

Do capital-goods incentive announcements move Italian industrial stocks?

An event study with LLM-coded regulatory acts, 2019–2026

Marco Belardi

Independent researcher. ORCID: 0009-0002-0471-3502. Disclosure: the author advises the Italian Ministry of Enterprises and Made in Italy (MIMIT) on the implementation of the incentive measures examined in this study and writes on them for Il Sole 24 Ore. The study was conducted independently, on public data only; it was not commissioned, reviewed or endorsed by MIMIT or by any other institution, and the views expressed are the author's own.

Version 2.1, 10 September 2026. Working paper prepared for SSRN. This document is a research study: it contains no investment recommendation. Earlier versions circulated under the title «Artificial intelligence and open data for anticipating price trends in equities and derivatives».

Keywords: event study; regulatory announcements; market efficiency; large language models; text coding; industrial policy; capital-goods tax incentives; Transizione 4.0; Transizione 5.0; Euronext Milan.

JEL codes: G14, G18, H25, L52, C58.

Abstract

This paper asks whether the Italian stock market prices the acts of capital-goods incentive policy (Transizione 4.0, Transizione 5.0, Nuova Transizione 5.0) when they are issued, and whether a large language model (LLM) reading those acts can extract a signal that the market incorporates with a lag. A corpus of 37 legislative and administrative acts between April 2019 and June 2026 is coded manually and, independently and blind, by an LLM; concordance is near-complete on category and date, substantial on intensity and moderate on expected sign. A market-model event study on 30 directional events and a two-tier basket of eight Euronext Milan stocks (domestic users of incentivised assets and suppliers of eligible capital goods) finds no abnormal return after the event under any coding: signed CAAR of +0.09% on [+1, +5] and +0.39% on [0, +5] under the manual coding, both indistinguishable from zero; a publication-day estimate of +0.29% that does not survive multiple-testing correction, the Corrado rank test, the Kolari-Pynnönen cross-correlation adjustment, the clustering of overlapping acts or the LLM coding. The result holds at the political-announcement date, is robust to estimation windows, basket composition and a survivorship control, and has a minimum detectable effect of about 1.5%. A post hoc hypothesis that the direction of acts follows the sector cycle is rejected on machine-tool orders and on official Eurostat data: the direction of an act is explained by its administrative nature (industrial policy versus spending control). The contribution is a documented negative result on a source of unstructured open data not covered by the financial text-mining literature.

1. Introduction

Italian industrial policy for capital goods has been implemented since 2016 through a sequence of tax measures (hyper-depreciation, the Transizione 4.0 tax credit, the Transizione 5.0 tax credit, the Nuova Transizione 5.0 depreciation surcharge) whose operating rules are set out in budget laws, decree-laws, ministerial implementing decrees, tax-authority circulars and notices published by the managing agency (GSE). These acts are public by definition and are open data in the sense of the study: freely accessible, unstructured, and written in a technical-legal register that is costly to interpret.

The starting hypothesis is that part of the information contained in such acts (implementing decrees, allocation of resources, opening of application platforms, interpretative clarifications) is incorporated into prices with a lag, because it requires regulatory expertise to be read. If the lag exists and is measurable, an LLM that extracts structured variables from such acts could turn it into a signal. The study limits itself to verifying the existence and statistical consistency of the signal. It does not build an investment strategy and makes no recommendation.

The study was designed with a sequence of phases fixed in advance: construction of the event corpus and event study (phase 1), walk-forward backtest with transaction costs (phase 2), paper trading (phase 3) and final write-up (phase 4). The decision rule was set in advance: if the first hypothesis (existence of a cumulative abnormal return after the event) is not verified, the study closes with a documented negative result and no backtest is run. This is what happened. The paper therefore reports the event-study results in full, the robustness checks, and the tests of an alternative hypothesis that emerged during the analysis, namely that policy acts are endogenous to the sector cycle.

The paper is organised as follows. Section 2 sets out the theoretical framework. Section 3 states the hypotheses. Section 4 describes the research design. Section 5 describes the data. Section 6 describes the methodology. Section 7 reports the results. Section 8 discusses them and lists the limits. Section 9 concludes. Appendix A lists the event corpus; Appendix B is the work log; Appendix C is the version history of the document.

2. Theoretical framework

2.1 Informational efficiency

In the semi-strong form of the efficient-market hypothesis (Fama, 1970) prices incorporate all public information. Open data are public by definition: any advantage does not derive from access but from the speed and quality of interpretation. The study therefore tests a local violation of semi-strong efficiency linked to the cognitive cost of information.

2.2 Alternative data and financial text mining

The literature on extracting signal from text (Tetlock, 2007; Loughran and McDonald, 2011) reports measurable but small effects, concentrated in short windows. Quantitative funds have applied similar techniques for more than a decade; the residual value on generalist public sources is largely arbitrated away. Large language models have recently been used as text-to-signal instruments: Lopez-Lira and Tang (2023) show that headline classification by a general-purpose LLM predicts next-day returns in the cross-section of US stocks, and Kim, Muhn and Nikolaev (2023) extract political, climate and technological risk measures from earnings-call transcripts that predict volatility and investment. Both work on corporate or news text; no published study applies an LLM to legislative and administrative acts as the information source. The niche examined here (Italian technical acts on capital-goods incentives) is less covered and is the element of originality of the study.

2.3 Limits of LLMs in forecasting

LLMs are reliable in classification, entity extraction and sentiment analysis; they are not calibrated numerical predictors. In this study the LLM is used only to transform text into structured variables (date, measure, expected sign, sectors affected, intensity). Prediction is left to conventional statistical methods.

2.4 Known methodological risks

The design addresses the following risks explicitly:

- overfitting and backtest overfitting (Bailey et al., 2014);
- look-ahead bias, that is the use of information not available at the event date;
- survivorship bias, that is the exclusion of stocks delisted during the period;
- non-stationarity of the series and regime changes;
- neglected transaction costs and slippage.

2.5 Regulatory event studies and evaluations of the Italian incentives

The use of stock prices to measure the effect of regulation goes back to Binder (1985), who shows that regulatory events are hard to date because information leaks through the legislative process, and to Lamdin (2001), who sets out the consequences for event-study design: multiple candidate dates, anticipation, and clustering of events on the same day, which inflates test statistics unless cross-sectional correlation is accounted for (Kolari and Pynnönen, 2010). The present design addresses each of these points: announcement and publication dates are tested separately, inference is drawn at event level and with the cross-correlation-adjusted test, and overlapping acts are clustered.

The real effects of the Italian measures have been evaluated on firm-level data. Bratta, Romano, Acciari and Mazzolari (2023) find, on tax returns, that hyper-depreciation raised investment in 4.0 assets among adopting firms; the scientific committee established by the Ministry of Economy, MIMIT and Banca d'Italia reports about 29 billion euro of 4.0 tax credits accrued in the first three years and positive effects on investment, employment and turnover (Comitato scientifico, 2024), and its final report estimates that each euro of credit activated 1.5 to 2 euro of tangible investment in 2020–2023, with employment effects concentrated in micro and small firms and no measurable effect for large firms (Banca d'Italia and MEF, 2026); Cefis, Scrofani and Tubiana (2025) find positive but heterogeneous productivity effects with a difference-in-differences design. These studies measure what the incentives did to the firms that used them. The present study asks a different question: whether the market priced the acts when they were issued, on the listed firms most exposed to them. The evaluation literature indicates that the effects are real but diffuse across a population dominated by small firms, which is consistent with a small footprint on the value of the handful of listed users and suppliers.

3. Hypotheses

Three hypotheses were stated, with their test windows, before any price was observed. They were not deposited on a public registry; «fixed ex ante» in this paper means that they are documented in version 0.1 of this working document, dated before the price data were obtained (Appendix C).

H1. Announcements of incentive measures for capital goods generate statistically significant cumulative abnormal returns (CAR) on the exposed stocks in the window $[0, +5]$ trading days.

H2. The CAR is concentrated in days $[+1, +5]$ rather than on day 0, indicating delayed incorporation of the information.

H3. The intensity of the event estimated ex ante (financial allocation, breadth of the eligible population, sign) is correlated with the size of the CAR.

If H1 is not verified, the study closes with a documented negative result. If H1 is verified but H2 is not, the signal exists but cannot be exploited after the event.

H4 (post hoc). A fourth hypothesis emerged from an exploratory finding in the political-announcement window (Section 7.3): the direction of incentive acts is endogenous to the state of the sector, in the sense that expansionary acts follow a contraction in domestic demand for capital goods and restrictive acts follow an expansion. H4 was formulated after observing the data and is treated accordingly: its tests (Section 7.9) are reported as confirmatory of a hypothesis generated in-sample, not as results of a hypothesis fixed ex ante.

4. Research design

4.1 Market

Euronext Milan (formerly Borsa Italiana). The choice is motivated by the direct causal link between Italian industrial-policy acts and the fundamentals of listed issuers. Benchmark: FTSE MIB index. Control envisaged: FTSE Italia Mid Cap.

4.2 Basket

The initial basket (suppliers of 4.0 capital goods, mostly exporters) was replaced in phase 1 by a two-tier basket designed to strengthen the causal channel. Tier A: users of incentivised assets with predominantly domestic revenues (criterion: Italian share of revenues not below 50% and plants in Italy, to be verified on the geographical note of the financial statements): Sesa, Wiit, Tinexta, Zignago Vetro. Tier B: suppliers of eligible capital goods with an Italian share not below 20%: Biesse, Sabaf, Datalogic, El.En. Interpump and Danieli, initially candidates, were excluded because of dominant exports. Prima Industrie, a tier B supplier of laser and sheet-metal machinery delisted on 7 March 2023 after a tender offer, is used as the survivorship-bias control (Section 7.8).

Table 1. Basket composition and verification status of the Italian revenue share

Ticker	Issuer	Activity	Tier and role	Italian share of revenues	Status
SES	Sesa	IT, software and digital services for firms	A, domestic user/supplier	~97% (FY to 30/04/2026: foreign 129.1 of 3,565 m; narrative disclosure)	A, verified
WIIT	Wiit	Cloud and data centres	A, domestic user	34.9% (FY2025: Italy 58.6 of 167.9 m)	Below threshold, reclassified
TNXT	Tinexta	Digital services, credit information, consulting	A, domestic user	83.7% (FY2025: Italy 381.1 of 455.6 m by company location; 81.8% by revenue area)	A, verified
ZV	Zignago Vetro	Glass containers	A, domestic user	69.5% (FY2025: 30.5% abroad)	A, verified
BSS	Biesse	Machinery for wood, glass, stone	B, supplier of eligible goods	Not disclosed (FY2025: EMEA 62.5%, Americas 23.8%, APAC 13.6%)	B, not verifiable

Ticker	Issuer	Activity	Tier and role	Italian share of revenues	Status
SAB	Sabaf	Components for household appliances	B, supplier/user	Not disclosed (FY2025: Europe excl. Turkey 29.3%, Turkey 23.4%, North America 23.4%)	B, not verifiable
DAL	Datalogic	Sensors, code reading, automation	B, supplier of eligible goods	8.9% by market (FY2025: Italy 44.6 of 500.1 m); 49.4% by invoicing company	Below threshold, reclassified
ELN	El.En.	Industrial and medical laser systems	B, supplier of eligible goods	22.2% (FY2025: Italy 131.2 of 590.8 m; industrial segment 53.7% Italy)	B, verified
PRI	Prima Industrie	Laser and sheet-metal machinery	B, supplier of eligible goods (control)	Not applicable (delisted 7 March 2023)	Survivorship control only

Source: Dataset_eventi_paniere_v0.12.xlsx, sheet Paniere; company annual reports FY2025 (Sesa: FY to 30 April 2026). Segmentation bases differ across issuers (customer location, company location, regional aggregates) and are reported as disclosed.

4.3 Instruments

The study uses spot equities only. The original work plan envisaged a simulated operating phase on index futures or single-stock options conditional on H1 and H2; since neither hypothesis is verified, that phase was not executed and no instrument other than spot equities enters the analysis.

4.4 Event universe

Period 2019–2026. Categories: publication of implementing decrees; allocation and refinancing of resources; opening or closure of application platforms (GSE, NPTR5 platform); circulars and rulings of the Italian Revenue Agency (Agenzia delle Entrate); MIMIT and Ministry of Economy (MEF) press releases; provisions in budget laws and decree-laws. The observation period was extended back to December 2019 to include the switch from hyper-depreciation to the 4.0 tax credit (Law 160/2019), an event of ambiguous sign that serves as a control, and to April 2019 to include the reintroduction of the 130% super-depreciation (Decree-Law 34/2019).

4.5 Windows

Estimation window: [-250, -21] trading days. Event window: [-5, +5], with separate analysis of [0], [+1, +5] and [-5, -1], the last to detect anticipation or leaks. The four windows were fixed before the analysis and are the only ones on which inference is drawn.

5. Data

5.1 Event corpus

A first version of the corpus (28 events, December 2019 to June 2026) was classified into seven categories: NORMA (primary legislation: budget law, decree-law, conversion law), CDM (approval by the Council of Ministers with press release, which precedes publication), DM (ministerial implementing decree, signature or publication), PIATTAFORMA (opening, closure or exhaustion of resources of GSE platforms), PRASSI (circulars, resolutions and rulings of the Revenue Agency, MIMIT operating circulars), PROROGA (extensions of delivery or use deadlines) and RESTRIZIONE (introduction of compliance obligations, suspensions, ceilings, exclusions). For each event the

expected sign on the basket (+1, 0, -1), the intensity (1 to 3) and, where known, the financial allocation were assigned ex ante, before observing prices, and were not modified afterwards.

Dates were verified on primary sources (Official Journal, government and MIMIT press releases, GSE notices). The verification produced several corrections with respect to the initial coding: Decree-Law 38/2026 was published in the Official Journal on 27 March 2026 with negative sign for users (recognition of only 35% of the credit to firms left in the unfunded waiting list, the so-called esodati 5.0), Decree-Law 42/2026 of 3 April 2026 with positive sign (restoration to 89.77%); the exhaustion of Transizione 5.0 resources (MIMIT decree of 6 November 2025, cut-off date 7 November), Decree-Law 175/2025 (Official Journal 21 November 2025) and Law 4/2026 (Official Journal 20 January 2026) were added. Events whose date could not be verified were excluded from the main analysis.

The corpus was then extended to 38 events (32 directional) with eight further records: Decree-Law 34/2019 (super-depreciation), Decree-Law 34/2020 (delivery extension), Decree-Law 73/2021 (one-year use of the credit), Decree-Law 4/2022 (intangible 4.0 assets), the Council of Ministers meeting and Law 213/2023 (2024 budget without measures for capital goods), Revenue Agency resolution 14/E/2026 and the GSE notice of 29 April 2026 (credit for the esodati). The attribution of event 6 was corrected: the extension of deliveries to 30 September 2023 is in Law 197/2022 and not in Decree-Law 198/2022, which does not contain it despite the announcement of the Council of Ministers of 21 December 2022. One event (Decree-Law 34/2020) is excluded because no price is available on day 0; 31 directional events and 248 event-stock pairs remain. The full corpus is listed in Appendix A.

A second verification round on primary sources (Official Journal, Normattiva, official decree texts, MIMIT website) and on specialised press, carried out for this version, found three miscoded events. Event 7: the further extension of the delivery deadline to 30 November 2023 is not in Decree-Law 132/2023 (29 September 2023) but in Law 14/2023, the conversion law of Decree-Law 198/2022, art. 12 c. 1-ter, Official Journal no. 49 of 27 February 2023. Event 19: Decree-Law 95/2025 (30 June 2025) contains no provision on 4.0 credits; the 2025 regime (2.2 bn ceiling, prior communication, priority by chronological order of transmission) derives from Law 207/2024, art. 1 c. 445-448, and was implemented by the MIMIT directorial decrees of 15 May 2025 (form) and 16 June 2025 (platform opening), the latter being the operative event. Event 30: the implementing decree of 7 May 2026 was not published in the Official Journal; it was registered by the Court of Auditors on 20 May 2026 and published on the MIMIT website on 11 June 2026, one trading day before the opening of the NPTR5 platform (event 29). Event 34 (Decree-Law 4/2022, art. 10) had a wrong content description: the article concerns tangible 4.0 assets (5% credit on the share of PNRR green investments above 10 m), not intangibles; sign and intensity are unchanged. Event 18 (ministerial decree of 24 April 2025) does not exist: the MIMIT register of Transizione 5.0 acts lists no such decree, and the changes introduced by Law 207/2024, art. 1 c. 427-429, were applied through revisions of the MIMIT FAQ (21 and 24 February, 10 April 2025) without an amending decree. The event is dropped. The date of event 19 was confirmed on the text of the directorial decree of 16 June 2025, which opened the platform from 14:00 of 17 June 2025. The three miscoded events are recoded in dataset v0.10 and the event study was re-estimated at the corrected dates (Section 7.5). The final corpus has 37 events; the directional events with a usable price history are 30 (240 event-stock pairs), all with a date verified on a primary source except event 16, whose Council of Ministers date rests on the government press-release register and specialised press.

A structural limit of the corpus was identified: budget laws are published on 30–31 December, in thin trading sessions, while their content is known from the Council of Ministers meeting in October and from the parliamentary votes. For this reason the announcement date was reconstructed for the NORMA events (Section 5.2).

5.2 Political-announcement dates

For ten NORMA events the date of the political announcement (approval by the Council of Ministers) was reconstructed: 2021 budget (16 November 2020), 2022 budget (28 October 2021), Decree-Law 50/2022 (2 May 2022), Decree-Law 198/2022 (21 December 2022), Decree-Law 132/2023 (28 September 2023), Decree-Law 19/2024 (26 February 2024), Decree-Law 39/2024 (26 March 2024), 2025 budget (15 October 2024), Decree-Law 175/2025 (20 November 2025), 2026 budget (17 October 2025) and 2024 budget (16 October 2023). In this version the dates of the Council of Ministers meetings for the 2024, 2025 and 2026 budget bills (press releases no. 54 of 16 October 2023, no. 100 of 15 October 2024, no. 146 of 17 October 2025) and for Decree-Law 39/2024 (no. 75 of 26 March 2024) were confirmed on the government press-release register and on specialised press. The announcement date of Decree-Law 132/2023 (Council of Ministers no. 52 of 27 September 2023, not 28 September) is no longer relevant because event 7 has been recoded (Section 5.1); for the extension introduced by parliamentary amendment in Law 14/2023 the announcement date is taken as the Senate vote of 15 February 2023 (final Chamber vote on 22 February, Official Journal on 27 February). All eleven announcement dates are therefore verified.

5.3 Prices

Adjusted daily closing prices from 1 June 2018 to 8 September 2026 for the eight stocks and the FTSE MIB index (2,099 daily observations), obtained via yfinance. Day 0 is the first trading day equal to or following the publication (or announcement) date. For the survivorship control, the daily series of Prima Industrie from 1 June 2018 to its last session on 3 March 2023 (1,211 observations) was added; from 12 July 2022 the price is pinned to the tender-offer consideration (€24.5–25.0), so the series is informative only for events up to June 2022.

5.4 Sector activity indicators

Two sources measure the state of the sector before each act. The first is the quarterly index of domestic machine-tool orders of the UCIMU Study Centre (the Italian machine-tool builders' association), reconstructed from press releases for 2022Q4–2026Q2 (15 quarters: 9 readings verified, 6 derived by difference from annual or half-yearly data and from absolute values; the change of base year from 2015 to 2021 introduced in 2024 makes the derivations approximate). The complete series is reserved to UCIMU members. Levels of the index (2021 = 100) were reconstructed for 2023Q1–2026Q2 and seasonally adjusted with a four-quarter moving average.

The second is the official Eurostat series of domestic turnover of the machinery industry (sts_intv_m, Italy, NACE C28, domestic component, seasonally adjusted, 2021 = 100) and the C28 production index (sts_inpr_m), monthly from January 2015 to June 2026. The series covers the whole study period. Three indicators are computed on the last month released at the reference date of each event, assuming a release lag of 45 days: 12-month change of domestic turnover (three-month moving average), 12-month change of production, and deviation of the three-month average from the twelve-month average.

5.5 Data sources

Table 2. Data sources

Source	Content	Nature
Gazzetta Ufficiale (Official Journal)	Primary legislation, implementing decrees	Unstructured (text)
MIT, MEF	Press releases, directorial decrees, FAQ	Unstructured
GSE	Platform openings, resource counters, notices	Semi-structured
Agenzia delle Entrate	Circulars, resolutions, rulings	Unstructured
Euronext Milan (via yfinance)	Prices, volumes, indices	Structured
Consob	Price-sensitive disclosures of issuers	Semi-structured

Source	Content	Nature
Confindustria (Study Centre)	Flash surveys, positions on the measures	Unstructured
UCIMU, Federmacchine	Quarterly machine-tool order surveys	Structured
Eurostat	Short-term business statistics, NACE C28 (sts_intv_m, sts_inpr_m)	Structured

6. Methodology

6.1 Pipeline

1. Corpus collection: Official Journal, MIMIT website, GSE, Revenue Agency, MEF.
2. Manual coding of each act (publication date, category, expected sign, intensity, allocation), performed before any price was observed.
3. Independent structured extraction with an LLM, blind to the manual coding, on the same acts (Section 6.5); measurement of the concordance between the two codings and re-estimation of the event study under the LLM coding.
4. Adjusted daily prices of stocks and indices.
5. Event study with the market model: abnormal returns (AR), CAR per event, cumulative average abnormal returns (CAAR) per category; t-tests and non-parametric tests (Corrado rank test).
6. Walk-forward out-of-sample backtest on a simple rule (entry at the close of day 0, exit at +3 or +5; costs of 15–20 bp per side, slippage): conditional on H1 and H2.
7. Paper trading on a demo account for 3–6 months: conditional on the backtest.

6.2 Event study

The script `event_study.py` (Python, pandas, scipy) estimates the market model on the window $[-250, -21]$, computes abnormal returns on $[-5, +5]$ and the CAR on the four windows $[-5, -1]$, $[0]$, $[+1, +5]$ and $[0, +5]$. CARs are multiplied by the expected sign so that expansionary and restrictive events can be summed; events with sign 0 are excluded from directional tests. Inference is at two levels: pooled over event-stock pairs and, as the main specification, on the mean per event, to correct for the cross-sectional correlation of stocks on the same day. Benjamini-Hochberg (BH) correction is applied across the four windows. The Corrado (1989) rank test is applied as a non-parametric check. Because all stocks in the basket share the same event day, abnormal returns are cross-sectionally correlated and pooled tests on pairs overstate significance; the standardised cross-sectional test of Boehmer, Musumeci and Poulsen (1991) is therefore reported together with its adjustment for the average cross-correlation of the estimation-window residuals (Kolari and Pynnönen, 2010).

6.3 Multiple-testing control and synthetic test

The script was tested on a synthetic series (random returns with beta 0.8 on the market). On data with no signal by construction, one of the 32 category-window combinations returned $p = 0.019$. The result confirms the risk of false positives from multiple tests and imposed, in the analysis phase, the Bonferroni or Benjamini-Hochberg correction and the restriction of inference to the windows and categories stated ex ante in Section 3.

6.4 Endogeneity tests

Each directional event is matched to the last sector reading released before its announcement (or publication) date. An event is «concordant» with H4 if an expansionary act follows a fall in domestic activity or a restrictive act follows a rise. Tests: one-sided binomial test on the share of concordant events; Spearman correlation between intensity and the activity indicator; Mann-Whitney test on the

indicator before expansionary versus restrictive acts; logit of the direction of the act on the indicator. Because several closely spaced events share the same sector reading, a clustering correction reduces the sample to distinct reading-sign combinations. Acts are also classified as industrial policy (PI: institution, refinancing, implementing decrees, platform openings) or spending control (CS: compliance obligations, ceilings, exhaustion of resources, cuts for budget coverage).

6.5 LLM extraction and concordance with the manual coding

The extraction was performed with a Claude model (Anthropic) accessed in September 2026 through an agent with web-retrieval tools. The model received, for each of the 37 acts, only the identifier used in the corpus (act reference and hypothesised date), the definitions of the seven categories, the sign rule (expected effect on a basket of Italian listed users and suppliers of incentivised capital goods) and a three-level intensity rubric (1 procedural or limited; 2 substantive change or moderate allocation; 3 new measure or major allocation, about 4 billion euro or more). It was instructed to retrieve the primary text or an official summary and to return, for each act, category, sign, intensity, allocation, date found, a one-sentence rationale, source and confidence. It had no access to the manual coding or to the results. By its own account, 13 acts were coded from a fetched official source, 17 from specialised press and 7 from the model's prior knowledge because the sources could not be fetched; no act was flagged as low confidence. The run is single; stability across repeated runs and across models was not tested.

Concordance is measured on the 37 acts with simple agreement and Cohen's kappa (linearly weighted for intensity). The event study is then re-estimated using the LLM sign and intensity in place of the manual ones, and on the consensus subset where the two codings agree on a non-zero sign.

6.6 Software

Python (pandas, scipy, statsmodels); Claude API for structured extraction; versioned notebooks. Price data via yfinance. The vectorbt or backtesting.py libraries and an Interactive Brokers demo account were reserved for phases 2 and 3, which were not executed.

7. Results

7.1 Baseline event study (25 verified events)

Sample: 25 events with verified date, 8 stocks, 200 event-stock pairs. Directional tests on 20 events with expected sign different from zero (160 pairs).

Table 3. Signed CAAR by window, baseline sample

Window	n pairs	CAAR %	p pooled	n events	p per event	p BH
[-5, -1]	160	-0.79	0.056	20	0.145	0.289
[0]	160	+0.35	0.017	20	0.125	0.289
[+1, +5]	160	-0.01	0.980	20	0.988	0.988
[0, +5]	160	+0.34	0.340	20	0.615	0.820

CAARs multiplied by the expected sign. p pooled: t-test on pairs; p per event: t-test on the mean per event; p BH: Benjamini-Hochberg correction across the four windows.

7.2 Reading of H1, H2 and H3

H1 (abnormal return in [0, +5]): not verified. The CAAR of +0.34% is not distinguishable from zero.

H2 (delayed incorporation in [+1, +5]): not verified. The post-event CAAR is nil (-0.01%). The only signal, +0.35% on day 0 with pooled p = 0.017, falls to p = 0.125 on the mean per event and to 0.289 after

correction for multiple tests. If an effect exists, it is priced within the publication session: consistent with semi-strong efficiency and contrary to the interpretative-lag thesis.

H3 (intensity): not verified, with sign opposite to expectation. The Spearman correlation between the intensity assigned ex ante and the signed CAR per event is -0.43 ($p = 0.061$): events with larger allocations show lower signed CARs. A plausible reading is that the most relevant measures are anticipated by the market in the preceding weeks (signed pre-event CAAR -0.79%) and that formal publication closes the information cycle rather than opening it.

By stock: no significant signed CAAR. Datalogic and Biesse (tier B, suppliers) show the highest values (+1.18% and +0.81% on [0, +5]); domestic users (tier A) show values close to zero or negative. The pattern is compatible with a supplier channel rather than a user channel, but the sample size does not allow inference.

By event: Law 234/2021 (first trading day 3 January 2022) shows a mean CAR of -9.1%, largely attributable to the sector rotation of early 2022 and not to the law. The Transizione 5.0 events of 2024 (Decree-Law 19, implementing decree, platform opening) all have negative CARs despite the positive expected sign. The two extensions (2022, 2023) have positive CARs consistent with expectation, but they are two observations.

7.3 Political-announcement window

For the NORMA events with a reconstructed announcement date, the event study was repeated on the announcement window (Council of Ministers) and compared with the publication window, for the same events and the same expected sign (9 directional events).

Table 4. Signed CAAR %, announcement versus publication window (9 directional events)

Window	[-5, -1]	[0]	[+1, +5]	[0, +5]
Council of Ministers announcement	-1.59 (p 0.003)	+0.05	-1.18	-1.13
Official Journal publication	-0.81	+0.21	+0.26	+0.47

Source: *car_annuncio_vs_publicazione.csv*; Dataset sheet *Annunci*.

Neither window shows a positive signed abnormal return after the event. The paired test between announcement CAR and publication CAR is not significant ($p = 0.24$). The thesis that the signal exists but is anticipated by the political announcement is not confirmed: even at the Council of Ministers date the stocks do not react in the expected direction.

An exploratory result, not among the hypotheses fixed ex ante and therefore to be treated with caution: in the week preceding the political announcement the signed CAAR is -1.59% ($p = 0.003$, 8 events out of 9 with negative sign). Exposed stocks underperform before an expansionary announcement and outperform before a restrictive one. This finding generated H4: incentives are announced when the sector is in difficulty, restrictions when resources are under pressure. If confirmed, the phenomenon would reverse the causal relation assumed at the start of the study: acts would not move prices; the condition of the sector would move the acts. Section 7.9 reports the tests.

7.4 Robustness and rank test

The robustness checks are run on the final corpus (30 directional events, 240 event-stock pairs; Section 7.5). The same checks on the baseline sample of 20 events (version 0.9 of this document) gave the same picture.

Table 5. Robustness of the event study on the extended corpus (signed CAAR %, mean per event, n = 30; p in brackets)

Specification	[-5, -1]	[0]	[+1, +5]	[0, +5]
Baseline: estimation [-250, -21], FTSE MIB, 8 stocks	-0.39 (0.42)	+0.29 (0.08)	+0.09 (0.83)	+0.39 (0.46)
Estimation [-120, -21]	-0.38 (0.43)	+0.30 (0.07)	+0.13 (0.76)	+0.44 (0.39)
Estimation [-500, -21]	-0.35 (0.47)	+0.30 (0.08)	+0.01 (0.98)	+0.32 (0.55)
Without Wiit (Italian share below threshold)	-0.24 (0.64)	+0.23 (0.19)	+0.03 (0.95)	+0.26 (0.61)
Without Wiit and Datalogic	-0.38 (0.42)	+0.21 (0.22)	-0.09 (0.85)	+0.13 (0.82)
Suppliers only (tier B)	+0.21 (0.78)	+0.06 (0.78)	+0.34 (0.47)	+0.40 (0.40)
Domestic users only (tier A)	-0.98 (0.03)	+0.53 (0.02)	-0.15 (0.80)	+0.37 (0.59)
Corrado rank test, 240 pairs (z; p)	-0.30 (0.76)	+1.18 (0.24)	-0.22 (0.83)	+0.28 (0.78)
BMP standardised test, 240 pairs (t; p)	-1.32 (0.19)	+2.38 (0.02)	-0.08 (0.94)	+0.97 (0.33)
Kolari-Pynnönen adjusted test, mean residual correlation 0.16 (t; p)	-0.83 (0.41)	+1.49 (0.14)	-0.05 (0.96)	+0.61 (0.54)

The rank test uses signed abnormal returns ranked over the estimation and event windows of each pair (Corrado, 1989; multi-day windows as in Campbell and Wasley, 1993). The BMP test standardises each pair's CAR by its estimation-window residual standard deviation; the Kolari-Pynnönen adjustment multiplies the BMP statistic by the square root of $(1 - r)/(1 + (n - 1)r)$, with r the average pairwise correlation of the residuals of the $n = 8$ stocks over the estimation windows. p -values are uncorrected for multiple tests.

The results do not depend on the estimation window or on the composition of the basket. The Corrado rank test, non-parametric and robust to the tails of return distributions, confirms the absence of post-event abnormal returns and does not support the publication-day reaction either ($z = 1.18$, $p = 0.24$): the marginal day-0 estimate of the parametric test rests on a few large observations. The pooled BMP test on pairs would call day 0 significant ($p = 0.02$), but the eight stocks react on the same day and their residuals are correlated (average 0.16); the Kolari-Pynnönen adjustment brings the statistic to $t = 1.49$ ($p = 0.14$), in line with the event-level inference. This is the expected pattern of a spurious pooled result under event-day clustering. The subset of domestic users shows a nominally significant day 0 (+0.53%, $p = 0.02$) and a negative pre-event window (-0.98%, $p = 0.03$): four stocks, no correction, and no post-event continuation; it is an indication for a larger sample, not a result.

The Italian share of revenues was verified on the FY2025 financial statements (Table 1). Tier A is confirmed for Sesa (about 97%), Tinexta (83.7%) and Zignago Vetro (69.5%); Wiit (34.9%) is below threshold. In tier B, El.En. (22.2%) is above the 20% threshold; Datalogic (8.9% by market of destination) is below it and is reclassified; Biesse and Sabaf do not disclose an Italian figure and publish only regional aggregates, so their tier assignment cannot be verified. The sensitivity of the extended-corpus results to the exclusion of Wiit and Datalogic is reported in Table 6.

7.5 Extended corpus and power analysis

On the final corpus (30 directional events, 240 event-stock pairs) the event study was repeated with the same specification. Table 6 reports the results after the recoding of events 7, 19 and 30 and the removal of event 18 (Section 5.1), re-estimated on the price series at the corrected dates; the implementation was validated by reproducing the previous CAR file on the unchanged events to numerical precision. The figures before the corrections are kept in Table 6b for comparison.

Table 6. Signed CAAR %, mean per event, final corpus (p; p BH)

Sample	[-5, -1]	[0]	[+1, +5]	[0, +5]
30 directional events, all	-0.39 (0.42; 0.61)	+0.29 (0.08; 0.33)	+0.09 (0.83; 0.83)	+0.39 (0.46; 0.61)
29 events with date verified on a primary source	-0.38 (0.45; 0.59)	+0.28 (0.11; 0.44)	+0.18 (0.69; 0.69)	+0.45 (0.39; 0.59)

Sample	[-5, -1]	[0]	[+1, +5]	[0, +5]
17 of these with intensity 2–3	-0.77 (0.19; 0.78)	+0.07 (0.72; 0.96)	-0.10 (0.87; 0.96)	-0.04 (0.96; 0.96)
30 events, 6 stocks (without Wiit and Datalogic)	-0.38 (0.42; 0.83)	+0.21 (0.22; 0.83)	-0.09 (0.85; 0.85)	+0.13 (0.82; 0.85)

Source: *car_corpus_finale_v1.0.csv*. Market model on [-250, -21] with log returns; CARs multiplied by the expected sign.

Table 6b. Same statistics before the corrections (dataset v0.8, 31 events including event 18), for comparison

Sample	[-5, -1]	[0]	[+1, +5]	[0, +5]
31 directional events, all	-0.41 (0.42; 0.55)	+0.40 (0.02; 0.07)	+0.22 (0.60; 0.60)	+0.62 (0.22; 0.44)
24 events with verified date	-0.67 (0.21; 0.28)	+0.43 (0.03; 0.13)	+0.37 (0.47; 0.47)	+0.81 (0.19; 0.28)
14 verified-date events with intensity 2–3	-0.97 (0.16; 0.62)	+0.08 (0.71; 0.98)	-0.10 (0.89; 0.98)	-0.02 (0.98; 0.98)
28 events, miscoded events 7, 19, 30 removed	-0.33 (0.54; 0.54)	+0.36 (0.04; 0.18)	+0.37 (0.42; 0.54)	+0.73 (0.19; 0.37)

The corrections change the size of the publication-day estimate but not the reading. At their corrected dates the three recoded events all have negative signed CARs on [0, +5] (event 7, -1.61%; event 19, -1.59%; event 30, -2.11%), against a positive contribution at the wrong dates: the day-0 CAAR falls from +0.40% ($p = 0.02$, corrected $p = 0.07$) to +0.29% ($p = 0.08$, corrected $p = 0.33$) and no longer survives the multiple-testing correction. The post-event windows are nil in every specification (+0.09% on [+1, +5]), and the day-0 estimate weakens further when the two stocks below their tier threshold are removed (+0.21%, $p = 0.22$). The signal is absent in the higher-intensity events, and the Spearman correlation between ex ante intensity and signed CAR per event on the final corpus is -0.34 ($p = 0.07$), the same sign as in the baseline sample (Section 7.2): the more relevant acts are priced before publication. H1 on [0, +5] and H2 remain not verified; the picture is that of semi-strong efficiency with, at most, a marginal information content on the publication day.

Power analysis: with 30 events and a standard deviation of the mean CAR per event of 2.79% on [0, +5], the minimum detectable effect (alpha 5%, power 80%) is 1.48%. A post-event effect below that threshold would not have been detected: the negative result should be read as «no effect above about 1.5%», not as absence of any effect.

7.6 LLM coding: concordance and robustness of the results to the coding

Table 6d. Concordance between manual and LLM coding, 37 acts

Field	Simple agreement	Cohen's kappa	Notes
Category (7 classes)	97.3% (36/37)	0.96	One disagreement: Decree-Law 4/2022 coded PROROGA instead of NORMA
Publication date	94.6% exact; 100% within one day	—	Two one-day differences (Decree-Law 42/2026; signature of the 2026 implementing decree)
Intensity (1–3)	81.1% exact; 100% within one level	0.75 (linear weights); Spearman 0.78	Model rates platform openings and decree publications one level lower
Expected sign (+1, 0, -1)	78.4% (29/37)	0.59	Eight disagreements, see text

Category and date are recovered almost perfectly; intensity is recovered up to one level. Sign is where the two codings differ, and the disagreements are informative. Five of the eight concern acts that the manual coding marks as restrictive because they omit an expected measure (the 2024 budget bill and law with no refinancing of 4.0, the 2025 budget bill with the 4.0 ceiling, the 2022 budget law with reduced rates) or as favourable because they complete an administrative step (the Revenue Agency resolution on the tax code for the waiting-list credit): the model reads them as neutral. Two are contested acts with mixed content: Decree-Law 38/2026, negative for the users in the waiting list (35% of the credit) and positive for suppliers and non-EU sourcing (removal of the EU-origin clause), and

the directorial decree of June 2025, which the manual coding reads as rationing (ceiling and chronological order) and the model as the opening of a booking window. The last is Law 160/2019, ambiguous in the manual coding and favourable for the model. The pattern is that the manual coder, an insider, prices disappointed expectations; the model prices the letter of the act.

Table 6e. Event study under alternative codings (signed CAAR %, mean per event; p; p BH)

Coding	[-5, -1]	[0]	[+1, +5]	[0, +5]
Manual coding: 30 directional events	-0.39 (0.42; 0.61)	+0.29 (0.08; 0.33)	+0.09 (0.83; 0.83)	+0.39 (0.46; 0.61)
LLM coding: 26 directional events	-0.17 (0.74; 0.79)	+0.09 (0.61; 0.79)	-0.21 (0.56; 0.79)	-0.12 (0.79; 0.79)
LLM coding, intensity 2–3: 15 events	-0.64 (0.37; 0.50)	+0.06 (0.79; 0.79)	-0.73 (0.07; 0.29)	-0.67 (0.18; 0.35)
Consensus: 23 events with the same non-zero sign	-0.42 (0.42; 0.64)	+0.14 (0.48; 0.64)	-0.30 (0.46; 0.64)	-0.15 (0.76; 0.76)

The negative result does not depend on who codes the acts. Under the LLM coding the publication-day reaction disappears (+0.09%, $p = 0.61$) and every window is indistinguishable from zero; on the consensus subset the picture is the same. The residual day-0 estimate of the manual coding comes from the acts on which the two coders disagree, that is from the sign the insider assigns to omissions and to mixed acts. Whichever reading is right, no coding produces a post-event abnormal return.

7.7 Overlapping events

Five pairs of events fall within five trading days of each other, so that the post-event window of the first overlaps the event window of the second: the Council of Ministers approval and the publication of Decree-Law 19/2024 (26 February and 2 March 2024); the publication of the Transizione 5.0 implementing decree and the opening of the GSE platform (6 and 7 August 2024); Decree-Laws 38/2026 and 42/2026 (27 March and 3 April 2026, of opposite sign); the GSE notice to the waiting-list firms and the signature of the Nuova Transizione 5.0 implementing decree (29 April and 4 May 2026); the publication of that decree and the opening of the NPTR5 platform (11 and 12 June 2026). Three treatments are applied: keeping only the first event of each cluster; averaging the CARs of the events in a cluster and treating the cluster as one observation; merging each same-sign cluster into a single event whose day-0 window runs from the first to the last event date, with [-5, -1] before the first and [+1, +5] after the last (the opposite-sign cluster of March–April 2026 is excluded).

Table 6c. Treatment of overlapping events (signed CAAR %, mean per observation; p; p BH)

Treatment	[-5, -1]	[0]	[+1, +5]	[0, +5]
No treatment: 30 events	-0.39 (0.42; 0.61)	+0.29 (0.08; 0.33)	+0.09 (0.83; 0.83)	+0.39 (0.46; 0.61)
First event of each cluster: 25 observations	-0.20 (0.72; 0.72)	+0.30 (0.11; 0.44)	+0.33 (0.50; 0.67)	+0.63 (0.29; 0.59)
Cluster mean: 25 observations	-0.15 (0.77; 0.77)	+0.32 (0.08; 0.34)	+0.28 (0.56; 0.75)	+0.60 (0.31; 0.61)
Merged windows, same-sign clusters: 24 observations	+0.03 (0.96; 0.96)	+0.20 (0.39; 0.90)	+0.21 (0.67; 0.90)	+0.42 (0.52; 0.90)

None of the treatments produces a post-event abnormal return. With merged windows the day-0 estimate falls to +0.20% ($p = 0.39$): when the market is given the whole span between the first and the last act of a cluster, the reaction attributable to the cluster is not distinguishable from zero. The clustering also removes the inflation of the sample by administratively distinct but economically identical acts (signature and publication of a decree, publication and platform opening).

7.8 Survivorship control

The basket contains only stocks listed throughout the period. Prima Industrie, a supplier of eligible laser and sheet-metal machinery with a business profile close to Biesse and El.En., was delisted on 7 March 2023 after the tender offer by Femto Technologies; it is the only exit from the candidate set during the study period. Its abnormal returns were estimated with the same market model for the five

directional events that precede the pinning of its price to the offer consideration (events 2, 4, 5, 33, 34; December 2020 to May 2022) and compared with the mean of the eight surviving stocks on the same events.

Table 6f. Prima Industrie versus the eight surviving stocks, signed CAR %, five events 2020–2022

Event	Beta	PRI [-5, -1]	PRI [0]	PRI [+1, +5]	PRI [0, +5]	8 stocks [0, +5]
2. Law 178/2020 (2021 budget)	0.67	-4.02	+0.99	+0.84	+1.83	-2.37
4. Law 234/2021 (2022 budget)	1.06	+2.33	-1.30	+3.41	+2.11	+9.13
5. Decree-Law 50/2022	1.22	-3.12	-2.11	+0.56	-1.56	+1.48
33. Decree-Law 73/2021	0.83	+6.29	-2.27	+3.83	+1.56	+4.52
34. Decree-Law 4/2022	1.07	+2.29	+3.78	-6.38	-2.60	-2.71
Mean		+0.75	-0.18	+0.45	+0.27	+2.01

The delisted stock reacts less than the survivors: mean signed CAR on [0, +5] of +0.27% against +2.01%, a difference that is not significant on five observations (paired t-test $p = 0.40$; Wilcoxon $p = 0.63$). Adding Prima Industrie as a ninth stock on these five events (245 pairs) leaves the corpus estimates unchanged: [-5, -1] -0.38%, [0] +0.29% ($p = 0.07$), [+1, +5] +0.06%, [0, +5] +0.35% ($p = 0.47$). Within tier B on the same five events, the mean CAR on [0, +5] moves from +1.82% (four stocks) to +1.51% (five stocks). Survivorship, if it biases the sample at all, biases it towards finding a reaction, not against it; the negative result does not depend on it.

7.9 Tests of the endogeneity hypothesis (H4)

Machine-tool orders, annual changes. Each directional event was matched to the last UCIMU reading released before its announcement date. Of 21 directional events since September 2023, 15 are concordant with H4 (one-sided binomial test $p = 0.039$). The mean change in domestic orders before expansionary acts is -10.3%, before restrictive acts +1.7%. Tests on intensity (Spearman 0.27, $p = 0.23$; Mann-Whitney $p = 0.24$) are not significant. The clustering correction removes the result: several closely spaced events share the same reading (four events on 2026Q1, three on 2023Q4). Reduced to the 13 distinct reading-sign combinations, 8 are concordant and significance disappears ($p \approx 0.29$). The discordant cases are instructive: the 2026 budget law (Nuova Transizione 5.0, expansionary) arrives after two quarters of strong growth in domestic orders, but the growth is a base effect on absolute levels among the lowest of the series (index 54 against a 2021 mean of 100); Decree-Law 39/2024 (restrictive) follows the worst quarter of the series, but it is a spending-control act, not an industrial-policy act.

Machine-tool orders, levels, and classification of acts. Two corrections were applied. First, the state of the sector is measured by the level of the domestic-orders index (2021 = 100), seasonally adjusted with a four-quarter moving average. Second, acts are classified as industrial policy (PI) or spending control (CS). The picture changes. The moving average of the level oscillates between 36 and 56 over the whole period: the domestic market never returned above half of its 2021 level. With the threshold at the period median, concordance falls to 10 events out of 20 ($p = 0.59$) and to 6 out of 12 after clustering; the mean level before expansionary acts (45.3) and restrictive acts (48.7) does not differ (Mann-Whitney $p = 0.18$). The result on annual changes was an artefact of base effects on very low levels. The classification, instead, produces a complete separation: the 14 industrial-policy acts are all expansionary, the 6 restrictive acts are all spending-control acts. The direction of an act is determined by its administrative nature, not by the state of the sector.

Official Eurostat series, whole period. The Eurostat series allows the test to be extended to all 25 directional events with a reference date between 2020 and 2026.

Table 7. Endogeneity test on Eurostat C28 indicators, 25 directional events

Indicator	Concordant	p binomial	After clustering	Mean, expansionary / restrictive	Mann-Whitney p
Domestic turnover, 12-month change	12/25	0.66	9/19	+3.1 / +4.3	0.51
Production, 12-month change	13/25	0.50	10/19	-0.6 / +1.0	0.59
Turnover, 3-month vs 12-month average	11/25	0.79	8/19	+1.1 / +0.9	0.76

The logit of the direction of the act on the change in domestic turnover has a coefficient of -0.016 ($p = 0.76$). No specification and no subset (industrial-policy acts only; events from 2023 only) shows an association. H4 is rejected also with official data over the whole period.

A descriptive pattern remains visible: the six Transizione 5.0 acts of 2024 (Decree-Law 19, conversion law, implementing decree, platform opening) were adopted during the only period of falling C28 domestic turnover (between -3.6% and -8.3%), while the expansionary acts of 2025–2026 (Nuova Transizione 5.0, its implementing decree, the NPTR5 platform) arrived with domestic turnover growing by 8–11%. With 25 events the pattern does not pass the statistical threshold and is reported as an observation, not as a result.

7.10 Machine tools versus the machinery industry

The comparison between the UCIMU and Eurostat series adds a finding of its own. The two series are not correlated in 2023–2026 (Spearman 0.10 on quarterly levels). Domestic turnover of the C28 industry in 2023–2026 lies between 107 and 124 (2021 = 100), while UCIMU domestic machine-tool orders oscillate between 13 and 95. The reading «persistently depressed sector» holds for the machine-tool segment and not for the machinery industry as a whole, which in nominal terms has grown. The divergence is itself a result: capital-goods incentives address a segment whose domestic demand has contracted while the industry that produces it has kept invoicing, mainly abroad.

8. Discussion and limits

The result is negative with respect to the starting thesis: on incentive-policy acts no incorporation lag exploitable with an LLM emerges. Three explanations were distinguished and tested where possible.

8. The signal does not exist: capital-goods incentives are a modest fraction of the value of the issuers and the market ignores them. The day-0 reaction of +0.29% on the final corpus, not significant after correction, not supported by the rank test or by the cross-correlation-adjusted test, and absent under the LLM coding, indicates that the content is at most marginally priced, and only within the publication session.
9. The signal exists but is anticipated: the relevant date is the political announcement (Council of Ministers, amendment, ministerial statement), not the Official Journal. Tested in Section 7.3 with a negative outcome: the announcement window shows no reaction in the expected direction either.
10. The signal exists but is covered by noise: 20 to 31 directional events are few, the period includes the pandemic shock, the 2022 rotation and the 2026 Iran war, and the selected stocks have heterogeneous betas and liquidity. The power analysis bounds this explanation: any post-event effect, if present, is below about 1.5% on $[0, +5]$.

The endogeneity hypothesis, which would have changed the nature of the paper from a study of a market signal to a study of the relation between the sector cycle and the timing of incentive policy, does not hold: the direction of acts does not depend on the state of the sector, measured in any of the ways tried (annual changes and levels of machine-tool orders; changes and deviations of official domestic

turnover and production). The pre-announcement signal of Section 7.3 therefore remains unexplained and is treated as a possible false positive on 9 events.

The reading the data support is more modest and verifiable: between 2019 and 2026 neither the publication nor the announcement of incentive-policy acts produced abnormal returns on the exposed stocks, and the sequence of acts is explained by public-finance constraints rather than by the sector cycle. In 2023–2026 incentive policy did not follow the cycle: it alternated expansionary announcements and accounting tightenings on a domestic machine-tool market that remained persistently depressed, while the machinery industry grew abroad. The separation between industrial-policy acts and spending-control acts is the key to reading that sequence.

Residual limits: Italian revenue share not disclosed by two issuers (Biesse, Sabaf) and below threshold for two others (Wiit, Datalogic), so that the two-tier design rests on four verified stocks; sample of 30 directional events (24 to 25 after clustering) with a minimum detectable effect of 1.5%; the UCIMU series reconstructed from press releases rather than obtained from the official source; the survivorship control limited to one delisted stock and five events, because the tender offer pinned its price from July 2022; the LLM coding obtained from a single run of a single model, with seven acts coded from the model's prior knowledge rather than from a fetched source, and no second human coder.

Operating consequence: no directional backtest is run, neither on the publication date nor on the announcement date. Phases 2 and 3 of the work plan are closed without execution, in accordance with the decision rule stated in Section 3.

9. Conclusions and further research

The study set out to measure whether an LLM reading Italian incentive legislation could anticipate the prices of the exposed stocks. The LLM reads the legislation well: it recovers category and date almost perfectly and intensity within one level, and it diverges from the expert coder only where the expert prices what the act fails to say. But there is nothing to anticipate on this sample: whatever information content the acts have is priced, at most marginally, within the publication session; the highest-intensity acts are priced before publication; no lag is left to exploit. The contribution is threefold. First, a documented negative result on a niche of unstructured open data not covered by the financial text-mining literature, with windows fixed *ex ante*, event-level inference, multiple-testing correction, non-parametric and cross-correlation-adjusted confirmation, a treatment of overlapping events, a survivorship control, an independent LLM coding and a power bound. Second, the rejection, on two independent data sources, of the hypothesis that Italian capital-goods incentive policy is endogenous to the sector cycle. Third, the classification of acts into industrial policy and spending control, which separates the sample completely and explains the direction of each act.

Three lines of further work follow from the limits. The corpus can be extended towards the threshold of 50 events with Revenue Agency rulings, GSE resource-counter notices and Nuova Sabatini decrees, and the announcement dates still open can be verified. The basket can be completed with the geographical verification of the two issuers that do not disclose an Italian revenue share. Finally, the same event-study apparatus is being applied, in a companion study, to the United States market for critical minerals, where the hypothesis is reversed: not whether the market reacts to presidential communications, but whether exposed stocks move before the first public appearance of the information. That extension tests pre-event run-up, abnormal volume and run-up ratios and will be reported separately.

Abbreviations and Italian terms

Acts are cited as in Italian practice: «Law 199/2025» is Law no. 199 of 2025; «Decree-Law 19/2024» a decree-law (decreto-legge) later converted into law; «DM» a ministerial or interministerial decree (decreto ministeriale); «DD» a directorial decree (decreto direttoriale) of a ministry department; «c.» a paragraph (comma) of an article.

Term	Meaning
AR, CAR, CAAR	Abnormal return; cumulative abnormal return of one event-stock pair; cumulative average abnormal return across pairs or events
BH	Benjamini-Hochberg false-discovery-rate correction for multiple tests
CdM	Consiglio dei ministri, the Italian Council of Ministers; its approval of a bill or decree-law precedes publication
CS / PI	Spending-control act / industrial-policy act, the classification of acts introduced in Section 6.4
Esodati 5.0	Firms whose Transizione 5.0 applications were left in the unfunded waiting list after the exhaustion of resources on 7 November 2025
FTSE MIB	Main index of Euronext Milan, used as the market factor
GSE	Gestore dei Servizi Energetici, the state-owned company that manages the Transizione 5.0 and Nuova Transizione 5.0 application platforms
GU	Gazzetta Ufficiale, the Italian Official Journal
IDEM	Italian derivatives market of Euronext Milan
LLM	Large language model
MDE	Minimum detectable effect at 5% significance and 80% power
MEF	Ministero dell'economia e delle finanze, the Italian Ministry of Economy and Finance
MIMIT	Ministero delle imprese e del made in Italy, the Italian Ministry of Enterprises and Made in Italy
NACE C28	Statistical classification of economic activities, division 28: manufacture of machinery and equipment
NPTR5	GSE platform for the Nuova Transizione 5.0 communications, opened on 12 June 2026
Nuova Transizione 5.0	Depreciation surcharge for capital goods (180%/100%/50%) introduced by Law 199/2025, art. 1 c. 427-436, for investments from 1 January 2026 to 30 September 2028
PNRR	Piano nazionale di ripresa e resilienza, the Italian National Recovery and Resilience Plan
Revenue Agency	Agenzia delle Entrate, the Italian tax authority; issues circulars (circolari), resolutions (risoluzioni) and rulings (interpelli)
Transizione 4.0	Tax credit for tangible and intangible capital goods introduced by Law 160/2019, successor of hyper-depreciation
Transizione 5.0	Tax credit for investments with energy-saving targets introduced by Decree-Law 19/2024, art. 38
UCIMU	UCIMU-Sistemi per Produrre, the Italian machine-tool builders' association; its Study Centre publishes the quarterly order index

References

- Bailey D.H., Borwein J., López de Prado M., Zhu Q.J. (2014), «Pseudo-mathematics and financial charlatanism: the effects of backtest overfitting on out-of-sample performance», *Notices of the AMS*, 61(5), 458–471.
- Banca d'Italia, Ministero dell'economia e delle finanze (2026), Piano Transizione 4.0: una valutazione degli effetti su investimenti, occupazione e produttività. Rapporto finale, 14 May 2026.
- Binder J.J. (1985), «Measuring the effects of regulation with stock price data», *RAND Journal of Economics*, 16(2), 167–183.
- Boehmer E., Musumeci J., Poulsen A.B. (1991), «Event-study methodology under conditions of event-induced variance», *Journal of Financial Economics*, 30(2), 253–272.

- Bratta B., Romano L., Acciari P., Mazzolari F. (2023), «Assessing the impact of digital technology diffusion policies. Evidence from Italy», *Economics of Innovation and New Technology*, 32(8), 1114–1137.
- Campbell C.J., Wasley C.E. (1993), «Measuring security price performance using daily NASDAQ returns», *Journal of Financial Economics*, 33(1), 73–92.
- Campbell J.Y., Lo A.W., MacKinlay A.C. (1997), *The econometrics of financial markets*, Princeton University Press.
- Cefis E., Scrofani S., Tubiana M. (2025), «Innovation policies and firms' productivity: the Italian Industry 4.0 Plan for digital transformation», *Journal of Technology Transfer*, 50(5), 2443–2475.
- Comitato scientifico per la valutazione dell'impatto economico degli interventi del Piano Transizione 4.0 (2024), *Gli incentivi in investimenti 4.0: una valutazione dell'impatto della misura. Rapporto intermedio*, MEF, MIMIT, Banca d'Italia, 29 November 2024.
- Corrado C.J. (1989), «A nonparametric test for abnormal security-price performance in event studies», *Journal of Financial Economics*, 23(2), 385–395.
- Fama E.F. (1970), «Efficient capital markets: a review of theory and empirical work», *Journal of Finance*, 25(2), 383–417.
- Kim A.G., Muhn M., Nikolaev V.V. (2023), «From transcripts to insights: uncovering corporate risks using generative AI», *Becker Friedman Institute Working Paper 2023-132*, University of Chicago.
- Kolari J.W., Pynnönen S. (2010), «Event study testing with cross-sectional correlation of abnormal returns», *Review of Financial Studies*, 23(11), 3996–4025.
- Kothari S.P., Warner J.B. (2007), «Econometrics of event studies», in Eckbo B.E. (ed.), *Handbook of Empirical Corporate Finance*, vol. 1, Elsevier, 3–36.
- Lamdin D.J. (2001), «Implementing and interpreting event studies of regulatory changes», *Journal of Economics and Business*, 53(2–3), 171–183.
- López de Prado M. (2018), *Advances in financial machine learning*, Wiley.
- Lopez-Lira A., Tang Y. (2023), «Can ChatGPT forecast stock price movements? Return predictability and large language models», SSRN 4412788; *Journal of Financial Economics*, forthcoming.
- Loughran T., McDonald B. (2011), «When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks», *Journal of Finance*, 66(1), 35–65.
- MacKinlay A.C. (1997), «Event studies in economics and finance», *Journal of Economic Literature*, 35(1), 13–39.
- Tetlock P.C. (2007), «Giving content to investor sentiment: the role of media in the stock market», *Journal of Finance*, 62(3), 1139–1168.

Appendix A. Event corpus

Sign: expected sign on the basket (+1, 0, -1). Int.: intensity assigned ex ante (1 low, 2 medium, 3 high). Status: V = date verified on a primary source; V (press) = verified on the government press-release register and specialised press; R = recoded in dataset v0.9 and re-estimated at the corrected date; Excl. = excluded. Events with sign 0 are excluded from directional tests; event 32 is excluded for lack of a price on day 0; event 18 is excluded because the act does not exist. Source: Dataset_eventi_paniere_v0.12.xlsx, sheet Eventi.

Id	Date	Category	Act and content	Sign	Int.	Status
1	30/12/2019	NORMA	Law 160/2019 (2020 budget), art. 1 c. 184-197: 4.0 tax credit replaces hyper- and super-depreciation	0	2	V
2	30/12/2020	NORMA	Law 178/2020 (2021 budget): 4.0 credit raised to 50% for 2021-2022, PNRR funding (13.4 bn)	1	3	V
3	23/07/2021	PRASSI	Revenue Agency circular 9/E/2021: clarifications on the 4.0 credit	0	1	V
4	31/12/2021	NORMA	Law 234/2021 (2022 budget): 4.0 extended to 2023-2025 with reduced rates	-1	2	V
5	17/05/2022	NORMA	Decree-Law 50/2022, art. 21: intangible 4.0 assets at 50% for 2022	1	1	V
6	29/12/2022	PROROGA	Law 197/2022 (2023 budget): delivery deadline extended to 30/09/2023	1	1	V
7	27/02/2023	PROROGA	Law 14/2023 (conversion of Decree-Law 198/2022), art. 12 c. 1-ter: further extension to 30/11/2023 (recoded in v0.9; was Decree-Law 132/2023 of 29/09/2023)	1	1	V, R
8	26/02/2024	CDM	Council of Ministers approves the PNRR decree with Transizione 5.0 (6.3 bn)	1	3	V
9	02/03/2024	NORMA	Decree-Law 19/2024, art. 38: Transizione 5.0 tax credit instituted	1	3	V
10	29/03/2024	RESTRIZIONE	Decree-Law 39/2024, art. 6: mandatory prior and final communications for 4.0; offsetting suspended	-1	2	V
11	29/04/2024	NORMA	Law 56/2024 (conversion of Decree-Law 19/2024)	1	1	V
12	24/07/2024	DM	MIMIT-MEF decree of 24/07/2024 (signature): Transizione 5.0 implementing decree	1	2	V
13	06/08/2024	DM	Decree of 24/07/2024 published in the Official Journal	1	2	V
14	07/08/2024	PIATTAFORMA	GSE Transizione 5.0 platform opened	1	2	V
15	16/08/2024	PRASSI	MIMIT operating circular of 16/08/2024	0	1	V
16	15/10/2024	CDM	Council of Ministers no. 100, 2025 budget bill: 4.0 ceiling, software excluded, T5.0 simplifications	-1	2	V (press)
17	31/12/2024	NORMA	Law 207/2024 (2025 budget): T5.0 simplifications; 4.0 ceiling 2.2 bn; intangibles excluded	0	2	V
18	24/04/2025	DM	MIMIT-MEF decree of 24/04/2025: amendments to the T5.0 implementing decree (act does not exist; the L. 207/2024 changes were applied through FAQ revisions; excluded)	1	1	Excl.
19	16/06/2025	PIATTAFORMA	MIMIT directorial decree of 16/06/2025: platform for 2025 4.0 communications, 2.2 bn ceiling, chronological priority, platform open from 17/06/2025 (Law 207/2024 c. 445-448; recoded in v0.9; was Decree-Law 95/2025 of 30/06/2025)	-1	1	V, R
20	17/10/2025	CDM	Council of Ministers no. 146, 2026 budget bill: Nuova Transizione 5.0 announced (4 bn)	1	3	V
21	07/11/2025	PIATTAFORMA	Exhaustion of Transizione 5.0 resources (MIMIT decree of 06/11/2025): unfunded waiting list	-1	2	V
22	21/11/2025	NORMA	Decree-Law 175/2025: mandatory option between T5.0 and 4.0; waiting-list rules	-1	1	V

Id	Date	Category	Act and content	Sign	Int.	Status
23	20/01/2026	NORMA	Law 4/2026 (conversion of Decree-Law 175/2025)	0	1	V
24	30/12/2025	NORMA	Law 199/2025 (2026 budget), art. 1 c. 427-436: Nuova Transizione 5.0 (180%/100%/50% surcharge)	1	3	V
25	27/03/2026	NORMA	Decree-Law 38/2026, art. 7-8: 35% of the credit to the waiting list; EU constraint removed	-1	2	V
26	03/04/2026	NORMA	Decree-Law 42/2026: waiting-list credit restored to 89.77% (1,302.3 m) plus MIMIT FER contribution	1	2	V
27	04/05/2026	DM	Nuova Transizione 5.0 implementing decree (DM IPER) signed by MIMIT	1	2	V
28	22/05/2026	NORMA	Law 88/2026 (conversion of Decree-Law 38/2026): waiting-list credit confirmed	1	1	V
29	12/06/2026	PIATTAFORMA	GSE NPTR5 platform opened	1	2	V
30	11/06/2026	DM	MIMIT-MEF decree of 07/05/2026 published on the MIMIT website (no Official Journal publication; recoded in v0.9; was 07/05/2026)	1	2	V, R
31	30/04/2019	NORMA	Decree-Law 34/2019, art. 1: 130% super-depreciation reintroduced for 2019	1	2	V
32	19/05/2020	PROROGA	Decree-Law 34/2020, art. 50: delivery deadline extended to 31/12/2020 (excluded: no price on day 0)	1	1	V
33	25/05/2021	NORMA	Decree-Law 73/2021, art. 20: 4.0 credit on ordinary assets usable in one year	1	1	V
34	27/01/2022	NORMA	Decree-Law 4/2022, art. 10: 5% credit on tangible 4.0 PNRR green investments above 10 m (content corrected in v0.9)	1	1	V
35	16/10/2023	CDM	Council of Ministers no. 54, 2024 budget bill: no measure for capital goods	-1	1	V
36	30/12/2023	NORMA	Law 213/2023 (2024 budget): no 4.0 refinancing	-1	1	V
37	16/04/2026	PRASSI	Revenue Agency resolution 14/E/2026: tax code for the waiting-list credit	1	1	V
38	29/04/2026	PIATTAFORMA	GSE notice of 29/04/2026: confirmations to waiting-list firms	1	1	V

Appendix B. Work log

Chronological record of activities, decisions and outcomes.

Date	Phase	Activities and decisions	Outcome
9 Sep 2026	0	Preliminary assessment of the thesis. Decisions: Euronext Milan; basket of industrial machinery and automation; spot equities in the study phase, derivatives only after verification of H1 and H2; LLM limited to structured extraction; market-model event study and walk-forward backtest.	Research design set. Document v0.1.
9 Sep 2026	1	Event corpus built (28 records, 7 categories, sign and intensity assigned ex ante), basket (10 stocks plus 1 delisted, 2 benchmarks) and script event_study.py. Test on a synthetic series. Decision: multiple-testing correction and windows fixed ex ante.	Phase 1 started. Files: Dataset v0.1, event_study.py.
9 Sep 2026	1	Decisions: two-tier basket; Confindustria and UCIMU among the sources. Price file received (yfinance, 8 stocks plus FTSE MIB, 2018–2026). Web verification of 7 dates; dataset v0.2 with 30 events (25 verified). Event study with event-level inference and BH correction.	Results compiled. H1, H2, H3 not verified. Phase 2 redirected to political-announcement dates.

Date	Phase	Activities and decisions	Outcome
9 Sep 2026	2	Ten Council of Ministers dates reconstructed. Event study on announcement versus publication window, 9 directional events, paired test. Signed pre-announcement CAAR -1.59% (p 0.003), exploratory.	Anticipation thesis not confirmed. New hypothesis: endogeneity of incentive policy. Files: Dataset v0.3, car_annuncio_vs_publicazione.csv.
9 Sep 2026	3	UCIMU domestic-orders series 2022Q4–2026Q2 reconstructed from press releases. Event matched to last reading; binomial, Spearman, Mann-Whitney tests; clustering correction.	15/21 concordant (p 0.039), 8/13 after clustering (n.s.). Hypothesis compatible, not confirmed. Files: Dataset v0.4, ucimu_interno.csv, endogeneita_eventi_ucimu.csv.
9 Sep 2026	4	UCIMU levels 2023Q1–2026Q2 reconstructed (2021 base), four-quarter moving average. Acts classified PI/CS. Concordance test repeated on levels.	Concordance 10/20, n.s.; result on changes was a base effect. Complete separation PI expansionary / CS restrictive. Cyclical hypothesis rejected. Files: Dataset v0.5, ucimulivelli.csv, endogeneita_livelli.csv.
10 Sep 2026	5	Eurostat API not reachable from the analysis environment; series sts_intv_m and sts_inpr_m (C28, Italy, 2015–2026) downloaded locally and uploaded. Endogeneity test on 25 events with three indicators, logit, subsets. UCIMU versus C28 comparison. Council of Ministers date of Decree-Law 175/2025 verified.	No association (12/25, logit p 0.76). Cyclical hypothesis rejected with official data. UCIMU and C28 uncorrelated. Files: Dataset v0.6, eurostat_c28_elaborato.csv, endogeneita_eurostat.csv.
10 Sep 2026	6	Robustness: alternative estimation windows, exclusion of Wiit, subsets A/B; Corrado rank test. Italian share verified for Wiit (34.9%) and Zignago Vetro (69.5%) on FY2025 statements.	Results stable; Corrado day 0 p 0.08. Wiit reclassified. File: Dataset v0.7.
10 Sep 2026	7	Corpus extended to 38 events (8 new), event 6 corrected, Council of Ministers dates 2021 and 2022 verified. Event study on 31 directional events; power analysis.	Day 0 +0.40% (p 0.02, BH 0.07); post-event nil; minimum detectable effect 1.4%. Files: Dataset v0.8, car_corpus_esteso.csv.
10 Sep 2026	8	Second verification round: Council of Ministers dates (events 16, 20, 35; announcements 10, 17, 36) confirmed; events 7, 19, 30 found miscoded and recoded; event 34 content corrected; event 18 not found. Italian revenue share verified on FY2025 reports for Sesa, Tinexta, Datalogic, El.En. (Biesse and Sabaf do not disclose it). Sensitivity of Table 6 to the corrections computed from the CAR file.	Table 6b; reading unchanged. Dataset v0.9 (recoded rows flagged for re-run). Cover: independent researcher with ORCID; JEL codes confirmed.
10 Sep 2026	9	Price series received (prezzi.csv, 2,100 observations to 8 September 2026). Event study re-implemented and validated against the v0.8 CAR file (log returns, market model on [-250, -21]); CARs re-estimated for events 7, 19, 30 at the corrected dates; Table 6 rebuilt; power analysis and intensity correlation updated.	Day 0 +0.31% (p 0.06, BH 0.25), post-event nil; the three recoded events have negative signed CARs. Files: car_corpus_esteso_v0.9.csv, Dataset v0.9.
10 Sep 2026	10	Robustness checks and Corrado rank test repeated on the extended corpus after recoding (31 events, 248 pairs): alternative estimation windows, exclusion of Wiit and Datalogic, tiers A and B.	Table 5 rebuilt on the final sample. Results stable; rank test day 0 z 1.35 (p 0.18).
10 Sep 2026	11	Third verification round: DD MIMIT 15 May and 16 June 2025 confirmed (platform open 17 June 2025); DM 24 April 2025 does not exist (event 18 dropped); L. 14/2023 votes 15 and 22 February 2023; Prima Industrie delisted 7 March 2023, prices available from vendors. Final corpus of 30 directional events; Tables 5 and 6 rebuilt; overlapping events treated (Table 6c); abbreviations section added.	Day 0 +0.29% (p 0.08, BH 0.33; Corrado p 0.24); post-event +0.09%; merged clusters day 0 +0.20% (p 0.39). Files: car_corpus_finale_v1.0.csv, Dataset v0.10.
10 Sep 2026	12	Prima Industrie daily series received (yfinance, 1 June 2018 to 3 March 2023). Price pinned to the offer from 12 July 2022; control run on the five directional events 2020–2022; corpus re-estimated with nine stocks on those events.	Table 6d: PRI +0.27% vs survivors +2.01% on [0, +5], n.s.; corpus estimates unchanged. Files: car_corpus_finale_v1.1.csv, Dataset v0.11.

Date	Phase	Activities and decisions	Outcome
10 Sep 2026	13	Referee-oriented revision: blind LLM coding of the 37 acts (Claude, September 2026) and concordance with the manual coding; event study under LLM and consensus coding; BMP and Kolari-Pynnönen tests; literature on regulatory event studies and on the evaluation of the Italian incentives; pre-registration wording replaced by «fixed ex ante»; conflict-of-interest disclosure on the cover.	Category kappa 0.96, sign kappa 0.59, intensity weighted kappa 0.75; under LLM coding day 0 +0.09% (p 0.61); KP-adjusted day 0 p 0.14. Dataset v0.12 (sheet Codifica_LLM).
10 Sep 2026	4 (final)	Final write-up in English. Restructuring as a paper; figures of Table 6 recomputed from car_corpus_esteso.csv and Dataset v0.8 (the third row is identified as the 14 verified-date events with intensity 2–3); H4 stated explicitly as post hoc; Corrado (1989) added to the references; event corpus listed in Appendix A.	Version 1.0.

Appendix C. Version history

Version	Date	Changes
0.1	9 Sep 2026	First draft: thesis, framework, hypotheses, design, methodology, work plan.
0.2	9 Sep 2026	Phase 1: event corpus, basket, script, synthetic test, constraints.
0.3	9 Sep 2026	Two-tier basket; Confindustria and UCIMU sources; date verification; preliminary results and discussion.
0.4	9 Sep 2026	Political-announcement window; endogeneity hypothesis; reorientation of the paper.
0.5	9 Sep 2026	Verification with UCIMU data; discussion updated.
0.6	9 Sep 2026	Levels and PI/CS classification; thesis redefined.
0.7	10 Sep 2026	Verification with the Eurostat C28 series over the whole period; UCIMU/C28 comparison.
0.8	10 Sep 2026	Robustness and rank test; Italian share of two stocks; limits updated.
0.9	10 Sep 2026	Corpus extended to 38 events; power analysis; limits updated.
1.0	10 Sep 2026	English version restructured as a paper: abstract, introduction, H4 stated as post hoc, results consolidated in Section 7, conclusions, event corpus in Appendix A. Figures unchanged and recomputed from the CAR file.
1.1	10 Sep 2026	Cover: independent researcher, ORCID. Second verification round (Sections 5.1, 5.2); events 7, 19, 30 recoded, event 34 corrected; Table 1 with verified revenue shares; Table 6b sensitivity; limits updated; dataset v0.9.
1.2	10 Sep 2026	Event study re-estimated at the corrected dates on the price series; Table 6 rebuilt (Table 6b keeps the pre-recoding figures); abstract, Sections 7.5, 8 and 9 and limits updated.
1.3	10 Sep 2026	Table 5 (robustness and rank test) rebuilt on the extended corpus; Campbell and Wasley (1993) added to the references.
1.4	10 Sep 2026	Open points closed: event 18 dropped, event 19 and all announcement dates verified; final corpus of 30 events (Tables 5, 6); overlapping events (Section 7.6, Table 6c); abbreviations section; limits reduced to revenue-share disclosure, sample size, UCIMU source and the Prima Industrie control.
1.5	10 Sep 2026	Survivorship control executed with Prima Industrie (Section 7.7, Table 6d, since renumbered); Table 1 and Sections 4.2, 5.3, 8, 9 updated. All points of the work plan closed.
2.0	10 Sep 2026	Version for submission: LLM coding and concordance (Sections 6.5, 7.6, Tables 6d–6e); BMP and Kolari-Pynnönen tests (Table 5); Section 2.5 and eleven references added; «pre-registered» replaced by «fixed ex ante»; disclosure on the cover; sections renumbered.
2.1	10 Sep 2026	Title aligned with content; abstract rewritten for SSRN; disclosure completed (not commissioned, reviewed or endorsed by any institution); Section 4.3 restated as spot equities only.