BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Corporate bond pricing in the AI era

by Marco Albori and Andrea Zaghini

September 2026

Number 1057

BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Corporate bond pricing in the AI era

by Marco Albori and Andrea Zaghini

Number 1057 – September 2026

The series Occasional Papers *presents studies and documents on issues pertaining to the institutional tasks of Banca d'Italia and the Eurosystem. The* Occasional Papers *appear alongside the* Working Papers *series which are specifically aimed at providing original contributions to economic research.*

The Occasional Papers *include studies conducted within Banca d'Italia, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.*

The series is available online at www.bancaditalia.it .

ISSN 1972-6643 (online)

Designed by the Printing and Publishing Division of Banca d'Italia

CORPORATE BOND PRICING IN THE AI ERA

by Marco Albori* and Andrea Zaghini*

Abstract

We study how the emergence of generative artificial intelligence (AI) affected corporate bond pricing. Using bond-level issuance data for U.S. non-financial corporations and a difference-in-differences design centered on the launch of ChatGPT, we examine whether debt investors differentiate firms according to their position in the AI value chain.

We find that hyperscalers experienced a significant decline in borrowing costs relative to comparable issuers, while software firms faced less favorable financing conditions. The results suggest that credit markets rapidly incorporated expectations regarding the distribution of gains and losses from generative AI.

More broadly, the evidence indicates that technological shocks affect not only firm valuations but also the cost of external finance, thereby influencing the allocation of capital across firms during periods of technological transition.

JEL Classification: G12, G32, C23.

Keywords: artificial intelligence, asset pricing, bonds, corporate finance.

DOI: 10.32057/0.QEF.2026.1057

* Banca d'Italia.

1 Introduction

The rapid expansion of artificial intelligence (AI) has triggered an unprecedented wave of investment in digital infrastructure.[1] The construction of the computational capacity required to train and deploy large AI models is already affecting aggregate investment and production dynamics (Carpinelli et al., 2026; Korinek and McKelvey, 2026), while also generating substantial implications for electricity demand and energy prices (Ferriani and Gazzani, 2026). At the same time, the growing size and importance of large technology firms in financial markets has stimulated an intense debate regarding whether current valuations are justified by fundamentals or reflect speculative dynamics associated with expectations about AI adoption (Albori et al., 2025; Caballero, 2026).

More recently, the impact of AI has expanded beyond equity markets and become increasingly visible in corporate debt markets. As planned capital expenditure on data centres, cloud infrastructure and AI services has risen sharply, technology firms have increasingly turned to bond markets to finance investment. If sustained, this development could extend the dominance of AI-related firms from equity markets to credit markets, potentially affecting market concentration, risk allocation and financing conditions for other borrowers.

Despite the growing literature studying the macroeconomic implications of AI and the effects of AI adoption on firm behaviour, little is known about how the AI boom is reflected in corporate bond markets. Most existing studies focus on stock prices and investor expectations, while relatively little attention has been devoted to the financing conditions faced by firms in primary debt markets.

[1]We thank Marco Taboga and Fabrizio Ferriani for valuable comments on the preliminary version of this paper. All opinions expressed here are our own and do not necessarily reflect those of Banca d'Italia.

Our analysis also contributes to the growing literature examining the financing implications of the AI investment cycle. While Aldasoro et al. (2026) document the increasing role of debt financing in supporting AI-related capital expenditure, existing contributions have largely focused on aggregate financing patterns rather than on the pricing of debt. We complement this literature by studying whether investors differentiate across firms according to their exposure to the opportunities and risks created by generative AI. In doing so, we connect the macro-financial implications of AI-related investment (Carpinelli et al., 2026; Korinek and McKelvey, 2026) with the asset-pricing literature that studies how financial markets incorporate technological shocks into security valuations (Andrews and Farboodi, 2025; Gaies, 2026).

More broadly, our paper relates to a growing body of work investigating how firms adjust investment and financing decisions in response to AI adoption. Babina et al. (2023) and Babina (2026) show that firms investing in artificial intelligence significantly modify organisational structures and production processes. To sustain these investments, however, firms require external financing. We examine whether debt markets facilitate or hinder this adjustment process and whether financing conditions differ across firms occupying different positions within the AI ecosystem. Indeed, bond markets play a crucial role in allocating capital, particularly for large-scale infrastructure projects requiring sizeable and long-dated funding. Understanding whether AI affects not only equity valuations but also the cost of debt is therefore important for understanding how capital is allocated during technological transitions.

This paper studies how the emergence of generative artificial intelligence affected corporate bond pricing. We focus on two groups of firms that occupy opposite positions within the AI value chain. The first group comprises the five large hyperscalers:

Alphabet, Amazon, Meta, Microsoft and Oracle. These firms are central to the provision of cloud computing services and AI infrastructure and stand to benefit directly from increasing demand for computational capacity. The second group consists of software firms, whose business models may be complemented but also disrupted by generative AI technologies.

To identify the impact of the AI boom on debt financing conditions, we exploit the public release of ChatGPT on 30 November 2022 as a plausibly exogenous shock that altered investors' expectations regarding the future distribution of profits within the technology sector. Using a difference-in-differences framework and bond-level issuance data, we show that hyperscalers experienced a substantial decline in borrowing costs following the emergence of generative AI, whereas software firms experienced a deterioration in bond pricing conditions.

Our contribution is twofold. First, we provide evidence that the AI boom has affected realized financing conditions in primary credit markets rather than merely investors' expectations reflected in equity prices. Second, we document substantial heterogeneity within the technology sector. The same technological shock that improved financing conditions for suppliers of AI infrastructure simultaneously worsened conditions for firms more exposed to the risk of AI-induced disruption. This finding suggests that credit markets have rapidly incorporated expectations regarding the distribution of gains and losses associated with generative AI.

The remainder of the paper is organised as follows. Section 2 presents the data and stylised facts. Section 3 introduces the empirical strategy. Section 4 discusses the main results. Section 6 provides the conclusions and flags possibile policy implications.

2 Data and Stylised Facts

Our analysis combines information from Dealogic DCM Analytics, LSEG Data & Analytics, and Bloomberg. Dealogic and Bloomberg provide issuance-level information including pricing dates, issuance size, currency denomination and issuer identity, while the annualised yield to maturity at issuance is obtained from LSEG. The final sample consists of approximately 14,800 bond placements by non-financial corporations in the United States between January 2020 and February 2026.

Unlike traditional software development, deploying frontier AI models requires substantial investments in data centres, networking equipment, specialised semiconductors, cooling systems and electricity infrastructure. In contrast to previous waves of software innovation, generative AI is highly capital intensive. Figure 1 put into perspective this remarkable acceleration in capital expenditures. Consistently with the emergence of this new investment regime characterised by large upfront expenditures and long payback horizons, firms competing for leadership in AI markets require large volumes of external finance, making access to bond markets increasingly important.

Indeed, the rapid expansion of AI-related investment has coincided with an extraordinary increase in bond issuance among the largest technology firms. Figure 2 reports issuance activity among US technology companies and highlights the sharp acceleration that occurred after the emergence of generative AI. During 2025, the five hyperscalers issued approximately USD 108 billion of bonds, compared with an annual average close to USD 20 billion during 2023 and 2024. Issuance activity accelerated further during the first months of 2026, with more than USD 56 billion raised in only a few weeks.

Figure 1: **Hyperscalers capital expenditure**

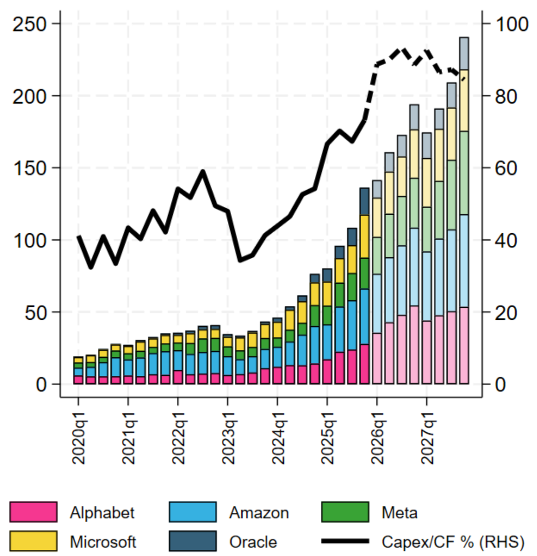

Source: LSEG. The shaded line and dashed bars represent forecasts.

Figure 2: **Issuance of U.S. technology firms**

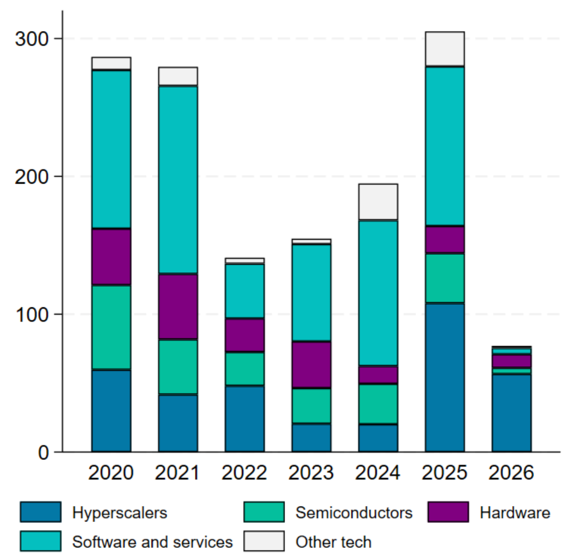

Source: Dealogic DCM Analytics.

The surge in borrowing documented in Figure 2 is consistent with the broader transformation in the financing structure of large technology firms. Aldasoro et al. (2026) document that AI-related investment increasingly exceeds internally generated cash flows, leading firms to rely more heavily on both public debt and private-credit markets.

Concerning the timing of the recourse to the bond market, Figure 2 suggests that the increase in issuance was not merely a continuation of pre-existing trends. While borrowing activity had already recovered after the pandemic period, the expansion became considerably stronger following the commercial breakthrough of generative AI. The divergence is particularly visible for hyperscalers, whose issuance expanded much more rapidly than that of the broader technology sector.

The maturity structure of issuance is equally noteworthy. Several hyperscalers placed bonds with maturities extending beyond thirty years, and in February 2026 Alphabet successfully issued a 100-year sterling-denominated bond. Such long-dated placements suggest exceptionally favourable demand conditions and indicate investors' willingness to lock in exposures to firms perceived as central actors in the AI transition.

From a broader corporate-finance perspective, the shift toward debt financing is remarkable because many of the firms leading the AI revolution historically relied predominantly on internal funds. Large technology companies have traditionally accumulated vast liquidity buffers and maintained limited dependence on external finance. Their growing use of debt markets therefore reflects both the scale of planned investments and a deeper structural transformation. To an increasing extent, hyperscalers resemble infrastructure providers supplying a critical input into the digital economy rather than the asset-light software firms that characterised earlier technological cycles.

Thus, the increase in issuance reflects a broader reconfiguration of the technology sector. While hyperscalers have committed unprecedented amounts of capital to AI infrastructure, software firms face growing uncertainty regarding the consequences of generative AI for existing products and business models. This contrast suggests that the AI boom may have influenced not only the quantity of debt issued by technology companies but also the price at which investors are willing to supply funding.

3 Empirical Strategy

The public release of ChatGPT on 30 November 2022 constitutes a natural reference point for studying the effect of generative artificial intelligence on corporate financing conditions. Although artificial intelligence had been developing for decades, the launch of ChatGPT represented a visible and widely recognised turning point in market perceptions regarding the commercial viability, scalability and profitability of large language models.

Within a few months, investors revised expectations concerning the future revenues associated with cloud computing, AI infrastructure and computational services. Importantly, the shock affected market beliefs rather than firms' contemporaneous fundamentals. The accounting characteristics of technology firms did not suddenly change after November 2022. What changed was investors' assessment of future growth opportunities and competitive positioning.

From the perspective of credit markets, ChatGPT generated a particularly useful source of cross-sectional variation. Some firms were expected to benefit directly from increased AI adoption because they controlled the cloud infrastructure, computational resources and engineering capabilities necessary to support future growth. Other firms faced greater uncertainty because advances in generative AI threatened to alter competitive dynamics within their industries.

Hyperscalers occupied the first category. Investors anticipated that growing demand for training and operating large-scale AI models would increase demand for cloud services, data-centre capacity and computational infrastructure. Software firms were closer to the second category. Investors faced considerable uncertainty regarding whether generative AI would complement or substitute for existing software products and whether traditional Software-as-a-Service business models would remain as profitable in the future. These considerations make the launch of ChatGPT an attractive treatment event for a difference-in-differences framework. The shock was highly visible, clearly dated and plausibly exogenous to bond-market conditions, while simultaneously generating heterogeneous effects across firms depending on their position in the AI value chain.

The choice of ChatGPT as the treatment event is also consistent with recent contributions that view the release of widely adopted generative-AI applications as a turning point in investors' expectations. Andrews and Farboodi (2025) argue that market responses to AI-related news contain valuable information regarding beliefs about the transformative potential of AI technologies. Similarly, Gaies (2026) shows that AI-related uncertainty became an increasingly important driver of financial-market volatility after the emergence of generative AI. The launch of ChatGPT therefore represents not only a technological event but also a shift in financial-market expectations regarding future production technologies and competitive structures.

We estimate the following specification:

$$\begin{aligned} Yield_{b,n,i,c,t} &= \beta_0 + \beta_1 Hyperscaler_b + \beta_2 PostChatGPT_t \\ &\quad + \beta_3 (Hyperscaler \times PostChatGPT_t) + \alpha\, \mathbf{X}_{b,i,c,t} + \gamma \mathbf{Z}_{b,i,c,t} + \lambda_{n,c,t} + \varepsilon_{b,n,i,c,t} \end{aligned} \tag{1}$$

where $Yield_{b,i,c,t}$ is the annualized yield to maturity at issuance of bond b, placed by firm i, from nation n, in currency c, at time t, $Hyperscaler_b$ is a dummy variable tracking the bonds placed by the five hyperscalers, and $PostChatGPT_t$ is the dummy splitting issuances occurred before and after the launch of ChatGPT. We control for bond $\mathbf{X}_{b,i,c,t}$ (value, maturity, rate type, coupon frequency, dummy for collateralized, callable, subordinated, green and investment grade bonds) and issuer characteristics $\mathbf{Z}_{b,i,c,t}$ (nationality, listing, rating, total asset, return on assets, debt ratio, equity). Lastly, $\lambda_{n,c,t}$ includes country per time and currency per time fixed effects. In particular, by interacting the country dummy with time, this specification non-parametrically controls for any unobserved country-specific macroeconomic trends. Most notably, it fully accounts for the FED monetary tightening cycle – characterizing most of the post-treatment period – ensuring that our estimates are not confounded by shifts in the aggregate interest rate environment.

The coefficient of interest is β_3 that assesses whether the yield on the bonds placed by the hyperscalers enjoyed a discount or a premium with respect to other NFCs' bonds in the aftermath of the launch of ChatGPT.

To improve comparability between treated and untreated observations, all regressions are weighted using entropy balancing following Hainmueller (2012). The balancing procedure reweights the control group so that the distribution of observable characteristics closely matches that of the treated group. This approach reduces concerns that differences in issuance yields arise from compositional effects rather than the emergence of generative AI.[2]

Starting from the universe of corporate bond placed in the US market from January 2020 to February 2026, we distinguish the 1,186 treated bonds (those issued by firms providing datacenters) from the 20,412 of the control group (those placed by all the other non-financial corporations). We selected the following variables to be matched in the two samples: quarter of issuance; value placed; maturity at issuance; dummy for collateralized bonds; dummy for subordinated bonds; dummy for callable bonds; bond

[2]See Tsang et al. (2024), Yu et al. (2024), Di Tommaso et al. (2025), Fornari et al. (2026) for recent contributions in the empirical finance literature relying on the entropic matching approach.

rating; currency of denomination; coupon type; issuer size. In particular, for dummies and discrete (coded) variables we required the first moment to be matched, while for the continuous variables also the second moment has to be matched (Table 1).

Table 1: Summary statistics before and after entropy balancing

	Treatment		Control before		Control after	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Quarter of issuance	28.14	7.53	28.23	7.33	28.14	7.31
Issue size ($ bn)	802.9	891.29	481.7	439.69	802.9	891.27
Maturity (days)	4,627	4,219	4,028	4,527	4,628	4,219
Collateralized	0.3179	0.466	0.3864	0.487	0.3179	0.466
Subordinated	0.2125	0.409	0.2005	0.400	0.2125	0.409
Callable	0.9519	0.214	0.8916	0.311	0.9519	0.214
Rating	1.938	0.244	1.754	0.460	1.938	0.278
Currency	93.53	5.68	93.42	5.93	93.53	5.43
Coupon type	3.293	2.06	3.200	2.01	3.293	2.06
Total assets ($ bn)	164.574	124.10	73.696	127.67	164.575	124.10

Prior to reweighting, treated and control bonds differ across several observable dimensions. Treated issuers are substantially larger, exhibit stronger credit quality, and issue considerably larger debt placements than firms in the control group. Their bonds also feature longer maturities and a higher incidence of callability, while collateralized securities are relatively less common. These systematic differences suggest that simple comparisons of issuance yields may confound treatment effects with underlying differences in issuer and bond characteristics. The entropy-balancing procedure addresses this concern by aligning the distribution of observables in the control group with that of treated observations.

A key assumption underlying the difference-in-differences design is that treated and control firms would have followed parallel trends in the absence of treatment. To assess this assumption we estimate a dynamic event-study specification around the launch of

ChatGPT. In particular, we ran a regression interacting the dummy $Hyperscalers_b$ with six-month periods before and after November 2022.

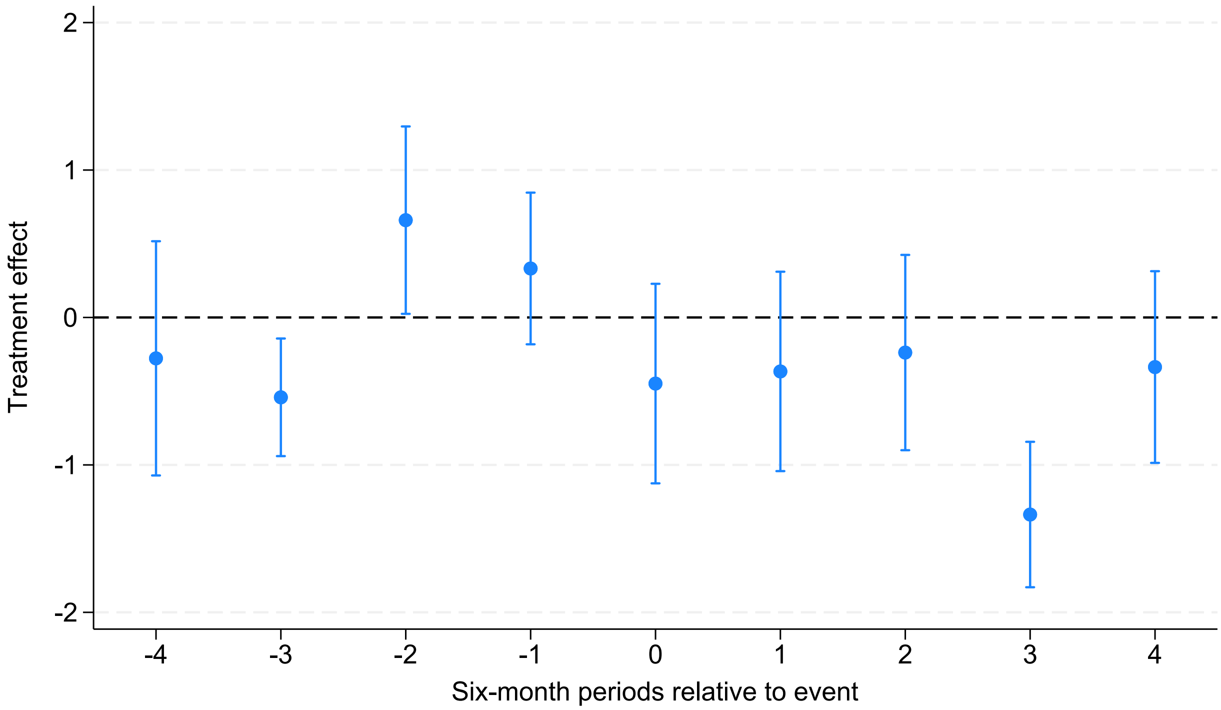

Figure 3: Dynamic event-study around the launch of ChatGPT.

The results are reported in Figure 3. The coefficients in the pre-treatment period are generally small and statistically indistinguishable from zero. Although one coefficient becomes significant shortly before the event, there is no evidence of a systematic divergence between treatment and control firms before November 2022. This absence of a clear pre-trend supports the validity of the parallel-trends assumption.

The post-treatment dynamics are equally informative. Following the launch of ChatGPT, the coefficients turn negative and remain predominantly below zero throughout the remainder of the sample. The magnitude of the effect gradually increases and reaches its maximum approximately eighteen months after the shock. Rather than displaying an immediate one-off reaction, debt markets appear to have progressively revised their assessment of firms' prospects as information about AI adoption, capital expenditure plans and infrastructure investment accumulated. Indeed, this gradual adjustment is economically plausible. While ChatGPT demonstrated the feasibility of generative AI, investors initially faced considerable uncertainty regarding its long-run commercial implications. Subsequent announcements regarding in particular cloud demand and data-centre expansion likely reinforced market perceptions that certain firms would emerge as major beneficiaries of the new technological environment.

The dynamic evidence therefore supports both the chosen treatment date and the causal interpretation of the main difference-in-differences results.

4 Results

Table 2 reports the main estimation results. Column (1) presents a baseline specification estimated over the entire sample period (January 2020 to February 2026) and taking into account all non-financial corporations. Hyperscalers do not exhibit a significantly different issuance yields with respect to other borrowers. However, this average effect masks substantial differences between the pre- and post-ChatGPT periods. When we move to the full model described in Equation 1 (Column 2), we find that hyperscalers have experienced a marked improvement in bond financing conditions of 47 basis points during the AI boom. Interestingly, before the launch of ChatGPT, the yield on their bonds, while not statistically significant, was even positive. Combining the treatment coefficient and interaction term yields a total post-treatment discount of approximately 29 basis points.

A similar pattern emerges when the control group is restricted to technology firms only. The interaction coefficient remains strongly negative and statistically significant, implying that the improvement in borrowing conditions cannot be attributed solely to favourable developments affecting the technology sector as a whole. Relative to other technology firms, hyperscalers experience a post-treatment discount of roughly 26 basis points.

Table 2: Difference-in-Differences Estimates of Bond Issuance Yields. Dependent variable is the yield at issuance. Robust standard errors clustered by country are reported in parentheses. Fixed effects include currency per time and nationality per time effects. Bond and issuer controls are included in all specifications. Statistical significance: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1) All NFCs	(2) All NFCs	(3) All Tech	(4) Tech no Soft	(5) Software
Hyperscaler	-0.0795	0.1791	0.0908	0.3444	-0.0343
	(0.2569)	(0.1997)	(0.1501)	(0.3014)	(0.0684)
Hyperscaler × PostGPT		-0.4690***	-0.3461***	-0.5728***	-0.3841***
		(0.1081)	(0.0641)	(0.1045)	(0.0983)
Post-treatment effect	-0.0795	-0.2899*	-0.2553**	-0.2283	-0.4184***
	(0.2569)	(0.1789)	(0.1063)	(0.3976)	(0.0535)
Observations	17,008	17,008	3,641	2,165	1,619
R^2	0.795	0.824	0.811	0.851	0.830
Bond controls	Yes	Yes	Yes	Yes	Yes
Issuer controls	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes

These findings indicate that investors substantially revised their assessment of hyperscalers after the emergence of generative AI. Firms perceived as central suppliers of cloud computing and AI infrastructure were able to issue debt at increasingly advantageous terms despite a macro-financial environment characterised by rising policy rates and tighter financing conditions. One possible interpretation is that investors viewed AI-related capital expenditure not as a source of additional financial risk but rather as an investment opportunity capable of generating substantial future cash flows. Under this interpretation, expectations regarding future profitability dominated concerns about rising leverage and borrowing needs. More broadly, the evidence suggests that debt investors rapidly integrated expectations regarding the AI revolution into bond pricing. The emergence of a sizeable and persistent advantageous borrowing-cost spread among hyperscalers indicates that the AI boom affected not only capital expenditure decisions but also firms' access to external finance. We label this yield spread the "hyperium", borrowing from the green-finance literature where the term "greenium" denotes the pricing advantage enjoyed by green securities.

The results reported in Columns (4) and (5) provide further insight into the distributional consequences of the AI shock within the technology sector. While the preceding specifications establish that hyperscalers experienced a significant decline in borrowing costs after the launch of ChatGPT, the final two columns examine whether this advantage persists when the comparison group is narrowed to firms that are themselves highly exposed to AI-related developments.

Column (4) excludes software companies from the technology sample and compares hyperscalers with the remaining technology universe. The interaction coefficient remains negative, economically large and highly statistically significant. Following the launch of ChatGPT, hyperscalers obtained an additional reduction in issuance yields of approximately 57 basis points relative to non-software technology firms, placing the hyperium at 23 basis points. This result confirms that the favourable repricing of hyperscaler debt cannot be attributed solely to differences between technology and non-technology issuers. Investors appear to assign a distinct credit premium to firms controlling the infrastructure layer of the AI ecosystem.

The contrast becomes even more informative in Column (5), which restricts the control sample to technology firms providing software. This specification places hyperscalers against a group of firms that, at least superficially, might also be expected to benefit from the diffusion of generative AI. Nevertheless, the interaction term remains strongly negative and statistically significant. The estimated post-treatment effect implies that hyperscalers enjoyed issuance yields approximately 42 basis points lower than comparable AI-oriented software firms after the emergence of ChatGPT, the largest of all regressions.

The persistence of the effect within this narrow sample suggests that debt investors did not view all AI-related firms as equal beneficiaries of the technological shift. Instead, investors appear to have distinguished sharply between firms supplying the critical infrastructure required for AI deployment and firms whose business models may themselves be disrupted by AI-enabled competition. This interpretation is broadly consistent with the emergence of what market participants increasingly described as the SaaSpocalypse: the concern that generative AI could weaken some of the fundamental sources of competitive advantage traditionally enjoyed by software companies.

Historically, software firms benefited from strong customer lock-in, high switching costs, proprietary data environments and recurring subscription revenues. Generative AI has the potential to erode these advantages. Large language models can increasingly replicate functionalities previously embedded in specialized software applications, reduce the effort required to switch providers, and lower barriers to entry for new competitors. As a result, although AI creates growth opportunities for many software companies, it simultaneously introduces substantial uncertainty regarding the long-term sustainability of incumbent business models.

From a credit-market perspective, such uncertainty is particularly important. Bond investors are concerned less with upside potential than with the predictability and resilience of future cash flows. A technological shock that increases uncertainty about future revenues, profit margins or competitive positioning may therefore raise required credit spreads even when current operating performance remains strong. The results in Column (5) suggest that investors attached significantly greater confidence to the future cash-flow generation of hyperscalers than to that of AI-exposed software firms.

Taken together, the evidence points to a growing bifurcation within the technology sector. The emergence of generative AI did not lead debt markets to reward technology firms indiscriminately. Rather, investors increasingly differentiated between firms expected to capture the infrastructure rents generated by AI and firms exposed to the risk of technological disruption. Even when compared exclusively with software companies, hyperscalers continued to enjoy significantly more favourable financing conditions. The results therefore suggest that corporate bond markets rapidly identified not only the likely beneficiaries of the AI transition, but also the segments of the technology sector most vulnerable to its disruptive effects.

5 Robustness checks

We provide a series of robustness exercises designed to assess whether the baseline findings depend on the particular measure of borrowing costs, the composition of the control group, or the definition of the treated firms.

First, Columns (1)–(3) of Table 3 replace the yield at issuance with the spread over a sovereign bond of comparable maturity. This specification isolates the firm-specific component of borrowing costs and removes the influence of changes in the risk-free term structure. The launch of ChatGPT coincided with a period of substantial monetary tightening in the United States and globally. Although our baseline specification already controls for time-varying country and currency effects, the spread measure provides an even stricter test by netting out movements in benchmark sovereign yields. The results remain remarkably similar to the baseline findings. The interaction coefficient is negative and statistically significant in all specifications, while the post-treatment effect continues to imply a meaningful borrowing-cost advantage for hyperscalers. This suggests that the estimated hyperium is not driven by changes in the level of interest rates but rather reflects a genuine reassessment of firm-specific credit risk. Moreover, Column (3) adds support to the interpretation that software companies are those that suffered the most from the AI boom: the yield spread being again the largest.

Second, Columns (4)–(6) restrict the analysis to publicly listed issuers. This exercise addresses concerns regarding comparability between hyperscalers and the broader population of bond issuers. Listed firms are generally larger, more transparent, more closely followed by financial analysts, and subject to stricter disclosure requirements. As a result, they constitute a more appropriate comparison group for the large publicly traded firms that make up the hyperscaler sample. The estimated effects remain economically large and statistically significant. Actually, the post-treatment discount enjoyed by hyperscalers becomes even more pronounced in the full non-financial corporation sample, reaching approximately 56 basis points. The persistence of the results within this more homogeneous set of issuers indicates that the baseline findings are not driven by differences between publicly traded and privately held firms.

Finally, Columns (7)–(9) reverse the perspective of the analysis and estimate the model from the viewpoint of software firms rather than hyperscalers. If the baseline results truly reflect a growing divergence between AI infrastructure providers and firms exposed to technological disruption, the coefficient estimates should mirror those obtained in the main specifications. This is precisely what emerges from the data. Column (7) compares software firms with the universe of non-financial corporations. The interaction term is positive but statistically insignificant, indicating that software issuers did not experience a significant change in financing conditions relative to the broader corporate sector following the launch of ChatGPT. This result is unsurprising given the heterogeneity of the control group, which includes firms operating in industries largely unrelated to the AI transition. When the comparison is restricted to technology firms in Column (8), the picture changes considerably. Software companies already exhibited a sizeable yield premium before the AI shock, and this premium widened significantly by further 25 basis points after the launch of ChatGPT. The the estimated post-treatment effect implies a very large premium of 87 basis points relative to the rest of the technology sector. This finding suggests that debt investors became substantially more cautious regarding the prospects of software firms following the emergence of generative AI.

Column (9) provides instead the most direct counterpart to the baseline analysis by restricting the sample to software firms and hyperscalers only. In this specification, the estimated post-treatment premium for software firms reaches approximately 42 basis points, with the interaction coefficient remaining highly significant. Note that this magnitude is numerically identical to the borrowing-cost discount obtained for hyperscalers in Table 2, but with the opposite sign.[3] Rather than rewarding all technology firms indiscriminately, debt investors appear to have differentiated sharply between firms expected to capture the rents associated with AI infrastructure and firms facing greater uncertainty regarding the future viability of their business models.

All in all, the provided robustness exercises reinforce the conclusion that the emergence of generative AI generated a substantial reallocation of credit-market pricing across firms occupying different positions in the AI value chain.

[3]Actually, the regressions in Columns (5) of Table 2 and Column (9) of Table 3 are exactly the same with switched treated and control groups.

Table 3: Robustness checks. Dependent variable is alternatively the spread over the sovereign benchmark of comparable maturity (Columns 1–3) or the issuance yield (Columns 4–9). Robust standard errors clustered by country are reported in parentheses. Bond controls, issuer controls and fixed effects are included in all specifications. Statistical significance: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	All NFCs	Tech Only	Software	All NFCs	Tech Only	Software	All NFCs	Tech Only	H5 Only
	Spread	Spread	Spread	Listed	Listed	Listed	Software POV	Software POV	Software POV
Hyperscaler	0.2413	0.1458	0.0445	-0.1759	0.1459	-0.3237			
	(0.2009)	(0.1585)	(0.0703)	(0.1449)	(0.1425)	(0.0728)			
Hyperscaler × PostGPT	-0.4932***	-0.3616***	-0.3761***	-0.3877***	-0.4530***	-0.3679***			
	(0.1194)	(0.0860)	(0.1008)	(0.1067)	(0.0463)	(0.0374)			
Software							0.1446	0.6209***	0.0343
							(0.2418)	(0.1066)	(0.0684)
Software × PostGPT							0.1346	0.2514***	0.3841***
							(0.4019)	(0.0704)	(0.0983)
Post-treatment effect	-0.2520	-0.2157**	-0.3316***	-0.5636***	-0.3071**	-0.6916***	0.2792	0.8723***	0.4184***
	(0.1791)	(0.1011)	(0.0527)	(0.1114)	(0.1167)	(0.0400)	(0.3388)	(0.0723)	(0.0535)
Observations	17,000	3,639	1,619	14,956	3,103	1,411	17,066	3,641	1,619
R^2	0.773	0.718	0.768	0.788	0.818	0.841	0.795	0.812	0.830
Bond controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Issuer controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

6 Conclusion

This paper examines how the emergence of generative artificial intelligence affected financing conditions in corporate bond markets. Using bond-level issuance data and a difference-in-differences framework centered on the launch of ChatGPT, we document a substantial divergence in borrowing costs across technology firms. Hyperscalers, which occupy a central role in the provision of cloud computing and AI infrastructure, experienced a significant decline in issuance yields. Software firms, by contrast, faced a deterioration in financing conditions, particularly relative to hyperscalers and among issuers whose business models appear more vulnerable to AI-related disruption.

These findings suggest that credit markets have rapidly incorporated expectations regarding the distribution of gains and losses associated with generative AI. Investors appear willing to provide cheaper financing to firms positioned to benefit from the AI infrastructure buildout, while demanding higher compensation from firms facing uncertainty concerning the future viability of existing business models. More broadly, the results reveal that the AI revolution is reshaping not only equity valuations but also debt-market pricing. The growing demand for infrastructure investment creates financing advantages for a relatively small group of firms controlling critical AI inputs, while simultaneously generating pressures for firms exposed to technological disruption. In this respect, our findings are consistent with the broader view that periods of technological transformation are accompanied by substantial capital reallocation toward firms perceived to be best positioned to exploit new technologies (Perez, 2002).

The results also carry broader policy implications. By lowering financing costs for AI infrastructure providers and at the same time increasing the relative cost of capital for firms more vulnerable to the AI boom, bond markets may reinforce the concentration of investment and market power within a small number of dominant technology companies. This mechanism echoes concerns in the literature regarding the emergence of "superstar firms" and the increasing concentration of economic activity among a small number of highly productive firms (Autor et al., 2020). The rapid differentiation in financing conditions highlights the growing role of credit markets in transmitting the economic consequences of technological change. While such repricing is a normal feature of efficient capital allocation, it may also create refinancing challenges for firms perceived to be on the losing side of the AI transition.

As AI-related investment continues to expand, the increasing importance of large technology firms in corporate bond markets may affect the distribution of financing costs across borrowers, increase market concentration, and alter the composition of corporate debt issuance. The rapid growth of debt financing linked to AI infrastructure also raises the possibility of future valuation adjustments. If expectations regarding AI adoption, profitability, or technological leadership were to change, the large stock of debt accumulated by AI-related firms could become subject to sharp repricing. Such dynamics are particularly relevant in sectors characterized by large investments in intangible and technology-related assets, whose financing capacity can be highly sensitive to revisions in growth expectations (Falato et al., 2022). Monitoring these developments is therefore relevant not only for corporate finance but also for financial stability. In addition, it is critical to actively evaluate whether these pricing effects persist as AI adoption becomes more widespread and whether the growing financing needs of AI-related firms generate spillovers to borrowing conditions in other sectors of the economy

References

ALBORI, M., V. N. LANDI, AND M. TABOGA (2025): "Is there a tech bubble in the US stock market? Evidence from an agnostic valuation procedure," Questioni di Economia e Finanza (Occasional Papers) 975, Bank of Italy, Economic Research and International Relations Area.

ALDASORO, I., S. DOERR, AND D. REES (2026): "Financing the AI boom: from cash flows to debt," BIS Bulletins 120, Bank for International Settlements.

ANDREWS, I. AND M. FARBOODI (2025): "Do Markets Believe in Transformative AI?" NBER Working Papers 34243, National Bureau of Economic Research, Inc.

AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 135, 645–709.

BABINA, T. (2026): "Understanding Firms' AI Efforts and Their Economic Impact," NBER Working Papers 35123, National Bureau of Economic Research, Inc.

BABINA, T., A. FEDYK, A. X. HE, AND J. HODSON (2023): "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition," NBER Working Papers 31325, National Bureau of Economic Research, Inc.

CABALLERO, R. J. (2026): "Speculative Growth and the AI Bubble," NBER Working Papers 34722, National Bureau of Economic Research, Inc.

CARPINELLI, L., F. NATOLI, AND M. TABOGA (2026): "Artificial intelligence and the US economy: an accounting perspective on investment and production," Questioni di

Economia e Finanza (Occasional Papers) 1006, Bank of Italy, Economic Research and International Relations Area.

Di Tommaso, C., S. Perdichizzi, S. Vigne, and A. Zaghini (2025): "Is the Government always greener?" *Energy Economics*, 152, 108976.

Falato, A., D. Kadyrzhanova, J. W. Sim, and R. Steri (2022): "Rising Intangible Capital, Shrinking Debt Capacity, and the U.S. Corporate Savings Glut," *Journal of Finance*, 77, 2799–2854.

Ferriani, F. and A. Gazzani (2026): "Powering the digital economy: the global expansion of data centres and its energy implications," Questioni di Economia e Finanza (Occasional Papers) 1014, Bank of Italy, Economic Research and International Relations Area.

Fornari, F., D. Pianeselli, and A. Zaghini (2026): "Environmental score and bond pricing: It better be good, it better be green," *Journal of International Money and Finance*, 161, 103498.

Gaies, B. (2026): "AI uncertainty and global stock market volatility: Any signals of a Dot-com 2.0?" *Economics Letters*, 264, 112954.

Hainmueller, J. (2012): "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 20, 25–46.

Korinek, A. and P. McKelvey (2026): "Measuring the AI economy," Working Paper Series WP26-9, Peterson Institute for International Economics.

PEREZ, C. (2002): *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Cheltenham, UK: Edward Elgar Publishing.

TSANG, A., Y. WANG, Y. XIANG, AND L. YU (2024): "The rise of ESG rating agencies and management of corporate ESG violations," *Journal of Banking & Finance*, 169, 107312.

YU, Q., E. C.-M. HUI, AND J. SHEN (2024): "The real impacts of third-party certification on green bond issuances: Evidence from the Chinese green bond market," *Journal of Corporate Finance*, 89, 102694.