The macroeconomic effects of technology news: high-frequency evidence from semiconductor conferences*

Naixin Zhang
George Washington University

November 2026

*** PRELIMINARY DRAFT — SUBJECT TO CHANGES ***

Abstract

This paper studies how technological disclosures shape financial markets' macroeconomic expectations. I combine semiconductor-sector returns around recurring technology conferences with large language model (LLM) classifications of keynote transcripts. The classification distinguishes technological capability from demand and adoption, and capability already available from capability yet to arrive. I find that positive conference-day surprises are associated with higher economic growth expectations for both delivered-capability-heavy and promised-capability-heavy keynotes, but their inflation implications differ. Expectation is inflationary for delivered-capability-heavy keynotes and disinflationary for promised-capability-heavy keynotes. These joint responses are consistent with the interplay between anticipated demand expansion and improvements in productive capacity. Subsequent increases in capital-goods orders in both cases, and lower core prices alongside higher production following promised-capability-heavy keynotes, provide complementary evidence. The findings suggest that the economic content of technological disclosures matters for explaining why technology news may have different inflation implications.

Keywords: Technology news, Semiconductor, Macroeconomic expectations, Large Language Model, High-frequency identification.

JEL Classification: E31, E44, G14, O33.

*Department of Economics, George Washington University, 2115 G Street NW, Washington, DC 20052, USA; email: nxzhang@gwu.edu. I am grateful to Thomas Drechsel, Roberto Samaniego, and Tara Sinclair for their invaluable discussions, comments, and suggestions. I also thank participants at the George Washington University Macro Lunch Seminar for helpful comments.

1 Introduction

Does good news about technology raise or lower inflation? This is a longstanding question with new stakes. Technological progress expands what an economy can produce, but learning about that progress can change spending before production improves. Households and firms may bring expenditure forward, while expectations of cheaper production can put downward pressure on prices. The balance between these forces matters for monetary policy, investment decisions, and the interpretation of technology-driven economic expansions.

Artificial intelligence (AI) makes this question particularly pressing. The prospect of cheaper computation and greater productive capacity comes alongside demand for chips, data centers, and complementary investment. The expenditure required to adopt a technology may arrive well before its productivity benefits spread through the economy. At the same time, announcements of future capabilities can change expectations about production costs before those capabilities become available. An important distinction is therefore whether the disclosed technological capabilities are already available or still forthcoming. Both types of news may improve growth prospects, but their implications for anticipated demand expansion and future productive capacity may differ.

Despite the importance of these competing forces, direct evidence on their relative role when technological information arrives remains limited. Existing technology-news measures typically infer information from productivity dynamics, financial markets, or patent-based measures without directly observing the content of the underlying disclosures. It is consequently difficult to determine whether different inflation responses reflect differences in information content, economic transmission, or identifying assumptions. Progress requires observing both when technological information reaches investors and what they learn.

This paper is the first, to my knowledge, to systematically study the macroeconomic effects of disclosures at recurring technology conferences. I examine how these disclosures shape financial markets' expectations of economic growth and inflation, and how the responses depend on announcement content. My setting is semiconductor conference keynotes, where firms announce processor performance, computing architectures, product roadmaps, and commercial applications. Their recurring schedules provide observable disclosure opportunities even when the announcements themselves remain uncertain. I combine conference calendars with semiconductor-sector returns, keynote transcripts, dividend futures, and inflation swaps to measure the information disclosed and the associated macroeconomic repricing.

The final baseline sample contains 114 keynotes on 112 trading days. Events are selected based on institutional characteristics rather than realized returns or subsequently extracted content. I exclude identified overlaps with FOMC statements,

CPI releases, and issuer earnings calls, as well as non-trading-day observations. The baseline surprise is the raw daily semiconductor-sector return on the retained event days. I classify keynote statements using a large language model (LLM) that does not observe asset prices or macroeconomic outcomes. The classification distinguishes technological capability from demand and adoption, and separates capabilities already achieved or available from those described as forthcoming. I refer to these categories as delivered and promised capability.

I estimate how dividend-futures and inflation-swap responses vary with these content shares. Following the existing literature, I use dividend-futures prices as a market-based proxy for economic growth expectations Gormsen and Koijen (2020); Gormsen, Koijen and Martin (2021); Hazell and Hobler (2024); Drechsel and Zhang (2026). Dividend futures reflect expected corporate dividends at different horizons, which are closely related to expectations of aggregate economic activity. I therefore interpret increases in dividend-futures prices as improvements in market-implied growth expectations, while recognizing that these prices also reflect inflation and risk premia. Inflation swaps measure inflation expectations across maturities. Their joint responses allow me to examine how markets revise their expectations of economic growth and inflation in response to technology news. Neither the event return nor the content classification alone identifies a structural technology shock. Instead, their combination allows me to examine whether macroeconomic repricing varies systematically with the information disclosed.

My central finding is that favorable technology news improves economic growth expectations under both disclosure profiles but generates opposite inflation-compensation responses. Without content interactions, conference-day semiconductor returns are positively associated with dividend-futures returns across maturities from one to ten years, whereas the inflation-swap response is statistically significant only at the one-year horizon. The content decomposition reveals a pattern obscured by this pooled result. At the delivered-capability profile, positive revisions to growth expectations accompany higher inflation expectations at medium and long maturities. At the promised-capability profile, growth expectations also improve, but inflation expectations declines.

These differences are economically and statistically meaningful. Evaluating the model at the mean content share of the top decile of delivered capability, a one-percentage-point semiconductor-sector return is associated with a 0.87-basis-point increase in five-year inflation expectations. At the corresponding top-decile promised-capability profile, the response is -0.91 basis points, and the difference is statistically significant. At ten years, the respective responses are +0.81 and -0.61 basis points, again significantly different. In contrast, dividend-futures responses do not differ significantly across the two profiles at any maturity. The contrast therefore lies more

clearly in inflation expectations than in growth expectations, which improve under both profiles.

These findings are consistent with the interplay between anticipated demand expansion and improvements in productive capacity. News about already available capabilities may stimulate complementary investment and technology adoption before the associated productivity gains are fully realized, consistent with higher inflation expectations. News about forthcoming capabilities may also encourage investment, but the accompanying decline in inflation expectations is consistent with expectations of future productive-capacity expansion and lower production costs. Thus, favorable technology news can improve growth prospects under both disclosure profiles while generating opposite inflation-compensation responses, potentially reflecting differences in the relative importance and timing of anticipated demand and supply adjustments. This interpretation relates to the differing inflation responses documented in the technology-news literature Barsky and Sims (2011); Barsky, Basu and Lee (2015); Ahmed and Farah (2022); Cascaldi-Garcia and Vukotić (2022).

The estimated responses also provide a scale for the macroeconomic information conveyed by individual disclosures. Following the illustrative dividend-to-GDP conversion in Hazell and Hobler (2024), drawing on Gormsen and Koijen (2020), I adjust nominal dividend responses for corresponding price-level revisions and apply a conversion factor of 0.67. At the promised-capability profile, a one-percentage-point semiconductor-sector return implies approximate real GDP-level revisions of 0.019, 0.049, and 0.105 percent at one, five, and ten years, respectively. For context, Acemoglu (2025) estimates ten-year GDP gains from AI of 0.93–1.16 percent, while Blumenfeld et al. (2026) estimates gains of 3.6–6.5 percent based on AI-related expectations accumulated through December 2025. A revision of approximately one-tenth of a percent need not be small for an individual disclosure, although these broader estimates concern different economic experiments. The conversion is intended to provide an illustrative magnitude rather than a directly comparable aggregate effect.

The joint specification separates two dimensions of disclosure content. Replacing demand statements with delivered-capability statements, holding promised capability fixed, raises longer-horizon inflation-compensation responses. Shifting capability from delivered to promised, holding total capability fixed, lowers inflation-compensation responses from two years onward. The observed profiles vary along both dimensions, while the conditional horizon interaction shows that the availability of capability matters even when its total share is held constant. This distinction concerns whether the underlying technological capability is already available or still forthcoming, rather than whether an announcement itself is anticipated or unanticipated. Either type of statement may contain information that investors did not previously expect.

To examine the proposed interpretation, I turn to subsequent prices, production,

investment, and labor inputs. Monthly local projections associate both content profiles with higher capital-goods orders and mostly positive equipment-production estimates, consistent with investment in response to technological opportunities. Following promised-capability-heavy disclosures, core prices decline while industrial production, manufacturing hours, and machinery-manufacturing employment increase. The semiconductor-industry producer price responses provide additional evidence: under the promised-capability profile, estimated PPI responses remain negative and trend downward over two years, whereas under the delivered-capability profile, they are initially negative before trending upward. Under the delivered profile, core-price and industrial-production point estimates are also positive, and manufacturing-hours responses are mostly positive. These patterns provide complementary evidence consistent with differences in anticipated demand expansion and productive-capacity improvements, although they do not separately identify the two channels.

I also examine the reach and interpretation of equity repricing. Across sectoral portfolios, responses are strongest in semiconductors, followed by AI-themed and broader technology funds. A presenter-excluded semiconductor portfolio also responds positively, indicating that repricing extends beyond the announcing firm. In a heterogeneous identification exercise, the delivered–promised inflation contrast remains statistically significant from two through ten years among events in which the issuer and its semiconductor peers have abnormal returns of the same sign. The central result therefore does not depend on opposing issuer and peer valuations.

Alternative return windows and outcome measures broadly preserve the baseline pattern. Market-adjusted daily returns and 90- and 120-minute intraday return surprises support the positive delivered-profile inflation-compensation response, while promised-profile estimates remain negative at most maturities. Intraday requirements retain 75 of the 112 baseline event days, so these comparisons change event coverage as well as the return window. Treasury breakeven inflation expectations also rises at the delivered profile and falls at the promised profile. These exercises support the broad direction of the findings while documenting differences in precision across specifications.

This paper contributes to four strands of literature. First,Existing technology-news studies news about future productivity from stock prices and long-run restrictions, estimated dynamic models, or patent-related market reactions (Beaudry and Portier, 2006; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012; Kurmann and Otrok, 2013; Kogan et al., 2017; Cascaldi-Garcia and Vukotić, 2022). These approaches have produced different inflation responses: productivity-based measures often imply disinflation (Barsky and Sims, 2011; Barsky, Basu and Lee, 2015), whereas patent-based news raises inflation (Cascaldi-Garcia and Vukotić, 2022). They do not directly observe the economic content of the underlying disclosures, leaving open whether informa-

tion composition helps explain these differences. This paper is the first to connect statement-level measures of technology-news content to the joint repricing of growth expectations and inflation expectations. I show that content matters: delivered and promised capability are associated with opposite inflation responses, while dividend responses remain positive under both profiles. The result identifies information composition as a possible source of heterogeneity within a common disclosure setting.

Second, text-based research extracts economic information from earnings calls, business news, and other corporate communications (Hassan et al., 2019; Bybee et al., 2024; Hassan et al., 2024). This paper is the first to use semiconductor conference transcripts to measure the economic content of technology disclosures and relate it to macroeconomic expectations. I distinguish capability from demand and delivered from promised capability. Linking these measures to dividend futures and inflation swaps shows that disclosure content matters.

Third, I extend event-based identification to recurring private-sector technology disclosures. High-frequency approaches have been applied to monetary policy, oil supply, fiscal policy, carbon pricing, and bank regulation (Gürkaynak, Sack and Swanson, 2005; Nakamura and Steinsson, 2018; Swanson, 2021; Känzig, 2021; Hazell and Hobler, 2024; Bauer, Känzig and Rudebusch, 2026; Drechsel and Miura, 2025). This paper is the first to construct a technology-news surprise series around recurring conference keynotes and use their transcripts to characterise the information accompanying each surprise.

Fourth, I contribute to the macroeconomics of AI. Theoretical work studies production, task allocation, growth, and the demand effects of anticipated technological change (Aghion, Jones and Jones, 2018; Acemoglu, 2025; Trammell and Korinek, 2024; Aldasoro et al., 2024; Faria-e Castro and Ozkan, 2026). Recent empirical work examines AI model releases through bond markets and estimates broader gains using equity-market exposures (Andrews and Farboodi, 2025; Blumenfeld et al., 2026). I study the computing technologies underlying AI, alongside other semiconductor disclosures, and examine the macroeconomic effects.

The remainder of the paper proceeds as follows. Sections 2 and 3 describe the institutional setting and data. Section 4 constructs the surprise series and content measures. Section 5 presents the main expectations results. Section 6 examines cross-sectoral equity repricing and subsequent economic outcomes. Section 7 relates the findings to existing evidence. Sections 8 and 9 report alternative identification exercises and robustness checks. Section 10 concludes.

2 Semiconductor tech conferences and information disclosure

Technology conferences provide recurring opportunities for semiconductor and computing firms to disclose developments in computing architectures, processor performance, product roadmaps, and commercial adoption. Unlike earnings announcements, which primarily focus on financial performance, conference keynotes offer insights into technological advances and their economic applications. The semiconductor industry is a particularly relevant setting for studying these disclosures because improvements in processors, accelerators, and computing infrastructure can reduce computing costs and expand the range of economically feasible applications, while information about customer adoption and product deployments can reveal shifts in technology demand. Such announcements may therefore convey information relevant to both the announcing firm and broader economic conditions.

The conferences in our setting take two main institutional forms. The first consists of firm-hosted events, such as the Intel Developer Forum, NVIDIA's GPU Technology Conference (GTC), and Qualcomm's Snapdragon Summit. These conferences provide recurring platforms through which firms communicate technological developments and product roadmaps to developers, customers, industry participants, and the public. The second consists of industry-wide conferences organized by independent institutions or industry associations, including the Consumer Electronics Show (CES), Computex, and the Open Compute Project Global Summit. These events bring together multiple technology firms and provide scheduled keynote sessions in which participating companies present technological developments and commercial applications.[1]

Across these institutional formats, the common feature is the keynote presentation: a scheduled public address in which a firm's representatives disclose information within an established conference program. This distinguishes the events from unscheduled product announcements, private investor meetings, and informal interviews. The recurring conference calendar provides an observable set of disclosure opportunities that can be identified independently of the information eventually presented and the subsequent financial-market response.

NVIDIA's GPU Technology Conference provides a concrete illustration of this institutional setting. In April 2021, NVIDIA organized a five-day virtual GTC, featuring approximately 1,300 sessions covering artificial intelligence, accelerated computing, data centers, networking, and related technologies. According to NVIDIA, the conference attracted more than 200,000 registrations from 195 countries, while its opening

[1] For descriptions of the institutional formats, see the official websites of CES (https://www.ces.tech/) and Computex (https://www.computextaipei.com.tw/), as well as the Open Compute Project's conference archive (https://www.opencompute.org/events/).

keynote subsequently accumulated over 14 million views.[2] On April 12, NVIDIA founder and CEO Jensen Huang delivered the opening keynote at 8:30 a.m. Pacific Time (11:30 a.m. Eastern Time), during regular U.S. equity trading hours.

The keynote featured several major technological announcements, including NVIDIA Grace, the company's first Arm-based data-center CPU, and BlueField-3, a new data processing unit designed to accelerate networking, storage, and security operations. NVIDIA projected that Grace-based systems could deliver approximately ten times the performance of its existing CPU-based systems for certain large-scale AI and high-performance computing workloads. The company also announced plans by the Swiss National Supercomputing Centre and Los Alamos National Laboratory to build supercomputers incorporating Grace processors.[3] The example illustrates how a scheduled keynote combines technological and commercial information. The Grace announcement described a new computing capability alongside prospective customer deployments.

Beyond the content of any individual keynote, a defining characteristic of the conferences in our setting is their recurrence. Many are organized annually or on another regular schedule, and their continued operation does not depend on a particular technological breakthrough being announced in a given year. For example, NVIDIA GTC and Qualcomm's Snapdragon Summit provide recurring firm-hosted disclosure opportunities, while CES and Computex operate as established industry-wide events. This institutional structure distinguishes recurring conferences from standalone product launches. A firm may arrange a special announcement in direct response to a newly developed technology or an unexpected commercial opportunity. In such cases, both the timing and the content of the announcement may be determined by the underlying news. Recurring conferences, by contrast, provide disclosure opportunities that are scheduled as part of an established event program.

The dates and programs of upcoming conference editions are commonly announced in advance. Online Appendix A documents this for the whole series sample. Table 7 gives, for each of the fifteen conference series, the organiser's or host's own words establishing that the series recurs on a fixed cycle, with the source for each. Table 8 then reports, series by series, the most recent edition for which a dated advance announcement could be located, how far ahead of the keynote that announcement fell, and the announcing sentence itself. For fourteen of the fifteen series such a document exists, and the documented lead times run from 80 to 357 days with a median of 180. These are examples rather than a census: they establish that advance scheduling is a regular feature of these series.

[2] "NVIDIA Announces Financial Results for First Quarter Fiscal 2022," May 26, 2021.

[3] "NVIDIA Announces CPU for Giant AI and High Performance Computing Workloads," April 12, 2021. See also "NVIDIA Extends Data Center Infrastructure Processing Roadmap with BlueField-3," April 12, 2021.

The distinction between the anticipated occasion and the potentially unexpected information is fundamental to our empirical design. Investors may anticipate that a company will discuss new processors, computing platforms, or product roadmaps at an upcoming keynote. They need not, however, know the precise performance specifications, architectural developments, implementation schedules, or commercial arrangements that will be revealed. Consequently, the conference calendar provides a set of identifiable information windows, while the market's response during these windows may reflect revisions to previously held expectations.

Within these identifiable information windows, the information disclosed may have different macroeconomic implications. Figure 1 illustrates their main technological themes, including data-centre computing, computational performance and efficiency, and artificial intelligence.

Figure 1: Technological themes in conference keynotes



Notes: Word clouds show four principal components from an unsupervised PCA of TF–IDF-weighted unigrams and bigrams in 30,572 sentences classified as news from 182 keynote editions between 2001 and 2026. Terms must appear in at least thirty sentences; larger words indicate higher component loadings. The leading frequency component and non-lexical fillers are omitted. Appendix C provides methodological details.

The central economic distinction is between *capability news*, concerning improvements in technological performance, efficiency, and architecture, and *demand news*, concerning customer adoption, orders, deployments, and investment. The former may expand productive capacity or reduce costs, while the latter may signal anticipated spending and economic activity; other content is classified separately. I also

distinguish realised from future developments, AI-related from other news, and different layers of the technology stack. Because equity-price movements alone cannot identify these dimensions, I combine high-frequency semiconductor-sector returns with LLM-based transcript classifications to measure unexpected market repricing and characterise the economic content of disclosures, respectively.

3 Data

The data combine a conference registry, keynote records, daily and intraday asset prices, and monthly macroeconomic series. Appendix A documents the conference registry and sample audits, Appendix B describes caption processing, and Appendix C evaluates the classifier.

The collection registry covers fifteen recurring tech conference series and 223 editions between 2001 and 2026. Its purpose is to identify scheduled public technology disclosures by listed semiconductor firms, allowing their content to be measured and matched to asset prices. The registry is assembled from organiser archives, programmes, press records, and video sources. Conference series must satisfy five criteria: relevance to semiconductor or computing products, processes, or roadmaps, established from the organiser's description and programme; recurrence on a repeating cycle; prepared keynote addresses delivered to the general conference audience and open to the press or public; a scheduled keynote slot within the organiser's published programme rather than an ad hoc appearance; and selection independent of realised returns, estimated macroeconomic responses, and speech-content measures, including LLM classifications, technology-news scores, and ex post assessments of announcements. Press conferences, show-floor interviews, technical sessions, and executive briefings are excluded. Section A.2 provides evidence of recurrence and advance scheduling. These features support anticipated disclosure opportunities without establishing random assignment of disclosure dates or content.

The speaking firm is identified separately for each keynote from the document's attribution or the host's introduction, rather than from firms mentioned in the speech. A keynote qualifies only if the speaker represents a firm on a 114-name list compiled from 6,977 dated constituent and holdings snapshots of the PHLX Semiconductor Sector Index (SOX), the iShares Semiconductor ETF (SOXX), and the VanEck Semiconductor ETF (SMH). These benchmarks cover US-listed securities, so keynotes by firms outside this list are excluded, including several large semiconductor and storage firms. Twenty-two eligible firms appear in the collected keynote records.

To cross-check sample coverage, I also downloaded all available records classified as *Company Conference Presentations* or *Special Calls* for these 114 firms over the same period from the S&P Capital IQ transcript database and searched them for technology

events. All technology conferences identified through this exercise that satisfied our selection criteria were already included in the conference registry.

I distinguish three units. An *edition* is one occurrence of a conference series in a calendar year. For touring series, only the flagship event is counted, because subsequent legs largely repeat the main keynote and would duplicate overlapping news. A *keynote session* is one firm's speech in one scheduled slot, with its own start time, and is the unit of intraday estimation. A *venue-day observation* combines all qualifying keynotes at the same venue on the same calendar day into one document and is the unit of daily estimation and text aggregation, because daily returns cannot separate sessions within the day. These counts describe different units and are not additive. Observations are dated to the keynote rather than the opening of the conference. For example, NVIDIA's CES and Computex keynotes may occur one or two days after the show opens.

Of the 223 editions, 171 have an identified constituent-firm keynote, 31 have no constituent keynote, and 21 have undetermined programmes. The 171 editions yield 175 keynote-day records. Requiring an independent qualifying speech record leaves 152; restricting to the inflation-swap study window leaves 125. The macroeconomic-release, non-trading-day, and issuer-earnings screens leave 114 keynotes on 112 trading days across twelve series. Outcome availability further determines the sample used in individual regressions.

Qualifying speech records come from four sources: verbatim transcripts or prepared remarks published by the organiser or speaking firm; human-provided captions; automatic captions; and professionally edited transcripts from Capital IQ. Press summaries, slides, and records of other speaking occasions are excluded. Both caption types are obtained from the official YouTube channels of the conference series or speaking firms. All four sources provide full-length records of the speaker's own remarks and are treated consistently in constructing the content measures.

As delivered, stenographic transcripts and human-provided captions carry 6.8 and 6.5 sentence terminators per hundred words; automatic captions carry 0.2, because machine transcription emits almost no punctuation. This matters because the content measures are applied to statement-level units, so text with no sentence boundaries has nothing for a unit to attach to — the distinctions being measured are clause-level, capability against demand and forward-looking against realised. The gap is therefore closed before any classification runs, by XLM-RoBERTa-based punctuation restoration via token classification, documented and validated in Appendix B: after it, the three sources carry 6.8, 6.5 and 6.7 terminators per hundred words, and the words themselves are provably unaltered.

The surprise is measured on the iShares Semiconductor ETF (SOXX), whose adjusted closing prices I take from 2 January 2002 to 20 October 2026. SPY serves as the

market benchmark wherever a market-adjusted variant is reported, and daily adjusted closes for 21 individual semiconductor names — the index's large constituents across the sample, listed in Section 6.1 — support the cross-sectional exercises of Sections 6.1 and 8.2. The intraday work of Section 8.1 uses minute-frequency last-trade prices for SOXX and SPY from WRDS TAQ, 09:30 to 16:00 New York time, with the accessible libraries beginning on 10 September 2003; SOXX traded as IGW before 15 October 2010, and the predecessor series is carried under its own ticker.

Expectations are read off two markets. Zero-coupon USD inflation swap rates at maturities of one to ten years come from Bloomberg (`USSWIT1`–`USSWIT10`) and are quoted daily from 21 July 2004. Annual dividend futures on the S&P 500 come from Bloomberg (`ASDZ17`–`ASDZ29`) and are quoted daily from 13 November 2015; because these are dated rather than constant-maturity claims — `ASDZ26` is a claim on calendar-2026 dividends — the reported regressions use horizon-indexed dividend outcomes. I also use FRED's constant-maturity series at five and ten years: nominal yields (`DGS5`, `DGS10`), real yields from inflation-indexed securities (`DFII5`, `DFII10`), and the implied break-even rates (`T5YIE`, `T10YIE`), which are published on a common daily grid.

The realised-outcome projections of Section 6 use monthly FRED series: core consumer prices (`CPILFESL`), industrial production (`INDPRO`), semiconductor and finished-goods producer prices (`PCU334413334413`, `PPIFGS`), core capital-goods orders (`NEWORDER`), business-equipment production (`IPBUSEQ`), machinery-manufacturing employment (`CES3133300001`), the quits rate (`JTSQUR`), and average weekly hours in manufacturing (`AWHMAN`).

4 Identifying tech-news surprises and its information content

4.1 Semiconductor index returns

For each keynote e I take the semiconductor index's close-to-close return over the trading day on which the keynote's information reaches the market:

$$z_e \;=\; 100\big[\log P_{I,\bar{t}_e} - \log P_{I,\underline{t}_e}\big], \tag{1}$$

where $\bar{t}_e$ is the close of the day the keynote was given, $\underline{t}_e$ the close of the trading day before it, and P_I the adjusted closing price of SOXX, the iShares semiconductor index. The surprise is measured in percentage points and is never standardised, so a coefficient estimated on it is the response to a one-per-cent move in the semiconductor index.

The daily return uses the keynote's recorded date. The baseline uses the daily return on the keynote's recorded date because keynotes vary in duration from 30 to 180 minutes and conferences typically span several days, with related information

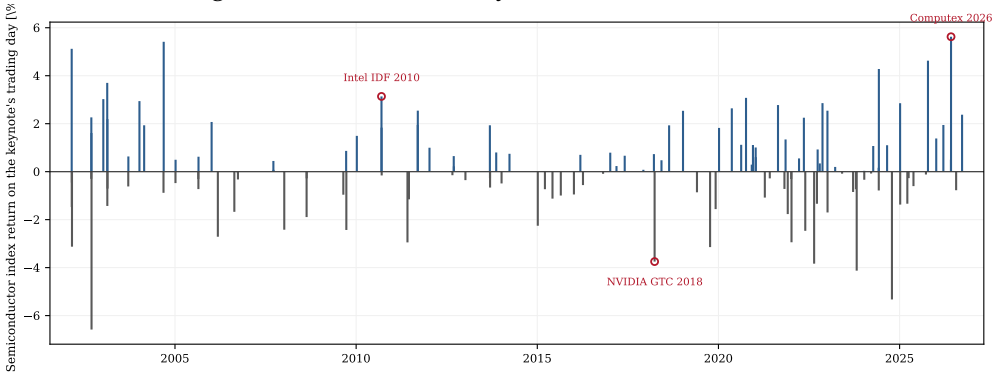
potentially released outside the keynote itself. A common narrow intraday window may therefore capture different fractions of the relevant disclosures across events. Section 8.1 uses keynote start times to construct narrower windows for presentations held during U.S. trading hours.

I exclude days with FOMC statements, CPI releases, and earnings calls by the firm giving the keynote. When keynotes share a pair of closes, they contribute one return observation.

Figure 2 plots the surprise series: 112 observations drawn from 114 keynotes on 112 trading days across twelve conference series, from 7 September 2004 to 22 September 2026, with a standard deviation of 1.88 percentage points and a mean of +0.14. The largest move in absolute value is 5.6 per cent. The series is plotted over the study window the inflation swaps define — it opens on 21 July 2004, the first day the swap curve is quoted at all ten maturities.

To assess whether event-day returns reflect continuation from earlier trading sessions, I correlate them with returns on the most recent keynote-free trading day. Across all 112 event days, the correlation is -0.08 ($t = -0.88$), indicating no significant positive association. If the immediately preceding trading day also contains a keynote, I look further back to the nearest trading day without one.

Figure 2: Conference-day semiconductor returns

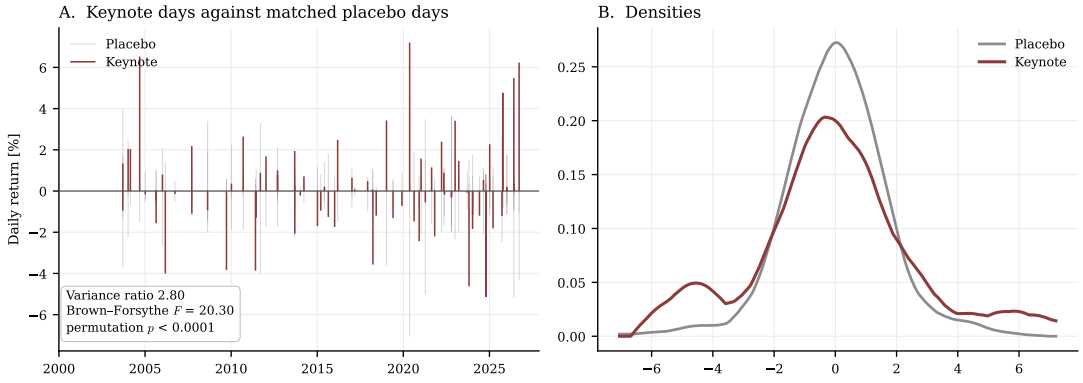


Notes: This figure shows conference-day semiconductor returns: for each pair of closes bracketing a keynote, SOXX's close-to-close log return over the day the keynote's information reaches the market, in per cent, as defined in equation (1). Circles mark three of the largest absolute surprises in the sample, discussed below.

To assess whether keynote days contain unusually large price revisions, I compare the variance of SOXX returns on keynote days with that on matched non-keynote days, following Känzig (2021). Controls are matched by weekday and week of the month over the preceding three months. The diagnostic starts from 132 keynotes, including eighteen before the inflation-swap window, and requires no swap quote. I assess significance using the Brown–Forsythe test and a permutation test with 20,000 draws, each replacing every event surprises with one of its matched control returns.

Figure 3 shows that keynote-day returns are more variable than their matched controls. Event-day variance is 6.78, compared with 2.42 on control days, a ratio of 2.80. The Brown–Forsythe statistic is 20.30, and the permutation probability is below 0.0001. The evidence supports elevated volatility relative to close calendar matches.

Figure 3: Keynote days versus matched control days



Notes: SOXX daily close-to-close log returns, in per cent, on keynote observations and non-keynote controls matched by weekday and week of the month. Left: observations over time. Right: Epanechnikov kernel densities with a common bandwidth equal to twice Silverman's event-sample rule. Returns are not standardized.

Narrative evidence. Three of the largest absolute surprises in the sample illustrate disclosures about technological capability, risks to commercial applications, and firms' commercial prospects.

Intel's IDF, 13 September 2010. Intel was then the largest index constituent. Paul Otellini's keynote introduced Sandy Bridge as the Second Generation Intel Core and demonstrated its integrated graphics against a similarly configured laptop with a discrete graphics card. The keynote also reported progress towards 22-nanometre production in the second half of 2011 and demonstrated a two-socket server platform running 32 threads, scheduled for production in 2011. These announcements concerned computing architecture, manufacturing technology, and performance; Sandy Bridge itself used a 32-nanometre process, while the 22-nanometre developments concerned subsequent products.[4] The semiconductor index gained 3.14 per cent, compared with 1.09 per cent for the S&P 500, while Intel's own return was 3.12 per cent. Such advances could affect demand for complementary semiconductor products by encouraging PC and server upgrades, while also changing competitive conditions. These are possible transmission channels rather than established explanations for the observed sector increase.

NVIDIA's GTC, 27 March 2018. Nine days after an Uber self-driving vehicle

[4] "Intel Details 2011 Processor Features, Offers Stunning Visuals Built-in," September 13, 2010.

killed a pedestrian in Arizona, Jensen Huang addressed the accident and discussed autonomous-driving safety during the keynote. Uber used NVIDIA GPU hardware. During the presentation, news emerged that NVIDIA had suspended its own public-road autonomous-vehicle testing; Huang subsequently explained the decision at a conference press briefing.[5] NVIDIA's own return was −8.0 per cent, while the S&P 500 fell 1.7 per cent. The episode illustrates how information about risks to a technology's commercial use can accompany negative repricing.

Computex, 2 June 2026. The semiconductor index outperformed the S&P 500 by approximately 5.5 percentage points. Marvell CEO Matt Murphy delivered a scheduled keynote on connectivity and the scaling of AI infrastructure. Jensen Huang joined him on stage and called Marvell "the next trillion-dollar company." Contemporary reporting linked the remark to Marvell's sharp share-price increase; the stock closed up 32.5 per cent.[6]

In addition to the three episodes above, Intel's Computex keynote on 28 May 2019 serves as a good illustration of unmet disclosure expectations. Intel presented its tenth-generation Ice Lake processors and updates to its PC product portfolio, but some anticipated developments remained unaddressed. Intel presented its tenth-generation Ice Lake processors and updates to its PC product portfolio, but some anticipated developments remained unaddressed. A contemporary Computex review in Tom's Hardware described Intel's showing as "a bit underwhelming," particularly noting the absence of expected information on its discrete graphics plans.[7] Intel's own return was −2.27 per cent; the semiconductor index fell 0.86 per cent and SPY fell 0.92 per cent. The case illustrates how substantive product announcements can coexist with disappointment about information that was not disclosed, although the commentary does not establish the cause of the price decline.

These cases distinguish the content of technological disclosures from their unexpected valuation effects and illustrate conference-day sector returns as market repricing around public disclosure opportunities.

4.2 The information content of technology news

The sector return measures valuation changes, while the transcript records disclosed information. I classify the text to distinguish technological capability from commercial demand and to measure whether capability is described as delivered or

[5] "NVIDIA Pauses Self-Driving Vehicle Testing; Confirms DRIVE Wasn't In Last Week's Uber Collision," March 29, 2018. See also "Jensen Huang on the Uber Tragedy and Why Nvidia Suspended Testing."

[6] For the keynote schedule, see "Marvell CEO Keynote," COMPUTEX 2026. For Huang's appearance and remarks, see "Jensen Huang joins Marvell CEO on Computex stage, hails it as 'the next trillion-dollar company'," June 2, 2026.

[7] "Community Questions: What Are Your Thoughts on Computex 2019?" The article is trade-press commentary on the conference.

promised. The classifier does not observe prices or macroeconomic outcomes.

Sentences enter the informational category when they contain a qualifying factual statement rather than framing or evaluation. The label does not measure novelty relative to investors' prior information. Informational statements receive an economic role: capability, demand, or context. Capability covers performance, efficiency, architecture, and product roadmaps. Demand covers customers, orders, shipments, market size, and commercial adoption. Context covers statements assigned to neither role. Horizon labels distinguish claims about existing or achieved developments from claims about future developments. Table 1 provides examples. Appendix C reports agreement with author annotations and the model-selection procedure. These comparisons use sentences consulted during prompt development; they are in-sample agreement measures. The role and horizon fields enter the regressions. Additional technology-stack do not enter the main analysis.

Table 1: Classification of technology-news statements

Statement	Role	Horizon
"This is the Xeon 5100 series platform . . . delivering up to 3X performance over single-core platforms, 3.5 times the performance per watt." *Intel IDF 2006*	Capability	Realised
"And even though we've introduced 45-nanometer technology with High-K/Metal Gate, we've been able to reduce throughput time by 50 percent." *Intel IDF 2009*	Capability	Realised
"There'll be a large number of SKUs that are launched on November 12th for servers and for high-end desktops." *Intel IDF 2007*	Capability	Forward
"And I'm thrilled to announce today that SEGA and Intel are working together to put SEGA's premium IP on to AppUp." *Intel IDF 2010*	Demand	Realised
"Seven OEMs are going to be shipping our notebooks with our Echo Peak technology in the first quarter of this year." *Intel IDF 2008*	Demand	Forward

Notes: Each row is a sentence as it appears in the official transcript, with the classification the model assigned to it.

Let N_e^{Side} count statements classified as capability or demand, excluding context. Define

$$C_e = \frac{N_e^{\text{Capability}}}{N_e^{\text{Side}}}, \qquad F_e = \frac{N_e^{\text{Capability,Future}}}{N_e^{\text{Side}}}, \qquad P_e = C_e - F_e. \tag{2}$$

C_e measures total capability content, F_e promised capability, and P_e the residual capability share. Interpreting P_e as delivered capability requires horizon labels to exhaust the capability category.

All three measures use the same denominator. The demand share is $1 - C_e$. The share of all forward-looking statements, which includes future demand, is a different measure and does not enter F_e in the main specification. Shares describe the composition of statements, not the composition of the unexpected information priced by investors.

The baseline content regressions use qualifying transcripts associated with the

screened event sample, subject to each outcome's availability. For days containing multiple keynotes I average the shares. Question-and-answer sessions appended to prepared keynotes are excluded from the content calculation. Table 2 reports the available transcript-level descriptive statistics. Capability statements account for a mean share of 0.734 among capability and demand statements. Promised capability has a mean share of 0.095.

Table 2: The information content of keynote news

	Mean	S.D.	p10	Median	p90	Max
Shares of a keynote's news sentences						
Capability, C_e	0.734	0.130	0.531	0.751	0.882	0.992
Demand	0.266	0.130	0.118	0.249	0.469	0.616
Context	0.173	0.091	0.075	0.151	0.284	0.611
Forward-looking, all roles	0.135	0.070	0.058	0.126	0.222	0.365
Promised capability, F_e	0.095	0.057	0.030	0.085	0.173	0.267
Underlying counts						
News sentences	169	104	63	150	288	688
Capability + demand	143	92	50	116	254	580
All sentences	610	330	250	540	976	2,510
News rate	0.275	0.082	0.187	0.273	0.372	0.517

Notes: Capability and demand shares sum to one. The all-roles forward-looking share uses capability plus demand as its denominator and includes both roles in its numerator; promised capability restricts the numerator to future capability. Context is reported as a share of all news statements.

5 The effect of tech news on macroeconomic expectations

5.1 The surprise on its own

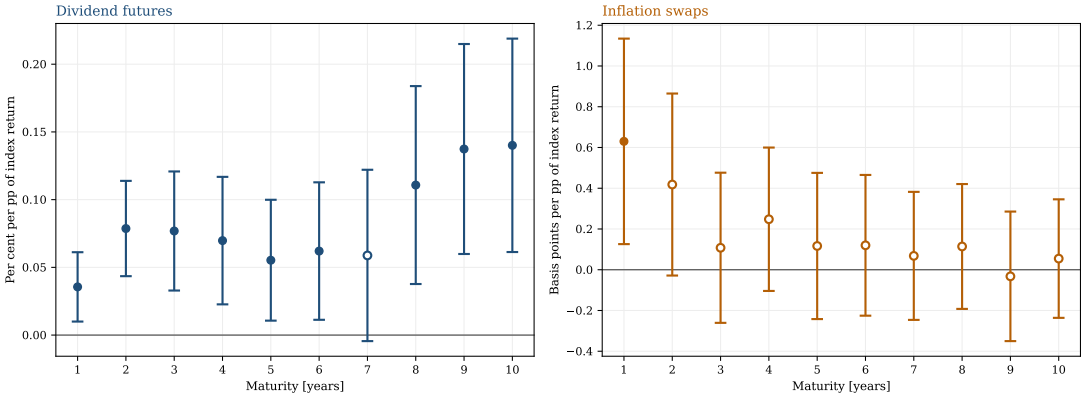
I first estimate the association between conference-day sector returns and changes in dividend-futures prices and inflation-swap rates. Let z_t equal the sector return on an event day and zero otherwise. For maturities $m = 1, \ldots, 10$,

$$\Delta y_{m,t} = \alpha_m + \beta_m z_t + \Gamma_m' X_t + \varepsilon_{m,t}, \tag{3}$$

where X_t contains weekday indicators and one lag of the outcome. I use heteroskedasticity-robust standard errors and each contract's available daily history. Inflation swaps begin in July 2004 and dividend futures in November 2015. The swap regressions contain 101–109 event days and the dividend regressions 64–68, depending on maturity.

β_m is the change in the outcome associated with a one-percentage-point sector return on an event day. It does not compare that loading with the corresponding loading on non-event days. Consequently, it can reflect ordinary comovement between the semiconductor index and the outcome as well as conference information.

Figure 4: The response to a keynote-day index return, with no content split.



Notes: Each panel plots β_m from equation (3) against the contract's maturity, with 90 per cent heteroskedasticity-robust confidence intervals; filled markers are the maturities at which the interval excludes zero. Left: dividend futures, in per cent per percentage point of index return, on 2,414 to 2,724 trading days of which 64 to 68 carry a keynote. Right: inflation swaps, in basis points per percentage point, on 5,439 to 5,660 trading days of which 101 to 109 carry a keynote. Of the 112 keynote trading days within the inflation-swap study window, 101–109 enter each maturity-specific regression, depending on the availability of current and lagged outcome data. Each contract is estimated over its own quoted history: the dividend futures from November 2015, the swaps from July 2004. The regressor is the series of Figure 2 placed on the full trading calendar with a value of zero on days without a keynote.

Table 3: The keynote-day return on its own

	Maturity (years)					
	1	2	3	5	7	10
Panel A. Dividend futures (per cent per percentage point)						
Keynote-day return z_t (β_m)	0.036**	0.079***	0.077***	0.055**	0.059	0.140***
	(0.016)	(0.021)	(0.027)	(0.027)	(0.038)	(0.048)
Observations	2,450	2,702	2,724	2,718	2,448	2,414
Keynote days	64	68	68	68	64	64
Panel B. Inflation swaps (basis points per percentage point)						
Keynote-day return z_t (β_m)	0.630**	0.418	0.108	0.117	0.068	0.055
	(0.306)	(0.272)	(0.224)	(0.218)	(0.191)	(0.177)
Observations	5,523	5,439	5,589	5,557	5,603	5,657
Keynote days	104	101	105	107	106	106

Notes: Estimates of β_m from equation (3), corresponding to Figure 4. Each maturity uses its available daily history, with weekday indicators and one lag of the outcome. The regressor equals zero on non-event days. Heteroskedasticity-robust standard errors are in parentheses; *, **, and *** denote significance at the 10, 5, and 1 per cent levels. Dividend-futures data begin in November 2015 and inflation-swap data in July 2004. Sample sizes vary across maturities because current and lagged outcome data must both be available. Six of ten maturities are reported.

Figure 4 and Table 3 show contrasting patterns across the two markets. Dividend-futures prices respond positively to conference-day semiconductor returns throughout the term structure: the estimated response rises from 0.036 per cent at one year to 0.140 per cent at ten years per percentage point of sector return. inflation expectations responds positively at the front end, with a one-year coefficient of 0.63 basis points,

significant at the five per cent level, but the estimates are smaller and not statistically significant at longer maturities. This attenuation does not simply reflect greater uncertainty: the standard error falls from 0.306 basis points at one year to 0.177 at ten years. The ten-year 90 per cent confidence interval, approximately $[-0.24, 0.35]$, excludes a response as large as the one-year point estimate, although a formal cross-maturity comparison requires a direct test. The pooled return therefore captures a broad positive dividend response without a comparable response along the inflation-swap curve. I next examine whether disclosure content helps explain the contrasting responses.

5.2 Adding what the keynote was about

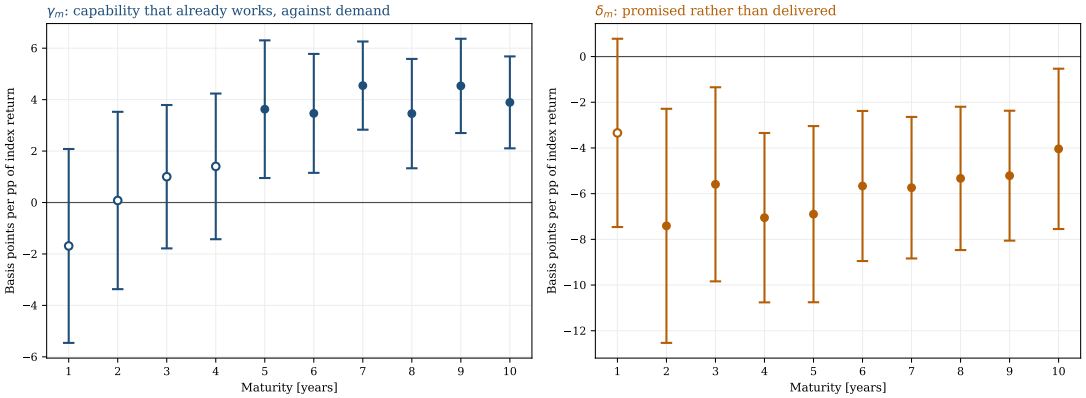
Section 4.2 assigns each classified keynote two shares over one denominator, the capability and demand statements it makes. The capability share C_e is the fraction of that denominator that is capability news rather than demand news. The forward-looking share F_e is the fraction that is capability the firm says it *will* have rather than capability it already demonstrates, so that $C_e = P_e + F_e$ splits total capability content into a delivered part and a promised part. Writing $Z_d^C = z_d C_d$ and $Z_d^F = z_d F_d$ for the two interacted regressors and C_d, F_d for the day's mean shares, equation (3) becomes

$$\Delta y_{m,t} = \alpha_m + \beta_m z_t + \gamma_m Z_t^C + \delta_m Z_t^F + \theta_m C_t + \phi_m F_t + \Gamma_m' X_t + \varepsilon_{m,t}, \qquad (4)$$

estimated on the same calendar, with the same controls and the same standard errors, so that the only difference from Section 5.1 is the four content terms. Setting $\gamma_m = \delta_m = \theta_m = \phi_m = 0$ returns equation (3) exactly. The content-level terms allow average outcome changes to vary with disclosure composition, separately from how composition changes the loading on the event surprises. All event-specific regressors equal zero on non-keynote days.

Since $C_d = P_d + F_d$, the return terms can be written as $\beta_m z_d + \gamma_m z_d P_d + (\gamma_m + \delta_m) z_d F_d$. Thus γ_m measures the change in the return response when delivered capability replaces demand statements, holding promised capability fixed. δ_m measures the shift from delivered to promised capability at a fixed total capability share. The promised-capability gradient relative to demand is $\gamma_m + \delta_m$. Estimating both terms jointly allows us to distinguish these margins and calculate the uncertainty of their contrasts.

Figure 5: Capability against demand, and delivered against promised.



Notes: γ_m (left) and δ_m (right) from equation (4), estimated jointly, one regression per maturity on the full daily trading calendar from July 2004. γ_m is the change in the response to one percentage point of index return when the capability share rises by one, holding the forward-looking share fixed, so it raises the delivered-capability share and lowers the demand share. $\delta_m = \gamma_m^F - \gamma_m^P$ holds total capability content fixed and shifts it from delivered to promised. Bars are 90 per cent heteroskedasticity-robust confidence intervals; filled markers are the maturities at which the interval excludes zero.

Delivered capability and the horizon shift have different implications for inflation expectations. Figure 5 shows that replacing demand with delivered capability raises the response at longer maturities. The gradient is 3.63 basis points at five years, 4.55 at seven, and 3.89 at ten. The response is concentrated beyond four years, where the confidence intervals exclude zero. Only the inflation swaps are drawn: on dividend futures γ_m is indistinguishable from zero at all ten maturities and δ_m at nine of the ten, and Panel A of Table 4 reports them.

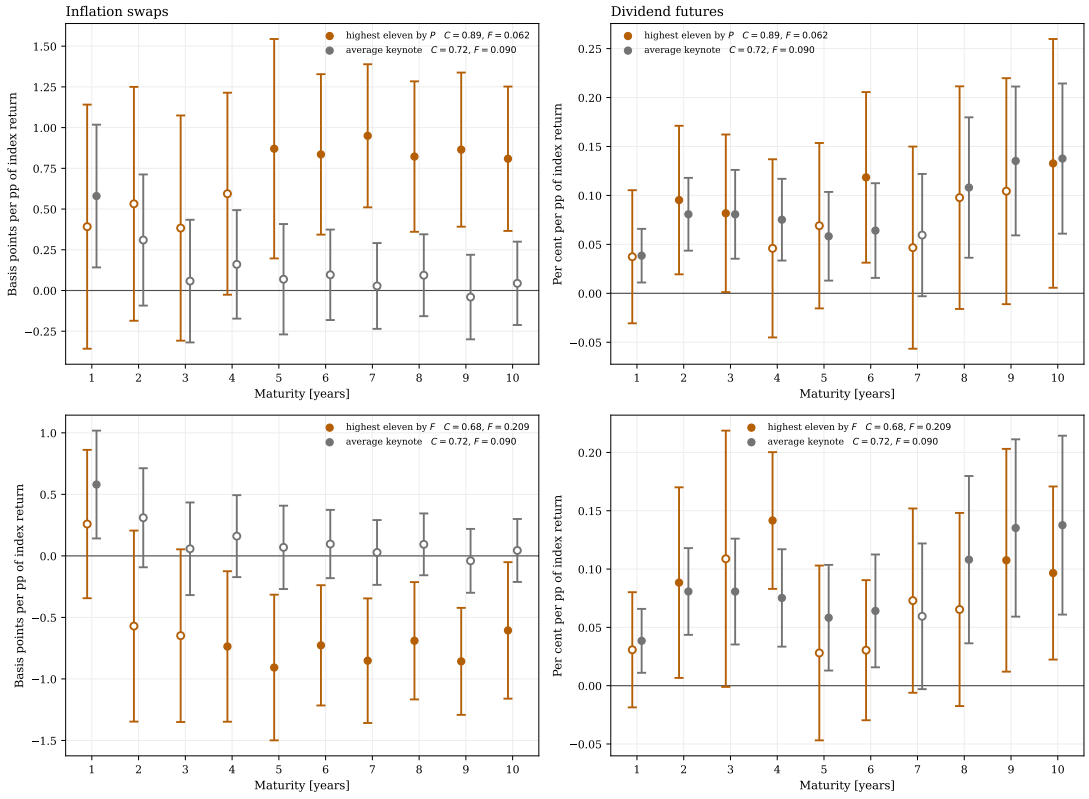
Shifting capability from delivered to promised moves inflation expectations in the opposite direction. The horizon contrast is negative throughout the curve and statistically significant from two years onward. At five years it is −6.90 basis points. Because these effects describe a unit change in a content share, I next translate them into responses at content profiles observed in the sample.

I then evaluate $R_m(C, F) = \beta_m + \gamma_m C + \delta_m F$ at two profiles drawn from the observed daily content distribution. The mean of the eleven days in the top decile of delivered capability is $(C, F) = (0.891, 0.062)$. The corresponding profile for promised capability is $(0.680, 0.209)$. Standard errors use the full covariance matrix of $(\beta_m, \gamma_m, \delta_m)$. I evaluate responses at observed content profiles rather than at the polar cases $P = 1$ or $F = 1$. These cases would describe keynotes composed entirely of delivered or promised capability statements, respectively, with no demand content. Neither occurs in the sample, and evaluating the fitted response there would require substantial extrapolation beyond the observed content distribution, particularly for promised capability. The profiles below instead provide empirically grounded comparisons.

These profiles differ in total capability content as well as its horizon. Their response difference is $R_m(0.891, 0.062) - R_m(0.680, 0.209) = 0.211\gamma_m - 0.147\delta_m$ up to rounding

of the evaluation points. It tests a difference between these two profiles; δ_m separately tests the horizon margin holding total capability fixed. The coefficient β_m cancels from the profile difference. Its level refers to $C = F = 0$, outside the observed capability range, so I do not interpret it on its own.

Figure 6: The net response at the top tenth of each cut.



Notes: Fitted response $\beta_m + \gamma_m C + \delta_m F$ with 90 per cent confidence intervals computed from the full covariance matrix. Top row: mean content of the eleven days in the top decile of delivered capability, $(C, F) = (0.891, 0.062)$. Bottom row: top decile of promised capability, $(0.680, 0.209)$. Each is compared with the sample-mean profile. The profiles select evaluation points; the regressions use all available event days for each contract. These comparisons change more than one content margin and are not separate estimates of γ_m or δ_m. Left: inflation swaps; right: dividend futures. Appendix F adds the bottom deciles.

Figure 6 translates the content effects into responses at observed content profiles. The left-hand panels show a clear separation in inflation compensation. At the delivered-capability profile, the response is positive and most pronounced at medium and long maturities, whereas at the promised-capability profile it turns negative beyond the shortest maturity. Relative to the sample-mean profile shown in grey, the two profiles move longer-horizon responses in opposite directions. At five years, a one-percentage-point sector return is associated with an increase of 0.87 basis points at the delivered-capability profile and a decline of 0.91 basis points at the promised-capability profile. Their difference is 1.78 basis points and significant at the one per

cent level (Table 4).

Table 4: The surprise, the capability share and the forward share

	Maturity (years)					
	1	2	3	5	7	10
Panel A. Dividend futures (per cent per percentage point)						
Surprise z_d	0.058	-0.004	0.022	0.064	0.096	0.240
	(0.133)	(0.149)	(0.192)	(0.191)	(0.240)	(0.314)
$\times Z_d^C$ (realised, against demand: γ_m^P)	-0.019	0.104	0.049	0.023	-0.061	-0.094
	(0.192)	(0.211)	(0.259)	(0.264)	(0.311)	(0.403)
$\times Z_d^F$ (promised less realised: δ_m)	-0.072	0.103	0.256	-0.247	0.091	-0.382
	(0.205)	(0.326)	(0.417)	(0.282)	(0.361)	(0.383)
Capability share C_d	-0.163*	-0.109	-0.165	-0.065	-0.051	0.082
	(0.086)	(0.112)	(0.131)	(0.184)	(0.166)	(0.188)
Forward share F_d	1.120*	0.573	0.876	-0.400	0.570	-0.457
	(0.602)	(0.897)	(1.101)	(1.671)	(0.917)	(1.315)
Derived: the net response $\beta_m + \gamma_m^P C + \delta_m F$ at each evaluation point, and their difference.						
At a realised-heavy keynote ($C = 0.89$, $F = 0.06$)	0.037	0.095**	0.082*	0.069	0.047	0.133*
	(0.041)	(0.046)	(0.049)	(0.051)	(0.063)	(0.077)
After a forward-heavy keynote ($C = 0.68$, $F = 0.21$)	0.031	0.088*	0.109	0.028	0.073	0.097**
	(0.030)	(0.050)	(0.067)	(0.046)	(0.048)	(0.045)
Difference (realised less forward)	0.007	0.007	-0.027	0.041	-0.026	0.036
	(0.055)	(0.069)	(0.082)	(0.080)	(0.080)	(0.096)
Observations	2,450	2,702	2,724	2,718	2,448	2,414
Keynote days	64	68	68	68	64	64
Panel B. Inflation swaps (basis points per percentage point)						
Surprise z_d	2.103	0.921	-0.164	-1.934*	-2.744***	-2.409***
	(1.723)	(1.554)	(1.165)	(1.095)	(0.696)	(0.728)
$\times Z_d^C$ (realised, against demand: γ_m^P)	-1.688	0.079	1.004	3.628**	4.546***	3.893***
	(2.291)	(2.096)	(1.694)	(1.627)	(1.042)	(1.086)
$\times Z_d^F$ (promised less realised: δ_m)	-3.341	-7.409**	-5.591**	-6.898***	-5.739***	-4.040*
	(2.505)	(3.114)	(2.582)	(2.344)	(1.883)	(2.133)
Capability share C_d	-0.516	-1.982*	-1.905**	-1.487*	-0.970	-0.893
	(1.275)	(1.080)	(0.945)	(0.847)	(0.716)	(0.692)
Forward share F_d	1.713	13.259**	14.357***	9.172**	5.667	5.936
	(7.229)	(6.100)	(5.334)	(4.649)	(4.460)	(4.116)
Derived: the net response $\beta_m + \gamma_m^P C + \delta_m F$ at each evaluation point, and their difference.						
At a realised-heavy keynote ($C = 0.89$, $F = 0.06$)	0.392	0.532	0.383	0.870**	0.950***	0.809***
	(0.456)	(0.436)	(0.420)	(0.409)	(0.267)	(0.269)
After a forward-heavy keynote ($C = 0.68$, $F = 0.21$)	0.259	-0.571	-0.648	-0.907**	-0.852***	-0.606*
	(0.366)	(0.472)	(0.427)	(0.360)	(0.308)	(0.337)
Difference (realised less forward)	0.133	1.103*	1.032*	1.778***	1.802***	1.415***
	(0.640)	(0.627)	(0.528)	(0.522)	(0.410)	(0.442)
Observations	5,523	5,439	5,589	5,557	5,603	5,657
Keynote days	104	101	105	107	106	106

Notes: Each maturity is estimated on its available daily history, with weekday indicators and one lag of the outcome. Event-built regressors equal zero on non-event days. Heteroskedasticity-robust standard errors appear in parentheses; *, **, and *** denote ten, five, and one per cent significance. $Z_d^C = z_d C_d$ and $Z_d^F = z_d F_d$. The coefficient on Z_d^C replaces demand with delivered capability holding promised capability fixed; the coefficient on Z_d^F shifts delivered to promised capability holding total capability fixed. The promised gradient relative to demand is the sum of the two. Derived responses and their differences use the full covariance of $(\beta_m, \gamma_m, \delta_m)$. Evaluation points are rounded in the row labels. The profile difference changes both content margins; it is distinct from the conditional horizon test. Six of ten maturities are reported.

The contrast extends across the term structure: the profile difference is signifi-

cant at ten per cent from two through ten years. The content interactions therefore uncover opposing responses that are obscured when the return enters alone. The maturity-specific estimates describe one correlated curve; their confidence intervals are pointwise.

The right-hand panels show a different pattern. Dividend-futures responses are positive at both profiles, and neither profile differs systematically from the sample-mean response across maturities. The directly estimated difference between the two profiles is not statistically significant at any maturity. Interpreting dividend-futures prices as indicators of growth expectations, both profiles are associated with a positive growth outlook; the distinction between them lies more clearly in inflation expectations.

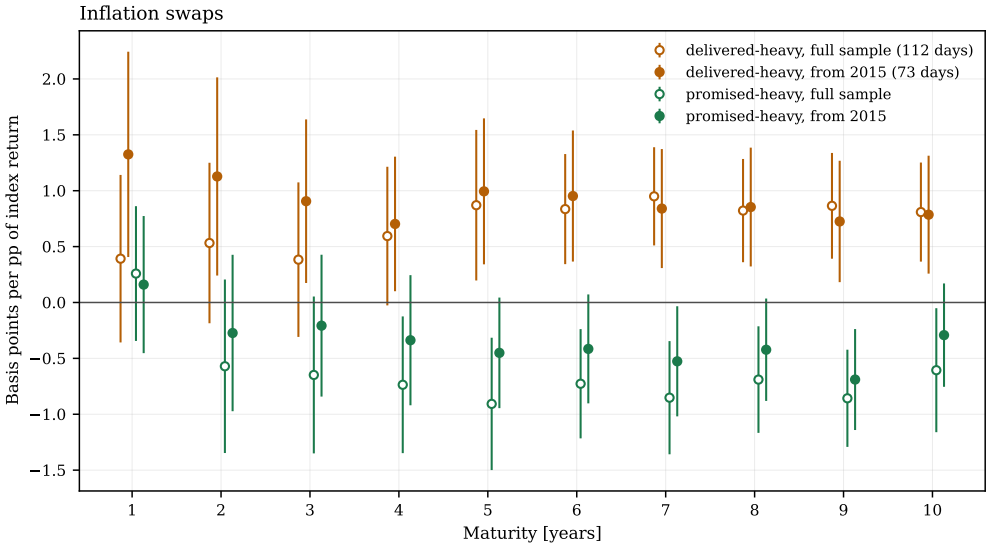
The joint responses suggest different economic interpretations. At the delivered-capability profile, positive growth expectation responses accompany higher inflation expectations, a pattern consistent with news about stronger anticipated demand. At the promised-capability profile, positive growth expectation responses accompany lower inflation expectations, consistent with news about future supply expansion or lower production costs. Section 7 discusses these interpretations further in relation to existing evidence.

5.3 The post-2015 subsample

To assess whether the main findings extend to the AI era, we re-estimate equation (4) using observations from 2015 onwards. This period captures the growing prominence of AI, accelerated computing, and data-centre infrastructure in semiconductor disclosures. The exercise examines whether the content-dependent inflation responses remain present in this environment, providing evidence relevant to the macroeconomic implications of AI-related technological change. The date restriction defines a broad technological period, rather than identifying AI-specific announcements.

We restrict the entire estimation panel, including both event and non-event days, to dates on or after 1 January 2015. The specification, controls, surprise measure, and content classifications remain unchanged. The later window contains 73 of the 112 baseline event days, before outcome-specific availability restrictions. We retain the full-sample delivered- and promised-capability evaluation points, $(C, F) = (0.891, 0.062)$ and $(0.680, 0.209)$, so that comparisons reflect changes in the estimated coefficients rather than changes in the content composition at which responses are evaluated. The corresponding subsample-specific top-decile averages are similar, at approximately $(0.88, 0.04)$ and $(0.69, 0.23)$, and yield the same qualitative interpretation.

Figure 7: Inflation-compensation responses: post-2015 and full samples



Notes: Fitted responses $\beta_m + \gamma_m C + \delta_m F$ at the baseline delivered- and promised-capability evaluation points, in basis points per percentage point of sector return. Hollow markers denote the full sample and filled markers the post-2015 sample, covering 112 and 73 event days before outcome matching. Bars denote 90 per cent confidence intervals based on the full robust coefficient covariance matrix.

Figure 7 shows that the later sample preserves the baseline direction of the inflation responses. Beyond the one-year maturity, the fitted response remains positive for delivered-capability-heavy keynotes and negative for promised-capability-heavy keynotes. Despite the smaller event sample, the difference between the two responses remains statistically significant from two through ten years. The content-dependent inflation pattern is therefore also present in the later period and is not solely driven by the earlier conferences.

The differences concern the shape and magnitude of the response across maturities. At the short end, the delivered-capability gradient γ_m becomes more positive: its post-2015 estimates are 3.34 at one year and 3.79 at two years. The delivered-minus-promised response difference rises from 0.13 to 1.17 basis points at one year and from 1.10 to 1.40 basis points at two years. At longer maturities, some content interaction coefficients are smaller in absolute value. The horizon contrast δ_m changes from -6.90 to -4.46 at five years, while γ_m changes from 4.55 to 3.17 at seven years. The central sign pattern thus persists, although the estimated term structure is flatter in the later sample. These comparisons describe point-estimate differences rather than formal tests of parameter stability.

Dividend-futures data begin in November 2015, so restricting the panel to dates from January 2015 does not shorten their available history. Figure 21 therefore provides a reference comparison rather than an independent subsample robustness test. The dividend results retain the baseline interpretation: fitted responses are positive at both

content evaluation points. Appendix D reports the underlying coefficients by maturity.

6 Financial repricing and subsequent economic outcomes

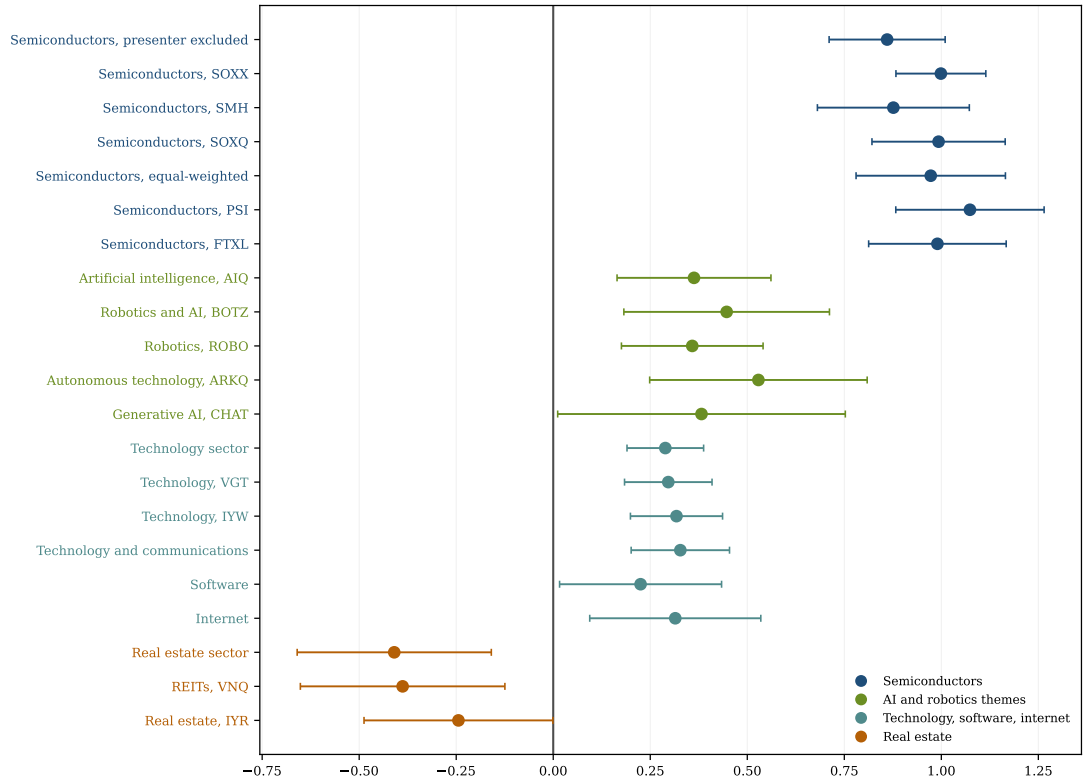
The preceding section documents how disclosure content shapes the repricing of dividend futures and inflation swaps and thus economic growth expectations and inflation expectations. I now extend the analysis in two directions. First, I examine how repricing is distributed across sectoral equity portfolios. Second, monthly local projections trace subsequent prices, production, investment, and labour-market outcomes over a two-year horizon.

6.1 Cross-sectoral equity repricing

This section examines the distribution of equity repricing across sectors. It uses a market-adjusted event surprises, and controls for the contemporaneous market return to reduce common market exposure. The baseline studies the macroeconomic implications of semiconductor technology news. Because such news can affect both semiconductor stocks and the broader market, subtracting the market return would remove part of the response of interest. I therefore use raw semiconductor returns in the baseline. This section asks a different question: which sectors reprice beyond their usual exposure to market movements? Therefore, I use abnormal semiconductor returns and control for the contemporaneous market return in each portfolio regression to address that question.

I collect daily price series for ETFs covering semiconductors, AI and robotics, software, internet, and real estate, spanning technology producers, downstream applications, and related sectors such as property markets connected to data-centre investment. I also constructed SEMIX, which is not a purchasable fund. it is the twenty-one listed semiconductor names — the universe this paper uses wherever a semiconductor peer group is needed — equal-weighted, with the presenting firm removed event by event. It is carried because every semiconductor index contains the presenter.

Figure 8: Cross-sectional equity repricing on keynote days



Notes: Each marker is the net response $b_p + \gamma_p C$ at $C = 0.891$ from a separate estimation of equation (5), in per cent per percentage point of abnormal index return. Bars are ninety per cent confidence intervals from heteroskedasticity-robust standard errors computed from the full covariance block. The surprise is the semiconductor index's abnormal return over the pair of closes the keynote is priced into and is zero on every other trading day; the market's own same-day return is a regressor. Colours group the funds by block.

I then estimate the same specification for each of the exchange-traded funds in the price dataset,

$$r_{p,d} = a_p + b_p z_d + \gamma_p (z_d C_d) + \vartheta_p C_d + \Lambda_p' D_d + \rho_p r_{p,d-1} + \beta_p r_d^{\text{mkt}} + u_{p,d}, \tag{5}$$

where $r_{p,d}$ is the fund's log price change in per cent, C_d is the day's capability share, D_d is four day-of-week indicators, and r_d^{mkt} is the market's own same-day return. The reported quantity is the net response $b_p + \gamma_p C$, evaluated at $C = 0.891$, the total capability share of the delivered-capability profile in Section 5.2. Standard errors are heteroskedasticity-robust, and the net response's standard error is computed from the full covariance block.

As I can see in Figure 8, equity repricing is largest and significant in portfolios closest to the semiconductor industry. All seven semiconductor portfolios have positiv responses, ranging from 0.86 to 1.07 per cent per percentage point of abnormal sector return. The presenter-excluded portfolio responds by 0.86 per cent, indicating that the association extends beyond the announcing firm's own equity.

The response is relatively smaller but still significantly positive in the five AI-themed funds, between 0.36 and 0.53 per cent. Broader technology funds form the next group: all six funds have significant responses between 0.23 and 0.33 per cent. Three real-estate funds move in the opposite direction, with significant declines between 0.24 and 0.41 per cent.

The estimated responses are largest in semiconductor portfolios, followed by AI-themed and broader technology funds. Some real-estate funds move in the opposite direction. This ordering is consistent with repricing concentrated near the technology producers. Data-centre investment links computing advances to property, equipment, and electricity demand. The broad fund comparison above describes national equity repricing rather than investment in particular facilities. It therefore leaves the local and longer-run effects of data-centre construction outside the scope of this cross-section.

6.2 Subsequent economic outcomes

The preceding section examines how markets reprice claims on future outcomes when technology information is disclosed. I now ask what follows in realized prices, production, and investment. Monthly local projections trace subsequent outcomes associated with the same event surprisess and content profiles. This exercise complements the daily evidence at a different frequency and over a two-year horizon.

Having examined contemporaneous financial repricing, I turn to subsequent economic outcomes. Monthly local projections trace the responses associated with the same event surprisess and content profiles over the following two years.

I firstly aggregate the baseline event surprises using the within-month weights of Ottonello and Winberry (2020), as applied to high-frequency surprises by Drechsel and Miura (2025). An event's weight w_d is the share of the month remaining after day d. The return interactions are $Z_m^P = \sum_{d\in m} w_d z_d P_d$ and $Z_m^F = \sum_{d\in m} w_d z_d F_d$. The 112 event days occur in 82 distinct months. The content-level controls P_m and F_m capture the month's content composition.

For each outcome y and each horizon $h = 0, \ldots, 24$ months,

$$\begin{aligned} 100\big(\log y_{m+h} - \log y_{m-1}\big) \;=\; & a_h + b_h z_m + g_h^P Z_m^P + g_h^F Z_m^F + \vartheta_h^P P_m + \vartheta_h^F F_m \\ & + \textstyle\sum_{j=1}^{3} \rho_{h,j}\, 100\big(\log y_{m-j} - \log y_{m-j-1}\big) + \lambda_h^{\mathrm{GFC}} D_m^{\mathrm{GFC}} + \lambda_h^{\mathrm{C}} D_m^{\mathrm{C}} + \kappa_h D_m^{\mathrm{Conf}} + \varepsilon_{m+h}, \end{aligned} \tag{6}$$

where $Z_m^P = \sum_{d\in m} w_d z_d P_d$ and $Z_m^F = \sum_{d\in m} w_d z_d F_d$ weight the surprise by the realised and forward-looking capability shares of Section 4.2, P_m and F_m are the month's mean shares, and D_m^{GFC} and D_m^{C} are indicators for 2008–09 and for the pandemic and D_m^{Conf} for a month containing a keynote. The logs are dropped, and the same equation estimated on level differences, for any series already quoted as a rate or a spread and for any series taking a non-positive value: the federal funds rate, Treasury yields, the

inventory-to-sales ratio and the survey and uncertainty indices are estimated that way, and core CPI, the producer price indices, industrial production, employment, orders and business-equipment production in logs. The lag block contains three monthly changes in the same transformation as the outcome. Crisis indicators enter additively rather than as slope interactions. Standard errors are Newey–West at $h + 1$ lags and the sample begins in January 2003. The event screens and content measures follow Sections 4 and 4.2.

Both capability shares enter jointly, so each interaction is conditional on the other. I report the net response

$$\text{net}_h(P, F) = b_h + g_h^P P + g_h^F F, \qquad \text{Var}(\text{net}_h) = a' V_h a, \quad a = (1, P, F)', \tag{7}$$

where V_h is the covariance block for the three return coefficients. The evaluation points are $(P, F) = (0.829, 0.062)$ and $(0.471, 0.209)$, as in Section 5. They lie within the observed content distribution. The figures report each profile's response and do not provide a direct test of the difference between the profiles.

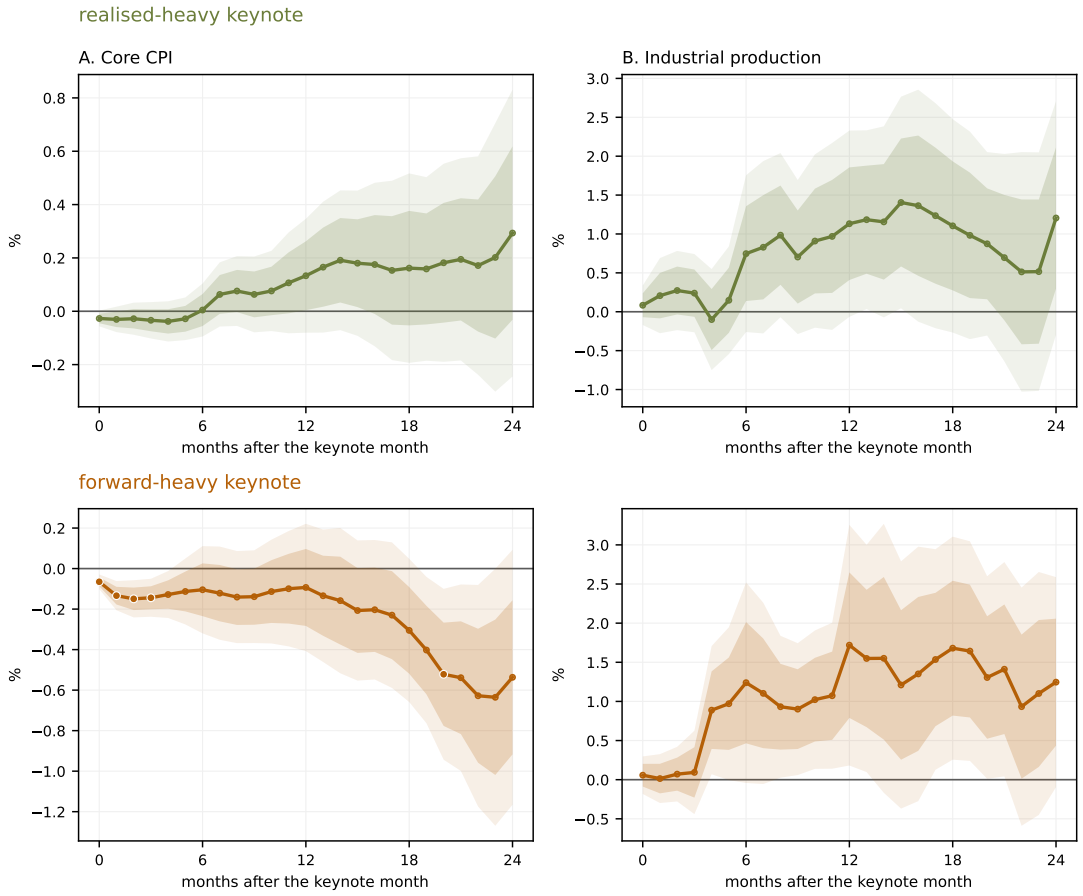
6.2.1 From expectations to realised inflation and economic growth

At the delivered-capability profile, both aggregate responses are positive. Core CPI is approximately unchanged at six months and rises by 0.13 per cent at one year and 0.29 per cent at two years. Industrial-production estimates are also positive. The monthly evidence therefore provides a clearer price–quantity pattern at the promised profile than at the delivered profile. Lower core prices alongside higher output are consistent with an expansion in productive capacity.

The promised-capability profile is associated with lower core prices and higher production. In Figure 9, core CPI declines gradually, by 0.11 per cent at six months and 0.54 per cent at two years. Industrial production rises by 1.25 per cent at six months and 1.72 per cent at one year before returning to 1.25 per cent at two years.

At comparable horizons, the monthly outcomes closely align with the earlier expectations evidence. At the delivered-capability profile, the one- and two-year inflation-swap responses and subsequent core-price responses are positive, although statistically insignificant. Dividend-futures and industrial-production point estimates are also positive, with stronger statistical evidence at second-year horizons. At the promised-capability profile, the negative two-year inflation-swap response and positive dividend-futures response align with subsequent declines in core prices and increases in production, with stronger statistical evidence at second-year horizons. These patterns provide complementary evidence of consistency between market expectations and subsequent economic outcomes.

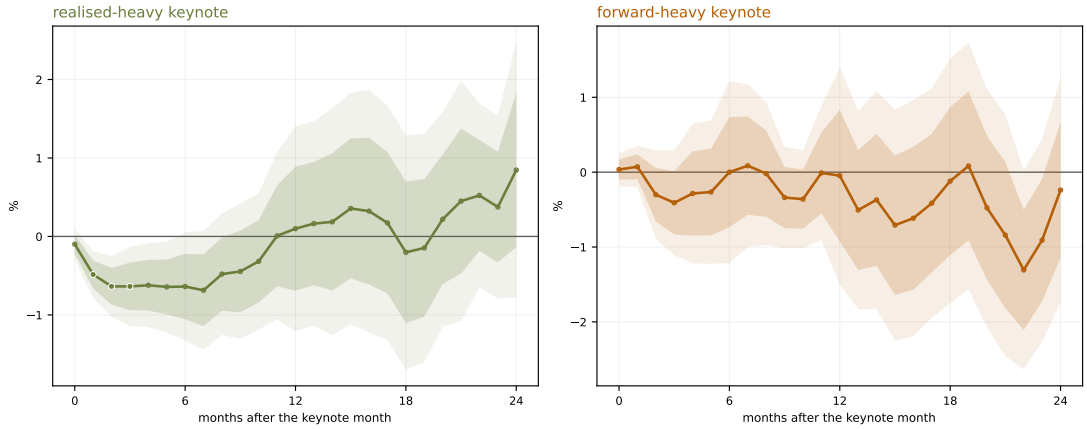
Figure 9: Core prices and output after the two kinds of keynote.



Notes: Monthly local projections from equation (6), with Newey–West standard errors at $h + 1$ lags. Responses are per percentage point of event surprises, with 68 and 90 per cent pointwise intervals. Top row: $(P, F) = (0.829, 0.062)$; bottom row: $(0.471, 0.209)$. Standard errors use the full covariance matrix. The rows show fitted responses at distinct content profiles; a test of their difference requires a direct covariance-based contrast. Filled markers indicate five per cent significance.

Figure 10 examines semiconductor producer prices. At the delivered-capability speeches, prices decline over the first several months before recovering. Together with the positive subsequent core-CPI point estimates, this pattern is consistent with sectoral price declines coexisting with broader price increases: improvements in computing capability may lower prices within the semiconductor industry while supporting demand elsewhere in the economy. At the promised-capability profile, semiconductor-price point estimates are negative at most horizons, alongside the decline in core CPI.

Figure 10: Semiconductor producer prices by contents



Notes: Monthly local projections from equation (6) for the semiconductor producer price index, in log changes expressed in per cent per percentage point of event return. Left: $(P, F) = (0.829, 0.062)$; right: $(P, F) = (0.471, 0.209)$. Shaded areas show 68 and 90 per cent pointwise confidence intervals, using Newey–West standard errors with $h + 1$ lags and the full covariance matrix of the return coefficients.

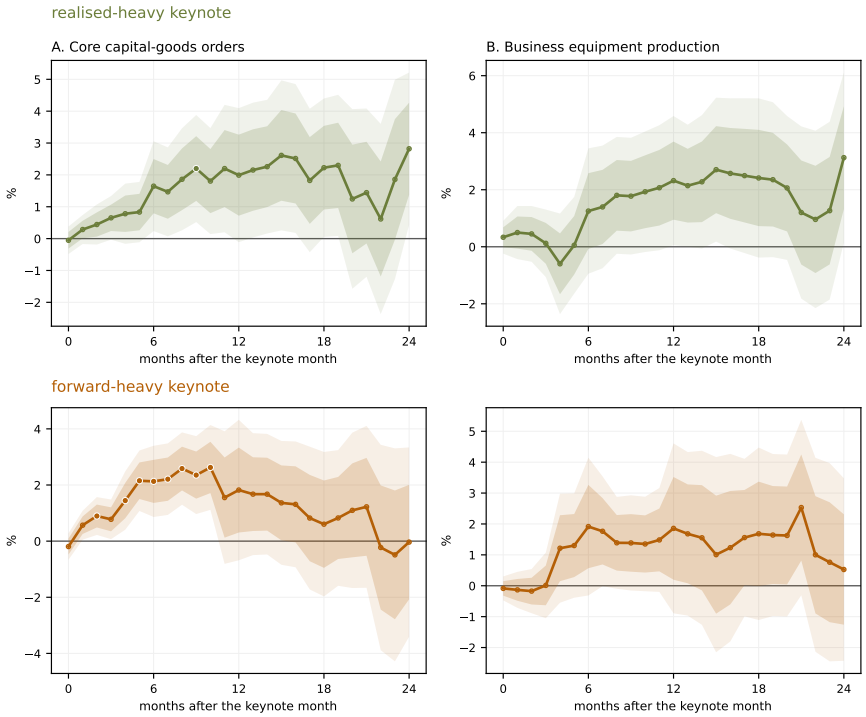
6.2.2 Inputs: orders, equipment and labour

To explore the mechanisms behind these responses, I examine capital-goods orders, equipment production, and labour-market activity. News about improvements in technology may stimulate investment as firms anticipate the benefits of adopting them. Firms may order complementary equipment and adjust labour inputs before these investments translate into greater productive capacity.

Capital-goods orders increase at both content profiles. Figure 11 shows an estimated rise of 1.66 per cent at six months after the delivered-capability profile, increasing to 2.01 per cent at one year and 2.82 per cent at two years. Orders also rise at the promised profile, by 2.15 per cent at six months. The positive responses are statistically significant over substantial portions of both paths. Higher orders are consistent with firms adjusting expenditure in response to technological opportunities, whether the capability is already available or forthcoming.

Business-equipment production is broadly consistent with the increase in core capital-goods orders, with positive point estimates at most horizons under both contents even though the effects are not significant. Some second-year responses are significant at the delivered-capability profile.

Figure 11: Investment after the two kinds of keynote.

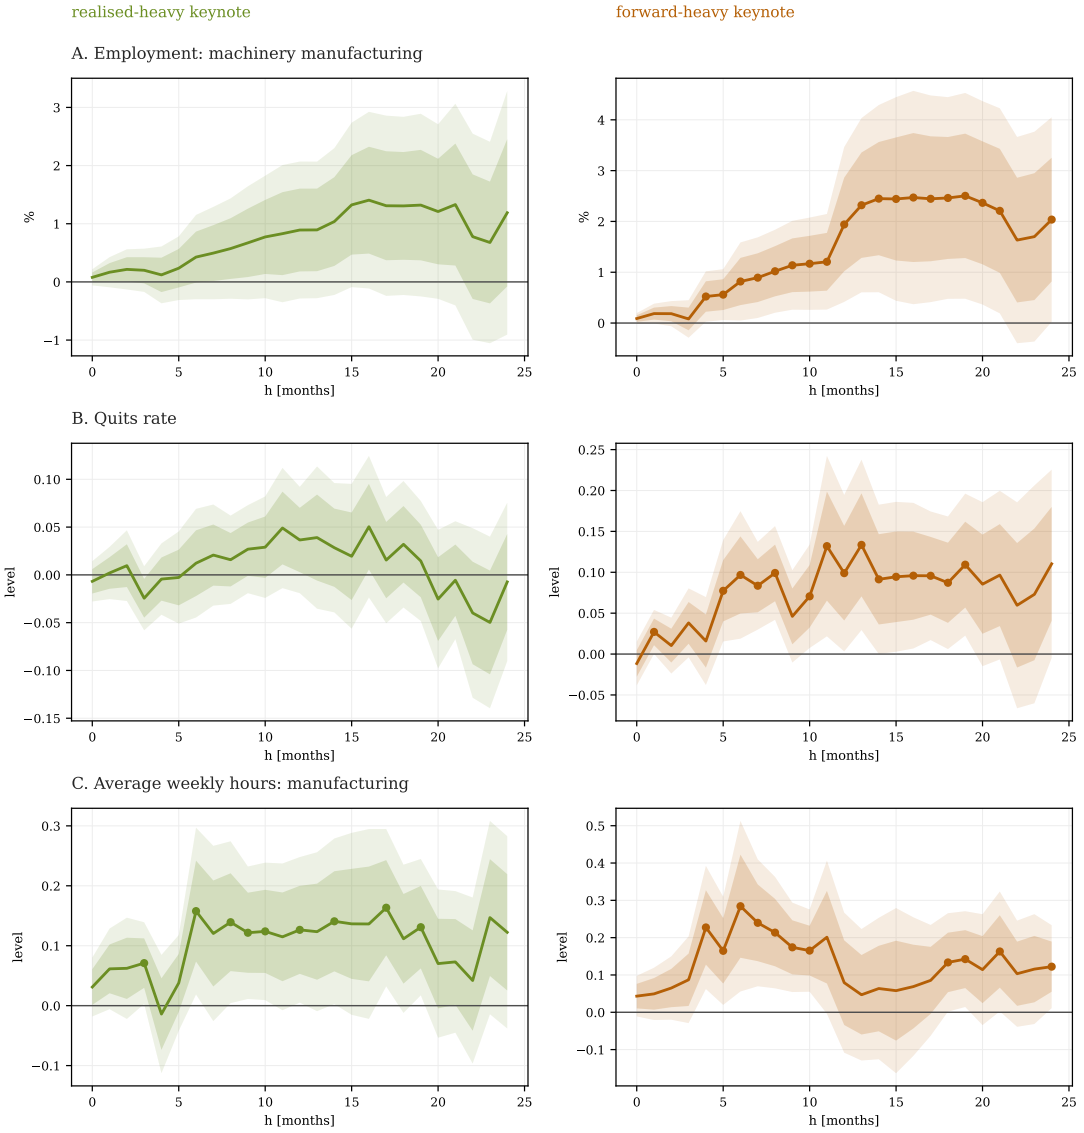


Notes: As Figure 9; top row realised-heavy, bottom row forward-heavy. New orders are for nondefense capital goods excluding aircraft and business-equipment production is an index in logs, both in log changes in per cent.

Figure 12 provides further evidence of adjustments in labour inputs. At the promised-capability profile, machinery-manufacturing employment rises gradually and remains elevated during the second year. The quits rate also increases, while manufacturing hours respond more quickly, rising during the first year before moderating. This timing is consistent with firms initially increasing hours and subsequently expanding employment. At the delivered-capability profile, employment and hours are also mostly positive, but employment is less significant estimated and the quits-rate response is small and statistically insignificant. But both of the two content profiles provide suggestive evidence of broader labour-market expansion.

Taken together, the increases in capital-goods orders at both profiles, alongside mostly positive equipment-production estimates, are consistent with an investment channel through which firms respond to technological opportunities. At the promised-capability profile, stronger labour-market activity provides further evidence of adjustments in productive inputs. These patterns support the investment and input-adjustment interpretation.

Figure 12: Labour responses by contents



Notes: Monthly local projections from equation (6). Left: delivered-capability profile, $(P, F) = (0.829, 0.062)$; right: promised-capability profile, $(P, F) = (0.471, 0.209)$. Responses are per percentage point of event return. Machinery-manufacturing employment is expressed as a log change in per cent; the quits rate and average weekly hours enter in level differences, measured in percentage points and hours, respectively. Shaded areas show 68 and 90 per cent pointwise confidence intervals, using Newey–West standard errors with $h + 1$ lags and the full covariance matrix of the return coefficients. Filled markers indicate significance at the five per cent level.

7 Interpreting the results relative to existing evidence

7.1 The sign of the inflation response to technology news

The existing literature documents mixed inflation responses to favorable technology news (Table 5). Barsky and Sims (2011) and Barsky, Basu and Lee (2015) find

Table 5: Macroeconomic Responses to Positive Technology News

Study	Inflation	Investment	Real GDP	Hours	Interpretation
Barsky and Sims (2011)	↓	↓ impact	↓ impact	↓ impact	Expected future productivity gains → Lower expected marginal costs → Disinflation.
Barsky, Basu and Lee (2015)	↓	↓ impact; ↑ later	↑ later	↓ impact; ↑ later	Expected future price declines → Delayed spending → Lower marginal costs → Disinflation.
Ahmed and Farah (2022)	↓	↑	↑	↑	Future IT news stimulates investment and activity; inflation falls in an extended specification.
Cascaldi-Garcia and Vukotić (2022)	↑	↑	↑	↑	Patent-based news → Economic expansion and higher inflation, consistent with NK models.
Aldasoro et al. (2024) (Unanticipated)	↓ initially	↑	↑	—	Productivity gains → Supply expansion → Lower inflation initially.
Aldasoro et al. (2024) (Anticipated)	↑ initially	Delayed; ↑ later	↑ gradually	—	Expected future income → Current consumption demand → Higher inflation.
Evgenidis and Fasianos (2025)	Insignificant	↑	↑	↑	AI news raises productivity and activity, with no significant inflation response.

Notes: Arrows summarize reported responses, not significance at every horizon. Aldasoro et al. (2024) present model simulations. Ahmed and Farah (2022) report inflation in an extended specification. The inflation measures and identification methods differ across studies.

disinflationary responses to news about future productivity improvements. Similarly, Ahmed and Farah (2022) show that future IT news raises investment, output, and labor input while reducing inflation in an extended specification. These findings closely resemble our promised-capability results, where growth prospects and subsequent capital-goods orders increase while inflation expectations declines. In contrast, Cascaldi-Garcia and Vukotić (2022) find that patent-based technology news raises inflation alongside economic activity, resembling our delivered-capability results. Importantly, patent news also contains information about future technological progress. Thus, even news about future productivity improvements can generate inflation responses of opposite signs.

One explanation concerns the relative strength and timing of anticipated demand expansion and improvements in productive capacity. Expectations of higher future productivity may lower anticipated marginal costs and inflation Barsky and Sims (2011); Barsky, Basu and Lee (2015). At the same time, favorable technology news may stimulate consumption, investment, and technology adoption before productivity gains materialize Beaudry and Portier (2006); Comin, Gertler and Santacreu (2009). Re-

cent AI models emphasize this tension: Aldasoro et al. (2024) show that unanticipated productivity improvements can initially lower inflation, whereas anticipated future improvements may stimulate current demand and raise inflation. Faria-e Castro and Ozkan (2026) similarly highlight inflationary pressures arising from optimism about future AI productivity. These mechanisms suggest that the net inflation response depends on the relative timing of demand and supply adjustments.

Our analysis introduces an observable dimension of technology news relevant to this balance: the availability horizon of disclosed capabilities. Using pre-scheduled semiconductor conferences, we distinguish between capabilities already available (delivered) and those expected to become available in the future (promised). Both may convey positive news about future aggregate productivity, but technological availability need not coincide with the realization of economy-wide productivity gains. Available technologies may require complementary investment, adoption, and diffusion, while forthcoming technologies may stimulate investment in advance.

Our results reveal contrasting inflation-compensation responses across these disclosure profiles. Delivered-capability-heavy surprises raise both dividend-futures prices and inflation expectations, consistent with adoption-related investment demand preceding the full realization of productivity gains. Promised-capability-heavy surprises also raise dividend-futures prices but reduce inflation expectations, consistent with expectations of future productive-capacity expansion and lower production costs. The semiconductor PPI responses in Table 10 reinforce this interpretation: under the promised-capability profile, estimated responses remain negative and trend downward over two years, whereas under the delivered-capability profile, they are initially negative before trending upward. Importantly, capital-goods orders increase under both profiles. Promised-capability disclosures are also followed by lower core prices, higher industrial production, and rising manufacturing hours and machinery employment, while hours responses under the delivered profile are mostly positive. The labor responses are broadly consistent with Ahmed and Farah (2022) and Cascaldi-Garcia and Vukotić (2022), potentially reflecting labor requirements associated with investment and technology adoption. Thus, investment and labor expansion need not imply higher inflation expectations.

Our contribution is not to establish that future technology news can be disinflationary, but to show that the availability horizon of disclosed technological capabilities is associated with opposite inflation-compensation responses within a common disclosure setting, even as both profiles support growth prospects and investment. These findings suggest that information composition may help explain heterogeneous inflation responses in the technology-news literature, potentially reflecting differences in the relative importance and timing of anticipated demand expansion and productive-capacity improvements.

7.2 Mapping the results to the AI–macro literature

The preceding subsection considers why technology news can have different inflation implications. I now compare the magnitude of the estimated responses with the AI–macro literature. I focus on one-, five-, and ten-year horizons, distinguishing the incremental information in an individual conference disclosure from the broader economic gains associated with AI adoption.

To express the dividend responses in macroeconomic units, I adopt the illustrative conversion used by Hazell and Hobler (2024), drawing on Gormsen and Koijen (2020). I subtract the revision to the expected price level from the nominal dividend response and multiply the resulting real-dividend revision by 0.67. Let $d_m(C, F)$ denote the fitted nominal dividend response in per cent and $\pi_m(C, F)$ the fitted annualised inflation-swap response in basis points, both per percentage point of sector return. The approximate GDP-equivalent revision is

$$\widehat{\Delta Y}_m^{\text{real}}(C, F) \approx 0.67 \left[d_m(C, F) - \frac{m\, \pi_m(C, F)}{100} \right]. \tag{8}$$

Multiplying the inflation response by maturity converts the annualised rate revision into an approximate cumulative price-level revision. The resulting figures describe revisions to the expected real GDP level at each horizon, rather than changes in annualised GDP growth.

Table 6 places these magnitudes alongside existing evidence. At the delivered-capability profile, a one-percentage-point sector return corresponds to approximate real GDP revisions of 0.022, 0.017, and 0.035 per cent at one, five, and ten years. At the promised-capability profile, the corresponding revisions are 0.019, 0.049, and 0.105 per cent. The larger medium- and long-horizon real revisions at the promised profile reflect its negative inflation response, which raises the inflation-adjusted value of the nominal dividend revision. This arithmetic does not establish statistically significant differences in real GDP responses.

Table 6: The estimated magnitudes relative to the AI–macro literature

Study	Method	Growth-related outcome	Inflation-related outcome
This paper	Event returns and disclosure content	Illustrative real GDP revisions at 1, 5, and 10 years: delivered profile +0.022, +0.017, and +0.035 per cent; promised profile +0.019, +0.049, and +0.105 per cent	inflation expectations at 1, 5, and 10 years: delivered profile +0.39, +0.87, and +0.81 basis points; promised profile +0.26, −0.91, and −0.61 basis points
Acemoglu (2025)	Task-based calibration	Ten-year TFP gains of 0.53–0.66 per cent; GDP gains of 0.93–1.16 per cent under modest investment adjustment	Not quantified
Blumenfeld et al. (2026)	Equity-market exposures and a general-equilibrium mapping	GDP-level gain of 3.6 per cent, or 6.5 per cent with an R&D channel, based on expectations accumulated through December 2025	Not quantified
Aldasoro et al. (2024)	DSGE model	Positive activity response	Direction depends on anticipation: unanticipated gains disinflationary, anticipated gains inflationary
Faria-e Castro and Ozkan (2026)	DSGE model with technology news	Positive activity response	Inflationary through an AI-optimism demand channel
Görtz, Gunn and Lubik (2025)	VAR with technology news	Positive activity response, peaking near ten years	Demand-like response
Andrews and Farboodi (2025), Wilding and Scott (2026)	Event studies of AI model releases	Declining long-dated yields; a different payoff from dividend futures	Not directly comparable

Notes: This paper's entries are per percentage point of sector return, evaluated at the delivered- and promised-capability profiles, $(C, F) = (0.891, 0.062)$ and $(0.680, 0.209)$. GDP-equivalent revisions use equation (8); inflation entries are revisions to annualised inflation expectations. The GDP conversion combines point estimates from regressions with different available samples and supplies no separate significance test. Other studies concern different outcomes, horizons, and economic experiments.

The implied revisions provide an economically meaningful scale for the information released around conference keynotes. At the promised-capability profile, a one-percentage-point sector return corresponds to an illustrative ten-year real GDP revision of approximately 0.11 per cent. For context, Acemoglu (2025) estimates ten-year GDP gains of 0.93–1.16 per cent, while Blumenfeld et al. (2026) estimate gains of 3.6–6.5 per cent. Our estimate concerns an incremental revision around a disclosure, whereas these studies quantify broader adoption scenarios or expectations accumulated over several years. A revision of around one-tenth of a per cent therefore need not be small in the context of an individual disclosure, even though it is below estimates of AI's overall economic gains.

The conversion is illustrative. It assumes that dividend and inflation risk premia are approximately stable over the event window and that the dividend-to-GDP relationship used by Hazell and Hobler (2024) applies to these disclosures. Using the same factor at five and ten years also extends its short-horizon application. Ac-

cordingly, the table reports GDP-equivalent magnitudes rather than directly estimated GDP responses.

The Treasury results provide a complementary comparison. At the promised-capability profile, five- and ten-year TIPS yields fall by approximately 0.70 and 0.75 basis points per percentage point of sector return (Appendix I). Andrews and Farboodi (2025) also document declining real yields around major AI releases, although differences in events and normalisation prevent a direct magnitude comparison. Lower safe real yields need not contradict higher growth expectations implied by dividend futures, investment, and production: these outcomes differ from the household consumption prospects relevant for safe-asset pricing, and long-term yields also reflect term and liquidity premia. The present analysis does not distinguish these channels.

8 Heterogeneous identification

In this section, I examine how the results depend on the return window, market adjustment, and joint issuer–peer repricing. These exercises change the information captured by the return and, in some cases, the event sample. For the intraday analysis, the timing and coverage requirements exclude 37 of the 112 baseline event days, leaving 75 days under both the ninety- and 120-minute windows. I first examine event-window variance under the alternative measures and then assess the corresponding macroeconomic-expectation responses.

8.1 Intraday identification

The intraday sample requires an observed return window within 09:30–16:00 Eastern Time. Both window lengths retain 75 event days. Daily–intraday comparisons therefore combine a change in the return window with a change in event composition. The number of events entering each outcome regression also depends on contract-price availability.

Keynotes vary substantially in duration, from approximately thirty minutes to three hours. I use ninety- and 120-minute windows to capture the main presentation and its immediate market reception while limiting exposure to other news during the trading day. The ninety-minute window provides the narrower measure, while the 120-minute window allows more time for announcements and market adjustment. Fixed windows facilitate comparison across events, although they may extend beyond shorter keynotes or end before longer presentations conclude.

The intraday surprise replaces the close-to-close return of Section 4.1 with the SOXX return over a window beginning at the keynote's recorded start time:

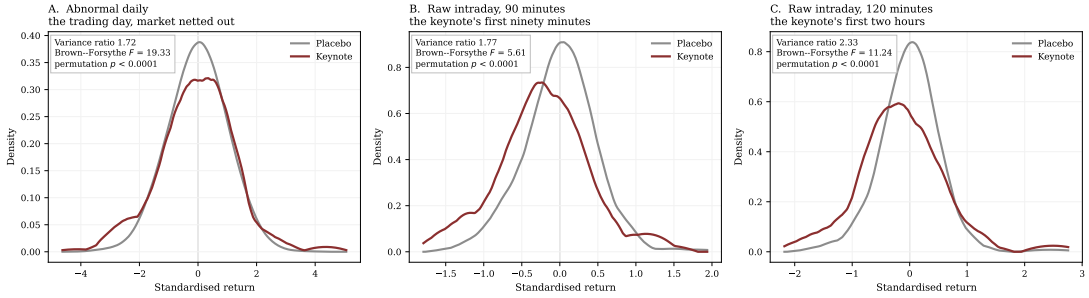
$$z_e^W = 100 \left[\log P_{\mathrm{SOXX},\, t_e+W} - \log P_{\mathrm{SOXX},\, t_e} \right], \qquad W \in \{90, 120\} \text{ minutes},$$

where t_e is expressed in Eastern Time. I retain raw sector returns to capture the broader equity response to the disclosure. A market-adjusted daily measure is reported separately, allowing us to examine market adjustment and window length as distinct choices. Narrow windows reduce the opportunity for unrelated information to arrive.

8.1.1 Event-window variance

I first apply the variance diagnostic to market-adjusted daily returns and raw ninety- and 120-minute returns. Non-event controls are matched by calendar position and, for intraday returns, by clock time. Before pooling, each event return and its controls are divided by the standard deviation of that event's matched control returns. This accounts for differences in normal volatility across events, so the comparison asks whether event returns are unusually large relative to their local benchmark.

Figure 13: Event-window variance under alternative return measures



Notes: Event and matched-control returns under market-adjusted daily, raw ninety-minute, and raw 120-minute measures. Controls match calendar position and, for intraday returns, clock time. Each event and its controls are scaled by the standard deviation of that event's matched control returns.

Event-window variance exceeds matched-control variance under all three alternatives. The variance ratio is 1.72 for market-adjusted daily returns, 1.77 for ninety-minute returns, and 2.33 for 120-minute returns. Figure 13 presents the comparisons and their inference. These results support elevated volatility within the respective designs.

8.1.2 Macroeconomic-expectation responses

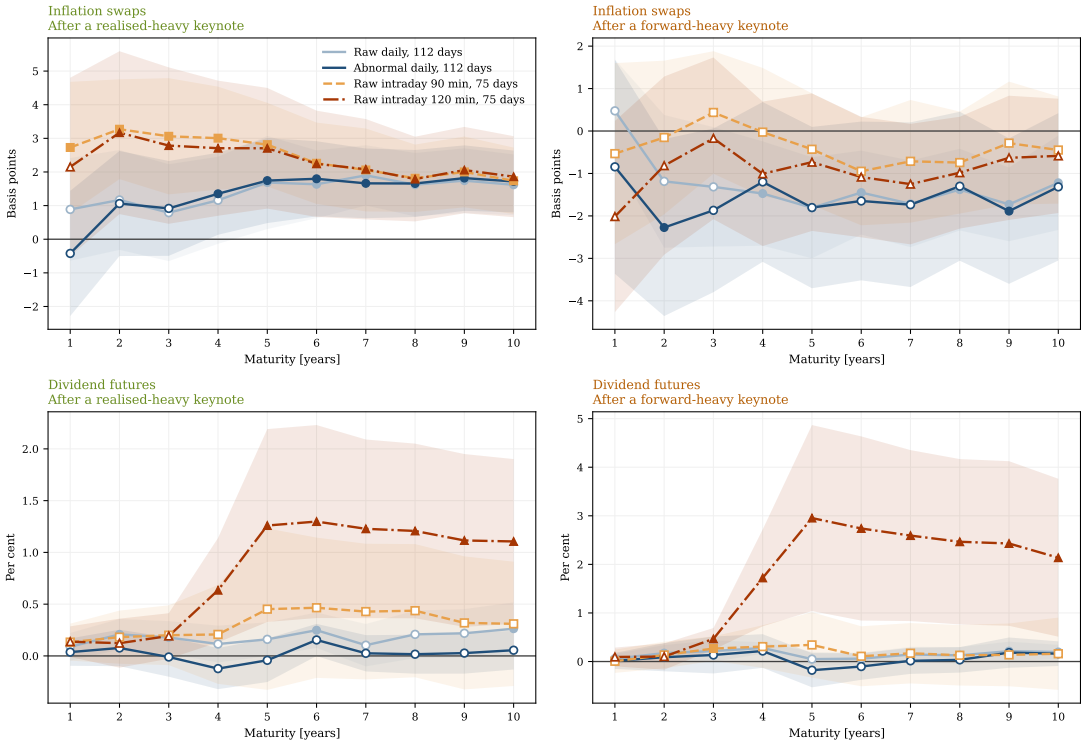
I next estimate equation (4) using each alternative return measure. The specification retains the content shares, the full trading calendar with zero event regressors on non-event days, four weekday indicators, one lag of the dependent variable, and HC1 standard errors.

For Figure 14, I standardize each return measure within its own event sample before constructing the interactions. Daily returns vary substantially more than ninety-

and 120-minute returns, so coefficients expressed per percentage point can differ simply because the measures have different scales. Standardization puts the responses on a common statistical scale: each is expressed per standard deviation of its own return measure. Specifically, I subtract the event-sample mean and divide by the event-sample standard deviation, leaving non-event values at zero. This normalization facilitates comparison across measures.

The delivered- and promised-capability evaluation points are also selected within each measure's event sample. Consequently, daily–intraday comparisons reflect changes in content evaluation points as well as return construction, coverage, and scaling. Confidence intervals for fitted responses use the full covariance matrix of the return and interaction coefficients.

Figure 14: Content-dependent responses under alternative surprise measures



Notes: Fitted responses from equation (4) using raw and market-adjusted daily returns and raw ninety- and 120-minute returns. Each measure is demeaned and scaled by its event-sample standard deviation; non-event values remain zero. This measure-level standardization differs from the event-specific control-volatility scaling in Figure 13. Responses are evaluated at the mean content of the top decile of delivered and promised capability within each event sample. Daily measures cover 112 event days and intraday measures 75, before outcome-specific availability restrictions. Confidence intervals use the full covariance matrix of the return and interaction coefficients. Top row: inflation swaps. Bottom row: dividend futures.

The positive inflation-compensation response at the delivered-capability profile appears under all four measures. Both daily measures produce positive and statisti-

cally significant responses over the longer part of the curve. The ninety- and 120-minute measures also produce positive responses, with confidence intervals excluding zero at most maturities. The association therefore appears within narrower keynote windows as well as after daily market adjustment.

At the promised-capability surprises, inflation-expectations point estimates remain negative at most maturities across the alternative measures, consistent with the baseline pattern. The alternative specifications therefore preserve the broad direction of the baseline response, although statistical significance is weaker. Changes in event coverage and content evaluation points also contribute to these comparisons. Dividend-futures responses likewise remain broadly positive, consistent with the baseline under both contents' surprises, although their magnitudes and statistical significance vary across measures. The raw 120-minute measure produces larger positive responses than the other specifications. Overall, the alternative measures broadly preserve the baseline pattern of positive dividend responses alongside content-dependent inflation responses, with some estimates becoming statistically insignificant.

8.2 Common versus relative equity repricing

A semiconductor-sector return can reflect information about the industry as a whole or changes in the competitive positions of individual firms. An announcement may improve the presenting firm's outlook while reducing its competitors' expected profitability. I therefore examine whether the baseline content-dependent inflation response remains present when the presenting firm and its semiconductor peers reprice in the same direction.

For each keynote, I compare the presenting firm's abnormal return with the equal-weighted abnormal return of its semiconductor peers over the same pricing window. The peer portfolio excludes the presenter. I classify an event as *common repricing* when the two returns have the same sign, retaining both joint increases and joint declines. Events with opposite signs are classified as *relative repricing*.

I adjust both returns for market exposure before taking their signs. For each firm and event, I estimate

$$r_{i,t} = a_i + b_i r_{m,t} + u_{i,t}$$

using consecutive non-conference trading-day pairs within sixty calendar days of the event, excluding the event day and requiring at least twenty usable pairs. The abnormal return is the event-window return less its market-model prediction. I then average the peer firms' abnormal returns to construct the peer leg.

This adjustment allows market exposure to differ across firms. The estimated market betas have a median of 1.52 and a tenth-to-ninetieth percentile range of 0.93 to 2.46. Simply subtracting the market return would impose a unit beta and could leave

a common market component in both legs. The market-model adjustment reduces the possibility that sign agreement reflects this shared exposure.

The classification covers 129 keynotes on 126 event days. Of these, 90 keynotes on 88 days exhibit common repricing, while 39 keynotes on 38 days exhibit relative repricing. A day containing multiple keynotes enters a group only if all its keynotes receive the same classification. No day is mixed under the market-model classification. I compare the full sample with the common-repricing subsample. Outcome availability leaves 98–106 event days in the full-sample regressions and 68–71 in the common-repricing regressions, depending on maturity.

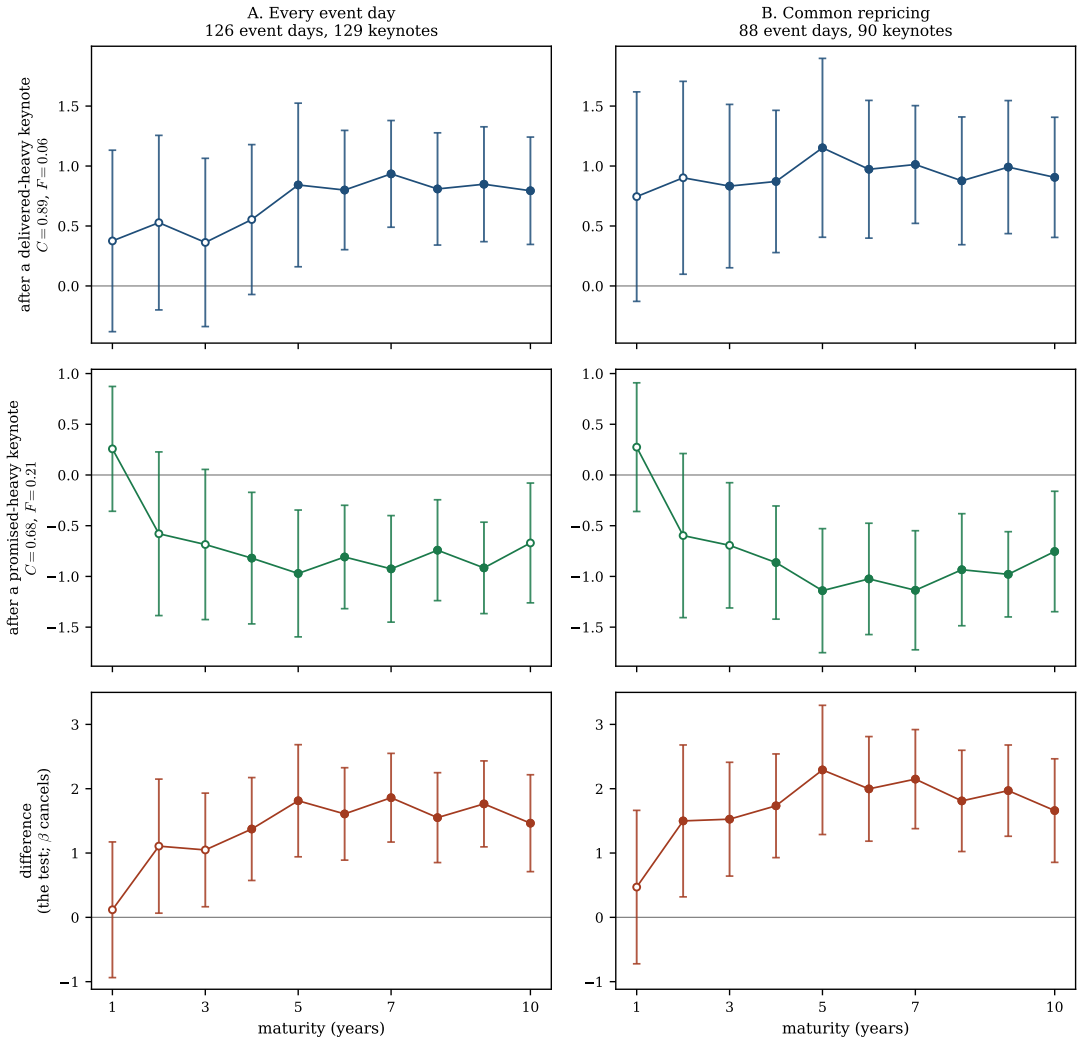
I estimate equation (4) in both samples, retaining the return interactions, content-level terms, and baseline controls. Figure 15 reports the fitted inflation-compensation response

$$R_m(C, F) = \beta_m + \gamma_m C + \delta_m F$$

at the delivered- and promised-capability profiles of Section 5.2. The evaluation points, $(C, F) = (0.891, 0.062)$ and $(0.680, 0.209)$, are held fixed across samples. Differences across columns therefore reflect changes in the estimated response rather than changes in the content profile at which it is evaluated.

The bottom row reports the delivered-minus-promised response difference, with confidence intervals calculated from the full coefficient covariance matrix. This contrast removes β_m and directly tests whether the two content profiles imply different responses. Because the profiles differ in both total capability content and its horizon, it is distinct from the conditional horizon contrast δ_m. Appendix H reports the underlying delivered-capability gradient γ_m, promised-capability gradient $\gamma_m + \delta_m$, and horizon contrast δ_m.

Figure 15: Inflation-compensation responses in the full and common-repricing samples



Notes: Equation (4) estimated on the full sample (left) and common-repricing subsample (right). Rows show the delivered-profile response, promised-profile response, and their difference, evaluated at the same content profiles across samples. Bars denote 90 per cent confidence intervals based on the full coefficient covariance matrix; filled markers indicate significance at five per cent. Common repricing requires issuer and peer abnormal returns to have the same sign. The regressions use raw daily sector returns.

Restricting the sample to common-repricing events preserves the baseline pattern. At the delivered-capability profile, inflation-compensation responses are positive across the maturity curve and statistically significant from three years onward. The estimated increase is 1.15 basis points at five years and 1.01 at seven years per percentage point of sector return. At the promised-capability profile, responses are negative beyond one year and statistically significant from four years onward. The decline reaches 1.14 basis points at five years and 0.98 at nine years.

The direct profile comparison confirms this separation. In the common-repricing sample, the delivered-minus-promised difference is positive and statistically signifi-

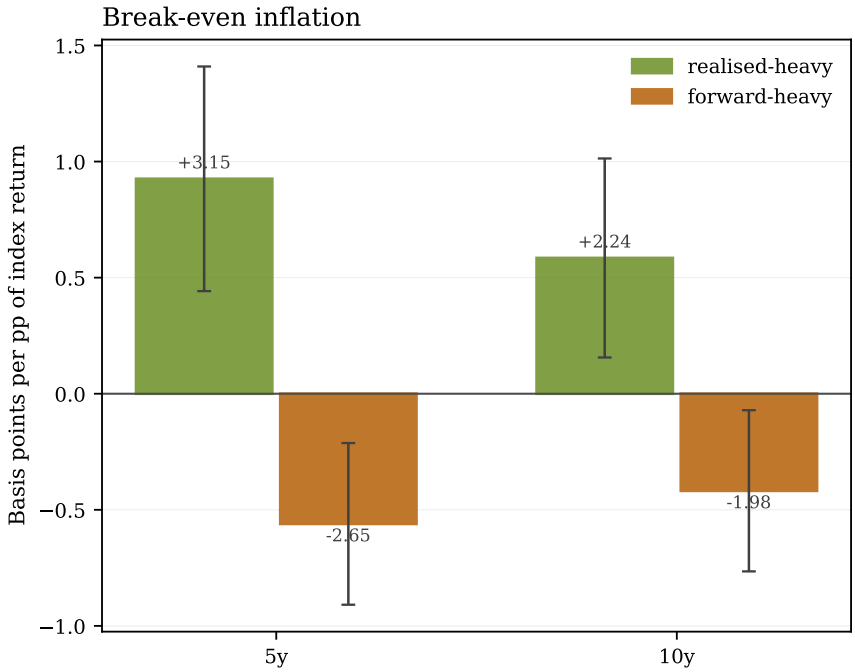
cant at the five per cent level at every maturity from two through ten years. It reaches 2.29 basis points at five years and 2.15 at seven years. The corresponding difference is significant at seven maturities in the full sample. The content-dependent inflation response therefore remains present when issuer and peer valuations move together.

This exercise shows that the baseline pattern does not depend on events with opposing issuer and peer returns. It does not establish that responses are larger under common repricing: pooled tests provide limited evidence of group differences in the capability gradient and no statistically significant differences in the horizon contrast. Nor does sign agreement alone identify an aggregate technology shock, since common demand or financial news can also move firms together. The sign restriction nevertheless addresses the concern that the sector return primarily captures competitive redistribution, complementing the timing and market-adjustment exercises above.

9 Robustness checks

9.1 Treasury break-even inflation

Figure 16: Treasury break-even inflation responses



Notes: Fitted responses from equation (4) at the baseline delivered- and promised-capability profiles, in basis points per percentage point of sector return. Bars denote 90 per cent confidence intervals based on the full robust coefficient covariance matrix. Labels report t statistics.

I assess whether the baseline inflation-compensation results extend beyond the inflation-swap market by applying equation (4) to daily changes in Treasury break-even

inflation. The content's effects reproduce the main pattern. Replacing demand with delivered capability, holding promised capability fixed, raises five-year break-even inflation by 4.03 basis points per percentage point of sector return. Shifting capability from delivered to promised, holding total capability fixed, lowers the response by 4.33 basis points at five years and 4.07 basis points at ten years. These estimates are statistically significant.

Figure 16 shows that the fitted responses also preserve the baseline direction at the same content profiles, $(C, F) = (0.891, 0.062)$ and $(0.680, 0.209)$. At the delivered-capability profile, five- and ten-year break-even inflation rise by 0.93 and 0.58 basis points, respectively. At the promised-capability profile, they fall by 0.56 and 0.42 basis points. All four responses are statistically significant at the five per cent level. Full Treasury results appear in Appendix I. The agreement across markets supports the robustness of the content-dependent inflation-compensation response.

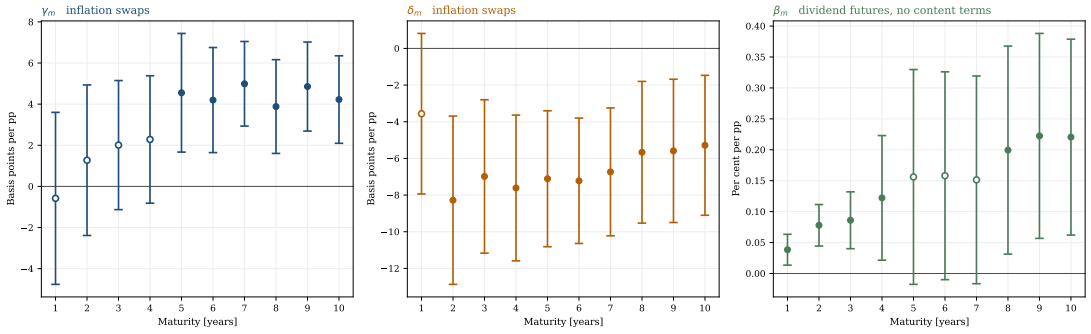
9.2 Alternative samples

I examine how the results change when I relax the sample restrictions in Table 9. The first alternative retains the study window and qualifying speech-record requirement but removes the event screens. The second includes all recorded keynote-days and estimates the specification without content terms.

9.2.1 Removing the event screens

Removing the macroeconomic-release, non-trading-day, and issuer-earnings screens yields 125 keynote-day records on 120 trading days. The screens exclude seven keynote-days overlapping FOMC statements or CPI releases, three opening keynotes held during US long weekends without a price on the recorded day, and one window containing the speaking firm's earnings announcement. I recompute returns from the price data using the same previous-close-to-close rule.

Figure 17: Results without the event screens



Notes: Estimates on the alternative event set of 125 keynote-day records on 120 trading days, before outcome-specific availability restrictions. Left and centre: γ_m and δ_m from equation (4) for inflation swaps. Right: the dividend-futures specification without content terms. Bars denote 90 per cent robust confidence intervals; filled markers indicate significance at ten per cent.

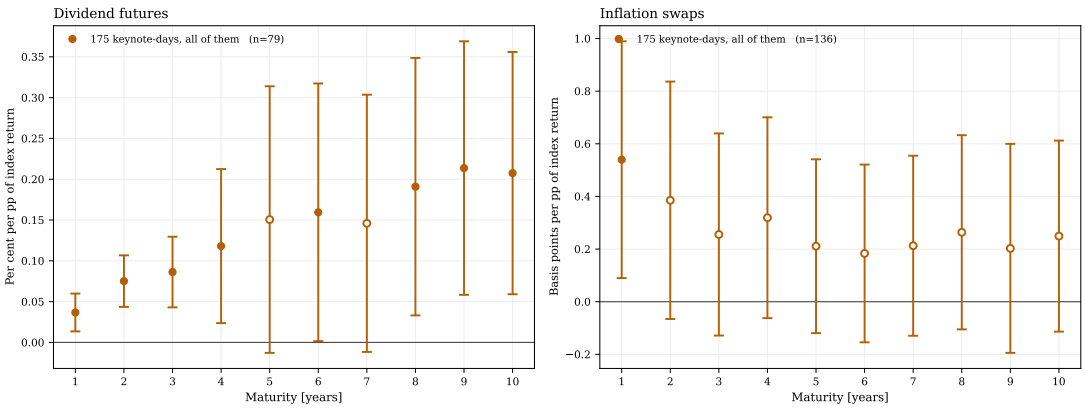
Figure 17 shows the sensitivity of the estimates to removing these screens. The exclusions are motivated by overlapping information rather than the resulting coefficient signs. For example, Intel's IDF on 18 September 2007 coincided with a fifty-basis-point FOMC rate cut and gains of approximately 2.9 per cent in both the semiconductor index and the S&P 500. The daily return therefore combines the keynote with a major monetary policy announcement.

The twenty-seven keynote-days before 21 July 2004 do not provide additional observations for the displayed regressions: they precede both the inflation-swap and dividend-futures histories. They can, however, enter exercises that do not require these outcomes, such as the return-variance diagnostic.

9.2.2 Including all recorded keynote-days

The broader collection contains 175 keynote-day records. Of these, twenty-three lack a qualifying speech record: seventeen contain material other than a transcript or equivalent speech record, and six contain a transcript of a different presentation. Content shares cannot be constructed for these observations. Setting C and F to zero would incorrectly treat missing content as an absence of capability news. I therefore estimate only the specification without content terms on this broader event set.

Figure 18: Results using all recorded keynote-days



Notes: The specification without content terms, estimated using the event set of 175 keynote-day records on 161 trading days, before outcome-specific availability restrictions. Left: dividend futures. Right: inflation swaps. Bars denote 90 per cent robust confidence intervals; filled markers indicate significance at ten per cent.

This broader event set contains forty-nine more trading days than the baseline. The dividend-futures coefficients move towards zero, and the number of significant

maturities falls from nine to eight. The inflation-swap response remains significant at only one maturity. The broader

9.3 Controlling for ordinary sector comovement

Equation (4) prices the keynote-day surprise only: z_t is the semiconductor index return on a conference day and zero on the remaining 98 per cent of the trading calendar. inflation expectations and the semiconductor sector may comove on ordinary days as well, through common exposure to growth or policy news. If conference days fall in periods when that comovement is unusually strong, or when the content shares happen to be unusually high or low, the interaction terms could reflect the comovement rather than the disclosures.

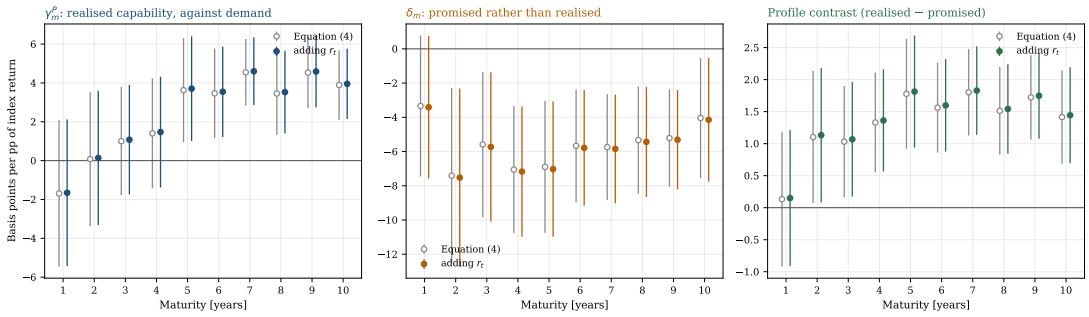
To address this I add the semiconductor index return on *every* trading day as a separate regressor,

$$\Delta y_{m,t} = \alpha_m + \psi_m r_t + \beta_m z_t + \gamma_m Z_t^C + \delta_m Z_t^F + \theta_m C_t + \phi_m F_t + \Gamma_m' X_t + \varepsilon_{m,t}, \quad (9)$$

where r_t is the index return over the same pair of closes as the outcome change. Because r_t is non-zero on all days while z_t is non-zero only on conference days, ψ_m is identified almost entirely off non-conference days. The content effects γ_m and δ_m are then measured net of whatever linear relationship the sector return bears to inflation expectations in ordinary times.

The semiconductor index calendar and the swap quotation calendar do not coincide: the swap curve is quoted on 191 dates between 2004 and 2026 on which the US equity market is closed. I therefore define r_t as the cumulative index return over the interval the swap change spans, so that an interval in which the equity market never opened contributes zero. This keeps the estimation sample identical across the two specifications, at 5,557 observations and 107 conference days, so that equation (9) is equation (4) with one column added rather than a different sample.

Figure 19: The content effects with and without the ordinary sector return



Notes: Inflation-swap changes in basis points per percentage point of index return. Hollow markers are equation (4); filled markers add the ordinary sector return r_t of equation (9). The profile contrast is the difference in fitted net response between $(C, F) = (0.89, 0.06)$ and $(0.68, 0.21)$, computed from the full coefficient covariance matrix. Bars denote 90 per cent robust confidence intervals. Appendix E reports the coefficients, including the fitted levels, which the figure omits because adding r_t shifts the intercept and therefore the levels but not their difference.

Ordinary comovement is present and precisely estimated: ψ_m is positive at every maturity, +0.46 at five years and +0.37 at seven, with t statistics near nine. It is absorbed by the intercept of the event-day response, not by the content terms: β_m falls from −1.93 to −2.43 at five years and from −2.74 to −3.14 at seven, and the fitted net response at the delivered-heavy profile falls with it, from +0.87 to +0.44 at five years. The profile contrast, from which β_m cancels, is unaffected. The sector return and inflation expectations do move together on ordinary days, and the specification absorbs that movement. The content effects are nonetheless unchanged. At five years γ_m moves from +3.63 to +3.71 and δ_m from −6.90 to −7.03; at seven years, from +4.55 to +4.60 and from −5.74 to −5.84. The profile contrast moves from +1.78 to +1.81 at five years and from +1.80 to +1.83 at seven. The counts of maturities significant at ten per cent are identical in the two specifications: six of ten for γ_m, nine of ten for δ_m, and nine of ten for the contrast. Controlling for ordinary comovement leaves the content-dependent responses where they were, and if anything slightly larger.

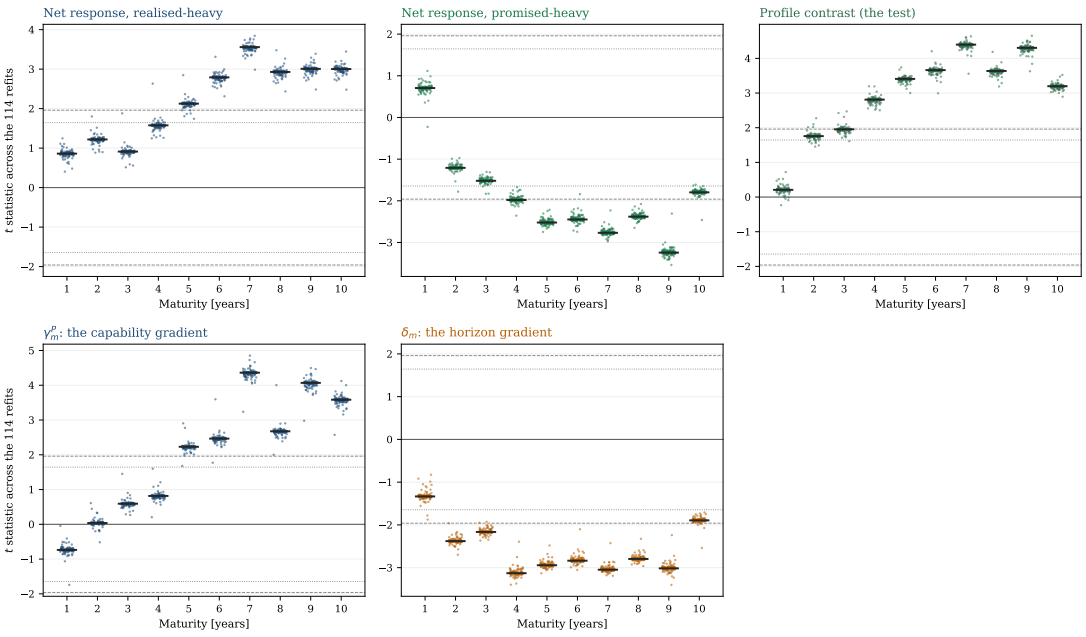
This check removes a common linear relationship between the sector return and inflation compensation. It does not remove a loading that varies over the sample period or across conference series, nor any non-linear relationship, and it is therefore evidence against one specific alternative rather than a general exclusion restriction.

9.4 Leaving out one keynote at a time

A hand-collected event sample of this size invites the question of whether one conference is carrying the result. We therefore re-estimate equation (4) once for every keynote in the sample, each time with that keynote removed, and ask how often the full-sample reading survives.

The exercise is at the level of the keynote rather than the trading day. Two days in the window carry two keynotes each, and a day's content shares are the unweighted mean over the keynotes on it, so dropping one keynote from such a day does not delete an observation: it recomputes that day's shares from the keynote that remains. A day leaves the sample only when the keynote dropped was the only one on it. Removing whole days instead would confound the influence of a keynote with the influence of its date and would discard the other firm that spoke. This gives 114 re-estimations on 112 trading days. The evaluation points are held at their full-sample values throughout; recomputing them inside each subsample would make the contrast a different quantity in every re-estimation and the 114 numbers would not be comparable with one another.

Figure 20: Equation (4) re-estimated with each keynote removed in turn



Notes: Each dot reports the t statistic from one of the 114 leave-one-keynote-out re-estimations, horizontally jittered within each maturity. Solid horizontal bars indicate the corresponding full-sample statistics. Dotted and dashed reference lines mark ±1.645 and ±1.96, respectively. The top row reports the fitted net responses, $\beta_m + \gamma_m^P C + \delta_m F$, evaluated at the delivered-heavy profile $(C, F) = (0.89, 0.06)$ and the promised-heavy profile $(0.68, 0.21)$, together with their difference. The bottom row reports the two interaction coefficients, γ_m^P and δ_m. Evaluation profiles are held fixed at their full-sample values throughout.

The interaction coefficients are similarly robust. The sign of γ_m^P is preserved in

every re-estimation at all maturities except two years, where the full-sample coefficient is close to zero (0.079) and statistically insignificant. From five years onward, γ_m^P remains significant at ten per cent in all 114 re-estimations and at five per cent in at least 113. The sign of δ_m is preserved across all maturities and all re-estimations. From two to nine years, δ_m remains significant at ten per cent in every case. At ten years, the full-sample t statistic is -1.89, and only five re-estimations reach five per cent significance, consistent with the full-sample result being significant at ten per cent rather than five per cent.

The influence of individual keynotes also differs across the two profiles. The largest changes in the delivered-capability response arise from excluding Computex on 5 June 2024 at longer maturities and Qualcomm's Uplinq on 1 June 2011 at intermediate maturities. For the promised-capability response, Hot Chips on 22 August 2022 is the most influential event across maturities. The latter keynote also produces the largest changes in the estimated interaction coefficient δ_m and the profile contrast. Nevertheless, its exclusion does not overturn the main finding: at seven years, the t statistic for the profile contrast falls from 4.40 to 3.56 but remains significant at one per cent.

10 Conclusions and remarks

Scheduled semiconductor conferences provide observable occasions on which firms disclose technological capability and commercial demand. Combining keynote transcripts with market returns allows us to study both the magnitude of repricing and the information associated with it. The central finding is that favourable technology news can have different inflation implications depending on when the disclosed capability becomes available. At five years, the delivered- and promised-capability profiles imply inflation-compensation responses that differ by 1.78 basis points per percentage point of sector return. Dividend-futures responses are positive under both profiles, with no statistically significant difference between them.

These patterns suggest a distinction between anticipated adoption demand and anticipated supply expansion. News about capability already available to firms may bring forward spending on adoption, complementary equipment, and infrastructure before the resulting productivity gains diffuse through the economy. Positive dividend responses, higher inflation expectations, and increased capital-goods orders at the delivered-capability profile are consistent with this anticipated-demand channel. News about promised capability may instead place greater weight on future productive capacity and lower production costs. At that profile, positive dividend responses accompany lower inflation expectations, followed by lower core prices and higher production, a pattern consistent with anticipated supply expansion. Capital-goods orders

increase under both profiles, suggesting that investment can respond to technological opportunities whether the capability is available or forthcoming.

The inflation contrast also remains present when the issuer and its peers reprice together. Market adjustment and intraday windows broadly preserve the baseline directions. The evidence therefore supports a content-dependent balance between expenditure and productive capacity, without identifying the two profiles as pure structural demand and supply shocks.

More broadly, the findings suggest that information composition can help explain why technology-news measures have different inflation implications. Technological progress can improve growth expectations while generating either upward or downward price pressure, depending on the spending and cost expectations associated with the disclosure. Applying comparable content measures to other settings would help establish how broadly these mechanisms extend beyond semiconductor conferences.

References

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy*, 40(121): 13–58.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. 2018. "Artificial Intelligence and Economic Growth." In *The Economics of Artificial Intelligence: An Agenda*. , ed. Ajay K. Agrawal, Joshua Gans and Avi Goldfarb, 237–282. Chicago:University of Chicago Press.

Ahmed, M Iqbal, and Quazi Fidia Farah. 2022. "On the macroeconomic effects of news about innovations of information technology." *Journal of Macroeconomics*, 71: 103389.

Aldasoro, Iñaki, Sebastian Doerr, Leonardo Gambacorta, and Daniel Rees. 2024. "The Impact of Artificial Intelligence on Output and Inflation." BIS Working Paper No. 1179.

Andrews, Isaiah, and Maryam Farboodi. 2025. "Do Markets Believe in Transformative AI?" NBER Working Paper No. 34243.

Barsky, Robert B., and Eric R. Sims. 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics*, 58(3): 273–289.

Barsky, Robert B., Susanto Basu, and Keyoung Lee. 2015. "Whither News Shocks?" *NBER Macroeconomics Annual*, 29(1): 225–264.

Bauer, Michael D., Diego R. Känzig, and Glenn D. Rudebusch. 2026. "Carbon Pricing and Inflation Expectations." Federal Reserve Bank of San Francisco, Working Paper.

Beaudry, Paul, and Franck Portier. 2006. "Stock Prices, News, and Economic Fluctuations." *American Economic Review*, 96(4): 1293–1307.

Blumenfeld, Alex, Jonathon Hazell, Chen Lian, and Andreas Schaab. 2026. "The Macroeconomic Effect of AI: Sizing the Software Engineering Channel." National Bureau of Economic Research NBER Working Paper 35793.

Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu. 2024. "Business News and Business Cycles." *Journal of Finance*, 79(5): 3105–3147.

Cascaldi-Garcia, Danilo, and Marija Vukotić. 2022. "Patent-Based News Shocks." *Review of Economics and Statistics*, 104(1): 51–66.

Comin, Diego A, Mark Gertler, and Ana Maria Santacreu. 2009. "Technology innovation and diffusion as sources of output and asset price fluctuations." National Bureau of Economic Research.

Drechsel, Thomas, and Ko Miura. 2025. "The Macroeconomic Effects of Bank Regulation: New Evidence from a High-Frequency Approach." University of Maryland Working Paper.

Drechsel, Thomas, and Naixin Zhang. 2026. "Do Judges Move Markets? Ideological Shifts in the Supreme Court and Macroeconomic Expectations." National Bureau of Economic Research.

Evgenidis, Anastasios, and Apostolos Fasianos. 2025. "AI News Shocks and the Macroeconomy: Evidence from UK Patent Data." Institute for Fiscal Studies Working Paper W25/48.

Faria-e Castro, Miguel, and Serdar Ozkan. 2026. "Can AI Optimism Raise Inflation? What a Standard Macro Model Says." Federal Reserve Bank of St. Louis, On the Economy Blog, March 26.

Gormsen, Niels Joachim, and Ralph S. J. Koijen. 2020. "Coronavirus: Impact on Stock Prices and Growth Expectations." *Review of Asset Pricing Studies*, 10(4): 574–597.

Gormsen, Niels J, Ralph SJ Koijen, and Ian WR Martin. 2021. "Implied dividend volatility and expected growth." Vol. 111, 361–365, American Economic Association, 2014 Broadway, Suite 305, Nashville, TN 37203.

Görtz, Christoph, Christopher Gunn, and Thomas A. Lubik. 2025. "What Can News Shocks Tell Us About the Effects of AI?" Federal Reserve Bank of Richmond, Economic Brief No. 25-16.

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

Hassan, Tarek Alexander, Stephan Hollander, Aakash Kalyani, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun. 2024. "Economic surveillance using corporate text." National Bureau of Economic Research.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2019. "Firm-Level Political Risk: Measurement and Effects." *Quarterly Journal of Economics*, 134(4): 2135–2202.

Hazell, Jonathon, and Stephan Hobler. 2024. "Do Deficits Cause Inflation? A High Frequency Narrative Approach." London School of Economics and Political Science Working Paper.

Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review*, 111(4): 1092–1125.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics*, 132(2): 665–712.

Kurmann, André, and Christopher Otrok. 2013. "News Shocks and the Slope of the Term Structure of Interest Rates." *American Economic Review*, 103(6): 2612–2632.

Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics*, 133(3): 1283–1330.

Ottonello, Pablo, and Thomas Winberry. 2020. "Financial Heterogeneity and the Investment Channel of Monetary Policy." *Econometrica*, 88(6): 2473–2502.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2012. "What's News in Business Cycles." *Econometrica*, 80(6): 2733–2764.

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

Trammell, Philip, and Anton Korinek. 2024. "Economic Growth under Transformative AI." National Bureau of Economic Research NBER Working Paper 32262.

Wilding, Tiffany, and J. R. Scott. 2026. "Does AI Raise or Lower Neutral Rates?" *PIMCO Macro Signposts*, Accessed 28 October 2026.

Online Appendix to

The macroeconomic effects of technology news: high-frequency evidence from semiconductor conferences

Naixin Zhang
George Washington University

October 2026

This appendix documents the event sample, punctuation restoration, and text classification, and provides supplementary analyses of content profiles, classification diagnostics, issuer–peer repricing, and Treasury yields.

Contents

A Construction and Validation of the Technology-Event Sample 55

B Punctuation Restoration for Automatic Captions 66

C Text Analysis and the Content Classifier 68

F Both Ends of Each Content Cut 77

G Is the Opposition of the Two Capability Halves Arithmetic? 78

H The Sign Split in the Cross-Section of Returns 79

I Additional Treasury-Yield Results 80

A Construction and Validation of the Technology-Event Sample

A.1 Text sources and measurement

Qualifying speech records come from three sources: an official verbatim transcript or prepared remarks, a human-provided caption track, and an automatic caption track. These sources record the keynote in the speaker's own words and provide a full speech-equivalent text. They are three archival representations of one economic object — the keynote as delivered — rather than three kinds of document, and the measures treat them as such. Prepared remarks are included on that basis and not on a claim of word-for-word identity with delivery: they are the speaker's own text of the speech, at full length, which is what the content measures need.

Presentation decks and speech records. The difference is visible in the archived records. Measured inside Hot Chips, where every keynote occupies the same sixty-minute slot and the same firms recur, **caption tracks have a median of 4,300 words and a presentation PDF of one runs 674**. A speaker covers roughly four thousand words in half an hour, so the captions are the speech and the deck is a sixth of it. Punctuation does not detect the difference — a deck scores 4.6 sentence terminators per hundred words against a stenographic transcript's 6.8, because bullet points end in full stops — which is why the test here is length against speaking rate and not punctuation. What the deck omits is also not random: it carries the conclusions and not the argument, and the demand- and supply-side reasoning that the content measure is built to detect is spoken rather than printed on the slide. Admitting decks would additionally make any word-denominated rate incomparable across years within a series, since the same thirty-minute keynote would contribute 4,300 words in one year and 674 in another.

An observation enters the content analysis only if EVERY keynote it carries has a speech record, not merely one of them. The reason is an identification one. The price response being explained is the whole day's disclosure, so the text has to account for all of it; a day whose AMD keynote is captioned and whose Intel keynote survives only as a deck would let the text explain part of the shock while the return reflects all of it, and the measured content would be a biased sample of the day's news. Of the 207 observations, **168 meet the condition and enter the content analysis**, with **1,904,810 words**; **39 do not** and are excluded from the text measures while remaining in the population and coverage counts, since they still establish that a keynote occurred and who gave it. As it happens the strict rule costs nothing relative to a weaker one requiring a single speech record: **no observation is partially covered**. That is a fact about this collection rather than a property of the design, and the condition is enforced rather than assumed. The exclusions are concentrated rather than scattered: all five Qualcomm BREW editions, documented only by the official show daily, which the

files themselves label "NOT a keynote transcript"; five early NVIDIA GTC editions documented by trade-press recaps; and the 2014 OCP summit, whose five constituent keynotes survive only as presentation PDFs.

Source differences and sentence segmentation. As delivered, stenographic transcripts and human-provided captions carry 6.8 and 6.5 sentence terminators per hundred words; automatic captions carry 0.2, because machine transcription emits almost no punctuation. This matters because the content measures are applied to statement-level units, so text with no sentence boundaries has nothing for a unit to attach to — the distinctions being measured are clause-level, capability against demand and forward-looking against realised. The gap is therefore closed before any classification runs, by the restoration procedure documented and validated in Appendix B: after it, the three sources carry 6.8, 6.5 and 6.7 terminators per hundred words, and the words themselves are provably unaltered. What remains is that one of the three had its boundaries placed by a model rather than a person, which is a measurement question and is treated as one there in the punctuation-validation exercise.

A.2 Predetermination and keynote timing

The event calendar consists of recurring conference programmes with scheduled keynote slots. Table 7 gives recurrence statements for fourteen of fifteen series; the remaining series have observed calendar histories but no located statement of this kind. Recurrence supports treating these events as regular disclosure opportunities, rather than occasions created for a particular announcement.

This evidence does not establish that the exact date of every edition was fixed a common number of months ahead, or that schedules were unaffected by product readiness. A dated advance announcement shows that dates were public by that time. The earliest document located is a lower bound on advance notice, not an estimate of the first public announcement. Table 8 reports selected documented examples on that basis.

Table 7: Recurrence evidence, series by series

Series	Statement that the series recurs	Source
Intel IDF	"Intel today announced the kick off of a semi-annual conference, the Intel Developer Forum ... being held in San Francisco from September 29 to October 1, 1997."	https://www.intel.com/pressroom/archive/releases/1997/CN91597B.HTM Intel pressroom archive, inaugural IDF announcement

Series	Statement that the series recurs	Source
CES	"Held in Las Vegas every year, it is the world's gathering place for all who thrive on the business of consumer technologies and where next-generation innovations are introduced to the marketplace."	https://www.cesweb.org/About-CES.aspx CTA/CEA, archived 2016
Computex	*no such statement located*	—
NVIDIA GTC	"NVIDIA had initially planned to host its 11th annual GTC on March 22-26 at the San Jose Convention Center, with an anticipated audience of 10,000."	https://nvidianews.nvidia.com/news/nvidia-gtc-news-to-be-shared-on-march-24-followed-by-investor-call
Hot Chips	"Since it started in 1989, HOT CHIPS has been known as one of the semiconductor industry's leading conferences on high-performance microprocessors and related integrated circuits. The conference is held once a year in August in the center of the world's capital of electronics activity, Silicon Valley."	https://www.hotchips.org/about/
OCP Global Summit	"The in-person event is held twice each year - in the Spring the OCP Regional Summit is held in Europe ...; while the OCP Global Summit is our largest, held in the Fall each year in San Jose, California."	https://www.opencompute.org/summit/global-summit archived 2023
Qualcomm Snapdragon Summit	"Qualcomm Technologies, Inc., a subsidiary of Qualcomm Incorporated (NASDAQ:QCOM), announced that the third annual Snapdragon Technology Summit keynotes will be broadcast live from Hawaii on December 4, 5 and 6 at 9:00 a.m. HST."	https://s204.q4cdn.com/645488518/files/doc_news/2018/11/2018-11-28_Qualcomm_Announces_Livestream_of_Keynotes_at_1223.pdf

Series	Statement that the series recurs	Source
Qualcomm BREW	"The third annual BREW 2003 Developers Conference is a leading educational and networking forum for developers of wireless software applications."	http://brew.qualcomm.com/brew/en/press_room/events/brew_2003/brew_2003.html Qualcomm BREW press room, archived
Arm TechCon / DevSummit	"Arm is announcing program changes and a new name for its annual flagship conference formerly known as Arm TechCon."	https://newsroom.arm.com/news/arm-hardware-and-software-ecosystems-will-converge-at-virtual-arm-devsummit-formerly-known-as-arm-techcon
Intel Innovation	"At its third annual Intel Innovation event, Intel unveiled an array of technologies to bring artificial intelligence everywhere and make it more accessible across all workloads, from client and edge to network and cloud."	https://www.intc.com/news-events/press-releases/detail/1648/intel-innovation-2023-empowering-developers-to-bring-ai
Intel Vision	"The annual event explores the latest in AI PC innovation with industry-leading security and manageability, and it introduces advanced data center, edge and networking solutions tailored to meet unique customer needs."	https://newsroom.intel.com/press-kit/intel-vision-2025 Intel newsroom press kit, via Internet Archive
AMD Fusion Dev. Summit	"The second annual AFDS will bring together top minds from the developer and computing communities as I focus on making heterogeneous computing accessible to mainstream developers and improving new end user experiences," said Manju Hegde, corporate vice president, Content, Applications and Solutions, AMD."	https://www.globenewswire.com/en/news-release/2012/04/12/1011747/0/en/AFDS-Returns-to-Pacific-Northwest-With-Focus-on-Heterogeneous-Computing-Gaming-Cloud-Computing-and-More.html

Series	Statement that the series recurs	Source
Mobile World Congress	"Each year, MWC Barcelona strengthens links between tens of thousands of senior leaders across the connectivity ecosystem, global business and government."	https://www.mwcbarcelona.com/about GSMA
Qualcomm Uplinq	"Qualcomm Life Unveils its 2net App SDK and Challenge During Annual Uplinq Conference"	https://www.qualcomm.com/news/releases/2012/06/27/qualcomm-life-unveils-its-2net-app-sdk-and-challenge-during-annual-uplinq Qualcomm newsroom, press-release headline
CeBIT	*no such statement located*	—

Notes: Each quotation is the organiser's or the host's own description of the series as a recurring event; sources are the organiser's press room, the organiser's site, or an Internet Archive capture of one, and the publisher is named rather than the archive. Where no such statement was located the row says so rather than substituting a weaker form of evidence. Full URLs and the capture dates are in the replication ledger `transcripts/_c1_declarations.csv`. Three rows carry a qualification. The Intel Developer Forum statement is the inaugural 1997 announcement, which calls the series *semi-annual*; that is the wording the data support, since IDF ran a spring and a fall edition in 2001, 2002, 2003 and 2006, and a later Intel release describing it as "annual" was not used for that reason. For Intel Vision the quoted sentence says "the annual event" without naming it, and is identified as Intel Vision only by the press kit it comes from. The Mobile World Congress statement describes MWC Barcelona and so does not cover the series' earlier 3GSM World Congress editions, for which no separate statement was located.

Advance notice. Table 8 uses the most recent edition in each series for which a dated advance-notice document was located. Fifteen series have such a record. Documented leads range from 44 to 357 days, with a median of 180. These examples establish advance scheduling for the series, rather than a common notice horizon for every estimation-sample keynote. A missing record does not establish short notice. The table distinguishes organizer publications, archived pages, secondary reports, and leads measured to the show's opening rather than a keynote.

Table 8: Advance notice: the most recent edition for which an announcement was located

Series	Edition	Days ahead	Announcement	Source
Intel IDF	2016	87	"IDF16 will be taking place in San Francisco, California, August 16-18."	http://web.archive.org/web/20160521172158/http://www.intel.com/content/www/us/en/intel-developer-forum-idf/san-francisco/2016/idf-2016-san-francisco.html website snapshot, published by the organiser
CES	2026	284	"CES will return to Las Vegas from January 6-9, 2026."	https://www.ces.tech/press-releases/ces-2025-audit-reveals-growing-attendance-from-executives-investors-and-media press release, published by the organiser
Computex	2026	355	"COMPUTEX to 06 05 – Taipei Nangang Exhibition Center, Hall 1 (TaiNEX 1), Hall 2 (TaiNEX 2), TWTC, TICC."	http://web.archive.org/web/20250611055139/https://www.computextaipei.com.tw/en/index.html website snapshot, published by the organiser
NVIDIA GTC	2026	162	"The GTC AI conference returns to San Jose, CA on March 16-19, 2026. Get notified when registration opens."	https://web.archive.org/web/20251005154019/https://www.nvidia.com/gtc/ website snapshot, published by the organiser
Hot Chips	2024	152	"Important Dates ... Conference Dates: August 25-27, 2024."	http://web.archive.org/web/20240328115249/https://hotchips.org/ website snapshot, published by the organiser

Series	Edition	Days ahead	Announcement	Source
OCP Global Summit	2025	220	“2025 OCP Global Summit – San Jose, California – October 13-16, 2025.”	http://web.archive.org/web/20250307201626/https://www.opencompute.org/summit/global-summit website snapshot, published by the organiser *the January captures of this page read October 14–16; the keynote day was added between 26 January and 7 March, and the lead is measured from the first document that covers it*
Qualcomm Snapdragon Summit	2026 (11th)	84	“Snapdragon Summit is September 22 - 24, 2026 in Maui, Hawaii.”	http://web.archive.org/web/20260630032434/https://www.qualcomm.com/company/events/snapdragon-summit website snapshot, published by the organiser *the page still showed the 2025 dates on 10 June, so the dates went up between then and 30 June and the lead is between 84 and 104 days*
Qualcomm BREW	2007 (7th)	198	“Winners will be revealed at an awards ceremony held at the BREW 2007 Developers Conference in San Diego June 20 - 22, 2007.”	http://web.archive.org/web/20061205093236/http://brew.qualcomm.com:80/brew/brew_2007/awards/awards.html website snapshot, published by the organiser

Series	Edition	Days ahead	Announcement	Source
Arm TechCon / DevSummit	2020	110	"Arm DevSummit 2020, October 6-8, Virtual Conference"	https://web.archive.org/web/20200618140507/https://devsummit.arm.com/ website snapshot, published by the organiser
Intel Innovation	2023	106	"Intel Innovation is scheduled for September 19-20th."	https://www.kitguru.net/channel/generaltech/joao-silva/intel-schedules-innovation-2023-event-for-september-19-20/ third-party report *medium confidence; organiser-issued record for this series: 2023, 11 days*
Intel Vision	2025	80	"Intel Vision 2025 event which is now slated to start on March 31."	https://wccftech.com/intel-vision-2025-event-scheduled-for-march-31/ third-party report *medium confidence; attributes the dates to Intel's own events page, which is not archived before this date*
AMD Fusion Dev. Summit	2013	201	"'APU13' will be held Nov. 11-14, 2013 in San Jose, California at the McEnery Convention Center."	https://ir.amd.com/news-events/press-releases/detail/1466/amd-announces-2013-developer-summit-apu13-set-for-san-jose-california press release, published by the organiser

Series	Edition	Days ahead	Announcement	Source
Mobile World Congress	2026	357	"See you next year, 2 - 5 March 2026, Fira Gran Via, Barcelona."	http://web.archive.org/web/20250310065741/https://www.mwcbarcelona.com/ website snapshot, published by the organiser *measured to the first day of the show: this project holds no keynote date for this edition, whose keynotes are not in the sample*
Qualcomm Uplinq	2012	245	"the Uplinq 2012 mobile developers conference will be held June 27-28, 2012 at the Hilton San Diego Bayfront hotel."	https://www.qualcomm.com/news/releases/2011/10/uplinq-2012-mobile-developers-conference-dates-location-set press release, published by the organiser
CeBIT*	—	—	*no dated announcement located*	—

Notes: Days ahead is the interval between the dated announcement and the keynote. *Source* names the document and the kind of document it is. Where an edition is documented by more than one such record the table reports the *earliest* of them, because a lead is a lower bound and the earlier document is the better bound. Intel Innovation is the case that settles the rule: two documents state the same dates, one of them three months ahead of the event and one eleven days out, and reporting the later would have credited the series with eleven days' notice for an edition already documented as public in June. * **A dash never means short notice.** It means no dated document could be located. The one starred series is CeBIT, which ceased in 2018 and whose organiser's site is archived too sparsely in the relevant years for a document carrying the dates to be recovered. The Mobile World Congress row is the one place where the lead is measured to the show's opening day rather than to a keynote: this project holds keynote dates for that series only for its 3GSM World Congress editions of 2001–2004, so the row documents advance scheduling for the series without being a lead for an observation in the estimation sample, and it says so. One row gives a window rather than a point, because the organiser's page is captured on both sides of the change: the Snapdragon Summit page brackets to between 84 and 104 days. One row carries a further qualification of its own: the dates of the 2025 OCP Global Summit were given as 14–16 October in January and as 13–16 October from March, and because the keynote in this sample falls on 13 October the lead is measured from the first document that covers that day rather than from the earlier one, which would credit the organiser with notice of a day it had not yet scheduled.

The table also flags records whose dating rests on a secondary report rather than on an archived

organiser publication. **Two rows do**: Intel Innovation 2023, which the report documents 106 days ahead, and Intel Vision 2025 at 80 days. In both cases the report cites the organiser's own event page, which is not archived early enough to be cited directly. A distinction is worth naming because the ledger turns on it: a release that announces the keynote *lineup* dates the speakers, not the event, and only a document stating the dates themselves is counted as advance notice of the occasion. **For the one starred series the only timing evidence is the recurrence statement above and the observed calendar history.**

The keynote hour, counted on sessions and not on days. An intraday design works on the keynote session, so that is the unit in which timing coverage has to be reported. Of **at least 214** reconstructed keynote sessions, **190** have an identified start, implying timing coverage of **no more than 89%**; **110** have an identified end, implying coverage of no more than 51%. The corresponding venue-day figure — 150 of 161 observations, 93% — is *not* the same statistic and should not be read as one: a venue-day counts as timed if *one* of the keynotes it pools has an hour, so a day running four firms' keynotes with one known slot contributes a timed observation and three untimed sessions. Reported properly the coverage is four points lower, and **the two measures move in opposite directions across the sample**: on days, the later period looks better (99% against 89%); on sessions, it looks worse (88% against 90%), because the later period is where the multi-firm formats sit — the OCP summit runs four or five constituent firms through one afternoon, and the organiser's programme names the keynote slots it timed, not every firm that spoke. Eleven venue-days carry no hour at all, each recorded with the route that failed. Session coverage is complete for six of the fourteen series that supply an observation, and 48 of Intel's 50 Developer Forum observations are timed.

A.3 Sample funnel, missingness and limitations

Table 9 is the whole funnel in one place, from the population frame to the days the regressions are estimated on. Three features of it decide how it should be read, and each is taken up below: documented zeros differ from missing programmes, the unit changes from editions to keynote-days, and the final restrictions define the estimation sample.

Table 9: From the population frame to the estimation sample

223	**editions**	Fifteen series, 2001–2026. The frame, not the sample.
−31		no index constituent on a full keynote programme (documented zeros)
−21		undetermined, almost all before 2008 (missing data)
171	**editions**	at which a constituent firm keynoted
The unit changes: an edition runs several days and a firm may keynote on more than one.		
175	**keynote-days**	produced by those 171 editions
−17		the material is not a verbatim record of a keynote: five conference show dailies, four journalists' write-ups, a slide deck, seven fragments
−6		a verbatim record, but not of *that* keynote: a watch party, a later year's keynote, two other occasions, an academic talk, one duplicate
152	**keynote-days**	with an independent qualifying verbatim record
−27		before 21 July 2004, when the swap curve is first quoted at all ten maturities; these have a price and a transcript and are used by nothing
125	**keynote-days**	inside the study window
−7		an FOMC statement day or a CPI release day
−3		an opening keynote on a United States long weekend: no price that day
−1		the speaking firm reported earnings inside the pricing window
114	**keynotes**	
112	**trading days**	twelve series, eighty-five series-years — **the estimation sample**

Notes: Counts above the first midrule are editions, defined as one occurrence of a series in one calendar year; below it they are keynote-days and then trading days, and the three must never be added together. Every edition in the first two deductions carries a recorded search path and outcome in `population_and_evidence/keynote_registry.csv`, and every entry in the −6 row carries its evidence in `population_and_evidence/transcript_defects_2026_10_26.md`.

The registry covers 223 editions: 171 with constituent-firm keynotes, 31 with no qualifying speaker on the recovered keynote programme, and 21 with unrecovered programmes. The latter are missing observations, whereas the documented zeros establish that the conference occurred without a qualifying speaker. Inclusion in the registry is determined independently of transcript availability. Text coverage is strongest for company-hosted developer conferences and Hot Chips, and weaker for large third-party shows in the 2000s, so the content analysis draws disproportionately on company-hosted and more recent events. A final audit applying the keynote criterion removed 15 observations and corrected 21 dates. Nine removals involved firm press conferences or product briefings mistaken for organiser-scheduled keynotes; the date corrections replaced conference opening dates with the actual keynote dates. The replication files document these corrections and the search routes used for each edition. Although no external historical database establishes complete coverage, the registry provides a documented sample of scheduled public technology disclosures by US-listed semiconductor firms at recurring conferences. Its concentration in computing and communications firms limits generalisation to equipment, materials, and EDA firms, as well as to other disclosure formats such as earnings calls, standalone launches, regulatory filings, and academic conference presentations.

B Punctuation Restoration for Automatic Captions

Table 10: Punctuation restoration against human-punctuated text

	Official transcripts	Human-provided captions	Pooled
Documents	69	23	92
Words	401,754	128,593	530,347
Period			
Precision	0.749	0.789	0.764
Recall	0.729	0.732	0.732
F1	0.743	0.738	0.740
Question-mark F1	0.583	0.720	0.609
End-of-statement			
Precision	0.758	0.797	0.759
Recall	0.724	0.729	0.728
F1	**0.742**	**0.740**	**0.741**
Random baseline	0.061	0.056	0.060
Periodic baseline	0.061	0.060	0.061
Mean sentence length, original	16.1	16.1	16.1
Mean sentence length, restored	16.9	16.6	16.9

Notes: Each cell is the median across documents. A document enters if it has at least 800 words and if removing punctuation leaves its word count unchanged; it is truncated at 6,000 words so that documents are comparable. *End-of-statement* pools periods, question marks and exclamation marks — the positions that define the unit of classification. *Random baseline* inserts the correct *number* of sentence ends at random positions, and *periodic baseline* inserts them at the correct average *frequency*; both are shown as reference levels for boundary prediction without syntax.

Automatic captions require punctuation restoration before sentence segmentation: they contain only 0.2 sentence terminators per hundred words, compared with 6.8 in official transcripts and 6.5 in human captions. I use a Transformer-based token classifier built on XLM-RoBERTa-large, starting from `oliverguhr/fullstop-punctuation-multilang-large` (revision `345e80ad`). It predicts six labels after each word: no mark, period, comma, question mark, dash, and colon. Processing uses overlapping 230-word windows advancing by 150 words, discarding outer predictions to reduce boundary effects, with sentence-initial capitalization. Official transcripts and human captions retain their existing punctuation.

To validate boundary detection, I remove punctuation from human-punctuated records, restore it, and compare the predicted boundaries with the originals. Table 10 reports results for the unadapted checkpoint on 69 official transcripts and 23 human-captioned records. The primary metric pools sentence-ending punctuation, which defines the downstream classification units. Median boundary F1 is 0.741, compared with approximately 0.06 for random and periodic baselines. These results evaluate the unadapted model only, since some records subsequently enter adaptation.

I then evaluate domain adaptation using the twenty-three human-captioned keynotes, split by whole keynote into fifteen for training, four for stopping-epoch selection, and four for evaluation. Fine-tuning uses a learning rate of 10^{-5} and freezes the lower half of the encoder; validation F1 selects the third epoch. On the four held-out keynotes, adaptation raises median boundary F1 from 0.781 to 0.844, mainly through higher recall (Table 11). The adapted checkpoint is used throughout the paper.

Table 11: Sentence-boundary accuracy before and after domain adaptation

	Held-out keynotes (4)		
Model	Precision	Recall	F1
Published punctuation checkpoint	0.840	0.752	0.781
+ semiconductor-keynote adaptation	0.838	0.847	**0.844**
Gain	−0.002	+0.095	+0.063
Model used in the paper	the adapted checkpoint		

Notes: End-of-statement precision, recall and F1, median across the four held-out keynotes — the ones used for neither adaptation nor the choice of stopping epoch. *Published punctuation checkpoint* is the model of Section B.1 applied without any adaptation; *adapted* is the same checkpoint after light fine-tuning on the fifteen adaptation keynotes. Both rows are scored on identical documents.

The adoption rule required a held-out F1 improvement of at least 0.02, a threshold set after baseline scoring but before adapted-model scores were obtained. The four held-out keynotes therefore inform model adoption and constitute a small evaluation sample rather than an independent final test of the selected pipeline. Human punctuation also permits alternative valid boundaries, so disagreement need not imply a linguistic error.

The adapted model produces 6.7 sentence terminators per hundred words, compared with 5.4 for the unadapted model and 6.5–6.8 for human-punctuated sources. This density comparison is descriptive rather than boundary-level validation. A separate check strips punctuation and case and verifies that restoration preserves the normalized word sequence, aborting on any token difference. All 77 stored files, combining 116 automatic-caption source documents, pass this check across 737,084 normalized words. The difference from the 737,141 source-word count consists of 57 punctuation-only tokens. This establishes word preservation, not transcription or boundary accuracy; the effect of segmentation errors on event-level content measures remains unquantified.

C Text Analysis and the Content Classifier

C.1 What the text produces on its own, before any classifier

Before asking a language model to label anything, I ask what structure the corpus has with no labels at all. I apply principal component analysis to the news-sentence corpus used in Figure 1. The news screen uses the classifier; component extraction is unsupervised conditional on that screen. Unlike supervised approaches that depend on predefined labels, PCA recovers the main axes of variation directly from the text. Three of the four components shown in Figure 1 are readily named: data-centre computing (for example, *data center*, *cloud*, *storage*, *training*), performance and efficiency (*per second*, *efficiency*, *high performance*, *bandwidth*), and artificial intelligence (*models*, *generative*, *nvidia*). They are dimensions of *subject matter*: what a keynote is talking about. The fourth is a weaker contrast between component-level vocabulary (*processor*, *core*) and system-level vocabulary (*network*, *devices*), overlapping the first three, and I do not rest anything on it.

The technology/demand distinction does not appear among the main components. This is not a failure of the method but a property of the object. Two sentences about the same accelerator, one announcing how fast it is and one announcing that a customer has bought it, use nearly the same words and are economically opposite; the reported leading word-frequency components do not directly separate those economic roles. That axis is therefore theory-driven and rests on annotation, which is what the rest of this appendix documents. The AI dimension, by contrast, the text does volunteer.

C.2 The classifier

Economic classification. The leading textual components describe subject matter, while the regressions require economic role and horizon. A performance announcement and a customer deployment can concern the same accelerator but convey different information. I use explicit annotation to distinguish these statements and evaluate the resulting labels against human annotation below.

What the classifier is asked. Each sentence is labelled on six fields. Three are screening questions about the text itself — whether it identifies something specific, whether it states a fact or evaluates, and whether it is a consumer-product feature or a passing remark about someone else's product. Three are about content: `layer`, where in the technology stack the statement sits (manufacturing, chip, platform, application, cross-cut); `sublayer`, which of eighteen kinds within that layer (process node, architecture, performance, model, software tooling, market size, and so on, each permitted under exactly one layer); and `role`, whether the statement is about capability, about demand,

or about neither. A fourth content field, `horizon`, records whether what is claimed has already happened. A second pass asks, lexically, whether the sentence ties its subject to artificial intelligence.

Three of these fields are not used by the main results, which need only `role` and `horizon`. The layer and sublayer are collected for two reasons: they let the content variable be cut by what kind of technology a keynote is about, and a sublayer that cannot occur under the layer the model chose is a detectable error, which a scheme with only broad categories does not provide. The AI flag is collected because whether AI content specifically moves expectations is the obvious question to ask of these events; it is in the replication package for anyone who wants to ask it. No result in this paper uses it.

Derived news indicator. The informational screen is computed from the model's factual and exclusion labels:

$$\texttt{is_news} = \mathbb{1}\{\texttt{claim_type} = \text{fact}\}\mathbb{1}\{\texttt{excluded} = \text{none}\}.$$

The rule identifies qualifying statements without asking the model whether investors found them surprising. The separate AI pass likewise uses explicit lexical cues rather than a judgment about whether a product is associated with AI.

Two passes, and why. The content fields and the AI field are asked in separate calls. The AI question is about the words in the sentence and nothing else, so it separates cleanly, and asking it on its own costs a tenth as much as asking it alongside the full taxonomy. The two passes are joined on the sentence identifier and the join is complete: all 98,440 sentences carry both.

The prompt. The following excerpt reproduces the content prompt, with elisions marked. The omitted material includes technology-stack categories, worked examples, and the output schema. The complete prompt is stored as `1_prompt_v2tax_c89b69551fb2.txt`; its hash is recorded with the labels. The separate AI-cue prompt is `1_prompt_aicue_072aa92db773.txt`. The role and horizon definitions used in the paper appear in the excerpt below.

System message, abridged:

```
You are a senior semiconductor-industry analyst. You will be given statements taken
verbatim from AI/semiconductor conference keynotes. Label each one on the dimensions below.
Judge every statement on its own; do not let neighbouring statements influence it.

[...]

1) SCREENING -- three questions about the TEXT ITSELF. Answer each one literally.
```

```
   DO NOT judge whether the statement is important, interesting, newsworthy or substantive
   ↪  enough.
   That judgement is not yours to make here: a one-line spec off a slide counts exactly as much
   ↪  as a
   product launch. Answer only what is asked.

   has_specific : yes/no -- does the text contain AT LEAST ONE of: a number or quantity, a
   ↪  product or
        platform name, a company or customer name, a process node, a date or year?
        Count it even if that is the whole text.

[...]

   claim_type   : fact | commentary | none -- WHAT KIND of statement is it? This is the field
   ↪  that
        decides everything, so read the three definitions carefully.

        `fact`       : it states an attribute, capability, state, event, quantity or date OF A
                       CONCRETE THING -- a chip, a node, a package, a piece of software, a
                       ↪  model, a
                       customer, a fab, a market, a transaction. Something a reader could in
                       principle check. Being brief, unquantified or fragmentary does not stop
                       ↪  it
                       from being a fact.
             "x86 performance with Arm-competitive energy efficiency"       -> fact
             "Abnormal traffic detection"                                    -> fact (a
             ↪  capability)
             "Uses Integrated Fan-Out technology"                            -> fact
             "Extension of N4P technology"                                   -> fact
             "From any cloud or on-prem you can monitor your fleet"          -> fact (a
             ↪  capability)
             "The amazing amount of high speed memory lets the GPU work on
              very large data sets"                                          -> fact (an
              ↪  attribute)
             "New dedicated fab announced for 3DFabric assembly and test"    -> fact (an event)
             "I expect Eos to be the fastest AI computer in the world"      -> fact (an
             ↪  expectation
                                                                                about a named
                                                                                ↪  system)
             "Only foundry capable of massive 7nm volume production"          -> fact (a
             ↪  checkable
                                                                                claim about
                                                                                ↪  capability)

        `commentary` : it evaluates, interprets or frames rather than states -- how important,
                       aggressive, transformative or well-positioned something is; where the
                       ↪  industry
                       is heading in general terms; what the speaker intends or believes; an
                       explanation of a concept. Nothing here is checkable against a fact.
             "Neoverse represented Arm's most aggressive move into Intel's
              traditional territory"                                        -> commentary
             "Agents agents is a new workload."                              -> commentary

[...]

   excluded     : none | consumer_feature | third_party
        Two categories are excluded BY RESEARCH DESIGN, regardless of how specific they are.
```

```
        consumer_feature : an on-device capability offered to the end user, EVEN IF performed
        ↪  by AI
             -- camera, photo/video editing, display, audio, gaming, connectivity features.
             ↪  What a
             phone does for its owner updates no belief about how strong or cheap compute is.
             BUT a number describing the silicon or the model is NOT excluded:
                  "using AI to remove unwanted objects from full resolution videos" ->
                  ↪  consumer_feature
                  "industry first to stream lossless CD bluetooth audio"            ->
                  ↪  consumer_feature
                  "45 TOPS of AI performance"                                       -> none
                  "the NPU is 4.5x faster than the previous generation"             -> none
        third_party : a third party's own situation with no transaction with the speaker
                  "GlobalFoundries still pursuing 7nm"                              ->
                  ↪  third_party

   is_news is then DERIVED, not asked: news iff claim_type = fact and excluded = none.
   Do not output is_news. has_specific is recorded for diagnostics only and does not enter the
   ↪  rule,
   so answer it honestly rather than to steer the outcome. If the derived value is not_news,
   layer/sublayer/role/horizon are ignored, so you may still fill them; they will be discarded.

   GRAMMAR IS NOT A CRITERION. These statements come from slides, spec tables and bullet lists
   ↪  as
   well as from speech, so many are fragments with no subject or verb. Fragments are usually
   ↪  facts:

[...]

3) role -- WHAT DOES IT LET A LISTENER UPDATE? (economics; this is the dimension that matters)
   Ask: after hearing this, do I revise my belief about HOW STRONG/CHEAP COMPUTE CAN GET, or
   ↪  about
   HOW MUCH OF IT PEOPLE WILL BUY?
   tech    : capability itself -- a new architecture, a faster node, a benchmark, a new
   ↪  library,
             a roadmap to a future generation. The speaker is saying "this is what I can now
             ↪  do".
   demand  : who wants it and how much -- a named customer deploying or running it, orders,
             shipments, revenue, customers' capital expenditure, market-size or growth
             ↪  forecasts,
             AND capacity expansion (building fabs and adding wafer capacity is a response to
             demand, so it updates beliefs about demand, not about capability).
   context : neither -- government policy and export controls, the speaker's strategy or
             positioning, industry background, competitive commentary.

[...]

   Rules that decide most hard cases:
     * A named customer DOING something with the technology is `demand`, whatever the layer --
       even when the sentence is phrased as a partnership.
     * Ecosystem and partner counts are `demand` only when they measure uptake (developer
     ↪  counts,
       customers shipped to). A partnership announced with no deployment behind it is
       ↪  `context`.
     * A NODE OR PRODUCT REACHING PRODUCTION IS `tech`, NOT `demand`. "N3 enters high-volume
       manufacturing in H2 2022", "5nm high-volume production: 2020", "risk production Q1 2020"
       ↪  say
```

```
      that the technology is ready and usable -- they update capability, so they are tech.
      ↪  Only a
      CAPACITY quantity is demand: wafers per month, a fab being built, a capex figure, a
      ↪  stated
      expansion ("CoWoS capacity doubles to 70k wafers/month"). Reaching production = tech;
      how much of it there is = demand.
    * Making a technology AVAILABLE to customers or partners is `tech` (it is the capability
    ↪  being
      released): "Holoscan is generally available in May", "I will make NVLink available to
      partners to build companion chips". It becomes demand only once someone is named as
      ↪  taking it.

[...]

4) horizon -- present | future | na
   The question is whether the fact is ALREADY REALISED or NOT YET, because only the latter
   revises expectations about the future.
   future  : explicitly not yet realised -- "will be", previewed, planned, a dated roadmap, a
             node or product scheduled for a later year ("N2 enters production in 2026";
             "in 2028 I will introduce Florence").
   present : everything else that has already happened or already exists. This includes facts
             with no explicit time marker: prices, cumulative shipments, revenue shares, the
             specifications of an existing product ("Price: $69,000"; "150+ billion Arm-based
             chips shipped to date"; "North America = 71% of TSMC revenue"; "TIV pitch:
             ↪  300um").
             Absence of a time word is NOT a reason to choose na.
   na      : reserve for the rare case where the statement is neither realised nor tied to any
             future point.
```

User message: a JSON array of the sentences to label, each as `{"id": <id>, "text": <the sentence>}`, with the preceding and following sentence given as context in the same object. Temperature is zero.

Validity of the sublayer field. The taxonomy permits each sublayer under exactly one layer, so a response pairing a sublayer with the wrong layer is detectably wrong without any reference to a human label. Across the corpus, 355 sentences of 98,440 — 0.36 per cent — come back that way. I keep the layer, which is the coarser and better-identified of the two, and void the sublayer rather than discard the sentence. This is a lower bound on the error rate in the sublayer field and says nothing about sentences where the pairing is internally consistent but wrong.

Models compared. I need the model whose labels track a human reading of the same sentences. I compare five OpenAI models on the 400 hand-annotated sentences under an identical prompt: `gpt-4.1-mini`, `gpt-4o`, `gpt-5-mini`, `gpt-4o-mini` and `gpt-4.1-nano`. Table 12 reports accuracy and F1 against the human labels for each of the four annotated dimensions. The quantity I select on is neither of those but the bias in the corpus share — the model's share of each category minus the annotator's, both

reweighted by label-cell population — and it is given in the text below, because every regression in the paper uses the shares rather than the sentence labels.

Table 12: Agreement with the human annotation

	is news		capability		horizon		summary	
Model	Acc.	F1	Acc.	F1	Acc.	F1	$\overline{\text{F1}}$	$\sum$ |bias|
gpt-4.1-mini	0.75	0.793	0.83	0.893	0.87	0.924	0.870	**0.182**
gpt-4o	0.77	0.804	0.86	0.907	0.86	0.914	0.875	0.339
gpt-5-mini	0.76	0.808	0.86	0.915	0.88	0.929	**0.884**	0.204
gpt-4o-mini	0.74	0.782	0.82	0.881	0.77	0.853	0.839	0.279
gpt-4.1-nano	0.61	0.672	0.73	0.820	0.60	0.710	0.734	0.487

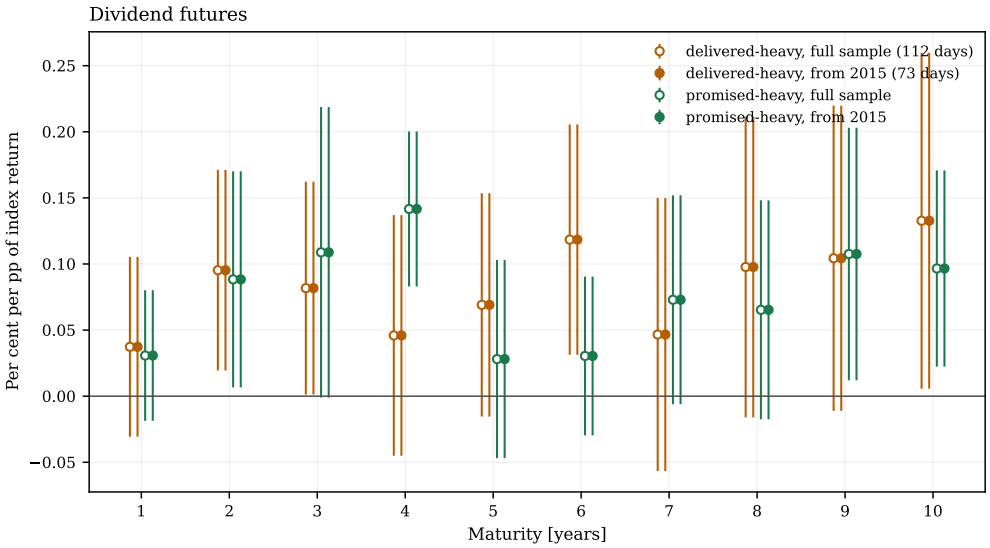
Notes: The author annotated 400 sentences sampled equally from five label cells and scored with population reweighting. Accuracy and F1 use the reported majority category in each dimension; $\overline{\text{F1}}$ averages the three F1 scores. $\sum$ |bias| is the sum of absolute discrepancies between model and human category shares. Role and horizon scores condition on both model and annotator assigning news, leaving 194, 193, 200, 190, and 159 observations by row. Prompt development used the annotated sentences, so these are in-sample agreement measures, not held-out accuracy. Layer and sublayer are not validated against human labels in this table.

Model selection. I use `gpt-4.1-mini`, which has the smallest total absolute share discrepancy in Table 12: 0.182, compared with 0.204 for `gpt-5-mini`. Its mean F1 is 0.870, below 0.884 for `gpt-5-mini` and 0.875 for `gpt-4o`. I emphasize share discrepancies because the regressions use content shares. However, low discrepancy at the pooled-corpus level need not imply low error in each event's share: offsetting errors across keynotes can still bias the estimated content interactions. The comparison does not quantify that error.

Human annotation. The author annotated 400 sentences sampled in equal numbers from five label cells and scored with population reweighting. The annotation sheet provided adjacent sentences but omitted event names, dates, and model labels. Twenty unmarked duplicates were included to assess within-annotator consistency. Agreement was 90 per cent for the news screen and 83.3 per cent for economic role; agreement on horizon was complete among the twelve duplicated news observations. These are small-sample repeatability measures, not ceilings on model accuracy or evidence of inter-annotator agreement.

D Post-2015 Subsample

Figure 21: Dividend-futures responses: post-2015 and full samples



Notes: As Figure 7, for dividend-futures responses in per cent per percentage point of sector return. Dividend-futures data begin in November 2015, so the calendar restriction does not remove observations from their baseline estimation sample.

Table 13: Equation (4) on the full sample and from 2015

	Maturity (years)					
	1	2	3	5	7	10
Panel A. Full sample (112 keynote-days, basis points per percentage point)						
Surprise z_d	2.103	0.921	-0.164	-1.934*	-2.744***	-2.409***
	(1.723)	(1.554)	(1.165)	(1.095)	(0.696)	(0.728)
$\times Z_d^C$ (realised, against demand: γ_m^P)	-1.688	0.079	1.004	3.628**	4.546***	3.893***
	(2.291)	(2.096)	(1.694)	(1.627)	(1.042)	(1.086)
$\times Z_d^F$ (promised less realised: δ_m)	-3.341	-7.409**	-5.591**	-6.898***	-5.739***	-4.040*
	(2.505)	(3.114)	(2.582)	(2.344)	(1.883)	(2.133)
Capability share C_d	-0.516	-1.982*	-1.905**	-1.487*	-0.970	-0.893
	(1.275)	(1.080)	(0.945)	(0.847)	(0.716)	(0.692)
Forward share F_d	1.713	13.259**	14.357***	9.172**	5.667	5.936
	(7.229)	(6.100)	(5.334)	(4.649)	(4.460)	(4.116)
Derived: the net response $\beta_m + \gamma_m^P C + \delta_m F$ at each evaluation point, and their difference.						
At a realised-heavy keynote ($C = 0.89$, $F = 0.06$)	0.392	0.532	0.383	0.870**	0.950***	0.809***
	(0.456)	(0.436)	(0.420)	(0.409)	(0.267)	(0.269)
After a forward-heavy keynote ($C = 0.68$, $F = 0.21$)	0.259	-0.571	-0.648	-0.907**	-0.852***	-0.606*
	(0.366)	(0.472)	(0.427)	(0.360)	(0.308)	(0.337)
Difference (realised less forward)	0.133	1.103*	1.032*	1.778***	1.802***	1.415***
	(0.640)	(0.627)	(0.528)	(0.522)	(0.410)	(0.442)
Observations	5,523	5,439	5,589	5,557	5,603	5,657
Keynote days	104	101	105	107	106	106
Panel B. From 2015 (73 keynote-days, basis points per percentage point)						
Surprise z_d	-1.452	-1.994	-1.130	-2.063	-1.685	-1.741
	(2.550)	(2.204)	(1.636)	(1.336)	(1.192)	(1.188)
$\times Z_d^C$ (realised, against demand: γ_m^P)	3.336	3.789	2.558	3.742**	3.166**	3.043*
	(3.345)	(2.946)	(2.220)	(1.830)	(1.595)	(1.597)
$\times Z_d^F$ (promised less realised: δ_m)	-3.141	-4.094	-3.909*	-4.463**	-4.759***	-2.967*
	(2.225)	(2.542)	(2.288)	(1.992)	(1.829)	(1.775)
Capability share C_d	-1.778	-2.514*	-2.103*	-1.502	-1.204	-0.977
	(1.377)	(1.331)	(1.116)	(0.942)	(0.845)	(0.819)
Forward share F_d	8.107	14.449**	15.222***	9.983**	7.622	7.314
	(6.597)	(6.407)	(5.887)	(4.953)	(4.964)	(4.552)
Derived: the net response $\beta_m + \gamma_m^P C + \delta_m F$ at each evaluation point, and their difference.						
At a realised-heavy keynote ($C = 0.89$, $F = 0.06$)	1.325**	1.128**	0.906**	0.994**	0.840***	0.786**
	(0.558)	(0.539)	(0.445)	(0.397)	(0.324)	(0.321)
After a forward-heavy keynote ($C = 0.68$, $F = 0.21$)	0.160	-0.273	-0.207	-0.451	-0.526*	-0.292
	(0.373)	(0.425)	(0.386)	(0.301)	(0.299)	(0.281)
Difference (realised less forward)	1.165	1.401*	1.114*	1.445***	1.367***	1.078**
	(0.786)	(0.732)	(0.570)	(0.468)	(0.416)	(0.421)
Observations	2,849	2,817	2,943	2,910	2,961	2,969
Keynote days	66	63	67	68	68	68

E Ordinary-Return Check

Section 9.3 reports the content interaction coefficients; this appendix reports every coefficient behind them, including the fitted net responses.

Table 14: Equation (4) with and without the ordinary sector return

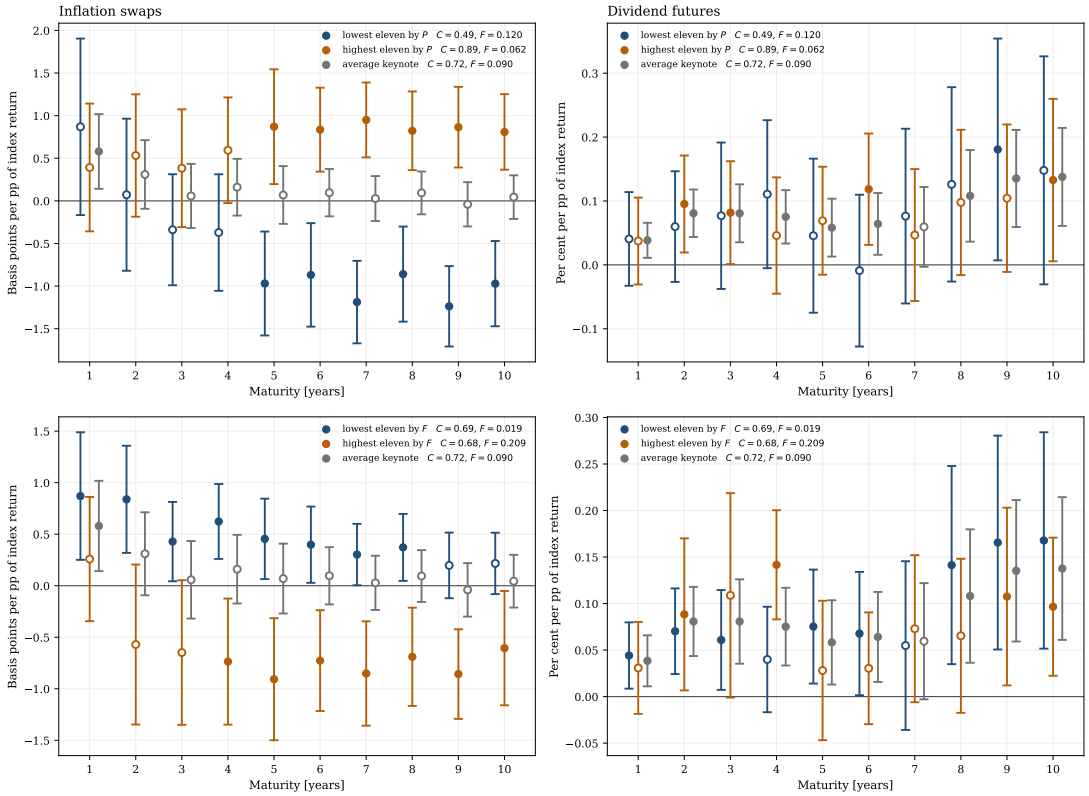
	Maturity (years)					
	1	2	3	5	7	10
Panel A. Equation (4) (basis points per percentage point)						
Surprise z_d	2.103	0.921	-0.164	-1.934*	-2.744***	-2.409***
	(1.723)	(1.554)	(1.165)	(1.095)	(0.696)	(0.728)
$\times Z_d^C$ (realised, against demand: γ_m^P)	-1.688	0.079	1.004	3.628**	4.546***	3.893***
	(2.291)	(2.096)	(1.694)	(1.627)	(1.042)	(1.086)
$\times Z_d^F$ (promised less realised: δ_m)	-3.341	-7.409**	-5.591**	-6.898***	-5.739***	-4.040*
	(2.505)	(3.114)	(2.582)	(2.344)	(1.883)	(2.133)
Capability share C_d	-0.516	-1.982*	-1.905**	-1.487*	-0.970	-0.893
	(1.275)	(1.080)	(0.945)	(0.847)	(0.716)	(0.692)
Forward share F_d	1.713	13.259**	14.357***	9.172**	5.667	5.936
	(7.229)	(6.100)	(5.334)	(4.649)	(4.460)	(4.116)
Derived: the net response $\beta_m + \gamma_m^P C + \delta_m F$ at each evaluation point, and their difference.						
At a realised-heavy keynote $(C = 0.89, F = 0.06)$	0.392	0.532	0.383	0.870**	0.950***	0.809***
	(0.456)	(0.436)	(0.420)	(0.409)	(0.267)	(0.269)
After a forward-heavy keynote $(C = 0.68, F = 0.21)$	0.259	-0.571	-0.648	-0.907**	-0.852***	-0.606*
	(0.366)	(0.472)	(0.427)	(0.360)	(0.308)	(0.337)
Difference (realised less forward)	0.133	1.103*	1.032*	1.778***	1.802***	1.415***
	(0.640)	(0.627)	(0.528)	(0.522)	(0.410)	(0.442)
Observations	5,523	5,439	5,589	5,557	5,603	5,657
Keynote days	104	101	105	107	106	106
Panel B. Adding the ordinary sector return r_t (basis points per percentage point)						
Ordinary sector return r_t (ψ_m)	0.224*	0.411***	0.499***	0.457***	0.375***	0.414***
	(0.126)	(0.080)	(0.065)	(0.053)	(0.042)	(0.044)
Surprise z_d	1.866	0.486	-0.692	-2.425**	-3.140***	-2.846***
	(1.729)	(1.556)	(1.173)	(1.098)	(0.704)	(0.732)
$\times Z_d^C$ (realised, against demand: γ_m^P)	-1.655	0.142	1.078	3.707**	4.602***	3.955***
	(2.296)	(2.103)	(1.712)	(1.636)	(1.059)	(1.099)
$\times Z_d^F$ (promised less realised: δ_m)	-3.410	-7.524**	-5.730**	-7.026***	-5.843***	-4.146*
	(2.528)	(3.163)	(2.654)	(2.402)	(1.930)	(2.199)
Capability share C_d	-0.511	-1.973*	-1.893**	-1.471*	-0.962	-0.875
	(1.276)	(1.082)	(0.949)	(0.849)	(0.719)	(0.695)
Forward share F_d	1.820	13.469**	14.607***	9.385**	5.845	6.080
	(7.241)	(6.114)	(5.355)	(4.665)	(4.472)	(4.132)
Derived: the net response $\beta_m + \gamma_m^P C + \delta_m F$ at each evaluation point, and their difference.						
At a realised-heavy keynote $(C = 0.89, F = 0.06)$	0.181	0.146	-0.086	0.442	0.597**	0.420
	(0.474)	(0.448)	(0.432)	(0.417)	(0.277)	(0.280)
After a forward-heavy keynote $(C = 0.68, F = 0.21)$	0.031	-0.987**	-1.154***	-1.371***	-1.232***	-1.023***
	(0.392)	(0.486)	(0.439)	(0.371)	(0.315)	(0.348)
Difference (realised less forward)	0.150	1.133*	1.068**	1.813***	1.829***	1.443***
	(0.646)	(0.638)	(0.544)	(0.532)	(0.420)	(0.455)
Observations	5,523	5,439	5,589	5,557	5,603	5,657
Keynote days	104	101	105	107	106	106

Notes: Inflation-swap changes in basis points per percentage point of index return. Baseline columns reproduce equation (4); the $+\,r_t$ columns add the ordinary sector return of equation (9). The profile contrast is the difference in the fitted response between the delivered-heavy and promised-heavy profiles, $(C, F) = (0.89, 0.06)$ and $(0.68, 0.21)$, computed from the full coefficient covariance matrix; the intercept β_m cancels from it. The last line reports ψ_m at selected maturities. HC1 standard errors. *, ** and *** denote significance at ten, five and one per cent.

F Responses across the Content Distribution

Figure 22 adds the bottom decile of each content ranking to the top-decile profiles in Figure 6. All responses are evaluations of equation (4); the regression sample and coefficients are unchanged. The additional profiles show how fitted responses vary over the observed content distribution.

Figure 22: The net response at both ends of both cuts.



Notes: As Figure 6, with the lowest eleven conference days on each ranking added. Top row, cut on the realised share: the low end carries $C = 0.494$ against 0.891 at the high end. Bottom row, cut on the forward share: $C = 0.694$ against 0.680, so the two ends differ almost only in the horizon.

The two rankings vary different aspects of content. Ranking on delivered capability changes the total capability share substantially, from 0.494 at the bottom to 0.891 at the top. At the demand-heavy bottom profile, seven-year inflation expectations falls by 1.19 basis points. Ranking on the forward share instead holds total capability approximately fixed, at 0.694 and 0.680. Its low-forward-share profile gives a positive response of 0.84 basis points at two years. The figures thus illustrate both the role and horizon margins. Dividend-futures estimates remain positive across the displayed profiles, without comparably clear content differentiation.

G Randomized Horizon-Label Diagnostic

Delivered and promised capability shares satisfy $P_e + F_e = C_e$. Holding total capability content fixed, increasing one share necessarily reduces the other. This raises the concern that the opposing inflation responses reflect the mechanical construction of the shares rather than economically meaningful differences in disclosure content. The adding-up identity does not itself require opposite regression coefficients, but it motivates checking whether a content-free division of capability news produces a similar pattern.

I therefore compare the observed horizon split with random assignments of capability sentences to delivered and promised categories. The randomization preserves total capability content and the overall expected forward-looking fraction while removing the observed relationship between sentence content and horizon. This provides a benchmark for assessing whether the observed opposition is typical of an arbitrary split. Because the exercise uses separate single-share regressions, it is a diagnostic of the horizon partition rather than a direct test of the joint horizon interaction in equation (4).

This appendix thus compares the observed horizon-label pattern with random assignments of capability sentences to horizons. The random split preserves $P_e + F_e = C_e$ while removing the observed sentence-level horizon assignment. It evaluates a specific random-label benchmark rather than a general algebraic restriction on jointly estimated coefficients.

The classifier assigns each of the 12,598 capability sentences in the sample to one of two horizons; 1,466 of them, or 11.64 per cent, are forward-looking. The placebo discards that assignment. For each keynote with n capability sentences it draws $k^* \sim \text{Binomial}(n, 0.1164)$ and sets

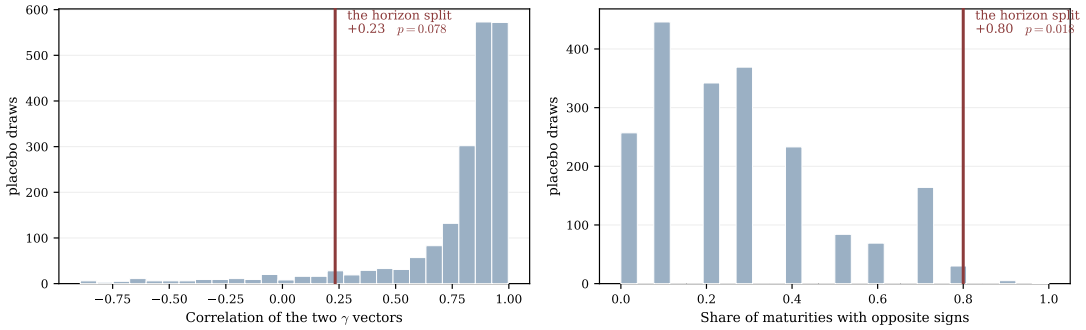
$$F_e^* = k^*/\text{den}_e, \qquad P_e^* = (n - k^*)/\text{den}_e,$$

so that $P_e^* + F_e^* = C_e$ exactly, the expected sentence-weighted forward fraction matches the observed fraction, and which sentences count as forward-looking is unrelated to what they say. Everything else is held fixed: the same keynotes, the same surprises, the same calendar, and the same single-share interaction regression, estimated separately for each share on the ten inflation-swap maturities. Each run yields a vector of ten t statistics on γ_m, and the two vectors are summarised by their correlation and by the share of maturities at which they disagree in sign. Two thousand draws give the distribution of those two statistics under a split that carries no information.

The observed horizon split gives a correlation of +0.233 between the coefficient-t vectors and opposite signs at eight maturities. Random splits give a mean correlation of +0.757 and opposite signs at 27 per cent of maturities. The observed pattern is less

positively correlated and more often opposite-signed than the typical random split. This comparison does not establish that the adding-up identity itself forces jointly estimated coefficients toward either sign pattern.

Figure 23: Two thousand random splits of capability news.



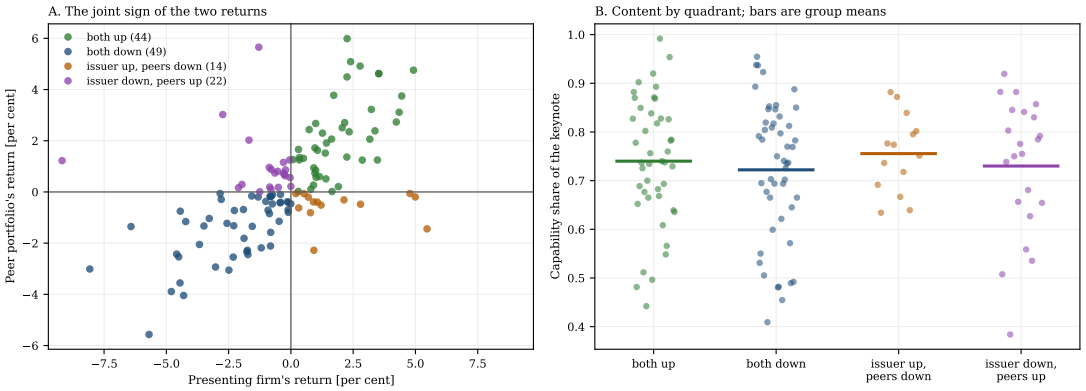
Notes: Each placebo draw reassigns every capability sentence to the forward-looking horizon with probability 0.1164, recomputes the two shares, re-estimates the separate single-share regressions on the ten inflation-swap maturities, and records the correlation between the two vectors of γ t statistics and the share of maturities at which they disagree in sign. The vertical line is the value the horizon split produces.

Of 2,000 draws, 36 produce at least as much sign disagreement ($p = 0.018$), and 156 produce a correlation as low as observed ($p = 0.078$). The random-split correlation has a standard deviation of 0.323 and a minimum of -0.893. The probabilities describe the specified random-label comparison.

H Issuer–Peer Return Patterns and Disclosure Content

Section 8.2 divides the events by the joint sign of the presenting firm's return and the return of its semiconductor peers. This appendix displays the issuer–peer return distribution and its relationship with keynote content.

Figure 24: The issuer–peer sign split and the content of the keynotes it separates



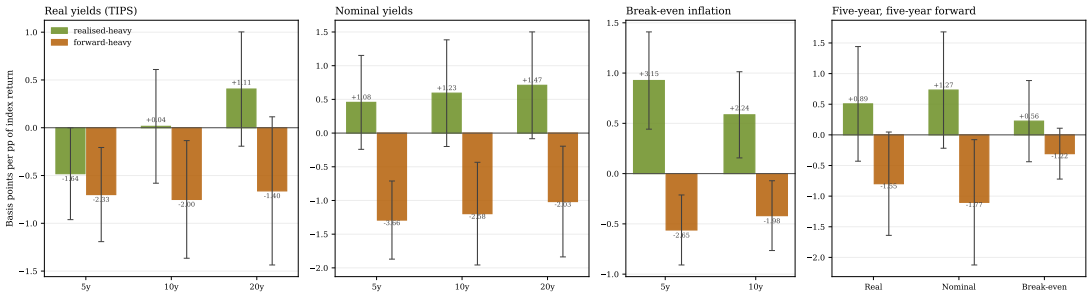
Notes: Both panels use the 129 keynotes of the event set for which the presenting firm and at least five of the other semiconductor names in the universe of Section 6.1 can be priced over the keynote's pair of closes; the peer portfolio is equal-weighted and excludes the presenter. Panel A plots the two returns against each other, coloured by quadrant. Panel B plots the capability share of each keynote by quadrant, with horizontal bars at the group means. Returns in this descriptive figure are raw. The regression classification in Section 8.2 instead uses firm-specific market-model residuals; the raw-return quadrant counts are therefore distinct from the common- and relative-repricing regression groups.

Two things are visible. The first is that the split has real variation: the two returns are positively related but far from collinear, and 36 of the 129 keynotes fall in the off-diagonal quadrants, 14 with the presenter up and its peers down and 22 the other way round. The second is that average capability content differs little across the raw-return quadrants. The capability share averages 0.740 where both legs rise, 0.722 where both fall, 0.756 where the presenter alone rises and 0.730 where the presenter alone falls; the common-minus-relative difference is −0.009. The return-sign grouping and the text shares describe distinct features of the same events. Similar group means do not establish statistical independence.

I Additional Treasury-Yield Results

This appendix reports the Treasury-yield estimates discussed in Sections 9.1 and 7. Daily changes in nominal yields, TIPS yields, and break-even inflation are estimated using equation (4).

Figure 25: Treasury yields responses by contents



Notes: The net response $\beta + \gamma C + \delta F$ from equation (4), in basis points per percentage point of index return, with standard errors from the full 3×3 covariance of the three coefficients. The two content points are the ends of the forward-share ranking: a low-forward-share keynote at $C = 0.694$, $F = 0.019$ and a forward-heavy one at $C = 0.680$, $F = 0.209$, the mean content of the eleven conference days at each end of the forward share. Whiskers are 90 per cent intervals and the number above each bar is the t statistic.