady state in all other quarters. The vertical axis reports a numerical approximation to the cumulative investment response over the entire transition, measured in percentage-point-quarters. The solid line shows the cumulative investment response under finite planning with the measured horizon distribution, weighted by firm counts; the dotted line shows the response under infinite planning.

Table E.1: Empirical planning-horizon distribution.

H	ω_H	H	ω_H	H	ω_H	H	ω_H
1	0	11	0.0925	21	0.0414	31	0.0052
2	0.0001	12	0.0889	22	0.0098	32	0.0367
3	0.0006	13	0.0637	23	0.0475	33	0.0037
4	0.0333	14	0.0379	24	0.0047	34	0.0017
5	0.0238	15	0.0902	25	0.0047	35	0.0042
6	0.0228	16	0.0166	26	0.0051	36	0.0020
7	0.0256	17	0.0334	27	0.0022	37	0.0011
8	0.0630	18	0.0247	28	0.0063	38	0.0012
9	0.0599	19	0.0258	29	0.0022	39	0.0037
10	0.0768	20	0.0107	30	0.0013	40	0.0248

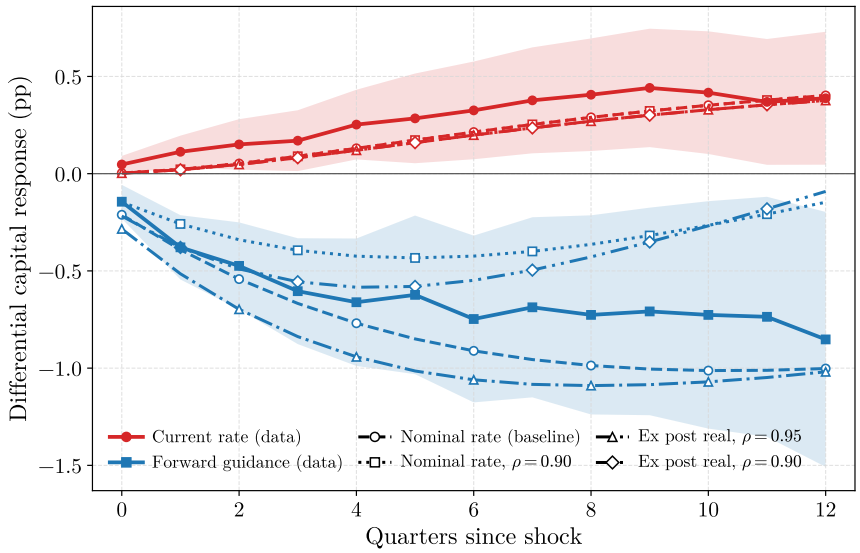
Notes: H is the planning horizon in quarters and ω_H the steady-state capital weight defined in (17). The displayed grid begins at $H = 1$, which has zero mass, and the weights sum to one. The final type collects every firm whose planning horizon exceeds forty quarters.

Figure E.2: Nominal interest-rate responses to current-rate and forward-guidance shocks



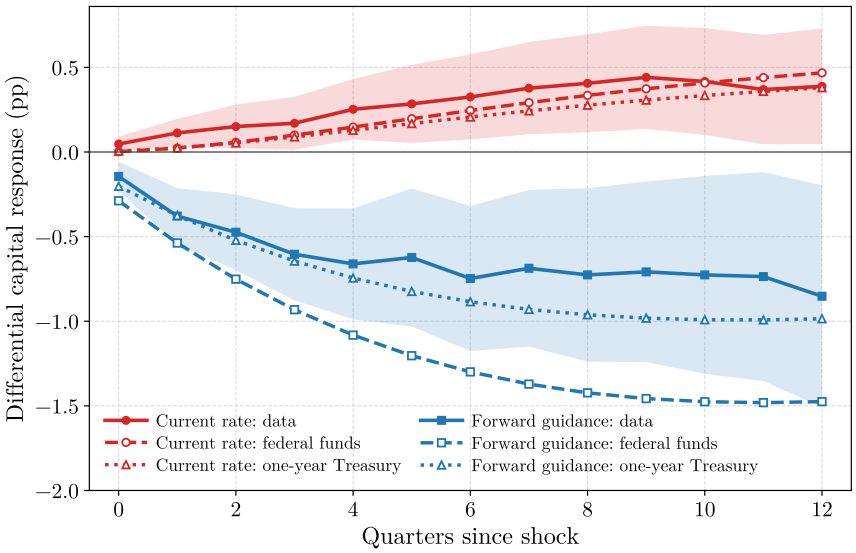
Notes: The figure reports nominal interest-rate responses to monetary tightening. Responses are measured in annualized basis points. Solid lines show estimated responses through quarter 12; dashed lines show extensions. The current-rate response is set to zero from quarter 13 onward. The forward-guidance response is extended beyond quarter 12 with quarterly persistence 0.95, broadly in line with empirical interest-rate persistence.

Figure E.3: Sensitivity to the interest-rate measure and persistence



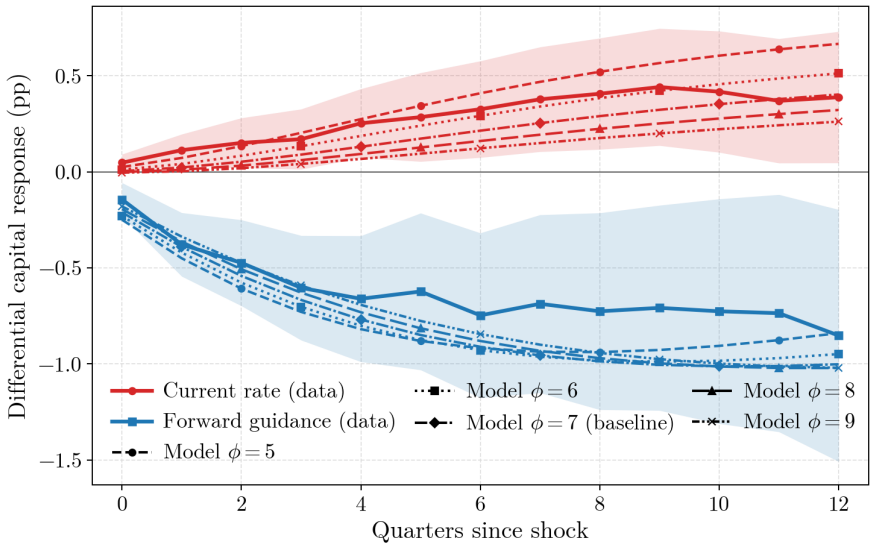
Notes: Solid lines show empirical cumulative capital-growth differentials for a one-standard-deviation increase in log planning horizon; shaded bands show 90 percent confidence intervals. Red denotes current-rate shocks and blue denotes forward-guidance shocks. Model curves use nominal-rate paths or ex post real-rate paths, defined as the nominal rate less next-quarter annualized GDP-deflator inflation. The forward-guidance extensions decay geometrically with quarterly persistence 0.95 or 0.90. The calibration and capital-weighted planning-horizon distribution are those used in Figure 12.

Figure E.4: Sensitivity of planning-horizon differentials to alternative interest-rate measures



Notes: Solid lines show empirical cumulative capital-growth differentials for a one-standard-deviation increase in log planning horizon; shaded bands show 90 percent confidence intervals. Red denotes current-rate shocks and blue denotes forward-guidance shocks. Dashed lines show model predictions using effective federal funds rate inputs; dotted lines use one-year Treasury yield inputs. Calibration, shock normalization, and extrapolation rules are those used in Figure 12.

Figure E.5: Sensitivity of the planning-horizon differential to investment adjustment costs



Notes: Solid lines show empirical differences in cumulative capital growth following current-rate and forward-guidance tightening, measured in percentage points for a one-standard-deviation increase in log planning horizon. Shaded bands show 90 percent confidence intervals for the empirical estimates. Model curves vary the investment adjustment-cost parameter over $\phi \in \{5, 6, 7, 8, 9\}$, with $\phi = 7$ the baseline. All other parameters, the capital-weighted horizon distribution, and the real-rate paths are held fixed at their values in Figure 12.

Annual planning-horizon distributions The comparison in Section 6.4 uses the all-country annual horizon series in Figure 14. Each firm's annual horizon is the equal-call average of its call-level horizons, and firms receive equal weight within each year. The series covers 2002–2025. For this series, we screen long-horizon observations for phone numbers interpreted as years, financial amounts interpreted as quarter or year dates, and ages or historical references interpreted as prospective durations. These screens address identified measurement errors without capping the resulting horizons.[25]

We approximate each annual distribution with a generalized-gamma distribution whose location, log scale, and shape parameters vary by year. Estimation minimizes the mean squared distance between fitted and empirical cumulative distribution functions, with equal weight on each year, plus a penalty on squared year-to-year changes in the parameters. Four-fold firm-level cross-validation selects the penalty by minimizing held-out CDF distance; all observations from a firm belong to the same fold.[26] This approach allows short-lived changes in the distribution while pooling information across years. Figure E.6 compares the fitted and empirical distributions for 2002, 2021, and 2025. The fitted medians in 2002 and 2025 are 1.50 and 2.04 years, compared with empirical medians of 1.50 and 2.05 years. The approximation places more weight in the far upper tail: fitted means are 2.86 and 3.57 years, compared with empirical means of 2.41 and 3.24 years.

For the model, we discretize the fitted continuous distributions to the nearest quarter, with a minimum horizon of one quarter. Equal steady-state capital per firm makes the firm share in each horizon group equal to its steady-state capital share. Households have rational expectations, and all other parameters are fixed across distributions. Each policy experiment announces at date zero a one-quarter, 50-basis-point annualized real-rate increase. Cumulative responses sum investment deviations over the full transition, including the convergent tail after implementation.[27]

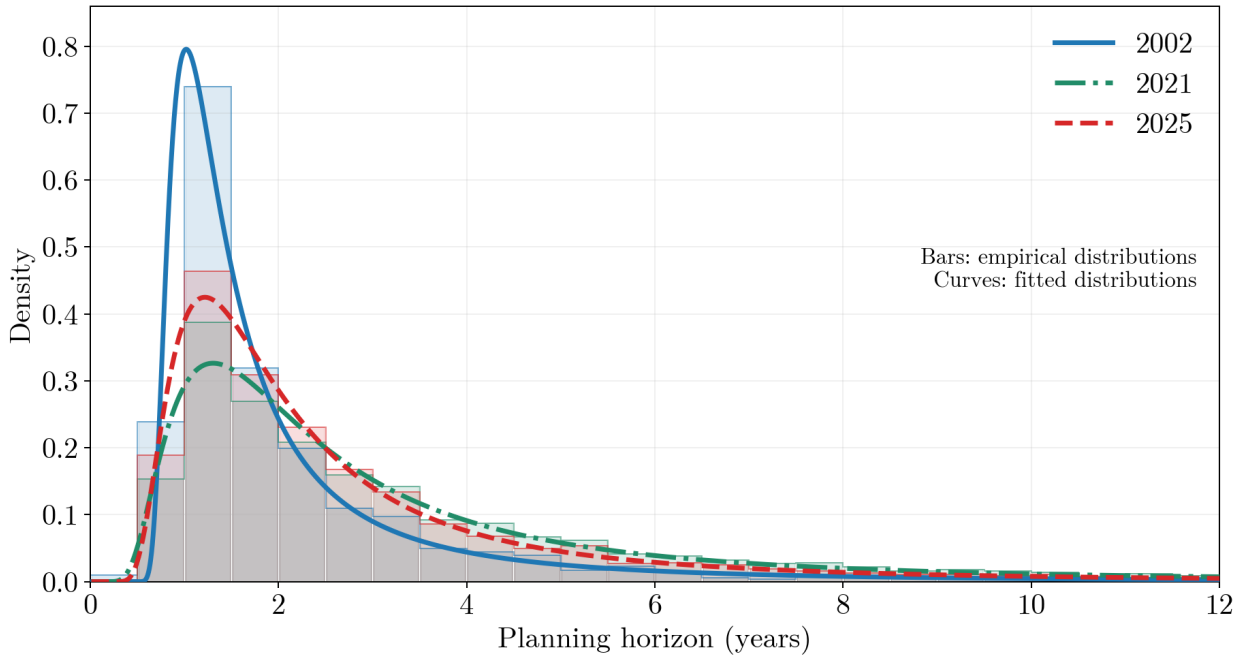
Table E.2 reports the responses under the fitted 2002, 2025, and 2021 distributions. Figure E.7 plots the investment responses at announcement for these distributions across implementation dates. The more elevated 2021 distribution strengthens the announcement response to distant increases further, while reducing cumulative contractions for near-term increases. At five years ahead, the cumulative contractions differ by less than one percent across the three distributions.

[25] The screens target sentences with horizons above twenty years and sentences that determine the call-level 95th percentile for firm-years with horizons above twenty years. They do not constitute an exhaustive verification of every horizon observation.

[26] CDF distance is evaluated at 161 equally spaced log-horizon points between 0.1 and 160 years. We use the Prentice parameterization, restrict the shape parameter to $[-2, -0.08]$, and select a penalty coefficient of 10^{-4} from $\{0, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$. The penalty is the mean squared first difference across all three annual parameter sequences.

[27] The final horizon group collects horizons of 480 quarters or more. Moving this cutoff from 320 to 480 quarters changes the reported model responses by less than 10^{-12}.

Figure E.6: Empirical and fitted planning-horizon distributions



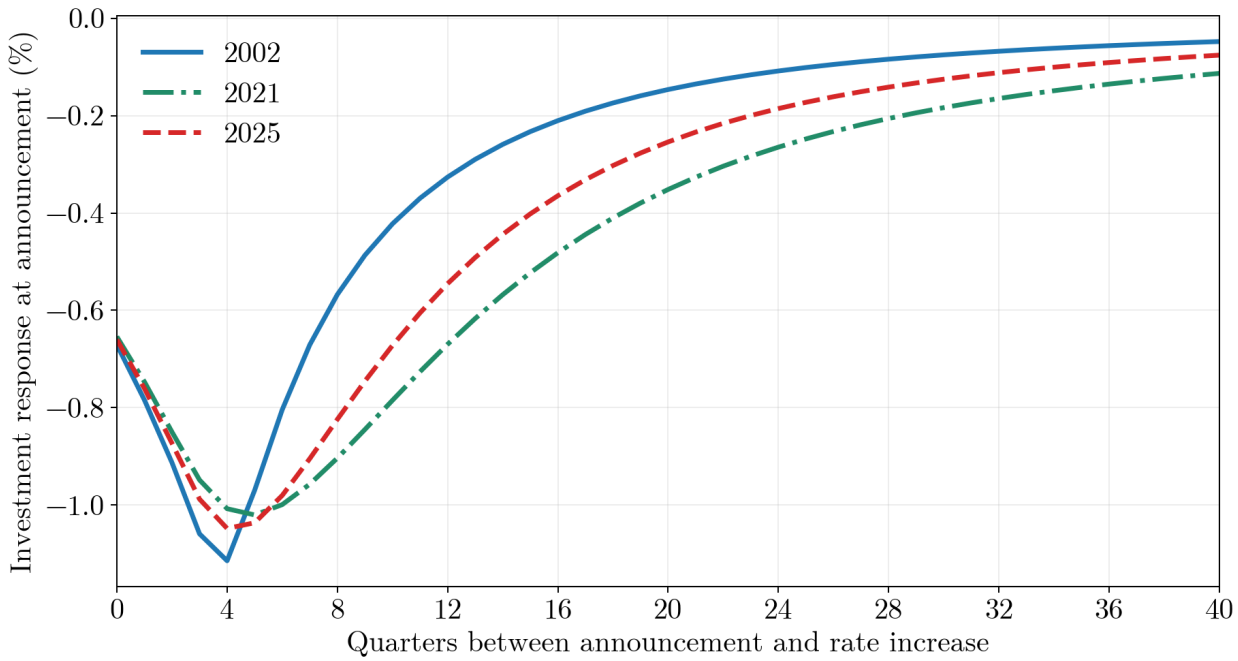
Notes: Bars show empirical densities, and curves show fitted generalized-gamma densities. Blue, green, and red denote 2002, 2021, and 2025, respectively. The figure displays horizons up to twelve years; estimation uses all observations. Firms receive equal weight within each year.

Table E.2: Investment responses under annual planning-horizon distributions

		Investment response		Change in contraction (%)	
Rate increase	Distribution	Announcement	Cumulative	Announcement	Cumulative
Five years ahead	2002	-0.146	-4.434	—	—
	2025	-0.254	-4.452	+73.5	+0.4
	2021	-0.352	-4.409	+140.2	-0.6
One year ahead	2002	-1.114	-1.860	—	—
	2025	-1.048	-1.427	-6.0	-23.3
	2021	-1.007	-1.246	-9.6	-33.0
Immediate	2002	-0.671	-0.275	—	—
	2025	-0.660	-0.218	-1.6	-20.8
	2021	-0.654	-0.194	-2.5	-29.5

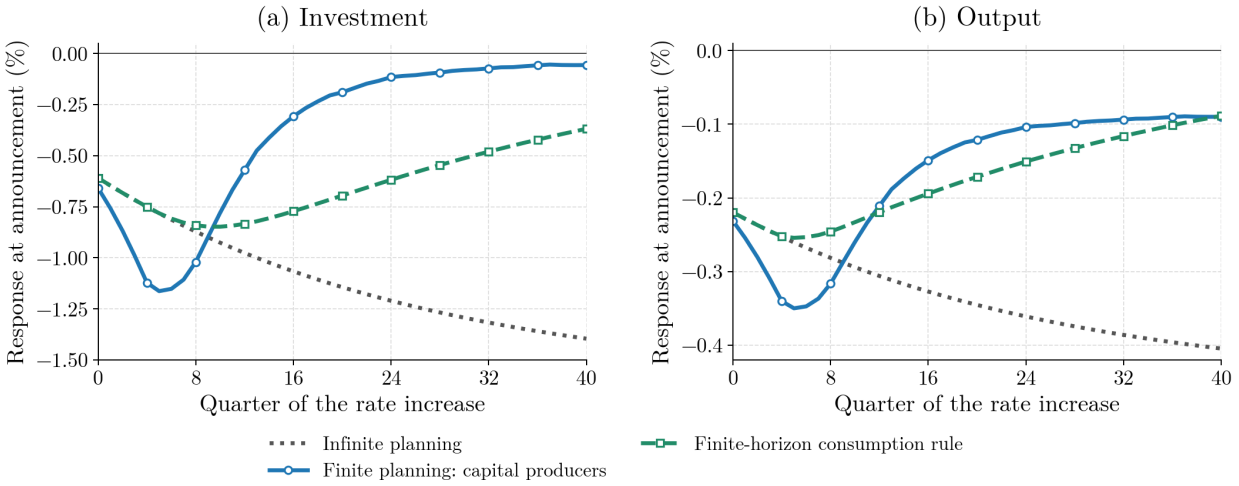
Notes: Each experiment announces at date zero a one-quarter, 50-basis-point annualized real-rate increase. Announcement responses are in percent; cumulative responses sum quarterly investment responses over the full transition. Percentage changes compare absolute contractions with those under the fitted 2002 distribution and are computed before rounding. A positive change denotes a larger contraction. Only the planning-horizon distribution varies across economies.

Figure E.7: Announcement-date investment responses under annual horizon distributions



Notes: The figure reports the investment response at announcement, in percent, as a function of the quarter in which the rate increase occurs. The policy experiment is that of Table 4. The solid blue, dash-dotted green, and dashed red lines use the fitted 2002, 2021, and 2025 distributions, respectively. All other model parameters and the policy experiment are held fixed.

Figure E.8: Finite planning by households versus capital producers



Notes: Each experiment announces at date 0 a one-quarter, 50-basis-point annualized real-rate increase, implemented in the quarter on the horizontal axis. Panels (a) and (b) report investment and output responses at announcement, in percent. Both finite-planning cases use the measured horizon distribution, weighted by firm counts. The solid blue line imposes finite planning on capital producers with rational households. The dashed green line imposes a finite-horizon household consumption rule with rational capital producers. The dotted gray line shows infinite planning.

E.6 Finite Planning by Households and Capital Producers

This appendix compares finite planning by households and capital producers and the implications of alternative horizon distributions. Three cases face the experiment of Section 6.3: a one-quarter, 50-basis-point annualized real-rate increase announced at date zero and implemented at quarter T, with rates at steady state at every other date. In the first, households and capital producers have infinite horizons. In the second, capital producers plan over finite horizons and households are rational, as in Section 5.1. In the third, we impose a finite-horizon household consumption rule while capital producers are rational. Investment is determined conditional on the resulting consumption path.[28] Both finite-horizon cases use the measured distribution with the firm-count weights of Figure 13. Figure E.8 plots the announcement-date responses of investment and output against the implementation date T.

Capital changes the implications of finite planning for the timing of monetary transmission. For the real-rate announcements considered here, finite planning in the New Keynesian model without capital considered by Woodford (2019) weakens the output contraction as implementation is delayed. With capital and infinite planning, delaying implementation deepens the investment contraction. In the model with capital, the measured distribution of finite planning horizons produces the hump in Figure 13.

[28] The consumption rule fixes the announced real-rate path at each date within households' perceived plans. This comparison does not derive households' wealth and asset choices in an economy with capital.

Households. A household whose plan reaches the implementation date carries the rate increase through the Euler equation (13) and, with consumption back at steady state once the increase has passed, cuts consumption at announcement by $\sigma^{-1}\Delta r$, where Δr is the quarterly increase. This is the response of every household in the first two cases, whatever the value of T. Let $c_0(T)$ denote aggregate consumption at announcement, in log deviation from steady state, when the increase is implemented at quarter T. Under the finite-horizon consumption rule, a household with horizon H includes the increase in its plan only when $T \leq H$, so

$$c_0(T) = -\sigma^{-1}\Delta r \Pr(H \geq T), \tag{94}$$

the rational response scaled by the share of households whose plan reaches T. This share cannot increase with T, so the consumption contraction can only weaken as implementation is delayed. Under the measured distribution, consumption is nearly unchanged through quarter 4 because fewer than one percent of firms plan less than four quarters ahead, and then approaches steady state as the increase leaves more plans. For a geometric distribution, the share is ρ^T, giving the same geometric attenuation of the prescribed-rate consumption response as cognitive discounting with coefficient $M = \rho$.

Capital producers. Producers take the consumption path as given and value capital through the recursion (10), with investment given by (11). Rational producers solve the recursion along the equilibrium path with the bounded continuation $q_t, k_t \to 0$. Let $\iota_0(T)$ denote the aggregate investment rate at announcement when the increase is implemented at quarter T. Since capital is predetermined at date zero, it is also the investment response, the object plotted in panel (a) of Figure E.8. With infinite horizons everywhere, the dotted line, the contraction deepens at every T in the plotted range. A more distant increase depresses consumption, and with it rental rates, for more quarters, and the value of capital at announcement falls by more.

Under the finite-horizon consumption rule, additional households respond as the rate increase enters their planning horizons. At date $t \leq T$, aggregate consumption is

$$c_t(T) = -\sigma^{-1}\Delta r \Pr(H \geq T - t),$$

and it returns to steady state after T. Rational capital producers anticipate this changing consumption path. Delaying implementation lengthens the demand shortfall but reduces the share of households responding at announcement. The resulting investment response is hump-shaped, with the deepest contraction at quarter 10. For every implementation date shown, the contraction is no larger than under infinite planning. Thus, imposing finite planning on consumption produces a hump in the model with capital, without producing the larger near-term investment contractions found when capital producers plan over finite horizons.

A producer with horizon H solves the recursion along its perceived path and imposes the terminal condition at the truncation date, and a producer with $H < T$ has no rate increase in its plan and

does not respond at announcement. Let $\iota_{0,H}(T)$ denote its investment rate at announcement. In the second case aggregate investment sums over the types whose plans reach the increase,

$$\iota_0(T) = \sum_{H \geq T} \mu_H \, \iota_{0,H}(T), \tag{95}$$

and moving the implementation date by one quarter changes it by

$$\iota_0(T+1) - \iota_0(T) = \sum_{H \geq T+1} \mu_H \left[\iota_{0,H}(T+1) - \iota_{0,H}(T)\right] - \mu_T \, \iota_{0,T}(T). \tag{96}$$

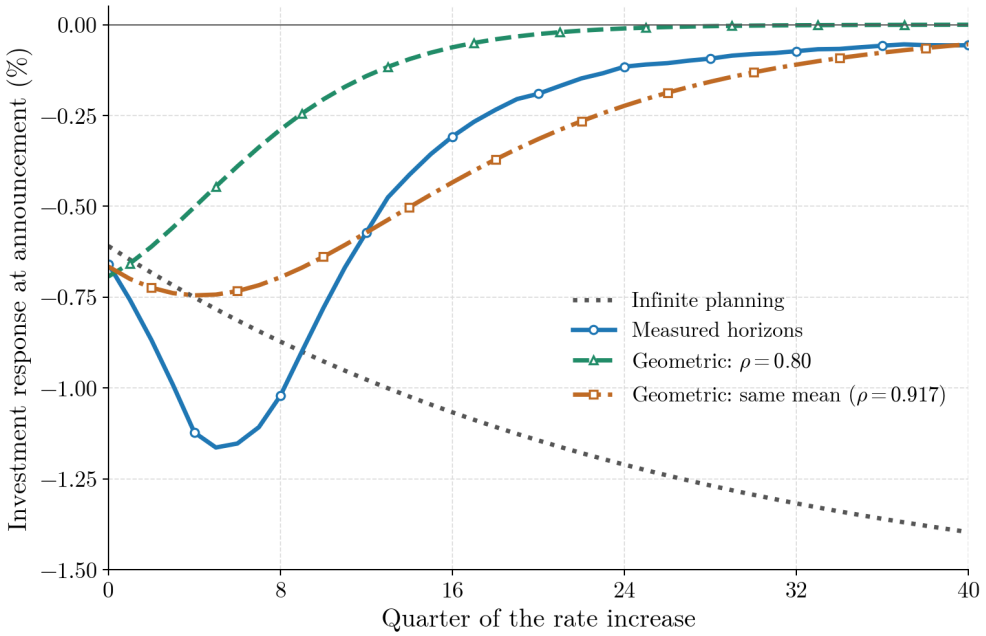
The first term is the change among producers that still see the increase, and the second removes the type whose plan no longer reaches it. These are the two forces of Section 6.3. Under the measured producer distribution, the first term dominates over the near-term dates because the distribution has almost no mass below four quarters and its mode at seven. The contraction is deepest at quarter 5, exceeds the infinite-planning contraction through quarter 8, and is weaker than the contraction under the household consumption rule from quarter 10 on.

An immediate increase lies inside every plan, so at $T = 0$ the household-consumption case coincides with infinite planning, while finite-planning producers contract investment more. A finite-horizon producer's plan contains less of the recovery in returns of capital that follows the increase, which lowers the value of capital at announcement. A fiscal analogue appears in Woodford and Xie (2022), where finite planning limits households' anticipation of the future taxes financing transfers and breaks Ricardian equivalence.

Output. Output combines the two responses through the resource constraint (20). Investment carries the nonmonotonicity and consumption determines how much of it reaches output. With finite-horizon producers, consumption is the same at every T, so output inherits the investment hump and is deepest at quarter 5. Under the finite-horizon consumption rule, the attenuation of consumption offsets part of the investment hump. Output also has its deepest contraction at quarter 5, but the hump is shallower, as panel (b) of Figure E.8 shows.

The distribution of planning horizons. Figure E.9 holds household expectations rational and varies the producer-horizon distribution. The measured distribution has a mean of 11.04 quarters, compared with four quarters under the geometric benchmark with $\rho = 0.80$. A geometric distribution with the same mean requires $\rho = \bar{H}/(1+\bar{H}) \simeq 0.917$, where $\bar{H}$ is the measured mean. This distribution also produces a hump, with the deepest contraction at quarter 4. Even with the same average planning horizon, the measured distribution generates a maximum investment contraction about 56 percent larger than the geometric benchmark (1.16 versus 0.74 percent). Thus, matching the average planning horizon is not enough to reproduce the strength of the investment response. The absence of a hump under the $\rho = 0.80$ benchmark is a result of that calibration, rather than a general property of geometric horizons.

Figure E.9: Investment responses under alternative planning-horizon distributions



Notes: Each experiment announces at date 0 a one-quarter, 50-basis-point annualized real-rate increase, implemented in the quarter on the horizontal axis. The vertical axis reports the investment response at announcement, in percent. Households are rational in every economy. The solid blue line uses the measured producer-horizon distribution, weighted by firm counts. The dashed green line uses geometric horizons with $\rho = 0.80$. The dash-dotted orange line uses geometric horizons with $\rho \simeq 0.917$, matching the measured mean of 11.04 quarters. The dotted gray line shows infinite planning. All other parameters are identical.

E.7 Departures from Rational Expectations with Capital

This appendix compares the two classes of departure from full-information rational expectations (FIRE) discussed in Section 6.3 in the experiment of that section. Attenuation of higher-order general equilibrium feedback weakens successive rounds of endogenous responses without reweighting dates. Attenuation across calendar time discounts or truncates dates within a perceived transition. Angeletos and Lian (2023) show that models under the two headings, despite their different informational and behavioral foundations, deliver the same qualitative prediction in standard macroeconomic environments. They attenuate perceived general equilibrium feedback.[29] Once investment and capital accumulation are introduced, the two classes operate differently.

A higher real-rate path initially lowers investment. Investment below replacement then leaves the capital stock below steady state, raising subsequent returns of capital and inducing investment above replacement until the stock is rebuilt. The dated investment-rate contributions to the indirect effect therefore change sign over the transition. They are negative during the contraction and positive during rebuilding. Attenuation of higher-order feedback does not change the relative weight on these two phases, whereas attenuation across calendar time places less weight on delayed rebuilding. The latter can make the investment-rate component more negative, or larger in absolute value, than under FIRE even when the consumption and capital-stock components are dampened.

E.7.1 The common investment equation

For model $\mathcal{M}$, let $\overline{\mathbb{E}}_t^{\mathcal{M}}$ denote the average expectation entering the capital producer's decision, and let $r_{t+j}^{K,\mathcal{M}}$ denote the rental rate generated by that model. Under this expectation, the shadow value (12) and the investment rule (11) give

$$\iota_t^{\mathcal{M}} = \underbrace{-\frac{1}{\phi\delta}\sum_{j=0}^{\infty}\beta^j\overline{\mathbb{E}}_t^{\mathcal{M}}[r_{t+j}]}_{\mathrm{PE}_t^{\mathcal{M}}} + \underbrace{\frac{\bar{\varrho}^K}{\phi\delta}\sum_{j=1}^{\infty}\beta^j\overline{\mathbb{E}}_t^{\mathcal{M}}[r_{t+j}^{K,\mathcal{M}}]}_{\mathrm{GE}_t^{\mathcal{M}}}. \tag{97}$$

The first term is the direct effect through discounting, denoted PE. The second is the indirect effect through expected returns of capital, denoted GE. Their sum is the total general equilibrium investment response. The sum for the direct effect starts at t and the sum for the indirect effect at $t+1$, because capital installed at t earns its first rental one period later. Under FIRE, $\overline{\mathbb{E}}_t^{\mathrm{FIRE}} = \mathbb{E}_t$, the full-information expectation.

We use reduced-form versions of the mechanisms below inside the same linear capital model. The first class weakens agents' expectations of other agents' responses. The second reduces the weight that agents place on distant outcomes. In every specification the same restriction is imposed on households and capital producers, so a common value of the relevant parameter, or a common planning horizon, governs both consumption and investment decisions. All other equilibrium conditions

[29] This does not necessarily attenuate the total response. The direct and indirect effects may have opposite signs, so weakening the latter can increase the total response in absolute value.

are those of the main text.

E.7.2 Attenuation of higher-order feedback

These models limit the depth of general equilibrium reasoning without discounting or truncating calendar dates. For any endogenous aggregate X, let X^λ denote its equilibrium path under the restriction. Define $\overline{\mathbb{E}}_t^{0,\lambda}[X] \equiv \mathbb{E}_t[X]$ and let $\overline{\mathbb{E}}_t^{h,\lambda}$ denote the average expectation of order h. The reduced form used in the comparison is

$$\overline{\mathbb{E}}_t^{h,\lambda}[X_{t+j}^\lambda] = \lambda^h \mathbb{E}_t[X_{t+j}^\lambda], \qquad h = 1, 2, \ldots \tag{98}$$

Here h is the order of beliefs, whereas j is the number of periods ahead. The parameter $\lambda \in [0, 1]$ measures how strongly higher-order beliefs respond to the shock. FIRE obtains at $\lambda = 1$. We keep first-order beliefs about the announced real-rate path at their full-information values, so the direct effect is unchanged. The capital producer's condition becomes

$$\iota_t^\lambda = \mathrm{PE}_t + \underbrace{\lambda \frac{\bar{\varrho}^K}{\phi\delta} \sum_{j=1}^{\infty} \beta^j \mathbb{E}_t[r_{t+j}^{K,\lambda}]}_{\mathrm{GE}_t^\lambda}. \tag{99}$$

Equation (99) contains one factor λ because the capital producer forms a first-order expectation of future rental rates. The power λ^h in (98) applies only to an expectation of order h. The rental rate $r_{t+j}^{K,\lambda}$ is the equilibrium rental rate under the restriction, not the FIRE rental rate. The restriction therefore weakens general equilibrium feedback without changing the relative weight on early and late dates.

Dispersed information. Angeletos and Lian (2018) derive attenuation from dispersed information under common priors. Agents remain rational but receive noisy signals, and λ is tied to the precision and commonality of their information. We use a constant-λ specialization without learning.

Heterogeneous priors. Angeletos and Sastry (2021) obtain the same reduced form from heterogeneous priors. Agents observe the policy message but assign probability $p \in [0, 1]$ that others observe it as well, so p measures perceived common knowledge and not signal precision. We set $p = \lambda$, so the two coincide by construction despite their different informational foundations.

Level-k thinking. Farhi and Werning (2019) anchor level-zero beliefs at the pre-announcement equilibrium. Level one responds only to the real-rate path, and each higher level best responds to the path conjectured at the preceding level. Successive levels add rounds of general equilibrium reasoning while retaining the full calendar path. In our capital model, however, every finite level leaves a permanent capital shortfall, so the resulting transition does not return to steady state.

Reflective equilibrium. García-Schmidt and Woodford (2019) replace a discrete reasoning level with a Poisson distribution over levels. Its mean n measures reflection depth. Agents then best respond to the Poisson-weighted average path. Like level-k thinking, reflective equilibrium limits rounds of general equilibrium reasoning but retains the complete calendar path.

E.7.3 Attenuation across calendar time

Cognitive discounting and finite planning horizons change the weight on different dates. Let G denote cognitive discounting with attention $M \in [0,1]$, and let (W,H) denote a plan of horizon H. For any variable Z along the relevant equilibrium path, the two rules are

$$\begin{aligned} \overline{\mathbb{E}}_t^{G}[Z_{t+j}] &= M^j \mathbb{E}_t[Z_{t+j}], \quad j \geq 0, \\ \overline{\mathbb{E}}_t^{W,H}[Z_{t+j}] &= \mathbb{E}_t[Z_{t+j}], \qquad 0 \leq j \leq H. \end{aligned} \tag{100}$$

Cognitive discounting gives a smaller weight to a more distant date. A finite-horizon plan omits dates after $t+H$; it does not predict that the variables at those dates are zero. Thus M and H operate on calendar distance j, whereas λ in (98) operates on belief order h. Substituting the two rules into (97) gives, for cognitive discounting,

$$\iota_t^G = \underbrace{-\frac{1}{\phi\delta}\sum_{j=0}^{\infty}(\beta M)^j \mathbb{E}_t[r_{t+j}]}_{\mathrm{PE}_t^G} + \underbrace{\frac{\bar{\varrho}^K}{\phi\delta}\sum_{j=1}^{\infty}(\beta M)^j \mathbb{E}_t[r_{t+j}^{K,G}]}_{\mathrm{GE}_t^G}, \tag{101}$$

and, for a finite-horizon plan with horizon H,

$$\iota_t^{W,H} = \underbrace{-\frac{1}{\phi\delta}\sum_{j=0}^{H}\beta^j \mathbb{E}_t[r_{t+j}]}_{\mathrm{PE}_t^{W,H}} + \underbrace{\frac{\bar{\varrho}^K}{\phi\delta}\sum_{j=1}^{H}\beta^j \mathbb{E}_t[r_{t+j}^{K,W,H}]}_{\mathrm{GE}_t^{W,H}}. \tag{102}$$

The superscripts on r^K identify the equilibrium generated by each rule. Unlike (99), both rules may change the direct effect, because they discount or truncate the real-rate path itself.

Cognitive discounting. Gabaix (2020) assigns cognitive weight M^j to a deviation j periods ahead. We impose the same attention parameter M on households and capital producers. FIRE obtains at $M = 1$.

Finite planning horizons. Woodford (2019) gives agents a finite planning window and a simple terminal continuation value. In each horizon-H economy, households and capital producers share the terminal date $t+H$. For the geometric comparison, we solve each such economy separately and

average its outcome using $\Pr(\mathsf{H} = H) = (1-\rho)\rho^H$, $\rho \in [0,1)$,

$$\iota_t^{W,\rho} = \sum_{H=0}^{\infty} (1-\rho)\rho^H \iota_t^{W,H}. \tag{103}$$

The same averaging is applied to the direct and indirect effects. This is an average of separately solved common-horizon economies. It is not the economy of Section 5.1, in which horizon types interact and households have rational expectations, and it is not the finite-household economy of Appendix E.6. When $\rho = M$, the fraction of plans that reach date $t+j$ equals the cognitive-discounting weight, $\rho^j = M^j$. The two models can nevertheless differ, because every finite horizon generates its own equilibrium path. FIRE is approached as $\rho \uparrow 1$.

E.7.4 Why attenuation across calendar time can overshoot

Equation (22) writes the rental rate as $r_t^K = A_c c_t + A_k k_t + A_\iota \iota_t$, with $A_c, A_\iota > 0$ and, under (71), $A_k < 0$. Lower consumption and a lower investment rate reduce the rental rate, while a capital shortfall raises it. Substituting this decomposition into the term for the indirect effect in (97) defines, for $x \in \{c, k, \iota\}$,

$$\mathrm{GE}_{x,t}^{\mathcal{M}} \equiv \frac{\bar{\varrho}^K}{\phi\delta} \sum_{j=1}^{\infty} \beta^j \overline{\mathbb{E}}_t^{\mathcal{M}}[A_x x_{t+j}^{\mathcal{M}}], \qquad \mathrm{GE}_t^{\mathcal{M}} = \mathrm{GE}_{c,t}^{\mathcal{M}} + \mathrm{GE}_{k,t}^{\mathcal{M}} + \mathrm{GE}_{\iota,t}^{\mathcal{M}}. \tag{104}$$

We say that a component overshoots FIRE when it is larger in absolute value than its FIRE counterpart. This comparison does not by itself order total investment, because the direct effect and the other components of the indirect effect may also change.

The following diagnostic isolates the role of calendar weights. Hold the FIRE transition fixed and let $w_j \in [0,1]$ be the weight placed on date j, with $w_j = M^j$ under cognitive discounting and $w_j = \mathbf{1}\{j \le H\}$ under a fixed horizon. For $x \in \{c, k, \iota\}$, reweighting the FIRE component gives

$$\underbrace{\frac{\bar{\varrho}^K}{\phi\delta} \sum_{j=1}^{\infty} \beta^j w_j A_x \mathbb{E}_0[x_j^{\mathrm{FIRE}}]}_{\text{calendar-reweighted FIRE path}} - \mathrm{GE}_{x,0}^{\mathrm{FIRE}} = -\frac{\bar{\varrho}^K}{\phi\delta} \sum_{j=1}^{\infty} \beta^j (1-w_j) A_x \mathbb{E}_0[x_j^{\mathrm{FIRE}}]. \tag{105}$$

If the dated terms $A_x \mathbb{E}_0[x_j^{\mathrm{FIRE}}]$ all have the same sign, lower weights move their sum toward zero. If the terms change sign, the direction is ambiguous, because early and late contributions are reweighted by different amounts.

The investment path changes sign because the capital stock adjusts gradually. Before and at implementation, the higher real-rate path lowers investment. The resulting capital shortfall raises returns of capital and, once the rate increase has passed, induces investment above replacement. Hence the dated investment-rate contributions are negative during the contraction and positive during rebuilding. The dated consumption and capital-stock contributions are respectively nonpositive and nonnegative throughout the transition studied below.

Under attenuation across calendar time, the delayed rebuilding phase receives less weight than the preceding contraction. The investment-rate component of the indirect effect can therefore become more negative than under FIRE and can increase in absolute value. The same force is absent from the consumption and capital-stock components, because their dated contributions do not change sign in this transition.

Equation (105) holds the FIRE path fixed only to make this logic transparent. The figure recomputes the equilibrium under every expectation rule. Its componentwise ordering is therefore a quantitative result, not a theorem for all capital models.

E.7.5 Quantitative comparison

Conditional on the announced shock, the transition is deterministic, and we suppress expectation signs in describing the figure. The economy begins in steady state. A one-quarter, 50-basis-point annualized real-rate increase is announced at date 0 and implemented at quarter $T \in \{0, \ldots, 40\}$, the experiment of Section 6.3. The structural parameters are those of Table 3. The attenuation parameters take the value of the household attention coefficient in Gabaix (2020), $M = \rho = \lambda = p = 0.85$, with reflection depth $n = 1$. Common numerical values aid comparison but do not equate the frictions.

Figure E.10 recomputes the equilibrium after imposing the same expectation restriction on households and capital producers and applies the decomposition in (97). The left column reports attenuation of higher-order feedback and the right column attenuation across calendar time. Panels (a) and (b) show total investment, panels (c) and (d) the direct effect, and panels (e) and (f) the indirect effect. Each row uses a common vertical scale across the two groups. The infinite-planning benchmark is the same as in Figures 13 and E.8. The geometric curve averages separately solved common-horizon economies, as described above; it differs from the heterogeneous-producer economy with rational households in Figure E.9.

Begin with FIRE. As implementation is postponed, the direct effect remains negative but shrinks slightly, because the one-period rate increase is discounted for more periods. Consumption remains weak for longer before implementation, so its contribution to the indirect effect becomes more negative. Over the same interval, investment below replacement leaves a larger capital shortfall. Since $A_k < 0$, this scarcity raises future returns of capital, and the capital-stock contribution becomes more positive. The same capital shortfall explains the investment-rate contribution. Investment is initially below replacement but later rises above replacement as the capital stock returns to steady state. For a near-term increase the discounted rebuilding phase dominates the initial contraction, so $\mathrm{GE}^{\mathrm{FIRE}}_{\iota,0}$ is positive. As implementation is delayed, rebuilding occurs farther in the future, and from quarter 21 the component is negative. The indirect effect as a whole is positive for an immediate increase, when only the recovery lies ahead, and increasingly negative as the increase moves out, so date-0 investment becomes more contractionary throughout the range.

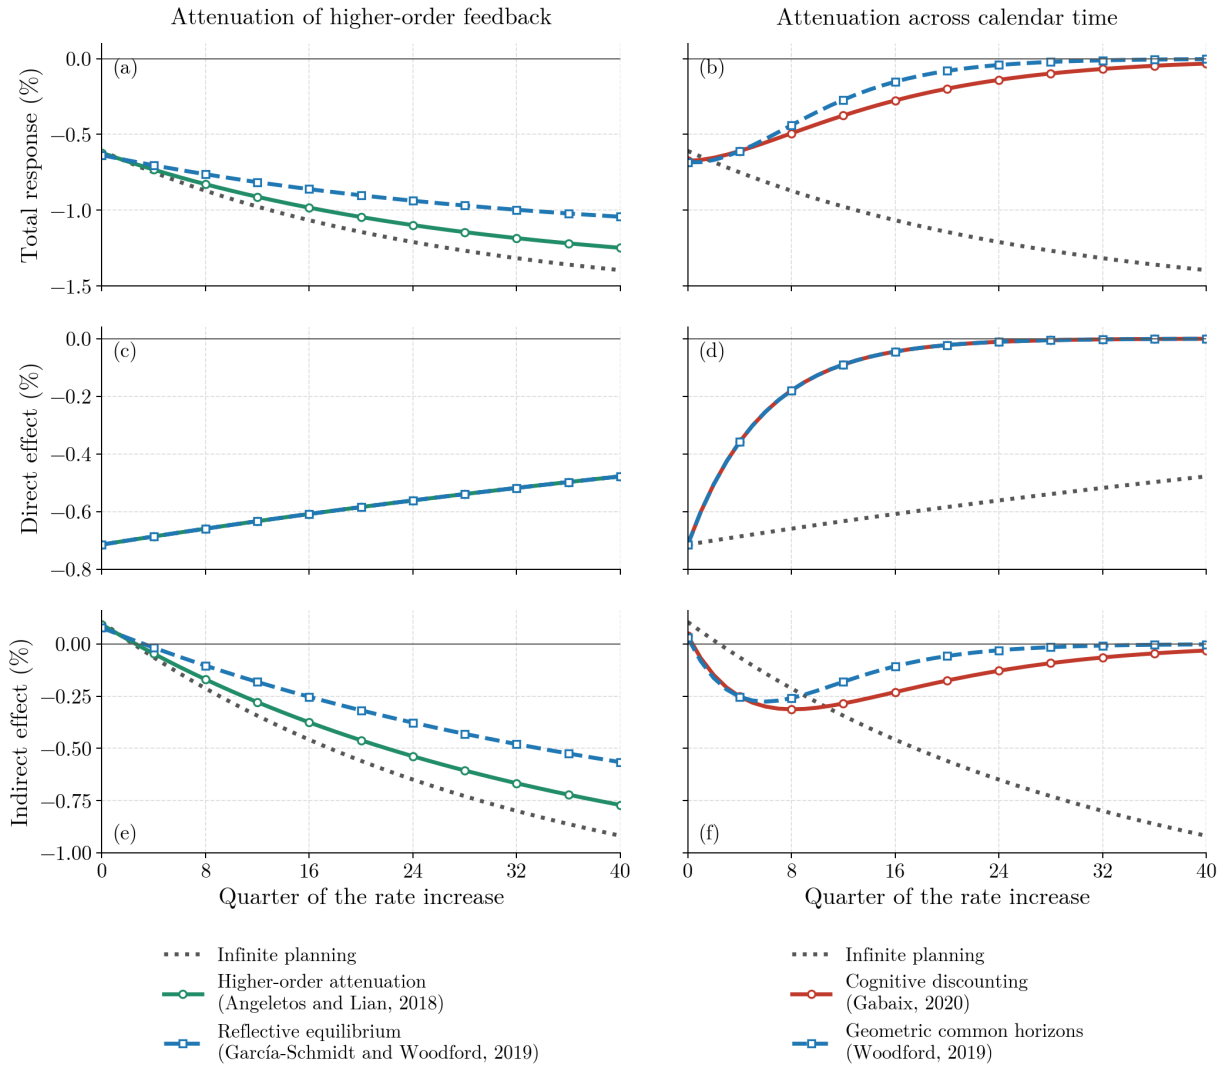
Attenuation of higher-order feedback and reflective equilibrium leave the direct effect unchanged, because agents assign the full-information weight to the announced real-rate path. Their direct-

effect curves therefore coincide with FIRE in panel (c). The indirect effect in panel (e) follows the profile of FIRE at a smaller scale. Date-0 investment still deepens with delay, only more slowly. For an immediate increase, the contraction is slightly larger than under infinite planning because the positive indirect effect is weaker. Under this calibration, the restrictions dampen the response to distant news but do not produce a hump.

Cognitive discounting and geometric finite planning impose attenuation across calendar time. They reduce the direct effect by placing less weight on distant rate increases, and with $M = \rho$ their direct effects coincide in panel (d). Their consumption contributions to the indirect effect also coincide, while the capital-stock and investment-rate contributions can differ. Limited anticipation of rebuilding can make the indirect effect more negative than under FIRE for near-term increases, as panel (f) shows. Nevertheless, date-0 investment weakens with delay because the direct effect fades geometrically and the indirect effect does not offset that decline. Neither benchmark produces a hump at this calibration. This result concerns the common-horizon comparison with restrictions on both households and capital producers; Figure E.9 shows that geometric producer horizons can produce a hump at a higher ρ with rational households.

Level-k thinking with a finite number of rounds is not on the figure. Level one contains only the direct effect, and further temporary-equilibrium iterations would add complete rounds of general equilibrium feedback. In this capital model those finite-level paths are not admissible bounded paths. The permanent capital shortfall generated at level one produces polynomial investment tails of increasing degree at later levels. The discounted date-0 value of any fixed level is finite, but the full path does not return to steady state and leaves the linear neighborhood. We therefore omit the curve rather than report a cutoff-dependent approximation.

Figure E.10: Investment responses under alternative expectation restrictions



Notes: Each experiment announces at date 0 a one-quarter, 50-basis-point annualized real-rate increase, implemented in the quarter on the horizontal axis. Expectation restrictions apply to both households and capital producers. The left column shows attenuation of higher-order feedback; the right column shows attenuation across calendar time. Rows report total investment responses, direct effects, and indirect effects, in percent. The calibration uses $M = \rho = \lambda = p = 0.85$ and reflection depth $n = 1$. The geometric curve averages separately solved economies in which households and capital producers share a common planning horizon. All direct-effect curves coincide in panel (c); cognitive discounting and geometric planning coincide in panel (d).

E.8 Learning Extension for Capital-Producing Firms

The benchmark finite-planning model closes every perceived countdown tree with the common steady-state Bellman tail in (5). That boundary condition is useful for isolating horizon heterogeneity, but it also assumes that firms keep using the common steady-state continuation value after the economy has moved away from that steady state. This extension relaxes only that margin. Firms still solve the same horizon-H_i countdown problem and still implement only the current-node choice, but each date-t planning exercise also produces an estimate of the continuation value of installed capital. Firms then use that estimate to update the continuation value associated with the producer's return-discounted objective, adapting the value-function updating procedure of Woodford (2019).

Belief Updating. Let $v_{K,t}(K)$ denote the common learned terminal value for installed capital used in date-t countdown trees, and let $v^{\text{est}}_{K,t}(K)$ denote the value function estimate produced by those date-t trees. The timing follows Woodford (2019)'s Section IV: firms use $v_{K,t}$ when solving at date t, compute $v^{\text{est}}_{K,t}$ from that solution, and update the terminal value for date $t+1$ according to

$$v_{K,t+1}(K) = (1-\gamma_K)v_{K,t}(K) + \gamma_K v^{\text{est}}_{K,t}(K), \qquad \gamma_K \in (0,1). \tag{106}$$

The gain γ_K governs how quickly firms correct the terminal value used after the explicit planning window ends.

The firm first-order conditions depend on the marginal terminal value. Around the deterministic steady state, write this marginal value for agent i's terminal stock as

$$\log\left(\frac{v'_{K,t}(K_i)}{\bar{R}}\right) = \nu_{K,t} + \chi_{K,t}\hat{K}^{\text{term}}_i, \tag{107}$$

where $\hat{K}^{\text{term}}_i \equiv \log(K_i/\bar{K}_i)$ and $\bar{R}$ is the steady-state marginal value of terminal capital from (5). The intercept $\nu_{K,t}$ is the learned level of the terminal marginal value. The slope $\chi_{K,t}$ captures any learned dependence of that marginal value on the terminal capital stock.

For agent i with horizon $H_i = H$, the root envelope derivative of the estimated value function is

$$V^{H,\text{est}}_{i,K,t} = r^{K,H}_{t|t} - \Psi_K(I^H_{i,t|t}, K^{H+1}_{i,t|t}) + (1-\delta)Q^H_{i,t|t}. \tag{108}$$

This is the marginal payoff from one more unit of capital at the root of the date-t tree: the extra rental income, the adjustment-cost envelope term, and the undepreciated continuation value. Linearizing (108) and using the quadratic identity from Appendix E.2 gives the agent-specific estimate

$$\nu^{H,\text{est}}_{K,i,t} = \beta\left(\bar{r}^K \hat{r}^{K,H}_{t|t} + \hat{Q}^H_{i,t|t}\right). \tag{109}$$

The coefficient on $\hat{Q}^H_{i,t|t}$ is one because the adjustment-cost envelope contributes $\delta\hat{Q}^H_{i,t|t}$ and the

undepreciated continuation value contributes $(1-\delta)\hat{Q}^H_{i,t|t}$. The reported learning exercise maintains a linear terminal value in capital, so the estimated marginal value is locally independent of the individual terminal stock:

$$\chi^{H,\text{est}}_{K,t} = 0. \tag{110}$$

This is a restriction on the terminal-value approximation, not a claim that equilibrium marginal products are independent of aggregate capital. Let $\mathcal{H}$ be the set of root planning horizons, and let ω^V_H be the weight placed on horizon-H estimates, with $\omega^V_H \geq 0$ and $\sum_{H\in\mathcal{H}} \omega^V_H = 1$. In the quantitative implementation, the natural capital-weighted choice is

$$\omega^V_H = \sum_{i\in\mathcal{I}:H_i=H} \omega_i.$$

For each horizon, define the capital-weighted estimate among agents with that horizon as

$$\nu^{H,\text{est}}_{K,t} = \frac{1}{\omega^V_H} \sum_{i\in\mathcal{I}:H_i=H} \omega_i \nu^{H,\text{est}}_{K,i,t},$$

with the convention that horizons with $\omega^V_H = 0$ are omitted. The aggregate estimated coefficients are

$$\begin{aligned} \nu^{\text{est}}_{K,t} &= \sum_{H\in\mathcal{H}} \omega^V_H \nu^{H,\text{est}}_{K,t}, \\ \chi^{\text{est}}_{K,t} &= \sum_{H\in\mathcal{H}} \omega^V_H \chi^{H,\text{est}}_{K,t} = 0. \end{aligned} \tag{111}$$

Equation (106) then implies the coefficient updates

$$\begin{aligned} \nu_{K,t+1} &= (1-\gamma_K)\nu_{K,t} + \gamma_K \nu^{\text{est}}_{K,t}, \\ \chi_{K,t+1} &= (1-\gamma_K)\chi_{K,t}. \end{aligned} \tag{112}$$

Thus an initial $\chi_{K,0} = 0$ remains zero, and any nonzero initial slope decays geometrically. The learning state that matters for the quantitative extension is the intercept $\nu_{K,t}$.

Within-Tree Decomposition. This subsection restores the forecast-origin notation $\tau|t$ to make the link to Woodford (2019)'s Section IV explicit. For any first-order forecast object $\hat{x}^h_{\tau|t}$ in a date-t tree with h explicit periods remaining, decompose

$$\hat{x}^h_{\tau|t} = \widetilde{x}^h_{\tau|t} + \bar{x}^h_{\tau|t}. \tag{113}$$

The tilde component is the solution of the same date-t tree with the same root capital stocks and the same exogenous paths, but with $\nu_{K,t} = 0$. The bar component is the belief-induced difference created by the learned terminal marginal value. Both $\widetilde{x}^h_{\tau|t}$ and $\bar{x}^h_{\tau|t}$ are components of the hatted

first-order deviation $\hat{x}^h_{\tau|t}$, not level variables. A bar on a time-indexed object denotes this belief component; an unindexed object such as $\bar{R}$, $\bar{r}^K$, or $\bar{K}$ remains a steady-state value.

With reciprocal-return valuation and the learned terminal intercept and slope, the zero-horizon boundary condition becomes

$$\hat{Q}^0_{i,\tau|t} = -E_\tau \hat{r}^0_{\tau|t} + \nu_{K,t} + \chi_{K,t} E_\tau \hat{K}^0_{i,\tau+1|t}. \tag{114}$$

The reported calibration imposes $\chi_{K,t} = 0$, so (114) reduces to the terminal intercept used below. For $h \geq 1$, the backward marginal-value recursion is the same as in Appendix E.2:

$$\hat{Q}^h_{i,\tau|t} = E_\tau \left[\beta \bar{r}^K \hat{r}^{K,h-1}_{\tau+1|t} - \hat{r}^h_{\tau|t} + \beta \hat{Q}^{h-1}_{i,\tau+1|t} \right]. \tag{115}$$

Combining (113)–(115) gives the zero-learning and belief-induced systems:

$$\begin{aligned} \widetilde{Q}^0_{i,\tau|t} &= -E_\tau \widetilde{r}^0_{\tau|t}, \\ \widetilde{Q}^h_{i,\tau|t} &= E_\tau \left[\beta \bar{r}^K \widetilde{r}^{K,h-1}_{\tau+1|t} - \widetilde{r}^h_{\tau|t} + \beta \widetilde{Q}^{h-1}_{i,\tau+1|t} \right], \qquad h \geq 1, \end{aligned} \tag{116}$$

and

$$\begin{aligned} \bar{Q}^0_{i,\tau|t} &= \nu_{K,t}, \\ \bar{Q}^h_{i,\tau|t} &= E_\tau \left[\beta \bar{r}^K \bar{r}^{K,h-1}_{\tau+1|t} + \beta \bar{Q}^{h-1}_{i,\tau+1|t} \right], \qquad h \geq 1. \end{aligned} \tag{117}$$

The real-rate path is the same in the learning and zero-learning trees, so the belief-induced recursion contains no direct real-rate term.

The root marginal value used for current investment in a horizon-H date-t tree is therefore

$$\hat{Q}^H_{i,t|t} = \widetilde{Q}^H_{i,t|t} + \bar{Q}^H_{i,t|t}. \tag{118}$$

The investment and capital laws split in the same way. At any node,

$$\begin{aligned} \hat{I}^h_{i,\tau|t} &= \hat{K}^{h+1}_{i,\tau|t} + \frac{1}{\phi\delta} \hat{Q}^h_{i,\tau|t}, \\ \hat{K}^h_{i,\tau+1|t} &= (1-\delta) \hat{K}^{h+1}_{i,\tau|t} + \delta \hat{I}^h_{i,\tau|t}. \end{aligned} \tag{119}$$

Hence the belief-induced investment and capital components are

$$\begin{aligned} \bar{I}^h_{i,\tau|t} &= \bar{K}^{h+1}_{i,\tau|t} + \frac{1}{\phi\delta} \bar{Q}^h_{i,\tau|t}, \\ \bar{K}^h_{i,\tau+1|t} &= (1-\delta) \bar{K}^{h+1}_{i,\tau|t} + \delta \bar{I}^h_{i,\tau|t}. \end{aligned} \tag{120}$$

Within a single date-t decomposition, the learning and zero-learning trees start from the same

inherited capital vector:

$$\widetilde{K}^{H+1}_{i,t|t} = \hat{K}_{i,t}, \qquad \bar{K}^{H+1}_{i,t|t} = 0. \tag{121}$$

Past learning can affect the actual inherited state $\hat{K}_{i,t}$ through earlier implemented investment, but the date-t tilde and bar trees do not start from different physical capital stocks.

Static Block. The belief component of the rental-rate path is pinned down by the same general equilibrium block used in the main model. In the learning comparison, consumption and productivity paths are the same in the learning and zero-learning trees, so their bar components are zero. The belief-induced static block is

$$\begin{aligned}
\bar{K}^{h+1}_{\tau|t} &= \sum_{i\in\mathcal{I}} \omega_i \bar{K}^{h+1}_{i,\tau|t}, \\
\bar{Y}^{h}_{\tau|t} &= \sum_{i\in\mathcal{I}} s_{I,i} \bar{I}^{h}_{i,\tau|t}, \\
\bar{N}^{h}_{\tau|t} &= \frac{\bar{Y}^{h}_{\tau|t} - \alpha \bar{K}^{h+1}_{\tau|t}}{1-\alpha}, \\
\bar{W}^{h}_{\tau|t} &= \varphi \bar{N}^{h}_{\tau|t}, \\
\overline{mc}^{h}_{\tau|t} &= \bar{W}^{h}_{\tau|t} - \alpha \bar{K}^{h+1}_{\tau|t} + \alpha \bar{N}^{h}_{\tau|t}, \\
\bar{r}^{K,h}_{\tau|t} &= \overline{mc}^{h}_{\tau|t} + (\alpha-1)\bar{K}^{h+1}_{\tau|t} + (1-\alpha)\bar{N}^{h}_{\tau|t}.
\end{aligned} \tag{122}$$

Equations (117), (120), and (122) are the capital-producing-firm analogue of Woodford's bar system. For a fixed date-t tree, they solve the effect of the current belief coefficient $\nu_{K,t}$ on forecast investment, capital, output, and rental rates. Future belief updates are not forecast inside the tree; they occur only after the date-t current-node choice has been implemented.

Calendar-Time Implementation. Only the current-node choice is implemented in calendar time. If agent i has planning horizon H_i, then

$$\hat{I}_{i,t} = \hat{I}^{H_i}_{i,t|t}, \qquad \hat{K}_{i,t+1} = (1-\delta)\hat{K}_{i,t} + \delta \hat{I}_{i,t}. \tag{123}$$

After implementation, the same date-t root objects determine the estimate in (109), the cross-horizon average in (111), and the update in (112). The sequence of state transitions is

$$\begin{aligned}
\nu_{K,t} \quad &\longrightarrow \quad \{\bar{Q}^{h}_{i,\tau|t}, \bar{I}^{h}_{i,\tau|t}, \bar{K}^{h}_{i,\tau|t}\} \\
\{\hat{K}_{i,t}\}_{i\in\mathcal{I}} \quad &\longrightarrow \quad \{\widetilde{Q}^{h}_{i,\tau|t}, \widetilde{I}^{h}_{i,\tau|t}, \widetilde{K}^{h}_{i,\tau|t}\}
\end{aligned} \quad \Longrightarrow \hat{I}_{i,t} = \widetilde{I}^{H_i}_{i,t|t} + \bar{I}^{H_i}_{i,t|t}$$

$$\Longrightarrow \begin{aligned}
&\nu^{\text{est}}_{K,t} \longrightarrow \nu_{K,t+1}, \\
&\hat{K}_{i,t+1} = (1-\delta)\hat{K}_{i,t} + \delta\left(\widetilde{I}^{H_i}_{i,t|t} + \bar{I}^{H_i}_{i,t|t}\right).
\end{aligned}$$

This timing preserves the receding-planning structure of the main model. The learned terminal value affects current Q through the finite-horizon recursion, current investment affects next period's physical capital, and both the belief state and the physical capital vector become state variables for the next calendar date.

Relation to Woodford's Learning System. The introduction of capital creates three differences relative to Woodford's baseline value-function learning system.

First, the bar block is static only with respect to beliefs. In Woodford's Section IV construction, agents solve the date-t planning problem holding the current value-function coefficients fixed throughout the whole date-t tree. They do not forecast future belief updates inside the tree. Future imagined subproblems are therefore copies of today's subproblem with fewer planning steps left, all evaluated under the same current beliefs. The capital-producing-firm extension keeps exactly this timing: firms use the current terminal-value coefficient $\nu_{K,t}$ at every imagined future node. The difference is that capital is a genuine state variable. Planned capital evolves forward inside the tree through planned investment, so the bar variables are solved conditional on $\nu_{K,t}$ and on the within-tree capital law (120).

Second, the tilde block is not dynamically autonomous across calendar time. In Woodford's no-capital model, the tilde component is the outcome that would arise if the value-function belief coefficients were fixed at their steady-state values. It depends only on shocks and the exogenous paths entering the planning problem, so the whole zero-learning sequence can be solved first:

$$\{\widetilde{x}_t^h\}_{t,h}.$$

The bar block is then static conditional on current beliefs. Starting from initial beliefs, the iteration is

$$\nu_t \longrightarrow \bar{x}_t^h \longrightarrow x_t^h = \widetilde{x}_t^h + \bar{x}_t^h \longrightarrow \nu_t^{\text{est}} \longrightarrow \nu_{t+1} \longrightarrow \cdots .$$

In the capital model, that separation holds only within a fixed date-t tree. Conditional on the current root capital vector $\{\hat{K}_{i,t}\}_{i\in\mathcal{I}}$, one can solve the tilde variables in the date-t tree; conditional on the current belief $\nu_{K,t}$, one can solve the bar variables in the same tree. Across calendar time, however, next period's root capital vector depends on full implemented investment, not only on the tilde component:

$$\begin{array}{rcl} \nu_{K,t} & \longrightarrow & \bar{x}_{\tau|t}^h \\ \{\hat{K}_{i,t}\}_{i\in\mathcal{I}} & \longrightarrow & \widetilde{x}_{\tau|t}^h \end{array} \Longrightarrow \hat{x}_{\tau|t}^h = \widetilde{x}_{\tau|t}^h + \bar{x}_{\tau|t}^h \Longrightarrow \begin{array}{l} \nu_{K,t}^{\text{est}} \longrightarrow \nu_{K,t+1}, \\ \hat{K}_{i,t+1} = (1-\delta)\hat{K}_{i,t} + \delta\left(\widetilde{I}_{i,t|t}^{H_i} + \bar{I}_{i,t|t}^{H_i}\right). \end{array}$$

Thus learning is not the only state-like object that links the bar block across calendar dates. Physical capital is an additional state, and because it is updated using full implemented investment, the zero-learning block is no longer autonomous over calendar time.

Third, the learned terminal value can weaken the saddle-path force. In the infinite-horizon firm problem, beginning-of-period capital is predetermined and Tobin's Q is a jump variable. The stable

Table E.3: Investment responsiveness after adding stable capital-producing-firm learning.

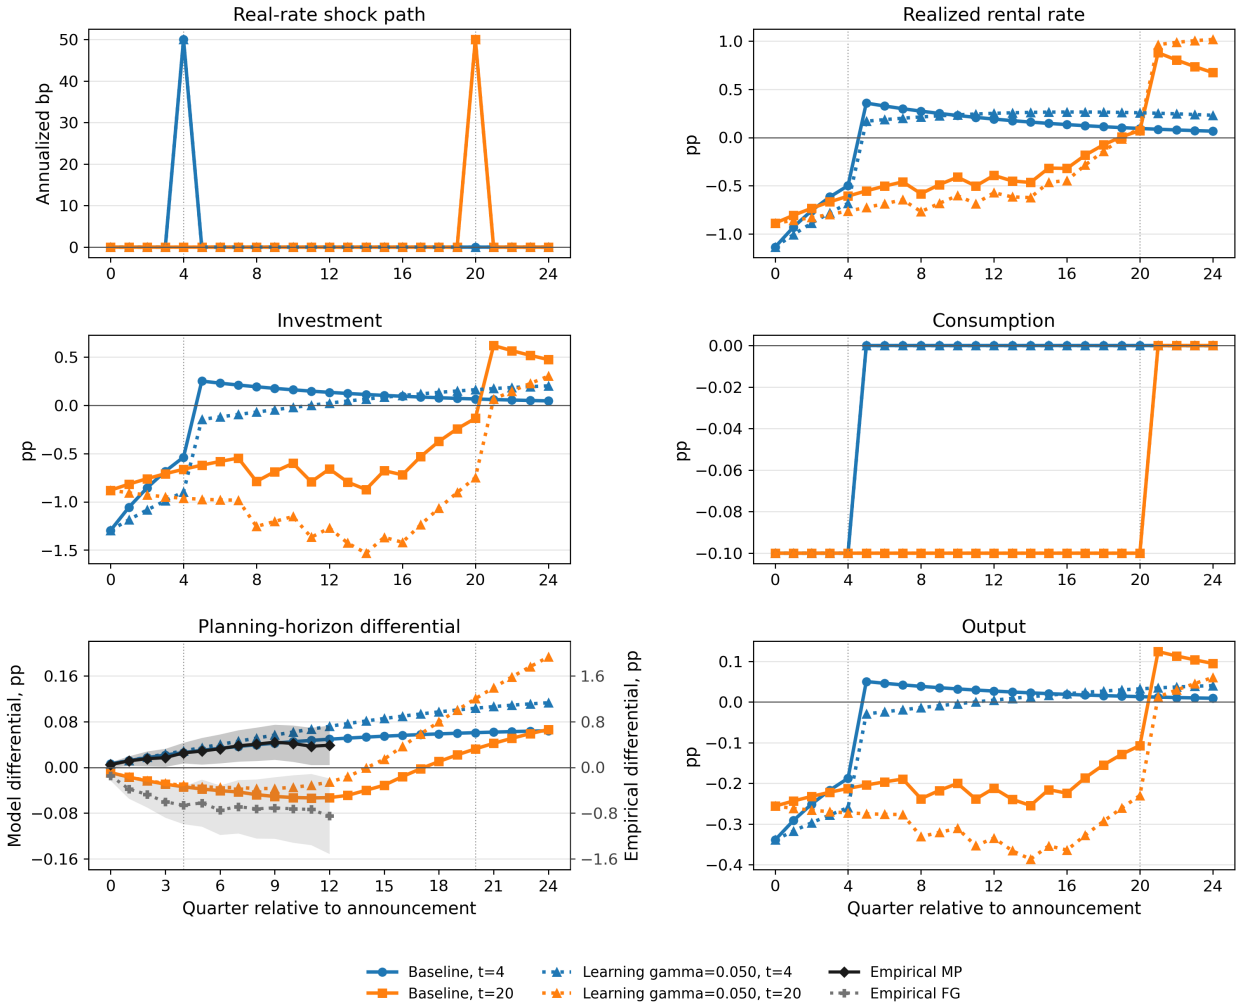
Learning vs. FPH benchmark	Short-term ($t = 4$)	Long-term ($t = 20$)
Impact	+0.0%	+0.0%
Cumulative, full path	+53.0%	+41.7%

Notes: Entries compare the absolute investment response in the benchmark-plus-learning economy with the standardized-shock finite-planning benchmark. The learning gain is $\gamma_K = 0.05$. Both one-quarter, 50 annualized-basis-point rate increases are announced at date 0. The short-term benchmark is implemented at $t = 4$ and the long-term benchmark at $t = 20$.

saddle path chooses Q_t so that a high inherited capital stock lowers expected future marginal products, reduces Q_t, weakens investment, and pulls capital back toward steady state. In this learning extension, Q_t is still forward-looking, but it is pinned down by a finite-horizon recursion with the learned terminal marginal-value intercept $\nu_{K,t}$ rather than by the true infinite-horizon transversality condition. A high $\nu_{K,t}$ raises terminal Q, propagates backward to current Q_t, and stimulates investment. Since implemented investment becomes next period's inherited capital, the belief-induced component is carried forward as a physical state, not only as a current wedge. High capital pushes Q_t down through lower marginal products, while a high learned continuation value pushes Q_t up through the terminal condition. If the gain is large enough, the belief channel can dominate the stabilizing marginal-product channel; investment then remains too high even when capital is already high, and the joint system for $(\{\hat{K}_{i,t}\}_{i\in\mathcal{I}}, \nu_{K,t})$ need not be stable. In the local linear implementation, stability means that the transition matrix for this joint state has spectral radius below one.

Quantitative Result. Table E.3 compares the standardized-shock finite-planning benchmark with a stable benchmark-plus-learning variant that sets $\gamma_K = 0.05$ and starts from $\nu_{K,0} = 0$. Because the initial learning state is zero, learning does not change impact investment. It changes the path only after firms update their terminal marginal values. The full-path cumulative investment response is 53.0 percent larger for the short-term benchmark and 41.7 percent larger for the long-term benchmark. Figure E.11 shows the corresponding response paths: the learning variant is close to the finite-planning benchmark at the announcement date, then departs as the learned terminal value feeds back into later plans.

Figure E.11: Responses with and without learning



Notes: Each experiment announces at date 0 a one-quarter, 50-basis-point annualized real-rate increase, implemented in quarter 4 or 20. Colored solid lines show finite planning without learning; colored dotted lines show the learning variant, which updates the terminal marginal value of capital with gain $\gamma_K = 0.05$. In the planning-horizon panel, model responses use the left axis and empirical current-rate and forward-guidance estimates use the right axis as scale references. Zeros are aligned, and at corresponding heights the right-axis value is ten times the left-axis value.

E.9 Heterogeneous Depreciation Rates

Consider infinite-horizon capital producers with depreciation rates $\delta_i \in (0,1)$ and common real interest rates R_t. Producer i uses the homogeneous installation technology $\Psi_i(I,K) = K\psi_i(I/K)$, normalized at its replacement rate: $\psi_i(\delta_i) = \psi_i'(\delta_i) = 0$ and $\psi_i''(\delta_i) = \phi > 0$, with common curvature ϕ. Its capital earns the common rental rate ϱ_t^K plus the fixed per-unit adjustment $a_i \equiv \delta_i - \delta$, where δ is the baseline depreciation rate. Rental income net of depreciation is therefore $\varrho_t^K + a_i - \delta_i = \varrho_t^K - \delta$ for every type.

The producer solves the infinite-horizon version of the capital-producer problem in Appendix E.1, with rental income $(\varrho_t^K + a_i)K_{i,t}$ and capital accumulation $K_{i,t+1} = (1-\delta_i)K_{i,t} + I_{i,t}$, starting from its steady-state capital stock. Its first-order conditions are

$$Q_{i,t} = 1 + \Psi_{i,I}(I_{i,t}, K_{i,t}), \tag{124}$$

$$Q_{i,t} = R_t^{-1}\left[\varrho_{t+1}^K + a_i - \Psi_{i,K}(I_{i,t+1}, K_{i,t+1}) + (1-\delta_i)Q_{i,t+1}\right]. \tag{125}$$

At the steady state, $\bar{I}_i = \delta_i \bar{K}_i$, $\bar{Q}_i = 1$, and $\bar{\Psi}_{i,K} = 0$. The fixed rental adjustment gives the common condition $\bar{\varrho}^K = \bar{R} - 1 + \delta$, with $\beta = \bar{R}^{-1}$.

Write $x_{i,t} \equiv I_{i,t}/K_{i,t}$ and let $q_{i,t}$ denote the first-order log deviation of $Q_{i,t}$ around one, so $Q_{i,t} = 1 + q_{i,t}$ to first order. Let r_t and r_t^K denote log deviations of R_t and the common unadjusted rental rate ϱ_t^K. The adjustment-cost identities in Appendix E.2 imply $x_{i,t} - \delta_i = q_{i,t}/\phi$ and $\Psi_{i,K}(I_{i,t}, K_{i,t}) = -\delta_i q_{i,t}$ to first order. Substituting into the linearization of (125) yields

$$\begin{aligned} q_{i,t} &= -r_t + \beta\bar{\varrho}^K r_{t+1}^K + \beta\big[\delta_i + (1-\delta_i)\big]q_{i,t+1} \\ &= -r_t + \beta\bar{\varrho}^K r_{t+1}^K + \beta q_{i,t+1}. \end{aligned} \tag{126}$$

The depreciation rate cancels from the recursion. Faster depreciation reduces the value carried forward by surviving capital but increases the installation costs saved by an additional unit of capital; the two contributions offset at first order.

Impose $\lim_{J\to\infty} \beta^J q_{i,t+J} = 0$ for every type and consider common rate and rental paths for which the forward sums converge. For two types i and m, (126) implies $q_{i,t} - q_{m,t} = \beta^J(q_{i,t+J} - q_{m,t+J})$, so the terminal condition gives $q_{i,t} = q_{m,t} \equiv q_t$. Writing $k_{i,t} \equiv \log(K_{i,t}/\bar{K}_i)$, capital accumulation then gives

$$k_{i,t+1} - k_{i,t} = x_{i,t} - \delta_i = \frac{q_t}{\phi} \tag{127}$$

to first order. Cumulative capital-growth responses are thus identical across types. The proportional deviation of the investment rate is $\iota_{i,t} = q_t/(\phi\delta_i)$ and need not be identical, because steady-state investment rates differ.

The equality holds along any common rate and rental paths, including those arising in general equilibrium. It is conditional on producers internalizing installation costs, common curvature, and the fixed rental adjustments; their financing and the equilibrium paths require a separate equilibrium specification.