

# Pricing Transformative AI

Isaiah Andrews and Maryam Farboodi\*

## Abstract

Is AI expected to transform the economy? We answer this question by examining bond yield responses to major AI model releases. Economic theory predicts that transformative technologies may influence interest rates by changing growth expectations, increasing uncertainty about growth, or raising concerns about existential risk. We find that bond yields show economically and statistically significant declines around AI model release dates. Under a simple model these declines indicate that new AI model releases have led investors to meaningfully reduce growth expectations, lower the perceived probability of existential risk, or both. By contrast, changes in investor uncertainty about future growth do not appear to explain observed yield changes. Financial markets appear to take AI growth impacts seriously, but recent model releases may have fueled pessimism about the direction of these effects.

## 1 Introduction

Since the debut of ChatGPT in November 2022, generative AI models have attracted intense interest from policymakers, researchers, and businesses. Some discussions of these models have raised the possibility they could lead to an increase, perhaps even a dramatic acceleration, in the rate of economic growth. Other discussions have suggested the possible gains from AI may be overstated, and argued that widespread AI adoption could potentially slow economic growth. Many authors and AI entrepreneurs have even raised the possibility that poorly understood- and controlled-AI could pose an existential risk to humanity.

It can be unclear to what extent the enthusiasm around AI reflects a genuine belief in its transformative potential, as opposed to belief in profit opportunities which may

---

\*This version: July 21, 2025. First version: July 14, 2025. Andrews: MIT Department of Economics and NBER; iandrews@mit.edu. Farboodi: MIT Sloan School of Business, NBER, and CEPR; farboodi@mit.edu. We thank Marcellus Andrews, Markus Brunnermeier, Basil Halperin, Lukasz Kowalik, Stefan Nagel, Jesse Shapiro, and Cheryl Smith for very helpful discussions of and comments on this project.

not translate into widespread or persistent growth. While the future impacts of AI are inherently unknown, understanding the beliefs of market participants is a potentially valuable input to both policy and research discussions. Financial market data has been used to infer market beliefs about novel technologies in other settings, but empirical evidence about market beliefs on transformative AI is limited. In this paper, our goal is to use financial market data to provide systematic evidence regarding the beliefs of market participants about the possibility of transformative AI, by which we will mean AI technologies with a large and sustained impact on living standards, particularly through impacts on consumption growth or existential risk. The premise of our analysis is that if economic actors take seriously the possibility of transformative AI, this should be reflected in a wide range of forward-looking behaviors and, consequently, in long-term asset prices, including assets such as US Treasury bonds which are not directly connected to AI.

That beliefs about transformative AI should affect agents' optimal choices is pointed out by e.g. Jones (2024). Chow et al. (2024) combine this observation with classic insights from consumption-based asset-pricing to relate risk-free interest rates to market beliefs about transformative AI. The intuition is simple: if we think AI will dramatically increase the rate of economic growth, then (on average across the economy) we must expect to be richer in the future than we are today. This should decrease the marginal value of future consumption relative to present consumption, so real interest rates must rise in equilibrium. On the other hand, if we think AI poses an existential risk, and so doubt that we will be alive in the future, this should also drive up interest rates. Thus, both higher growth expectations and more concern for existential risk should increase real interest rates.

Motivated by this prediction, we study the behavior of US Treasury yields around major model release dates for five major AI labs (OpenAI, Anthropic, Google Deepmind, xAI, and DeepSeek) in calendar years 2023 and 2024. As shown in Figure 1, we find that US Treasury yields substantially *decline* around model release dates, with an average decline across model release dates exceeding 12 basis points, or 0.12 percentage points. These declines are economically large and persist through 15 trading days after the model release. We find similar results for Treasury Inflation Protected Securities (TIPS) and corporate bond yields, and our results for longer-maturity bonds are statistically significant under an auxiliary assumption that AI model release dates are as good as random. Yield impacts appear to begin before the release of the model, which may not be surprising given that for at least some releases, we know that models were made available to outside experts prior to the release date.

We interpret these results through a simple equilibrium model of asset prices. We

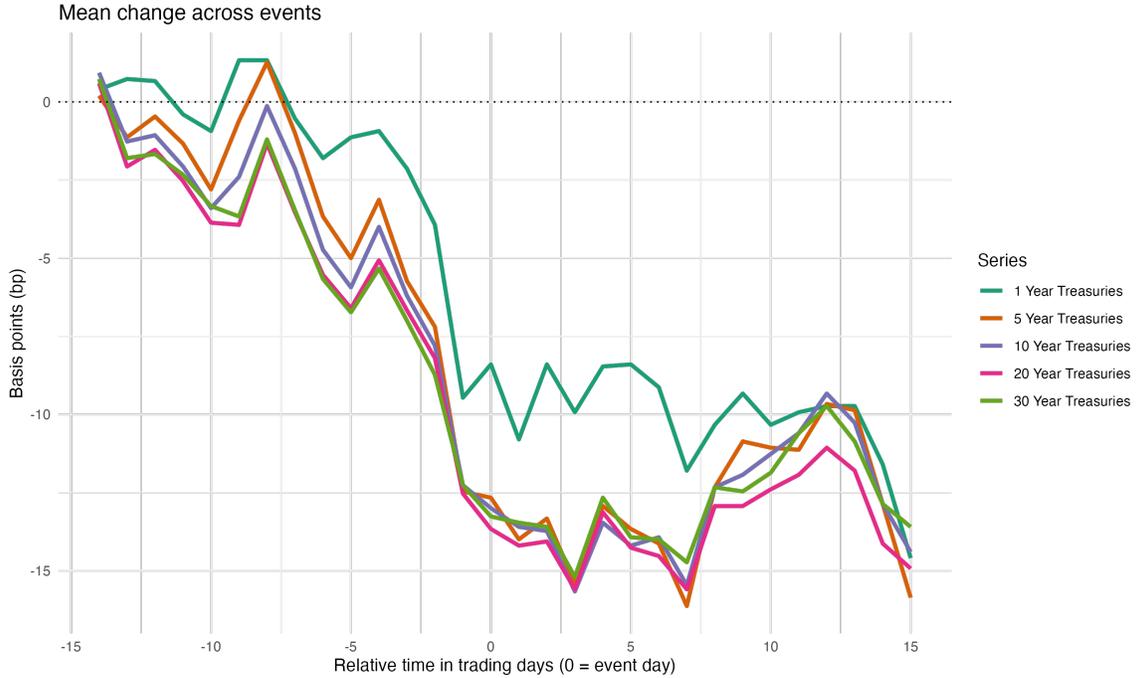


Figure 1: Average change in yields relative to fifteen trading days before event for constant-maturity US Treasury Bonds. Average taken across 15 AI release events in the 2023 and 2024 calendar years.

first show that “doom” (i.e. existential risk) and “bliss” (i.e. extremely fast growth which eliminates material scarcity, and which e.g. Jones 2024 discusses as a singularity) have equivalent asset pricing implications, since both drive the marginal utility of consumption to zero. Thus, we cannot hope to disentangle beliefs about “bliss” and “doom” via asset prices.

We then use a restrictive version of this model (assuming, inter alia, a representative agent with CRRA utility, as commonly used in consumption-based asset pricing literature) to go further and quantitatively interpret our empirical results. This restricted model implies that the average AI model release in our sample led investors to think that (i) expected consumption growth is lower, (ii) extreme “doom” or “bliss” outcomes are less likely, or some combination of the two. In particular, this restricted version of our model implies that the average model release led to an approximately 0.132 percentage point drop in the annual probability of “bliss” or “doom,” or a  $0.132/\gamma$  percentage point decline in the annual rate of consumption growth (for  $\gamma$  the risk aversion of the representative agent). By contrast, changes in consumption growth uncertainty do not appear to explain our results.

Taken together, these results suggest that investors, in aggregate, do take seriously the

possibility of transformative AI, since new information about AI models has an economically and statistically significant impact on non-AI-related asset prices. However, the primary direction of updating, on average across the model release dates we study, was towards lower consumption growth, a lower probability of “bliss” and “doom,” or some combination of the two.<sup>1</sup> Of course, there are many ways that reality may deviate from our simple model, and these deviations may suggest alternative explanations for our empirical findings. Any alternative explanation must, however, account for large, sustained yield declines in one of the most liquid markets in the world around AI model release dates, and so may be of interest in its own right.

**Literature Review** A number of recent papers discuss the theoretical possibility that transformative AI could increase the growth rate of the world economy (Nakamura et al., 2013; Brynjolfsson et al., 2019; Trammell and Korinek, 2023; Acemoglu and Lensman, 2024; Jones, 2024; Korinek and Suh, 2024; Erdil et al., 2025). On the other hand, some papers have suggested that widespread adoption of AI could lead to a growth slowdown, (Gordon, 2016; Acemoglu and Restrepo, 2020), or even pose an existential risk to humanity (Acemoglu and Lensman, 2024; Jones, 2024; Kokotajlo et al., 2025).

On the empirical front, although there is little evidence about the aggregate impact of AI on the economy and market perception of this impact, there is a small but growing literature that uses data on job postings or asset prices to study the impact of AI on labor outcomes and compensation, as well as on firm growth (Webb, 2019; Acemoglu et al., 2022; Babina et al., 2023, 2024; Eisfeldt et al., 2024).

The influence of growth prospects on financial markets is a widely discussed topic in the consumption-based asset pricing literature (see Mehra 2012 for a summary and Duffie 2010 for a textbook treatment). An important observation in this literature is that agents’ discount factors, expected growth, and perceived growth uncertainty all influence prevailing interest rates. Chow et al. (2024) abstract away from consumption growth uncertainty and show that discounting (e.g. due to existential risk) and growth expectations impact interest rates the same way in the context of transformative AI. Chow et al. (2024) further show that, consistent with their theoretical analysis, real interest rates are positively correlated with both growth expectations and realized growth in a cross-country analysis. Our main

---

<sup>1</sup>This is not necessarily incompatible with rising stock prices for AI-related firms, since one could think these firms will be highly profitable for reasons, e.g. substitution for labor thanks to marginally lower costs, which need not imply large or sustained growth effects. See Section 6 of Chow et al. (2024) for further discussion of why the relationship between AI expectations and equity prices may be ambiguous.

contribution is to document how news about AI progress impacts interest rates. We then interpret these impacts through the lens of a consumption-based asset pricing model.

Section 2 introduces a simple equilibrium asset pricing model and uses it to predict the effects of transformative AI on bond yields. Section 3 describes our data and empirical strategy, including the permutation tests we use to assess statistical significance. Section 4 reports our main empirical results, while Section 5 quantitatively interprets these results in a simplified version of our model. Finally, Section 6 provides additional discussion.

## 2 Transformative AI in a Dynamic Economy

This section lays out a simple model of a dynamic stochastic economy, and shows that this model makes stark predictions about how investor beliefs about the possibility of transformative AI translate to asset prices. As we discuss in the introduction, by *transformative AI* we mean AI technology that substantially changes the future trajectory of the economy. Specifically, following Jones (2024) and Chow et al. (2024) we consider the possibility that AI may (i) substantially change the rate of economic growth or even (ii) lead to a more radical shift after which asset holdings become irrelevant for utility.

For changes in the rate of economic growth, we allow that the growth effects of AI could be either positive or negative. For instance, AI technology could lead to a slowdown in economic growth due to so-so AI and rent-seeking, as argued in Acemoglu and Restrepo (2020). We thus regard the impact of AI as one source of uncertainty about future consumption growth. As new information is revealed about the development of AI, this will change agents' consumption expectations.

More extreme changes due to AI could likewise take different forms. On the one hand, as again discussed in e.g. Jones (2024) and Chow et al. (2024) AI could pose an existential risk. If AI were to wipe out humanity, it is clear that asset returns beyond that point are irrelevant for utility. Interestingly, happier scenarios in which AI leads to an explosive increase in the rate of economic growth and the elimination of material scarcity lead to the same conclusion: if consumption possibilities become unconstrained, asset returns are once again irrelevant. We shorthand these possibilities as “doom” and “bliss,” respectively.

To model these possible impacts from AI, following e.g. Chapter 2 of Duffie (2010), we consider a discrete-time economy over periods  $t=0,1,\dots,\bar{T}$ , with uncertainty described by a probability space  $(\Omega,\mathcal{F},\mathbb{P})$  where the  $\mathcal{F}_t \subseteq \mathcal{F}$  denotes the set of events which are known at period  $t$ . For simplicity we assume a finite number of states  $\omega \in \Omega$  and agents  $i$ . To

capture the possibility of “doom” and “bliss,” similar to Jones (2024) we assume each agent  $i$  has time-separable utility

$$\mathbb{E}_0 \left[ \sum_{t=1}^{\bar{T}} \beta^t (1\{t \leq T\} u_i(C_{i,t}) + 1\{t > T\} U_{i,t}^*) \right],$$

where  $\mathbb{E}_t$  denotes the conditional expectation given  $\mathcal{F}_t$ , and  $T \leq \bar{T}$  denotes the (random) date after which “doom” or “bliss” occurs. We assume that  $u_i$  is increasing and concave for all  $i$  with  $\lim_{c \rightarrow \infty} u'_i(c) = 0$ , while flow-utilities  $U_{i,t}^*$  after  $T$  are independent of asset holdings and we henceforth normalize these to zero and write agent  $i$ 's utility as  $\mathbb{E}_0 \left[ \sum_{t=1}^T \beta^t u_i(C_{i,t}) \right]$ .

We assume complete markets and absence of arbitrage. By standard arguments (Duffie, 2010), this implies that there exists a stochastic discount factor (SDF) that prices all assets. In particular if we consider an asset that pays  $Y_{t+h}$  units of consumption in period  $t+h$  and nothing at any other time, its period  $t$  price is given by

$$V_t(Y_{t+h}) = \mathbb{E}_t[M_{t,t+h} Y_{t+h}], \quad (1)$$

where  $M_{t,t+h}$  is the SDF  $t$  to  $t+h$ , for simplicity we write  $M_{t+1} \equiv M_{t,t+1}$ , and  $M_{t,t+h} = \prod_{s=1}^h M_{t+s}$  cumulates the one-step-ahead SDFs. More generally, let  $Y = \{Y_{t,h}\}_{h=0}^{\bar{T}-t}$  denote a general stream of payoffs  $Y_{t,h}$  for periods  $h=0, \dots, \bar{T}-t$ . The asset with this stream of payoffs has time- $t$  price  $V_t(Y) = \sum_{h=0}^{\bar{T}-t} V_t(Y_{t+h})$ .

Standard arguments further imply that in equilibrium the SDF coincides with the marginal rate of substitution for a representative agent with utility  $\mathbb{E}_0 \left[ \sum_{t=1}^T \beta^t u(C_t) \right]$ , where  $C_t = \sum_i C_{i,t}$  is aggregate consumption and  $u(C_t) = \sum_i \lambda_i u_i(C_{i,t})$  for  $\lambda_i \geq 0$ , where  $u$  is increasing and concave by construction. That is, we can write the SDF as

$$M_{t,t+h} = \beta^h \frac{u'(C_{t+h})}{u'(C_t)} 1\{t+h \leq T\}, \quad (2)$$

so there is a direct relationship between aggregate consumption  $C_t$ , the “doom” or “bliss” date  $T$ , and the SDF.<sup>2</sup>

Equation (2) has two important implications. First, note that the extreme possibilities

---

<sup>2</sup>Indeed, this follows from the fact that each agent’s marginal utility obeys the same equality,

$$M_{t,t+h} = \beta^h \frac{u'_i(C_{i,t+h})}{u'_i(C_{i,t})} 1\{t+h \leq T\}. \quad (3)$$

of “doom” and “bliss” both enter only through the date  $T$  after which asset holdings are irrelevant. Consequently, beliefs about “doom” and “bliss” have identical asset-pricing implications. Hence, under this model we have no hope of telling the two apart based on asset prices. Second, note that since the representative agent’s flow utility  $u$  is increasing and concave, increases in future aggregate consumption  $C_{t+h}$  lead to a decrease in the SDF. Hence, if agents expect AI to lead to an acceleration in the rate of aggregate consumption growth this will, ceteris paribus, lead to a drop in the SDF and additional discounting of future payoffs. As noted by e.g. Gil (2024), however, since  $u$  is concave even news which increases expected future consumption  $E_t[C_{t+h}]$  could lead to an increase in the mean of the SDF and thus a *decrease* in discounting if it implies a sufficiently large increase in uncertainty.

**Bond Pricing Implications** While the analysis above applies to general payoff streams  $Y$ , our empirical analysis will focus on bond prices. To study the implications of AI beliefs for bond prices, let  $1_{t+h}$  denote a risk-free,  $h$ -period ahead zero-coupon bond (i.e. the risk-free bond which pays one unit of consumption  $h$  periods in the future and nothing at any other time). By Equation (1) this bond’s price is given by

$$V_t(1_{t+h}) = \mathbb{E}_t[M_{t,t+h}] = \mathbb{E}_t \left[ \beta^h \frac{u'(C_{t+h})}{u'(C_t)} 1_{\{t+h \leq T\}} \right],$$

and so is simply the expected  $h$ -period ahead SDF. Since it is more common to work with bond yields than with prices, note that the period- $t$  yield on the risk-free bond  $1_{t+h}$  can be written as  $y_{t,t+h} \equiv V_t(1_{t+h})^{-\frac{1}{h}} = \mathbb{E}_t[M_{t,t+h}]^{-\frac{1}{h}}$ , which can be further re-written as

$$y_{t,t+h} = \frac{1}{\beta \mathbb{P}(t+h \leq T)^{\frac{1}{h}} \mathbb{E}_t \left[ \frac{u'(C_{t+h})}{u'(C_t)} \mid t+h \leq T \right]^{\frac{1}{h}}}.$$

Thus, zero-coupon bond yields are decreasing in the discount factor  $\beta$ , increasing in the probability that  $T$  arrives before the bond pays off  $\mathbb{P}(t+h > T)$ , and decreasing in expected marginal utility in period  $t+h$  conditional on  $T$  not yet having arrived,  $\mathbb{E}_t \left[ \frac{u'(C_{t+h})}{u'(C_t)} \mid t+h \leq T \right]$ . Since  $u$  is concave, yields are thus increasing in consumption growth (in the sense of first order stochastic dominance). Thus, as noted by Chow et al. (2024) both an increase in anticipated consumption growth and a closer expected arrival for  $T$  lead to higher risk-free yields.<sup>3</sup>

<sup>3</sup>We note, however, that beliefs about consumption growth and about  $T$  have distinct implications for the prices of *risky* assets. In particular, if we consider the ratio of risky and risk-free asset prices for a given future period,  $V_t(Y_{t+h})/V_t(1_{t+h}) = \mathbb{E}_t \left[ Y_{t+h} \frac{u'(C_{t+h})}{u'(C_t)} \mid t+h \leq T \right]$ , this ratio depends only on behavior

While the Treasuries which are the focus of our analysis below are multi-period rather than zero coupon bonds, the comparative statics are much the same. In particular, if we consider a  $h$ -period risk-free bond with coupon  $c$  and face value  $d$ , this corresponds to payoff stream  $B = \{c1_{t+1}, c1_{t+2}, \dots, c1_{t+h-1}, (c+d)1_{t+h}\}$  and so has price

$$V_t(B) = d\mathbb{E}_t[M_{t,t+h}] + c \sum_{s=1}^h \mathbb{E}_t[M_{t,t+s}] = \frac{d}{y_{t+h}^h} + c \sum_{s=1}^h \frac{1}{y_{t+s}^s}.$$

**Empirical Strategy** The model above suggests an empirical strategy for learning about changes in AI beliefs from asset prices: if we have a date  $t$  at which we believe information arrived about the future course of AI, then changes in long-dated asset prices around this date should incorporate the impact of the new information about AI.

To fix ideas, again consider the price for an asset that pays  $Y_{t+h}$  units in period  $t+h$ . If we think new information about AI arrived at  $t$ , we may compare prices at  $t_-$  and  $t_+$  for  $t_- < t < t_+ \ll h$ , and use the fact that  $V_{t_-}(Y_{t+h}) = \mathbb{E}_{t_-}[M_{t_-,t_+} V_{t_+}(Y_{t+h})]$  to write

$$V_{t_+}(Y_{t+h}) - V_{t_-}(Y_{t+h}) = V_{t_+}(Y_{t+h}) - \mathbb{E}_{t_-}[V_{t_+}(Y_{t+h})] - \mathbb{E}_{t_-}[(M_{t_-,t_+} - 1)V_{t_+}(Y_{t+h})].$$

If the time difference  $t_+ - t_-$  is reasonably small we expect the final term to be negligible.<sup>4</sup> Hence, by (1) and the law of iterated expectations we can approximate

$$V_{t_+}(Y_{t+h}) - V_{t_-}(Y_{t+h}) \approx \mathbb{E}_{t_+}[M_{t_+,t+h} Y_{t+h}] - \mathbb{E}_{t_-}[M_{t_+,t+h} Y_{t+h}]. \quad (4)$$

Thus the change in prices between  $t_-$  and  $t_+$  gives us, approximately, the difference in conditional expectations for the discounted payoff  $M_{t_+,t+h} Y_{t+h}$  at information sets  $\mathcal{F}_{t_-}$  and  $\mathcal{F}_{t_+}$ . In particular, if we consider the risk-free asset  $Y_{t+h} = 1_{t+h}$ , changes in prices reveal the change in the conditional mean of the SDF  $M_{t_+,t+h}$ .

For a given pair  $t_-$  and  $t_+$  the difference  $V_{t_+}(Y_{t+h}) - V_{t_-}(Y_{t+h})$  reflects all information that arrives between those dates, not just information about AI. Hence, in our empirical analysis we will average across a series of AI news dates. So long as there is not other price-relevant information which systematically arrived at the same time as AI news, comparing behavior at AI dates to that at other dates will estimate the average effect of

---

conditional on  $T$  not yet having arrived. We hope to use this fact to distinguish changes in beliefs about consumption from changes in beliefs about  $T$  in a future version of this paper.

<sup>4</sup>By Cauchy-Schwarz,  $\mathbb{E}_{t_-}[(M_{t_-,t_+} - 1)V_{t_+}(Y_{t+h})] \leq \sqrt{\mathbb{E}_{t_-}[(M_{t_-,t_+} - 1)^2]} \sqrt{\mathbb{E}_{t_-}[V_{t_+}(Y_{t+h})^2]}$ .

AI news, though it will be important to account for the possibility of other news when assessing statistical uncertainty. As already noted, we will also use data on multi-period bonds rather than zeros. Since our primary focus will be on long-maturity bonds, however, most bond payoffs will be in the future and the intuition provided above for zeros will again translate to the bonds we study.

### 3 Data and Methods

As the theory above suggests, if market participants think that AI may have large growth effects then new information about the trajectory of AI should impact long-term asset prices, including for assets that are not directly related to AI such as long-term risk-free bonds. We examine this prediction empirically, describing the data and methods we employ in this section and our empirical results in the next.

#### 3.1 AI News Events

To look for asset prices changes around the arrival of AI news, we need to know a set of dates at which AI information arrived. While there are a variety of reasonable approaches one might take to this problem, we focus on release dates for new generative AI models from five major AI laboratories: OpenAI, Google DeepMind, Anthropic, xAI, and DeepSeek.<sup>5</sup> For each lab, we focus on major updates to the lab’s flagship model series (e.g. ChatGPT in the case of OpenAI), and use the release date from the lab’s website.<sup>6</sup> We limit attention to releases in calendar years 2023 and 2024, a period that (i) follows the November 2022 release of ChatGPT, which saw a significant increase in public attention to AI capabilities, and (ii) precedes the tariff announcements and other US macroeconomic policy changes that began in 2025. For OpenAI we include the “reasoning model” o1, since other labs included such models as part of their flagship series rather than numbering them separately (e.g. Gemini 2.5 from Google and Claude 3.7 from Anthropic, both released after our main analysis window).<sup>7</sup> Table 1 collects the resulting release dates.

We use new AI model releases as our event dates in order to capture new, forward-looking

---

<sup>5</sup>These are the laboratories appearing in the top 10 style-adjusted rankings on the Chatbot Arena leaderboard as of June 29, 2025 (Chiang et al., 2024).

<sup>6</sup>For DeepSeekV2, we were unable to find an announcement on the lab’s website, and so instead use an announcement date from DeepSeek’s X account.

<sup>7</sup>Appendix A reports results which drop the o1 release from our date series.

Table 1: AI Model Release Dates

<b>Date</b>	<b>AI Laboratory</b>	<b>Model</b>
<i>2023 Releases</i>		
02/06/2023	Google	Bard
03/14/2023	OpenAI	ChatGPT 4
03/14/2023	Anthropic	Claude 1
07/11/2023	Anthropic	Claude 2
11/03/2023	xAI	Grok 1
11/21/2023	Anthropic	Claude 2.1
12/06/2023	Google	Gemini Pro 1.0
<i>2024 Releases</i>		
02/15/2024	Google	Gemini Pro 1.5
03/04/2024	Anthropic	Claude 3
03/28/2024	xAI	Grok 1.5
05/06/2024	DeepSeek	DeepSeek V2
05/13/2024	OpenAI	ChatGPT 4-o
06/20/2024	Anthropic	Claude 3.5 Sonnet
08/13/2024	xAI	Grok 2
09/05/2024	DeepSeek	DeepSeek 2.5
12/05/2024	OpenAI	o1
12/11/2024	Google	Gemini 2.0
12/26/2024	DeepSeek	DeepSeek V3

*Notes:* This table presents the major AI model releases used in our event study analysis.

information about AI capabilities, rather than other aspects of technology or financial performance of firms. Put differently, our hypothesis is that major model releases provide information not only about the current state of AI capabilities but also about the rate of technological progress, potentially causing market participants to update their beliefs about future AI development. These events are also less directly linked to financial outcomes than some other plausible event dates, such as earnings announcements. At the same time, it is clear that information about AI system capabilities arrives outside of new model releases for these particular AI labs. There are many other AI researchers and firms, and even the firms we study make many announcements and incremental model releases outside the set of major releases we consider. So long as some information is arriving around the dates we study, such alternative information sources do not pose a threat to our empirical strategy.

More directly relevant for us, for at least some model releases we know that certain experts were given early access to the model prior to the official release.<sup>8</sup> To partially cap-

<sup>8</sup>See for instance Mollick (2024).

ture such information “leakage” our empirical specifications will include a window of dates prior to the model release (15 trading days, or approximately 3 weeks, for our preferred specifications). While this extended window is still unlikely to capture all information leakage, uncaptured leakage should reduce the amount of information arriving in our event windows. We expect this will bias us against finding yield responses.

### 3.2 Financial Market Data

Motivated by the theory in Section 2, to look for effects of AI information on long-run consumption expectations we examine the behavior of bond yields of different maturities around major model release dates. We consider two bond series in the main text.

1. **Nominal Treasury Yields:** We use constant-maturity Treasury yields from the Federal Reserve Economic Data (FRED) database for maturities of 1, 5, 10, 20, and 30 years (Board of Governors of the Federal Reserve System, US, 2025b).
2. **Corporate Bond Indices:** We use ICE BofA corporate bond effective yield indices broken out by maturity (1-3 year, 3-5 year, 5-7 year, 7-10 year, 10-15 year, and 15+ year indices – Ice Data Indices, LLC 2025a).

All yield data are measured in percentage points and recorded at daily frequency. Appendix C.1 considers a third data series, Treasury Inflation-Protected Securities (TIPS).

### 3.3 Event Study Methodology

We use an event-study approach to look for changes in yields around our event dates. For each AI event date  $t \in \mathcal{T}$  and each yield series, we calculate the change in yields relative to a pre-event baseline, defined as  $b$  trading days before the event. Thus, the change from the baseline date to relative date  $s$  is:

$$\Delta y_{t,s} = y_{t+s} - y_{t-b}. \tag{5}$$

This gives us a yield change for each event date  $t \in \mathcal{T}$ . We next aggregate these changes across event dates to obtain a single summary statistic, considering both the mean change

$$\text{MeanChange}_s = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \Delta y_{t,s}$$

and the mean absolute change

$$\text{MeanAbsChange}_s = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} |\Delta y_{t,s}|.$$

The mean measures whether there were systematic patterns in the direction of yield changes around our event dates, while the mean absolute change measures whether there were systematic patterns in the magnitude of yield changes. Since means can be sensitive to behavior of a single observation, especially given our small sample size, Appendix C.3 provides results for median and median absolute changes, which are by construction more robust to outliers. These prove to be qualitatively very similar to our main results.

### 3.4 Permutation Inference

To gauge whether markets are responding to AI model releases, we need a way to judge whether the yield movements we observe around model releases are larger than one would expect due to chance. Given our very limited sample size, it is important to use a method that is valid in small samples. To this end, we assess statistical significance via permutation inference, under the assumption that our AI release dates are as good as randomly assigned (and, in particular, can be treated as a uniform random draw from the trading dates in our analysis window).

Our procedure works as follows:

1. We define the set of potential “placebo” event dates as all trading days in our sample period (subject to the full event window from  $t_- = t - b$  to  $t_+ = t + s$  being within the sample).
2. For each  $m \in \{1, \dots, 5000\}$  we randomly sample (without replacement)  $K$  placebo dates from this set, where  $K$  equals the number of actual AI events in our sample (again restricted to events where the full event window is within the sample), and compute our test statistics using these placebo dates.
3. We compare the test statistics computed using the actual model release dates to the empirical distribution across placebo samples. If markets did not react to AI events, the event dates were selected as good as randomly, and yields were continuously distributed, then the probability that our observed test statistics would exceed the  $p$ -th percentile of the placebo distribution would equal  $p$  up to simulation error. In

reality our yield data are only measured up to level of basis points, so in cases of ties we round away from statistical significance.

This approach gives a test for the “sharp” null hypothesis of no impact on yields, which is valid in finite samples under the auxiliary assumption that the AI model release dates can be treated as a random draw. While this is a strong assumption, it may be partly justified by the fact that model releases are substantially driven by technical development timelines rather than financial market conditions. Examining the release dates in Table 1 we do not observe very strong calendar patterns, though e.g. Fridays appear somewhat underrepresented (with only 1 of the 17 unique dates in the sample), and there are more dates in 2024 than 2023 (with 11 of the 17 unique dates). If one wanted to replace our assumption that release dates are drawn uniformly at random with some other specific assumption about their distribution, our approach generalizes directly. To explore sensitivity to our assumptions, we discuss several robustness checks following our empirical results.

## 4 Empirical Results

We next report our empirical results. We begin by examining whether there are statistically significant changes in yields around our event dates, evaluating statistical significance relative to the placebo distribution as described in the previous section.

Recall that p-values measure the probability of observing a more extreme outcome were the null hypothesis true. Hence, small p-values correspond to outcomes which are unlikely to arise under the null (in our case, if AI model releases have no effect on yields, and release dates are as good as random). Consequently, a 10% test of the null rejects when the p-value is less than 0.1, and a 5% test rejects when the p-value is less than 0.05. Tables 2 and 3 report two-sided p-values for the fixed-income yield series we consider (US Treasuries and corporate bond indices), reporting results for both mean and mean absolute changes, and comparing yields either five or fifteen days before and after each model release (that is, setting  $b=s=5$  or  $b=s=15$  in the notation of Equation 5).

The results in Tables 2 and 3 paint a consistent picture. First considering mean changes in bond yields, we see suggestive evidence of changes in yields for longer maturity bonds for the  $\pm 5$  trading day specifications, though the p-values fall short of significance at conventional levels.<sup>9</sup> For the  $\pm 15$  day specifications, we see statistically significant changes

---

<sup>9</sup>The effects for 20 and 30 year Treasuries are significant at the 10% level if one instead adopts the

Table 2: Two-sided p-values based on constant-maturity US Treasury yields

Maturity	Mean Change		Mean Absolute Change	
	$\pm 5$ days	$\pm 15$ days	$\pm 5$ days	$\pm 15$ days
1 Year	0.231	0.144	0.795	0.428
5 Year	0.195	0.054*	0.884	0.309
10 Year	0.120	0.032**	0.726	0.341
20 Year	0.110	0.020**	0.605	0.344
30 Year	0.121	0.025**	0.582	0.333

*Notes:* The “Mean Change” columns consider the mean change in yields across event dates, while the “Mean Absolute Change” columns consider mean absolute changes. For each statistic, we compare yields 5 or 15 days before and after each model release (in  $\pm 5$  the  $\pm 15$  columns, respectively). P-values are computed based on drawing placebo event dates 5000 times (uniformly at random from trading days in the sample with sufficient window on either side) and comparing resulting placebo distributions to observed changes around AI model releases. \*\* (\*) denotes statistical significance at the 5% (10%) level.

in yields for longer-maturity bonds. This holds true whether we consider Treasuries or corporate bonds. By contrast, when we consider mean absolute changes we do not find statistically significant effects at conventional significance levels for any of the maturities studied. These patterns again hold across both Treasuries and corporate bonds. Appendix C.1 finds analogous results for TIPS.

This pattern is different than we, at least, anticipated before analyzing the data: if market participants took seriously the possibility of transformative AI, and learned more than usual about AI’s future trajectory around model release dates, we would expect larger-than-average yield changes around model release dates (and hence, potentially, statistical significance for mean absolute changes) but not necessarily a consistent direction of change (and hence, potentially, no statistical significance for mean changes). Our results show the opposite: there do not appear to be yield changes of statistically different magnitude around AI model release dates (since we do not see statistical significance for mean absolute changes). However, there are statistically significant patterns in the direction of yield changes, especially at longer maturities, as revealed by our results on mean changes.

**Event Study Plots** To further explore what is happening around our AI events, Figures 2 and 3 plot, for each yield series and each horizon  $s \in \{-14, \dots, 15\}$ , the mean and mean

---

stricter event-date definition considered in Appendix A which omits OpenAI’s o1 model release, or considers median changes in yields as in Appendix C.3.

Table 3: Two-sided p-values based on ICE corporate bond index yields

Maturity	Mean Change		Mean Absolute Change	
	$\pm 5$ days	$\pm 15$ days	$\pm 5$ days	$\pm 15$ days
5 Year	0.354	0.069*	0.298	0.788
10 Year	0.208	0.021**	0.296	0.942
20 Year	0.184	0.014**	0.394	0.810
30 Year	0.206	0.014**	0.421	0.623

*Notes:* The “Mean Change” columns of consider the mean change in yields across event dates, while the “Mean Absolute Change” columns consider mean absolute changes. For each statistic, we compare yields 5 or 15 days before and after each model release (in  $\pm 5$  the  $\pm 15$  columns, respectively). P-values are computed based on drawing placebo event dates 5000 times (uniformly at random from trading days in the sample with sufficient window on either side) and comparing resulting placebo distributions to observed changes around AI model releases. \*\* (\*) denotes statistical significance at the 5% (10%) level.

absolute change in yields (relative to  $b = 15$  days before the event) across the observed AI model releases. For comparison, at each horizon we also plot the mean of the placebo distribution and bands which contain, 90%, 95%, and 99% of the placebo draws pointwise at each horizon (with equal mass assigned to the two tails). These bands are another way to express our placebo tests. For instance, our placebo test rejects the null of no effect at the 10% level at a given horizon if and only if the mean change at that horizon lies outside the 90% placebo band.

Examining these plots, we see that for both mean and mean absolute changes there is limited (and largely statistically insignificant) departure from the placebo distribution between  $t - 15$  and  $t - 7$ . Bond yields, especially for long-maturity bonds, show declines starting around  $t - 7$ . These declines largely persist through  $t + 15$ . The apparent anticipatory effects (i.e. effects before the model release date  $t$ ) are consistent with the fact, discussed above, that some information about new models may become available to market participants prior to the official model release.

The overall fall in yields around model releases is quantitatively large, exceeding 12 basis points by the end of the window considered. Moreover, consistent with our previous findings these changes are statistically significant relative to the placebo distribution at conventional significance levels. By contrast, our series for mean absolute changes again do not show strong deviations from the placebo distribution. Thus, we find economically and statistically significant declines in long-maturity bond yields around AI model releases,

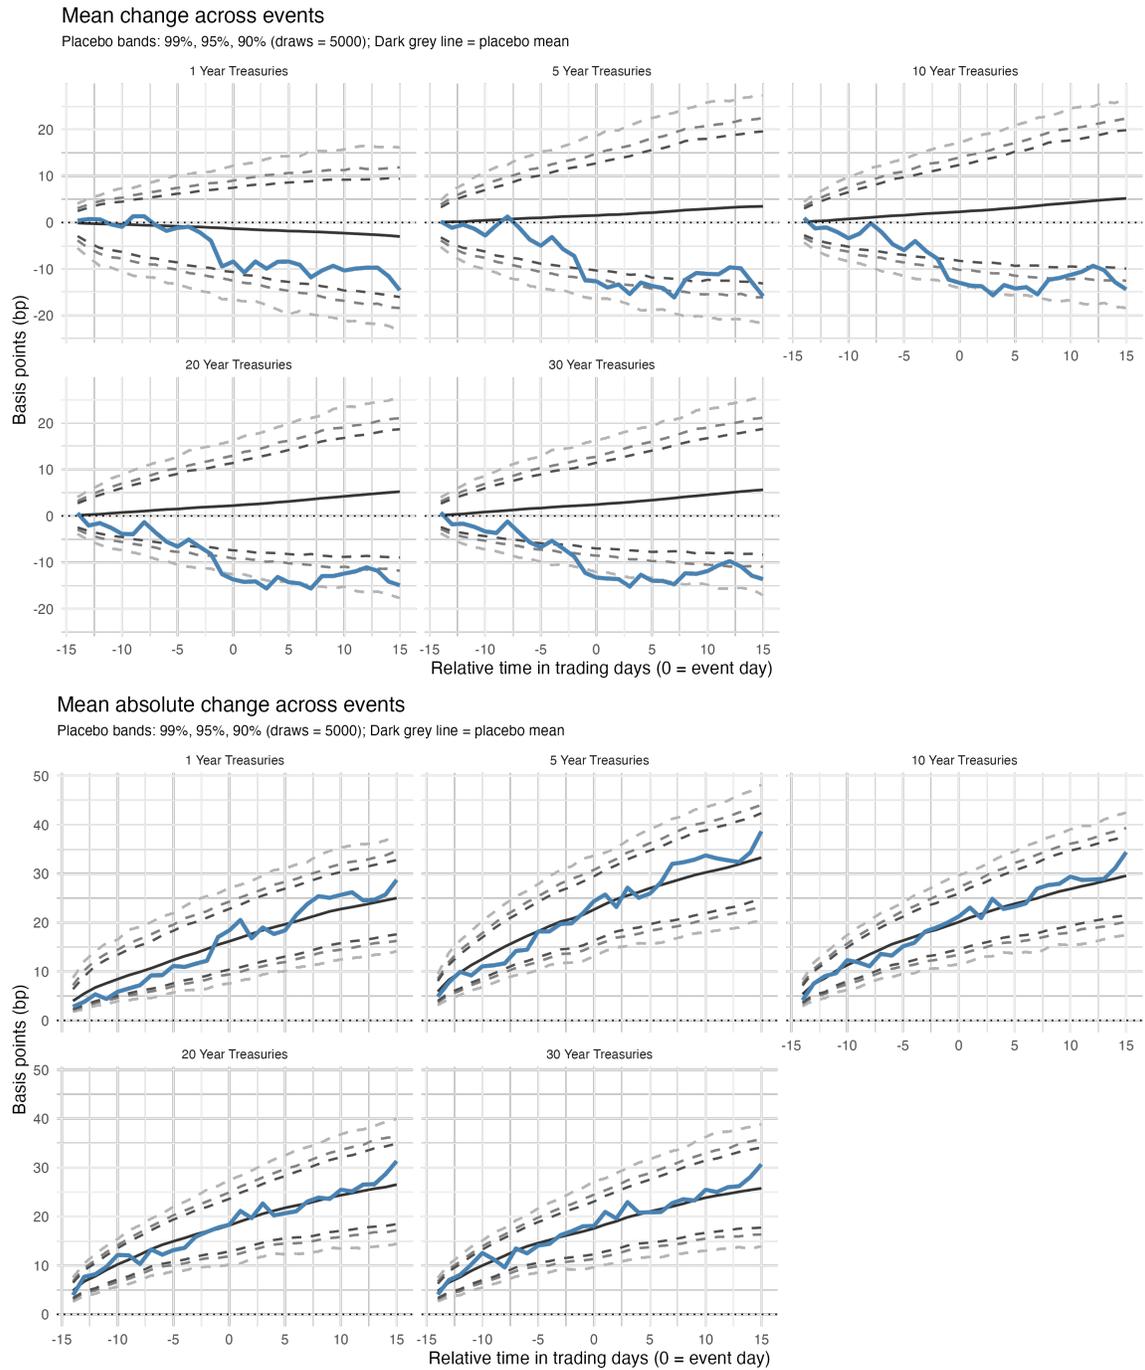


Figure 2: Mean and mean absolute change in yields (relative to fifteen trading days before event) for constant-maturity US Treasury Bonds. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

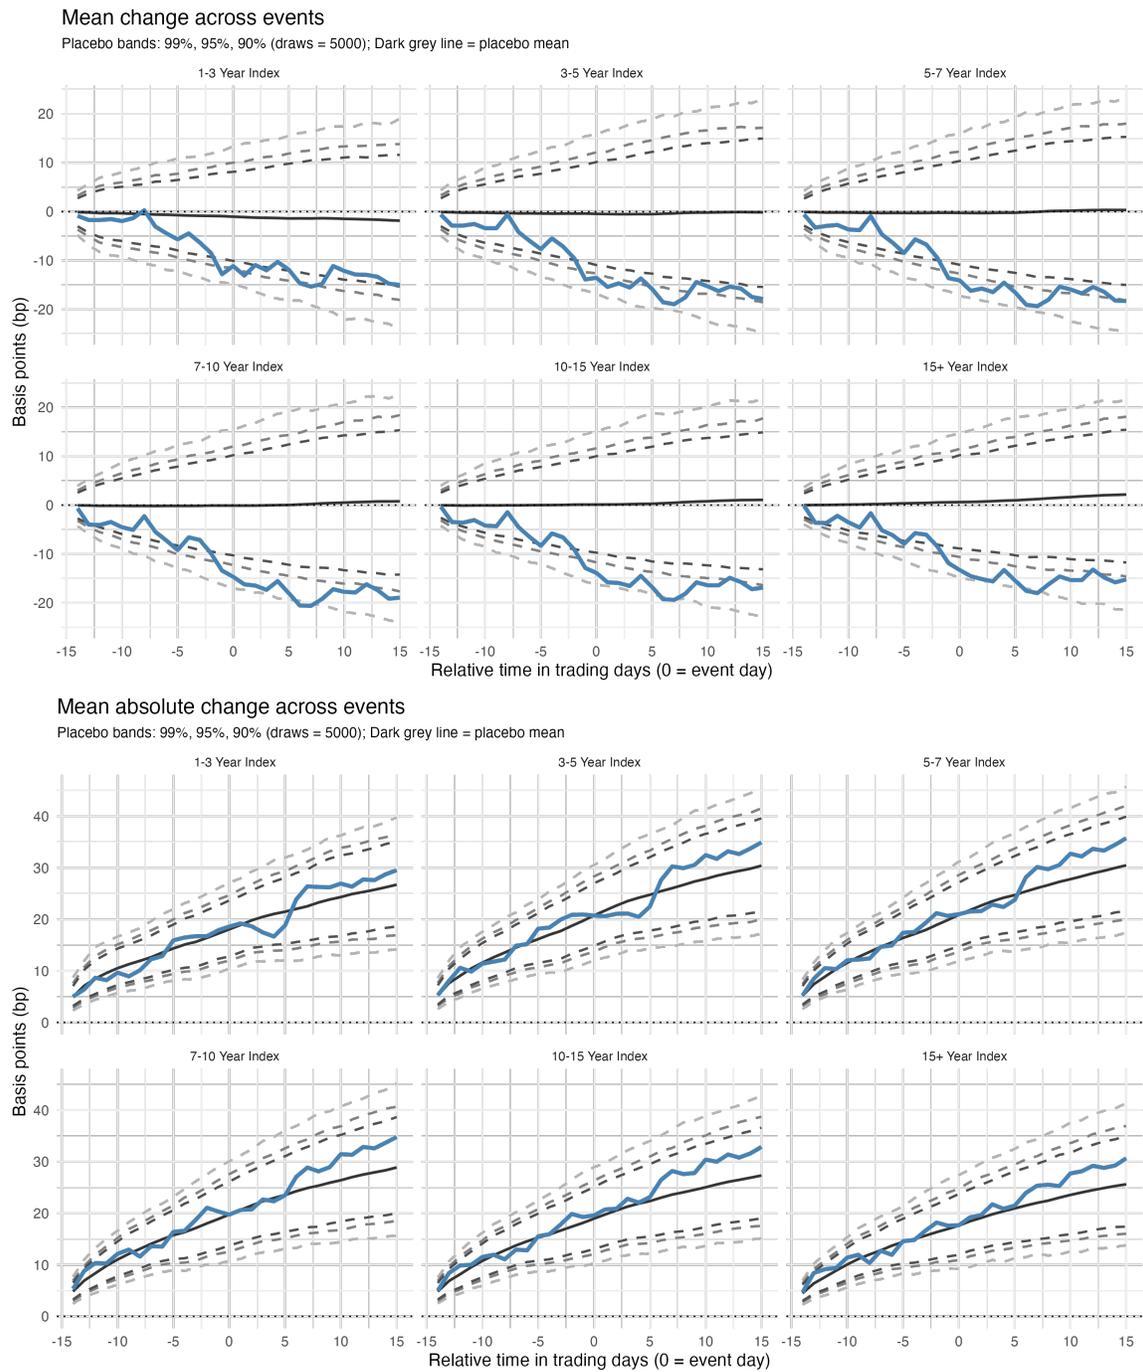


Figure 3: Mean and mean absolute change in yields (relative to fifteen trading days before event) for corporate bond indices. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

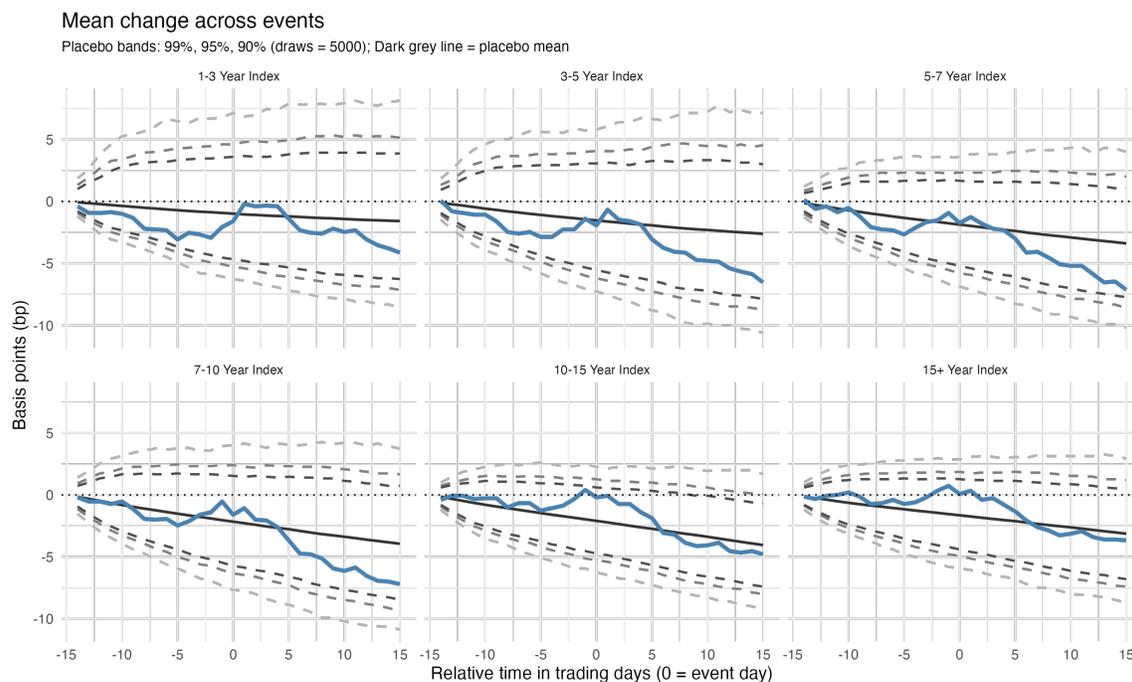


Figure 4: Mean change in option-adjusted spreads (relative to fifteen trading days before event) for corporate bond indices. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

where these declines persist for at least three weeks after the release date.

**Corporate Bond Spreads** Figures 2 and 3 show a significant decline in both Treasury and corporate yields around AI model release dates, especially at the long end of the yield curve. These observations raise an immediate question: is there any change in the corporate yield above and beyond the change in Treasury yields? Put differently, is the impact in corporate yield curve fully explained by the change in Treasury yields, or does AI news have an additional impact on corporate bond yields?

To answer this question, Figure 4 plots the event study for the ICE BofA Option-Adjusted Spread index, where spreads are measured relative to US Treasuries (Ice Data Indices, LLC, 2025b). Comparing the observed changes in spreads to placebo bands we find no statistically significant changes in spreads.<sup>10</sup>

<sup>10</sup>We also find no significant effects on spreads when we look at corporate bond indices broken out by credit rating, though for the sake of brevity we do not report those results.

**Exchange Rates** Given our findings on bond yields, one might wonder whether AI model releases are leading to international capital flows. To provide some evidence on this point, in Appendix B we plot the event study for a broad trade-weighted US dollar exchange rate index around our model release dates (Board of Governors of the Federal Reserve System, US, 2025d). We find that AI model releases are associated with a marginally statistically significant depreciation of the dollar, which starts a few days before the model release and persists through 15 trading days after. These declines are more gradual than the bond yield changes we find above, but appear consistent with e.g. a depreciation of the dollar following a drop in interest rates.

**Robustness Checks** We report several robustness checks in Appendix C. To ensure that our results are not driven by changes in inflation expectations, Appendix C.1 reports versions of our main results using TIPS instead of regular Treasury bonds, and the empirical patterns are the same. To explore sensitivity to the choice of placebo dates and the possibility that AI model releases could be correlated with other market-relevant events, Appendix C.2 reports versions of the mean change plots in Figures 2 which use one of (i) FOMC meetings (ii) major tech-firm annual events (iii) major tech firm earnings releases, (iv) CPI release dates, (v) jobs report release dates (vi) retail sales release dates, and (vii) Treasury auction dates for 10, 20, and 30 year bonds as the source of our placebo dates. None of the resulting placebo series explain what we find for the AI model release dates. Finally, Appendix C.3 reports versions of our main results instead considering medians and median absolute changes. Our findings there are similar to those reported above, so the choice between means and medians does not appear to be an important factor in our results.

## 5 Interpretation

Our empirical analysis shows that major AI model releases led to a reduction in long-term bond yields. As discussed in Section 2, viewed through the lens of the complete-market, representative agent model, falling yields on the risk-free asset imply that the expected future marginal utility of consumption is rising, because expected future consumption is falling, uncertainty is increasing, or the date  $T$  after which asset holdings are irrelevant is believed to be shifting further into the future (or is less likely to arrive at all).

A natural follow-up question is how much investor beliefs must have updated to rationalize observed changes in yields. Providing a quantitative answer to this question unfortu-

nately requires imposing additional assumptions beyond those in Section 2. Since this interpretive exercise nevertheless appears worthwhile, in the remainder of this section we consider a more restrictive version of our model which we use to quantitatively interpret our findings.

**A Simplified Model** As discussed in Section 2, the assumption of complete markets implies the existence of a representative agent, so in this section we focus on that agent’s consumption and utility. Following e.g. Jones (2024) we assume that the representative agent has CRRA flow utility from consumption,  $u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}$ .<sup>11</sup> Under this assumption, the SDF simplifies to

$$M_{t,t+h} = \beta^h \left( \frac{C_{t+h}}{C_t} \right)^{-\gamma} 1\{t+h \leq T\}.$$

To obtain interpretable expressions for the SDF, we make the further simplifying assumption that there exists a horizon  $k \geq 0$  such that at each date  $t$  the representative agent thinks that for all horizons  $h \geq k \geq 0$  periods in the future, aggregate consumption evolves according to

$$C_{t+h+1} = (1+g)X_{t+h+1}C_{t+h}$$

where  $g$  captures the consumption growth impact of AI and  $\{X_s\}_{s=t+k+1}^{\bar{T}}$  is a stochastic process capturing the non-AI determinants of consumption growth. We make this assumption starting  $k$  periods in the future, rather than immediately, to allow the possibility of richer dynamics in short-term consumption, e.g. because the growth impacts of AI could take some time to “kick in.” Together with CRRA utility, this implies that for  $h \geq k$  the  $h$ -period ahead SDF is

$$M_{t,t+h} = \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \beta^h (1+g)^{-(h-k)\gamma} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} 1\{t+h \leq T\}.$$

We next impose additional restrictions to further simplify expressions of the expected SDF and thus bond yields. We assume that conditional on information available at  $t$  and the event  $t+k \leq T$ , (i)  $\left( \{X_s\}_{s=t+k+1}^{t+h}, \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \right)$ ,  $T$ , and  $g$  are mutually independent (ii)

---

<sup>11</sup>Jones (2024) includes an additional location term in the utility, but this term will be irrelevant for our purposes so we drop it.

$T$  is thought to arrive with probability  $\delta_t$  in each period following  $t+k$ ,

$$\mathbb{P}_t(t+h \leq T | t+k \leq T) = \prod_{s=k+1}^h \mathbb{P}_t(t+s \leq T | t+s-1 \leq T) = (1-\delta_t)^{h-k},$$

and (iii)  $1+g$  is believed to be log-normally distributed,  $\log(1+g) | \mathcal{F}_t, t+k \leq T \sim N(\mu_t, \sigma_t^2)$ . These assumptions are restrictive, and appear unlikely to hold exactly. For instance, one might expect that more effective AI (i.e. AI yielding a higher  $g$ ) would be associated with a closer arrival date for  $T$ . Similarly, if the growth effects of AI may “kick in” strictly before period  $t+k$  then a higher  $g$  should lead to a higher  $C_{t+k}$ . Nevertheless, additional assumptions are needed to quantitatively interpret our results, and those above are the least objectionable assumptions we have thus far found that suffice to yield tractability.

Under these assumptions, let us consider the the period  $t$  forward yield from  $t+k$  to  $t+h$ , i.e. the per-period yield earned by, in period  $t$ , selling a period  $t+k$  zero and buying a period  $t+h$  zero,

$$f_{t+k,t+h} = \left( \frac{y_{t,t+h}^h}{y_{t,t+k}^k} \right)^{\frac{1}{h-k}}.$$

Appendix D shows that under our assumptions (i)-(iii), the log forward can be written as

$$\begin{aligned} \log(f_{t+k,t+h}) &= \frac{1}{h-k} \log \left( \frac{\mathbb{E}_t[M_{t,t+k} | t+k \leq T]}{\mathbb{E}_t[M_{t,t+h} | t+k \leq T]} \right) = \frac{h}{h-k} \log(y_{t,t+h}) - \frac{k}{h-k} \log(y_{t,t+k}) = \\ &= -\log(\beta) - \log(1-\delta_t) + \gamma\mu_t - \frac{\gamma^2}{2}(h-k)\sigma_t^2 - \frac{1}{h-k} \log \left( \mathbb{E}_t \left[ \frac{\left( \frac{C_{t+k}}{C_t} \right)^{-\gamma}}{\left[ \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \right]} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right] \right). \end{aligned}$$

Consequently, if we consider the difference in log forward yields at two dates  $t_- < t < t_+$  we have

$$\begin{aligned} &\log(f_{t_+,t_+,t_++h}) - \log(f_{t_-,t_-,t_++h}) = \\ &= -\log \left( \frac{1-\delta_{t_+}}{1-\delta_{t_-}} \right) + \gamma(\mu_{t_+} - \mu_{t_-}) - \frac{\gamma^2}{2}(h-k)(\sigma_{t_+}^2 - \sigma_{t_-}^2) - \eta_{t_-,t_+,k,h} \end{aligned}$$

where

$$\eta_{t_-,t_+,k,h} = \log \left( \frac{\mathbb{E}_{t_+} \left[ \left( \frac{C_{t_++k}}{C_{t_+}} \right)^{-\gamma} \left( \prod_{s=k+1}^h X_{t_++s} \right)^{-\gamma} \right]}{\mathbb{E}_{t_-} \left[ \left( \frac{C_{t_-+k}}{C_{t_-}} \right)^{-\gamma} \left( \prod_{s=k+1}^h X_{t_-+s} \right)^{-\gamma} \right]} \cdot \frac{\mathbb{E}_t \left[ \left( \frac{C_{t_-+k}}{C_{t_-}} \right)^{-\gamma} \right]}{\mathbb{E}_t \left[ \left( \frac{C_{t_++k}}{C_{t_+}} \right)^{-\gamma} \right]} \right).$$

Let us again consider our set of event dates  $t \in \mathcal{T}$ , and for each  $t$  consider  $t_+ = t + s$  and  $t_- = t - b$ .<sup>12</sup> Let  $\mathcal{A}$  denote the set of all dates  $t$  such that  $t_+$  and  $t_-$  are both in the sample. We assume that for all  $h \geq k$  the residuals  $\eta_{t_-,t_+,k,h}$  have approximately the same mean across our event dates  $\mathcal{T}$  as across  $\mathcal{A}$ ,

$$\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \eta_{t_-,t_+,k,h} \approx \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \eta_{t_-,t_+,k,h}. \quad (6)$$

For instance, if we assumed that  $\eta_{t_-,t_+,k,h}$  were stationary across time conditional on our event dates  $\mathcal{T}$  and regularity conditions held, this would follow from the law of large numbers as  $|\mathcal{T}| \rightarrow \infty$ . Motivated by this assumption, we consider the difference in differences of log forward rates  $\log(f_{t+k,t+h})$  across times  $t \in \mathcal{T}$  and  $t \in \mathcal{A}$ :

$$\begin{aligned} \text{DID}(\log(f_{t+k,t+h}); \mathcal{T}, \mathcal{A}) &\equiv \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \log \left( \frac{f_{t_++k,t_++h}}{f_{t_++k,t_++h}} \right) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \log \left( \frac{f_{t_++k,t_++h}}{f_{t_++k,t_++h}} \right) \approx \quad (7) \\ &\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \left( -\log \left( \frac{1-\delta_{t_+}}{1-\delta_{t_-}} \right) + \gamma(\mu_{t_+} - \mu_{t_-}) - \frac{\gamma^2}{2}(h-k)(\sigma_{t_+}^2 - \sigma_{t_-}^2) \right) - \\ &\frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \left( -\log \left( \frac{1-\delta_{t_+}}{1-\delta_{t_-}} \right) + \gamma(\mu_{t_+} - \mu_{t_-}) - \frac{\gamma^2}{2}(h-k)(\sigma_{t_+}^2 - \sigma_{t_-}^2) \right) \end{aligned}$$

Thus, if we consider the slope of  $\text{DID}(\log(f_{t+k,t+h}))$  with respect to the horizon  $h$ , this approximately recovers the difference in differences for the variance  $\sigma_t^2$ , scaled by  $-\frac{\gamma^2}{2}$ .

$$-\frac{\gamma^2}{2} \text{DID}(\sigma_t^2; \mathcal{T}, \mathcal{A}) \equiv -\frac{\gamma^2}{2} \left( \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\sigma_{t_+}^2 - \sigma_{t_-}^2) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} (\sigma_{t_+}^2 - \sigma_{t_-}^2) \right). \quad (8)$$

If we think that dates in  $\mathcal{A} \setminus \mathcal{T}$  have little news relevant to the growth impacts of AI, we might expect the second term to be small relative to the first. However, our event dates are also included in the second term and we moreover do not want to rule out the

<sup>12</sup>Thus,  $t_+$  and  $t_-$  are implicitly functions of  $t$ , though we suppress this dependence for readability.

possibility that AI-relevant news arrives at dates outside of  $\mathcal{T}$ . Hence, we focus on the difference-in-differences interpretation.

Similarly, the intercept measures the difference in differences for expected log growth, scaled by  $\gamma$ , less the difference in differences in log probability that  $T$  does not arrive in a given year,

$$\gamma \text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) - \text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) \quad (9)$$

for

$$\begin{aligned} \text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) &\equiv \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\mu_{t_+} - \mu_{t_-}) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} (\mu_{t_+} - \mu_{t_-}) \\ \text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) &\equiv \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \log\left(\frac{1 - \delta_{t_+}}{1 - \delta_{t_-}}\right) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \log\left(\frac{1 - \delta_{t_+}}{1 - \delta_{t_-}}\right). \end{aligned}$$

**Taking the Model to the Data** The simplified model discussed above predicts the behavior of yields on risk-free zero coupon bonds, so to take these predictions to the data, we use daily Treasury yield curves from FRED (Board of Governors of the Federal Reserve System, US, 2025a), which are based on a three-factor term structure model due to Kim and Wright (2005). These data cover maturities up to 10 years.

To apply the above results, we must choose a horizon  $k$  from after which to start considering forward yields. To guide this choice, Figure 5 plots the difference in differences in one period-ahead log forward yields,  $\text{DID}(\log(f_{t+h, t+h+1}); \mathcal{T}, \mathcal{A})$  for  $h \in \{0, \dots, 9\}$ . Equation (7) implies that for  $h \geq k$  this curve should be approximately linear in  $h$ . This does not appear to hold exactly in our data, but for  $h \geq 4$  it appears a good approximation. Motivated by this finding, for the remainder of our analysis we take  $k=4$ .

After selecting the initial horizon  $k=4$ , we regress  $\text{DID}(\log(f_{t+k, t+h}); \mathcal{T}, \mathcal{A})$  for horizons  $h \in \{5, \dots, 10\}$  on the difference  $h-k$  relative to the initial horizon. This yields a slope of 0.0016 log points (corresponding to the 64.8th percentile of the placebo distribution) and an intercept of approximately -0.132 log points (corresponding to the 4.3rd percentile of the placebo distribution).<sup>13</sup> Our finding that the slope of the yield curve is not substantially changing around model release dates is consistent with our finding in Section 4 above that the yield impacts of model releases are quite similar for the various bond maturities above 5 years. Thus, it appears that the changes we observe around event dates are driven by shifts in the level of the forward curve, rather than the slope.

To further interpret these results through the lens of the simplified model developed

---

<sup>13</sup>Our estimates qualitatively similar for other  $k \in \{4, \dots, 9\}$ .

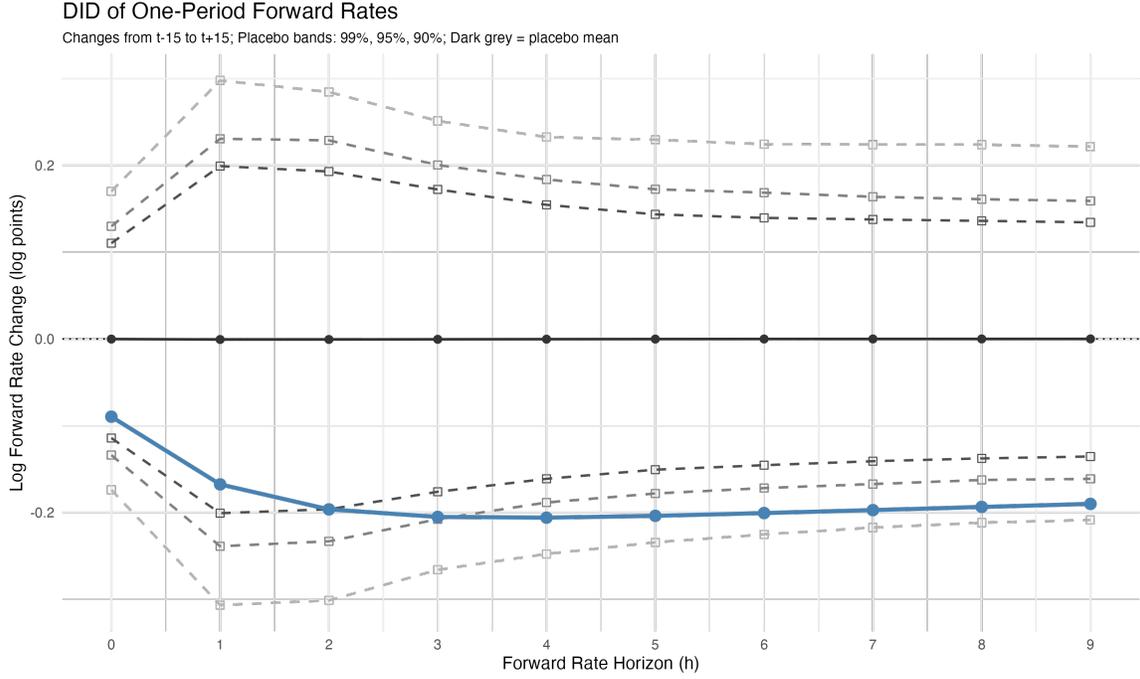


Figure 5: Difference in differences for one-period ahead log forward rate starting at period  $h$ ,  $\text{DID}(\log(f_{t+h,t+h+1}); \mathcal{T}, \mathcal{A})$ . Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

above, we separately consider the interpretation of the slope and intercept.

**Interpreting the Slope** First considering the slope (8) of the forward curve difference in differences (7) with respect to the maturity difference  $h-k$ , recall that the slope coefficient estimates the scaled average variance change. Thus, our estimated DID coefficient for the variance of  $\log(1+g_t)$  around model releases, relative to other dates in the sample, is

$$\text{DID}(\widehat{\sigma_t^2}; \mathcal{T}, \mathcal{A}) = -\frac{3.2}{\gamma^2} \cdot 10^{-5},$$

for  $\gamma$  the CRRA coefficient of the representative agent. Hence, the simplified model considered in this section suggests that consumption growth uncertainty actually *fell* slightly on average around model release dates relative to the average day in our sample.

These estimates are small, and are not statistically different from zero according to our placebo distribution, even using generous thresholds for statistical significance.<sup>14</sup> This

<sup>14</sup>To interpret the magnitude of our estimated variance reduction, note that it is equivalent to, on the average event date, removing a noise component from  $\log(1+g_t)$  with standard deviation equal to  $0.57/\gamma$

finding of little evidence for growth uncertainty changes around our event dates is consistent with our finding in Section 4 that there does not appear to be a clear trend in yield changes across 10, 20, and 30 year Treasuries. That said, given our limited sample size we do not have much power to detect small slope changes (the 5th and 95th percentiles of our placebo distribution correspond to slopes of approximately  $\pm 1.1 \cdot 10^{-3}$ , respectively).

Overall, our simplified model suggests that, if anything, consumption growth uncertainty may have slightly fallen around the model release dates we study, though our estimates are imprecise. Nevertheless, we have sufficient evidence to conclude that, through the lens of our simplified model, changes in consumption growth uncertainty do not explain the yield decreases we observe around AI event dates.

**Interpreting the Intercept** We next turn to the intercept (9) in the forward curve difference in differences (7). Recall that under our simplified model this term captures two forces: changes in the anticipated arrival rate  $\delta_t$  of  $T$  (where a closer expected arrival for  $T$  increases yields) and changes in the mean  $\mu_t$  of the log growth rate  $\log(1+g_t)$  (where a higher value of  $\mu_t$  again increases yields).

If our estimated intercept were due entirely to a change in beliefs about  $T$ , the model implies that the average model release in our sample led to a roughly 0.132 percentage point increase in  $\log\left(\frac{1-\delta_{t+}}{1-\delta_{t-}}\right)$  relative to the average in the sample

$$\text{DID}(\log(1-\delta_t); \mathcal{T}, \mathcal{A}) \approx 0.132\%.$$

If we assume  $\delta_t$  is close to zero, it follows that  $\text{DID}(\delta_t; \mathcal{T}, \mathcal{A}) \approx -0.132\%$ , so the average AI event in our sample is associated with a 0.132 percentage point larger reduction in  $\delta_{t+}$  than the average date in the sample. Cumulated over the 15 model releases in our analysis sample, this corresponds to a 1.98 percentage point decrease in the probability that  $T$  arrives in a given year, which seems like a large effect.

If observed changes in yields were instead due entirely to changes in consumption growth expectations, the model implies that the average model release in our sample led to an approximately  $0.132/\gamma$  percentage point larger decrease in  $\mu_t$  than the average date in the sample,

$$\text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) = -\frac{0.132\%}{\gamma},$$

---

percentage points. While this is not a negligible uncertainty reduction for e.g.  $\gamma \in [1, 5]$ , it is difficult to interpret given its statistical insignificance.

for  $\gamma$  the CRRA coefficient of the representative agent. If we assume that  $\sigma_t^2 = \text{Var}_t(\log(1+g))$  is small for all  $t$ , and further assume that  $\mu_t$  is close to zero, this implies that

$$\text{DID}(\mathbb{E}_t[g]; \mathcal{T}, \mathcal{A}) \approx -\frac{0.132\%}{\gamma},$$

so the average model release in our sample implies a  $0.132/\gamma$  percentage point reduction in expected consumption growth, relative to the average date in the sample. Thus, under  $\gamma=1$  (i.e. log utility) our results imply a 0.132 percentage point, or 13.2 basis point, drop in expected consumption growth, while under  $\gamma=2$  they imply a roughly 0.066 percentage point drop, and under  $\gamma=5$  they imply an approximately 0.026 percentage point drop. Even at the lower end, these are meaningful effects for an event that occurred 15 times in our analysis sample.

Of course, it could be that beliefs about both  $T$  and  $g$  update in response to AI model releases. To explore this broader range of possible interpretations, Figure 6 depicts the  $\text{DID}(\mu_t; \mathcal{T}, \mathcal{A})$  and  $-\text{DID}(\log(1-\delta_t); \mathcal{T}, \mathcal{A}) \approx \text{DID}(\delta_t; \mathcal{T}, \mathcal{A})$  combinations compatible with an intercept (9) of -0.132 percentage points, for different levels of CRRA parameter  $\gamma$ . There is downward-sloping relationship between the implied effects on the arrival rate of  $T$  and  $g$ : the larger the decrease in the arrival rate of  $T$ , the more positive the growth effects which rationalize observed yield changes, and vice versa.

Overall, our simplified model implies that the changes in bond yields we observe around AI model release dates are primarily driven by some combination of decreases in growth expectations (i.e.  $\mu_t$ ) and decreases in the perceived arrival rate of  $T$  (i.e.  $\delta_t$ ) rather than changes in growth uncertainty (i.e.  $\sigma_t$ ).

## 6 Discussion

We have found evidence of economically and statistically large declines in long-term bond yields around major AI model releases. Viewed through the lens of the simple asset pricing model in Section 2, this suggests that investors are updating their beliefs towards some combination of (i) lower expected consumption growth (ii) higher uncertainty about future consumption or (iii) a lower probability of extreme “doom” or “bliss” scenarios. We can roughly quantify the extent of belief updating under the additional assumptions laid out in Section 5, and find that the implied magnitudes are substantial, with (i) and/or (iii) playing a much more important role than (ii) in explaining our results.

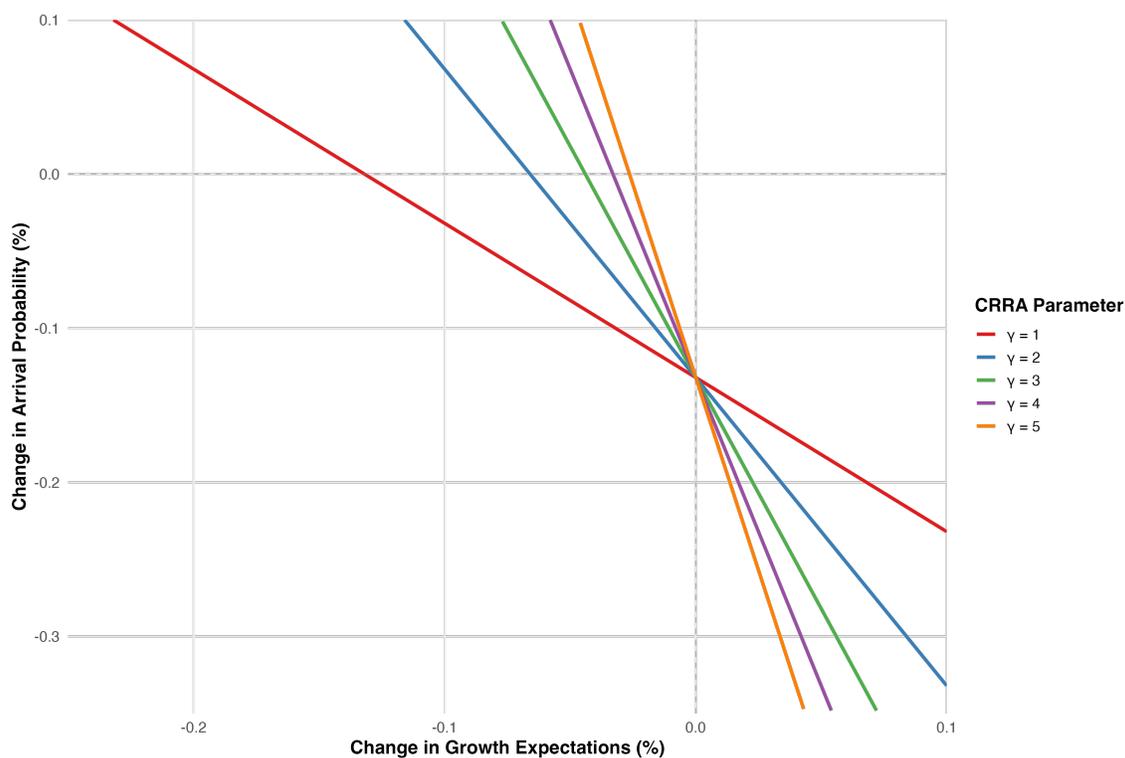


Figure 6: Values for  $\text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) \approx \text{DID}(\mathbb{E}_t[g]; \mathcal{T}, \mathcal{A})$  and  $-\text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) \approx \text{DID}(\delta_t; \mathcal{T}, \mathcal{A})$  compatible with an intercept value (9) equal to -0.132 percentage points under the simplified model and different levels of CRRA parameter  $\gamma$ .

These conclusions are subject to several important caveats. First, it may be that none of the bonds we consider is a reasonable proxy for a risk-free asset. Second, updates to investor beliefs around the model release dates we study could be non-representative of overall investor beliefs about AI, and third, the simple complete-market, representative agent model might imply a misleading interpretation of market responses. We discuss each possibility in turn.

On the first possibility, it is very plausible that investors do not think US Treasuries are approximately risk-free. Treasuries are subject to inflation risk, and potentially to default risk given the large and growing budget deficits run by the US government.<sup>15</sup> If market participants think there is a non-trivial probability of a US default over the period we study it could be that news about AI raises expected future tax revenue, and thus lowers Treasury

<sup>15</sup>The TIPS series covered in Appendix C.1 are designed to reduce inflation risk, but remains subject to default risk.

yields by lowering the embedded risk premium rather than by changing growth expectations.

Our data can provide some limited evidence on this possibility. One would expect that if the US government were to default this might increase the risk of many companies defaulting as well, so the mere fact that corporate bond yields also fall around AI model releases does not rule out this explanation. However, to the extent that not all US companies would necessarily default if the US government did, we would expect a drop in US government default risk to increase the spread between corporate bond yields and Treasury yields. To examine this possibility, recall that Figure 4 plots the event study for the ICE BofA Option-Adjusted Spread index, and shows no statistically or economically significant increase in spreads. While this does not fully rule out that the effects we observe could be driven by changes in risk premia on US Treasuries, the risk premia on corporate bonds would need to move essentially in tandem.

A second explanation for our results could be that, while we are obtaining valid estimates for the impact of AI news at the dates we study, our event dates are in some sense non-representative. That is, it could be that the net effect of investor beliefs about AI has been to increase bond yields over the 2023-4 period, but that the particular event dates we've selected saw updates in the opposite direction. While we cannot rule this out, it is not clear to us why it would be the case: we include all dates from a well-defined universe (all major model release dates from a set of prominent AI firms), and it is not clear to us why the impact of information arriving at these dates should be directionally different, in aggregate, than that of AI information arriving at other dates in the same two year window.

A third possibility is that while we are accurately capturing market responses to AI news, the model in Section 2 implies a misleading interpretation of these results. In particular, there are a wide variety of reasons why reality may deviate from the fully-optimizing, complete market benchmark, including market incompleteness, a wide array of market frictions and constraints, behavioral deviations from rationality and optimization, and many more. To explain our results, however, an alternative story needs to explain economically large and apparently persistent yield changes in one of the deepest financial markets in the world. This suggests that alternative explanations could themselves be of considerable interest.

## References

- Acemoglu, Daron and Pascual Restrepo**, “The wrong kind of AI? Artificial intelligence and the future of labour demand,” *Cambridge Journal of Regions, Economy and Society*, 2020, *13* (1), 25–35.
- **and Todd Lensman**, “Regulating transformative technologies,” *American Economic Review: Insights*, 2024, *6* (3), 359–376.
- , **David Autor, Jonathon Hazell, and Pascual Restrepo**, “Artificial intelligence and jobs: Evidence from online vacancies,” *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson**, “Artificial intelligence, firm growth, and product innovation,” *Journal of Financial Economics*, 2024, *151*, 103745.
- , — , **Alex X He, and James Hodson**, “Firm investments in artificial intelligence technologies and changes in workforce composition,” 2023, *31325*.
- Board of Governors of the Federal Reserve System, US**, “Fitted Yield on 1-10 Year Zero Coupon Bond, Quoted on an Investment Basis, Inflation-Indexed [THREEFY1,THREEFY2,THREEFY3,THREEFY4,THREEFY5,THREEFY6,THREEFY7,THREEFY8,THREEFY9,THREEFY10],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “Market Yield on U.S. Treasury Securities at 1,5,10,20,30-Year Constant Maturity, Quoted on an Investment Basis [DGS1, DGS5, DGS10, DGS20, DGS30],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “Market Yield on U.S. Treasury Securities at 1,5,10,20,30-Year Constant Maturity, Quoted on an Investment Basis, Inflation-Indexed [DFII5, DFII10, DFII20, DFII30],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “Nominal Broad U.S. Dollar Index [DTWEXBGS], retrieved from FRED, Federal Reserve Bank of St. Louis,” <https://fred.stlouisfed.org/series/DTWEXBGS> 2025.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson**, “Artificial intelligence and the modern productivity paradox,” *The economics of artificial intelligence: An agenda*, 2019, *23* (2019), 23–57.

- Chiang, Wei-Lin, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica**, “Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference,” 2024.
- Chow, Trevor, Basil Halperin, and J Zachary Mazlish**, “Transformative AI, existential risk, and real interest rates,” Technical Report, Working Paper 2024.
- Duffie, Darrell**, *Dynamic asset pricing theory*, Princeton University Press, 2010.
- Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska**, “The labor impact of generative AI on firm values,” *Available at SSRN*, 2024, 4436627.
- Erdil, Ege, Andrei Potlogea, Tamay Besiroglu, Edu Roldan, Anson Ho, Jaime Sevilla, Matthew Barnett, Matej Vrzla, and Robert Sandler**, “GATE: An Integrated Assessment Model for AI Automation,” *arXiv preprint arXiv:2503.04941*, 2025.
- Gil, Hamilton Galindo**, “Asset Pricing Impacts of AI in a Model with Heterogeneous Risk Aversion,” Technical Report, Working Paper 2024.
- Gordon, Robert**, *The Rise and Fall of American Growth*, Princeton University Press, 2016.
- Ice Data Indices, LLC**, “ICE BofA 1-3, 3-5, 7-10, 10-15, 15+ Year US Corporate Index Effective Yield, [BAMLC1A0C13YEY, BAMLC2A0C35YEY, BAMLC3A0C57YEY, BAMLC4A0C710YEY, BAMLC7A0C1015YEY, BAMLC8A0C15PYEY],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “ICE BofA 1-3, 3-5, 7-10, 10-15, 15+ Year US Corporate Index Option-Adjusted Spread, [BAMLC1A0C13YY, BAMLC2A0C35YY, BAMLC3A0C57YEY, BAMLC4A0C710YY, BAMLC7A0C1015YY, BAMLC8A0C15PYEY],” FRED, Federal Reserve Bank of St. Louis 2025.
- Jones, Charles I**, “The ai dilemma: Growth versus existential risk,” *American Economic Review: Insights*, 2024, 6 (4), 575–590.
- Kim, Don H. and Jonathan H. Wright**, “An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates,” Technical Report, Working Paper 2005.

**Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean**, “AI 2027,” 2025.

**Korinek, Anton and Donghyun Suh**, “Scenarios for the Transition to AGI,” Technical Report, National Bureau of Economic Research 2024.

**Mehra, Rajnish**, “Consumption-based asset pricing models,” *Annu. Rev. Financ. Econ.*, 2012, 4 (1), 385–409.

**Mollick, Ethan**, “Google’s Gemini Advanced: Tasting Notes and Implications,” 2024.

**Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa**, “Crises and recoveries in an empirical model of consumption disasters,” *American Economic Journal: Macroeconomics*, 2013, 5 (3), 35–74.

**Trammell, Philip and Anton Korinek**, “Economic growth under transformative AI,” Technical Report, National Bureau of Economic Research 2023.

**Webb, Michael**, “The impact of artificial intelligence on the labor market,” *Available at SSRN 3482150*, 2019.

# Appendix

## A Alternative Date Series

Figure 7 shows results for US Treasury bonds when we adopt a narrower definition of the “flagship” series for each firm and so exclude the release of OpenAI’s o1 model from our event series. These results are similar to those in Figure 2, but show somewhat larger declines in yields.

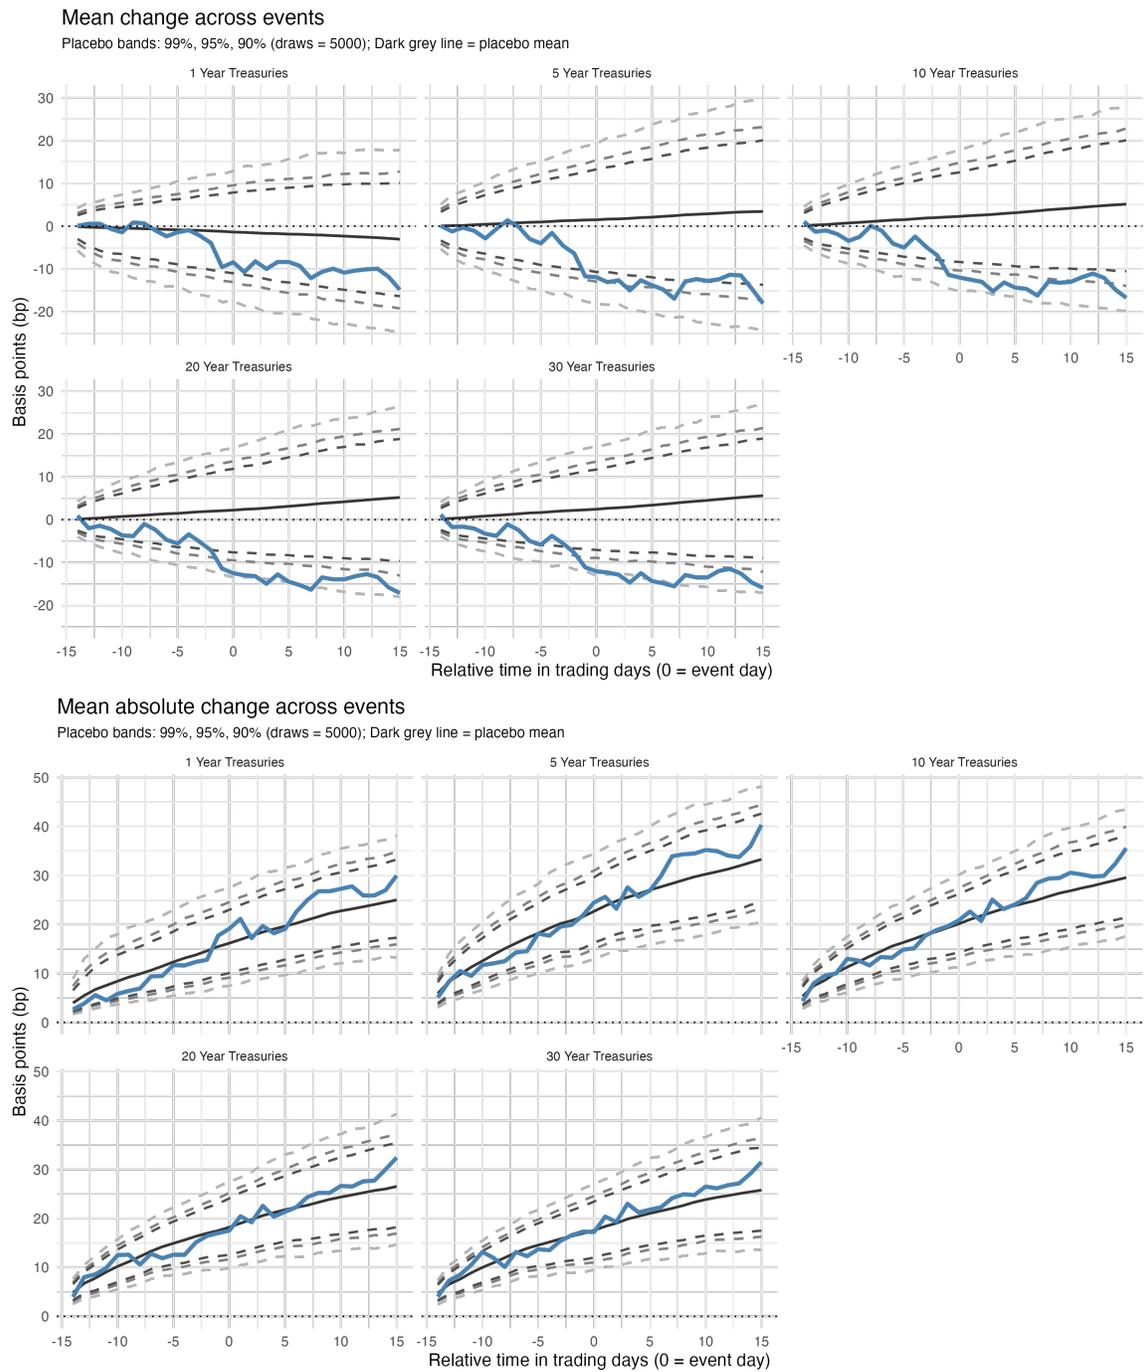


Figure 7: Mean and mean absolute change in yields (relative to fifteen trading days before event) for constant-maturity US Treasury Bonds. Mean taken across AI release events in the 2023 and 2024 calendar years, excluding release of OpenAI’s o1 model. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

## **B Exchange Rate Responses**

Figure 8 shows an event-study for a broad trade-weighted US dollar index from FRED (Board of Governors of the Federal Reserve System, US, 2025d). As this plot shows, on average the model releases in our sample saw a weakening of the dollar, consistent with lower demand for the dollar following the fall in interest rates estimated in the main text. These declines are significant relative to our placebo distribution, and again persist through 15 trading days after the model release.

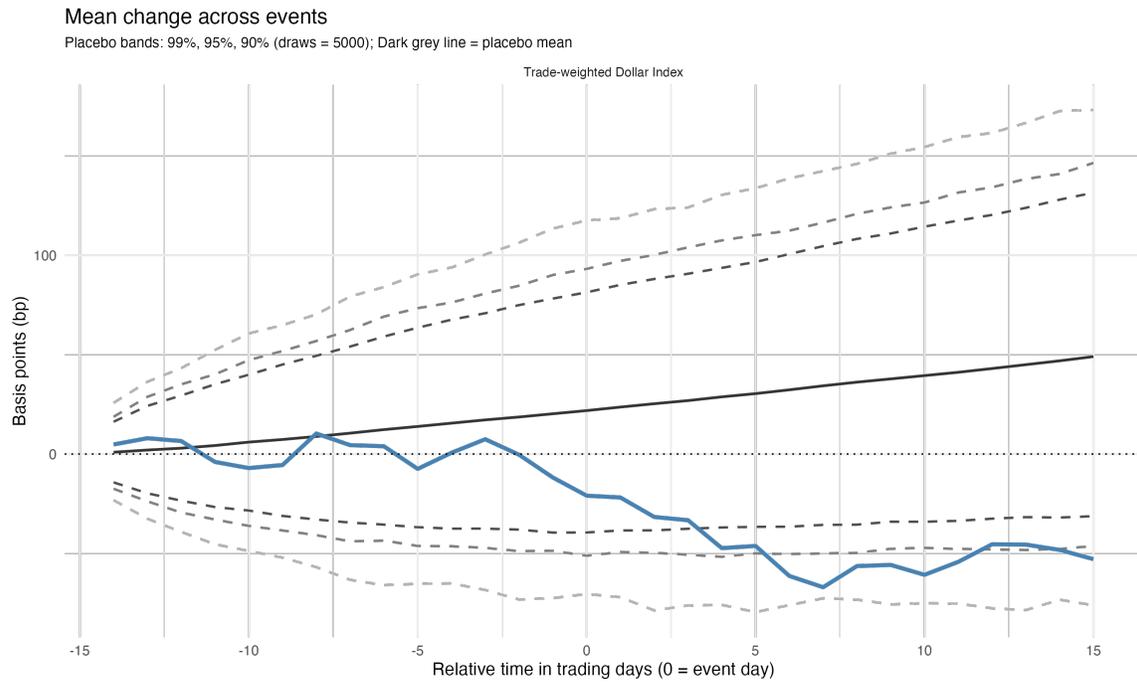


Figure 8: Mean change in trade-weighted US Dollar index (relative to fifteen trading days before event). Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

## C Robustness Checks

### C.1 TIPS

Here we extend the results in Section 4 to an alternative bond series, **Treasury Inflation-Protected Securities (TIPS)**. We use constant-maturity real yields for TIPS with maturities of 5, 10, 20, and 30 years, again taken from FRED (Board of Governors of the Federal Reserve System, US, 2025c).

Table 4 report two-sided p-values from the exercise explained in Section 4 for TIPS data series, reporting results for both mean and mean absolute changes, and comparing yields either five or fifteen days before and after each model release. It paints a very similar picture to Tables 2 and 3 in the main text.

Table 4: Two-sided p-values based on constant-maturity TIPS yields

Maturity	Mean Change		Mean Absolute Change	
	$\pm 5$ days	$\pm 15$ days	$\pm 5$ days	$\pm 15$ days
5 Year	0.354	0.069*	0.298	0.788
10 Year	0.208	0.021**	0.296	0.942
20 Year	0.184	0.014**	0.394	0.810
30 Year	0.206	0.014**	0.421	0.623

*Notes:* The “Mean Change” columns of consider the mean change in yields across event dates, while the “Mean Absolute Change” columns consider mean absolute changes. For each statistic, we compare yields 5 or 15 days before and after each model release (in  $\pm 5$  the  $\pm 15$  columns, respectively). P-values are computed based on drawing placebo event dates 5000 times (uniformly at random from trading days in the sample with sufficient window on either side) and comparing resulting placebo distributions to observed changes around AI model releases. \*\* (\*) denotes statistical significance at the 5% (10%) level.

Figure 9 plots, for the TIPS series and each horizon  $s \in \{-14, 15\}$ , the mean and mean absolute change in yields (relative to 15 days before the event,  $t-15$ ) across the observed AI model releases. For comparison, at each horizon we also plot the mean of the placebo distribution, and bands which contain, 90%, 95%, and 99% of the placebo draws pointwise at each horizon, with equal mass assigned to the two tails. The results are again parallel to Figures 2 and 3 in the main text and confirm the robustness of our empirical results.

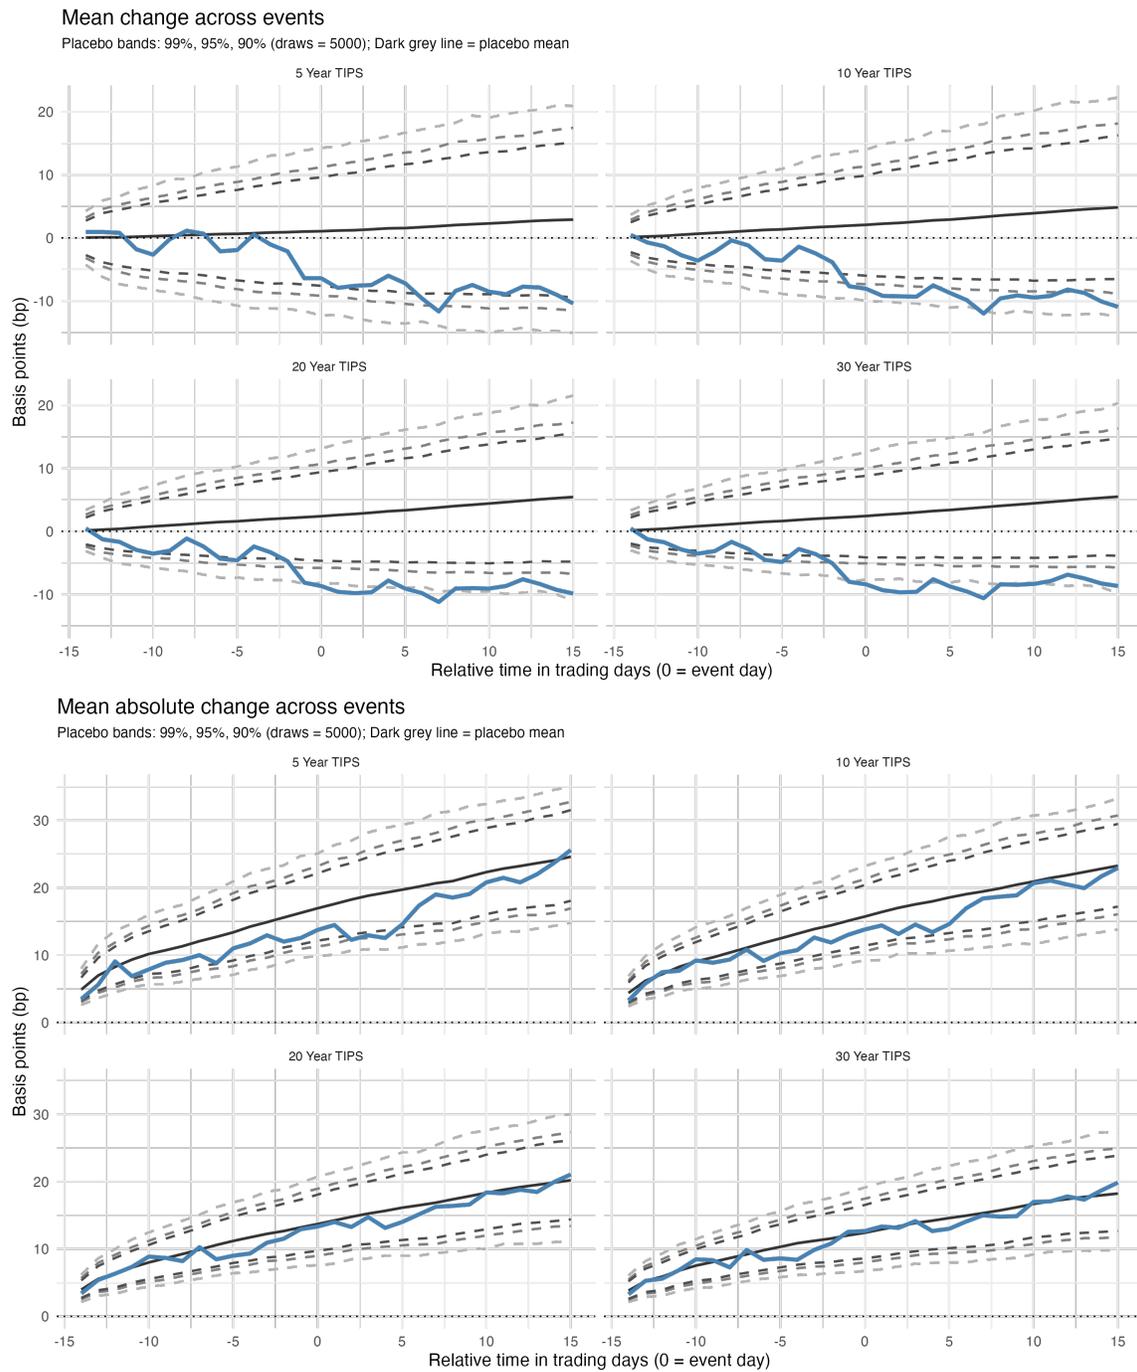


Figure 9: Mean and mean absolute change in yields (relative to fifteen trading days before event) for constant-maturity inflation-protected US Treasury Bonds. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

Table 5: Major Tech Conference Dates

<b>Date</b>	<b>Company</b>	<b>Conference</b>
<i>2023 Conferences</i>		
6/9/2023	Apple	Worldwide Developer Conference
5/10/2023	Google	Google I/O
9/28/2023	Meta	Meta Connect
12/1/2023	Amazon	AWS re:Invent
5/25/2023	Microsoft	Microsoft Build
3/23/2023	Nvidia	Nvidia GTC
3/1/2023	Tesla	Tesla AI Day
<i>2024 Conferences</i>		
6/14/2024	Apple	Worldwide Developer Conference
5/14/2024	Google	Google I/O
9/26/2024	Meta	Meta Connect
12/6/2024	Amazon	AWS re:Invent
5/23/2024	Microsoft	Microsoft Build
3/21/2024	Nvidia	Nvidia GTC
4/22/2024	Tesla	Tesla AI Day

*Notes:* This table presents the major annual tech conferences from the “magnificent seven” companies used in our event study analysis. For multi-day conferences, we use the final day.

## C.2 Alternative Placebos

In this appendix, we explore the extent to which other date series could potentially explain the results we find around AI model release dates by employing these alternative series as placebo dates. Specifically, we consider placebo series corresponding to (i) FOMC meetings (ii) major conferences series held by the so-called “magnificent seven” major tech firms (listed in Table 5) (iii) earnings release dates for the same seven firms, (iv) Consumer Price Index release dates, (v) Bureau of Labor Statistics Employment Situation Series release dates (vi) US Census Bureau Advance Monthly Sales for Retail and Food Services release dates and (vii) US Treasury auctions for 10, 20, or 30 year bonds.

For the first two series, once we limit attention to our analysis window of calendar years 2023 and 2024, we have no more events than we do AI model releases. Hence, we lack sufficient dates to generate placebo distributions, and simply plot the mean change in yields, based on these alternative date series, against the mean change in yields across the model release dates used in our analysis. To limit the number of plots we focus on mean yield changes for US Treasuries, and plot the results in Figure 10. Across both series we see that these events are not associated with a decline in long-maturity treasury yields

of the sort observed around the model release dates.

For the remaining placebo date series we have sufficiently many events to produce non-trivial placebo distributions. There is considerable clustering in earnings announcements, so given our long event windows we group earnings announcements within a single calendar week into a single event, with the event date recorded as the final earnings announcement in that week. For each placebo series we draw dates at random (without replacement) from the corresponding date lists, and obtain the placebo distributions shown in Figures 11-13. Once again, we find that behavior of US Treasury yields around AI model release dates, particularly for long-maturity bonds, appears statistically and economically different from behavior around these alternative date series.

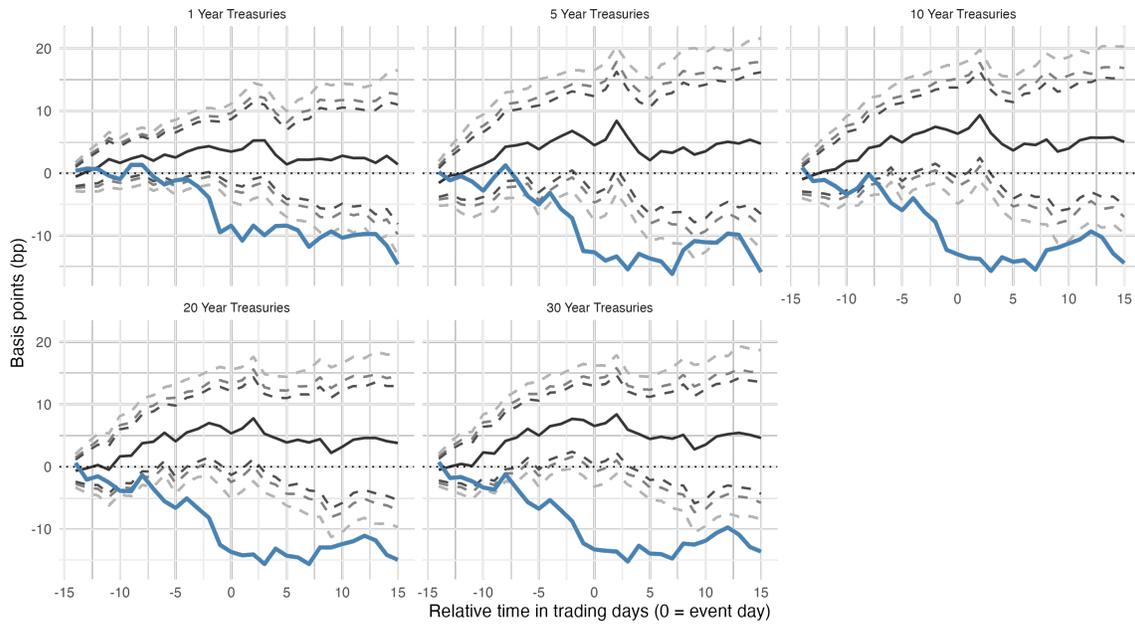
It is worth emphasizing that statistical insignificance of our AI events relative to these alternative placebo series would *not* suffice to imply that these alternative event series explain our results, since these alternative event dates would also need to be sufficiently correlated with our events (though there is a question of how to correctly measure this correlation). On the other hand, statistical significance for our results relative to these alternative placebos implies that even if our event dates were a subset of these alternative dates (which they are not, for any of these series), the alternative series would not explain our results.



Figure 10: Mean change in yields (relative to fifteen trading days before event) for US treasuries. Mean taken across AI release events in the 2023 and 2024 calendar years. First panel placebo takes mean over FOMC meetings in the 2023 and 2024 calendar years. Second panel placebo takes mean over “magnificent seven” tech company conference dates.

### Mean change across events - Tech Earnings Dates

Placebo bands: 99%, 95%, 90% (draws = 5000); Dark grey line = placebo mean - Placebo from tech dates only



### Mean change across events - CPI Dates

Placebo bands: 99%, 95%, 90% (draws = 5000); Dark grey line = placebo mean - Placebo from CPI only

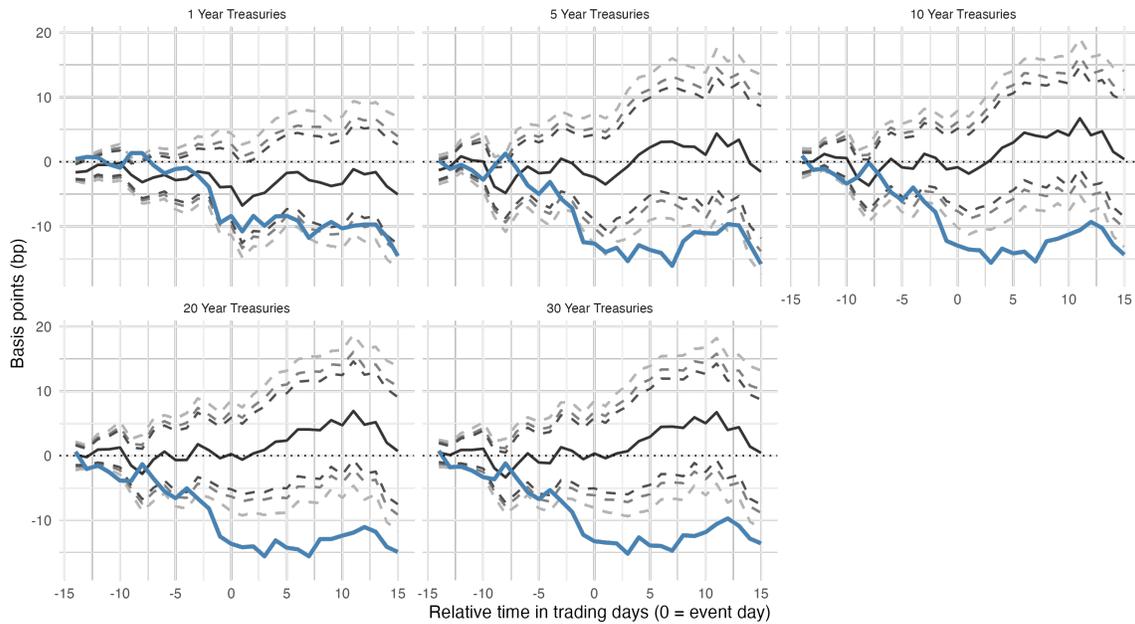


Figure 11: Mean change in yields (relative to fifteen trading days before event) for US Treasuries. Mean taken across AI release events in the 2023 and 2024 calendar years. First panel placebo takes mean over “magnificent seven” earnings announcement dates. Second panel placebo takes mean over CPI release dates.

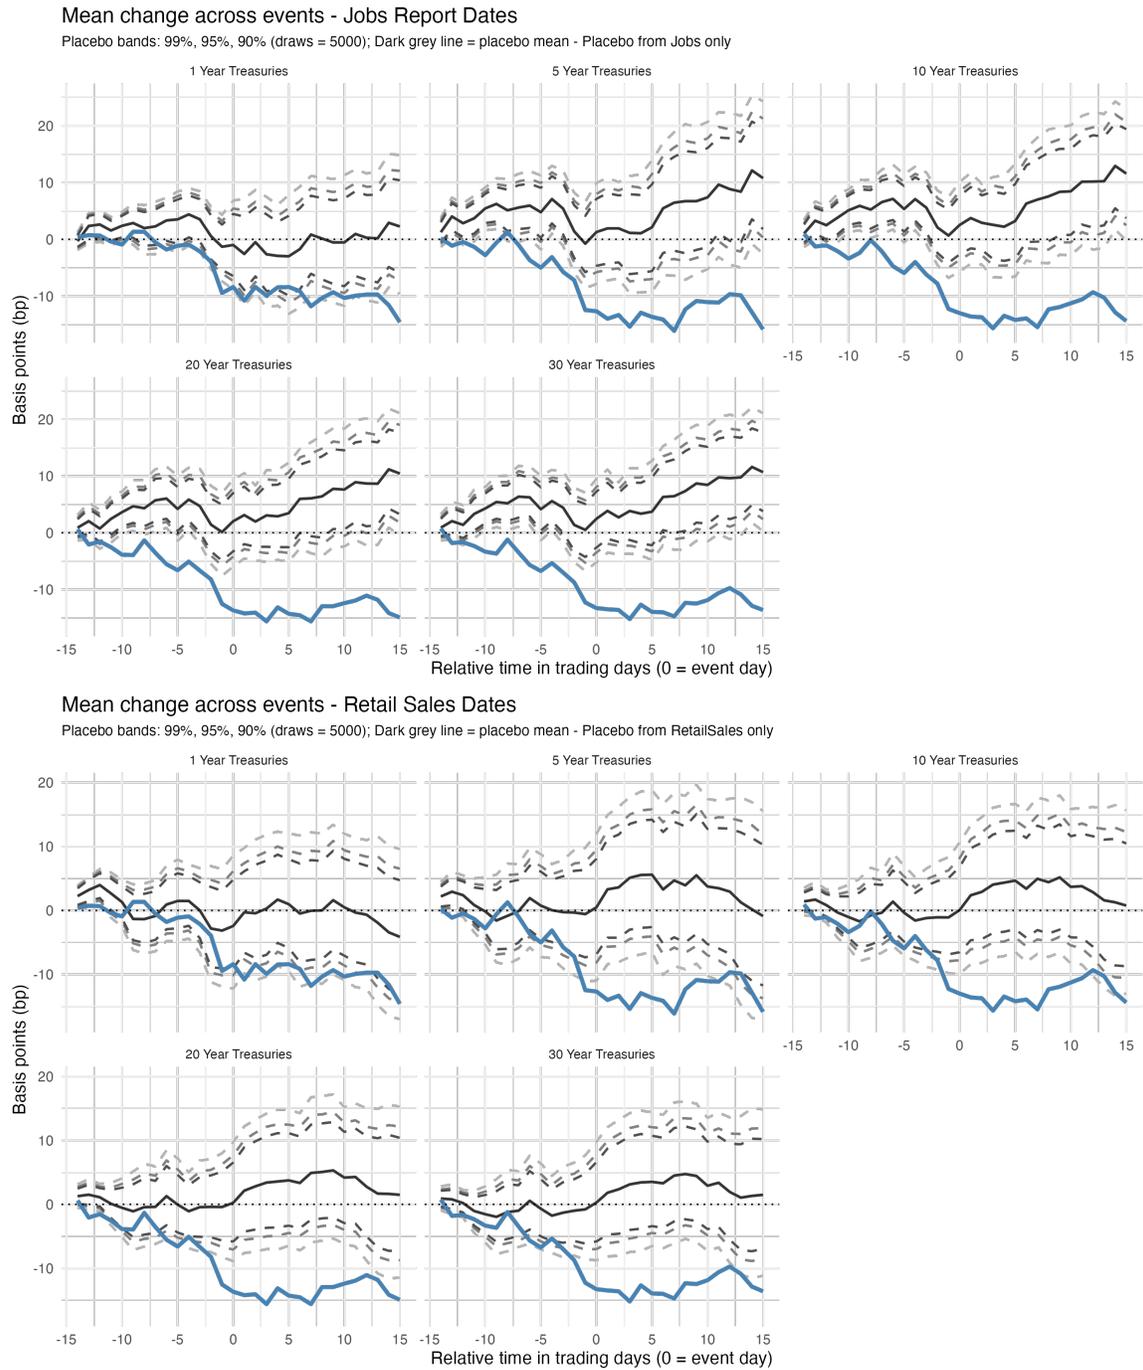


Figure 12: Mean change in yields (relative to fifteen trading days before event) for US Treasuries. Mean taken across AI release events in the 2023 and 2024 calendar years. First panel placebo takes mean over BLS employment situation report release dates. Second panel placebo takes mean over retail sales report release dates.

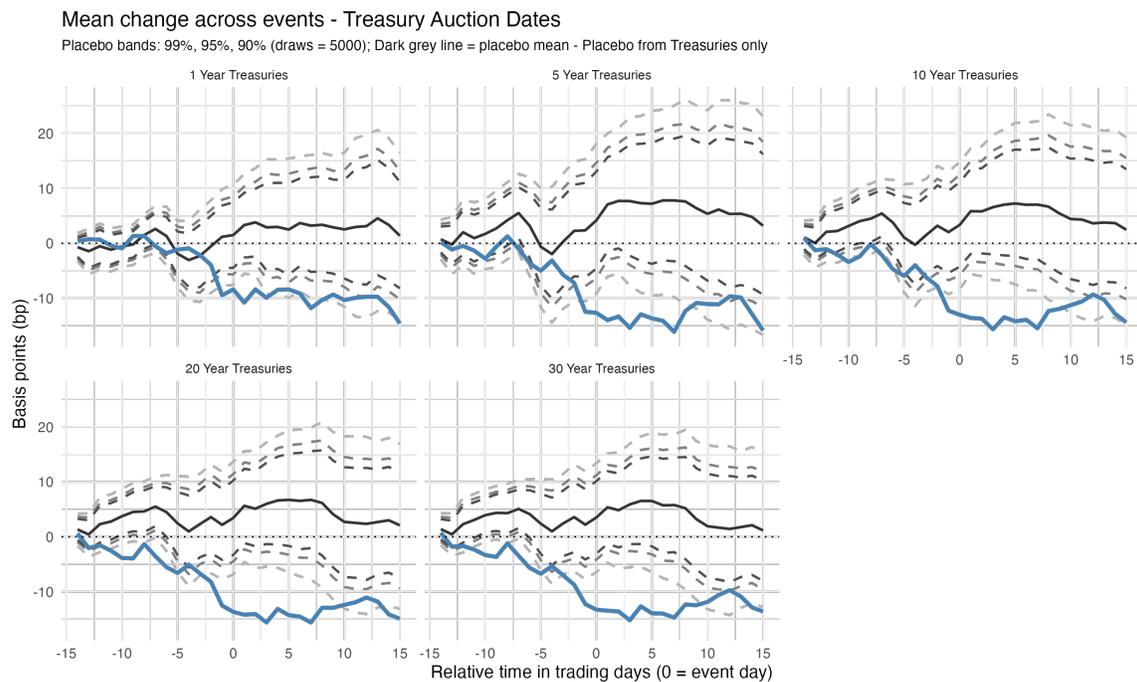


Figure 13: Mean change in yields (relative to fifteen trading days before event) for US Treasuries. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo takes mean over auction dates for 10, 20, and 30 year US Treasuries.

### C.3 Median and Median Absolute Changes

In the main text we focus on results for mean and mean absolute changes in yields around major AI model releases. To complement those results, Figures 14 and 15 report versions of Figure 2 and 3 instead using the median change

$$\text{MedianChange}_s = \text{Med}_{t \in \mathcal{T}}(\Delta y_{t,s})$$

and the median absolute change

$$\text{MedianAbsChange}_s = \text{Med}_{t \in \mathcal{T}}(|\Delta y_{t,s}|),$$

motivated by the fact that medians are more robust to outliers. The findings are qualitatively similar to those based on means, with strong declines in yields, particularly at longer maturity, after the model release dates. Relative to the results based on means, we see even less evidence of declines more than 5 days before the model release date.

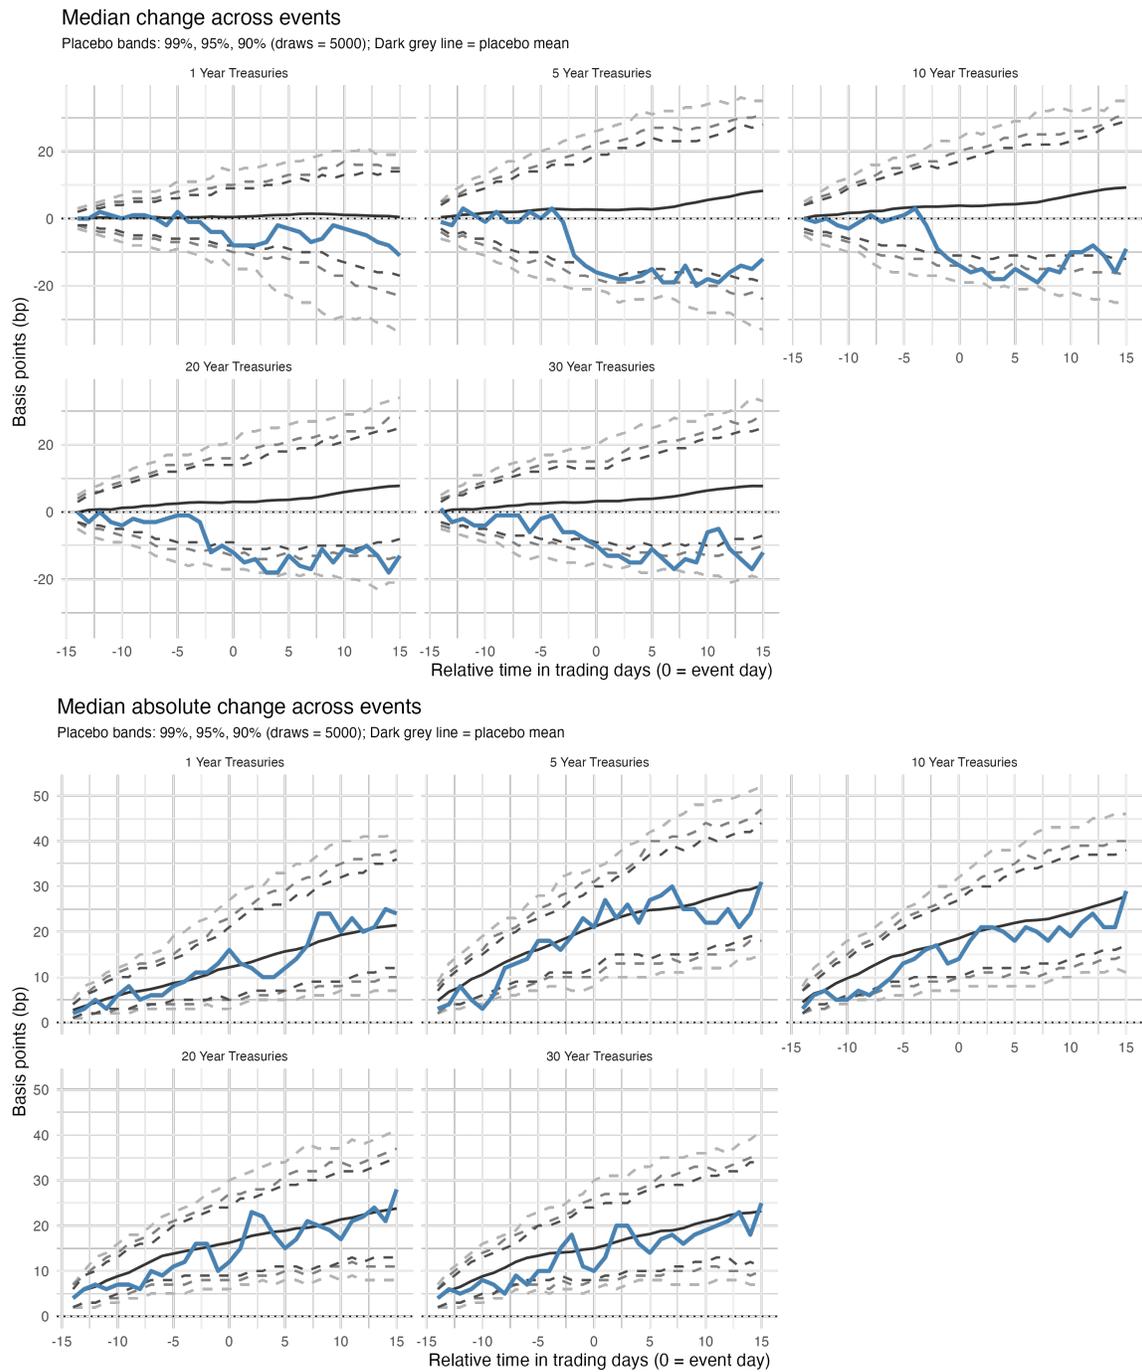


Figure 14: Median and median absolute change in yields (relative to fifteen trading days before event) for constant-maturity US Treasury Bonds. Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

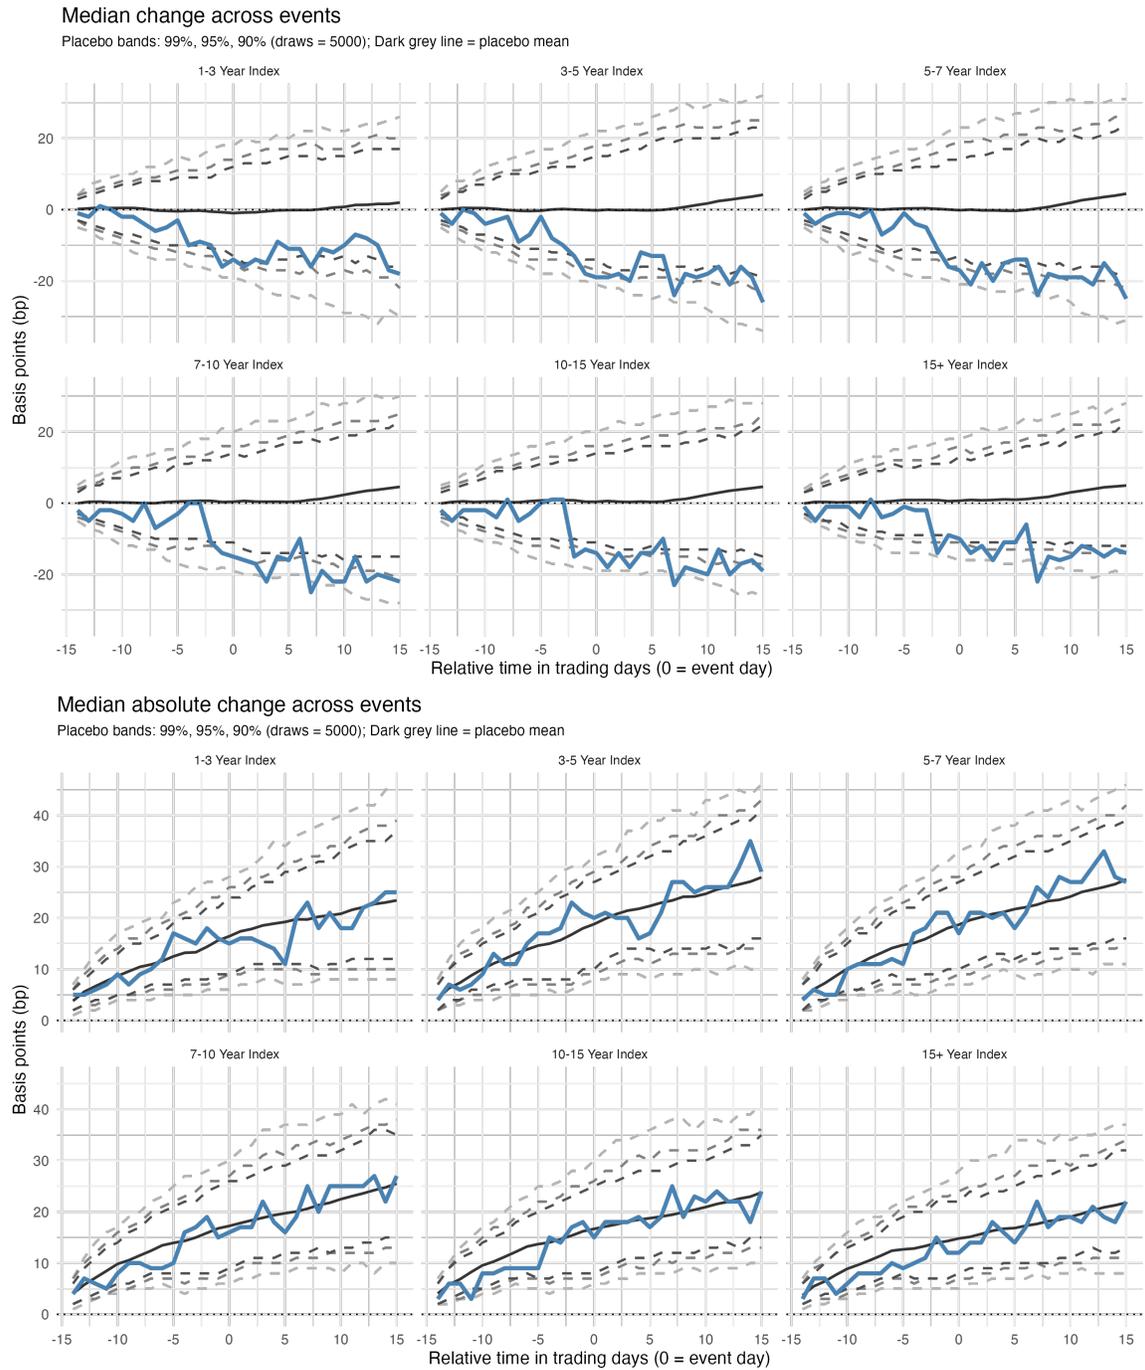


Figure 15: Median and median absolute change in yields (relative to fifteen trading days before event) for corporate bond index. Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

## D Additional Derivations for Section 5

Note that by the law of iterated expectations and our conditional independence assumptions

$$\begin{aligned}\mathbb{E}_t[M_{t,t+h}] &= \mathbb{P}_t(t+k \leq T) \mathbb{E}_t[M_{t,t+h} | t+k \leq T] = \\ &= \mathbb{P}_t(t+k \leq T) \mathbb{E}_t \left[ \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \beta^h (1+g)^{-(h-k)\gamma} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \mathbb{1}_{\{t+h \leq T\}} \middle| t+k \leq T \right] = \\ &= \mathbb{P}_t(t+k \leq T) \beta^h \mathbb{E}_t \left[ \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \middle| t+k \leq T \right] \times \\ &= \mathbb{E}_t \left[ (1+g)^{-(h-k)\gamma} \middle| t+k \leq T \right] \mathbb{P}_t(t+h \leq T | t+k \leq T).\end{aligned}$$

By the formula for the mean of a lognormal, however,

$$\mathbb{E}_t \left[ (1+g)^{-(h-k)\gamma} \middle| t+k \leq T \right] = \exp \left( -\gamma(h-k)\mu_t + \frac{\gamma^2(h-k)^2\sigma_t^2}{2} \right),$$

so we can re-write the expected  $h$ -period ahead SDF as

$$\begin{aligned}\mathbb{E}_t[M_{t,t+h}] &= \\ &= \beta^h (1-\delta_t)^{(h-k)} \exp \left( -\gamma(h-k)\mu_t + \frac{\gamma^2(h-k)^2\sigma_t^2}{2} \right) \mathbb{P}_t(t+k \leq T) \mathbb{E}_t \left[ \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right].\end{aligned}$$

Consequently, if we consider the difference of log expected SDFs,  $\log(\mathbb{E}_t[M_{t,t+h}]) - \log(\mathbb{E}_t[M_{t,t+k}])$  we get

$$\begin{aligned}(h-k)\log(\beta) + (h-k)\log(1-\delta_t) - \gamma(h-k)\mu_t + \frac{\gamma^2}{2}(h-k)^2\sigma_t^2 \\ + \log \left( \frac{E_t \left[ \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right]}{E_t \left[ \left( \frac{C_{t+k}}{C_t} \right)^{-\gamma} \right]} \right).\end{aligned}$$

Note, however, that since  $y_{t,t+h} = \mathbb{E}_t[M_{t,t+h}]^{-\frac{1}{h}}$ , we have

$$\log(f_{t+k,t+h}) = \frac{\log(\mathbb{E}_t[M_{t,t+k}]) - \log(\mathbb{E}_t[M_{t,t+h}])}{h-k} =$$

$$-\log(\beta) - \log(1 - \delta_t) + \gamma\mu_t - \frac{\gamma^2}{2}(h-k)\sigma_t^2 - \frac{1}{h-k} \log \left( E_t \left[ \frac{\left(\frac{C_{t+k}}{C_t}\right)^{-\gamma}}{E_t \left[ \left(\frac{C_{t+k}}{C_t}\right)^{-\gamma} \right]} \left( \prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right] \right),$$

as stated in the main text.