

Capital Market Consequences of Generative AI: Early Evidence from the Ban of ChatGPT in Italy

Jeremy Bertomeu
Washington University in Saint Louis

Yupeng Lin
National University of Singapore

Yibin Liu
National University of Singapore

Zhenghui Ni
National University of Singapore

Oct 6th, 2023

Abstract

On March 31, 2023, the Italian Data Protection Authority found ChatGPT in violation of privacy laws and banned the service in Italy, providing a natural experiment to evaluate the economic impact of generative AI. Italian firms with higher exposure to the technology suffered a negative market reaction of about 9% during the ban period compared to firms with lower exposure. The negative impact was greater for smaller and more newly established businesses, supporting a link with creative destruction. We also document that the ban affected the information environment. Italian-based analysts issued fewer forecasts than foreign analysts covering the same Italian firm. Moreover, bid-ask spreads widened during the ban, especially for firms with fewer institutional investors, limited analyst coverage, and fewer foreign investors. Overall, our results suggest a dual role for generative AI in enhancing firm productivity and information processing.

Keywords: *AI, Generative, ChatGPT, privacy, disruption, innovation, information processing capacity.*
JEL Codes : O00, O33, O38, E44, E65, G10, G14.

Jeremy Bertomeu is an Associate Professor at the Olin School of Business at Washington University in St Louis. Yupeng Lin is an Associate Professor at the National University of Singapore. Yibin Liu is an Assistant Professor at the National University of Singapore. Zhenghui Ni is a Ph.D. candidate at the National University of Singapore. We thank John Barrios, Ben Charoenwong, Edwige Cheynel, Shushu Jiang, Ranjani Krishnan, Oliver Zhen Li, Chenkai Ni, and other seminar participants at Washington University in Saint Louis, Baruch College, Fudan University, Rutgers University, and the Shanghai Lixin Accounting Conference for many helpful suggestions. Contact information: bj Jeremy@wustl.edu.

“ChatGPT is a development on par with the printing press, electricity and even the wheel and fire”

— Lawrence H. Summers, former US Treasury Secretary

I. Introduction

The advent of generative AI has revolutionized access to computer-assisted productivity tools. Leading contenders in this field, such as Google’s Bard and OpenAI’s ChatGPT, have enabled organizations around the globe to automate tasks that once required trained employees and extensive infrastructure, creating a powerful force for accelerated creative destruction (Schumpeter, 1942). However, the many applications of the technology make it challenging to characterize and analyze its economic value across a wide range of tasks. Ideally, a randomized experiment would be required to evaluate its economy-wide impact: In such an experiment, a sample of representative firms would be randomly treated with a plausibly exogenous shock affecting access. Unfortunately, conducting such an experiment, which would deprive firms of the potentially large benefits of the technology, seems improbable.

Nevertheless, on March 31, 2023, Italy banned ChatGPT within its geographic borders and blocked access to servers for all IP addresses from Italy due to concerns about consent and personal data protection. After nearly a month of discussions and improvements, Italian authorities re-authorized the service on April 28, 2023.

This natural experiment provided a remarkable setting to assess the effect of generative AI on an entire economy. First, the ban was enforced by a regulator to comply with privacy laws, not in response to economic factors related to ChatGPT, such as perceived productivity gains or lobbying efforts by professional groups directly affected by the technology. Second, there was no advance warning from Italian regulators of their intention to impose the ban, nor was there any ongoing public discourse with OpenAI, implying that the ban came as a surprise to financial markets. Moreover, the abrupt nature of the ban in terms of timing and geography allows us to better control

for trends toward the adoption of related technologies correlated to the use of generative AI. Third, the ban took effect after ChatGPT was already widely used in Italy. This provides a unique opportunity to observe the impact of the technology across a range of market participants.

Drawing on Acemoglu et al. (2018), our approach to assessing the impact of the ChatGPT ban is to examine the market reaction for different firms. This approach allows us to estimate the value of generative AI by analyzing changes in the relative stock market valuations of firms with varying levels of exposure (Eloundou et al., 2023; Eisfeldt et al., 2023).¹ On the one hand, greater exposure implies that certain tasks within a firm's operations can be automated or facilitated. Consequently, a ban is likely to result in a relative reduction in productivity for these firms.² On the other hand, exposure may facilitate replacing or disrupting outdated business models while new businesses emerge (Caballero et al., 1998, 2001, 2005; Acemoglu and Restrepo, 2018). From this perspective, the ban's impact on market values may depend on the nature of their businesses, as the ban may either protect or impede pre-existing operations.

To analyze the potential impact of creative destruction, we adopt the generative AI exposure measure proposed by Eisfeldt et al. (2023), which assesses the level of exposure to ChatGPT by

¹ In a supplementary appendix, we find that the ban leads to insignificant change in the aggregate Italian equity market value. One possible explanation is that listed firms are larger on average and, therefore, the ban, reducing creative destruction targeted toward larger incumbents, marginally benefited larger (more likely to be traded) incumbents (Caballero et al., 2008) at the expense of smaller private rivals. Unfortunately, there are significant empirical challenges in measuring the expected return of an entire economy in an event study, so we cannot accurately estimate a benchmark counter-factual market return if the ban had not taken place. Consequently, we follow Acemoglu et al. (2018) and focus on firms with different levels of exposure in our main tests and examine different methodologies to correct for market-level return in a supplementary appendix.

² As an illustrative example, we show in Figure A1 how the stock price of Expert.AI, a listed firm on the Milan Stock Exchange, responded to the ChatGPT ban. Expert.AI specializes in providing organizations with data analysis and business solutions. In its 2022 annual report, Expert.AI emphasized the importance of technologies such as "GPT," "Machine Learning," "Natural Language Processing," "Artificial Intelligence," "Machine Translation," and "Speech Recognition," mentioning the keywords more than 40 times. Additionally, on February 15, 2023, Expert.AI announced the launch of its Hybrid AI Platform that integrates GPT, highlighting its reliance on Generative AI (source: Refinitiv/PR Newswire). The stock price of Expert.AI declined by approximately -8% on the announcement day of the ban.

examining its ability to perform specific tasks.³ We build value-weighted long portfolios consisting of industries with high AI exposure, specifically firms in the Professional, Scientific, and Technical Services (54) and Information (51) sectors. Simultaneously, we create value-weighted short portfolios of industries with low exposure. Figure 1 shows that the long-short portfolio strategy results in a negative cumulative return of -1.50% over three days following the announcement of the ban on March 31, 2023. Over the entire ban, the long-short portfolio exhibits a negative cumulative return of -5.48%.⁴

An alternative explanation is that the underperformance of ChatGPT/AI-exposed firms reflects the overall trends of AI-related industries rather than the ban.⁵ To alleviate this concern, we construct a matched sample based on European listed firms with similar size, return on assets (ROA), and from the same industry. Using a difference-in-differences research design, our results in Table 2 suggest that Italian firms with greater exposure experience a decline in firm value of about 9% throughout the ban.

Another possibility is that the exposure measure developed by Eloundou et al. (2023) and Eisfeldt et al. (2023) captures systematic risks in the economy because the technology sector is cyclical. To alleviate this concern, we conduct a falsification test utilizing Italy's negative sovereign rating outlook issued by Moody's on August 6, 2022. If the underperformance of firms with higher exposure to ChatGPT is solely driven by an unmodeled country-level risk exposure, we expect similar effects. Our findings are inconsistent with this alternative interpretation.

³ Eisfeldt et al. (2023) use ChatGPT to evaluate whether the tasks performed by different occupations can be executed or improved by ChatGPT's capabilities. This measure captures the potential of generative AI to enhance or replace tasks performed by pre-existing labor.

⁴ In Figure A2, we find no notable reversal in the market value of high AI-exposed firms in the two months following the reactivation of ChatGPT on April 28, 2023. The absence of reversal suggests that the short-term ban could have damaged investors' confidence in the long-run prospects of AI-exposed firms in Italy, perhaps due to uncertainties about regulations on data privacy. Investors may still remain uncertain about the temporary nature of the reactivation and refrain from making significant changes to their investments.

⁵ For example, there could be increased regulatory scrutiny on AI-related industries in EU, leading to a decline in the stock prices of these industries in the entire EU.

We then explore an alternative approach to ChatGPT exposure by assessing shareholders' search behavior for AI-related information. This method allows us to capture the importance of the technology for the firm's shareholder base. We find that the ban lowers the stock value of firms with greater attention. In addition, we define firms as having high exposure if the number of mentions of AI topics in the most recent earnings conference call is above the median. This measure seeks to capture the exposure to AI using firms' disclosure. Italian firms with high attention (more mentions of AI topics) experience a negative return of -6.8% (-8.7%) throughout the ban.

Lastly, firms at different stages of development may be differentially affected. On the one hand, larger firms and firms with extensive operational experience may achieve higher returns from adopting new scalable data-based technology (Farboodi et al., 2019; Babina et al., 2022; Chen and Srinivasan, 2023), making them more vulnerable to the negative impact of the ban. On the other hand, unlike other AI technologies, ChatGPT enables smaller and new firms to take advantage of OpenAI's non-rivalrous information production, helping them reduce the information advantage of larger incumbents. Our results support the latter hypothesis of a more stronger impact on newer and smaller firms, suggesting that ChatGPT differs from other data-based technologies and enable new and smaller firms to narrow the competitive gap. The ban may hinder the adoption of this new technology, slow down creative destruction, and negatively affect smaller and new firms (Caballero et al., 2008; Acemoglu and Restrepo, 2018).

The ChatGPT ban may also have far-reaching effects on users and producers of information in the capital markets. Accessing and analyzing relevant textual and financial information to make informed investment decisions may become more complex and time-consuming without ChatGPT (e.g., Li, 2008; Dong et al., 2016; Blankespoor et al., 2019). To accurately identify this effect of ChatGPT on information processing capacity, we first focus on financial analysts, who are an

important group of information producers and could potentially be affected by the ban. To isolate the impact of the ban on analysts, we use the country codes of analysts' office phone numbers to infer their locations and therefore identify treated analysts located in Italy. Comparing to analysts in other countries covering the same firm, we document a decrease in the number of analyst forecasts. Our results thus provide early evidence of the impact of AI on information intermediaries.

The use of ChatGPT can affect the capital market by changing the information processing capacity of a wide range of investors (Blankespoor et al., 2020). ChatGPT changes this landscape by providing a conversational and interactive platform for various investors, especially local retail investors, to gather insights and navigate the complexity of financial markets. Consistent with this conjecture, we show that Italian stocks experience an increase in bid-ask spreads during the ban. We expect and confirm that firms with fewer institutional investors and lower analyst coverage experience a more significant increase in bid-ask spreads. Since the ban is limited to the Italian geographic boundaries, foreign investors are unlikely to be affected. We find that firms with a higher share of foreign investors experience a smaller increase in bid-ask spreads during the ban. Taken together, our findings support the conclusion that the ban leads to higher information acquisition costs and thus greater information asymmetry (Verrecchia, 1982).

Our research contributes to multiple strands of literature. First, we add to the literature on measuring the impact of disruptive technologies on firm value. Prior work has examined how technological advances affect firm valuations both analytically and empirically, such as investment-specific technological changes (Papanikolaou, 2011; Kogan and Papanikolaou, 2014), routine automation (Zhang, 2019), and artificial intelligence (Babina et al., 2022, 2023).⁶ Recent

⁶In addition to firm value, the literature has examined the economic effects of revolutionary technological advancements. As examples, prior studies estimate the impacts of the printing press in European cities from 1450 to 1600 (Dittmar, 2012), the telegraph in China's banking industry during the 19th century (Lin et al., 2021), the extension of high-speed transportation (Chen et al., 2022) and, recently, the expansion of 3G in the US (Fang et al. 2023). Bloom et al. (2021) examine the diffusion of 29 disruptive technologies across firms and geographical locations in the U.S.

work by Einfeldt et al. (2023) measures U.S. public firms' labor force exposure to ChatGPT following Eloundou et al. (2023) and demonstrates that firms in the highest exposure quintile outperformed those in the lowest quintile by 0.4% on a daily basis following the release of ChatGPT in November 2022.

Second, we add to the information processing literature reviewed by Blankespoor et al. (2020) by exploiting exogenous variation resulting from the ban. We further contribute to the recent literature in finance and accounting demonstrating the capacity of machine learning to extract information from unstructured sources, such as determining the sentiment of news headlines (Lopez-Lira and Tang, 2023), predicting future earnings-per-share (Li et al., 2023), forming economic expectations based on news articles (Bybee, 2023), improving the readability of management disclosures and conference calls (Frankel et al., 2022; Kim et al., 2023), quantifying new information conveyed by managers in conference calls (Bai et al., 2023), measuring managers' anticipated capital expenditure (Jha et al., 2023), measuring business complexity (Bernard et al., 2023), and predicting innovation success (Yang, 2023).⁷ Our results provide evidence of how the adoption of disruptive AI technologies affects the information processing capacity of capital market participants.⁸

⁷ In addition to text analysis, ChatGPT can potentially complete a broader range of tasks. For example, an emerging literature shows that ChatGPT can assist in auditing (Emett et al., 2023; Eulerich and Wood, 2023; Gu et al., 2023), test-taking (OpenAI, 2023; Eulerich et al., 2023; Katz et al., 2023), answering common and accounting questions (Wood et al., 2023; O'Leary, 2023a, 2023b), and conducting research (de Kok, 2023; Korinek, 2023).

⁸ To our knowledge, Kreitmer and Raschky (2023) is the only other study that examines the impact of an exogenous shock to generative AI, and it also examines the context of the Italian ban. Kreitmer and Raschky (2023) find a short-term effect of the ban on the output of coding engineers. They also document that this sophisticated user base quickly found remedial strategies, as evidenced by a recovery of coding activity after a few days and an increase in search for proxy servers (VPN). Italian businesses could, in principle, use VPN to bypass IP restrictions. However, not all employees would want to set up a VPN, which might lower connection speed, or would be allowed to do so from their company network, especially for public firms with network security restrictions. In contrast, many coding engineers contributing open source code are self-employed or work independently or for smaller private firms, which might involve fewer restrictions.

II. Institutional Background and Data

2.1. Background on Ban and Restoration in Italy

OpenAI is a non-profit laboratory founded in 2015 with the goal of promoting AI development. The organization first released GPT-1 as open source on June 11, 2018, and GPT-2 on February 14, 2019 (Radford et al. 2018, 2019), earning general acclaim. GPT-1 and GPT-2 were primarily designed as chatbots with capabilities for text generation rather than as reliable productivity tools or for accurate research or problem-solving. On November 30, 2022, OpenAI released ChatGPT, built on the GPT-3 architecture, first to select beta testers and then to the general public on January 30, 2023. The increased size of GPT-3, approximately 100 times the size of GPT-2, made GPT-3 versatile at performing a wide range of complex natural language tasks. GPT-4 was released on March 14, 2023, with about 1,000 times the number of parameters of GPT-3 with vastly improved capabilities to perform complex tasks (OpenAI 2023).

On March 20, 2023, OpenAI paused the service for four days after discovering a security vulnerability allowing access to the chat history and personal information of other users over a nine-hour window. Although OpenAI patched the vulnerability, shortly after the breach, the Italian data protection authority (Garante per la protezione dei dati personali) determined that OpenAI had violated EU privacy laws. On March 30, Italian regulators issued an initial decision (Registro dei provvedimenti n. 112) requiring “temporary limitation of the processing of personal data of data subjects established in the Italian territory” (translated from Italian) and then clarified the decision on a statement dated March 31, noting that “It will have to notify the Italian SA within 20 days of the measures implemented to comply with the order, a fine of up to EUR 20 million EUR or 4% of worldwide sales may be imposed.”

The decision was largely unexpected and prompted OpenAI to discontinue its service for IP addresses originating from Italy on March 31, 2023. The news received mixed reactions, with

Deputy Prime Minister Matteo Salvini stating, “This is a work tool that many young people, many companies, many start-ups were using and I hope that they can use it again as soon as possible, because otherwise Italy will have a gap compared to all other European countries” (Reuters, April 4, 2023). As OpenAI met with Italian regulators, it was widely expected that a reactivation would occur, but no timetable was set (AP News “OpenAI to offer remedies to resolve Italy’s ChatGPT ban,” April 6).

On April 28, the service was reactivated after OpenAI introduced age verification to exclude users below 13 years of age and allow users to exercise their right to object to the training of models with their personal data (“ChatGPT: OpenAI reopens the platform in Italy, guaranteeing more transparency and more rights for European users and non-users,” press release, April 28, 2023). However, while the service has been uninterrupted since its reactivation, there are still concerns about the long-term regulation of AI in the European Union: the rival service Google Bard was unavailable in Europe as of May 31 2023 (it launched on July 13) and EU regulators have been revising laws to regulate AI (“EU lawmakers vote for tougher AI rules as draft moves to final stage”, Reuters, June 14, 2023).

2.2. Data Sources and Variable Construction

Stock returns, trading and ownership information of companies in the European Union and the United Kingdom are obtained from the Refinitiv Eikon and Datastream database. Since the first decision was released on March 31, we define the ban period from March 31 to April 28 (the restoration date). Table 1 presents variable definitions and summary statistics.⁹ A typical firm in

⁹ We check the accuracy of the stock returns data obtained from Refinitiv by comparing it with Yahoo Finance’s stock prices. Specifically, we randomly select 200 European securities in our sample. We download the data on adjusted closing prices from Yahoo Finance and calculate the daily stock returns. The absolute value of stock return difference between Refinitiv and Yahoo Finance is 0 for over 90% of the sample and is less than 1.3% for 98% of the sample. After checking the data quality, we follow Ince and Porter (2006) and Amihud, Hameed, Kang, and Zhang (2015) to

our sample has a median monthly return of 0.9%, a book-to-market ratio of 0.60, a market capitalization of EUR 315 million, and a ROA of 0.05. As shown in Table A1 Panel A, the average market capitalization is EUR 2.26 billion for Italian stocks and EUR 1.97 billion for matched Euro stocks, and the average ROA is 0.04 for both Italian and matched firms. There is no statistical difference between Italian and matched firms in terms of market capitalization and ROA.

III. The Impact of ChatGPT Ban on Firm Value

3.1. The Heterogeneous Impacts of ChatGPT Ban Across Firms

We assess the value of ChatGPT by estimating changes in the relative stock market valuation of firms with different exposure. The impact of the technology on different business models varies widely. For example, adopting generative AI in information and professional services industries may provide more benefits, including increased efficiency, reduced costs, and higher customer satisfaction (Eisfeldt et al., 2023). These benefits can lead to creative destruction, where less productive processes are phased out and replaced with more productive ones (Caballero et al., 1998, 2001, 2005; Acemoglu and Restrepo, 2018). This dynamic restructuring drives productivity gains and contributes to overall growth and competitiveness. We first focus on stock returns because it is a forward-looking measure. In this way, we can capture market perceptions about the current and future evolution of the technology.

An important first step is to identify the generative AI exposure of each industry. Eisfeldt et al. (2023) first measure each occupation's exposure to ChatGPT based on the occupation's task descriptions. Then, they aggregate exposures to the firm level using occupation employment data. Along the same lines, Eisfeldt et al. (2023) further develop industry-level exposure to ChatGPT

address return outliers. We set daily returns as missing if they are above 200% or if $(1 + r_{i,t}) \times (1 + r_{i,t+1}) - 1 \leq 50\%$ and either $r_{i,t}$ or $r_{i,t+1}$ is above 100%, where $r_{i,t}$ is the return of stock i on day t .

for 2-digit NAICS industries. We follow their approach of dividing the sample based on this task-based ChatGPT exposure. Specifically, we obtain Industry-level (2-digit NAICS) exposure to generative AI from Eisfeldt et al. (2023). The industries with the highest exposure are Finance and Insurance (52), Professional, scientific, and technical services (54), Management of companies and enterprises (55), and Information (51). There are no Italian firms in our sample in the Management of companies and enterprises (55). We exclude the Finance and Insurance industry for our baseline analysis due to its unique characteristics and the possible influence of overall instability in the EU financial market in early 2023.¹⁰

Following Eisfeldt et al. (2023), we form value-weighted portfolios with high exposure (i.e., firms from Professional, scientific, and technical services (54) and Information (51)) versus low exposure (all other industries except Finance and Insurance). We plot the cumulative return over the ban period for high and low-exposure portfolios as well as for a hedge portfolio that is long high exposure and short low exposure. Figure 1 shows that high exposure industries underperform low exposure industries during the ban. The long-short portfolio has a negative cumulative return of -1.5% over three days after the announcement and -5.48% over the entire ban period, suggesting a drift after the announcement as more investors price in the effect of the ban. Our results suggest that industries that potentially benefit the most from adoption are harmed the most after the ban is announced.

¹⁰ <https://www.bloomberg.com/news/articles/2023-03-15/credit-suisse-is-fueling-a-broader-rout-in-european-bank-stocks>.

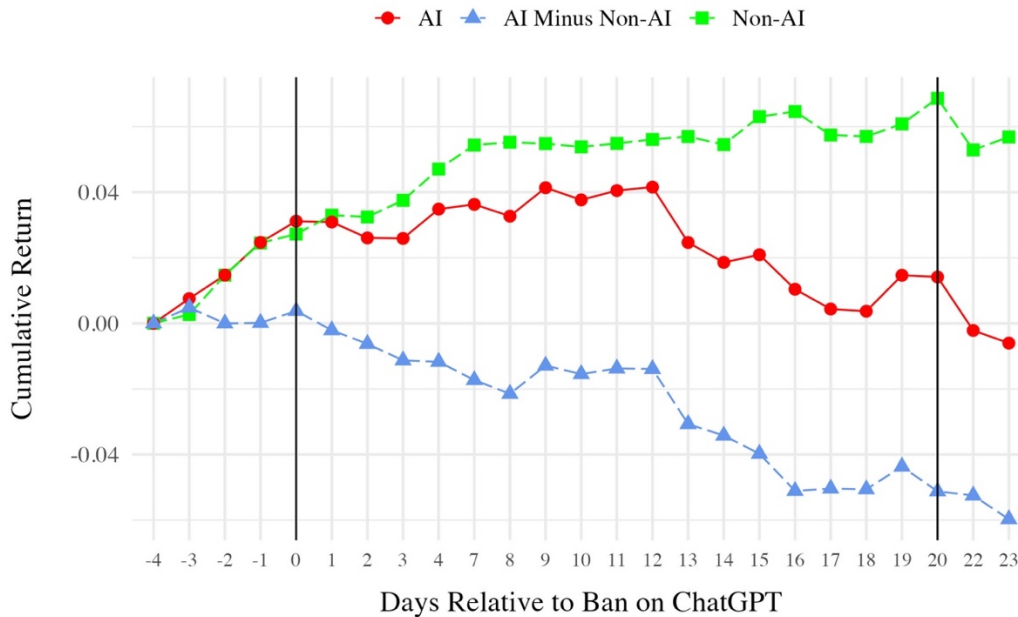


Figure 1: Value-weighted cumulative return for portfolios of Italy stocks in high generative AI exposure and low exposure industries, and a long-short portfolio that longs high exposure and short low exposure stocks. Day 0 (20) is the day of the ban (restoration) of ChatGPT in Italy.

One identification concern is that the ban may primarily reflect growing public concern within the European Union (EU) about data privacy issues associated all AI-related industries. From this perspective, the decline in market value observed in exposed industries could be driven by the public’s apprehension over privacy and AI, rather than the inaccessibility of the productivity gain brought by the technology. To alleviate the concern, we employ propensity score matching to construct a set of control firms.¹¹ Each Italian stock is matched to a non-Italian stock based on industry, classified by a 2-digit NAICS code, market capitalization, and profitability measured by return on assets (ROA) without replacement. The market capitalization for each firm is calculated as of March 24, 2023 (i.e., one week before the ban). After matching, Italian and matched firms

¹¹ Our results are robust to using different matching methods such as matching each Italian firm to three control firms, matching based on Mahalanobis distance matrix, or matching using Coarsened Exact Matching (CEM). Details of alternative matching methods are reported in Table A1.

have insignificant differences across market cap and ROA. We also require firms to have non-missing value on returns from March 24 to March 31.

For our primary analysis, we use a control sample of firms in the European Union (excluding Italy) and the United Kingdom, resulting in a sample of 228 Italian firms and 228 matched European firms.¹² It is conceivable that investors may have anticipated a ban in other European countries, as evidenced by announcements of investigations in France (April 11), Spain (April 13), and Germany (April 24), muting the documented effect if control firms are also treated. To address this concern, we show in Table A6 that our results are similar using a matched control sample with only UK firms.

Based on the matched sample, we test whether the ban affects firms with different levels of exposure by running a multivariate regression to control for other confounding factors. We form a sample of Italian firms and their matched counterparts and examine the effect of exposure by estimating the following equation:¹³

$$Return_{i,t} = \beta_1 Italy_i \times D_i + \beta_2 D_i + \beta_3 Italy_i + X_{i,t-1} + Industry FE + \epsilon_{i,t}, \quad (1)$$

where *Italy* equals one for all Italian stocks and zero for the matched non-Italian stocks. D_i is a dummy variable that equals one for either high exposure firms, new firms, or small firms. We use three definitions of exposure. First, as mentioned above, firms are classified as high exposure if they belong to the professional, scientific, and technical services industry and the information industry (Eloundou et al., 2023; Eisfeldt et al., 2023).

¹² The total market capitalization of Italian firms in our sample is 514.8 billion EUR, approximately 73% of the total Italian market capitalization in February 2023 estimated by CEIC (Source: <https://www.ceicdata.com/en/italy/market-capitalization/market-capitalization-total>).

¹³ All the regressions using stock return as the dependent variable are estimated by weighted least squares where the weight is the market cap of the stock on March 24, 2023 (one week before the announcement of the ban). In doing so, the regression estimates resemble the value-weighted portfolio returns and better capture the economic significance.

Second, we alternatively define high exposure firms if their shareholders' attention to ChatGPT is above the median value (High AI Att). This definition aims to capture investors' perceptions of the importance of AI to firm value. To construct firm-level AI attention, we download from Refinitiv a country-level breakdown of the shareholder base for all firms in our sample. We then determine investor attention to ChatGPT based on a country-level Google search volume (GSI) of AI/ChatGPT from December 1, 2022 (i.e., after the release of ChatGPT) to March 24, 2023 (i.e., one week before the ban). As mentioned in Google Trends and other studies using GSI (Eichenauer, et al., 2022; Carrière, et al., 2013), Google employs a sampling procedure that introduces sampling variation, affecting the results. In order to address this issue, we draw 50 random samples of the GSI data and average the outputs. After obtaining the average country-level attention, we calculate the firm-level ex-ante investor attention by taking the country-level Google search index average weighted by the percentage of shares held by investors from each country.¹⁴

We examine whether ChatGPT, a low-cost technology, is a force for creative destruction and levels the playing field for smaller and younger firms by augmenting their capabilities and narrowing the competitive gap with their larger and more established counterparts. If ChatGPT accelerates creative destruction, the impact of the ban will be more pronounced for younger firms. These firms are disproportionately affected as they are hindered from using ChatGPT to close the resource gap between themselves and incumbents. We define younger firms as those whose age is below the median within the industry (New).¹⁵

Finally, we classify firms as small if their market capitalization is below the median within the industry on March 24 (i.e., one week before the ban), and large otherwise (Small Mkt.). $X_{i,t-1}$

¹⁴ We require the sum of country-level ownership to be above 80% to ensure reasonable coverage of the country-level ownership data and below 110% to eliminate obvious errors.

¹⁵ We calculate firms' age by taking the difference between 2023 and each firm's founding year. We manually collect the data on firms' founding years from various public sources.

consists of the lagged book-to-market ratio (Book-to-Market), idiosyncratic volatility (Idio Vol), market capitalization (Market Cap), and stock return over the last month (Lag Return). We control for industry fixed effects to account for time-invariant industry-level pricing effects, and robust standard errors are reported in parentheses.

Table 2 shows that firms with higher exposure to generative AI perform worse under all three definitions of exposure during the ban. In particular, the estimated β_1 in Column 1 indicates that Italian AI stocks experience an average negative return of -8.7% relative to their matched European counterparts. Consistent with the finding in Eisfeldt et al. (2023) that AI industries benefit from the introduction of ChatGPT, the estimated β_2 in Column 1 implies that European AI stocks outperform non-AI stocks by 4.7%.¹⁶

The incremental decline in firm value ranges from -6.8% to -9.2%, across specifications. This result is consistent with investors perceiving the ban as a more significant reduction in productivity for firms with higher exposure. Young (small) firms underperform old (large) firms by -6.8% (-7.1%) during the ban. This finding suggests that the ban benefits incumbent firms, which tend to be older and larger, by hindering competitors' access to productivity tools.

3.2. Alternative AI Exposure Measure using Corporate Disclosure

Corporate disclosure plays an essential role in revealing individual firms' AI exposure. We follow Acemoglu et al. (2022) to develop the list of AI-related keywords (see Tables A7 and A8 for more details) and analyze Italian firms' 2022 annual reports to identify firms whose number of mentions of AI-related keywords is above the median. In Table 3, we find consistent results, suggesting that firms with a higher number of AI-related keywords experience a price decline

¹⁶ Furthermore, we re-run Column 1 using four other matching methods and the results are consistent (see Figure A4).

during the ban period.

3.3. Additional Analysis and A Placebo Test

We investigate whether the effect of the ChatGPT on stock returns differs across firms with different investor bases. Using data on Italian firms' ownership from Refinitiv, Table A3 highlights significant differences in stock market reactions across different types of investors, consistent with our findings that market-wide returns remain largely unaffected. First, Italian firms with higher government (domestic/insider) ownership show a positive (negative) return. Second, we do not observe significant return patterns related to retail or institutional ownership.

A legitimate concern in interpreting the cross-sectional tests is that firms with higher exposure may be more exposed to market risk. Consequently, the underperformance could be due to their market risk exposure rather than the ban. Unfortunately, because the ban affected the entire Italian market, it is not possible to adjust returns for the market return of Italian firms without eliminating the effect of the ban. To address this concern, we show in Figure A3 and Table A2 that the underperformance of AI industries in Italy is qualitatively similar when returns are adjusted for market performance in other EU countries excluding Italy. Second, in Table A4, we conduct a placebo test using Moody's downgrade of Italy's sovereign rating outlook from stable to negative on August 6, 2022. If our results are driven by AI industries reacting more strongly to a country-level market shock, we expect exposed industries to react to the downgrade. However, we do not find underperformance for high-exposure firms relative to low-exposure firms.

3.4. The Impact of ChatGPT Ban on Market-wide Valuations

Last, we examine the impact of the ChatGPT ban on market-wide stock reactions in Italy. The market-wide stock response test is a joint test of a) whether investors paid attention to the

announcement of the ban and its restoration, b) whether investors perceive ChatGPT as a productive tool, c) whether investors agree on the effect of ChatGPT adoption on firm value. Even if all three conditions are met, we still might not see a significant market-wide price reaction because ChatGPT could benefit a subset of firms while harming the value of other firms. For example, we find that Italian firms with higher AI exposure perform worse than firms with lower AI exposure. The heterogeneous effects of the ChatGPT ban across firms could cancel out and lead to an insignificant market-wide response.

Table B1 and Figure B1 document an insignificant reaction of the Italian stock market to the ChatGPT ban for different time windows after the announcement date. There is no significant market-wide response from the ChatGPT restoration date to one day after (i.e., [+20, +21]) and three days after (i.e., [+20, +23]). We further show that this result is robust to a) comparing the difference between Italian stock returns and the matched European stock returns during the ban (Figure B2), b) using five different matched samples (Figure B3), and c) testing the return difference in a multivariate regression that controls for other risk factors and industry fixed effects (Figure B4).

IV. The Impact of the ChatGPT Ban on the Information Environment

In this section, we first examine the effect of the ban on the information environment. First, we examine analysts' forecasting behavior to assess how the ban influences information production. Second, we use the bid-ask spread to assess whether the ban constrains investors' information processing capabilities, increasing information asymmetry. Third, to further explore the impact of the ban on investors' response to public information, we study the market reaction to earnings announcements during the ban.

4.1. The Impact of the ChatGPT Ban on Analysts' Information Production

We investigate whether the ban affects financial analysts' information production. On the one hand, prior research shows that generative AI has the potential to assist financial market participants in information processing and content creation (e.g., Li et al., 2023). On the other hand, the impact of the ban on analysts may be negligible for several reasons. First, the treatment effect depends on how widely ChatGPT was adopted by analysts before the ban. Second, analysts may have more resources (e.g., VPN) to circumvent the ban. Therefore, our results on financial analysts may provide a lower bound of the technology's impact on information production.

The first empirical challenge arises from the direct impact of the ban on economic fundamentals, as discussed in our earlier section (Table 2). Consequently, the change in analyst behavior results from changes in information processing costs and firm fundamentals. To mitigate this problem, within the group of analysts following a particular Italian firm, we compare analysts based in Italy with their counterparts in other countries. This research design allows us to control for changes in firm fundamentals, ensuring that differences in forecast frequency can be attributed to the ban. We determine the location of analysts using the country code of their telephone number.

Data on analysts' names¹⁷ and their forecasts is obtained from the I/B/E/S international file. The forecasts are matched to Italian firms using stock ticker.¹⁸ This procedure yields a sample of 177 analysts covering 112 Italian firms, with one third of the analysts are likely to be based in Italy based on their office numbers' country code. We hypothesize that analysts in Italy are less inclined to issue forecasts or decrease the frequency of their forecasts compared to analysts in non-Italian

¹⁷ While the identity of analysts is masked in I/B/E/S Detail History file, I/B/E/S recommendation file provides the surname and the first letter of the given name of the analysts.

¹⁸ Given that tickers can be reused and may overlap with each other across different stock exchanges, we manually check the matching accuracy and delete incorrect matches.

countries following the same firm. To test this hypothesis, we estimate the following specification:¹⁹

$$\begin{aligned} Forecast_{i,g,t} = & \beta_1 Italy Analyst_{i,g} \times Ban_t + \beta_2 Italy Analyst_{i,g} + \\ & Firm \times Month FE + \epsilon_{i,g,t}, \end{aligned} \quad (3)$$

where $Forecast_{i,g,t}$ is the forecast outcomes of analyst group g covering firm i in month t . We focus on three different measures of forecast propensity at the firm-group-month level and include: a) a dummy variable set to 1 if at least one forecast is issued by group g covering firm i in month t , and 0 otherwise (forecast likelihood); b) the number of forecasts by each group g covering firm i in month t (forecast frequency); c) the number of forecasts scaled by the pre-ban analyst coverage for each group g covering firm i in month t (forecast frequency per analyst).²⁰ The first two outcomes shed light on the extensive and intensive margin of the impact of the ChatGPT ban. The third measure captures the average number of forecasts per analyst. Ban is a time dummy that equals one for the ban period and zero for the pre-periods, i.e., January, February, and March (before March 31). Our data's three-dimensional (firm-group-month level) nature allows us to control for firm-by-month fixed effects.

The results are shown in Table 4. We evaluate the extensive margin of analyst forecast (forecast likelihood) in Column 1 of Table 4. The forecast likelihood decreases by 10.9% for the group of analysts located in Italy compared to analysts in other countries following the same Italian firm during the ban period. As for the intensive margin, the number of forecasts and forecasts per analyst decrease by 0.29 and 0.14 (6.97% and 10.31% of the sample standard deviation) for

¹⁹ The baseline regressions are estimated using OLS. In untabulated supplementary analyses, we use Poisson regressions either with or without fixed effects to avoid biased estimates and find similar results.

²⁰ There are one treatment group and one control group for each firm – analysts in Italy and analysts in other countries. In addition to analyzing analyst forecasts at firm-group-month level, we perform a robustness check at firm-analyst-month level and find qualitatively similar results.

analysts in Italy relative to their non-Italian counterparts, consistent with ChatGPT assisting analysts in their day-to-day operations. Note that we cannot identify from these tests how analysts use the technology because we only observe the occurrence of a forecast; in particular, it could be that ChatGPT helps analysts process information, or assists them to better explain and present their forecast. Both channels are consistent with a reduction in the information available to investors.

4.2. The Impact of the ChatGPT Ban on Bid Ask Spread

We hypothesize that generative AI can help investors process and summarize unstructured information (e.g., Kim et al., 2023), leading to better information flow, more trading, and lower information asymmetry. To test this, we employ a difference-in-differences design by estimating the following equation:

$$Bid - ask_{i,t} = \beta_1 Italy_i \times Ban_t + \beta_2 Italy_i + X_{i,t-1} + Time/Industry FE + \epsilon_{i,t}, \quad (2)$$

where Ban is a time dummy that equal to one for the ban period and zero for the pre-periods, i.e., January, February, and March (before March 31). We mainly focus on the bid-ask spread and the bid-ask spread adjusted for trading volume. Table 5 shows that both the raw (column 1) and adjusted (column 2) bid-ask spreads increase significantly by about 0.03, which is about 1.5 times the standard deviation.²¹ This result is consistent with the ban hindering information processing and increasing information asymmetry across investors.

We also conduct several cross-sectional tests on the adjusted bid-ask spreads (Table 6) to shed light on the mechanism driving the increased information asymmetry.²² Firms with limited analyst coverage are more prone to information processing problems and experience decreased liquidity

²¹ We show in Table A5 the results on stock return total volatility, idiosyncratic volatility, and turnover. Turnover declines significantly during the ban period for Italian stocks, consistent with our hypothesis.

²² Our cross-sectional results remain consistent using unadjusted bid-ask spreads.

during the ban. Firms with higher foreign ownership (column 2) exhibit a smaller increase in information asymmetry because the ban only applies to Italian geographic boundaries and does not directly affect foreign investors. Consistent with the notion that smaller and unsophisticated investors rely more on ChatGPT to parse through financial information than institutional investors, firms with higher institutional ownership (High Institution, column 3) exhibit a smaller increase in bid-ask spreads. Overall, our results point to ChatGPT's role as a productivity tool that facilitates investors' information processing.

4.3. The Impact of ChatGPT Ban on Market Reaction to Earnings Announcements

The ChatGPT ban may reduce or amplify the stock market reaction to earnings announcements. On the one hand, if the ban hampers investors' information processing capacity, they may struggle to digest earnings information efficiently during the ban period. This reduced capacity could lead to delayed or less accurate assessments of earnings reports, resulting in a delayed market response.

On the other hand, the ChatGPT ban may increase the reaction because the news is not efficiently integrated into stock prices before the earnings announcement. To illustrate this, consider the price movement of a hypothetical stock with and without prior information acquisition before an earnings announcement or another scheduled public news event. In a market where traders are able to collect private information, part of the information is acquired and integrated into the price before its public disclosure (e.g., Kyle 1985, Back 1992), dampening the market reaction to the earnings information. In contrast, without informed traders, no market participant is aware of the information content of the upcoming public disclosure. Therefore, the immediate response of stock prices to earnings announcements is substantial. Another possible channel is that the ban results in fewer analyst forecasts before the earnings announcements, as we document

earlier. Consequently, there is a higher inflow of information during earnings-announcement windows, potentially resulting in more pronounced market reactions (Ball and Shivakumar 2008).

In summary, the ban's impact on reactions to earnings announcements is ex-ante unclear, and we test these two opposing hypotheses empirically. Following Livnat and Mendenhall (2006) and Mian and Sankaraguruswamy (2012), we calculate the unexpected earning (UE) which represents the unexpected part of the earnings announcement, as follows:

$$UE_{i,t} = \frac{Earnings_{i,t} - Earnings_{i,t-1}}{P_{i,t}}, \quad (4)$$

where $Earnings_{i,t}$ is the earnings per share of firm i in year t and $P_{i,t}$ the stock price two days before the earnings announcement date in year t to avoid a contamination effect of announcements on prices. We use prior year earnings ($Earnings_{i,t-1}$) as investors' expected earnings for year t . We also use the median value of analyst forecasts as the market's expected earnings and find similar results.²³ UE is winsorized at 1 percent and 99 percent to mitigate the effect of outliers.²⁴

We measure the stock market's reaction to earnings announcements by regressing stock returns on unexpected earnings. Specifically, we estimate the following regression:

$$Ret_{i,t} = \beta_1 Italy_i \times Ban_t \times UE_t + \beta_2 Italy_i \times Ban_t + \beta_3 Italy_i \times UE_t + \beta_4 Ban_t \times UE_t + \beta_5 UE_t + \beta_6 Italy_i + Time/Industry FE + \epsilon_{i,t}, \quad (5)$$

where $Ret_{i,t}$ is the cumulative abnormal return for firm i in three different time windows -three, five, or seven days centered around the earnings announcement dates.²⁵ Ban is a time dummy that equals one for the ban period and zero for the pre-periods, i.e., January, February, and March

²³ In untabulated results, we construct an alternative measure of UE by replacing $Earnings_{i,t-1}$ with a measure of analysts' expectation. We calculate analysts' expectations as the median of 1-year ahead forecasts in the 90 days prior to the earnings announcements. The results are qualitatively similar.

²⁴ We focus on annual earnings announcements because the data on quarterly earnings announcements is scarce for Italian firms.

²⁵ We take the difference between daily stock return and market return as daily abnormal return. Our results are robust to alternative measures of abnormal returns adjusted based on the CAPM model.

(before March 31).

Table 7 shows that the market reaction to earnings announcements is stronger for Italian firms than for comparable non-Italian firms during the ban. With a one standard deviation (0.13) increase in unexpected earnings, Italian firms experienced an additional 5.42 (0.13×0.417) basis point increase in stock prices relative to their peers during a three-day window around the earnings announcement. The additional increase is 11.71 and 12.26 basis points for the five-day and seven-day time windows around the earnings announcement date, respectively. As hypothesized above, the ban may reduce the market reaction by slowing investors' processing of earnings news, or it may increase the reaction by preventing the news from being incorporated into prices before the earnings announcement. Our result suggests a net effect in favor of the second hypothesis: The ban negatively affects investors' information processing prior to the earnings announcement when the information is publicly disclosed.

V. Conclusion

Despite the overwhelmingly positive response to the availability of AI tools such as OpenAI's ChatGPT, a comprehensive assessment of the benefits of AI remains unclear. This study aims to contribute to this question by examining the capital market's response to the ban of ChatGPT in Italy, which created a unique environment in which access to the technology was hindered for a wide range of companies. The impact of the ban affected industries that were more affected by AI, as well as smaller and younger companies. These results are consistent with the redistributive effects of generative AI through creative destruction, which challenges established firms and benefits smaller, younger, and more technology-oriented firms. In addition, our results suggest that AI may favor trading, as evidenced by an increased bid-ask spread during the ban, especially for firms with fewer institutional investors, limited analyst coverage, and a lower presence of foreign

investors. Our analyses on financial analysts' behavior further suggest benefits of AI in facilitating information processing.

Our results are also subject to the limitations inherent to the nature of the event. Our results focus primarily on Italian listed companies, and caution should be exercised in generalizing these results to private companies or other countries. Moreover, the empirical framework is based on investors' beliefs, which may not always be consistent with actual outcomes. While our results suggest persistent underperformance of Italian exposed industries, it is important to acknowledge that evaluating the long-term impact of the ban on the economy necessitates a broader time frame and further evidence from post-ban outcomes.

Our results are also subject to caveats. First, Italian companies could circumvent the ban by using proxy servers with foreign IP addresses. As far as we know, OpenAI does not have a policy to block IP addresses used by proxy servers, and sophisticated users such as coding engineers, who likely benefited the most from ChatGPT's coding capabilities, use proxy servers more extensively and recovered quickly within a few days of the ban (Kreitmeir and Raschky, 2023). In a counterfactual where the ban had been prolonged and extended to other countries, it is possible that a greater share of the user base would have switched to such services. ChatGPT users may also have switched to alternative services during the ban, especially considering that Google's Bard was launched shortly before the March 21 ban. However, Bard was initially limited to select users on a waitlist. It was made available to the general public on May 10 and launched in the European Union on July 13, after the ban.

Second, if the market had anticipated a short-lived ban, we would underestimate the value destruction of a long-run ban. Therefore, our results can only be interpreted to speak about the potential long-term costs of regulations causing limited access to generative AI, rather than actual damages suffered by the Italian economy due to the ban. We also note that our results cannot

measure welfare and are intended only as inputs to policies on generative AI. Market values do not capture all the social value created by private firms or the value appropriated by other stakeholders, such as consumers, governments, or workers. In addition, there may be trade-offs between privacy and economic efficiency that governments are currently trying to address when evaluating approaches to regulating AI.

Finally, other interesting events may provide complementary evidence about the value of the technology. For example, OpenAI did not permit individuals to register in China, although it was not fully blocked and was compatible with Mandarin. However, in February 2023, Chinese authorities took steps to restrict access by ordering social media platforms to remove workarounds and enforcing stricter controls on VPN access. The banning of ChatGPT in China is a significant event that could further expand our understanding of the benefits of AI.²⁶

²⁶ The ban in China poses its own empirical challenges and would likely require a study exploiting the institutional features of this economy. First, since ChatGPT was never officially allowed in China, only a select group of users who used workarounds could access the service. These users faced the risk of enforcement actions, which likely limited the widespread adoption of the technology within the country. Second, the ban of ChatGPT in China was not a singular event with a specific date. Instead, it occurred gradually through various informal channels that closed down workarounds. This decentralized approach makes it difficult to pinpoint a precise moment when the ban took effect.

References

- Acemoglu, Daron, Tarek A. Hassan, and Ahmed Tahoun. 2018. "The Power of the Street: Evidence from Egypt's Arab Spring." *Review of Financial Studies* 31 (1): 1–42.
- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488-1542.
- Amihud, Y., Hameed, A., Kang, W., & Zhang H. (2015). The illiquidity premium: International evidence. *Journal of Financial Economics*, 117(2), 350–368.
- Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2023). Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics*, Forthcoming.
- Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2022). Firm investments in artificial intelligence technologies and changes in workforce composition. *Available at SSRN 4060233*.
- Back, K. (1992). Insider Trading in Continuous Time. *The Review of Financial Studies*, 5(3), 387–409.
- Bai, J. J., Boyson, N. M., Cao, Y., Liu, M., & Wan, C. (2023). Executives vs. Chatbots: Unmasking Insights through Human-AI Differences in Earnings Conference Q&A.
- Bernard, D., Blankespoor, E., de Kok, T., & Toynbee, S. (2023). Confused readers: A modular measure of business complexity. *Available at SSRN 4480309*.
- Blankespoor, E., Dehaan, E., Wertz, J., & Zhu, C. (2019). Why do individual investors disregard accounting information? The roles of information awareness and acquisition costs. *Journal of Accounting Research*, 57(1), 53-84.
- 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), 101-344.
- Bloom, N., Hassan, T. A., Kalyani, A., Lerner, J., & Tahoun, A. (2021). The diffusion of disruptive technologies (No. w28999). *National Bureau of Economic Research*.
- Bybee, L. (2023). Surveying Generative AI's Economic Expectations. *arXiv preprint arXiv:2305.02823*.
- Caballero, Ricardo J. & Mohamad L. Hammour. (1998). The macroeconomics of specificity. *Journal of Political Economy*, 106(4), 724-767.
- Caballero, Ricardo J. & Mohamad L. Hammour. (2001). Creative destruction and development: institutions, crises, and restructuring. *Annual World Bank Conference on Development Economics 2000*, 213-241.
- Caballero, Ricardo J. & Mohamad L. Hammour. (2005). The cost of recessions revisited: A reverse liquidationist view. *Review of Economic Studies*, 72, 313-341.
- Caballero, Ricardo J., Takeo Hoshi, and Anil K. Kashyap. (2008). Zombie lending and depressed restructuring in Japan. *American Economic Review*, 98 (5): 1943-77.
- Carrière-Swallow Y, Labbé F. Nowcasting with Google Trends in an emerging market[J]. *Journal of Forecasting*, 2013, 32(4): 289-298.
- Chen, D., Ma, Y., Martin, X., & Michaely, R. (2022). On the fast track: Information acquisition costs and information production. *Journal of Financial Economics*, 143(2), 794-823.
- Chen, W., & Srinivasan, S. (2023). Going digital: Implications for firm value and performance. *Review of Accounting Studies*, 1-47.
- Dittmar, J. E. (2011). Information technology and economic change: the impact of the printing press. *The Quarterly Journal of Economics*, 126(3), 1133-1172.

- de Kok, T. (2023). Generative LLMs and Textual Analysis in Accounting:(Chat) GPT as Research Assistant? *Available at SSRN*.
- Dong, Y., Li, O. Z., Lin, Y., & Ni, C. (2016). Does information-processing cost affect firm-specific information acquisition? Evidence from XBRL adoption. *Journal of Financial and Quantitative Analysis*, 51(2), 435-462.
- Eichenauer V Z, Indergand R, Martínez I Z, et al. Obtaining consistent time series from Google Trends[J]. *Economic Inquiry*, 2022, 60(2): 694-705.
- Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). *Generative AI and Firm Values* (No. w31222). National Bureau of Economic Research.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*.
- Emett, S. A., Eulerich, M., Lipinski, E., Prien, N., & Wood, D. A. (2023). Leveraging ChatGPT for Enhancing the Internal Audit Process—A Real-World Example from a Large Multinational Company. *Available at SSRN 4514238*.
- Eulerich, M., Sanatzadeh, A., Vakilzadeh, H., & Wood, D. A. (2023). Can Artificial Intelligence Pass Accounting Certification Exams? ChatGPT: CPA, CMA, CIA, and EA?. *ChatGPT: CPA, CMA, CIA, and EA*.
- Eulerich, M., & Wood, D. A. (2023). A Demonstration of How ChatGPT Can be Used in the Internal Auditing Process. *Available at SSRN 4519583*.
- Fang, B., Huang, S., Roychowdhury, S., & Sletten, E. Mobile Internet and Analyst Forecast Performance.
- Farboodi, Maryam, Roxana Mihet, Thomas Philippon, and Laura Veldkamp. 2019. “Big Data and Firm Dynamics.” *AEA Papers and Proceedings*, 109: 38-42.
- Felten, E. W., Raj, M., & Seamans, R. (2018, May). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings* (Vol. 108, pp. 54-57).
- Felten, E., Raj, M., & Seamans, R. (2023). How will Language Modelers like ChatGPT Affect Occupations and Industries?. *arXiv preprint arXiv:2303.01157*.
- Frankel, R., Jennings, J., & Lee, J. (2022). Disclosure sentiment: Machine learning vs. dictionary methods. *Management Science*, 68(7), 5514-5532.
- Gu, H., Schreyer, M., Moffitt, K., & Vasarhelyi, M. A. (2023). Artificial Intelligence Co-Piloted Auditing. *Available at SSRN 4444763*.
- Jha, M., Qian, J., Weber, M., & Yang, B. (2023). ChatGPT and corporate policies. *Chicago Booth Research Paper*, (23-15).
- Katz, D. M., Bommarito, M. J., Gao, S., & Arredondo, P. (2023). Gpt-4 passes the bar exam. *Available at SSRN 4389233*.
- Kim, A. G., Muhn, M., & Nikolaev, V. V. (2023). Bloated disclosures: Can ChatGPT help investors process information? *Chicago Booth Research Paper*, (23-07).
- Kogan, L., & Papanikolaou, D. (2014). Growth opportunities, technology shocks, and asset prices. *Journal of Finance*, 69(2), 675-718.
- Kogan, L., Papanikolaou, D., Schmidt, L. D., & Seegmiller, B. (2021). Technology, vintage-specific human capital, and labor displacement: Evidence from linking patents with occupations (No. w29552). *National Bureau of Economic Research*.
- Kogan, L., Papanikolaou, D., Schmidt, L. D., & Song, J. (2020). Technological innovation and labor income risk (No. w26964). *National Bureau of Economic Research*.
- Korinek, A. (2023). Language models and cognitive automation for economic research (No. w30957). *National Bureau of Economic Research*.

- Kreitmeir, D. H., & Raschky, P. A. (2023). The Unintended Consequences of Censoring Digital Technology--Evidence from Italy's ChatGPT Ban. arXiv preprint arXiv:2304.09339.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1336.
- Li, E. X., Tu, Z., & Zhou, D. (2023). The Promise and Peril of Generative AI: Evidence from ChatGPT as Sell-Side Analysts. Available at SSRN 4480947.
- Li, F. (2008). Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics*, 45(2-3), 221-247.
- Lin, C., Ma, C., Sun, Y., & Xu, Y. (2021). The telegraph and modern banking development, 1881–1936. *Journal of Financial Economics*, 141(2), 730-749.
- Livnat, J., & Mendenhall, R. R. (2006). Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts. *Journal of Accounting Research*, 44(1), 177–205.
- Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? return predictability and large language models. arXiv preprint arXiv:2304.07619.
- Meindl, B., Frank, M. R., & Mendonça, J. (2021). Exposure of occupations to technologies of the fourth industrial revolution. arXiv preprint arXiv:2110.13317.
- Mujtaba Mian, G., & Sankaraguruswamy, S. (2012). Investor sentiment and stock market response to earnings news. *Accounting Review*, 87(4), 1357–1384.
- Ince, O.S., & Porter, R.B. (2006). Individual Equity Return Data From Thomson Datastream: Handle With Care! *Journal of Financial Research*, 29, 463-479.
- OpenAI. "GPT-4 technical report." arXiv (2023).
- O'Leary, D. E. (2023a). An analysis of Watson vs. BARD vs. ChatGPT: The Jeopardy! Challenge. *AI Magazine*.
- O'Leary, D. E. (2023b). An analysis of three chatbots: BlenderBot, ChatGPT and LaMDA. *Intelligent Systems in Accounting, Finance and Management*, 30(1), 41-54.
- Papanikolaou, D. (2011). Investment shocks and asset prices. *Journal of Political Economy*, 119(4), 639-685.
- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8), 9
- Ball, R., & Shivakumar, L. (2008). How much new information is there in earnings? *Journal of Accounting Research*, 46(5), 975–1016.
- Restrepo, Felipe, Lina Cardona-Sosa, and Philip E. Strahan. (2019). Funding Liquidity without Banks: Evidence from a Shock to the Cost of Very Short-Term Debt. *Journal of Finance* 74 (6): 2875–2914.
- Schumpeter, J. A. (1942). *Capitalism, Socialism, and Democracy*. New York: Harper & Brothers.
- Spitale, Giovanni, Nikola Biller-Andorno, and Federico Germani. "AI model GPT-3 (dis) informs us better than humans." arXiv preprint arXiv:2301.11924 (2023).
- Veldkamp, Laura L. 2006. "Information Markets and the Comovement of Asset Prices." *Review of Economic Studies* 73(3): 823–45.
- Verrecchia, R. E. (1982). Information acquisition in a noisy rational expectations economy. *Econometrica: Journal of the Econometric Society*, 1415-1430.
- Webb, M. (2019). The impact of artificial intelligence on the labor market. Available at SSRN 3482150.

- Wood, D. A., Achhpilia, M. P., Adams, M. T., Aghazadeh, S., Akinyele, K., Akpan, M., ... & Kuruppu, C. (2023). The ChatGPT Artificial Intelligence Chatbot: How Well Does It Answer Accounting Assessment Questions?. *Issues in Accounting Education*, 1-28.
- Yang, S. (2023). Predictive Patentomics: Forecasting Innovation Success and Valuation with ChatGPT. *Available at SSRN 4482536*.
- Zhang, M. B. (2019). Labor-technology substitution: Implications for asset pricing. *Journal of Finance*, 74(4), 1793-1839.

Variable Definitions

Variable Name	Definition
<i>Return</i>	The monthly buy-and-hold stock return.
<i>Market-Adj. Return</i>	<p>Mkt-Adj Ret is obtained by cumulating the daily abnormal return over a month. The daily abnormal return is adjusted based on the CAPM model:</p> $R_t^{mkt\ adj} = R_t - \beta_t R_{MKT-RF,t},$ <p>where β_t is estimated by running the following regression with daily returns ranging from [-252, -1]:</p> $R_t - R_t^f = \alpha_t + \beta_t R_{MKT-RF,t} + \varepsilon_t$ <p>where R_t^f is the risk-free rate and $R_{MKT-RF,t}$ is the value-weighted portfolio return of all stocks in EU plus UK excluding those in Italy in excess of the risk-free rate.</p>
<i>Volatility</i>	The monthly standard deviation of daily raw return.
<i>Idio Vol</i>	The volatility of daily abnormal return calculated based the CAPM model over a month
<i>Turnover</i>	The monthly average of the ratio of traded shares to the total shares outstanding, multiplied by 100.
<i>Bid-Ask</i>	The monthly average of daily bid-ask spreads, which is the bid price minus the ask price divided by the closing price.
<i>Bid-Ask Adj.</i>	The monthly average of daily bid-ask spreads weighted by daily trading volume.
<i>Book-to-Market Ratio</i>	The book value of equity per share divided by the stock price at the month's end.
<i>Market Cap</i>	The market capitalization in EUR billions at the end of a month.
<i>ROA</i>	Earnings before interests and taxes (EBIT) divided by total assets
<i>Forecast Likelihood</i>	A dummy variable set to 1 if at least one forecast is issued by analysts in Italy (analysts in other countries) and 0 otherwise.
<i>Forecast Frequency</i>	The number of forecasts issued by analysts in Italy (analysts in other countries).
<i>Forecast Frequency Per Analyst</i>	The number of forecasts issued by analysts in Italy (analysts in other countries) scaled by the ex-ante number of analysts in the pre-ban period
<i>Unexpected Earnings</i>	<p>Similar to Livnat and Mendenhall (2006) and Mian and Sankaraguruswamy (2012), we calculate the unexpected earning, UE, which represents the unexpected part of the earnings announcement, as follows:</p> $UE_{i,t} = \frac{Earnings_{i,t} - Earnings_{i,t-1}}{P_{i,t}},$ <p>where $Earnings_{i,t}$ is the earnings per share of firm i in year t and $P_{i,t}$ is the price per share two days before the announcement date of earnings in year t to avoid the contamination effects of announcements on prices. We winsorize the distribution of UE at 1 percent and 99 percent to avoid the impact of extreme values. We focus on annual earnings announcement due to data coverage issue. We also define UE as indicated by the above equation replacing with a measure of analysts' expectation based on analysts' forecasts. The measure of analysts' expectations is constructed as the median of 1-year ahead forecasts in the 90 days prior to the earnings announcements.</p>

Table 1: Summary Statistics

Variable	Obs.	Mean	St. Dev.	10%	25%	Median	75%	90%
Return	1824	0.020	0.113	-0.087	-0.038	0.009	0.064	0.132
Volatility	1824	2.169	1.349	1.085	1.387	1.844	2.477	3.439
Turnover	1814	19.322	25.014	2.290	5.074	11.351	23.625	41.269
Bid-Ask	1820	0.016	0.019	0.001	0.002	0.010	0.021	0.038
Bid-Ask Adj.	1810	0.016	0.020	0.001	0.002	0.009	0.020	0.040
Book-to-Market	1797	0.850	0.857	0.181	0.320	0.602	1.139	1.693
Idio Vol	1824	5.367	10.388	0.867	1.411	2.635	4.803	10.244
Market Cap	1824	2.019	4.776	0.019	0.081	0.315	1.577	5.305
ROA	1824	0.044	0.102	-0.011	0.016	0.051	0.088	0.127
UE	233	0.001	0.13	-0.079	-0.021	0.003	0.018	0.054
Analysts in Italy								
Forecast Likelihood	642	0.805	1.310	0.000	0.000	0.000	1.000	2.000
Forecast Frequency	642	0.408	0.492	0.000	0.000	0.000	1.000	1.000
Forecast Frequency Per Analyst	642	0.397	0.528	0.000	0.000	0.000	1.000	1.000
Analysts in Other Countries								
Forecast Likelihood	642	0.671	1.507	0.000	0.000	0.000	1.000	2.000
Forecast Frequency	642	0.259	0.438	0.000	0.000	0.000	1.000	1.000
Forecast Frequency Per Analyst	642	0.256	0.481	0.000	0.000	0.000	0.000	1.000

This table shows the summary statistics of the variables in the ban period (i.e., 21 active trading days from March 31st to April 28th). Return is the monthly buy-and-hold stock return. Volatility is the monthly standard deviation of daily raw return. Turnover is the monthly average of the ratio of traded shares to the total shares outstanding, multiplied by 100. Bid-Ask is the monthly average of daily bid-ask spreads, which is the bid price minus the ask price divided by the closing price. Bid-Ask Adj. is the monthly average of daily bid-ask spreads weighted by daily trading volume. Book-to-Market Ratio is the book value of equity per share divided by the stock price at the month's end. Idio Vol is the volatility of daily abnormal return calculated using the 1-factor market CAPM model over a month. Market Cap is the market capitalization in EUR billions at the end of a month. Return on assets (ROA) is calculated as earnings before interests and taxes (EBIT) divided by total assets. UE is unexpected earnings. Forecast likelihood is a dummy variable set to 1 if at least one forecast is issued by analysts in Italy (analysts in other countries) and 0 otherwise. Forecast frequency is the number of forecasts issued by analysts in Italy (analysts in other countries). Forecast frequency per analyst is the number of forecasts issued by analysts in Italy (analysts in other countries) scaled by the ex-ante number of analysts in the pre-ban period. To minimize the effects of outliers, we winsorize all continuous variables at the 1st and 99th percentiles except for stock return.

Table 2: The Heterogeneous Effects of the Ban on Stock Returns

Dependent variable	Ret	Ret	Ret	Ret	Ret	Ret
Dummy (D) =	ChatGPT- Industry	ChatGPT - Industry	High ChatGPT/AI Att.	High AI Mention	New	Small Mkt.
<i>Italy</i> × <i>D</i>	-0.087** (0.036)	-0.092*** (0.030)	-0.068*** (0.025)	-0.087** (0.040)	-0.068** (0.032)	-0.071** (0.029)
D	0.047* (0.026)		0.039* (0.022)	0.045 (0.029)	0.053* (0.030)	0.038 (0.024)
Italy	0.027* (0.016)	0.028* (0.017)	0.029 (0.018)	0.026 (0.020)	0.038** (0.017)	0.022 (0.017)
Lag B-to-M	0.0007 (0.011)	-0.004 (0.011)	-0.010 (0.011)	-0.011 (0.017)	-0.030** (0.012)	-0.008 (0.011)
Lag Idio Vol	0.0006 (0.0004)	0.0006 (0.0004)	0.0007* (0.0004)	0.0009 (0.0006)	0.0009** (0.0004)	0.0006 (0.0004)
Lag Market Cap	0.002** (0.0007)	0.001* (0.0007)	0.001* (0.0007)	0.001 (0.0009)	0.002*** (0.0007)	0.001* (0.0008)
Lag Return	0.042 (0.065)	0.037 (0.071)	0.014 (0.069)	0.061 (0.111)	-0.017 (0.070)	0.013 (0.070)
Industry FE.	No	Yes	Yes	Yes	Yes	Yes
Observations	389	389	387	94	313	389
R ²	0.17	0.23	0.22	0.25	0.30	0.20
Within R ²	--	0.13	0.12	0.15	0.20	0.10

This table reports the heterogeneous effects of the Ban on stock returns. We exploit five ex-ante firm-level measures: AI industry, attention to ChatGPT, number of AI topics mentioned in conference call transcripts, age, and firm size. Firms are classified as having high AI exposure if they are in the professional, scientific, and technical services sector and information sector and low AI exposure otherwise. High AI Att. equals one for firms with above the median level ex-ante investor attention and zero otherwise. We calculate the firm-level ex-ante investor attention by taking the average of the country-level Google search index weighted by the percentage of shares held by the investors from each country. We require the sum of country-level ownership to be above 80% and below 110%. Firms are defined as new if the founding year of the firm is more recent than the median founding year within the industry. We download the incorporation date of each firm from Refinitiv and replace the missing value with manually-collected founding years of Italian firms from Wikipedia. Firms are defined as small if their market capitalization on March 24 is below the median value within the industry and large otherwise. Italy is a dummy variable set to 1 for Italian stocks and 0 otherwise. We include industry fixed effects in Columns 2 to 6. We report standard errors robust to heteroskedasticity.

Dependent Var.=	Ret	Market-Adj Ret
<i>Italy × ChatGPT-Industry × High AI Topic</i>	-0.109* (0.056)	-0.146*** (0.053)
ChatGPT-Industry × High AI Topic	0.036 (0.047)	0.058 (0.041)
Italy × ChatGPT-Industry	-0.009 (0.045)	0.017 (0.039)
Italy × High AI Topic	0.031 (0.030)	0.058* (0.031)
High AI Topic	-0.013 (0.026)	-0.030 (0.025)
Italy	0.009 (0.020)	-0.011 (0.019)
Other risk factors	Yes	Yes
Industry FE.	Yes	Yes
Month FE.	Yes	Yes
Observations	291	291
R ²	0.25	0.26
Within R ²	0.16	0.16

This table reports the heterogeneous effects of the ChatGPT Ban on stock returns. We measure Italian firms' exposure to ChatGPT by their mentions of AI topics in 2022 annual reports. Firms are classified as having high AI exposure if their frequency of mentioning the AI-related keywords is above the median. Italy is a dummy variable set to 1 for Italian stocks and 0 otherwise. Standard errors are robust to heteroskedasticity.

Table 4: The Effects of the ChatGPT Ban on Analyst Forecasts

Dependent Variable	Forecast Likelihood	Forecast Frequency	Forecast Frequency Per Analyst
<i>Italy Analyst</i> × <i>Ban</i>	-0.109** (0.050)	-0.290** (0.108)	-0.141*** (0.024)
Italy Analyst	0.165*** (0.027)	0.140 (0.129)	0.174*** (0.023)
Controls	Yes	Yes	Yes
Firm-by-Month FE.	Yes	Yes	Yes
Observations	856	856	856
(Pseudo)R ²	0.67	0.69	0.67
Within R ²	0.07	0.008	0.06

This table reports the effects of the ChatGPT Ban on analyst forecasts. The dependent variables are the forecast outcomes in group g covering firm i in month t , and include: a) a dummy variable set to 1 if at least one forecast is issued by group g covering firm i in month t , and 0 otherwise (forecast likelihood); b) number of forecasts by each group g covering firm i in month t (forecast frequency); c) number of forecasts scaled by ex-ante (i.e., pre-ban period) number of analyst coverage for each group g covering firm i in month t (forecast frequency per analyst). *Ban* is a time dummy that equals one for the ban period and zero for the pre-periods, i.e., January to March. We compare the outcome for the group of analysts located in Italy with those in other countries covering the same Italian firm, controlling for firm-by-month fixed effects. Standard errors are clustered by industry.

Table 5: The Effect of the ChatGPT Ban on Bid-Ask Spread

Dependent variable	Bid-ask spread	Bid-ask Adj
<i>Italy</i> × <i>Ban</i>	0.028*** (0.002)	0.031*** (0.002)
Italy	0.004*** (0.001)	0.004*** (0.0010)
Controls	Yes	Yes
Industry FE.	Yes	Yes
Month FE.	Yes	Yes
Observations	1,553	1,543
R ²	0.45	0.44
Within R ²	0.34	0.34

This table reports the effects of the ChatGPT Ban on the bid-ask spread. The sample period is from January to April 2023. Bid-Ask is the monthly average of the daily bid-ask spread, which is the bid price minus the ask price divided by the closing price. Bid-Ask Adj. is the monthly average of the daily bid-ask spread weighted by daily trading volume. We control for industry and month-fixed effects. Standard errors are clustered at the industry level.

Table 6: The Heterogeneous Effects of the ChatGPT Ban on Bid-Ask Spread

Dependent variable	Bid-ask Adj	Bid-ask Adj	Bid-ask Adj
Dummy (D) =	More Analysts	High Foreign	High Institution
<i>Italy</i> × <i>Ban</i> × <i>D</i>	-0.016***	-0.010***	-0.010**
	(0.002)	(0.002)	(0.004)
<i>Italy</i> × <i>Ban</i>	0.039***	0.035***	0.036***
	(0.002)	(0.002)	(0.003)
Ban×D	1.75e-5	0.0001	-0.0001
	(0.0005)	(0.0007)	(0.0007)
Italy×D	-0.009***	-0.007**	-0.008***
	(0.003)	(0.002)	(0.002)
Italy	0.008***	0.007***	0.008***
	(0.002)	(0.001)	(0.002)
D	-0.0001	-0.0004	-0.0001
	(0.001)	(0.001)	(0.002)
Controls	Yes	Yes	Yes
Industry FE.	Yes	Yes	Yes
Month FE.	Yes	Yes	Yes
Observations	1,543	1,467	1,531
R ²	0.50	0.46	0.47
Within R ²	0.40	0.37	0.37

This table reports the heterogeneous effects of the ChatGPT Ban on the bid-ask spread. The sample period is from January to April 2023. Bid-Ask is the monthly average of the daily bid-ask spread, which is the bid price minus the ask price divided by the closing price. Bid-Ask Adj. is the monthly average of the daily bid-ask spread weighted by daily trading volume. More analysts is a dummy variable set to 1 if the ex-ante number of analysts covering the Italian firm in 2023 Q1 is above median and 0 otherwise. High foreign (institution) is a dummy variable set to 1 if the foreign (institution) ownership level of the Italian firm by the end of 2023 Q1 is above median and 0 otherwise. These dummies are the same for the matched firms as those for the Italian firm. We control for industry and month-fixed effects. Standard errors are clustered at the industry level.

Table 7: The Effects of the ChatGPT Ban on the Market Reaction to Earnings Announcements

Dependent variable	Ret [-1,+1]	Ret [-2,+2]	Ret [-3,+3]
<i>Italy</i> × <i>Ban</i> × <i>UE</i>	0.417** (0.164)	0.901*** (0.155)	0.943*** (0.146)
<i>Ban</i> × <i>UE</i>	-0.152* (0.081)	-0.285*** (0.092)	-0.305*** (0.099)
<i>Italy</i> × <i>UE</i>	-0.037 (0.052)	-0.009 (0.082)	-0.071 (0.127)
<i>Italy</i> × <i>Ban</i>	0.081** (0.036)	0.105** (0.041)	0.128*** (0.035)
<i>UE</i>	0.070 (0.061)	0.078 (0.069)	0.104 (0.097)
<i>Italy</i>	0.005 (0.011)	0.003 (0.013)	-0.0007 (0.016)
Industry FE.	Yes	Yes	Yes
Month FE.	Yes	Yes	Yes
Observations	230	230	230
R ²	0.14	0.15	0.17
Within R ²	0.05	0.07	0.07

This table reports the effects of the ChatGPT Ban on market reaction to earnings announcements. *Ret* is the cumulative abnormal return for firm *i* during three different time windows - three, five, or seven days centered on the earnings announcement dates. *Ban* is a time dummy that equals one for the ban period and zero for the pre-periods, i.e., January, February, and March (before March 31st). We include industry and month-fixed effects, and the standard errors are clustered by industry.

A. Supplementary Appendix

Figure A1 An Illustrative Example of the ChatGPT Ban on Firm Value

Expert.AI, a listed firm on the Milan Stock Exchange headquartered in Modena, Italy, specializes in providing organizations with data analysis and business solutions. In its 2022 annual report, Expert.AI emphasized the importance of technologies such as “GPT,” “Machine Learning,” “Natural Language Processing,” “Artificial Intelligence,” “Machine Translation,” and “Speech Recognition,” mentioning the keywords more than 40 times. Additionally, on February 15, 2023, Expert.AI announced the launch of its Hybrid AI Platform that integrates GPT, highlighting its reliance on Generative AI (source: Refinitiv/PR Newswire). As an illustrative example, we examine the stock price reaction of Expert.AI to the ban announcement and find that its stock price experienced a decline of approximately 8% on the announcement day of the Ban.




Figure A2: Value-weighted cumulative return for portfolios consisting of Italian stocks in AI industries and Non-AI industries, and a portfolio that longs Italian AI stocks and shorts Italian non-AI stocks. Returns are value-weighted. Day 20 is the day of the restoration of ChatGPT in Italy. Day 60 is about 2 months after the restoration of ChatGPT in Italy.




Figure A3: Value-weighted cumulative return for portfolios consisting of Italian stocks in AI industries and Non-AI industries, and a portfolio that longs Italian AI stocks and shorts Italian non-AI stocks. Returns are value-weighted and adjusted for the EU (excluding Italy) market factor. Day 20 is the day of the restoration of ChatGPT in Italy.

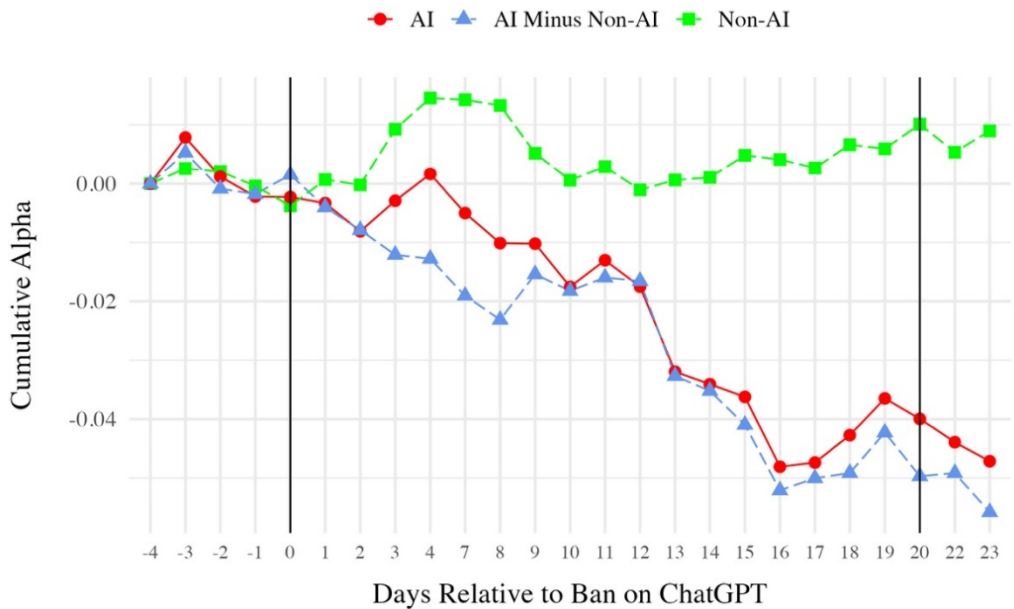


Figure A4: The effect of the Ban on the cumulative returns throughout the ban period, for Italian AI stocks versus non-AI stocks. We use four alternative matching methods for robustness, namely CEM, PSM 1-to-3, PSM Based on Mahalanobis Distance, and PSM using UK firms only. We examine the heterogeneous effects across different firms by estimating the following equation using the four matched samples:

$$Return_{i,t} = \alpha_i + \beta_1 Italy_i \times AI_i + \beta_2 Italy_i + X_{i,t-1} + Industry FE + \epsilon_{i,t},$$

where *Italy* equals one for all Italian stocks and zero for the matched non-Italian stocks. *AI_i* is a dummy variable that equals one for firms in AI industries. We include industry fixed effects and a vector of controls. The dots represent the estimated value of β_1 and the bars correspond to 95% confidence intervals. Standard errors are robust to heteroskedasticity.

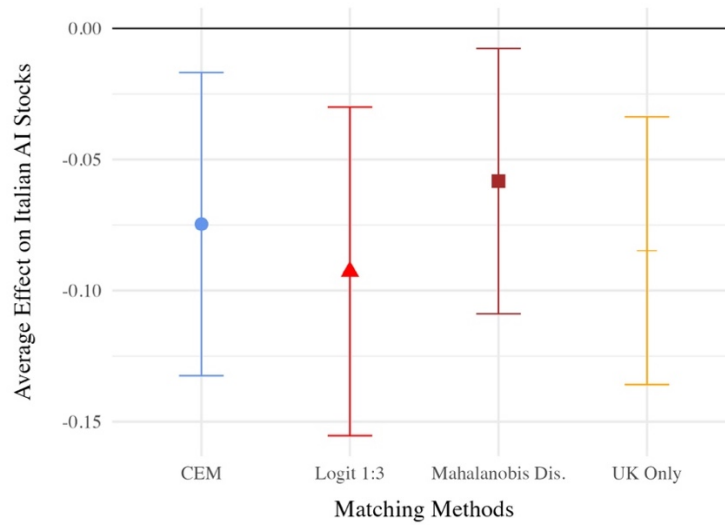


Table A1: Matching Accuracies for Different Matching Methods

Panel A: PSM (Logistic Link) – EU Plus UK			
	Italian Stocks (N=228)	Matched European Stocks (N=228)	Difference (Italy-Europe)
Mkt Cap	2258.04 (5.23)	1974.96 (6.2)	283.09 (0.73)
ROA	0.04 (4.09)	0.04 (3.48)	0 (0.94)
Panel B: PSM (Mahalanobis Distance) – EU Plus UK			
	Italian Stocks (N=217)	Matched European Stocks (N=217)	Difference (Italy-Europe)
Mkt Cap	2156.12 (5.24)	2066.78 (5.34)	89.34 (0.16)
ROA	0.04 (4.35)	0.04 (4.21)	0 (0.10)
Panel C: CEM – EU Plus UK			
	Italian Stocks (N=291)	Matched European Stocks (N=291)	Difference (Italy-Europe)
Mkt Cap	1586.12 (5.47)	1936.97 (6.64)	-350.85 (-0.33)
ROA	0.04 (4.33)	0.03 (3.86)	0 (0.33)
Panel D: PSM (Logistic; 1:3 Match) – EU Plus UK			
	Italian Stocks (N=228)	Matched European Stocks (N=677)	Difference (Italy-Europe)
Mkt Cap	2258.04 (5.23)	2042.12 (7.44)	215.93 (0.67)
ROA	0.04 (4.09)	0.04 (3.64)	0 (0.99)
Panel E: PSM (Logistic) - UK			
	Italian Stocks (N=260)	Matched UK Stocks (N=260)	Difference (Italy-Europe)
Mkt Cap	1829.98 (5.64)	2113.26 (5.87)	-283.27 (-4.33)
ROA	0.03 (3.92)	0.03 (3.7)	0.01 (0.9)

This table reports the matching accuracy between Italian stocks and stocks in the European Union and the United Kingdom, excluding Italy. We find a matched European stock in the same industry for each Italian firm with the closest market capitalization and ROA without replacement. In Panel A, we use the propensity score method to compute the distance between units as the absolute difference between the propensity scores of pairs of units. The propensity scores are estimated using logistic regression. Finally, we yield a sample of 456 firms in total (228 Italian stocks and 228 European stocks). In Panel B, we compute the Mahalanobis distance matrices from the covariates. In Panel C, we use coarsened exact matching (CEM). In Panel D, we follow the same method as in Panel A, except that each Italian stock is matched to 3 European stocks. In Panel E, we use the propensity score method to find a matched firm with the closest market cap in the same industry from the subsample consisting of only UK firms. Mkt Cap is in EUR millions. We report the t-values in parentheses.

Table A2: The Heterogeneous Effects of ChatGPT Ban on Market-adjusted Returns

D=	AI- Industry	High Attention	High AI Topic	New	Small Mkt.	Low Foreign	High Govern	High Insider
<i>Italy</i> × <i>D</i>	-0.093*** (0.032)	-0.074*** (0.027)	-0.139*** (0.049)	-0.070** (0.034)	-0.067** (0.030)	-0.079*** (0.026)	0.074** (0.031)	-0.058* (0.031)
D		0.048** (0.022)	0.071** (0.029)	0.054* (0.031)	0.043* (0.023)	0.067*** (0.022)	-0.048* (0.029)	0.038 (0.029)
Italy	0.025 (0.019)	0.027 (0.020)	0.025 (0.023)	0.036* (0.020)	0.019 (0.020)	0.036* (0.021)	-0.029 (0.021)	0.025 (0.021)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	389	387	94	313	389	387	200	212
R ²	0.21	0.21	0.26	0.27	0.19	0.26	0.26	0.20
Within R ²	0.11	0.10	0.14	0.16	0.07	0.16	0.15	0.09

This table reports the heterogeneous effects of ChatGPT Ban on market-adjusted returns. We define the ban period as the period from the announcement date of the Ban to the announcement date of the Reactivation, namely [2023-03-31,2023-04-28]. We then set the benchmark period as January, February, and March. Ban is a dummy variable set to 1 if it is the ban period and 0 otherwise. Mkt-Adj Ret is obtained by cumulating the daily abnormal return over the whole ban period. The daily abnormal return is adjusted based on the CAPM model:

$$R_t^{mkt\ adj} = R_t - \beta_t R_{MKT-RF,t},$$

where β_t is estimated by running the following regression with daily returns ranging from [-252, -1]:

$$R_t - R_t^f = \alpha_t + \beta_t R_{MKT-RF,t} + \varepsilon_t$$

where R_t^f is the risk-free rate and $R_{MKT-RF,t}$ is the value-weighted portfolio return of all stocks in EU plus UK excluding those in Italy in excess of the risk-free rate. All control variables are lagged by one month. We control for industry and month fixed effects and standard errors are robust to heteroskedasticity.

Table A3: The Heterogeneous Effects of the Ban on Stock Returns Across Firms with Different Ownership Structures

D=	High Domestic	High Govern	High Insider	High Retail	High Institutional
<i>Italy</i> × <i>D</i>	-0.078*** (0.024)	0.062** (0.029)	-0.065** (0.031)	-0.034 (0.034)	0.015 (0.030)
D	0.063*** (0.022)	-0.044 (0.031)	0.038 (0.031)	0.018 (0.031)	-0.005 (0.027)
Italy	0.039** (0.018)	-0.018 (0.020)	0.036 (0.022)	0.035 (0.024)	0.005 (0.024)
Controls	Yes	Yes	Yes	Yes	Yes
Industry FE.	Yes	Yes	Yes	Yes	Yes
Observations	387	200	212	245	386
R ²	0.26	0.25	0.24	0.22	0.19
Within R ²	0.17	0.16	0.15	0.11	0.09

This table reports the heterogeneous effects of the Ban on stock returns. We exploit five ex-ante firm-level ownership measures - domestic, government, retail, institutional, and insider ownership. We download the ownership level of domestic, retail, government, institutional, and insider investors for each firm from Refinitiv. We set the dummy variables to one if the ownership level in March for a particular ownership type is above the median and 0 otherwise. Italy is a dummy variable set to 1 for Italian stocks and 0 otherwise. We include industry fixed effects and report standard errors robust to heteroskedasticity in parentheses.

Table A4: A Placebo Test using Italy's Outlook Downgrade by Moody's on Aug. 6 2022

Panel A: Announcement Effect on August 6, 2022					
	Italian Stocks	Matched	Difference		
[0, +1]	0.002 (0.655)	0.008** (2.08)	-0.006 (-1.15)		
[0, +2]	-0.003 (-0.812)	-0.003 (-0.870)	-0.0005 (-0.121)		
[0, +3]	0.003 (0.811)	0.008* (1.95)	-0.005 (-0.941)		

Panel B: Heterogeneous Effects on Stock Return					
D=	AI-Industry	High AI Att	High AI Topic	New	Small Mkt.
<i>Italy</i> × <i>D</i>	-0.025 (0.017)	0.037* (0.019)	0.096*** (0.020)	0.003 (0.021)	0.096*** (0.020)
D		0.004 (0.018)	-0.038* (0.021)	0.011 (0.020)	-0.029* (0.018)
Italy	-0.024** (0.010)	-0.031*** (0.010)	-0.035*** (0.012)	-0.030*** (0.011)	-0.030*** (0.010)
Controls	Yes	Yes	Yes	Yes	Yes
Industry FE.	No	Yes	Yes	Yes	Yes
Observations	387	385	94	311	387
R ²	0.45	0.46	0.56	0.48	0.46
Within R ²	0.27	0.28	0.38	0.29	0.28

This table reports the announcement effects and heterogeneous effects of Italy's outlook cut by Moody's on stock returns. On August 6, 2022, Italy's sovereign rating outlook was lowered to negative from stable by Moody's. We expect the Italian market to respond negatively but without differential effects on firms with high vs. low AI exposures. In Panel A, we focus on the cumulative returns around the announcement window [0,+1], [0,+2], and [0,+3]. We focus on the cumulative returns in [0,+3] in Panel B and exploit five ex-ante firm-level measures - AI industry, attention to ChatGPT, AI topics mentioned in conference call transcripts, age, and firm size. Firms are classified as having high AI exposure if they are in the professional, scientific, and technical services sector and information sector and low AI exposure otherwise. High AI Att equals one for firms with above the median level ex-ante investor attention and zero otherwise. We calculate the firm-level ex-ante investor attention to ChatGPT by taking the average of the country-level Google search index weighted by the percentage of shares held by the investors from each country. To be conservative, we require the sum of country-level ownership to be above 80% and below 110%. Firms are defined as new if the founding year of the firm is more recent than the median founding year within the industry. We download the incorporation date of each firm from Refinitiv and replace the missing value with manually-collected founding years of Italian firms from Wikipedia. Firms are defined as small if their market capitalization on March 24 is below the median value within the industry and large otherwise. Italy is a dummy variable set to 1 for Italian stocks and 0 otherwise. All control variables are lagged by one month. We control for industry fixed effects in Panel B. We report standard errors robust to heteroskedasticity in both panels.

Table A5: The Effects of the ChatGPT Ban on Information Processing: Volatility and Turnover

Dependent Var.:	Volatility	Idio_Vol	Turnover
<i>Italy</i> × <i>Ban</i>	-0.110 (0.117)	-0.142 (0.107)	-3.68** (1.67)
Italy	-0.220*** (0.074)	-0.212*** (0.072)	2.29 (1.65)
Lag Book-to-	-0.023 (0.057)	-0.009 (0.054)	2.67 (2.32)
Lag Idio Vol	0.046*** (0.007)	0.042*** (0.005)	0.527*** (0.105)
Lag Market Cap	-0.032*** (0.008)	-0.036*** (0.011)	0.674** (0.267)
Lag Return	-0.287 (0.629)	-0.675 (0.451)	25.7*** (6.61)
Industry FE.	Yes	Yes	Yes
Month FE.	Yes	Yes	Yes
Observations	1,557	1,417	1,547
R ²	0.20	0.22	0.13
Within R ²	0.15	0.17	0.08

This table reports the effects of the ChatGPT Ban on stock return volatility and turnover. Volatility is a monthly measure of the standard deviation of daily raw return. Idio_Vol is the standard deviation of daily abnormal return based on the CAPM model. Turnover is the monthly average of the ratio of traded shares to the total shares outstanding. We control for industry and month-fixed effects. Standard errors are clustered at the industry level.

Table A6: Information Processing Results using only UK firms as Matched Sample

Effects on Micro-Structure Outcomes					
Dependent Var.:	Bid-Ask	Bid-Ask Adj	Volatility	Idio_Vol	Turnover
<i>Ban</i> × <i>Italy</i>	0.031*** (0.002)	0.034*** (0.003)	-0.188 (0.154)	-0.069 (0.147)	0.330 (3.72)
Controls	Yes	Yes	Yes	Yes	Yes
Industry FE.	Yes	Yes	Yes	Yes	Yes
Month FE.	Yes	Yes	Yes	Yes	Yes
Observations	1,776	1,775	1,792	1,454	1,779
R ²	0.61	0.60	0.44	0.44	0.44
Within R ²	0.15	0.14	0.15	0.12	0.09

This table reports the effects of the ChatGPT ban using matched UK firms. We define the ban period as the period from the announcement date of the Ban to the announcement date of the Reactivation, namely [2023-03-31,2023-04-28]. We then set the benchmark period as January, February, and March 2023. Ban is a dummy variable set to 1 if it is the ban period and 0 otherwise. Italy is a dummy variable set to 1 for Italian stocks and 0 otherwise. We control for industry and month-fixed effects, and standard errors are clustered at the industry level.

Table A7: Number of AI Topic Mentions in Italian Firms' Conference Call Transcripts

	Obs.	Mean	St. Dev.	50%	Max.
GPT	80	0	0	0	0
Machine Learning	80	0.013	0.112	0	1
Computer Vision	80	0	0	0	0
Machine Vision	80	0	0	0	0
Deep Learning	80	0	0	0	0
Virtual Agents	80	0	0	0	0
Image Recognition	80	0	0	0	0
Natural Language Processing	80	0	0	0	0
Speech Recognition	80	0	0	0	0
Pattern Recognition	80	0	0	0	0
Object Recognition	80	0	0	0	0
Neural Networks	80	0	0	0	0
AI Chatbot	80	0	0	0	0
Supervised Learning	80	0	0	0	0
Text Mining	80	0	0	0	0
Unsupervised Learning	80	0	0	0	0
Image Processing	80	0	0	0	0
Mahout	80	0	0	0	0
Recommender Systems	80	0	0	0	0
Support Vector Machines	80	0	0	0	0
Random Forests	80	0	0	0	0
Latent Semantic Analysis	80	0	0	0	0
Sentiment Analysis	80	0	0	0	0
Opinion Mining	80	0	0	0	0
Latent Dirichlet Allocation	80	0	0	0	0
Predictive Models	80	0	0	0	0
Kernel Methods	80	0	0	0	0
Keras	80	0	0	0	0
Gradient Boosting	80	0	0	0	0
Opencv	80	0	0	0	0
XGboost	80	0	0	0	0
Libsvm	80	0	0	0	0
Word2vec	80	0	0	0	0
Machine Translation	80	0	0	0	0
Sentiment Classification	80	0	0	0	0
Artificial Intelligence	80	0.138	0.443	0	2
Sum	80	0.15	0.506	0	3
Total Words	80	8634.488	2844.527	8350	15001

This table reports the frequency of Italian firms mentioning ChatGPT and AI in their earnings calls' transcripts in Q1 2023. We utilize a keyword list compiled by Acemoglu et al. (2022), which use the keywords to construct firm-level AI job vacancies. We include terms such as "GPT" in our analysis as well. We download Italian firms' earnings calls' transcripts from Refinitiv Eikon. We can retrieve the data for 80 firms. We translate the keywords to Italian for a small set of annual reports written in Italian and do the search.

Table A8: Number of AI Topic Mentions in Italian Firms' 2022 Annual Reports

	Obs.	Mean	St. Dev.	50%	Max.
GPT	264	0.034	0.342	0	5
Machine Learning	264	0.174	0.670	0	5
Computer Vision	264	0.015	0.150	0	2
Machine Vision	264	0.004	0.062	0	1
Deep Learning	264	0.030	0.211	0	2
Virtual Agents	264	0.008	0.123	0	2
Image Recognition	264	0.008	0.087	0	1
Natural Language Processing	264	0.008	0.087	0	1
Speech Recognition	264	0.011	0.137	0	2
Pattern Recognition	264	0.000	0.000	0	0
Object Recognition	264	0.000	0.000	0	0
Neural Networks	264	0.015	0.122	0	1
AI Chatbot	264	0.000	0.000	0	0
Supervised Learning	264	0.000	0.000	0	0
Text Mining	264	0.000	0.000	0	0
Unsupervised Learning	264	0.000	0.000	0	0
Image Processing	264	0.004	0.062	0	1
Mahout	264	0.000	0.000	0	0
Recommender Systems	264	0.000	0.000	0	0
Support Vector Machines	264	0.000	0.000	0	0
Random Forests	264	0.000	0.000	0	0
Latent Semantic Analysis	264	0.000	0.000	0	0
Sentiment Analysis	264	0.023	0.194	0	2
Opinion Mining	264	0.000	0.000	0	0
Latent Dirichlet Allocation	264	0.000	0.000	0	0
Predictive Models	264	0.042	0.251	0	3
Kernel Methods	264	0.000	0.000	0	0
Keras	264	0.000	0.000	0	0
Gradient Boosting	264	0.000	0.000	0	0
Opencv	264	0.000	0.000	0	0
XGboost	264	0.000	0.000	0	0
Libsvm	264	0.000	0.000	0	0
Word2vec	264	0.000	0.000	0	0
Machine Translation	264	0.004	0.062	0	1
Sentiment Classification	264	0.000	0.000	0	0
Artificial Intelligence	264	0.973	3.105	0	33
Sum	264	0.705	3.864	0	42
Total Words	264	86361.655	79718.256	66404.5	559419

This table reports the frequency of Italian firms mentioning ChatGPT and AI in their 2022 annual reports. We utilize a keyword list compiled by Acemoglu et al. (2022), which use the keywords to construct firm-level AI job vacancies. We include terms such as “GPT” in our analysis as well. We download Italian firms’ 2022 annual reports from Refinitiv Eikon. We can retrieve annual reports for 264 firms. We translate the keywords to Italian for a small set of annual reports written in Italian and do the search.

B. Supplementary Appendix on Market-Wide Response

Table B1 Returns during the Announcement Window and the Whole Ban Period

	Italian Stocks	Matched	Difference
[0,+1]	0.006* (1.66)	0.002 (0.468)	0.004 (0.772)
[0,+3]	0.004 (0.884)	-0.0004 (-0.098)	0.005 (0.961)
Ban Period [0, +20]	0.034*** (3.98)	0.018* (1.78)	0.016 (1.05)
[+20, +21]	-0.003 (-0.755)	-0.005 (-1.36)	0.002 (0.317)
[+20, +23]	-0.012*** (-3.12)	-0.006 (-1.27)	-0.006 (-1.05)

This table reports the cumulative returns of the portfolios of Italian stocks, matched European stocks, and their difference across different time windows. We report t-statistics adjusted for heteroskedasticity in parentheses.

Figure B1: Value-weighted announcement return around ChatGPT ban (March 31st) and restoration dates (April 28th) for Italy stocks versus EURO STOXX 50, MSCI Euro, or matched EU stocks




Figure B2: Italian stock return minus the matched European stock return over the ban period. The red line is the value-weighted cumulative returns difference. The grey lines represent 95% confidence intervals based on heteroskedasticity-robust standard errors.




Figure B3: Cumulative return around the ban period, for Italian stocks versus matched EU stocks. We use four different matching methods for robustness. The first one is PSM using logistic regression to estimate propensity scores. Second, we match firms using the Mahalanobis distance matrix from the covariates. Third, we use Coarsened Exact Matching (CEM), a semi-parametric matching method. We conduct one-to-one matching for the first three methods to find a European stock in the same industry with the closest market capitalization and ROA. Forth, we use the PSM with 1-to-3 matching. Last, we use PSM to find a matched firm in the same industry with the closest size from a subsample of UK firms only.

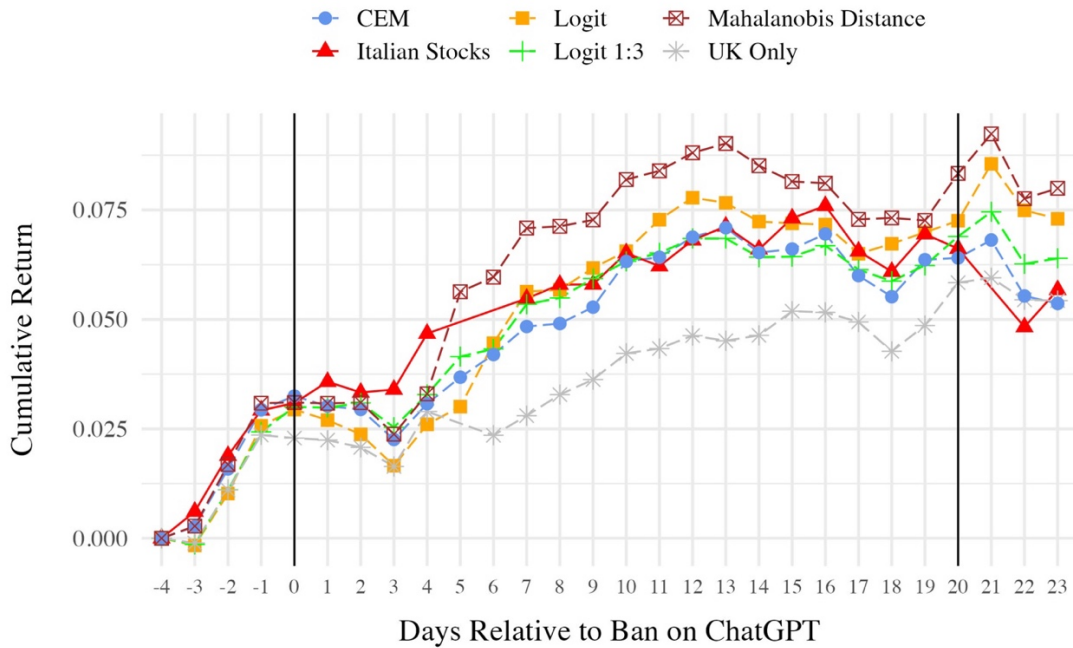


Figure B4: The effect of the Ban on the cumulative returns in the ban period, for Italian stocks versus matched EU stocks. We use four alternative matching methods for robustness, namely CEM, PSM 1-to-3, PSM Based on Mahalanobis Distance, and PSM using UK firms only. In Panel A, we examine the effects of the Ban on ChatGPT on Italian stock market by estimating the following equation using each of the five matched samples:

$$Return_{i,t} = \alpha_i + \beta_1 Italy_i + X_{i,t-1} + Industry\ FE + \epsilon_{i,t},$$

where *Italy* equals one for all Italian stocks and zero for the matched non-Italian stocks. We include industry fixed effects and a vector of controls. The dots represent the estimated value of β_1 and the bars correspond to 95% confidence intervals. Standard errors are robust to heteroskedasticity.

