Contextualizing R&D Capitalization

Yiwei Dou
yd18@stern.nyu.edu

Ziqing Wang
zw2361@stern.nyu.edu

Paul Zarowin
pz1@stern.nyu.edu

Department of Accounting
Stern School of Business
New York University

June 2026

Contextualizing R&D Capitalization

ABSTRACT
Capitalizing R&D outlays requires estimates of amortization rates over their useful lives. Prior research derives such rates from a lag structure estimated within an industry-year group that links past R&D expenditures to future benefits, presuming firms in that group share the same R&D payoff timeline. We use R&D disclosures in firms' 10-K filings to contextualize the lag structure estimation and derive a firm-year-specific amortization schedule. Specifically, for each firm-year, we convert R&D disclosures through BERT (Bidirectional Encoder Representations from Transformers) into a high-dimensional context vector that represents their semantic meanings. The relationships between past R&D expenditures and future earnings is modeled as functions of the context vector using artificial neural networks. The resulting R&D capital estimates outperform those from prior research in explaining R&D firms' contemporaneous stock prices and returns, as well as their future stock returns over the next 12 months.

JEL Classifications: C45, D84, G14, M41, O32
Keywords: R&D; Capitalization; Intangible Assets; BERT; Artificial Neural Network; Market Valuation

1. Introduction

The U.S. Generally Accepted Accounting Principles (GAAP) require the recognition of nearly all R&D expenditures as expenses rather than assets, leaving no quantitative trace of value on the financial statements. Yet the stock market capitalizes firms' R&D expenditures, as evidenced by R&D-intensive firms' high market-to-book ratios, and academics have conducted extensive research to estimate the value of firms' R&D capital. Until now, however, research on R&D capitalization has relied exclusively on quantitative information (e.g., R&D expenditures and future benefits). In this paper, we use firms' narrative R&D disclosures to contextualize capitalization of their R&D expenditures.

Our approach builds on Lev and Sougiannis (1996), who apply an Almon lag model to estimate the lag structure linking past R&D expenditures to current operating income, compute R&D amortization rates from the lag-structure coefficients, and use those rates to capitalize R&D expenditures. They adjust reported earnings and book values for R&D capitalization and show that these adjustments help explain stock prices and returns. Their model is estimated annually for firms within a two-digit SIC code industry to obtain an industry-year lag structure, implicitly assuming that the lag structure is the same for all firms in that industry. However, firms in the same industry-year often describe different R&D processes and payoff horizons in their 10-K filings. For example, in the industry of "Scientific Instrument" (two-digit SIC = 38) in year 2006, Mentor Corp. and Keithley Instruments Inc. disclose disparate processes linking their R&D expenditures to their final products (see Appendix B). Mentor explicitly refers to its long R&D timelines, whereas Keithley discloses a shorter R&D horizon as their engineering development activities are directed to partly complement the existing product line. We incorporate the R&D-related *qualitative* information in firms' 10-K disclosures to estimate a *firm-year-specific* lag

structure and associated amortization rates for R&D capitalization. We expect this contextualized R&D capitalization to better capture the underlying innovation value than the Almon lag approach, as reflected in greater explanatory power for stock prices and returns.

While it might seem obvious that adding information improves the estimate of the lag structure (linking R&D expenditures to future earnings) and thereby yields a better estimate of R&D capital, our research is not without tension for several reasons. First, regulators have expressed concern that 10-K disclosures have become redundant, unreadable, and boilerplate (KPMG, 2011; SEC, 1998, 2013), and Dyer, Lang, and Stice-Lawrence (2017) find large sample evidence consistent with the concern. As such, the textual disclosures in 10-K filings may just add noise and offer a limited signal beyond the Almon lag approach. Second, Bertomeu, Cheynel, Liao, and Milone (2025) demonstrate that more input information does not automatically yield better measures, and that machine-learning measures of accounting conservatism based on more input are complements rather than substitutes to existing approaches. Third, disclosure can be too costly to process; the costs of monitoring, acquiring, and analyzing disclosures affect investors' information choices, trading, and market outcomes (Blankespoor, deHaan, and Marinovic, 2020). For all these reasons, it is an empirical question whether contextual information improves the estimates of R&D capital.

We construct a sample of U.S. public firms with meaningful R&D expenditures from 1996-2021, representing over 14,000 firm-year observations. We follow Merkley (2014) to identify sentences in 10-K filings containing R&D-related keywords, like "R&D", "product development", "new patent", etc. Using BERT (Bidirectional Encoder Representations from Transformers), these sentences in a 10-K filing are converted into a 768-dimensional context vector that represents their semantic meanings for a firm-year. The context vector is mapped through two-layer artificial

neural networks into context coefficients that moderate the relationships between past R&D expenditures and future earnings (i.e., the lag structure).

We derive firm-year-specific amortization schedules using the contextualized lag structure and adjust the reported earnings and book values of sample firms for the R&D capitalization (i.e., misstatement due to recognizing all R&D outlays as expenses). Compared with the R&D capitalization under the Almon lag method, we find the following results. First, both the in-sample and out-of-sample performance in explaining current earnings with lagged R&D expenditures improve significantly. The root mean square error (RMSE) declines by 54%. Second, the adjustments to earnings and book values under the contextualized R&D capitalization dominate those under the Almon lag method in explaining concurrent stock prices and returns. Finally, the R&D capital estimated under the contextual approach better explain stock returns in the next 12-months than that under the Almon lag approach.[1] In summary, our evidence indicates that context-based amortization rates are superior to purely quantitative based-rates. It is important to take advantage of narrative disclosures to contextualize the quantitative information in firms' annual reports.

The rest of the paper is organized as follows. Section 2 reviews the literature on estimating R&D capital and using context to interpret quantitative information. Section 3 discusses our methodology. Section 4 describes our data and sample. Section 5 discusses our results. Section 6 concludes.

2. Literature Review

[1] Regardless of whether explaining future returns is due to market inefficiency or to Context reflecting risk (the famous Fama, 1970 joint test), the greater explanatory shows the efficacy of our methodology.

Our paper is at the intersection of two streams of research, that have not been linked before: estimating unrecognized R&D capital and using context to interpret quantitative information.

A robust result in research with U.S. data is that despite almost all R&D costs being expensed for accounting purposes, the stock market capitalizes R&D costs on its own (Sougiannis 1994, Lev and Sougiannis 1996, Chambers, Jennings, and Thompson 1998). For example, firms with high R&D expenditures have high market-to-book ratios, and capitalization-adjusted book values are more closely related to stock market values than actual book values, indicating that market values of R&D firms reflect assets that are missing from their accounting book values. Given this, many studies have estimated the market value of the unrecognized R&D capital. However, most studies with U.S. data assume 100% capitalization and a life of five years (Chan, Lakonishok, and Sougiannis, 2001; Falato, Kadyrzhanova, Sim, and Steri, 2022). Lev and Sougiannis (1996) and Iqbal, Rajgopal, Srivastava, and Zhao (2025) estimate industry-year specific R&D capitalization. Importantly, no paper uses information other than reported R&D expenditures, or accounts for cross-firm heterogeneity. Accounting for such heterogeneity is crucial: Ryu and Zarowin (2025), using UK data under IFRS, which mandates capitalization for development costs that meet certain conditions, show that capitalization rates (the ratio of capitalized R&D expenditures to total R&D expenditures) vary over time and across firms, even within an industry; i.e., R&D capitalization is firm-year-specific.

Merkley (2014) is the first paper to study qualitative R&D disclosures. He finds that firms adjust their R&D disclosures to provide more information about their earnings, and the market finds these disclosures to be informative. His focus is what drives the disclosures, rather than their outcome, so he does not address the issue of R&D capitalization. Aside from the difference in focus, our paper uses current technology that was not available to Merkley.

Recent research uses contemporary techniques to "translate" narrative disclosures in firms' annual reports into conditioning information that can be used for understanding quantitative accounting data. Using recent advances in deep learning Kim and Nikolaev (2024) use MD&A text as context to interpret reported accounting numbers, where context is based on the interaction between the numeric and textual information.[2] Interaction (between quantitative and qualitative information) is crucial, because the stand-alone informational value of each one, and even the combined but non-interactive information value, is limited.

For example, Kim and Nikolaev (2024) find that models that use only quantitative or qualitative information alone accurately predict the direction of future earnings changes 55% - 56% of the time, while models that combine both sets of information but not interactively (contextually), have a prediction accuracy of 58%. Their contextual (interactive) model has an accuracy of 65%. Thus, while narrative information by itself is informative, its primary information value is from interactions with the numbers; i.e., context.

Importantly for our research, Kim and Nikolaev find that context matters most when quantitative data are least informative. This is exactly the case with R&D, because U.S. GAAP's full expensing makes R&D a "black box". Moreover, R&D investments are primarily firm-specific, often with little marketable value outside of the firm conducting the R&D itself, making the firm's narrative disclosures especially meaningful. Ours is the first paper to exploit context in the R&D setting.

We start with Lev and Sougiannis' (1996) industry-year model, which is till the archetype for R&D capitalization, and we modify their Almon lag structure (i.e., weights on the lagged R&D expenditures) by conditioning the weights on the context. We find that our context-based measure

[2] We discuss their methods in more detail in Section 3.

dominates quantitative based measures in explaining the link between R&D expenditures and future earnings, and in how those expenditures map into R&D capital. Thus, our results can be useful to academics, analysts, and regulators, who care about the value relevance of R&D expenditures.

3. Methodology and Research Design

In this section, we describe our methodology for estimating R&D capitalization measures. We first discuss the benchmark R&D capitalization model, which follows Lev and Sougiannis (1996) and estimates the relation between current operating income and lagged R&D expenditures using an Almon lag structure. We then introduce our contextualized R&D capitalization model, which allows firm-year-specific R&D disclosure context to shape the estimated coefficients on lagged R&D expenditures. Next, we describe the estimation procedure used to train, validate, and test the contextualized model. Finally, we explain how the estimated amortization schedules are used to construct R&D amortization rates, R&D capital, adjusted earnings, and capitalization-based misstatement measures for the empirical tests.

3.1. Benchmark R&D Capitalization Model

We begin with the benchmark R&D capitalization model developed by Lev and Sougiannis (1996). The model estimates R&D capital from the relation between current operating income and past R&D expenditures. The underlying idea is that R&D expenditures generate benefits over multiple future periods; therefore, the pattern in which lagged R&D expenditures explain current operating income can be used to infer the useful life and amortization pattern of R&D capital.

Following Lev and Sougiannis (1996), we estimate the following R&D–earnings relation:

$$OI/Sale_t = \beta_0 + \beta_1 TA/Sale_{t-1} + \beta_2 AD/Sale_{t-1} + \sum_{k=0}^{9} \gamma_k RD/Sale_{t-k} \quad (1)$$

where OI_{it} is operating income before depreciation, advertising expense, and R&D expense for firm i in year t; $Sale_{it}$ is sales; $TA/Sale_{t-1}$ is the beginning-of-year tangible assets scaled by sales; $AD/Sale_{t-1}$ is the beginning-of-year advertising expense scaled by sales; $R\&D/Sale_{t-k}$ is R&D expenditure in year $t-k$ scaled by sales. The model is estimated at the two-digit SIC industry-year level. Because lagged R&D expenditures are highly persistent, we follow Lev and Sougiannis (1996) and use an Almon lag structure[3] to impose a polynomial restriction on the R&D lag coefficients, thereby reducing the number of parameters to estimate. To address potential simultaneity between R&D expenditures and current operating performance, we estimate the model using instrumental variables, following Lev and Sougiannis (1996).[4]

For each test year t, we compute the benchmark Almon lag amortization schedule using IV regressions estimated over the prior three years, $t-3$ through $t-1$. Specifically, we first estimate the IV Almon lag model separately for each of the prior three years and then average the estimated R&D lag coefficients across these years. Let $\widehat{\gamma_{k,t}}$ denote the average estimated R&D lag coefficients across these years from lag k used for test year t. These averaged coefficients capture the estimated contribution of R&D expenditures from each lag to current operating income.

[3] An Almon lag is a distributed-lag approach that represents the sequence of lag coefficients as a smooth function of the lag length.

[4] Following Lev and Sougiannis (1996), we use a two-stage instrumental-variable procedure to address potential endogeneity in firm-level R&D expenditures. In the first stage, firm-level R&D intensity, $R\&D/Sale_{t-k}$ is regressed on the average R&D intensity of other firms in the same four-digit SIC industry. In the second stage, the fitted R&D intensity from the first stage replaces the actual R&D intensity in the Almon lag regression in equation (1). The instrument is intended to capture industry-level R&D activity that is correlated with a firm's own R&D spending but less likely to reflect firm-specific shocks to current operating income.

We then convert the averaged R&D lag coefficients into amortization rates. For each lag k, the amortization rate is calculated as the share of total estimated R&D benefits attributed to that lag:

$$\delta_{k,t} = \frac{\widehat{\gamma_{k,t}}}{\sum_{k=0}^{9} \widehat{\gamma_{k,t}}} \quad (2)$$

Thus, each amortization rate represents the share of total estimated benefits from R&D spending that is realized in a particular year.[5] Lev and Sougiannis use this ratio to convert estimated R&D lag coefficients into annual amortization rates, which are then used to compute R&D amortization and R&D capital. This procedure ensures that the benchmark R&D capitalization measures are constructed using only information available before the test year.

The key limitation of this benchmark model is that it produces an industry-year amortization schedule. Firms in the same two-digit SIC industry and year are assigned the same R&D payoff pattern, regardless of differences in their innovation strategies, development timelines, regulatory environment, or product life cycles. This restriction is central to our research design. In the next subsection, we introduce a contextualized R&D capitalization model that relaxes this assumption by allowing firm-year-specific R&D disclosure context to shape the estimated coefficients on lagged R&D expenditures.

3.2. Contextualized R&D Capitalization Model

The benchmark model assumes that firms in the same industry-year share the same R&D amortization schedule. Our maintained assumption is that firms' narrative R&D disclosures

[5] In deriving the benchmark amortization schedule, we follow Lev and Sougiannis (1996) and exclude statistically insignificant R&D lag coefficients from the calculation of total R&D benefits.

contain qualitative information about the nature, timing, and expected payoff horizon of their R&D activities. This information may not be sufficiently informative on its own, but it can help interpret the economic meaning of reported R&D expenditures. For example, two firms may report similar levels of R&D spending, but the future benefits of those expenditures may differ substantially if one firm describes long clinical trials and regulatory approval processes while another describes near-term product development and commercialization. In this sense, the R&D disclosure context can help determine how past R&D expenditures map into current and future operating income.

Earlier textual-analysis approaches in accounting and finance commonly use readability measures, dictionary-based word lists, word counts, or document similarity to capture specific textual attributes, such as annual report readability, financial tone and uncertainty, changes in MD&A disclosure, or competition-related disclosure (Li, 2008; Loughran and McDonald, 2011; Brown and Tucker, 2011; Li, Lundholm, and Minnis, 2013). Although these approaches are useful, they may not fully capture how words are used together in context or the multidimensional semantic content of narrative disclosures. For example, references to "development," "approval," or "commercialization" may have different implications for R&D payoff horizons depending on whether the firm discusses clinical trials, regulatory review, product upgrades, or near-term product launches.

Recent language models enable the numerical representation of this contextual meaning. An embedding is a vector representation of text that summarizes information about the words and phrases in a passage. Earlier embedding models, such as Word2Vec[6], assign words numerical

[6] Word2Vec uses a neural network trained on a large text corpus to learn word vectors, or embeddings. These embeddings represent words in a numerical space where words used in similar contexts are located closer to one another. A common illustration is that vector operations over embeddings can capture semantic relationships; for example, the vector for "king" minus "man" plus "woman" is approximately close to the vector for "queen" (Mikolov, Sutskever, Chen, Corrado, and Dean, 2013)

representations based on their usage in a large corpus. BERT, viewed as a breakthrough in Natural Language Processing and conceptually similar to ChatGPT (both rely on transformer architecture), extends this idea by generating context-dependent representations. The same word can receive different representations depending on the words around it, as BERT uses a self-attention mechanism that encodes a word's meaning depending on where and how these words are used (Devlin, Chang, Lee, and Toutanova, 2019). BERT is pretrained on a large corpus of English Wikipedia texts and can process relatively long text of up to 512 tokens. Recent accounting and finance studies apply BERT-based or related language-model approaches to corporate disclosures and financial text, including sentiment analysis and the contextual interpretation of accounting numbers (Siano and Wysocki, 2021; Huang, Wang, and Yang, 2023; Kim and Nikolaev, 2024). Following this approach, we pass each firm-year's R&D disclosure context through BERT and use mean pooling to construct a 768-dimensional embedding. Mean pooling averages the BERT-generated embeddings across the tokens in the disclosure, producing a single vector that summarizes the disclosure's contextual meaning. We use this embedding not as a stand-alone disclosure measure, but as input that enables the model to estimate how the relationship between past R&D expenditures and current operating income varies across firms and years.

To allow disclosure context to shape the interpretation of R&D expenditures, we estimate a contextualized version of the R&D–earnings relation. Let $R\&D\ Context_{it}$ denote the firm-year-specific R&D disclosure context vector extracted from firm i's 10-K filing in year t. Specifically, $Context_{i,t}$ is constructed from R&D-related sentences identified using R&D keywords following Merkley (2014), and then encoded using BERT. We describe the extraction of R&D disclosures and the construction of this context vector in Section 4. The contextualized model allows the coefficients on lagged R&D expenditures to vary with $Context_{i,t}$.

$$OI/Sale_t = \beta_0 + \beta_1 TA/Sale_{t-1} + \beta_2 AD/Sale_{t-1} + \gamma_0(R\&D\ Context_{it}) + \sum_{k=0}^{9} \gamma_{1k}(R\&D\ Context_{it}) RD/Sale_{t-k} \quad (3)$$

The key difference between Equation (3) and the benchmark model in Equation (1) is that the R&D lag coefficients are firm-year-specific. In the benchmark model, γ_k is estimated at the industry-year level and is the same for all firms in the same industry-year. In the contextualized model, $\gamma_{1k}(R\&D\ Context_{it})$ varies with the firm's R&D disclosure context, allowing the contribution of lagged R&D expenditures to operating income to differ across firms and years. The model also includes $\gamma_0(R\&D\ Context_{it})$, which captures the direct context-dependent component of operating income. Thus, the same amount of lagged R&D expenditure can receive different weights depending on the information contained in the firm-year R&D disclosure. Figure 1 Panel A illustrates the overall structure of Equation (3). As in the benchmark specification, the model includes tangible assets and advertising expense with fixed linear coefficients. The key difference is that the R&D disclosure context is used to generate both the direct context component and the coefficients on lagged R&D expenditures.

To estimate these context-dependent coefficients, we use an artificial neural network (ANN). A neural network is a flexible function that maps inputs into outputs through layers of connected units, often called neurons. Information from the input layer propagates through hidden layers, where each layer combines information from the previous layer and applies nonlinear transformations before passing the signal forward. An ANN refers to the computational implementation of this structure, with the weights connecting layers learned from the data. Following Kim and Nikolaev (2024), we use ANNs to allow narrative disclosure context to shape the interpretation of quantitative accounting information. In our setting, the quantitative information is the firm's history of R&D expenditures, and the ANNs serve as coefficient-

generating functions. They map the R&D context embedding into the context-dependent coefficients $\gamma_0(R\&D\ Context_{it})$ and $\gamma_{1k}(R\&D\ Context_{it})$ in Equation (3). Specifically, we estimate a coefficient-generating ANN with 11 parallel output functions, each with the structure shown in Figure 1, Panel B. One output generates the direct context component $\gamma_0(R\&D\ Context_{it})$ and ten outputs generate the R&D lag coefficients $\gamma_{1k}(R\&D\ Context_{it})$ for $k = 0, \ldots, 9$. In each neural network, the R&D context embedding (768 dimensions) enters the input layer, passes through a hidden layer with a ReLU[7] activation function, and produces the corresponding context-dependent coefficient in the output layer. The output layer uses a linear activation function for $\gamma_0(R\&D\ Context_{it})$ and an exponential activation function for the R&D lag coefficients $\gamma_{1k}(R\&D\ Context_{it})$, which imposes a positivity restriction on the estimated R&D payoff weights. These R&D lag coefficients then weight the firm's lagged R&D expenditures and form the R&D component of predicted operating income.

This design embeds the economic structure of Lev and Sougiannis (1996)'s model into a flexible contextual model. Instead of using the ANN as a black box to predict operating income directly, we use it to estimate the coefficients that determine how past R&D expenditures contribute to current operating income. The resulting coefficients, therefore, have a direct accounting interpretation. They describe the firm-year-specific R&D payoff pattern implied by the R&D disclosure context. We then use these estimated payoff patterns to construct firm-year-specific R&D amortization schedules.

3.3. Estimation Procedure

[7] ReLU, or rectified linear unit, is a commonly used nonlinear activation function in neural networks, defined as $f(x) = max\ (0, x)$. This nonlinearity allows the network to capture flexible relations between the input embedding and the output coefficient.

We estimate the contextualized R&D capitalization model using a rolling out-of-sample design. Following the rolling-window approach in Chen, Cho, Dou, and Lev (2022), for each test year $t = 1999, \ldots, 2021$, we estimate a year-specific model using only information available before year t. Specifically, we train the ANN model using observations from years $t-3$ and $t-2$, validate the model using observations from year $t-1$, and then apply the optimal model to observations in year t. This rolling design ensures that the context-dependent coefficients used in the empirical analysis are strictly generated out of sample. This timing also makes the ANN-based estimates comparable to the benchmark Almon lag estimates, because both methods construct year-t capitalization measures using information from years $t-3$ through $t-1$, though they use this prior information differently.

For each test year, the ANN is estimated using the global sample rather than separately by industry. In principle, estimating the contextualized model within each industry could allow the mapping between R&D disclosure context and R&D payoff patterns to differ across industries, and would be more closely parallel to the benchmark Almon lag approach, which is estimated at the two-digit SIC industry-year level. In practice, however, industry-specific estimation would substantially reduce the training sample size in each rolling window, making the high-dimensional contextual model difficult to estimate reliably. Therefore, we estimate a global ANN that uses the full training sample to learn how R&D disclosure context maps into R&D payoff patterns, while still allowing the resulting coefficients to vary across firms and years through $Context_{it}$.

The model is trained to minimize the mean squared error (MSE) between actual operating income scaled by sales and the predicted value implied by Equation (3). For each test year, we use the validation year to select the optimal model specification and training epoch. Specifically, we conduct a grid search over hidden dimensions, dropout rates, and learning rates, and select the

model with the lowest validation MSE.[8] We then apply the selected model to the test-year observations to generate the context-dependent coefficients $\gamma_0(R\&D\ Context_{it})$ and $\gamma_{1k}(R\&D\ Context_{it})$ for $k = 0, \dots, 9$. The coefficients generated from this rolling estimation procedure serve as the inputs to the construction of the ANN-based capitalization measures described in Section 3.4.

The model is trained using mini-batch gradient-based optimization. Gradient-based optimization updates the model parameters in the direction that reduces the prediction error. In our setting, after the ANN predicts operating income for a group of firm-year observations, the algorithm computes how the prediction error changes with respect to each model parameter and updates the parameters to reduce the mean squared error. We use mini-batches rather than the full training sample in each update. Specifically, each gradient update is computed using 256 firm-year observations. This batch size is commonly used in neural-network training because it balances computational efficiency with stable gradient updates.

To mitigate overfitting, we use dropout and early stopping. Dropout is a regularization technique that randomly sets a fraction of hidden-layer units to zero during training, thereby preventing the model from relying too heavily on any particular hidden unit and helping the learned relation generalize out-of-sample (Srivastava, Hinton, Krizhevsky, Sutskever, and Salakhutdinov, 2014). Similar regularization is used in finance applications of deep learning; for example, Chen, Pelger, and Zhu (2024) use dropout to regularize deep neural networks in their asset-pricing model. In our setting, dropout is useful because the model maps high-dimensional BERT-based R&D disclosure embeddings into R&D payoff coefficients. For example, a dropout rate of 0.1 randomly

[8] The grid search considers hidden dimensions of $\{16, 32, 64\}$, dropout rates of $\{0, 0.1, 0.2\}$, and learning rates of $\{10^{-3}, 5 \times 10^{-4}, 10^{-4}\}$. For each hyperparameter combination, the model is trained for a maximum of 200 epochs and stops early if validation MSE does not improve by at least (10^{-6}) for 10 consecutive epochs.

omits 10% of the hidden-layer units during each training update. We select the optimal dropout rate using the validation sample.

For each hyperparameter combination, the model is trained for up to 200 epochs. An epoch is one pass through the training sample. In practice, the model usually stops before reaching this maximum, often around 100 to 150 epochs. Specifically, after each epoch, we compute the validation MSE. If validation MSE does not improve by at least (10^{-6}) for 10 consecutive epochs, training stops. We retain the model checkpoint from the epoch with the lowest validation MSE. We then select the hyperparameter combination with the lowest validation MSE and apply the selected model to the test-year observations. Appendix C provides an illustrative training diagnostic for the 2021 test-year model, with 2021 being the last year in our test sample. The figure shows that both training and validation MSE decline sharply in early epochs and then stabilize.

3.4 Constructing R&D Capitalization Measures

We use the estimated R&D payoff patterns from the benchmark Almon lag model and the contextualized ANN model to construct parallel R&D capitalization measures. This allows us to compare the benchmark industry-year amortization schedule with the firm-year-specific amortization schedule implied by the R&D disclosure context. Let $m \in \{Almon\ lag, ANN\}$ denote the amortization method. For each method, we use the estimated R&D payoff coefficients to compute amortization rates, $\delta_{i,t,k}^{m}$, for firm i and year t and for each R&D lag $k = 0, \ldots, 9$. For the Almon lag method, the estimated coefficients are the industry-year coefficients described in Section 3.1. For the ANN method, the estimated coefficients are the firm-year-specific coefficients generated by the contextualized ANN model described in Sections 3.2 and 3.3.

Specifically, for each method m, we calculate the amortization rate for lag k as the fraction of total estimated R&D benefits attributable to that lag:

$$\delta_{i,t,k}^{m} = \frac{\widehat{\gamma_{i,t,k}^{m}}}{\sum_{l=0}^{9} \widehat{\gamma_{i,t,l}^{m}}} \tag{4}$$

where $\gamma_{i,t,k}^{\widehat{Almon\ lag}}$ is the benchmark industry-year coefficient assigned to firm i in year t, and $\widehat{\gamma_{i,t,k}^{ANN}} = \widehat{\gamma_{1k}}(Context_{it})$. Thus, for the Almon lag method, amortization rates vary at the industry-year level. For the ANN method, amortization rates vary at the firm-year level because they are functions of firm-year R&D disclosure context.

Following Lev and Sougiannis (1996), annual R&D amortization is calculated as the weighted sum of current and lagged R&D expenditures:

$$RA_{it}^{m} = \sum_{k=0}^{9} \delta_{i,t,k}^{m} RD_{it-k} \tag{5}$$

where RA_{it}^{m} is R&D amortization under method m, and RD_{it-k} is the firm i's R&D expenditure in year $t-k$. We then define earnings under R&D capitalization as reported earnings plus current R&D expense, less estimated R&D amortization:

$$X_{it}^{C,m} = X_{it}^{E} + RD_{it} - RA_{it}^{m} \tag{6}$$

where $X_{i,t}^{E}$ is the reported earnings under GAAP and $X_{i,t}^{C,m}$ is the earnings under R&D capitalization using method m. The capitalization-based earnings misstatement is defined as the difference between earnings under capitalization and reported earnings:

$$X_{it}^{C,m} - X_{it}^{E} = RD_{it} - RA_{it}^{m} \tag{7}$$

This measure captures the difference between the current R&D expense recognized under GAAP and the amortization expense implied by the estimated R&D capitalization schedule. When $RD_{it} > RA_{it}^{m}$, reported earnings are lower than earnings under R&D capitalization, when $RD_{it} <$

RA_{it}^{m}, reported earnings are higher than earnings under R&D capitalization. R&D capital (book value misstatement) is calculated as the unamortized portion of current and past R&D expenditures:

$$RDC_{it}^{m} = BV_{it}^{C,m} - BV_{it}^{E} = \sum_{k=0}^{9} RD_{it-k}\left(1 - \sum_{j=0}^{k} \delta_{i,t,j}^{m}\right) \tag{8}$$

After constructing these R&D capitalization measures, we examine their value relevance using contemporaneous return and price tests. Following Lev and Sougiannis (1996), and consistent with Kothari and Zimmerman (1995), we use both return and price specifications because they provide complementary evidence. Return models examine whether the capitalization-based earnings misstatement explains changes in firm value, while price models examine whether the misstatement and R&D capital are reflected in share prices. We use these tests to assess whether the ANN-based capitalization measures are more informative than the benchmark Almon lag-based measures. The return specification relates annual stock returns to reported earnings and the R&D misstatement measure:

$$R_{it} = \alpha_1 + \beta_1 X_{it}^{E} + \gamma_1 (X_{it}^{C,m} - X_{it}^{E})/P_{it-1} + \phi_k + \psi_t + u_{it} \tag{9}$$

where R_{it} is the annual stock return from nine months before the fiscal t year-end through three months after it, P_{it-1} is the share price of firm i measured three months after the prior fiscal year-end. ϕ_k denotes SIC2 fixed effects, and ψ_t denotes year fixed effects. We also estimate a return specification that includes changes in reported earnings and changes in the R&D misstatement measure. This specification follows the return-model intuition that stock returns may reflect both the level of earnings-related information and changes in that information over the return window:

$$\begin{aligned} R_{it} = \alpha_2 + \beta_2 X_{it}^{E}/P_{it-1} + \gamma_2 \Delta X_{it}^{E}/P_{it-1} + \delta_2 (X_{it}^{C,m} - X_{it}^{E})/P_{it-1} \\ + \Omega_2 \Delta (X_{it}^{C,m} - X_{it}^{E})/P_{it-1} + \phi_k + \psi_t + u_{it} \end{aligned} \tag{10}$$

For the price tests, we estimate specifications that relate share prices to reported earnings and the capitalization-based earnings misstatement. Equation (12) further includes ($BV^{C}_{it} - BV^{E}_{it}$), which captures the book-value adjustment from capitalizing past R&D expenditures. This allows the price test to examine whether share prices reflect both the earnings adjustment and the unrecognized R&D capital implied by the capitalization procedure:

$$P_{it} = \alpha_3 + \beta_3 X^E_{it} + \gamma_3 (X^{C,m}_{it} - X^E_{it}) + \phi_k + \psi_t + u_{it} \quad (11)$$

$$P_{it} = \alpha_4 + \beta_4 X^E_{it} + \gamma_4 (X^{C,m}_{it} - X^E_{it}) + \Omega_4 (BV^{C,m}_{it} - BV^E_{it}) + \phi_k + \psi_t + u_{it} \quad (12)$$

where P_{it} is the share price of firm i three months after the fiscal year-end.[9]

Equations (9)-(12) provide the main validation tests for the capitalization-based misstatement measures. If the ANN-based amortization schedule better captures the timing of R&D benefits, then the ANN-based misstatement measure should have stronger explanatory power for returns and prices than the Almon lag-based measure. Consistent with Lev and Sougiannis (1996), we also expect the coefficient on the misstatement measure to be positive, because a larger capitalization-based misstatement reflects greater net R&D investment that is expected to generate future benefits.

We also examine whether the constructed R&D capital measures are associated with subsequent stock returns. This analysis follows the intertemporal analysis of Lev and Sougiannis (1996), which evaluates whether investors fully incorporate the value relevance of R&D capitalization contemporaneously. If R&D capital predicts subsequent returns, it may suggest that investors do not fully recognize its implications when the information becomes available, or that

[9] All right-hand-side variables in the return specifications are deflated by P_{it-1}, the share price measured three months after the prior fiscal year-end. In the price specifications, earnings, misstatement, and R&D capital are expressed on a per-share basis.

R&D capital captures an additional risk factor. In our setting, we estimate the following specification:

$$R_{it+1} = \alpha_5 + \theta \ln((RDC/M)_{it}^{m}) + \Gamma_1\beta_{it} + \Gamma_2 \ln(M_{it}) + \Gamma_3 \ln(B/M_{it}) + \Gamma_4 \ln(A/B_{it}) + \Gamma_5 (E(+)/M_{it}) + \Gamma_6 ((E/M\ dummy)_{it})\phi_k + \psi_t + u_{it} \quad (13)$$

where R_{it+1} is the future 12-month stock return measured from months 4 through 15 after the fiscal year-end. Three months are allowed for the public disclosure of the annual financial statements by all sample firms. The control variables include systematic risk CAPM beta (β_{it}), market capitalization ($ln(M_{it})$), book-to-market ratio ($ln(B/M_{it})$), leverage ($ln(A/B_{it})$), earnings-to-price ratio ($(E(+)/M)_{it}$), and a negative earnings indicator ($(E/M\ dummy)_{it}$). This test evaluates whether the ANN-based R&D capital measure better captures economically relevant R&D assets associated with subsequent stock returns. Consistent with Lev and Sougiannis (1996), we expect θ to be positive if R&D capital captures future benefits that are not fully reflected in contemporaneous prices or captures risk associated with R&D-intensive firms. Across Equations (9)–(13), all regressions include SIC2 fixed effects and year fixed effects, and t-statistics are based on standard errors clustered at the firm level.

Finally, we compare the ANN-based and Almon lag-based measures using two approaches. First, we use Vuong tests to compare the relative fit of specifications estimated separately with the ANN-based and Almon lag-based measures. Prior accounting studies use Vuong tests to compare the explanatory power of alternative accounting measures or valuation models in non-nested specifications (e.g., Dechow, 1994; Jennings, Simko, and Thompson, 1996). Second, we estimate horse-race specifications in which the Almon lag-based and ANN-based measures are included jointly in the same regression. This design allows us to examine whether the ANN-based measure contains incremental information beyond the benchmark Almon lag-based measure, similar to prior accounting studies that compare competing predictors within a common empirical

specification (e.g., Ball and Nikolaev, 2022). Together, these tests allow us to assess whether incorporating R&D disclosure context improves the measurement of R&D capital and capitalization-based earnings misstatement.

4. Data and Sample

4.1 Data Sources and Variable Construction

We construct our sample by combining accounting data, stock return and price data, and textual disclosure data from firms' 10-K filings. Accounting variables are obtained from Compustat. We use Compustat to construct operating income, sales, tangible assets, advertising expense, R&D expenditures, earnings per share, book value, and other control variables. We use CRSP to construct stock returns, share prices, and CAPM betas, and match CRSP to Compustat using firm identifiers.

We obtain 10-K filing text from BeanCounter, a large-scale dataset of business-oriented disclosures constructed from EDGAR filings (Wang and Levy, 2024). BeanCounter provides cleaned filing text and metadata, which facilitates large-scale textual analysis of corporate disclosures. We use BeanCounter rather than extracting filings directly from EDGAR because its pipeline applies a consistent extraction and cleaning procedure that removes tables and repetitive header content while preserving relevant narrative formatting. This reduces noise from raw HTML and variation in filing formats.

As discussed in Section 3.2, the contextualized ANN model requires a firm-year-specific R&D disclosure context vector. To construct this vector, we first identify R&D-related sentences in 10-K filings using the narrative R&D disclosure keywords from Merkley (2014).[10] We then

[10] Appendix A of Merkley (2014) provides the narrative R&D disclosure keyword list used to identify R&D-related disclosure. We apply these keywords to firms' 10-K filings to extract sentences discussing R&D activities.

aggregate the extracted R&D-related sentences at the firm-year level. These R&D-related sentences are intended to capture managers' narrative discussion of the nature, stage, and expected payoff horizon of firms' R&D activities. On average, firms have 36.9 R&D-related sentences in the training sample and 39.7 in the test sample, suggesting that the disclosures provide sufficient textual context for estimating firm-year-specific R&D payoff patterns.

We convert the extracted R&D disclosure text into a numerical representation using BERT. Specifically, we pass each firm-year's R&D disclosure through BERT and use mean pooling to construct a 768-dimensional embedding. Mean pooling averages the BERT-generated token embeddings across the input text, producing a vector that summarizes the contextual meaning of the firm-year R&D disclosure. This embedding is used as the $Context_{i,t}$ input in the contextualized ANN model described in Section 3.2.

4.2 Sample Selection and Descriptive Statistics

Table 1 summarizes the sample selection procedure. We begin with Compustat firm-year observations from 1996 to 2021 and sequentially exclude observations with negative sales, missing four-digit SIC codes, and large mergers that account for 50% or more of annual sales. We then require firms to have non-missing R&D expenditures over the past 10 years and R&D-to-sales of at least 2%, because constructing R&D capitalization measures requires a history of R&D investment. We further require non-missing R&D disclosure in 10-K filings and non-missing earnings and other control variables. These filters result in a sample of 14,535 firm-year observations.

From this sample, we construct the 1996–2020 training sample for model estimation and training and the 1999–2021 test sample for out-of-sample evaluation and regression analyses. The

test sample begins in 1999 because the construction of year t measures requires information from prior years. For the benchmark Almon lag method, the amortization schedule for year t is constructed using coefficient estimates from the prior three years. For the ANN method, the rolling estimation procedure uses years $t-3$ and $t-2$ for training and year $t-1$ for validation before generating coefficients for year t. The resulting training sample contains 13,973 firm-year observations, and the resulting test sample contains 13,306 firm-year observations.

5. Results

Table 2, Panel A presents descriptive statistics for the variables used in the model estimation and training sample. The sample consists of 13,973 firm-year observations from 1996 to 2020. The median operating income scaled by sales is 0.21, and the median R&D-to-sales ratio is 0.10 across the current and lagged R&D variables. Panel B reports model performance in explaining operating earnings, as measured as root mean square error (RMSE). The ANN model consistently yields lower RMSE than Lev and Sougiannis' (1996) Almon lag model in both the training periods (i.e., in-sample) and the test periods (i.e., out-of-sample). On average, the out-of-sample RMSE is 1.5784 for the ANN model, which is 54% lower than 3.4670 of the Almon Lag model. This is our first piece of evidence on the importance of context in the R&D setting.

Next, we use the Almon and ANN lag structures to compare reported (full expensing) R&D expense vs R&D expense based on capitalization and amortization, in explaining current prices and returns. Like Lev and Sougiannis (1996), we define *Misstatement* as the difference between the reported R&D expense and "capitalize and amortize" R&D expense. Table 3, Panel A presents descriptive statistics for the variables used in the return and price analyses within the 1999–2021 test sample. The median annual stock return is 0.05. Because the return

specifications use variables deflated by beginning-of-year share price, we focus on the deflated earnings and misstatement variables when describing the return-test variables. The deflated misstatement measures, $(X^{C}_{it} - X^{E}_{it})/P_{it-1}$, are positive on average under both amortization methods, with a mean of 0.02 for both the Almon lag and ANN methods. This indicates that, on average, earnings under R&D capitalization exceed reported earnings, consistent with current R&D expense exceeding estimated R&D amortization. The impact of these adjustments on reported earnings is substantial. For observations with positive reported earnings, reported earnings are, on average, understated by 23.6% under the Almon lag method and 26.2% under the ANN method, where understatement is measured as the earnings adjustment from R&D capitalization scaled by reported earnings. The understatement varies across industries. Under the Almon lag method, the earnings understatement ranges from 11.0% in Scientific Instruments to 62.0% in Transportation Vehicles. Under the ANN method, the earnings understatement ranges from 21.6% in Other Industries to 36.7% in Business Services. These results indicate that the capitalization adjustment has a meaningful effect on reported earnings and that the ANN-based amortization schedule generally implies a larger earnings adjustment than the benchmark Almon lag schedule.[11]

For the price-test variables, Table 3, Panel A reports per-share earnings, misstatement, and book-value adjustments. The mean R&D capital per share, $(BV^{C}_{it} - BV^{E}_{it})$, is 3.79 under the Almon lag method and 5.56 under the ANN method; the corresponding medians are 1.66 and 3.81, respectively. These differences suggest that the ANN method generally assigns a larger unamortized R&D capital balance than the benchmark Almon lag method for test-sample observations.

[11] When we retain all observations and scale the total earnings adjustment by the absolute value of total reported earnings, the average corresponding magnitudes are 61.6% under the Almon lag method and 68.8% under the ANN method.

Using the two measures of *Misstatement* from the two different models, we include both reported R&D expense and the two *Misstatement* measures as explanatory variables in our regression models, and we run a "horse race" to see which *Misstatement* measure better explains the dependent variables. Results are shown in Table 3, Panels B and C. For all regressions, *Misstatement* based on the ANN model dominates *Misstatement* based on the Amon lag model in terms of both statistical and economic significance.

To provide more supporting evidence on the efficacy of the ANN method, we perform the Vuong (year) test to compare the incremental R^2's on the two *Misstatement* variables. Results are shown in Table 3 Panels B and C. For both the return and prices regressions, the ANN measures provide significantly more explanatory power than the Almon measures.

We conduct intertemporal analysis to compare the explanatory power of R&D capital estimated under the Almon Lag vs. ANN models with respect to stock returns in the next 12 months. Table 4, Panel A reports summary statistics of regression variables. Panel B shows that R&D capital derived from the ANN model dominates that from the Almon Lag model in explaining future stock returns.

Finally, we check the robustness of the results using longer windows to estimate Almon Lag models. In the primary analyses, we use past three years to estimate both the ANN and Almon lag models to ensure a fair comparison. Nevertheless, Lev and Sougiannis (1996) estimate the Almon lag amortization rates using coefficient estimates averaged over 1975 through year *t*-1. As shown in Table 5, strictly following their estimation window to estimate amortization rates does not affect our inferences.

6. Conclusion

Capitalizing R&D outlays requires estimates of amortization rates over their useful lives. Prior research derives such rates from a lag structure estimated within an industry-year group that links past R&D expenditures to future benefits, presuming firms in that group share the same R&D payoff timeline. We use R&D disclosures in firms' 10-K filings to contextualize the lag structure estimation and derive a firm-year-specific amortization schedule. Specifically, for each firm-year, we convert R&D disclosures through BERT (Bidirectional Encoder Representations from Transformers) into a high-dimensional context vector that represents their semantic meanings. The relationships between past R&D expenditures and future earnings is modeled as functions of the context vector using artificial neural networks. The resulting R&D capital estimates outperform those from prior research in explaining R&D firms' contemporaneous stock prices and returns, as well as their future stock returns over the next 12 months. Our paper is the first to introduce qualitative R&D information to estimate firms' R&D capital unrecognized in their financial statements. The findings suggest that qualitative R&D information is useful for improving the estimation of firms' R&D capital that is unrecognized in their financial statements.

References:

Ball, R., & Nikolaev, V. V. (2022). On earnings and cash flows as predictors of future cash flows. *Journal of Accounting and Economics*, 73(1), 101430.

Bertomeu, J., Cheynel, E., Liao, Y., & Milone, M. (2025). Using machine learning to measure conservatism. *Management Science*, 71(2), 1504-1522.

Blankespoor, E., deHaan, E., & Marinovic, I. (2020). Disclosure processing costs, investors' information choice, and equity market outcomes: A review. *Journal of Accounting and Economics*, 70(2-3), 101344.

Brown, S. V., & Tucker, J. W. (2011). Large-sample evidence on firms' year-over-year MD&A modifications. *Journal of Accounting Research*, 49(2), 309-346.

Chambers, D. J., Jennings, R., & Thompson, R. B. (1998). Evidence on the usefulness of capitalizing and amortizing research and development costs. *Available at SSRN 58661*

Chan, L. K., Lakonishok, J., & Sougiannis, T. (2001). The stock market valuation of research and development expenditures. *Journal of Finance*, 56(6), 2431-2456.

Chen, L., Pelger, M., & Zhu, J. (2024). Deep learning in asset pricing. *Management Science*, 70(2), 714-750.

Chen, X., Cho, Y. H., Dou, Y., & Lev, B. (2022). Predicting future earnings changes using machine learning and detailed financial data. *Journal of Accounting Research*, 60(2), 467-515.

Dechow, P. M. (1994). Accounting earnings and cash flows as measures of firm performance: The role of accounting accruals. *Journal of Accounting and Economics*, 18(1), 3-42.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. *In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies*, volume 1 (long and short papers) (pp. 4171-4186).

Dyer, T., Lang, M., & Stice-Lawrence, L. (2017). The evolution of 10-K textual disclosure: Evidence from Latent Dirichlet Allocation. Journal of Accounting and Economics, 64(2-3), 221-245.

Falato, A., Kadyrzhanova, D., Sim, J., & Steri, R. (2022). Rising intangible capital, shrinking debt capacity, and the US corporate savings glut. *Journal of Finance*, 77(5), 2799-2852.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383-417.

Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2), 806-841.

Iqbal, A., Rajgopal, S., Srivastava, A., & Zhao, R. (2025). A better estimate of internally generated intangible capital. *Management Science*, 71(1), 731-752.

Jennings, R., Simko, P. J., & Thompson, R. B. (1996). Does LIFO inventory accounting improve the income statement at the expense of the balance sheet?. *Journal of Accounting Research*, 34(1), 85-109.

Kim, A. G., & Nikolaev, V. V. (2024). Context-Based Interpretation of Financial Information. *Journal of Accounting Research.*

Kothari, S. P., & Zimmerman, J. L. (1995). Price and return models. *Journal of Accounting and Economics*, 20(2), 155-192.

KPMG. (2011). The challenges of the evolving nature of disclosures. In the Headlines, Issue 2011/06. KPMG IFRG Limited.

Lev, B., & Sougiannis, T. (1996). The capitalization, amortization, and value-relevance of R&D. *Journal of Accounting and Economics*, 21(1), 107-138.

Li, F. (2008). Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics*, 45(2-3), 221-247.

Li, F., Lundholm, R., & Minnis, M. (2013). A measure of competition based on 10-K filings. *Journal of Accounting Research*, 51(2), 399-436.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35-65.

Merkley, K. J. (2014). Narrative disclosure and earnings performance: Evidence from R&D disclosures. *The Accounting Review*, 89(2), 725-757.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26.

Ryu, H. S., & Zarowin, P. A. (2025). Can the Stock Market Capitalize R&D Expenditures?*(* When Firms Aren't Mandated to). *Available at SSRN 4888485*.

Siano, F., & Wysocki, P. (2021). Transfer learning and textual analysis of accounting disclosures: Applying big data methods to small (er) datasets. *Accounting Horizons*, 35(3), 217-244.

Sougiannis, T. (1994). The accounting-based valuation of corporate R&D. *The Accounting Review*, 44-68.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929-1958.

Securities and Exchange Commission. (1998). A plain English handbook: How to create clear SEC disclosure documents. Office of Investor Education and Assistance, U.S. Securities and Exchange Commission.

Securities and Exchange Commission. (2013). Report on review of disclosure requirements in Regulation S-K. U.S. Securities and Exchange Commission.

Wang, S., & Levy, B. (2024). BeanCounter: A low-toxicity, large-scale, and open dataset of business-oriented text. *Advances in Neural Information Processing Systems*, 37, 91653-91690.

Appendix A: Variable Definitions

Variable	Definition
Dependent Variables	
$OI/Sale_t$	Annual operating income, before depreciation, advertising and R&D expenses, divided by sales.
R_{it}	Annual stock return from nine months before fiscal year-end through three months after fiscal year-end.
P_{it}	Share price measured three months after fiscal year-end.
R_{it+1}	Future 12-month buy-and-hold stock return from four months after fiscal year-end through fifteen months after fiscal year-end.
Independent Variables	
$(X^C_{it} - X^E_{it})$ *(Almon Lag)*	The difference between earnings under R&D capitalization and reported earnings, equal to current R&D expense minus R&D amortization, where the amortization schedule is estimated using the Almon lag method.
$(X^C_{it} - X^E_{it})/P_{it-1}$ *(Almon Lag)*	The difference between earnings under R&D capitalization and reported earnings, constructed using the Almon lag method, divided by P_{it-1}.
$(X^C_{it} - X^E_{it})$ *(ANN)*	The difference between earnings under R&D capitalization and reported earnings, equal to current R&D expense minus R&D amortization, where the amortization schedule is estimated using the ANN method.
$(X^C_{it} - X^E_{it})/P_{it-1}$ *(ANN)*	The difference between earnings under R&D capitalization and reported earnings, constructed using the ANN method, divided by P_{it-1}.
$\Delta(X^C_{it} - X^E_{it})$ (Almon Lag)	Annual change in the earnings misstatement, $(X^C_{it} - X^E_{it})$, constructed using the Almon lag method.
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)	Annual change in the earnings misstatement constructed using the Almon lag method, divided by P_{it-1}.
$\Delta(X^C_{it} - X^E_{it})$ (ANN)	Annual change in the earnings misstatement, $(X^C_{it} - X^E_{it})$, constructed using the ANN method.
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)	Annual change in the earnings misstatement constructed using the ANN method, divided by P_{it-1}.
$ln((RDC/M)_{it})$ (Almon Lag)	Natural logarithm of the ratio of estimated R&D capital under the Almon lag method to market value of equity at fiscal year-end.
$ln((RDC/M)_{it})$ (ANN)	Natural logarithm of the ratio of estimated R&D capital under the ANN method to market value of equity at fiscal year-end.
P_{it-1}	Share price measured three months after fiscal year-end at year t−1.
X^E_{it}	Reported earnings per share before extraordinary items.
X^E_{it}/P_{it-1}	Reported earnings per share before extraordinary items, divided by P_{it-1}.

ΔX^{E}_{it}	The annual change in reported earnings per share before extraordinary items.
$\Delta X^{E}_{it}/P_{it-1}$	Annual change in reported earnings per share before extraordinary items, divided by P_{it-1}.
$(BV^{C}_{it} - BV^{E}_{it})$ (Almon Lag)	The per-share book-value adjustment due to R&D capitalization, measured as estimated R&D capital constructed using the Almon lag method, divided by common shares outstanding.
$(BV^{C}_{it} - BV^{E}_{it})$ (ANN)	The per-share book-value adjustment due to R&D capitalization, measured as estimated R&D capital constructed using the ANN method, divided by common shares outstanding.
$TA/Sale_{t-1}$	The sum of plant and equipment, inventory, and investment in unconsolidated subsidiaries and goodwill, measured in current dollars at beginning-of-year values, divided by sales.
$AD/Sale_{t-1}$	Annual advertising expense, measured at beginning-of-year values, divided by sales.
$R\&D/Sale_{t}$	Research and development expense divided by sales.
β_{it}	CAPM beta of the firm, estimated using the prior 60 monthly stock returns, with a minimum of 24 monthly observations required.
$ln(M_{it})$	Natural logarithm of the market value of equity of the firm, calculated as share price times shares outstanding at fiscal year-end.
$ln(B/M_{it})$	Natural logarithm of the book-to-market ratio, defined as common equity plus deferred taxes divided by the market value of equity at fiscal year-end.
$ln(A/B_{it})$	Natural logarithm of the ratio of book value of total assets to book value of common equity at fiscal year-end.
$(E(+)/M)_{it}$	The ratio of earnings before extraordinary items, plus income-statement deferred taxes, minus preferred dividends, to the market value of equity at fiscal year-end; set to zero when the numerator is negative.
$(E/M\ dummy)_{it}$	An indicator equal to 1 if firm i's fiscal-year earnings are negative, and 0 otherwise.

Notes: All continuous variables are winsorized at 1% and 99% to reduce the influence of outliers.

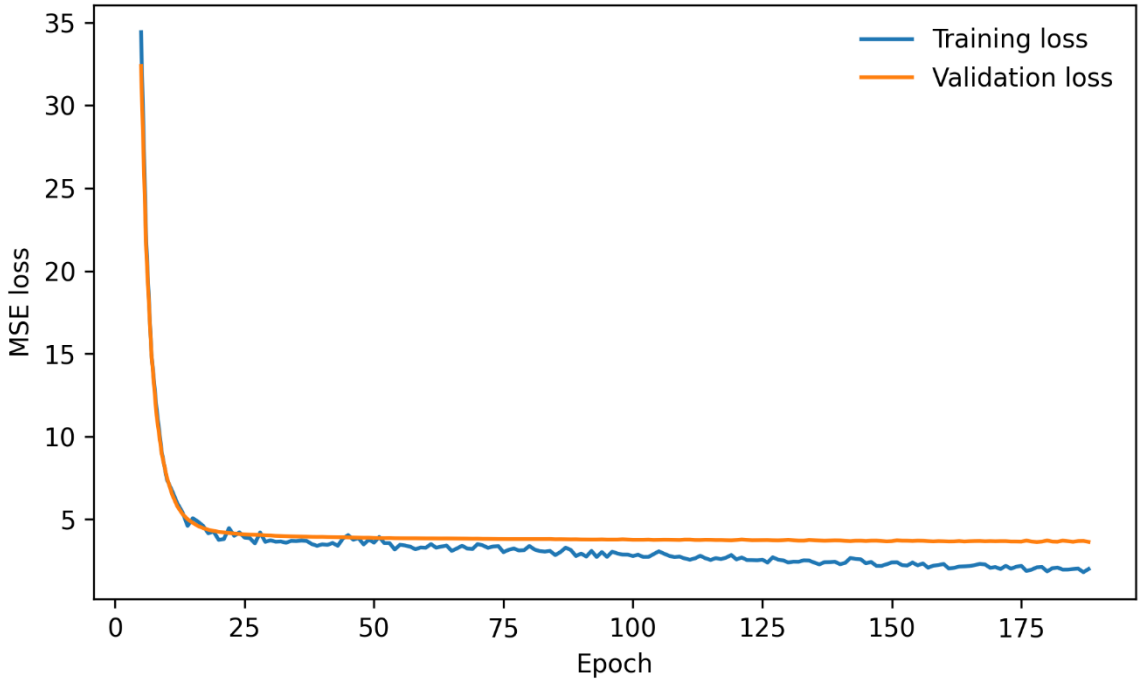
Appendix B: Examples of Within-Industry Variation in R&D Disclosure Context

SIC2	Firm	Disclosure-Implied R&D Horizon	R&D Context
28 (Chemicals and Allied Products)	Sensient Technologies Corp	Long	*[...] The Company's research, development and quality assurance personnel make significant contributions toward improving existing products and **developing new products tailored to customer needs**, while providing on-going technical support and know-how to the Company's manufacturing activities. [...]*
	Eastman Chemical Co.	Short	*[...] The Company's Fibers segment R&D efforts focus on **process and product improvements, as well as cost reduction**, with the objectives of increasing sales and reducing costs. The Fibers segment also conducts research to **assist acetate tow customers in the effective use** of the segment's products and in the customers' product development efforts. [...]*
35 (Industrial and Commercial Machinery and Computer Equipment)	SPX Technologies Corp	Long	*[...] This new technology is **intended over the long term**, to provide performance and efficiencies in transmitting electrical power. The acquisition of Copes-Vulcan in September 2000 provided new technology, complementary products and services, and an expanded international presence. [...]*
	Crossroads Impact Corp	Short	*[...] Our product development strategy includes the creation of code modules that can be used across different products, **enabling fast development** and high quality and reliability. [...]*
36 (Electronic & Other Electric Equipment)	IPG Photonics Corp	Long	*[...] We may not be able to respond by decreasing our expenses quickly enough or sufficiently, due in part, to our fixed overhead structure related to our vertically integrated operations and our commitments to **continuing investment in research and development and infrastructure for long term growth. [...]***
	Intel Corp	Short	*[...] We are a**ccelerating the process technology development** for our Intel Atom product line to deliver increased battery life, performance, and feature integration. [...]*
37 (Transportation Vehicles)	GE Aerospace	Long	*[...] In many instances, the competitive climate is characterized by changing technology that **requires continuing research and development**. We continue to invest in advanced technology development that will provide more value to our customers and more efficient solutions that comply with today, strict environmental regulations. **Product development cycles are long and product quality and efficiency are critical to success**. [...]*
	Wabtec Corp	Short	*[...] **Accelerate new product development.** We continue to emphasize research and development funding to create new and improved products. 4 million, respectively, on product development and improvement activities. Sometimes we conduct specific research projects in conjunction with universities, customers and other railroad product suppliers. [...]*

38 (Scientific Instruments)	Mentor Corp	Long	*[...] The production and marketing of our products and our ongoing research and development activities, including pre-clinical testing and clinical trial activities,* ***are subject to extensive regulation and review by numerous governmental authorities,*** *there has been a continuing trend toward more stringent FDA requirements in the areas of product approval and enforcement, causing medical device and biologics manufacturers to* ***experience longer research and development timelines****, longer approval cycles, greater risk and uncertainty, and higher expenses. [...]*
	Keithley Instruments Inc.	Short	*[...] New products are an important factor in our sales growth strategy, and we have increased our investment in product development activity spending levels to expand our product offering and accelerate the introduction of new products.* ***Our engineering development activities are directed toward the development of new products that will complement****, replace or* ***add to the products currently included in our product line****. [...]*
73 (Business Services)	Telos Corp/MD	Long	*[...] In contrast to traditional cybersecurity businesses with a focus on government customers, we own the* ***intellectual property*** *developed through our R&D initiatives and can* ***deploy our technology solutions across our entire public and private sector customer set****. [...]*
	PDF Solutions Inc.	Short	*[...] For example, our foundry customers generate and analyze key manufacturing data using our solutions to* ***shorten the time necessary for technology development*** *and to provide their fabless customers with a higher yielding process with improved electrical performance, which are both critical metrics for market success. [...]*
99 (Others)	Autoscope Technologies Corp	Long	*[...]* ***In addition to market penetration, our long-term success*** *depends on our ability to expand our business through new product development, mergers and acquisitions, and/or geographic expansion. [...]*
	Azenta Inc.	Short	*[...] As discussed above, we have made acquisitions over the years which* ***accelerated our product development cycle****, broadened our installed base and added customer relationships to our business. [...]*

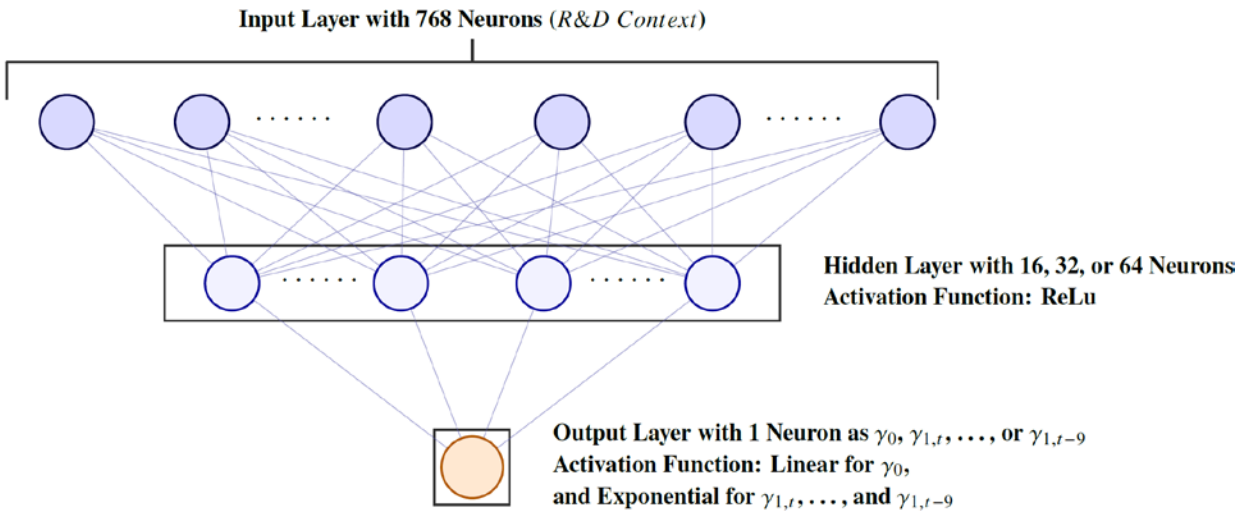
This appendix provides examples of R&D disclosures from firms in the same two-digit SIC industry. The examples illustrate within-industry variation in disclosure-implied R&D horizons. For each industry, we present one firm whose disclosure suggests a longer R&D payoff horizon and one firm whose disclosure suggests a shorter R&D payoff horizon. The bolded text highlights phrases that indicate differences in R&D timelines, regulatory approval cycles, product development processes, or new product introductions. These examples illustrate why firm-year-specific R&D disclosure context may help refine R&D capitalization estimates beyond an industry-level amortization schedule.

Appendix C: ANN Training and Validation Loss for the 2021 Test-Year Model



This figure plots the training and validation mean squared error (MSE) loss by epoch for the ANN model used to generate coefficients for the 2021 test year. The model is trained using observations from 2018 and 2019 and validated using observations from 2020.

Figure 1: Artificial Neural Network Architecture for Contextual Coefficients



This figure illustrates how R&D context maps into context coefficients (γ_0 and $\gamma_{1,t} \dots \gamma_{1,t-9}$) through an artificial neural network architecture. For each firm-year, R&D disclosures in 10-Ks are transformed into a 768-dimensional embedding vector (*R&D Context*) using BERT and then passed through 11 artificial neural networks to generate 11 latent contextual coefficients. One coefficient captures the main effect of *R&D context* (γ_0) and 10 ($\gamma_{1,t} \dots \gamma_{1,t-9}$) interact with the corresponding lagged R&D expenditures. The 11 elements combined with advertising expenses and tangible assets are used to explain operating income scaled by sales, $OI/S_{i,t}$ (see Equation (3)).

Table 1: Sample Selection

Sample Selection	Obs.
Firm-years from COMPUSTAT from 1996-2021	244,461
Drop firm-years with negative sales	221,638
Drop firm-years with a missing 4-digit SIC code	220,842
Drop firm-years with large mergers (those contributing 50% or more to annual sales)	216,289
Firm-years with non-missing R&D in the past 10 years and R&D/Sale>=2%	23,622
Firm-years with non-missing R&D disclosures in 10-K filings	15,988
Firm-years with non-missing EPS and other control variables	14,535
Training and Validation Sample (1996-2020)	13,973
Test Sample (1999-2021)	13,306

This table summarizes the sample selection procedure. Starting with firm-year observations in COMPUSTAT from 1996 to 2021, we sequentially exclude observations with negative sales, missing 4-digit SIC codes, and large mergers that account for 50% or more of annual sales. We then require non-missing R&D over the past 10 years, R&D-to-sales of at least 2%, non-missing R&D disclosures in 10-K filings, and non-missing earnings and other control variables. Starting in 1999, for each test year, we use the past three years to estimate the R&D-earnings relation and derive R&D amortization rates. In total, the training and validation sample comprises 13,973 firm-years from 1996-2020, and the test sample comprises 13,306 firm-years from 1999-2021.

Table 2: Estimation of the R&D-Earnings Relation

Panel A: Summary statistics

	Obs.	Mean	SD	P25	Median	P75
$OI/Sale_t$	13,973	-0.04	1.61	0.12	0.21	0.32
$TA/Sale_{t-1}$	13,973	0.68	1.07	0.29	0.46	0.70
$AD/Sale_{t-1}$	13,973	0.01	0.02	0.00	0.00	0.01
$R\&D/Sale_t$	13,973	0.57	2.51	0.05	0.10	0.19
$R\&D/Sale_{t-1}$	13,973	0.58	2.51	0.05	0.10	0.19
$R\&D/Sale_{t-2}$	13,973	0.56	2.28	0.05	0.10	0.19
$R\&D/Sale_{t-3}$	13,973	0.58	2.41	0.05	0.10	0.19
$R\&D/Sale_{t-4}$	13,973	0.60	2.49	0.04	0.10	0.19
$R\&D/Sale_{t-5}$	13,973	0.58	2.33	0.04	0.10	0.19
$R\&D/Sale_{t-6}$	13,973	0.59	2.33	0.04	0.10	0.20
$R\&D/Sale_{t-7}$	13,973	0.60	2.34	0.04	0.09	0.20
$R\&D/Sale_{t-8}$	13,973	0.66	2.67	0.04	0.09	0.20
$R\&D/Sale_{t-9}$	13,973	0.81	3.42	0.04	0.09	0.20

Panel B: Model performance based on root mean square error (RMSE)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Almon Lag Training	Almon Lag Test	ANN Training	ANN Validation	ANN Test	Diff	Diff
Year	RMSE	RMSE	RMSE	RMSE	RMSE	(1) - (3)	(2) - (5)
1999	0.6230	0.7929	0.7847	1.0477	0.8212	-0.1642	-0.0283
2000	0.6770	0.9326	0.6769	0.7038	0.8069	0.0001	0.1256
2001	0.6880	1.5196	0.9632	0.7106	1.3208	-0.2752	0.1988
2002	0.9131	0.9684	0.5660	1.2866	1.1787	0.3471**	-0.2103
2003	0.8790	1.0880	0.8037	1.0850	1.4466	0.0754	-0.3586**
2004	0.9122	1.4934	0.9094	1.3541	1.1626	0.0028	0.3308**
2005	0.8420	3.6335	0.5267	0.9808	1.6502	0.3152***	1.9833***
2006	1.1858	3.9881	0.8140	1.4414	1.2797	0.3718***	2.7084***
2007	1.1879	5.8215	0.7594	1.0629	1.3406	0.4285***	4.4809***
2008	1.2937	7.8954	0.8052	1.2435	1.6937	0.4885***	6.2017***
2009	1.3447	6.3564	0.8426	1.6525	1.5988	0.5022***	4.7576***
2010	1.4972	5.4511	1.1199	1.4758	1.5805	0.3773***	3.8706***
2011	1.6247	4.4195	1.0643	1.4707	1.6476	0.5604***	2.7720***
2012	1.6846	3.4567	1.0768	1.5690	1.7633	0.6078***	1.6934***
2013	1.8137	2.2942	1.1995	1.6504	1.6928	0.6142***	0.6014***
2014	1.8784	5.4483	1.0529	1.4108	1.8380	0.8255***	3.6103***
2015	1.8967	5.0882	1.0913	1.7169	1.8013	0.8054***	3.2869***
2016	1.9192	5.9388	1.0184	1.6633	1.9547	0.9008***	3.9840***
2017	1.9738	3.5037	1.0650	1.7901	1.9216	0.9088***	1.5820***
2018	2.0107	2.7859	1.1398	1.6926	1.7294	0.8709***	1.0565***
2019	1.9901	2.0574	1.1675	1.5796	1.8861	0.8226***	0.1713
2020	1.9105	2.1203	1.0903	1.5441	2.0257	0.8303***	0.0945
2021	1.8454	2.6867	1.4320	1.8911	2.1617	0.4134***	0.5250***
Average	1.4170	3.4670	0.9552	1.3923	1.5784	0.4618***	1.8886***

Panel A reports summary statistics for the variables (excluding the 768 contextual features) used to estimate the R&D-earnings relation. *OI/Sale* is operating income divided by sales. *AD/Sale* is advertising expenses divided by sales. *TA/Sale* is tangible assets divided by sales. *R&D/Sale* is R&D expenditures divided by sales. Panel B reports model performance, measured as root mean square error (RMSE), by test year from 1999-2021 and the average across test years for the Almon lag model (Equation (1)) and the ANN model (Equation (3)) used to explain *OI/Sale*. For each test year, we use the past three years to estimate an Almon Lag model and apply it to the test year. Column (1) reports the in-sample RMSE in the past three years, and column (2) reports the out-of-sample RMSE in the test year. For each test year, we also use the second and third preceding years to estimate 27 ANN models with different hyperparameter combinations and the preceding year to select the model with the lowest out-of-sample RMSE. Column (3) reports the in-sample RMSE of the chosen model in the second and third preceding years, and Column (4) reports the out-of-sample RMSE of the chosen model in the preceding year. Column (5) reports the out-of-sample RMSE of the chosen model in the test year. Columns (6) and (7) report the differences in training and test RMSE between the Almon lag and ANN models, respectively. Statistical significance for RMSE differences is assessed using bootstrap t-statistics with 500 replications. In each replication, we randomly resample firms with replacement within each test year and recompute the RMSE differences between the two models. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively, using two-tailed tests.

Table 3: Contemporaneous Analyses

Panel A: Summary statistics

	Obs.	Mean	SD	P25	Median	P75
R_{it}	13,306	0.17	0.68	-0.24	0.05	0.38
P_{it}	13,306	30.89	46.45	4.65	14.49	37.75
X^E_{it}	13,306	0.71	2.27	-0.34	0.29	1.44
X^E_{it}/P_{it-1}	13,306	-0.06	0.26	-0.07	0.02	0.05
ΔX^E_{it}	11,110	0.09	1.54	-0.33	0.07	0.49
$\Delta X^E_{it}/P_{it-1}$	11,110	0.01	0.22	-0.02	0.00	0.03
$(X^C_{it} - X^E_{it})$ (Almon Lag)	13,306	0.38	0.68	0.00	0.11	0.46
$(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)	13,306	0.02	0.05	0.00	0.01	0.03
$(X^C_{it} - X^E_{it})$ (ANN)	13,306	0.43	0.62	0.05	0.24	0.62
$(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)	13,306	0.02	0.05	0.00	0.01	0.04
$\Delta(X^C_{it} - X^E_{it})$ (Almon Lag)	11,110	0.01	0.40	-0.09	0.01	0.11
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)	11,110	-0.00	0.04	-0.01	0.00	0.01
$\Delta(X^C_{it} - X^E_{it})$ (ANN)	11,110	0.01	0.37	-0.12	0.01	0.14
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)	11,110	-0.00	0.04	-0.01	0.00	0.01
$(BV^C_{it} - BV^E_{it})$ (Almon Lag)	13,306	3.79	5.14	0.60	1.66	4.63
$(BV^C_{it} - BV^E_{it})$ (ANN)	13,306	5.56	5.47	1.83	3.81	7.18

Panel B: Return tests

	R_{it}					
	(1)	(2)	(3)	(4)	(5)	(6)
X^E_{it}/P_{it-1}	0.157***	0.156***	0.156***	-0.042	-0.045	-0.043
	(4.73)	(4.69)	(4.69)	(-1.02)	(-1.06)	(-1.00)
$\Delta X^E_{it}/P_{it-1}$				0.614***	0.620***	0.618***
				(12.56)	(12.46)	(12.45)
$(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)	0.387**		-0.013	0.528***		0.045
	(2.09)		(-0.06)	(2.59)		(0.20)
$(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)		0.811***	0.816***		0.962***	0.934***
		(4.71)	(4.23)		(5.12)	(4.40)
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)				-0.378*		-0.181
				(-1.77)		(-0.77)
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)					-0.288	-0.205
					(-1.25)	(-0.81)
Constant	0.169***	0.160***	0.160***	0.143***	0.133***	0.133***
	(31.02)	(28.72)	(27.31)	(25.17)	(22.94)	(21.52)
Observations	13,306	13,306	13,306	11,110	11,110	11,110
SIC2_FE	Yes	Yes	Yes	Yes	Yes	Yes
Year_FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.229	0.231	0.231	0.272	0.274	0.274
Cluster	Firm	Firm	Firm	Firm	Firm	Firm
Mean VIF	--	--	1.12	--	--	1.34
Vuong Test of Adjusted R^2	(2) v.s. (1)			(5) v.s. (4)		
Difference in Adjusted R^2	0.0020			0.0025		
Vuong Z-statistics	2.084			2.007		
p_value	0.0372			0.0448		

Panel C: Price tests

	P_{it}					
	(1)	(2)	(3)	(4)	(5)	(6)
X^E_{it}	15.921***	14.526***	14.440***	15.918***	13.837***	13.893***
	(7.96)	(7.70)	(7.76)	(7.92)	(7.40)	(7.35)
$(X^C_{it} - X^E_{it})$ (Almon Lag)	25.563***		8.748*	23.462***		4.124
	(4.36)		(1.93)	(3.81)		(1.03)
$(X^C_{it} - X^E_{it})$ (ANN)		34.305***	28.875***		28.721***	26.876***
		(7.11)	(8.71)		(5.93)	(6.16)
$(BV^C_{it} - BV^E_{it})$ (Almon Lag)				0.016		0.380
				(0.04)		(0.63)
$(BV^C_{it} - BV^E_{it})$ (ANN)					1.275***	0.860
					(4.62)	(1.59)
Constant	13.218***	8.495***	7.519***	13.197***	4.249*	4.232*
	(6.58)	(4.33)	(3.24)	(6.60)	(1.94)	(1.94)
Observations	13,306	13,306	13,306	13,306	13,306	13,306
SIC2_FE	Yes	Yes	Yes	Yes	Yes	Yes
Year_FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.303	0.326	0.329	0.303	0.330	0.331
Cluster	Firm	Firm	Firm	Firm	Firm	Firm
Mean VIF	--	--	1.35	--	--	4.9
Vuong Test of Adjusted R^2	(2) v.s. (1)			(5) v.s. (4)		
Difference in Adjusted R^2	0.0233			0.0273		
Vuong Z-statistics	4.012			4.02		
p_value	0.0001			0.0001		

This table reports summary statistics and regression analyses for the return and price tests within the 1999–2021 test sample. Panel A presents descriptive statistics. Panels B and C report return and price tests, respectively, using misstatement measures based on the Almon lag and ANN methods. The table reports both separate specifications, with Vuong tests comparing model fit, and horse-race specifications that include the Almon lag and ANN measures jointly. The dependent variables are annual stock return, measured from nine months before the fiscal year-end through three months after the fiscal year-end, and share price, measured three months after the fiscal year-end. The main independent variable is Misstatement, $(X^C_{it} - X^E_{it})$. The unit of observation is the firm-year. All regressions include SIC2 fixed effects and year fixed effects. t-statistics based on standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively, using two-tailed tests. Detailed definitions of all variables are provided in Appendix A.

Table 4: Intertemporal Analyses

Panel A: Summary statistics

	Obs.	Mean	SD	P25	Median	P75
R_{it+1}	11,624	0.15	0.65	-0.24	0.05	0.36
β_{it}	11,624	1.35	0.74	0.86	1.24	1.74
$ln(M_{it})$	11,624	6.42	2.31	4.70	6.31	7.96
$ln(B/M_{it})$	11,624	-1.01	0.86	-1.51	-0.95	-0.41
$ln(A/B_{it})$	11,624	0.59	0.52	0.22	0.45	0.79
$(E(+)/M)_{it}$	11,624	0.03	0.04	0.00	0.02	0.05
$(E/M\ dummy)_{it}$	11,624	0.40	0.49	0.00	0.00	1.00
$ln((RDC/M)_{it})$ (Almon Lag)	11,624	-2.20	1.43	-3.21	-2.24	-1.25
$ln((RDC/M)_{it})$ (ANN)	11,624	-1.50	1.16	-2.30	-1.54	-0.74

Panel B: Future return rests

	R_{it+1}			
	(1)	(2)	(3)	(4)
β_{it}	-0.006	-0.009	-0.014	-0.013
	(-0.66)	(-1.09)	(-1.62)	(-1.55)
$ln(M_{it})$	-0.007**	0.000	0.005	0.005
	(-2.30)	(0.06)	(1.60)	(1.56)
$ln(B/M_{it})$	0.074***	0.061***	0.057***	0.056***
	(3.44)	(2.86)	(2.74)	(2.68)
$ln(A/B_{it})$	0.004	0.004	0.004	0.004
	(1.64)	(1.59)	(1.54)	(1.55)
$(E(+)/M)_{it}$	-0.031	-0.145	-0.213	-0.213
	(-0.15)	(-0.73)	(-1.06)	(-1.07)
$(E/M\ dummy)_{it}$	-0.001	-0.026*	-0.038**	-0.038**
	(-0.06)	(-1.65)	(-2.33)	(-2.37)
$ln((RDC/M)_{it})$ (Almon Lag)		0.035***		0.013*
		(5.05)		(1.88)
$ln((RDC/M)_{it})$ (ANN)			0.053***	0.041***
			(7.61)	(4.24)
Constant	0.160***	0.218***	0.205***	0.216***
	(5.19)	(6.67)	(6.38)	(6.64)
Observations	11,624	11,624	11,624	11,624
SIC2_FE	Yes	Yes	Yes	Yes
Year_FE	Yes	Yes	Yes	Yes
Adjusted R-squared	0.223	0.226	0.229	0.229
Cluster	Firm	Firm	Firm	Firm
Mean VIF	--	--	--	1.65
Vuong Test of Adjusted R^2	(3) v.s. (1)		(3) v.s. (2)	
Difference in Adjusted R^2	0.0066		0.0030	
Vuong Z-statistics	2.89		1.997	
p_value	0.0039		0.0458	

This table reports descriptive statistics and regressions of future 12-month stock returns on R&D capital constructed using alternative R&D amortization schedules. The dependent variable is (R_{it+1}), the 12-month stock return measured from months 4 through 15 after the fiscal year-end. Panel A presents descriptive statistics for the variables used in the intertemporal analysis. Panel B reports specifications using no R&D capital measure, R&D capital based on the Almon lag method, R&D capital based on the ANN method, and both R&D capital measures included jointly in the same regression. The main independent variables are $ln((RDC/M)_{it})$ (Almon Lag) based on the Almon lag method and

$ln((RDC/M)_{it})$ (ANN) based on the ANN method. Vuong tests reported at the bottom of Panel B compare the relative fit of the ANN specification with the benchmark specifications. The unit of observation is the firm-year. All regressions include SIC2 fixed effects and year fixed effects. t-statistics based on standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively, using two-tailed tests. Detailed definitions of all variables are provided in Appendix A.

Table 5: Robustness Analyses Using Longer Windows to Estimate Almon Lag Models

Panel A: Return tests

	R_{it}					
	(1)	(2)	(3)	(4)	(5)	(6)
X^E_{it}/P_{it-1}	0.151***	0.156***	0.160***	-0.055	-0.045	-0.049
	(4.55)	(4.69)	(4.76)	(-1.29)	(-1.06)	(-1.15)
$\Delta X^E_{it}/P_{it-1}$				0.632***	0.620***	0.625***
				(12.82)	(12.46)	(12.58)
$(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)	0.621*		-0.697*	1.530***		0.265
	(1.85)		(-1.77)	(4.16)		(0.66)
$(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)		0.811***	1.002***		0.962***	0.754***
		(4.71)	(4.20)		(5.12)	(3.10)
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (Almon Lag)				-0.545		-0.127
				(-1.59)		(-0.27)
$\Delta(X^C_{it} - X^E_{it})/P_{it-1}$ (ANN)					-0.288	-0.274
					(-1.25)	(-0.86)
Constant	0.169***	0.160***	0.164***	0.135***	0.133***	0.134***
	(29.37)	(28.72)	(29.92)	(22.80)	(22.94)	(23.35)
Observations	13,306	13,306	13,306	11,110	11,110	11,110
SIC2_FE	Yes	Yes	Yes	Yes	Yes	Yes
Year_FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.229	0.231	0.231	0.273	0.274	0.274
Cluster	Firm	Firm	Firm	Firm	Firm	Firm
Mean VIF	--	--	1.52	--	--	1.88
Vuong Test of Adjusted R^2	(2) v.s. (1)			(5) v.s. (4)		
Difference in Adjusted R^2	0.0021			0.0015		
Vuong Z-statistics	2.452			1.627		
p_value	0.0142			0.0948		

Panel B: Price tests

	P_{it}					
	(1)	(2)	(3)	(4)	(5)	(6)
X^E_{it}	15.240***	14.526***	14.496***	15.048***	13.837***	13.815***
	(7.83)	(7.70)	(7.73)	(7.66)	(7.40)	(7.43)
$(X^C_{it} - X^E_{it})$ (Almon Lag)	45.969***		7.281	40.868***		5.066
	(5.51)		(1.03)	(4.70)		(0.65)
$(X^C_{it} - X^E_{it})$ (ANN)		34.305***	30.603***		28.721***	26.203***
		(7.11)	(7.70)		(5.93)	(6.07)
$(BV^C_{it} - BV^E_{it})$ (Almon Lag)				0.917**		-0.313
				(2.11)		(-0.33)
$(BV^C_{it} - BV^E_{it})$ (ANN)					1.275***	1.433**
					(4.62)	(2.06)
Constant	11.064***	8.495***	8.255***	9.674***	4.249*	4.141*
	(5.52)	(4.33)	(3.94)	(4.64)	(1.94)	(1.84)
Observations	13,306	13,306	13,306	13,306	13,306	13,306

SIC2_FE	Yes	Yes	Yes	Yes	Yes	Yes
Year_FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.311	0.326	0.326	0.312	0.330	0.33
Cluster	Firm	Firm	Firm	Firm	Firm	Firm
Mean VIF	--	--	1.81	--	--	4.76
Vuong Test of Adjusted R^2	(2) v.s. (1)			(5) v.s. (4)		
Difference in Adjusted R^2	0.0151			0.0181		
Vuong Z-statistics	4.374			4.389		
p_value	0.0000			0.0000		

This table reports robustness tests of return and price regressions in the 1999–2021 test sample. Panels A and B report return and price tests, respectively, using misstatement measures based on the Almon lag and ANN methods. The table reports both separate specifications, with Vuong tests comparing model fit, and horse-race specifications including both measures jointly. The dependent variables are annual stock return, measured from nine months before fiscal year-end through three months after fiscal year-end, and share price, measured three months after fiscal year-end. The main independent variable is ($X^C_{it} - X^E_{it}$). In this robustness test, the Almon lag amortization rates are calculated using coefficient estimates averaged over 1975 through year t-1. The unit of observation is the firm-year. All regressions include SIC2 and year fixed effects. t-statistics based on standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively, using two-tailed tests. Detailed definitions of all variables are provided in Appendix A.