

Stage-Specific Generative AI Disclosure: Mention, Planned Adoption, and Implementation

Zirui Huo

Nanjing Normal University Macquarie University Joint Institute, Nanjing Normal University, Nanjing, China

Abstract

As generative AI diffuses across industries, corporate disclosure becomes a strategic signal not only to investors, but also to regulators, employees, and product-market rivals. A central measurement problem is that firms often use similar AI language to communicate very different realities, including broad narrative positioning, forward-looking commitments, and claims of operational deployment. This paper argues that treating these statements as equivalent can distort both managerial inference and capital-market interpretation. We therefore construct stage-specific disclosure measures that separate mention, planned adoption, and implementation in U.S. 10-K filings, rather than relying on undifferentiated keyword intensity. Methodologically, we fine-tune a BERT-based classifier for sentence-level stage classification and then aggregate predicted stages to the filing level. The classifier attains 87.3% accuracy and 86.9% macro-F1 on a held-out test set, supporting scalable measurement in a large panel. We then estimate OLS regressions of abnormal stock price on disclosure composition, controlling for firm characteristics, year and industry fixed effects. Planning language is associated with significantly lower CAR, while implementation language is positive in point estimates but imprecisely estimated. Cross-industry patterns illustrate construct validity: information technology combines high GenAI intensity with high action-oriented disclosure, whereas finance exhibits broad GenAI discussion but comparatively low action shares. Overall, the evidence shows that AI disclosure is multi-dimensional and that composition, not only volume, is crucial for evaluating disclosure credibility, execution risk, and market response.

Keywords: Generative AI, Corporate Disclosure, Text Mining