Internet Appendix

Kris Shen

August 26, 2026

This Internet Appendix accompanies the paper "AI Valuations: Bubble or Fundamentals?" Sections, tables, and figures numbered with an `IA` prefix refer to this appendix. Plain numbers refer to the main paper, which is available at https://krisshen-finance.github.io/research/.

Internet Appendix Contents

IA1 AI Exposure and Summary Statistics

IA1.1 AI Vocabulary

To construct an *Artificial Intelligence Vocabulary (AIV)*, I study a corpus of Web of Science texts on the subject of AI. More specifically, I collect 207,717 academic paper titles from Web of Science via ProQuest. Following the literature (see, e.g., Engle et al., 2020), I extract the unigrams and bigrams that are directly related to *artificial intelligence*, and compute their TF-IDF scores. I exclude terms that are ambiguous. Figure IA1.1 presents word clouds that summarize the terms extracted from academic paper titles in the Web of Science. The larger the term size, the more frequent the term appears in the corpus (measured in TF-IDF score). The years range from 2009 to 2024. Panel (a) shows most frequent bigrams. Panel (b) shows most frequent unigrams. Panel (a) shows that most terms are largely related to machine learning and deep learning topics, which is consistent with the fact that these topics are at the core of AI. Other than those, I also see a lot of data-related terms like data mining and big data. Panel (b) is largely consistent with Panel (a) but extracts the unigrams that are also focused on the similar topics in machine learning and deep learning. Overall, these terms, including both unigrams and bigrams, largely capture the main intuitive terms that we could think of in the topics of AI. Here, I am not claiming these are the best keywords that are representative to AI. However, these represents well the terms that are most studied in academic research for this given period from 2009 to 2024.

For robustness, I construct the AIV using a larger and more general sample of 744,044 Web of Science academic articles or books that are directly related to: *artificial intelligence, machine learning, natural language processing, and computer vision.* These are the terms used in Babina et al. (2024a) when they study the systematic risk of AI through the labor channel. Appendix Figure IA1.2 shows the word cloud of bigrams. The top bigrams are consistent with the Figure IA1.1 Panel (a) in capturing mostly machine learning and deep learning topics.

I select the top 50 unigrams and top 50 bigrams as my final AIV in Figure IA1.1 for tractability. Specifying a different number, such as 25 or 100, would alter the quantitative outcomes, though the main qualitative results would remain consistent. Determining the optimal number of unigrams and bigrams to include in the final AIV is beyond the scope of this paper. Table IA1.1 lists the top 50 bigrams and top 50 unigrams by their TF-IDF scores.

To see whether the firm-level *AI Exposure* measure captures intuitive variation across firms, Figure IA1.3 presents six example firms and their quarterly firm-level AI Exposure measures from 2009 to 2024. AI chipmakers such as Nvidia and AMD, as well as the high-tech firm Apple, have higher AI Exposure on average than firms in more traditional industries, including Chevron in oil and gas, AAR in aviation services, and Pepsi in food and beverages. We can also see that AI Exposure increases more noticeably for Nvidia and AMD in the second half of the sample period, consistent with AI-related attention becoming more prominent in the 2020s. By contrast, the average AI Exposure of the traditional firms does not change much, even in the 2020s.

To further understand the context in which AI is discussed, I construct topic-specific versions of AI Exposure, including opportunity, regulation, and risk, as well as tone-based measures (positive and negative) following Loughran and McDonald (2011). These measures are defined by conditioning on the presence of topic-related terms in the same sentence as AI-related terms. For example, the opportunity-related measure is defined as

$$AIExposure_{i,t}^{Opp} = \frac{1}{U_{i,t} + B_{i,t}} \left(\sum_{u}^{U_{i,t}} (\mathbb{1}[u \in AIV] * \mathbb{1}[u, o \in S]) + \sum_{b}^{B_{i,t}} (\mathbb{1}[b \in AIV] * \mathbb{1}[b, o \in S]) \right), \tag{IA1.1}$$

where $\mathbb{1}[u/b, o \in S]$ indicates that AI-related terms appear in the same sentence as opportunity-related terms. Other topic-specific measures are constructed analogously. The topic dictionaries are based on Sautner et al. (2023) and extended using ChatGPT4, while tone dictionaries follow Loughran and McDonald (2011). Table IA1.3 documents the comparison between the AI Exposure measure and topic-based measures.

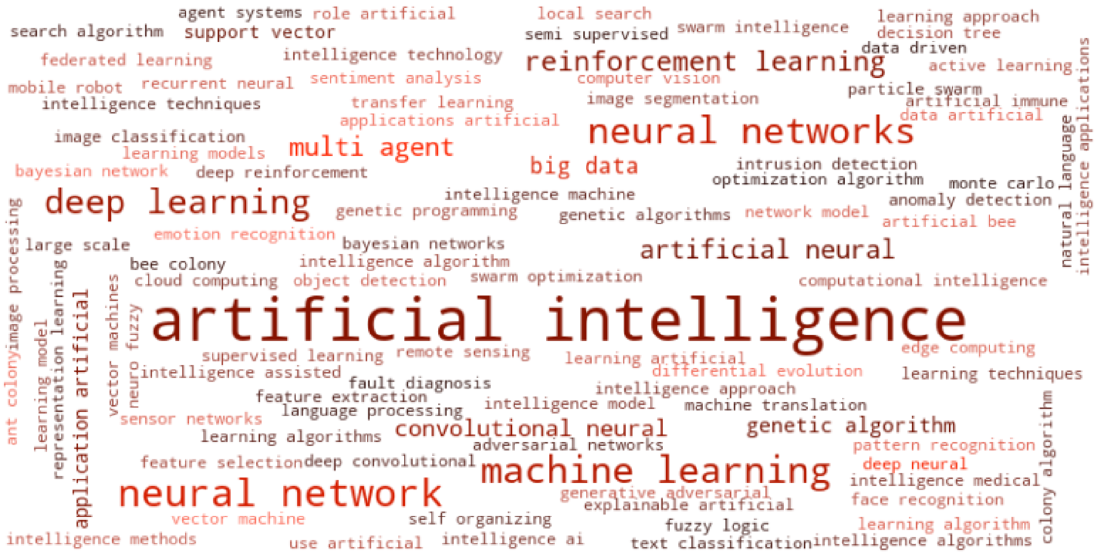

(a) **Artificial Intelligence Bigrams**

(b) **Artificial Intelligence Unigrams**

Figure IA1.1. Artificial Intelligence Vocabulary WordCloud. The figures present word clouds that summarize the *Artificial Intelligence Vocabulary (AIV)* extracted from academic paper titles in the Web of Science. Following the literature, I focus on unigrams and bigrams. The larger the term size, the more frequent the terms appear in the corpus (measured in TF-IDF score). The Web of Science paper titles are downloaded from ProQuest, and cover 207,717 papers specifically related to *artificial intelligence*. The years range from 2009 to 2024. Panel (a) shows most frequent bigrams. Panel (b) shows most frequent unigrams.

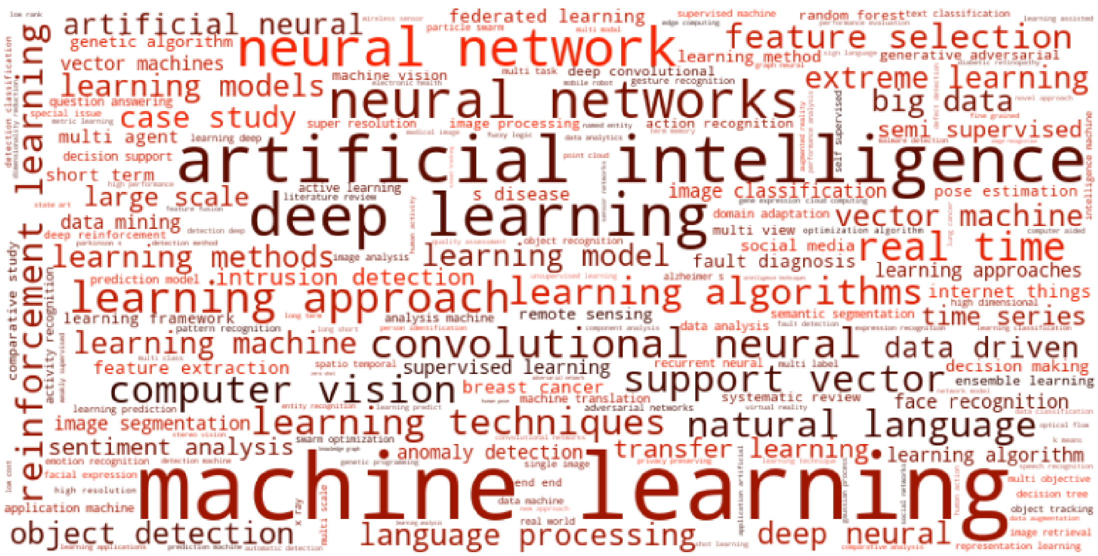

Figure IA1.2. Alternative Artificial Intelligence Vocabulary WordCloud. The figures present word clouds that summarize the bigrams extracted from 744,044 academic paper titles in the Web of Science that are directly related to *artificial intelligence, machine learning, natural language processing, and computer vision*. These are the keywords used in Babina et al. (2024a) when they study the systematic risk of AI through the labor channel. The larger the term size, the more frequent the terms appear in the corpus (measured in TF-IDF score). The Web of Science paper titles are downloaded from ProQuest. The years range from 2009 to 2024.

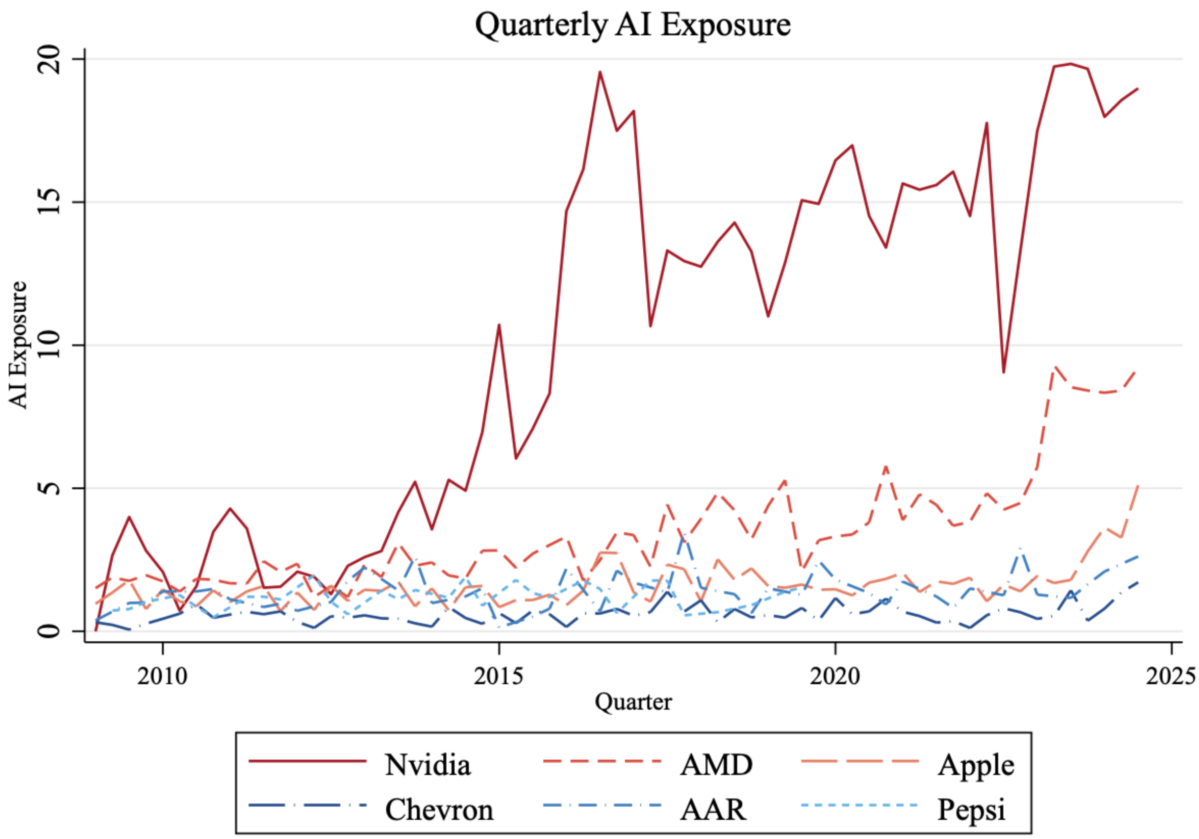

Figure IA1.3. Examples of Firm-Level AI Exposure. This figure presents quarterly firm-level *AI Exposure* for six example firms. The measure captures how much attention managers and investors have devoted to AI-related topics in quarterly conference call transcripts. The sample period is from 2009 to 2024. The examples include AI chipmakers NVIDIA and AMD and the technology firm Apple, as well as firms in more traditional industries: Chevron in oil and gas, AAR in aviation services, and Pepsi in food and beverages.

Table IA1.1
Top Artificial Intelligence Terms from Academic Titles

This table presents the top 50 AI bigrams (Panel A) and unigrams (Panel B) derived from 207,717 academic paper titles related to AI from the Web of Science. Terms are selected based on TF-IDF scores, after excluding ambiguous terms.

Panel A: Top 50 Artificial Intelligence Bigrams

artificial intelligence	neural network	machine learning
neural networks	deep learning	reinforcement learning
artificial neural	multi agent	convolutional neural
big data	genetic algorithm	application artificial
support vector	deep neural	feature selection
learning approach	intelligence techniques	fuzzy logic
natural language	intelligence ai	particle swarm
use artificial	swarm optimization	bee colony
deep reinforcement	fault diagnosis	artificial bee
large scale	intelligence machine	semi supervised
transfer learning	object detection	genetic algorithms
intrusion detection	intelligence technology	face recognition
data driven	agent systems	optimization algorithm
sentiment analysis	computational intelligence	bayesian networks
learning techniques	learning algorithms	vector machine
anomaly detection	colony algorithm	generative adversarial
intelligence approach	neuro fuzzy	

Panel B: Top 50 Artificial Intelligence Unigrams

artificial	intelligence	learning
neural	data	network
model	networks	algorithm
analysis	detection	deep
systems	machine	ai
classification	intelligent	optimization
prediction	recognition	fuzzy
models	agent	algorithms
technology	adaptive	support
dynamic	diagnosis	techniques
feature	selection	automatic
clustering	robot	digital
reinforcement	computing	mining
convolutional	estimation	automated
bayesian	monitoring	programming
distributed	simulation	computer
optimal	training	

IA1.2 AI Exposure by Year and Industry

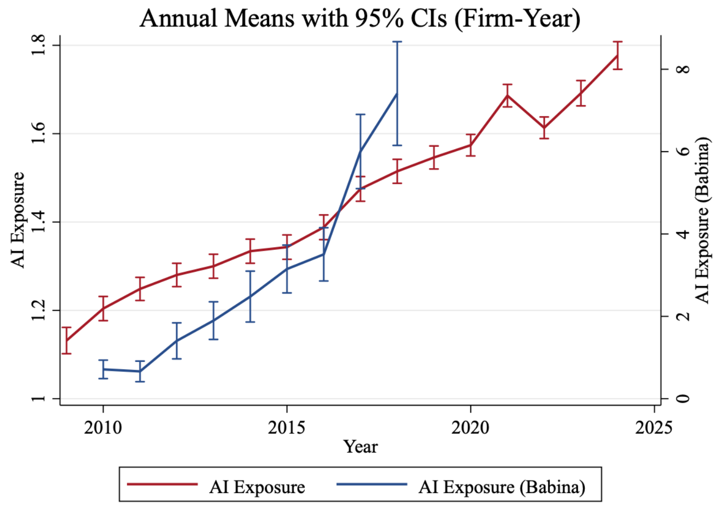

(a) **AI Exposure by Year**

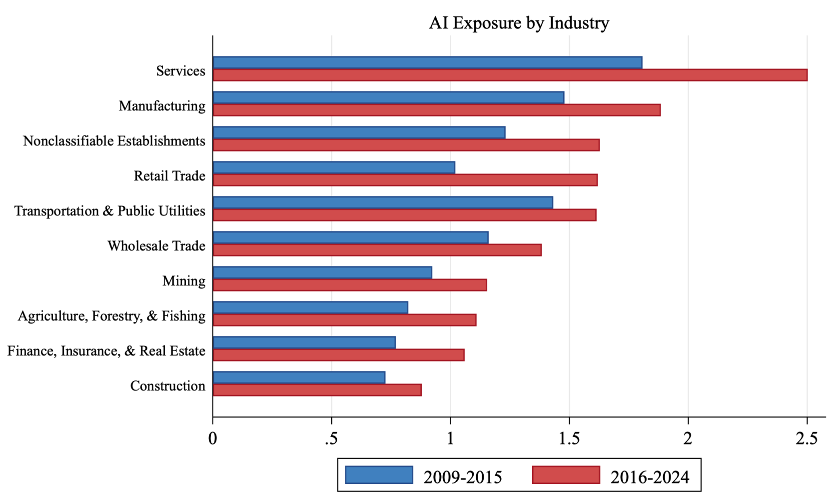

(b) **AI Exposure by Industry**

Figure IA1.4. AI Exposure by Year and Industry. This figure presents annual firm level *AI Exposure* measures by year and by industry. The annual firm level AI Exposure measure is defined as the maximum of the corresponding quarterly measures within each year. Panel (a) reports the mean and 95% confidence interval of AI Exposure by year, where *AI Exposure (Babina)* refers to the firm level fraction of AI related labor from Babina et al. (2024b) during 2010 to 2018. Panel (b) reports the mean of AI Exposure by Fama and French 10 industry. The sample covers the period from 2009 to 2024.

IA1.3 Summary Statistics

Table IA1.2
Characteristics of AI Exposure-Sorted Portfolios (Quintile)

This table reports the summary statistics for characteristics of *AI Exposure*–sorted quintile portfolios. At the end of June each year, stocks are ranked by their AI Exposure and sorted into quintiles. AI Exposure is defined as the fraction of AI-related unigrams and bigrams in a firm's call transcripts, multiplied by 1,000 for ease of exposition. Firm size is the market equity in June of year t. *BE/ME* is the ratio of book equity for the fiscal year ending in calendar year $t-1$ to market equity at the end of December of year $t-1$. *ROE* is profitability, measured as income before extraordinary items divided by book equity. *ROA* is another profitability measure, defined as income before extraordinary items divided by total assets. *R&D/Assets* is the ratio of research and development to lagged total assets. *Firm Age* is the number of years since the firm's initial listing in the CRSP database. *Asset Growth* is the change in total assets from year $t-2$ to year $t-1$, scaled by total assets in year $t-2$. Following Fahlenbrach, Rageth, and Stulz (2021), three measures of financial flexibility are constructed: *Cash/Assets*, *St Debt/Assets*, and *Lt Debt/Assets*. *Tangibility* is property, plant, and equipment divided by total assets. *Book Leverage* is the sum of current liabilities and long-term debt divided by total assets. All firm characteristics are winsorized at the 1st and 99th percentiles to mitigate the influence of outliers. The sample period is from 2009 to 2024.

Quintile Portfolio	1	2	3	4	5
	mean	mean	mean	mean	mean
AIExposure	0.64	1.17	1.74	2.60	4.57
Size (millions)	4524.89	6908.72	7627.06	7163.33	6256.13
BE/ME	1.49	1.37	1.24	1.12	1.17
ROA	0.05	0.04	0.03	-0.03	-0.10
ROE	0.02	0.01	-0.04	-0.16	-0.31
R&D/Assets	0.00	0.01	0.03	0.08	0.14
Age	24.08	26.20	24.82	21.66	18.55
Asset Growth	0.12	0.13	0.16	0.19	0.22
Cash/Assets	0.09	0.11	0.17	0.26	0.36
St Debt/Assets	0.05	0.04	0.04	0.04	0.03
Lt Debt/Assets	0.28	0.26	0.24	0.21	0.18
Tangibility	0.31	0.29	0.25	0.18	0.14
Book Leverage	0.34	0.30	0.28	0.24	0.21

Table IA1.3
Summary Statistics and Correlation: AI Exposure Measures

This table presents the summary statistics and correlation of the overall firm-level *AI Exposure* and different versions of topic-specific or tone-specific AI Exposure measures at the firm-year level from 2009 to 2024. Panel A presents the summary statistics of different versions of AI Exposure measures. Panel B presents the correlation among different versions of AI Exposure measures. I denote the overall AI Exposure measure using cleaned AI Vocabulary (excluding ambiguous terms), the overall AI Exposure measure using raw AI Vocabulary (not excluding ambiguous terms), the AI Exposure regarding opportunity based on the raw measure, the AI Exposure regarding regulation based on the raw measure, the AI Exposure regarding risk based on the raw measure, the AI Exposure with positive tones based on the raw measure, and the AI Exposure with negative tones based on the raw measure as $AIExposure$, $AIExposure^{Raw}$, $AIExposure^{Raw,Opp}$, $AIExposure^{Raw,Reg}$, $AIExposure^{Raw,Risk}$, $AIExposure^{Raw,Pos}$, and $AIExposure^{Raw,Neg}$, respectively. $AIExposure^{Babina}$ refers to the firm level fraction of AI related labor from Babina et al. (2024b) during 2010 to 2018. For the ease of exposition, I multiply 1,000 for each measure.

Panel A: Summary Statistics

Variables ($\times$ 1000)	Observation	Mean	SD	10%	25%	50%	75%	90%	Max
$AIExposure_{i,t}$	57105	2.08	1.58	0.61	0.97	1.65	2.76	4.14	20.86
$AIExposure_{i,t}^{Raw}$	57105	5.37	1.94	3.31	4.01	5.02	6.35	7.89	22.92
$AIExposure_{i,t}^{Raw,Opp}$	57105	4.14	2.71	0.00	2.79	4.33	5.77	7.29	22.92
$AIExposure_{i,t}^{Raw,Reg}$	57105	1.56	2.54	0.00	0.00	0.00	3.49	5.45	22.92
$AIExposure_{i,t}^{Raw,Risk}$	57105	1.64	2.43	0.00	0.00	0.00	3.64	5.27	22.42
$AIExposure_{i,t}^{Raw,Pos}$	57105	2.28	2.78	0.00	0.00	0.00	4.44	6.12	22.58
$AIExposure_{i,t}^{Raw,Neg}$	57105	0.45	1.43	0.00	0.00	0.00	0.00	2.13	19.63
$AIExposure_{i,t}^{Babina}$	19636	3.13	16.63	0.00	0.00	0.00	0.00	5.00	600.00

Panel B: Correlation

	$AIExposure_{i,t}$	$AIExposure_{i,t}^{Raw}$	$AIExposure_{i,t}^{Raw,Opp}$	$AIExposure_{i,t}^{Raw,Reg}$	$AIExposure_{i,t}^{Raw,Risk}$	$AIExposure_{i,t}^{Raw,Pos}$	$AIExposure_{i,t}^{Raw,Neg}$	$AIExposure_{i,t}^{Babina}$
$AIExposure_{i,t}$	1							
$AIExposure_{i,t}^{Raw}$	0.618	1						
$AIExposure_{i,t}^{Raw,Opp}$	0.440	0.625	1					
$AIExposure_{i,t}^{Raw,Reg}$	0.119	0.224	0.231	1				
$AIExposure_{i,t}^{Raw,Risk}$	0.024	0.115	0.157	0.330	1			
$AIExposure_{i,t}^{Raw,Pos}$	0.272	0.324	0.283	0.118	0.062	1		
$AIExposure_{i,t}^{Raw,Neg}$	0.001	0.021	0.050	0.060	0.069	0.017	1	
$AIExposure_{i,t}^{Babina}$	0.216	0.134	0.117	0.036	0.014	0.095	-0.010	1

IA2 Variables Definition

Variable	**Definition**
Book Leverage	The sum of current liabilities (DLC) and long-term debt (DLTT) divided by total assets (AT).
Cash/Assets	The ratio of cash (CHE) to total assets (AT).
St Debt/Assets	The ratio of debt in current liabilities (DLC) to total assets (AT).
Lt Debt/Assets	The ratio of total long-term debt (DLTT) to total assets (AT).
ROE	The measure of profitability. It is calculated as the income before extraordinary items (IB) divided by book equity.
ROA	Profitability measure, defined as income before extraordinary items divided by total assets.
R&D/Assets	The ratio of R&D to lagged total assets.
Firm Age	The age of a firm starting from its initial listing in the CRSP database.
Asset Growth	The ratio of the change in total assets from year $t-2$ to year $t-1$ to total assets in year $t-2$.
Sales	Total sales, calculated as the total income generated by an organization through the sale of products and services. Preferably, consolidated sales data is used when available.
Firm Size	*ln(ME)*, measured as the log of market equity in June of year t.
BE/ME	The ratio of book equity for the fiscal year ending in calendar year $t-1$ to market equity at the end of December of year $t-1$.
SIC4	Four-digit SIC Code, categorizing businesses based on their primary activity per the US 1987 SIC classification.
Tangibility	Property, plant, and equipment (PPENT) divided by total assets (AT).

Variable	Definition
TF-IDF	TF-IDF score is the Term Frequency (TF) times Inverse Document Frequency (IDF). Term Frequency (TF) measures how frequently a term occurs in a document. $$\text{TF}(t, d) = \frac{\text{Number of times term } t \text{ appears in document } d}{\text{Total number of terms in document } d}$$ Inverse Document Frequency (IDF) measures how important a term is. It downweights terms that appear more frequently across multiple documents. $$\text{IDF}(t, D) = \log\left(\frac{\text{Total number of documents } D}{\text{Number of documents with term } t}\right)$$

IA3 Expected Returns Details

I include two versions of spread when estimating *AI premium gap*. The first version is the long-term spread which is the difference between objective ICC as in Gonçalves (2021) and subjective ICC as in Eskildsen et al. (2026). The objective ICC is constructed by forecasting future firm fundamentals using a VAR and solving for the discount rate that equates current prices to the expected present value of future payouts. The subjective ICC is the average of four different implied cost of capital measures in the accounting literature, see, Gebhardt, Lee, and Swaminathan (2001), Claus and Thomas (2001), Ohlson and Juettner-Nauroth (2005), and Easton (2004). The second version is the short-term spread which is the difference between a machine learning based expected returns using fundamentals as in Bastianello (2022) and analyst one-year-ahead price-target-implied expected returns as in Loudis (2024). The machine learning based one-year-ahead expected returns are estimated using random forest to map firm fundamentals into expected returns. The analyst one-year-ahead price-target-implied expected returns are estimated using one-year-ahead analyst price target and current price.

Table IA3.5 shows the summary statistics of the four expected returns measure in the period 2009-2024. Figure IA3.5 shows the density distribution for them. The following sections document the details on how to construct each of these expected returns.

Table IA3.5
Summary Statistics - Expected Return Measures

	N	Mean	Std. Dev.	Min	25%	50%	75%	Max
Subjective ICC (%)	519,659	8.99	4.36	-57.37	6.53	8.48	10.67	190.57
Objective ICC (%)	351,737	13.24	1.66	10.39	12.43	13.02	13.74	96.19
Analyst Expected Returns (%)	214,513	11.08	32.52	-209.51	4.52	13.41	23.07	146.95
Machine Learning Expected Returns (%)	482,581	12.93	31.64	-65.15	-6.78	5.09	24.74	201.37

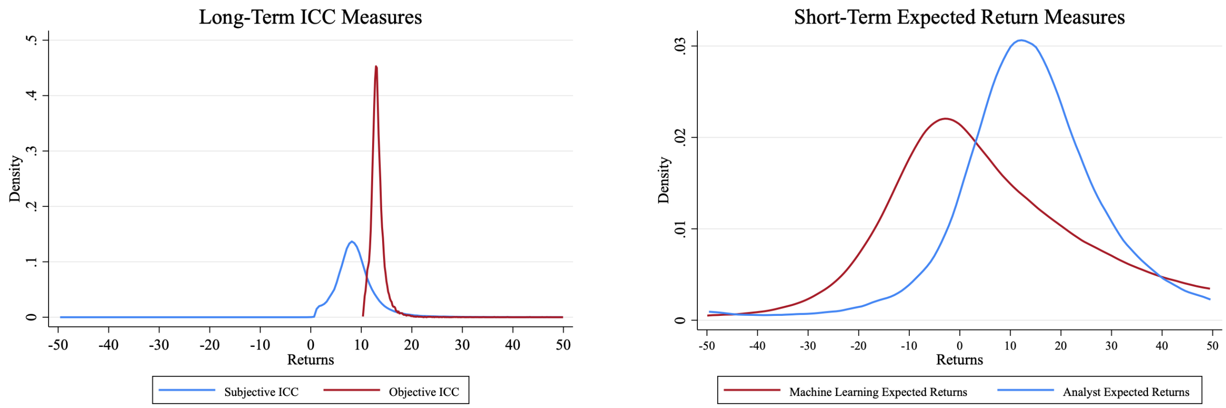

Figure IA3.5. Density Distribution of Expected Returns (%)

IA3.1 Objective ICC

The objective ICC is from Gonçalves (2021), and the data contain one observation per firm-fiscal year. The basic idea is to back out the implied cost of capital from current prices and fundamentals-implied future payouts. Different from subjective ICC measures, the cash flow forecasts here are not based on analyst earnings forecasts. They are generated from historical firm fundamentals through a cross-sectional Vector Autoregression (VAR). For firm i at time t, the ICC, denoted $r_{i,t}^{\text{ICC}}$, solves:

$$P_{i,t} = \sum_{j=1}^{T} \frac{\mathbb{E}_t\left[\text{Payout}_{i,t+j}\right]}{(1 + r_{i,t}^{\text{ICC}})^j} + \frac{\mathbb{E}_t[\text{TV}_{i,t+T}]}{(1 + r_{i,t}^{\text{ICC}})^T}, \tag{IA3.1}$$

where $P_{i,t}$ is the current market equity, T is the forecast horizon, and $\text{TV}_{i,t+T}$ is the terminal value. The estimation sets $T = 15$ years. The first step is to construct payouts. Payouts are not simply observed dividends. Instead, they are based on the clean-surplus accounting identity, which links earnings, book equity, and total distributions to shareholders:

$$\text{Payout}_{i,t} = \text{Earnings}_{i,t} - \Delta\text{BE}_{i,t}, \tag{IA3.2}$$

where $\text{BE}_{i,t}$ is the book value of equity at the end of year t and $\Delta\text{BE}_{i,t} \equiv \text{BE}_{i,t} - \text{BE}_{i,t-1}$. Intuitively, earnings can either be retained inside the firm and increase book equity, or paid out to shareholders. Therefore, the part of earnings that is not retained is treated as payout. This payout definition is closely related to:

$$\text{Payout}_{i,t} = \text{Dividends}_{i,t} + \text{Net Repurchases}_{i,t}, \tag{IA3.3}$$

where net repurchases equal share repurchases minus equity issuances. The advantage of using the clean-surplus identity is that future payouts can be recovered from forecasts of earnings and book equity, instead of separately forecasting dividends, repurchases, and equity issuances. To forecast future payouts, Gonçalves (2021) uses firm-level fundamentals. Flow variables are scaled by lagged book equity. The key variables are profitability, book-equity growth, and payout:

$$\text{ROE}_{i,t} = \frac{\text{Earnings}_{i,t}}{\text{BE}_{i,t-1}}, \tag{IA3.4}$$

$$\text{AG}_{i,t} = \frac{\text{BE}_{i,t} - \text{BE}_{i,t-1}}{\text{BE}_{i,t-1}}, \tag{IA3.5}$$

$$\text{PO}_{i,t} = \frac{\text{Payout}_{i,t}}{\text{BE}_{i,t-1}}. \tag{IA3.6}$$

By the clean-surplus identity, these variables satisfy:

$$\text{PO}_{i,t} = \text{ROE}_{i,t} - \text{AG}_{i,t}. \tag{IA3.7}$$

This means that payout is determined by how profitable the firm is and how much of that profitability is retained in book equity. A firm with high profitability but also high investment may have low current payout, while a firm with high profitability and low book-equity growth has high payout. The log book-to-market ratio is also included as a state variable:

$$bm_{i,t} = \log\left(\frac{\text{BE}_{i,t}}{P_{i,t}}\right). \tag{IA3.8}$$

This variable is useful because it contains information about long-run mean reversion in profitability and investment. The state vector can be written as:

$$X_{i,t} = \begin{pmatrix} \mathrm{ROE}_{i,t} \\ \mathrm{AG}_{i,t} \\ bm_{i,t} \end{pmatrix}, \tag{IA3.9}$$

with the payout rate recovered from equation (IA3.7). Future fundamentals are then forecast using a first-order cross-sectional VAR:

$$X_{i,t+1} = \alpha + A\, X_{i,t} + \varepsilon_{i,t+1}, \tag{IA3.10}$$

where A is the matrix of slope coefficients. The coefficients are estimated from the pooled cross-section of firms and years, rather than from separate firm-level time series. This is useful because many firms do not have long enough histories to estimate firm-specific dynamics. The pooled regression imposes common average dynamics, while the current firm-level state vector $X_{i,t}$ still allows expected future fundamentals to differ across firms. Given the first-order VAR, multi-period forecasts are obtained by iterating the system forward:

$$\mathbb{E}_t[X_{i,t+j}] = \alpha \sum_{k=0}^{j-1} A^k + A^j\, X_{i,t}. \tag{IA3.11}$$

The expected payout rate at horizon j is then:

$$\mathbb{E}_t[\mathrm{PO}_{i,t+j}] = \mathbb{E}_t[\mathrm{ROE}_{i,t+j}] - \mathbb{E}_t[\mathrm{AG}_{i,t+j}] = e_1^\top \mathbb{E}_t[X_{i,t+j}] - e_2^\top \mathbb{E}_t[X_{i,t+j}], \tag{IA3.12}$$

where e_1 and e_2 select the profitability and book-equity growth components of the state vector. Expected book equity is built recursively from the forecasted book-equity growth rate:

$$\mathbb{E}_t[\mathrm{BE}_{i,t+1}] = \mathrm{BE}_{i,t}\left(1 + \mathbb{E}_t[\mathrm{AG}_{i,t+1}]\right), \tag{IA3.13}$$

$$\mathbb{E}_t[\mathrm{BE}_{i,t+j}] = \mathbb{E}_t[\mathrm{BE}_{i,t+j-1}]\left(1 + \mathbb{E}_t[\mathrm{AG}_{i,t+j}]\right), \quad j \geq 2. \tag{IA3.14}$$

Therefore, the expected dollar payout at horizon j is:

$$\mathbb{E}_t[\mathrm{Payout}_{i,t+j}] = \mathbb{E}_t[\mathrm{PO}_{i,t+j}] \times \mathbb{E}_t[\mathrm{BE}_{i,t+j-1}]. \tag{IA3.15}$$

Since the present-value formula cannot be extended to infinity in practice, a terminal value is added at the end of the forecast horizon. The terminal value captures the value of payouts after year $T = 15$:

$$\mathrm{TV}_{i,t+T} = \frac{\mathbb{E}_t[\mathrm{Payout}_{i,t+T+1}]}{r_{i,t}^{\mathrm{ICC}} - \bar{g}}, \tag{IA3.16}$$

where $\bar{g}$ is the long-run steady-state growth rate. This terminal value is especially important for firms with low near-term payouts, because a large part of their value comes from cash flows far in the future. Finally, substituting the forecasted payouts and terminal value into equation (IA3.1), the objective ICC is the internal rate of return that equates the current market equity to the expected discounted value of future payouts:

$$r_{i,t}^{\mathrm{ICC}} = \arg\min_{r} \left| P_{i,t} - \sum_{j=1}^{T} \frac{\mathbb{E}_t[\mathrm{Payout}_{i,t+j}]}{(1+r)^j} - \frac{\mathbb{E}_t[\mathrm{TV}_{i,t+T}]}{(1+r)^T} \right|. \tag{IA3.17}$$

As is clear to see, the measure is long-term and fundamentals-implied because the expected payout stream is generated from accounting fundamentals, mainly profitability, investment, and book-to-market, rather than from analyst forecasts. This makes the measure useful in my setting because it provides an fundamentals-implied expected return benchmark that can be compared to subjective ICC. The next section describes the subjective ICC.

IA3.2 Subjective ICC

The subjective ICC is the average of four different implied cost of capital measures in the accounting literature, following Eskildsen et al. (2026). As Mohanram and Gode (2013) notes, taking the average across several ICC models can help reduce model-specific biases and measurement error. This is useful because each model makes slightly different assumptions about forecast horizon, payout, abnormal earnings, and terminal value. I use four standard ICC measures: GLS in Gebhardt, Lee, and Swaminathan (2001), CT in Claus and Thomas (2001), OJ in Ohlson and Juettner-Nauroth (2005), and PEG in Easton (2004).

All four measures are built from the same firm-month panel, so differences across them come from the valuation model rather than from the data. The panel is dated at the monthly IBES consensus summary dates, so forecasts and prices are always known at the same point in time. Prices are from CRSP. I take the last daily close on or before each forecast date, allowing up to a six-day gap for holidays, and adjust it for splits. The forecast inputs are the median consensus earnings per share forecasts for the next two fiscal years and the median long-term growth forecast. When the two-year-ahead forecast is missing, which is common in the early part of the sample, I set it to 1.15 times the one-year-ahead forecast.

The accounting inputs are annual. Book value per share is from Compustat and is supplemented by the IBES reported value where Compustat is unavailable. The payout ratio is dividends divided by income before extraordinary items, with the denominator replaced by 6% of total assets for loss firms, and it is winsorized at the 1st and 99th percentiles. Both enter with a four-month lag after the fiscal year end so that they are known to investors. The industry benchmark used by GLS is the median return on equity within each of the 49 Fama–French industries over the trailing ten years, computed from positive-ROE firm-years only. The long-run growth anchor used by CT and OJ is the ten-year constant-maturity Treasury yield minus 3%. A firm-month enters the panel only if it has both a price and a one-year-ahead earnings forecast, and I keep firms followed by at least three analysts on average over the previous twelve months. GLS and CT have no closed-form solution, so I solve them numerically over a range of 1% to 50% and drop firm-months without a root in that range.

The first measure is the GLS ICC from Gebhardt, Lee, and Swaminathan (2001). This model is based on residual income valuation. Let M_t denote market equity in year t, B_t denote book equity, and R denote the implied cost of capital. The GLS ICC solves:

$$M_t = B_t + \sum_{\kappa=1}^{11} \frac{E_t[(ROE_{t+\kappa} - R)B_{t+\kappa-1}]}{(1+R)^{\kappa}} + \frac{E_t[(ROE_{t+12} - R)B_{t+11}]}{R(1+R)^{11}}. \tag{IA3.18}$$

The main object here is residual income, $(ROE_{t+\kappa} - R)B_{t+\kappa-1}$, which measures earnings in excess of the required return on book equity. Near-term ROE is forecast using analyst earnings forecasts and book equity implied by clean-surplus accounting:

$$B_{t+\kappa} = B_{t+\kappa-1} + E_{t+\kappa} - D_{t+\kappa}, \tag{IA3.19}$$

where $E_{t+\kappa}$ is forecasted earnings and $D_{t+\kappa}$ is forecasted dividends. Forecasted dividends are forecasted earnings times the firm's payout ratio, which is held fixed over the horizon. After the explicit forecast period, expected ROE is assumed to mean-revert toward its industry median. Concretely, the first two years of expected ROE come from the one-year-ahead and two-year-ahead consensus forecasts, each scaled by beginning-of-year book equity. From year three through year twelve, expected ROE fades linearly to the industry median in ten equal steps, and book equity rolls forward each year by the earnings that are retained. Residual income in year twelve is then capitalized as a perpetuity, so the explicit horizon is twelve years. The ICC is the value of R that makes the residual-income valuation equal to current market equity.

The second measure is the CT ICC from Claus and Thomas (2001). It is also a residual-income model, but it uses a shorter explicit forecast horizon and a terminal value based on long-run growth:

$$M_t = B_t + \sum_{\kappa=1}^{5} \frac{E_t[(ROE_{t+\kappa} - R)B_{t+\kappa-1}]}{(1+R)^\kappa} + \frac{E_t[(ROE_{t+5} - R)B_{t+4}]\,(1+g)}{(R-g)(1+R)^5}. \qquad \text{(IA3.20)}$$

The first part of the formula discounts expected residual income over the first five years. The second part is the terminal value of residual income after year five, where g is the long-run growth rate. Years one and two use the same two consensus earnings forecasts as GLS. Years three through five grow forecasted earnings at the analyst long-term growth forecast, and book equity again rolls forward through clean-surplus accounting at the firm's payout ratio. Following the standard implementation, g is set to the current risk-free rate minus 3%, measured by the ten-year Treasury yield. Because the perpetuity requires $R > g$, the lower bound of the numerical search is raised just above g. Compared with GLS, the CT model puts more structure on the terminal value and therefore can produce different ICC estimates.

The third measure is the OJ ICC from Ohlson and Juettner-Nauroth (2005). This model gives a closed-form expression for the implied cost of capital based on near-term earnings forecasts, expected dividends, and expected earnings growth:

$$R = A + \sqrt{A^2 + \frac{E_t[E_{t+1}]}{M_t}\,(g - (\gamma - 1))}, \qquad \text{(IA3.21)}$$

where

$$A = \frac{1}{2}\left((\gamma - 1) + \frac{E_t[D_{t+1}]}{M_t}\right), \qquad \text{(IA3.22)}$$

and

$$g = \max\left\{\sqrt{\frac{E_t[E_{t+2}] - E_t[E_{t+1}]}{E_t[E_{t+1}]} \times g^{LT}},\ g^{LT}\right\}. \qquad \text{(IA3.23)}$$

Here $E_t[E_{t+1}]$ is the expected earnings in year $t+1$, $E_t[D_{t+1}]$ is the expected dividend in year $t+1$, g captures short-run earnings growth, and $\gamma - 1$ is the long-run growth rate in abnormal earnings. Short-run growth is the geometric average of the growth implied by the two consensus earnings forecasts and the analyst long-term growth forecast g^{LT}, floored at g^{LT} so that it never falls below long-run growth. The expected dividend is the one-year-ahead earnings forecast times the payout ratio, so OJ uses no book value at all. The long-run rate $\gamma - 1$ is the ten-year Treasury yield minus 3%. The intuition is that higher expected earnings relative to price and higher short-run earnings growth imply a higher implied cost of capital.

The fourth measure is the PEG ICC from Easton (2004). This model is also based on abnormal earnings growth, and it is a stripped-down version of OJ that drops the dividend term and the long-run growth term:

$$R = \sqrt{\frac{E_t[E_{t+1}]}{M_t} \times g}, \tag{IA3.24}$$

where g is the same short-run growth rate as in equation (IA3.23). Equivalently, R is the discount rate that links current market equity to expected short-run earnings growth. The PEG model is the most parsimonious of the four. It uses only the two consensus earnings forecasts, the long-term growth forecast, and the current price. It needs neither book value nor a risk-free rate, so it survives for the widest cross-section of firms.

For each firm-month, I compute the four ICC estimates separately and then take their average:

$$ICC_{i,t}^{\text{Analyst}} = \frac{1}{4}\left(ICC_{i,t}^{GLS} + ICC_{i,t}^{CT} + ICC_{i,t}^{OJ} + ICC_{i,t}^{PEG}\right). \tag{IA3.25}$$

Each estimate is an annualized cost of equity. This average is less dependent on any single valuation model. In my setting, this is important because I want the subjective ICC to capture the expected return implied by analysts' earnings forecasts, rather than mechanical assumptions from one particular ICC model.

IA3.3 Machine Learning Expected Returns

The machine-learning expected returns are estimated using a Random Forest (RF) regression following Bastianello (2022). The method is well established in the literature (see, e.g. van Binsbergen, Han, and Lopez-Lira, 2023; Zhang, Zhu, and Linnainmaa, 2025). The idea is to predict the total payoff over the next 12 months using firm characteristics and then convert it into expected returns.

The target variable is the total dollar payoff over the next 12 months, expressed in terms of the current (unadjusted) price scale:

$$Y_{i,t+12} = \left(\tilde{P}_{i,t+12} + \tilde{D}_{i,t,t+12}\right) \times CFACPR_{i,t-1}, \tag{IA3.26}$$

where $\tilde{P}_{i,t+12} = P_{i,t+12}/CFACPR_{i,t+12}$ is the split-adjusted price, and $\tilde{D}_{i,t,t+12}$ is the cumulative split-adjusted dividend:

$$\tilde{D}_{i,t,t+12} = \sum_{s=1}^{12}\left(r_{i,t+s} - r^{x}_{i,t+s}\right)\tilde{P}_{i,t+s-1}. \tag{IA3.27}$$

Multiplying by the lagged adjustment factor $CFACPR_{i,t-1}$ scales the payoff back to the price level at time t, so $Y_{i,t+12}/P_{i,t}$ directly gives the 12-month total return.

The predictors include four groups of variables observed at time t as in Zhang, Zhu, and Linnainmaa (2025). First, I include 67 accounting ratios from the WRDS Financial Ratios Suite, covering valuation, profitability, leverage, liquidity, and efficiency. These are aligned using their *public_date* and forward-filled for up to six months to avoid look-ahead bias.

Second, I include four CRSP market variables: log price $\ln P_{i,t}$, monthly return $r_{i,t}$, log market capitalization $\ln ME_{i,t}$, and trailing 12-month return:

$$r^{12}_{i,t} = \prod_{s=1}^{12}(1 + r_{i,t-s+1}) - 1. \tag{IA3.28}$$

Third, I construct 28 per-share variables from Compustat by scaling by implied shares

$Q_{i,t} = ME_{i,t}/P_{i,t}$:

$$x^{ps}_{i,t} \;=\; \frac{X^{Compustat}_{i,t}}{Q_{i,t}} \times 1{,}000. \tag{IA3.29}$$

Balance-sheet variables use the most recent quarter with a three-month lag, while income and cash-flow variables are aggregated over the trailing four quarters.

Finally, I include four macro variables lagged by one month: GDP growth, consumption growth, industrial production growth, and the unemployment rate.

Before estimation, I preprocess all features. Missing values are first filled using the industry-month median (based on Fama–French 49 industries), and then by the cross-sectional median if needed. Within each month, all features are winsorized at the 1st and 99th percentiles and then rank-normalized to $[-1, 1]$:

$$\tilde{x}_{i,k,t} \;=\; 2 \cdot \widehat{F}_{k,t}(x_{i,k,t}) - 1, \tag{IA3.30}$$

where $\widehat{F}_{k,t}$ is the empirical CDF. The target $Y_{i,t+12}$ is also winsorized within each month.

I then estimate a Random Forest model:

$$\widehat{Y}_{i,t+12} \;=\; \mathcal{F}\Big(\mathbf{x}_{i,t};\, \widehat{\theta}_t\Big)\,, \tag{IA3.31}$$

using a rolling out-of-sample window. For each month t, the training sample consists of all firm-months in the 36-month window ending 12 months before t:

$$\mathcal{T}_t \;=\; \{(i, s) : t - 48 < s \le t - 12,\; Y_{i,s+12} \text{ observed}\}. \tag{IA3.32}$$

This ensures that all realized outcomes used in training are available at time t. The model follows Zhang, Zhu, and Linnainmaa (2025) in using 1,000 trees, maximum depth 15, 55% subsampling, $\lfloor \sqrt{K} \rfloor$ features per split, and a minimum of five observations per leaf.

Given the predicted payoff, the 12-month expected return is:

$$\widehat{R}^{RF}_{i,t,t+12} \;=\; \frac{\widehat{Y}_{i,t+12}}{P_{i,t}} - 1. \tag{IA3.33}$$

I convert this to a monthly expected return:

$$\widehat{R}^{RF}_{i,t,t+1} \;=\; \left(1 + \widehat{R}^{RF}_{i,t,t+12}\right)^{1/12} - 1. \tag{IA3.34}$$

As a final step, 12-month returns are winsorized cross-sectionally at the 1st and 99th percentiles within each month.

IA3.4 Analyst Expected Returns

I estimate analyst-implied expected returns in the spirit of Loudis (2024). The measure is constructed from analyst 12-month price targets in IBES linked to CRSP stock prices at the firm-month level.

I obtain split-adjusted price targets from the IBES Detail file (*ptgdet*), restricting to 12-month horizon targets for U.S. firms with positive values. When an analyst issues multiple price targets for a given firm within a month, I retain only the most recent one. The consensus price target is then constructed as the cross-sectional median across analysts:

$$\tilde{P}_{i,t,t+12} \;=\; \text{Median}_j\left\{\tilde{P}^{j}_{i,t,t+12}\right\}, \tag{IA3.35}$$

where j indexes analysts.

To merge IBES with CRSP, I link IBES tickers to CRSP *permno* identifiers using the eight-digit CUSIP. Specifically, I match IBES *cusip* to the CRSP historical name file (*msenames*) field *ncusip*, and restrict matches to valid name date ranges.

Given the consensus price target, the 12-month analyst-implied expected return is:

$$\widehat{R}^A_{i,t,t+12} \;=\; \frac{\tilde{P}_{i,t,t+12} - P_{i,t-1}}{P_{i,t-1}}, \tag{IA3.36}$$

where $P_{i,t-1}$ is the end-of-previous-month stock price from CRSP. I use the lagged price rather than the current price to avoid mechanically incorporating the price reaction to the analyst report itself as in Loudis (2024). I also drop extreme values outside $[-0.90, 3.00]$ as likely data errors.

I convert this to a monthly expected return using geometric de-compounding:

$$\widehat{R}^A_{i,t,t+1} \;=\; \left(1 + \widehat{R}^A_{i,t,t+12}\right)^{1/12} - 1. \tag{IA3.37}$$

As a robustness check, I also construct a stale-target version, where the most recent consensus price target is carried forward for up to six months. The staleness is tracked as the number of months since issuance, and targets older than six months are set to missing. This version increases coverage at the cost of using somewhat older analyst information, but the qualitative estimation results largely remain the same. Loudis (2024) shows that including dividend yields does not materially affect the quantitative or qualitative properties of expected return measures. Accordingly, I focus on the version excluding dividend yields.

IA4 AI Premium Gap Estimation

IA4.1 Spread of Z-Scores

This subsection reports the *AI premium gap* estimated using the spread of z-scores, which standardizes each expected return series within month before differencing and therefore controls for differences in scale and dispersion across the objective and subjective measures. The estimates line up with the raw-spread results in Figure 1 of the main paper. After adding controls and industry fixed effects, the estimated coefficient is consistently around 0.05 for both the long-term and the short-term spread, so a one-standard-deviation increase in AI Exposure is associated with a 0.05 standard deviation increase in the expected return *spread*.

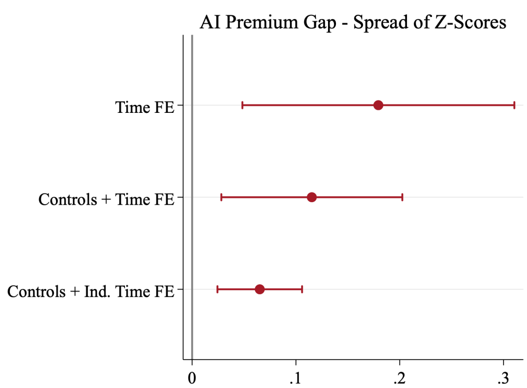

(a) **Long-term Spread of Z-Scores**

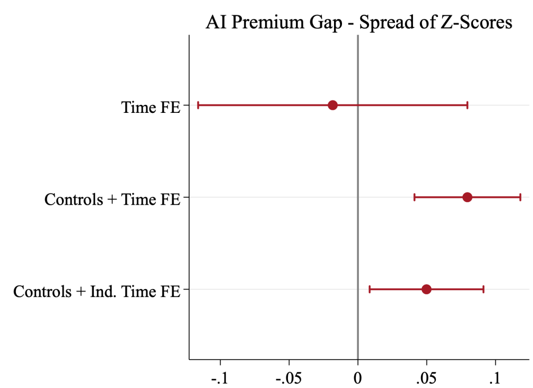

(b) **Short-term Spread of Z-Scores**

Figure IA4.6. AI Premium Gap (Spread of Z-Scores). The figure reports the estimated annual *AI premium gap* from 2009 to 2024, based on panel regressions of the spread between objective expected returns and subjective expected returns on lagged, standardized *AI Exposure*, following a similar approach to Eskildsen et al. (2026). Each expected return series is standardized within month before differencing. Each panel reports three specifications: Time FE only, Time FE with controls, and Time FE with controls and Industry FE. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry. Panel (a) uses the long-term spread, defined as econometric-based objective ICC following Gonçalves (2021) minus analyst-based subjective ICC following Eskildsen et al. (2026). Panel (b) uses the short-term spread, defined as machine-learning-implied expected returns following Bastianello (2022) minus analyst one-year-ahead target-price-implied returns following Loudis (2024).

IA4.2 Winsorized Raw Spread and Spread of Z-Scores

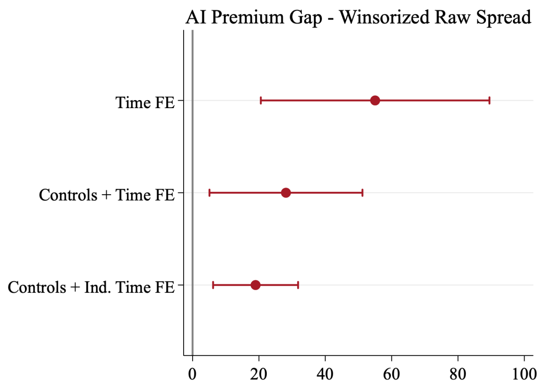

(a) **Long-term Raw Spread (bps)**

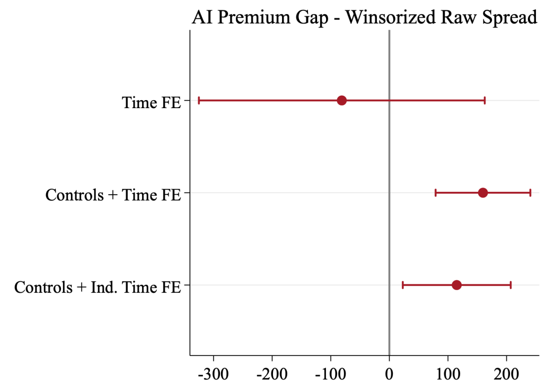

(b) **Short-term Raw Spread (bps)**

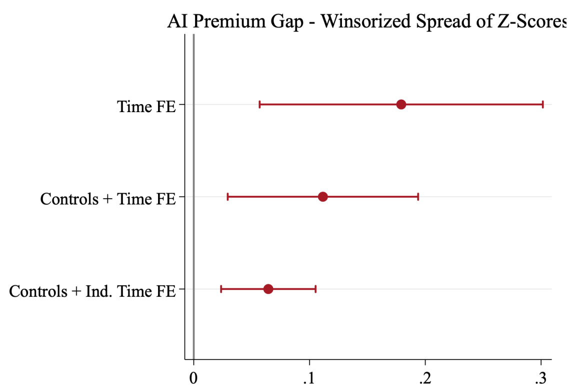

(c) **Long-term Spread of Z-Scores**

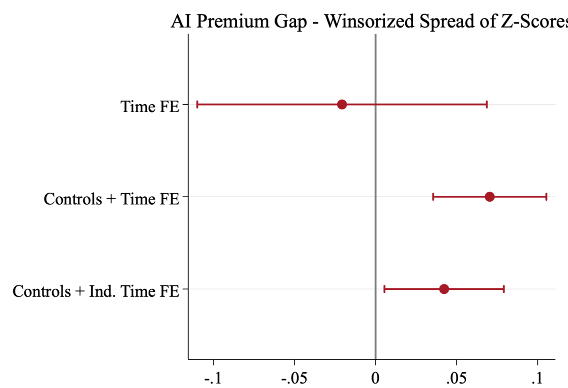

(d) **Short-term Spread of Z-Scores**

Figure IA4.7. AI Premium Gap (Winsorized). The figure reports the estimated annual *AI premium gap* from 2009 to 2024, based on panel regressions of the spread between objective expected returns and subjective expected returns on lagged, standardized *AI Exposure*, following a similar approach to Eskildsen et al. (2026). Each panel reports three specifications: Time FE only, Time FE with controls, and Time FE with controls and Industry FE. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry. Panels (a) and (c) use the long-term spread, defined as econometric-based objective ICC following Gonçalves (2021) minus analyst-based subjective ICC following Eskildsen et al. (2026). Panels (b) and (d) use the short-term spread, defined as machine-learning-implied expected returns following Bastianello (2022) minus analyst one-year-ahead target-price-implied returns following Loudis (2024). Panels (a)–(b) use winsorized raw spreads, and Panels (c)–(d) use winsorized spreads of z-scores. All spreads are winsorized at the 1st and 99th percentiles within month.

IA4.3 Individual Model ICC

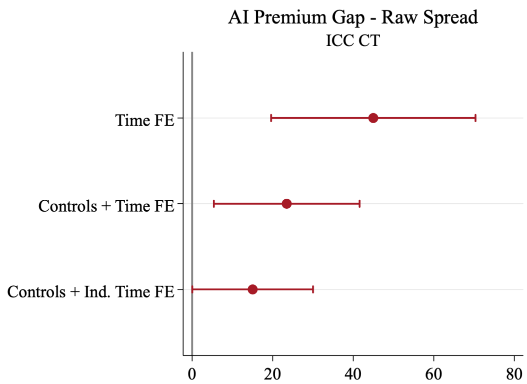

(a) **Raw Spread: ICC CT**

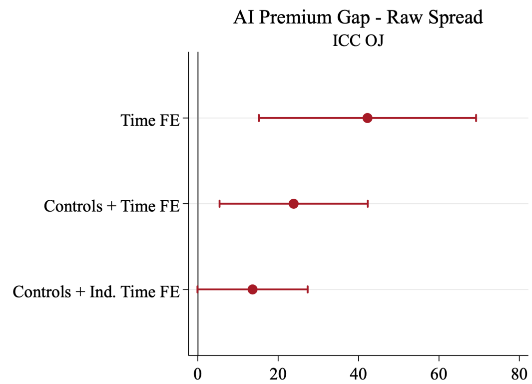

(b) **Raw Spread: ICC OJ**

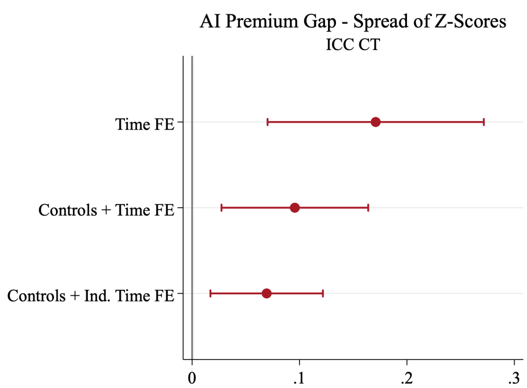

(c) **Spread of Z-Scores: ICC CT**

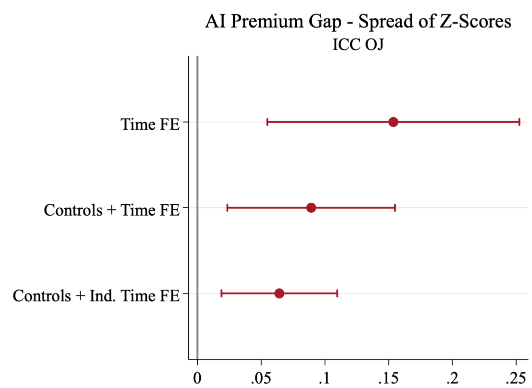

(d) **Spread of Z-Scores: ICC OJ**

Figure IA4.8. AI Premium Gap (ICC: CT and OJ). The figure reports the estimated annual *AI premium gap* from 2009 to 2024, based on panel regressions of the spread between objective expected returns and subjective expected returns on lagged, standardized *AI Exposure*, following a similar approach to Eskildsen et al. (2026). Each panel reports three specifications: Time FE only, Time FE with controls, and Time FE with controls and Industry FE. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry. The expected returns spread is defined as the econometric-based objective ICC following Gonçalves (2021) minus analyst-based subjective ICC. Panels (a) and (c) use the CT ICC following Claus and Thomas (2001). Panels (b) and (d) use the OJ ICC following Ohlson and Juettner-Nauroth (2005). Panels (a)–(b) use raw spreads, and Panels (c)–(d) use spreads of z-scores.

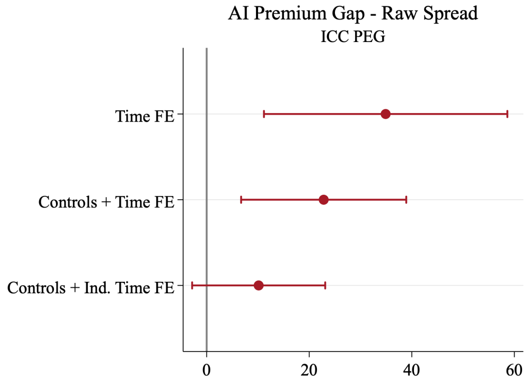

(a) **Raw Spread: ICC PEG**

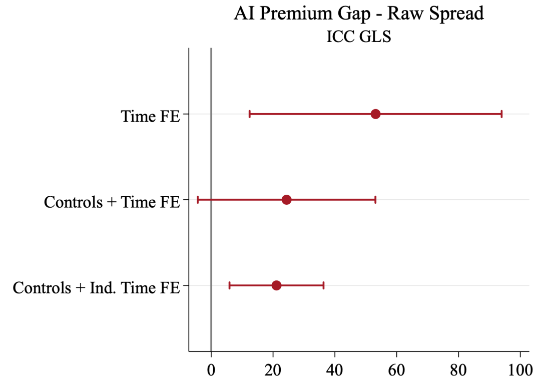

(b) **Raw Spread: ICC GLS**

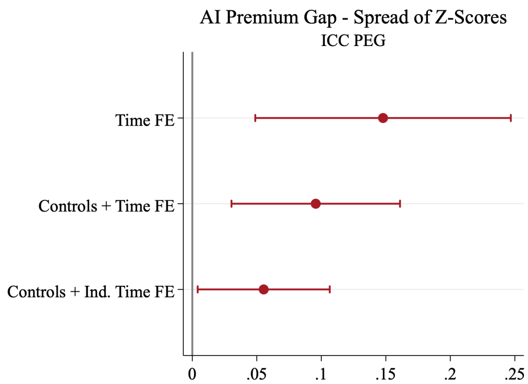

(c) **Spread of Z-Scores: ICC PEG**

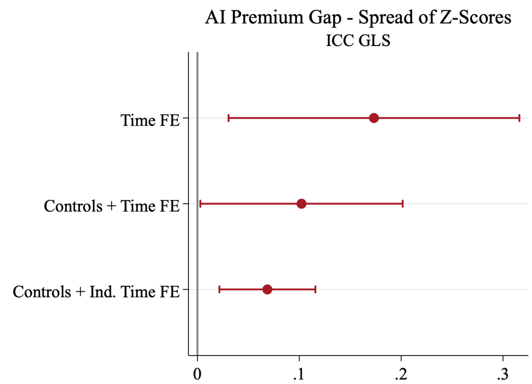

(d) **Spread of Z-Scores: ICC GLS**

Figure IA4.9. AI Premium Gap (ICC: PEG and GLS). The figure reports the estimated annual *AI premium gap* from 2009 to 2024, based on panel regressions of the spread between objective expected returns and subjective expected returns on lagged, standardized *AI Exposure*, following a similar approach to Eskildsen et al. (2026). Each panel reports three specifications: Time FE only, Time FE with controls, and Time FE with controls and Industry FE. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry. The expected return spread is defined as the econometric-based objective ICC following Gonçalves (2021) minus analyst-based subjective ICC. Panels (a) and (c) use the PEG model following Easton (2004). Panels (b) and (d) use the GLS model following Gebhardt, Lee, and Swaminathan (2001). Panels (a)–(b) use raw spreads, and Panels (c)–(d) use spreads of z-scores.

IA4.4 Controlling for AI Producer

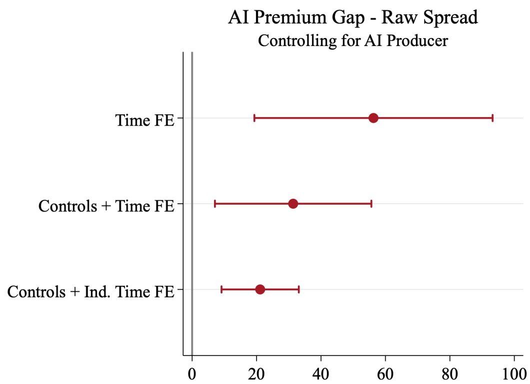

(a) **Long-term Raw Spread (bps)**

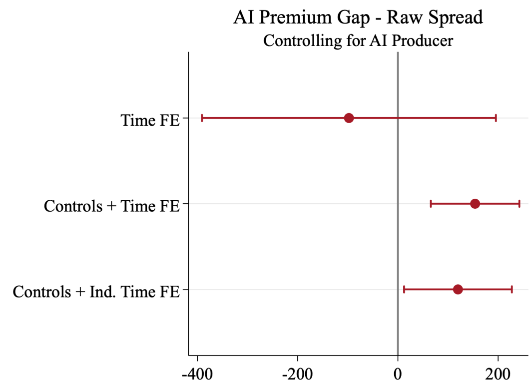

(b) **Short-term Raw Spread (bps)**

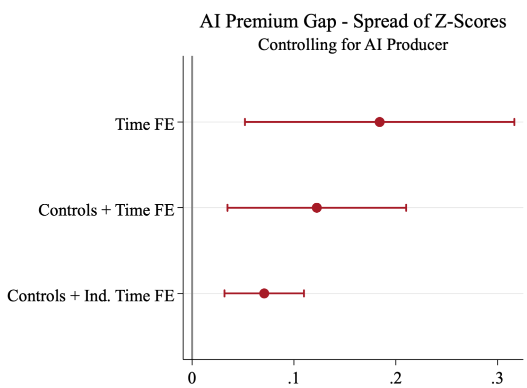

(c) **Long-term Spread of Z-Scores**

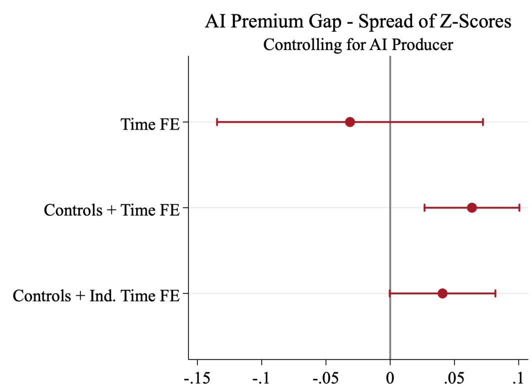

(d) **Short-term Spread of Z-Scores**

Figure IA4.10. AI Premium Gap (Controlling AI Producer). The figure reports the estimated annual *AI premium gap* from 2009 to 2024, based on panel regressions of the spread between objective expected returns and subjective expected returns on lagged, standardized *AI Exposure*, following a similar approach to Eskildsen et al. (2026), adding a control variable, *isAIProducer*, a firm-level dummy that identifies whether a firm produces AI-related products. Each panel reports three specifications: Time FE only, Time FE with controls, and Time FE with controls and Industry FE. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry. Panels (a) and (c) use the long-term spread, defined as econometric-based objective ICC following Gonçalves (2021) minus analyst-based subjective ICC following Eskildsen et al. (2026). Panels (b) and (d) use the short-term spread, defined as machine-learning-implied expected returns following Bastianello (2022) minus analyst one-year-ahead target-price-implied returns following Loudis (2024). Panels (a)–(b) use raw spreads, and Panels (c)–(d) use spreads of z-scores.

Table IA4.6
AI Premium Gap

This table reports panel regressions of the raw spread between objective expected returns and subjective expected returns on lagged, standardized *AI Exposure*, following a similar approach to Eskildsen et al. (2026). For each horizon, three specifications are reported: Time FE only, Time FE with controls, and Time FE with controls and Industry FE. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry and are reported in parentheses. The long-term spread is defined as econometric-based objective ICC following Gonçalves (2021) minus analyst-based subjective ICC following Eskildsen et al. (2026). The short-term spread is defined as machine-learning-implied expected returns following Bastianello (2022) minus analyst one-year-ahead target-price-implied returns following Loudis (2024). Both spreads are reported in basis points. *, **, *** denote significance at 10%, 5%, and 1% levels.

	Long-Term Spread			**Short-Term Spread**		
	(1)	(2)	(3)	(4)	(5)	(6)
$AIE_{i,t-1}$	55.415***	28.948**	19.071***	−67.020	189.263***	140.348**
	(18.534)	(12.307)	(6.497)	(138.943)	(47.397)	(55.616)
$Size_{i,t-1}$		−3.216	4.928		−46.588	−35.934
		(5.553)	(5.509)		(56.703)	(62.319)
$BE/ME_{i,t-1}$		−309.665***	−235.170***		−68.167	1.467
		(86.776)	(72.333)		(253.597)	(265.723)
$ROE_{i,t-1}$		34.759***	21.676***		−3.120	−56.368
		(12.788)	(7.713)		(74.872)	(79.868)
$R\&D_{i,t-1}$		86.070***	81.491***		−684.927***	−425.452**
		(23.335)	(13.451)		(195.913)	(190.183)
$Age_{i,t-1}$		0.460	0.681		−17.855	−16.108
		(9.883)	(7.296)		(61.265)	(63.293)
$Asset\ Growth_{i,t-1}$		−23.103***	−16.862***		−105.079***	−89.915***
		(5.334)	(3.641)		(37.635)	(28.829)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	Yes	No	No	Yes
Observations	286,937	283,850	283,850	164,077	157,983	157,983
Adjusted R^2	0.043	0.110	0.202	0.099	0.115	0.139

IA5 IT Premium Gap Estimation

IA5.1 Spread of Z-Scores

This subsection reports the *IT premium gap* estimated using the spread of z-scores. Each expected-return series is standardized within month before differencing, which controls for differences in scale and dispersion between the objective and subjective measures. The estimates are consistent with the raw-spread results in Figure 7 of the main paper. After adding controls and industry fixed effects, the estimated coefficients for the IT producer premium gap are consistently around -0.2, while the IT adopter premium gap remains statistically insignificant.

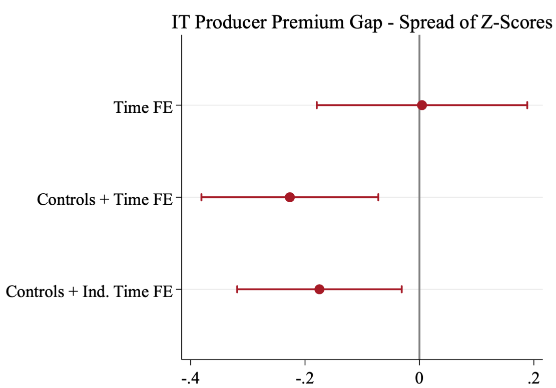

(a) **IT Producer Premium Gap**

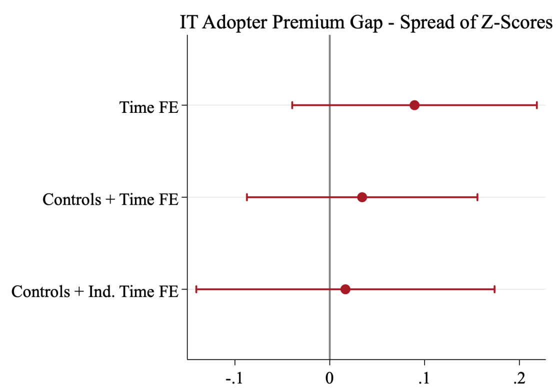

(b) **IT Adopter Premium Gap**

Figure IA5.11. IT Premium Gap (Spread of Z-Scores). The figure reports estimated annual *IT premium gaps* from 1990 to 1999, based on panel regressions in which the dependent variable is the long-term spread of z-scores, defined as the z-score of objective ICC following Gonçalves (2021) minus the z-score of analyst ICC as in Eskildsen et al. (2026). Panel (a) reports the *IT producer premium gap*, estimated using the IT sector dummy, while Panel (b) reports the *IT adopter premium gap*, estimated using standardized IT Intensity. Controls include size, BE/ME, ROE, R&D, age, and asset growth. Standard errors are clustered by firm, month, and Fama–French 49 industry.

IA6 LLMs: AI Producer Classification

This section documents the LLM prompt I use to classify firms as *AI producers* or *AI users* based on their conference call transcripts. Each transcript is processed by `gpt-5-nano`, which is instructed to rely on its knowledge of firms and industries to determine whether the firm builds or sells core AI technologies (AI producer), uses AI in products or operations (AI user), or both. The model follows a conservative rule that defaults to classifying firms as AI users when the evidence is ambiguous. The exact prompt is shown below:

```
You are an expert analyst classifying firms based on their AI involvement.
↪ You may use your knowledge of companies,
industries, and publicly available information to help make informed
↪ classifications.

Analyze this earnings call transcript and classify the firm's AI involvement
↪ into TWO categories:

**AI PRODUCER** (firms that build/sell AI technology):
1. Compute and hardware: GPUs, accelerators, memory chips, foundries,
↪ semiconductor manufacturing for AI workloads
   - Examples: NVIDIA, AMD, Intel, TSMC, Micron
2. Foundation models and toolchains: LLM development, model training
↪ platforms, fine-tuning tools, inference engines, AI APIs
   - Examples: OpenAI, Anthropic, Cohere, Hugging Face, Databricks
3. AI infrastructure software: MLOps platforms, orchestration tools, data
↪ platforms, vector databases, model evaluation, AI safety tools
   - Examples: Snowflake, Databricks, Scale AI, Weights & Biases
```

```
**AI USER** (firms that use AI for products or operations):
1. Vertical AI applications: Sector-specific AI products sold to customers
   - Examples: Healthcare AI (PathAI), legal AI (Harvey), marketing AI
    ↪ (Jasper), financial AI (Bloomberg GPT)
2. Internal adopters: Using AI to optimize own operations, enhance products,
 ↪ improve ads, logistics, customer support
   - Examples: Most retailers, banks, manufacturers, service companies

**CLASSIFICATION RULES:**
- AI Producer = Develops/sells foundational AI technology, infrastructure,
 ↪ or chips FOR OTHERS to build AI
- AI User = Uses AI tools/models internally OR sells AI-enhanced products
 ↪ (but doesn't build core AI infrastructure)
- A firm can be BOTH if it produces AI infrastructure AND uses AI for other
 ↪ purposes
- If unsure, use your knowledge of the company's primary business model
- **Conservative default**: When in doubt, classify as AI User (most firms
 ↪ use AI in some capacity today)

**Transcript excerpt:**
{t}

**OUTPUT FORMAT** (respond with ONLY these two lines, no explanation):
Line 1: isAIProducer: [YES/NO]
Line 2: isAIUser: [YES/NO]

**Examples:**
```

```
- NVIDIA (chip maker): "isAIProducer: YES\nisAIUser: YES"

- Walmart (retailer using AI): "isAIProducer: NO\nisAIUser: YES"

- OpenAI (foundation models): "isAIProducer: YES\nisAIUser: YES"

- Local restaurant chain: "isAIProducer: NO\nisAIUser: YES"
```

IA7 Analyst Forecast Bias

Table IA7.7
Analyst Earnings Forecast Bias: AI Producer

This table reports analyst earnings forecast bias across forecast horizons following the method in Gómez-Cram and Lawrence (2025). Forecast bias is defined as realized earnings per share minus the analyst forecast. Rows report results by forecast horizon. Columns (1)–(3) report forecast bias estimated separately for AI Producer firms, non-AI Producer firms, and the difference between the two groups, respectively. The classification of AI Producer firms is based on large language models reading conference call transcripts to identify whether firms produce AI-related products. For each group, forecast bias is estimated as

$$e_{i,t+h|t} \equiv y_{i,t+h} - \mathbb{F}_t y_{i,t+h} = a_h + u_{i,t+h}, \tag{IA7.1}$$

where $y_{i,t+h}$ denotes realized earnings per share for firm i at horizon h, $\mathbb{F}_t y_{i,t+h}$ is the analyst forecast formed at time t, and a_h captures average forecast bias at horizon h. Standard errors are clustered by firm and calendar year-quarter. The sample period is 2009 to 2024.

	AI Producers I (1)	Non-AI Producers II (2)	AI − Non-AI I - II (3)
Forecast horizon: one quarter			
Bias	8.95	-4.50	13.44
t-statistic	[3.09]	[-3.54]	[4.25]
Observations	4,211	147,520	151,731
Forecast horizon: two quarters			
Bias	7.18	-11.65	18.82
t-statistic	[2.41]	[-9.04]	[5.80]
Observations	4,134	143,285	147,419
Forecast horizon: three quarters			
Bias	4.86	-19.16	24.03
t-statistic	[1.58]	[-14.00]	[7.14]
Observations	4,035	138,170	142,205
Forecast horizon: four quarters			
Bias	2.94	-23.93	26.87
t-statistic	[0.93]	[-16.12]	[7.67]
Observations	3,914	133,120	137,034

IA8 Model Proofs

This appendix provides the detailed derivations for the two-period illustrative model and the dynamic model.

IA8.1 Two-Period Illustrative Model

Posterior mean and variance. The prior and signal are

$$\eta_{i,0} \sim \mathcal{N}(0, \sigma_{\eta i}^2), \qquad s_{i,0} = \eta_{i,0} + u_{i,0}, \qquad u_{i,0} \sim \mathcal{N}(0, \sigma_{ui}^2). \tag{IA8.1}$$

The posterior precision is the sum of the prior and signal precisions:

$$\widehat{\mathbb{V}}_{i,0}^{-1} = \sigma_{\eta i}^{-2} + \sigma_{ui}^{-2}. \tag{IA8.2}$$

Therefore,

$$\widehat{\mathbb{V}}_{i,0} = \left(\sigma_{\eta i}^{-2} + \sigma_{ui}^{-2}\right)^{-1} = (1 - G_{i,0})\sigma_{\eta i}^2, \tag{IA8.3}$$

$$\widehat{\mathbb{E}}[\eta_{i,0} \mid \widehat{\mathcal{I}}_{i,0}] = \widehat{\mathbb{V}}_{i,0}\sigma_{ui}^{-2} s_{i,0} = G_{i,0} s_{i,0}, \tag{IA8.4}$$

where

$$G_{i,0} = \frac{\sigma_{\eta i}^2}{\sigma_{\eta i}^2 + \sigma_{ui}^2}. \tag{IA8.5}$$

This proves (9) through (11). The posterior moments of the terminal dividend are

$$\widehat{\mathbb{E}}[D_{i,1} \mid \widehat{\mathcal{I}}_{i,0}] = \bar{D}_i + G_{i,0} s_{i,0}, \tag{IA8.6}$$

$$\widehat{\mathbb{V}}(D_{i,1} \mid \widehat{\mathcal{I}}_{i,0}) = \widehat{\mathbb{V}}_{i,0} + \sigma_{\varepsilon i}^2 = \widehat{\Omega}_{i,0}. \tag{IA8.7}$$

Portfolio choice and equilibrium price. The investor chooses $Q_{i,0}$ to solve

$$\max_{Q_{i,0}} \quad \widehat{\mathbb{E}}\Big[Q_{i,0}(D_{i,1}-P_{i,0}) \mid \widehat{\mathcal{I}}_{i,0}\Big] - \frac{\gamma}{2}\widehat{\mathbb{V}}\Big(Q_{i,0}(D_{i,1}-P_{i,0}) \mid \widehat{\mathcal{I}}_{i,0}\Big). \tag{IA8.8}$$

Because $P_{i,0}$ and $Q_{i,0}$ are known at date 0, the first-order condition is

$$\widehat{\mathbb{E}}[D_{i,1}-P_{i,0} \mid \widehat{\mathcal{I}}_{i,0}] - \gamma Q_{i,0}\widehat{\Omega}_{i,0} = 0. \tag{IA8.9}$$

The demand is therefore

$$Q_{i,0} = \frac{\widehat{\mathbb{E}}[D_{i,1}-P_{i,0} \mid \widehat{\mathcal{I}}_{i,0}]}{\gamma\widehat{\Omega}_{i,0}}. \tag{IA8.10}$$

Imposing market clearing, $Q_{i,0}=1$, gives

$$\begin{aligned} P_{i,0} &= \widehat{\mathbb{E}}[D_{i,1} \mid \widehat{\mathcal{I}}_{i,0}] - \gamma\widehat{\Omega}_{i,0} \\ &= \bar{D}_i + G_{i,0}s_{i,0} - \gamma\widehat{\Omega}_{i,0}, \end{aligned} \tag{IA8.11}$$

which proves (13). Substitution into the investor's expected return proves (14).

Objective and subjective expected returns. The full-information benchmark additionally observes $\eta_{i,0}$. Therefore,

$$\begin{aligned} \mathbb{E}[D_{i,1}-P_{i,0} \mid \mathcal{I}_{i,0}] &= \bar{D}_i + \eta_{i,0} - P_{i,0} \\ &= \gamma\widehat{\Omega}_{i,0} + \eta_{i,0} - G_{i,0}s_{i,0}, \end{aligned} \tag{IA8.12}$$

which proves (15). For completeness, the objective and subjective AI premiums are

$$\mathrm{Prem}_0^{obj} \equiv \mathbb{E}[D_{A,1}-P_{A,0} \mid \mathcal{I}_{A,0}] - \mathbb{E}[D_{N,1}-P_{N,0} \mid \mathcal{I}_{N,0}], \tag{IA8.13}$$

$$\mathrm{Prem}_0^{sub} \equiv \widehat{\mathbb{E}}[D_{A,1}-P_{A,0} \mid \widehat{\mathcal{I}}_{A,0}] - \widehat{\mathbb{E}}[D_{N,1}-P_{N,0} \mid \widehat{\mathcal{I}}_{N,0}]. \tag{IA8.14}$$

Subtracting the subjective premium from the objective premium gives

$$\mathrm{Prem}_0^{obj} - \mathrm{Prem}_0^{sub} = (\eta_{A,0} - G_{A,0}s_{A,0}) - (\eta_{N,0} - G_{N,0}s_{N,0}), \tag{IA8.15}$$

which is (16).

Forecast errors. Define

$$FE_{i,1} \equiv D_{i,1} - \widehat{\mathbb{E}}[D_{i,1} \mid \widehat{\mathcal{I}}_{i,0}]. \tag{IA8.16}$$

Using (6), (8), and (IA8.6),

$$\begin{aligned} FE_{i,1} &= \eta_{i,0} - G_{i,0}s_{i,0} + \varepsilon_{i,1} \\ &= (1 - G_{i,0})\eta_{i,0} - G_{i,0}u_{i,0} + \varepsilon_{i,1}, \end{aligned} \tag{IA8.17}$$

which proves (17). Differencing across assets gives

$$FE_{A,1} - FE_{N,1} = \mathrm{Gap}_0 + \varepsilon_{A,1} - \varepsilon_{N,1}, \tag{IA8.18}$$

which proves (18).

The two channels. Substituting $s_{i,0} = \eta_{i,0} + u_{i,0}$ into the AI premium gap gives

$$\mathrm{Gap}_0 = (1 - G_{A,0})\eta_{A,0} - (1 - G_{N,0})\eta_{N,0} - G_{A,0}u_{A,0} + G_{N,0}u_{N,0}. \tag{IA8.19}$$

If $\eta_{A,0} = \eta_{N,0} = \eta_0 > 0$ and the realized signal noise is negligible, this expression becomes (19). If $G_{A,0} = G_{N,0} = G_0$ and the realized signal noise is negligible, it becomes (20). These two special cases show why uncertainty and positive expected cash flow shocks interact in generating the AI premium gap.

IA8.2 Dynamic Model: Timing and Kalman Recursion

Suppress the asset index i temporarily. The state and signal equations are

$$\eta_{t+1} = \rho\eta_t + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{IA8.20}$$

$$s_{t+1} = \eta_t + \nu_{t+1}, \qquad \nu_{t+1} \sim \mathcal{N}(0, U). \tag{IA8.21}$$

At date t, before observing s_{t+1},

$$\eta_t \mid \widehat{\mathcal{I}}_t \sim \mathcal{N}(\widehat{\eta}_t, P_t). \tag{IA8.22}$$

Because s_{t+1} provides a noisy signal about η_t, the Kalman gain is

$$G_t = \frac{P_t}{P_t + U}. \tag{IA8.23}$$

After observing s_{t+1}, the investor's updated beliefs about η_t are

$$\begin{aligned} \widehat{\eta}_{t|t+1} &\equiv \widehat{\mathbb{E}}[\eta_t \mid \widehat{\mathcal{I}}_{t+1}] = \widehat{\eta}_t + G_t(s_{t+1} - \widehat{\eta}_t), \\ P_{t|t+1} &\equiv \widehat{\mathbb{V}}(\eta_t \mid \widehat{\mathcal{I}}_{t+1}) = (1 - G_t)P_t. \end{aligned} \tag{IA8.24}$$

The state transition then gives

$$\begin{aligned} \widehat{\eta}_{t+1|t+1} &= \rho\widehat{\eta}_{t|t+1}, \\ P_{t+1|t+1} &= \rho^2 P_{t|t+1} + \sigma_\epsilon^2. \end{aligned} \tag{IA8.25}$$

Using the shorthand

$$\widehat{\eta}_t \equiv \widehat{\eta}_{t|t}, \qquad P_t \equiv P_{t|t}, \tag{IA8.26}$$

gives

$$\widehat{\eta}_{t+1} = \rho \left[\widehat{\eta}_t + G_t (s_{t+1} - \widehat{\eta}_t) \right], \tag{IA8.27}$$

$$P_{t+1} = \rho^2 (1 - G_t) P_t + \sigma_\epsilon^2, \tag{IA8.28}$$

which proves (31) and (32). The order is important: the signal observed at $t+1$ first updates beliefs about η_t, and the updated belief is then propagated to η_{t+1}.

Substituting $s_{t+1} = \eta_t + \nu_{t+1}$ into the mean recursion gives

$$\widehat{\eta}_{t+1} = \rho(1 - G_t)\widehat{\eta}_t + \rho G_t \eta_t + \rho G_t \nu_{t+1}. \tag{IA8.29}$$

At the steady state,

$$G^* = \frac{P^*}{P^* + U}, \qquad P^* = \rho^2 (1 - G^*) P^* + \sigma_\epsilon^2, \tag{IA8.30}$$

which proves (35). Substituting the first expression into the second gives

$$(P^*)^2 + \left[U(1 - \rho^2) - \sigma_\epsilon^2 \right] P^* - \sigma_\epsilon^2 U = 0. \tag{IA8.31}$$

The nonnegative solution is

$$P^* = \frac{1}{2} \left\{ \sigma_\epsilon^2 - U(1 - \rho^2) + \sqrt{[U(1 - \rho^2) - \sigma_\epsilon^2]^2 + 4\sigma_\epsilon^2 U} \right\}. \tag{IA8.32}$$

IA8.3 Dynamic Model: Log Stochastic Discount Factor

The log stochastic discount factor is

$$m_{t+1} = -r_f - \frac{1}{2}\gamma^2 \sigma^2 - \gamma \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim \mathcal{N}(0, \sigma^2). \tag{IA8.33}$$

Its normalization follows from

$$\widehat{\mathbb{E}}_t[\exp(m_{t+1})] = \exp\left(-r_f - \frac{1}{2}\gamma^2\sigma^2 + \frac{1}{2}\gamma^2\sigma^2\right)$$
$$= \exp(-r_f). \tag{IA8.34}$$

Therefore, a risk-free gross return $R_f = \exp(r_f)$ satisfies the Euler equation. This verifies the normalization in (21).

IA8.4 Dynamic Model: Equilibrium Price-Dividend Ratio

Suppress the asset index temporarily and conjecture

$$z_t = A_0 + A_1\widehat{\eta}_t. \tag{IA8.35}$$

Because

$$\widehat{\mathbb{E}}_t[s_{t+1}] = \widehat{\eta}_t, \tag{IA8.36}$$

the Kalman recursion implies

$$\widehat{\mathbb{E}}_t[\widehat{\eta}_{t+1}] = \rho\widehat{\eta}_t. \tag{IA8.37}$$

The expected dividend growth is

$$\widehat{\mathbb{E}}_t[\Delta d_{t+1}] = \mu_d + \widehat{\eta}_t. \tag{IA8.38}$$

Substituting these expressions into the Campbell–Shiller identity gives

$$\widehat{\mathbb{E}}_t[r_{t+1}] = \kappa_0 + (\kappa_1 - 1)A_0 + \mu_d$$
$$+ (1 + \kappa_1 A_1\rho - A_1)\,\widehat{\eta}_t. \tag{IA8.39}$$

In the steady state, the conditional risk adjustment in the Euler equation is independent of $\widehat{\eta}_t$. Therefore,

$$1 + \kappa_1 A_1 \rho - A_1 = 0, \tag{IA8.40}$$

which gives

$$A_1 = \frac{1}{1 - \kappa_1 \rho}. \tag{IA8.41}$$

Restoring the asset index proves (36).

To derive the return innovation, substitute (IA8.29) into the affine price-dividend ratio:

$$z_{t+1} = A_0 + A_1 \rho (1 - G_t) \widehat{\eta}_t + A_1 \rho G_t \eta_t + A_1 \rho G_t \nu_{t+1}. \tag{IA8.42}$$

Substitution into the Campbell–Shiller identity gives

$$\begin{aligned} r_{t+1} = \kappa_0 &+ (\kappa_1 - 1) A_0 + \mu_d \\ &+ \left[\kappa_1 A_1 \rho (1 - G_t) - A_1\right] \widehat{\eta}_t \\ &+ \left(1 + \kappa_1 A_1 \rho G_t\right) \eta_t + \left(1 + \kappa_1 A_1 \rho G_t\right) \nu_{t+1}. \end{aligned} \tag{IA8.43}$$

Define

$$B_t \equiv 1 + \kappa_1 A_1 \rho G_t. \tag{IA8.44}$$

Using $\eta_t = \widehat{\eta}_t + (\eta_t - \widehat{\eta}_t)$, the return becomes

$$\begin{aligned} r_{t+1} = \kappa_0 &+ (\kappa_1 - 1) A_0 + \mu_d \\ &+ \left(1 + \kappa_1 A_1 \rho - A_1\right) \widehat{\eta}_t \\ &+ B_t (\eta_t - \widehat{\eta}_t) + B_t \nu_{t+1}. \end{aligned} \tag{IA8.45}$$

The coefficient on $\widehat{\eta}_t$ is zero under (36). In the steady state, $B_t = B^*$, so restoring the asset index proves (37) and (38). The signal surprise used in this derivation is

$$s_{i,t+1} - \widehat{\eta}_{i,t} = (\eta_{i,t} - \widehat{\eta}_{i,t}) + \nu_{i,t+1}. \tag{IA8.46}$$

Constant restriction for $A_{0,i}$. Under the investor's information set,

$$\widehat{\mathbb{V}}_t(r_{i,t+1}) = (B_i^*)^2(P_i^* + U_i), \qquad -\widehat{\mathrm{Cov}}_t(m_{t+1}, r_{i,t+1}) = \gamma B_i^* \beta_i \sigma^2. \tag{IA8.47}$$

Under conditional normality, the Euler equation implies

$$\kappa_0 + (\kappa_1 - 1)A_{0,i} + \mu_d - r_f = \gamma B_i^* \beta_i \sigma^2 - \frac{1}{2}(B_i^*)^2(P_i^* + U_i). \tag{IA8.48}$$

Solving for $A_{0,i}$ gives

$$A_{0,i} = \frac{\kappa_0 + \mu_d - r_f - \gamma B_i^* \beta_i \sigma^2 + \frac{1}{2}(B_i^*)^2(P_i^* + U_i)}{1 - \kappa_1}. \tag{IA8.49}$$

If the filter is not in its steady state, the intercept is deterministic and time varying. In that case,

$$A_{0,i,t} = \kappa_0 + \kappa_1 A_{0,i,t+1} + \mu_d - r_f - \gamma B_{i,t} \beta_i \sigma^2 + \frac{1}{2}B_{i,t}^2(P_{i,t} + U_i). \tag{IA8.50}$$

The main text focuses on the steady-state solution because the GMM estimation uses steady-state moments.

IA8.5 Dynamic Model: Expected Log Excess Returns and the Gap

Start from the investor's Euler equation,

$$\widehat{\mathbb{E}}_t[\exp(m_{t+1} + r_{i,t+1})] = 1. \tag{IA8.51}$$

Conditional normality implies

$$\widehat{\mathbb{E}}_t[m_{t+1} + r_{i,t+1}] + \frac{1}{2}\widehat{\mathbb{V}}_t(m_{t+1} + r_{i,t+1}) = 0. \tag{IA8.52}$$

Expanding the variance and using the risk-free-rate normalization gives

$$\widehat{\mathbb{E}}_t[r_{i,t+1} - r_f] = -\widehat{\mathrm{Cov}}_t(m_{t+1}, r_{i,t+1}) - \frac{1}{2}\widehat{\mathbb{V}}_t(r_{i,t+1}). \tag{IA8.53}$$

From the return decomposition,

$$-\widehat{\mathrm{Cov}}_t(m_{t+1}, r_{i,t+1}) = \gamma B_i^* \beta_i \sigma^2, \tag{IA8.54}$$

and the subjective and objective conditional return variances are

$$\widehat{\mathbb{V}}_t(r_{i,t+1}) = (B_i^*)^2(P_i^* + U_i), \qquad \mathbb{V}_t(r_{i,t+1}) = (B_i^*)^2 U_i. \tag{IA8.55}$$

Substitution proves (39).

The objective benchmark does not satisfy a separate pricing Euler equation because it is not the expectation of the investor who sets prices. Instead, condition the equilibrium return decomposition on the full information set. The belief error is then known, while the measurement noise has conditional mean zero. Therefore,

$$\mathbb{E}_t[r_{i,t+1} - r_f] = \widehat{\mathbb{E}}_t[r_{i,t+1} - r_f] + B_i^*(\eta_{i,t} - \widehat{\eta}_{i,t}), \tag{IA8.56}$$

which proves (40) and (41). Differencing across $i = A$ and $i = N$ proves (44). A necessary and sufficient condition for a positive AI premium gap is

$$B_A^*(\eta_{A,t} - \widehat{\eta}_{A,t}) > B_N^*(\eta_{N,t} - \widehat{\eta}_{N,t}). \tag{IA8.57}$$

Because $\widehat{\eta}_{i,t}$ is the conditional expectation of $\eta_{i,t}$,

$$\mathbb{E}[\eta_{i,t} - \widehat{\eta}_{i,t}] = 0. \tag{IA8.58}$$

The expected-return gap is therefore zero unconditionally, although it can vary over time along a realized latent-state and signal path. If $U_i = 0$, the date-$(t+1)$ signal reveals $\eta_{i,t}$ exactly. However, it does not make $\eta_{i,t}$ known when the date-t price is set, and $\eta_{i,t+1}$ still contains the new state innovation $\epsilon_{i,t+1}$. Thus, $G_i^* = 1$ does not by itself eliminate the date-t expected-return gap.

IA8.6 Dynamic Model: Forecast Errors

The investor's dividend-growth forecast is

$$\widehat{\mathbb{E}}_t[\Delta d_{i,t+1}] = \mu_d + \widehat{\eta}_{i,t}. \tag{IA8.59}$$

It follows that

$$FE_{i,t+1} = (\eta_{i,t} - \widehat{\eta}_{i,t}) + \beta_i \varepsilon_{t+1} + u_{i,t+1}. \tag{IA8.60}$$

Under the investor's information set,

$$\widehat{\mathbb{E}}_t[FE_{i,t+1}] = 0. \tag{IA8.61}$$

Under the full information set,

$$\mathbb{E}_t[FE_{i,t+1}] = \eta_{i,t} - \widehat{\eta}_{i,t}. \tag{IA8.62}$$

This proves (48) and (49). Because the belief error and measurement noise are independent,

$$\widehat{\mathbb{E}}_t[FE_{i,t+1}^2] = P_i^* + U_i, \tag{IA8.63}$$

$$\text{FEGap}^{MSE} = (P_A^* + U_A) - (P_N^* + U_N), \tag{IA8.64}$$

which proves (50) and (51).

IA8.7 Dynamic Model: Carter–Kohn Simulation Smoother

The lag between the state and signal indices must also be preserved in the smoothing step. Suppress the asset index for clarity. The state-space system is

$$\eta_{t+1} = \rho\eta_t + \epsilon_{t+1}, \qquad s_{t+1} = \eta_t + \nu_{t+1}. \tag{IA8.65}$$

Suppose the observed signal path is $s_{1:T}$. These signals directly measure $\eta_{0:T-1}$, while η_T is a one-step-ahead state conditional on the available signals.

Forward pass. For $t = 0, \ldots, T-1$, after observing s_{t+1}, store

$$\widetilde{\eta}_t \equiv \widehat{\mathbb{E}}[\eta_t \mid s_{1:t+1}] = \widehat{\eta}_t + G_t(s_{t+1} - \widehat{\eta}_t), \tag{IA8.66}$$

$$\widetilde{P}_t \equiv \widehat{\mathbb{V}}(\eta_t \mid s_{1:t+1}) = (1 - G_t)P_t. \tag{IA8.67}$$

Then predict

$$\widehat{\eta}_{t+1} = \rho\widetilde{\eta}_t, \qquad P_{t+1} = \rho^2\widetilde{P}_t + \sigma_\epsilon^2. \tag{IA8.68}$$

At the end of the forward pass,

$$\eta_T \mid s_{1:T} \sim \mathcal{N}(\widehat{\eta}_T, P_T). \tag{IA8.69}$$

Backward sampling. First draw

$$\eta_T^{(j)} \sim \mathcal{N}(\widehat{\eta}_T, P_T). \tag{IA8.70}$$

Then iterate backward for $t = T-1, \dots, 0$. Define

$$J_t = \frac{\widetilde{P}_t \rho}{P_{t+1}}. \tag{IA8.71}$$

The conditional draw is

$$\eta_t^{(j)} \mid \eta_{t+1}^{(j)}, s_{1:T} \sim \mathcal{N}\Big(\widetilde{\eta}_t + J_t \left(\eta_{t+1}^{(j)} - \widehat{\eta}_{t+1}\right),\ \widetilde{P}_t - J_t^2 P_{t+1}\Big). \tag{IA8.72}$$

Repeating this recursion gives an exact draw from

$$\eta_{0:T}^{(j)} \sim p(\eta_{0:T} \mid s_{1:T}). \tag{IA8.73}$$

The model-implied objective-minus-subjective spread at date t compares the smoothed state draw with the investor's date-t belief, which uses signals only through s_t:

$$\Pi_t^{gap,(j)} = B_A^*(\eta_{A,t}^{(j)} - \widehat{\eta}_{A,t}) - B_N^*(\eta_{N,t}^{(j)} - \widehat{\eta}_{N,t}), \qquad t = 1, \dots, T. \tag{IA8.74}$$

It must not be compared with $\widetilde{\eta}_t$, because $\widetilde{\eta}_t$ additionally uses s_{t+1}, which is not available when the date-t price is set. The latent-state innovation and its AI-minus-non-AI difference are

$$\epsilon_{i,t}^{(j)} = \eta_{i,t}^{(j)} - \rho_i \eta_{i,t-1}^{(j)}, \qquad t = 1, \dots, T, \tag{IA8.75}$$

$$\epsilon_{A,t}^{(j)} - \epsilon_{N,t}^{(j)} = (\eta_{A,t}^{(j)} - \rho_A \eta_{A,t-1}^{(j)}) - (\eta_{N,t}^{(j)} - \rho_N \eta_{N,t-1}^{(j)}). \tag{IA8.76}$$

These draws are the inputs to the model-implied AI premium gap and expected cash flow shock gap reported in the simulation analysis.

IA9 GMM Estimation of Structural Parameters

I estimate the structural parameters of the dynamic learning model using GMM. The estimation uses U.S. data from 2009 to 2024 and compares AI and non-AI firms, defined as firms in the highest and lowest *AI Exposure* quintiles. The model is disciplined by seven calibrated parameters and eight empirical moments.

I first calibrate the parameters that can be directly measured from the data. The log risk-free rate is calibrated to the average one year Treasury bill return from the CRSP fixed term indices, giving

$$r_f = 0.0117. \tag{IA9.1}$$

The aggregate shock volatility is calibrated to the annualized volatility of CRSP value-weighted market returns, giving

$$\sigma = 0.1804. \tag{IA9.2}$$

The mean dividend growth rate is calibrated from four-year firm-level dividend growth,

$$G_{n,t} = \sum_{s=1}^{4} \Delta d_{n,t+s}, \tag{IA9.3}$$

winsorized at five percent. I compute firm level annual geometric mean growth rates and then take the pooled average, giving

$$\mu_d = 0.0683. \tag{IA9.4}$$

The Campbell-Shiller constants are calibrated from the sample mean log price-dividend ratio, giving

$$\kappa_0 = 0.0743, \qquad \kappa_1 = 0.9858. \tag{IA9.5}$$

Finally, I estimate market betas separately for AI and non-AI firms:

$$\beta_A = 0.9414, \qquad \beta_N = 0.8835. \tag{IA9.6}$$

The remaining parameters are estimated by GMM. The parameter vector is

$$\theta = \left(\gamma, \rho_A, \rho_N, \sigma^2_{A,\epsilon}, \sigma^2_{N,\epsilon}, \sigma^2_{A,u}, \sigma^2_{N,u}\right)'. \tag{IA9.7}$$

Here γ is the risk aversion coefficient, ρ_i is the persistence of the latent dividend growth state, $\sigma^2_{i,\epsilon}$ is the innovation variance of the latent state, and $\sigma^2_{i,u}$ is the idiosyncratic observation noise variance.

For each candidate parameter vector, I solve the steady-state Kalman filter under the timing of the dynamic model. The total observation noise is

$$U_i = \beta_i^2 \sigma^2 + \sigma^2_{i,u}. \tag{IA9.8}$$

Here, P_i^* is the steady-state variance of $\eta_{i,t}$ conditional on signals observed through date t, before $s_{i,t+1}$ arrives. The steady-state Kalman gain is

$$G_i^* = \frac{P_i^*}{P_i^* + U_i}. \tag{IA9.9}$$

After $s_{i,t+1}$ arrives, the conditional variance of $\eta_{i,t}$ is $(1 - G_i^*)P_i^*$. Propagating this variance through the state transition gives

$$P_i^* = \rho_i^2 (1 - G_i^*) P_i^* + \sigma^2_{i,\epsilon}. \tag{IA9.10}$$

The Kalman gain G_i^* measures how much investors update their beliefs after observing dividend growth. Higher observation noise lowers the gain, so beliefs respond less to new information.

Given the steady-state Kalman objects, the Campbell-Shiller valuation coefficient is

$$A_{1,i} = \frac{1}{1 - \kappa_1 \rho_i}, \tag{IA9.11}$$

and the return transmission coefficient is

$$B_i^* = 1 + \kappa_1 A_{1,i} \rho_i G_i^*. \tag{IA9.12}$$

The term B_i^* captures both the direct cash flow effect and the price effect through belief revision.

I match eight empirical moments. The first two moments are subjective expected excess returns:

$$\mathbb{E}^{\text{sub}}[r_{i,t+1} - r_f] = \gamma B_i^* \beta_i \sigma^2 - \frac{1}{2}(B_i^*)^2 (U_i + P_i^*). \tag{IA9.13}$$

The next two moments are the unconditional variances of the log price-dividend ratio:

$$\mathbb{V}(z_{i,t}) = A_{1,i}^2 \left(\frac{\sigma_{i,\epsilon}^2}{1 - \rho_i^2} - P_i^* \right). \tag{IA9.14}$$

The next two moments are the variances of the *objective-minus-subjective spread*:

$$\mathbb{V}(\text{Gap}_{i,t}) = (B_i^*)^2 P_i^*. \tag{IA9.15}$$

The last two moments are forecast error variances:

$$\mathbb{E}[FE_{i,t+1}^2] = P_i^* + U_i. \tag{IA9.16}$$

The empirical target vector is

$$\widehat{m} = [0.0682, 0.0737, 0.9813, 1.0560, 0.0285, 0.0151, 0.1805, 0.1223]'. \tag{IA9.17}$$

These moments correspond to subjective expected excess returns, log price-dividend ratio variances, objective-minus-subjective spread variances, and forecast error variances for AI and non-AI firms. Let $m(\theta)$ denote the model-implied moment vector. The GMM estimator solves

$$\widehat{\theta} = \arg\min_{\theta \in \Theta} g(\theta)' W g(\theta), \qquad g(\theta) = m(\theta) - \widehat{m}. \tag{IA9.18}$$

I use a diagonal weighting matrix,

$$W = \text{diag}(2, 2, 1, 1, 1, 1, 1, 1). \tag{IA9.19}$$

I put more weight on the two expected return moments because the expected return gap is the main object of interest in the paper. Changing to the efficient weighting matrix, defined as the inverse of the variance-covariance matrix of the sample moments evaluated at a first-step estimate, does not change the qualitative results in the AI and value settings. However, the baseline weights do not always deliver an interior estimate. When they drive the estimator to a corner of the parameter space, as in the dot-com period, I use the efficient weighting matrix instead. The other moments discipline the learning and noise parameters. The parameter bounds are

$$\gamma \in [1, 10], \qquad \rho_i \in [0.50, 0.9999], \tag{IA9.20}$$

$$\sigma^2_{i,\epsilon} \in [10^{-5}, 0.5], \qquad \sigma^2_{i,u} \in [10^{-5}, 0.5]. \tag{IA9.21}$$

These bounds keep the estimates in a reasonable range and ensure that the Kalman filter is well behaved. Since the GMM objective can be non-convex, I use multiple starting values drawn within the parameter bounds and keep the estimate with the lowest objective value. This reduces the concern that the final estimates are driven by a bad local minimum.

Standard errors are computed using the delta method. I approximate the sampling covariance matrix of the eight moments, as diagonal because the moments are constructed from different underlying samples and their joint sampling covariance is unavailable. For the two expected return moments, the diagonal entries are the squared Newey–West standard errors from the corresponding portfolio regressions. For the other six moments, I approximate the sampling variances by $(2m_j^2/(T-1))$, with $(T = 16)$. This approach ignores covariances across moments. The approximation for the six second moments assumes independent and normally distributed annual observations, and the calculation holds the calibrated parameters fixed. The reported standard errors should therefore be interpreted as approximate.

IA10 AI Premium Shock Gap: Actual vs Model

This section explains the steps used to plot Figure 4 of the main paper, which shows the 5-year moving average of the actual *AI premium shock gap* and model-implied cash-flow shocks. For both series I fit an AR(1) with an intercept to the annual gap itself and standardize the residual innovation, so the plotted objects are innovations in the gap rather than levels.

I begin by sorting firms each month into AI Exposure quintiles and keeping only the highest quintile, denoted by $Q5$, and the lowest quintile, denoted by $Q1$. For each month t, I compute value-weighted objective expected excess log returns and subjective expected excess log returns within each quintile. Excess log returns are defined as $\ln(1+r) - \ln(1+r_f)$, and I drop firm-months with non-positive gross returns or non-positive market equity. I use log returns to match the model setup. A firm-month enters only if both the objective and the subjective measure are non-missing, and both sides use the same market-value weights and the same denominator, so the two portfolios are always built from the same firms. As in the AI premium gap estimation, I use both short-term spreads (machine-learning-implied objective expected returns as in Bastianello, 2022 minus target-price-implied subjective expected returns as in Loudis, 2024), and long-term spreads (objective ICC following Gonçalves, 2021 minus subjective ICC as in Eskildsen et al., 2026). Let $r^o_{Q5,t}$ and $r^o_{Q1,t}$ denote the objective expected excess log returns for the AI portfolio (fifth quintile sorted by *AI Exposure*) and the non-AI portfolio (first quintile sorted by AI Exposure), and let $r^s_{Q5,t}$ and $r^s_{Q1,t}$ denote the corresponding subjective expected excess log returns.

In the empirical implementation, I also consider two versions of these expected returns. For the short-term spread, r^o is measured using machine-learning-implied objective expected returns and r^s is measured using subjective expected returns implied by analyst target prices. For the long-term specification, r^o is measured using objective ICC and r^s is measured using subjective ICC. The two versions differ in one further respect. The short-term objective measure is available monthly, so both sides of that spread are built from monthly portfolios and then averaged within the year. The long-term objective ICC is available once per fiscal year, so I build the objective portfolio directly at the annual level from one observation per firm-year, while the subjective side is still built monthly and averaged within the year. The risk-free rate is common to $Q5$ and $Q1$ and cancels in the objective spread, so I do not subtract it on the annual objective side.

I then define the monthly objective AI premium and the monthly subjective AI premium as

$$g^o_t = r^o_{Q5,t} - r^o_{Q1,t}, \qquad g^s_t = r^s_{Q5,t} - r^s_{Q1,t}. \tag{IA10.1}$$

The monthly actual AI premium gap is then

$$g_t = g^o_t - g^s_t. \tag{IA10.2}$$

Because the model-implied series is constructed at the annual frequency, I aggregate the monthly actual AI premium gap to the annual level by taking simple yearly averages. For year y, this gives

$$g^o_y = \frac{1}{12}\sum_{t\in y} g^o_t, \qquad g^s_y = \frac{1}{12}\sum_{t\in y} g^s_t, \qquad g_y = g^o_y - g^s_y. \tag{IA10.3}$$

To isolate the unexpected component, I estimate a single AR(1) process with an intercept for the annual actual AI premium gap:

$$g_y = a + \rho g_{y-1} + \varepsilon_y. \tag{IA10.4}$$

Here, a captures the average level, ρ captures persistence, and ε_y is the innovation. The actual AI premium shock gap is the standardized innovation

$$z_y^{gap} = \frac{\varepsilon_y}{\sigma_\varepsilon}, \tag{IA10.5}$$

where σ_ε is the residual standard deviation of that regression. This captures whether the AI premium gap in a given year is larger or smaller than its own level and persistence would predict.

I fit the AR(1) to the gap rather than to g_y^o and g_y^s separately. Standardizing each innovation by its own residual standard deviation and then differencing them would mix a movement in the objective premium with a differently scaled movement in the subjective premium, so the resulting series would not be the innovation in any single object. Fitting one AR(1) to the gap keeps the shock on one scale and removes the level and the persistence of the object I care about.

I then compare this actual AI premium shock gap, z_y^{gap}, with the model-implied cash flow shock gap. The model-implied series conditions on the realized annual portfolio dividend growth signals for the AI and non-AI portfolios rather than on the simulated signal paths used in the Monte Carlo exercise. I run the Kalman filter at the estimated parameters, initialized at its steady-state conditional variance, and stop after the 2024 signal, so no post-sample dividend enters. I then draw 500,000 Carter–Kohn backward paths of the latent states, compute $\epsilon_{i,t} = \eta_{i,t} - \rho_i \eta_{i,t-1}$ along each path, and take the posterior mean across paths in each year. The model-implied gap is the difference between the AI and non-AI posterior means in each year, and I process it exactly as the actual series, that is, an AR(1) with an intercept over the comparison years followed by standardization of the innovations.

The actual series is standardized by the residual standard deviation of its AR(1) regression and the model series by the sample standard deviation of its innovations. Correlations are invariant to this choice, and the slope changes only by the factor $\sqrt{n/(n-2)}$. Since both objects isolate innovations, this comparison focuses on matching unexpected movements in the actual data with those implied by the model, rather than comparing raw levels that may differ mechanically due to trends or persistence. Given the short annual sample and the noise in year-to-year innovations, I plot the 5-year moving averages of both series to focus on their smoother medium-run comovement. Because the AR(1) specification requires a lag, the first year, 2009, is omitted.

I report 5-year moving averages for two reasons. First, the annual actual AI premium shock gap is noisy because it inherits estimation error from the implied cost of capital measures and portfolio construction. Averaging helps reduce this noise. Second, the model implication concerns a sustained sequence of expected cash flow shocks rather than year-by-year comovement. I first remove an AR(1) component from each annual series and standardize the resulting innovations, and then construct full-window moving averages, keeping only windows that are complete. Because these moving averages overlap, I report Newey–West standard errors for the slope, with the lag length set to the moving-average window minus one. That correction handles the overlap the moving average induces, but it rests on asymptotics that a sample of this length does not deliver, so I read the longest windows as descriptive rather than as inference. These calculations hold the estimated structural parameters fixed.

In Table IA10.8, I compare the actual AI premium shock gap with the model-implied cash flow shock gap over moving-average windows from one to ten years. The correlation is already positive and sizable at the annual frequency, 0.657 for the long-term measure and 0.690 for the short-term measure, and it is highest at the longest windows, reaching 0.931 and 0.938 at the ten-year window. At the five-year window plotted in Figure 4 of the main paper, the correlations are 0.718 for the long-term expected return measure and 0.664

for the short-term measure, and the slope is close to one in both panels, 0.941 and 0.993. These results are consistent with the model capturing a sustained sequence of expected cash flow shocks rather than year-by-year movements in the AI premium shock gap. The longest windows leave six to eight overlapping observations, so I do not read their t statistics as inference.

Table IA10.8
AI Premium Shock Gap: Actual vs Model

This table reports the relation between the actual AI premium shock gap and the model-implied cash flow shock gap. The actual gap is the annual AI premium gap, purged of AR(1) persistence with an intercept and standardized by the residual standard deviation of that regression. The model-implied gap is the posterior mean expected cash flow shock gap from the Carter–Kohn smoother, processed in the same way. I then construct full-window moving averages of both series. The slope coefficient is from a regression of the actual AI premium shock gap on the model-implied cash flow shock gap, and the t statistic uses Newey–West standard errors with the lag length set to the moving-average window minus one. Panel A uses the long-term expected return measure, and Panel B uses the short-term expected return measure. The calculations hold the estimated structural parameters fixed. The moving averages overlap and the longest windows leave six to eight observations, so the t statistics at those windows should be read as descriptive.

Panel A: Long-Term Expected Return Measure

Moving-average window	Observations	Sample period	Correlation	Slope	Newey–West t
One year	15	2010–2024	0.657	0.612	4.252
Two years	14	2011–2024	0.728	0.645	4.415
Three years	13	2012–2024	0.763	0.664	5.424
Four years	12	2013–2024	0.722	0.704	6.091
Five years	11	2014–2024	0.718	0.941	9.396
Six years	10	2015–2024	0.780	0.918	8.096
Seven years	9	2016–2024	0.880	0.974	8.869
Eight years	8	2017–2024	0.895	1.029	21.994
Nine years	7	2018–2024	0.926	1.056	9.755
Ten years	6	2019–2024	0.931	0.845	9.802

Panel B: Short-Term Expected Return Measure

Moving-average window	Observations	Sample period	Correlation	Slope	Newey–West t
One year	15	2010–2024	0.690	0.642	3.536
Two years	14	2011–2024	0.600	0.562	1.845
Three years	13	2012–2024	0.620	0.661	2.146
Four years	12	2013–2024	0.720	0.926	2.931
Five years	11	2014–2024	0.664	0.993	2.001
Six years	10	2015–2024	0.615	0.716	1.668
Seven years	9	2016–2024	0.624	0.657	1.938
Eight years	8	2017–2024	0.706	0.738	2.876
Nine years	7	2018–2024	0.850	0.913	5.831
Ten years	6	2019–2024	0.938	0.704	14.630

References

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, 2024a, Artificial Intelligence and Firms' Systematic Risk, Working paper.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, 2024b, Artificial Intelligence, Firm Growth, and Product Innovation, *Journal of Financial Economics* 151, 103745.

Bastianello, Federico, 2022, Time-Series and Cross-Section of Risk Premia Expectations: Evidence from Financial Analysts, Working paper.

Claus, James, and Jacob Thomas, 2001, Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets, *The Journal of Finance* 56, 1629–1666.

Easton, Peter D., 2004, PE Ratios, PEG Ratios, and Estimating the Implied Expected Rate of Return on Equity Capital, *The Accounting Review* 79, 73–95.

Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel, 2020, Hedging Climate Change News, *The Review of Financial Studies* 33, 1184–1216.

Eskildsen, Marc, Markus Ibert, Theis Ingerslev Jensen, and Lasse Heje Pedersen, 2026, In Search of the True Greenium, *Journal of Financial Economics* forthcoming.

Fahlenbrach, Rüdiger, Kevin Rageth, and René M. Stulz, 2021, How Valuable Is Financial Flexibility When Revenue Stops? Evidence from the COVID-19 Crisis, *The Review of Financial Studies* 34, 5474–5521.

Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan, 2001, Toward an Implied Cost of Capital, *Journal of Accounting Research* 39, 135–176.

Gómez-Cram, Roberto, and Alastair Lawrence, 2025, The Value of Software, *American Economic Review* 115, 3514–3558.

Gonçalves, Andrei S., 2021, The Short Duration Premium, *Journal of Financial Economics* 141, 919–945.

Loudis, Johnathan A., 2024, Stock Price Reactions to the Information and Bias in Analyst-Expected Returns, *The Accounting Review* 99, 281–313.

Loughran, Tim, and Bill McDonald, 2011, When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks, *The Journal of Finance* 66, 35–65.

Mohanram, Partha S., and Dan Gode, 2013, Removing Predictable Analyst Forecast Errors to Improve Implied Cost of Equity Estimates, *Review of Accounting Studies* 18, 443–478.

Ohlson, James A., and Beate E. Juettner-Nauroth, 2005, Expected EPS and EPS Growth as Determinants of Value, *Review of Accounting Studies* 10, 349–365.

Sautner, Zacharias, Laurence van Lent, Grigory Vilkov, and Ruishen Zhang, 2023, Firm-Level Climate Change Exposure, *The Journal of Finance* 78, 1449–1498.

van Binsbergen, Jules H., Xiao Han, and Alejandro Lopez-Lira, 2023, Man versus Machine Learning: The Term Structure of Earnings Expectations and Conditional Biases, *The Review of Financial Studies* 36, 2361–2396.

Zhang, Yingguang, Yandi Zhu, and Juhani T. Linnainmaa, 2025, Man versus Machine Learning Revisited, *The Review of Financial Studies* 38, 3768–3790.