

Politecnico di Torino

Politecnico di Torino

Department of Management and Production Engineering

Master's Degree in Engineering and Management

The Effects of a Geopolitical Shock on Stock Market Performance in the Defence Sector: A Difference-in-Differences and Synthetic Control Analysis

Evidence from the Russian Invasion of Ukraine

Supervisor:
Prof. Francesco Nicoli

Candidate:
Giacomo De Salve

Academic Year 2025-2026

September 2026

Abstract

Russia's invasion of Ukraine generated a geopolitical shock that altered expectations concerning military expenditure, procurement and the economic position of the defence industry. This thesis examines its total effect on the stock-market performance of listed defence producers relative to firms without a direct conflict-specific military-demand channel. It also investigates geographical heterogeneity, pre-invasion arms-sales exposure and the temporal evolution of the response.

The analysis uses 26,663 monthly firm-level observations from January 2010 to June 2026, covering 36 defence firms and 100 screened controls in Europe and the United States. The empirical strategy combines a two-way fixed-effects Difference-in-Differences design with separate synthetic controls for the aggregates of fourteen European and seventeen United States defence producers.

The Difference-in-Differences analysis estimates an average post-invasion differential of 0.0166 monthly log points. The estimated post-invasion return differential is concentrated in Europe, where it reaches 0.0323, compared with a statistically undetectable estimate of 0.0043 in the United States. The pre-treatment dynamics are inconsistent with the parallel-trends assumption, limiting the causal interpretation of these estimates. The synthetic-control evidence supports a causal interpretation of the relative repricing observed for the aggregated European defence producer within the counterfactual employed. No comparable causal effect is established for the United States aggregate.

The pre-invasion arms-sales share does not systematically explain cross-firm variation in the post-invasion return differential. The return differential remains observable at annual frequency beyond 2022 and becomes statistically undetectable from June 2025 onwards. The findings document a marked transatlantic asymmetry in the financial transmission of the geopolitical shock.

Acknowledgements

I would like to express my deepest gratitude to Professor Francesco Nicoli for supervising this thesis and for his valuable guidance.

My heartfelt thanks go to my family: to my mother, for the love and strength she has always given me, to my sister Bianca, for teaching me courage and determination and to my brother Tommaso, whose presence and affection I have always felt.

I am especially grateful to my girlfriend, Maria Letizia, for standing by my side with love and for her thoughtful advice. I would also like to thank her family for welcoming me so warmly and for the trust they have shown me.

I sincerely thank my aunt and uncle, Marina and Ettore, for their support throughout these years and for always believing in me.

My sincere gratitude also goes to Giancarlo for his invaluable hospitality in Turin over the past two years. Thank you for opening your home to me, for your constant friendship and for never letting me feel alone.

Finally, I dedicate this achievement to the memory of my father and my brother Luigi, who remain forever in my heart.

Table Of Contents

Abstract **1**

1 Introduction **6**

2 Context and Positioning within the Literature **9**

2.1 The Pre-2022 Defence Baseline....9

2.2 The 2022 Geopolitical Shock....11

2.3 From Immediate Response to Long-Term Defence Commitments....12

2.4 From Public Demand to Firm Value....13

2.5 Positioning within the Literature and Research Gap....14

2.5.1 Empirical Evidence on Geopolitical Risk and Defence Firms....14

2.5.2 Research Gap, Contribution and Scope....15

3 Theoretical Framework and Hypothesis Development **18**

3.1 Geopolitical Shocks, Expectations and Firm Valuation....18

3.1.1 Stock Prices as Forward-Looking Valuations....19

3.1.2 Geopolitical Risk, Expected Cash Flows and Discount Rates....19

3.1.3 The Economic Specificity of the Defence Industry....20

3.1.4 Expectations Before Accounting Realisation....22

3.2 Sectoral Exposure and the Average Post-Invasion Differential....23

3.2.1 Broad Market Channels of the Geopolitical Shock....23

3.2.2 The Defence-Specific Demand Channel....24

3.2.3 The Expected Sectoral Differential....25

3.3 Geographical Exposure and Transatlantic Heterogeneity....25

3.3.1 Geographical Proximity and Perceived Security Risk....26

3.3.2 Different Pre-Invasion Defence Equilibria....26

3.3.3 The Geography of the Procurement Response....27

3.3.4 The Expected Transatlantic Differential....28

3.4 Firm-Level Exposure to the Defence Business....29

3.4.1 Sector Membership and Economic Exposure....29

3.4.2 Pre-Invasion Arms-Sales Share as an Exposure Proxy....30

3.4.3 The Expected Firm-Level Gradient....31

3.5 Information Arrival and Persistence....32

3.5.1 The Immediate Incorporation of Geopolitical Information 32
3.5.2 Subsequent Policy Information and Continued Updating 33
3.5.3 The Expected Persistence of the Return Differential 33

4 Research Design **35**

4.1 Overview and Empirical Approach 35
4.2 Data 36
4.2.1 Data Sources 36
4.2.2 Sample Construction 37
4.2.3 Treatment and Control Groups 40
4.2.4 Outcome Variables 42
4.3 Empirical Strategy 43
4.3.1 Baseline Difference-in-Differences Specification 43
4.3.2 Specification for the Heterogeneity across Markets 44
4.3.3 Specification for the Heterogeneity across Firms 46
4.3.4 Identification Assumptions and Parallel Trends 47
4.3.5 Dynamic Specification of the Post-Treatment Effect 49
4.3.6 Robustness Design and Placebo Tests 52
4.3.7 Econometric Considerations 54
4.4 A Complementary Identification Strategy: The Synthetic Control Method 55
4.4.1 Motivation and Method 56
4.4.2 Treated Units and Donor Pool 57
4.4.3 Inference and Stability 58

5 Results **61**

5.1 Average Post-Treatment Return Differential 61
5.1.1 The Baseline Estimate 61
5.1.2 Robustness of the Estimate 63
5.2 Geographical Heterogeneity of the Return Differential 66
5.2.1 The Difference between the Two Markets 66
5.2.2 Robustness of the Estimate 69
5.3 The Intensity of Exposure 71
5.4 The Evolution of the Return Differential over Time 72
5.4.1 The Profile of the Return Differential 73

5.4.2 The Decomposition by Event 76
5.5 The Identifying Assumption 78
5.5.1 The Departure from Parallel Trends 79
5.5.2 Anticipation 83
5.5.3 The Placebo Dates 83
5.6 The Synthetic Counterfactual 85
5.6.1 The Average European Producer 85
5.6.2 The Three Individual Producers 91
5.6.3 The United States Aggregate 92
5.6.4 Robustness, Limitations and Verdict 92
6 Conclusions **95**
6.1 Research Question and Analytical Scope 95
6.1.1 Research Question and Analytical Dimensions 95
6.1.2 Empirical and Geographical Scope 96
6.2 Findings and Answer to the Research Question 97
6.2.1 H1: The Average Post-Invasion Differential 97
6.2.2 H2: Transatlantic Heterogeneity 98
6.2.3 H3: Pre-Invasion Defence Revenue Exposure 99
6.2.4 H4: Persistence of the Return Differential 99
6.2.5 Research Question Answer 100
6.3 Limitations of the Analysis 101
6.3.1 Sample and External Validity 101
6.3.2 Limits to Causal Identification 102
6.3.3 Measurement of Outcomes and Economic Exposure 103
6.3.4 Concurrent Events and Post-Treatment Attribution 104
6.4 Financing and Industrial Implications 106
6.5 Future Research 107
6.6 Final Remarks 108
References **110**
APPENDIX A **114**
APPENDIX B **119**

1 Introduction

Russia's invasion of Ukraine in February 2022 generated a major geopolitical shock with consequences extending across financial markets and industries. The conflict increased uncertainty, disrupted trade and production, and changed expectations concerning economic activity and risk. It also produced a substantial political and military response involving defence budgets, military assistance, procurement programmes, and longer-term security commitments. Defence producers occupied a distinctive position within this response because their products formed part of the capabilities that governments pursued to replenish or expand. The event therefore provides a setting in which to examine how a geopolitical shock is transmitted to the valuation of firms operating in a strategically exposed industry.

Previous research associates geopolitical risk with weaker economic activity, lower investment and changes in financial valuations [1, 2]. Research on the defence industry also shows that military expenditure, procurement decisions and geopolitical events contain information relevant to the valuation of listed producers [3]. Evidence concerning Russia's invasion documents an immediate positive response among major defence firms [4]. These findings leave open the longer-term evolution of the response, its distribution between European and United States producers, its relationship with firm-level exposure to military activity and the scope of its causal attribution. This thesis examines these dimensions within a common empirical framework.

The research question is:

"To what extent did the geopolitical shock generated by Russia's invasion of Ukraine causally affect the stock-market performance of listed defence firms in Europe and the United States relative to firms operating in sectors without a direct conflict-specific military-demand channel?"

The thesis targets the total effect of the geopolitical shock on the relative performance of defence firms. This effect includes the political and military response through which the shock was transmitted to financial markets. The empirical design does not identify the direct effect of the invasion separately from this response. The analysis considers four dimensions of the resulting stock-market differential: its average magnitude, its

distribution between Europe and the United States, its relationship with firms' pre-invasion dependence on arms sales and its temporal evolution.

The empirical analysis uses 26,663 monthly firm-level observations from January 2010 to June 2026. The sample contains 36 defence firms identified through the SIPRI Arms Industry Database and 100 screened control firms drawn from the S&P 500 and the STOXX Europe 600. The geographical perimeter comprises the United States, the member states of the European Union, the United Kingdom and Norway. The control firms remain exposed to the broad macroeconomic and financial consequences of the geopolitical shock. Their defining characteristic is the absence of a direct demand channel arising from military procurement.

The empirical strategy combines two complementary approaches. A two-way fixed-effects Difference-in-Differences design estimates the average post-invasion differential and examines its geographical, firm-level and temporal distribution. Alternative specifications, robustness checks and dynamic estimates assess the stability of the results and the identifying assumptions. The pre-treatment dynamics are inconsistent with the parallel-trends assumption. The Difference-in-Differences estimates consequently document post-invasion differentials but do not independently establish their causal attribution to the geopolitical shock. The synthetic-control analysis constructs separate regional counterfactuals for the aggregates of fourteen European and seventeen United States defence producers with complete price histories.

The baseline Difference-in-Differences estimate identifies a positive average differential of 0.0166 monthly log points, corresponding to approximately 1.67 percentage points in simple-return terms. The main finding is a marked difference between Europe and the United States. The monthly differential reaches 0.0323 log points in Europe, while the corresponding United States estimate is 0.0043 and statistically indistinguishable from zero. The difference between the two geographical estimates is significant at the one per cent level. At the firm level, the analysis finds no systematic association between pre-invasion arms-sales shares and post-invasion return differentials. Regarding its evolution over time, the differential becomes pronounced around the invasion, remains positive at annual frequency beyond 2022 and becomes statistically undetectable from June 2025 onwards.

The synthetic European producer reproduces the twelve-year pre-invasion trajectory of the treated aggregate, including the pandemic period, and separates from its counterfactual in February 2022. The European divergence is larger than that observed for every placebo unit, with a rank-based p-value of 0.0196, and remains stable when the principal donors are removed. Over the fifty-three post-treatment months, the treated series records a cumulative divergence of 119.5 per cent relative to the synthetic counterfactual. This evidence supports a causal interpretation of the European aggregate divergence within the counterfactual and donor pool employed by the design. A comparable causal effect is not established for the United States aggregate.

This thesis contributes to the literature by combining long-horizon firm-level evidence with a constructed regional counterfactual. It distinguishes the estimated post-invasion differential from the narrower scope of the available causal evidence and examines geographical heterogeneity, firm-level exposure and temporal persistence within the same framework. Chapter 2 reviews the relevant literature. Chapter 3 develops the theoretical framework and hypotheses. Chapter 4 presents the data and empirical methodology. Chapter 5 reports the results. Chapter 6 summarises the findings, limitations, implications and directions for future research.

2 Context and Positioning within the Literature

An assessment of the financial effects of Russia's full-scale invasion of Ukraine requires a clear account of the conditions preceding the shock and of the institutional mechanisms through which it could affect defence producers. This chapter establishes that interpretive foundation. It identifies the historical, strategic and industrial context necessary to understand the expected change in military demand and situates the empirical question within the existing literature. The hypotheses and empirical results are developed in the subsequent chapters.

2.1 The Pre-2022 Defence Baseline

With the end of the Cold War, the prospect of large-scale conventional conflict in Europe appeared to recede. Maintaining the force structures and spending levels inherited from bipolar confrontation consequently became more difficult to justify, opening the possibility of reallocating public resources towards civilian priorities. This potential gain is captured by the concept of the peace dividend: sustained reductions in military expenditure can release resources for civilian uses and support higher economic growth [5].

The post-Cold War adjustment varied substantially across the Atlantic Alliance. The United States retained broader global military commitments and bore a larger share of the Alliance's defence burden. Many European allies maintained smaller defence budgets and relied to varying degrees on expenditure and capabilities supplied by other members. George and Sandler associate this configuration with persistent burden-sharing asymmetry and free-riding incentives within NATO [6].

A related asymmetry concerned procurement. United States military demand was organised through federal defence budgets and acquisition institutions. European procurement remained distributed across national budgets and procurement authorities. This structure sustained domestic procurement preferences and limited European market integration [7, 8]. These differences are relevant here because they shaped the channels through which a subsequent increase in military demand could reach defence firms.

Russia's annexation of Crimea and the outbreak of armed conflict in eastern Ukraine in 2014 prompted a reassessment of the European security environment. NATO's response combined changes in its deterrence posture with a renewed commitment to defence investment. At the Wales Summit, allies already meeting the guideline of spending 2 per cent of GDP on defence agreed to maintain it. Allies below the guideline committed to halt further reductions, increase expenditure in real terms as their economies grew and move towards the 2 per cent level within a decade [9].

The Defence Investment Pledge also addressed the composition of expenditure. Allies already allocating at least 20 per cent of their defence budgets to major equipment, including the associated research and development, committed to maintaining that level. Allies below the benchmark agreed to move towards it. The Wales commitments established common objectives for spending and investment. Their implementation remained nationally determined and consequently varied in speed and extent [9].

Official NATO expenditure series show that the direction of aggregate spending changed after 2014. European allies and Canada recorded seven consecutive years of real increases between 2015 and 2021. Implementation remained uneven. In 2021, eight allies met the guideline of spending 2 per cent of GDP on defence, compared with three in 2014 [10].

George and Sandler's alliance-membership-weighted specification produces a positive post-2014 interaction. The remaining weighting schemes produce statistically insignificant or negative coefficients. Their models using Russian military expenditure also estimate a negative association in several specifications. The evidence indicates limited and specification-dependent progress in burden sharing during 2015-2020 [6].

The expenditure increase coincided with persistent institutional and industrial constraints. The European Commission and the High Representative identified persistent capability shortfalls, depleted equipment stocks, limited production capacity and procurement organised predominantly on a national basis. In 2020, collaborative investment accounted for 11 per cent of the defence investment reported by eleven Member States, the agreed benchmark was 35 per cent [11]. The post-2014 period combined rising aggregate expenditure with uneven national implementation, fragmented procurement and persistent capability and burden-sharing gaps.

2.2 The 2022 Geopolitical Shock

February 2022 defines the empirical boundary between the pre-invasion period and the realisation of full-scale war. Financial markets had already begun incorporating information concerning the risk of military escalation. The following sections examine this anticipatory adjustment and the military-demand channel generated by the invasion.

Russia's troop build-up and the public discussion of a possible attack made military escalation increasingly salient before 24 February 2022. Investors could revise their expectations concerning its economic and financial consequences during this period. The invasion converted this prospective risk into observed full-scale war.

Granát et al. examine adjusted returns for 62 country-level equity markets to estimate when information concerning the conflict began to be incorporated into prices. Their results suggest that this process started at least 50 days before the invasion. Approximately one-third of the estimated information incorporation occurred on the event day itself [12]. The evidence therefore indicates a partial anticipation.

The unit of analysis is important for interpreting this result. Country-level indices aggregate firms operating in different industries and with different forms of exposure to the conflict. The estimates document partial anticipation in country-level equity indices. The thesis therefore examines the timing and magnitude of information incorporation among listed defence producers.

The date of 24 February remains an economically meaningful event boundary because it marks the beginning of observed large-scale military action. February 2022 is the first post-treatment observation. Its month-end adjusted price incorporates the market reaction to the invasion. The corresponding monthly return covers the complete interval from the end of January to the end of February, including trading days preceding and following 24 February. Chapter 4 examines the sensitivity of the results to the exclusion of this observation.

The invasion generated a common macro financial exposure across listed firms through heightened uncertainty, disruptions to trade and commodity markets, and changes in financial conditions. Geographical proximity to Ukraine contributed to the cross-country heterogeneity of this exposure [1, 13]. The control firms provide the comparison used to account for the market-wide and macroeconomic component of the shock.

Defence producers faced an additional channel associated with the political and military response. Assistance to Ukraine, the replacement of transferred equipment and the reassessment of national defence requirements created expectations of higher military demand. European military expenditure increased by 13 per cent in real terms in 2022, the largest annual increase in the region during the post-Cold War period. Russian and Ukrainian expenditure accounted for a substantial share of this regional increase. Several other European countries also raised their military budgets [14]. The firm-level transmission of this expenditure depends on its allocation across programmes, suppliers and contracts, as well as on producers' capacity to expand output.

Defence producers were consequently exposed both to the adverse economic effects shared with other firms and to an additional military-demand channel. The empirical comparison concerns this differential exposure within a common shock. Because the invasion generated the political and military responses through which the demand channel operated, these responses form part of the treatment. The estimand is the total invasion-centred geopolitical shock, including the assistance, budgetary and procurement decisions generated by the invasion.

2.3 From Immediate Response to Long-Term Defence Commitments

The initial response to the invasion addressed urgent requirements, including military assistance to Ukraine and the replacement of transferred equipment. These measures were progressively accompanied by longer-term planning. Inventory replenishment, capability development and production expansion generally extend across multiple budget cycles. The policy response combined higher expenditure with an extension of the expected duration of military demand.

In May 2022, the European Commission and the High Representative translated the identified gaps into a policy agenda centred on coordinated replenishment, joint procurement and reinforcement of the European defence industrial base. Coordination was presented as the mechanism for converting higher expenditure into capability development and industrial capacity [11].

At the Vilnius Summit in July 2023, NATO converted the 2 per cent guideline into an enduring minimum and adopted a Defence Production Action Plan addressing demand aggregation, joint procurement and industrial capacity [15]. The decision connected the expansion of defence budgets with the productive capacity required to implement them.

The subsequent evolution extended these immediate concerns into a longer planning horizon. By 2025, NATO allies had agreed to invest 5 per cent of GDP annually in defence and defence-related security expenditure by 2035. At least 3.5 per cent was allocated to core defence requirements. Expenditure on critical infrastructure, resilience, innovation and the defence industrial base could account for up to a further 1.5 per cent. Allies also agreed to submit annual plans setting out a credible incremental path towards the objective [16].

These commitments define the intended scale and duration of future resources. Their conversion into realised industrial demand proceeds through national budgets, procurement programmes, contract awards and productive capacity. The longer planning horizon signals a more persistent policy commitment to military demand. Firm-level revenues depend on the subsequent allocation and implementation of expenditure.

2.4 From Public Demand to Firm Value

Higher military expenditure reaches defence firms through a sequence of budgetary and procurement decisions. This sequence connects strategic requirements, budgetary authorisation, programme definition, supplier selection, contract award, production and delivery. The share available to defence manufacturers depends on the allocation of military budgets among personnel, operations, infrastructure, equipment and research. Aggregate expenditure therefore represents potential demand whose industrial effect depends on its composition and implementation [7].

The allocation of this demand depends on the correspondence between procurement requirements and firms' products, technological capabilities, programme eligibility, contract access and productive capacity [7]. National-security and industrial-policy considerations can also influence supplier selection. Kluth documents the persistence of domestic procurement bias following cross-border industrial consolidation [8]. The arms-sales

share captures firms' predetermined dependence on military activity and serves as an approximate firm-level exposure indicator based on the SIPRI Arms Industry Database [17].

Financial markets can react before procurement decisions produce observable revenues. Under semi-strong market efficiency, publicly available information is incorporated into security prices when it becomes available [18]. Unexpected stock returns reflect news concerning expected future cash flows and discount rates [19]. For defence firms, relevant information includes earnings, analysts' recommendations, geopolitical events, procurement bids and contract announcements [3].

Stock returns measure changes in the adjusted price series over the relevant interval and therefore capture total shareholder returns. Contracts, production and recognised revenue measure realised operating activity. Implementation affects subsequent returns when it changes expectations concerning future cash flows or discount rates.

2.5 Positioning within the Literature and Research Gap

The existing literature can be organised according to its units of analysis, observation horizons and empirical objectives. It covers the general economic consequences of geopolitical uncertainty, short-window market reactions and longer-run relationships between geopolitical risk and defence-stock performance. Each strand addresses a specific outcome and estimand.

2.5.1 Empirical Evidence on Geopolitical Risk and Defence Firms

The broader literature establishes that both geopolitical threats and their realisation can influence economic activity and asset prices. Caldara and Iacoviello construct a news-based measure that distinguishes threats from realised geopolitical events. Higher geopolitical risk is associated with lower investment and employment, a higher probability of economic disasters and greater downside risks. Their industry- and firm-level evidence also indicates that the consequences vary with exposure [1].

Asset prices may adjust during periods of uncertainty concerning the occurrence and consequences of a conflict. Section 2.2 documents anticipatory information incorporation before the invasion. Federle et al. identify geographical heterogeneity in the same episode, with lower equity returns in countries located closer to Ukraine [13]. These

findings establish the aggregate and cross-country setting for the firm-level defence-control differential examined in the thesis.

Research focused directly on defence companies shows that their valuations respond to firm-specific, financial, procurement and geopolitical information. Capelle-Blancard and Couderc conduct a systematic event study of 58 major listed defence companies between 1995 and 2005. The study identifies earnings announcements, analysts' recommendations, geopolitical events, procurement bids and contract awards as significant drivers of defence-stock price movements [3].

Other studies examine the relationship over different samples and horizons. Gurdgiev et al. analyse five major US defence companies using quarterly data from 1990 to 2019 and find that military involvement and defence-budget announcements have dynamic effects extending beyond the announcement period [20]. Zhang et al. use wavelet coherence to examine the co-movement between geopolitical risk and the returns and volatility of 36 defence and aerospace companies from ten countries [21]. Together, these designs characterise time-varying associations and lead-lag relationships at specified frequencies.

Martins et al. focus specifically on the beginning of the 2022 invasion. Their event-study design identifies favourable short-window reactions among the 100 largest publicly listed defence firms and stronger performance among portfolios with greater dependence on defence sales [4]. The thesis extends this evidence by estimating the multi-year persistence and transatlantic distribution of the defence-control return differential.

2.5.2 Research Gap, Contribution and Scope

The first dimension of the research gap is temporal. Existing evidence covers the initial repricing of listed defence firms and the relationship between geopolitical risk and defence-stock performance across investment horizons [4, 21]. The thesis estimates the multi-year evolution of the defence-control return differential from January 2010 to June 2026. The observation window supports an assessment of the pre-treatment relationship, the initial repricing and the subsequent evolution of the differential.

The second dimension is geographical. Market, macroeconomic and institutional conditions vary between Europe and the United States. The thesis constructs region-specific defence-control comparisons. The transatlantic estimand measures the difference

between the change in the European defence-control differential and the corresponding change in the United States.

The third dimension concerns firm-level exposure. The thesis extends the documented short-window relationship between defence-sales dependence and stock performance by estimating the association between median pre-invasion arms-sales share and post-invasion return changes [4, 17]. Programme eligibility, procurement access and production capacity constitute additional firm-level attributes.

The fourth dimension concerns counterfactual design. The defence-firm literature reviewed above applies event-study, panel and time-frequency methods to the analysis of geopolitical events and defence-sector valuations [3, 4, 20, 21]. The thesis estimates a firm-level defence-control comparison and region-specific synthetic counterfactuals for aggregate European and United States defence producers. These strategies produce method-specific outcomes and estimands.

The empirical strategy comprises four fixed-effects specifications. The baseline Difference-in-Differences model estimates the post-versus-pre-invasion change in the average monthly return differential between defence and control firms. The geographical interaction estimates the difference between the European and United States changes. The treated-only arms-sales model estimates the exposure gradient associated with firms' predetermined dependence on military activity. The dynamic specifications trace the defence-control differential through calendar-time and event-phase decompositions.

The arms-sales coefficient measures the change in post-invasion returns associated with a unit difference in pre-invasion arms-sales dependence. The calendar-time coefficients measure the defence-control differential relative to the stated reference period. In the event-based decomposition, each coefficient measures the incremental change associated with the beginning of a new phase, and the differential within a phase is obtained by summing the relevant coefficients. Each phase includes the geopolitical and macroeconomic developments occurring within its boundaries.

The synthetic-control analysis covers two regional aggregates and three illustrative European firms. The regional analyses compare the European and United States defence price indices with counterfactuals constructed from their respective donor pools. Each treated unit has its own placebo distribution and estimand. The Difference-in-Differences models quantify firm-level monthly return differentials, and the synthetic-control models

quantify price-index divergence from constructed counterfactuals. Together, the specifications evaluate the average post-invasion differential, its transatlantic heterogeneity, its association with pre-invasion arms-sales exposure, its temporal evolution and the valuation paths examined through synthetic controls. Chapter 4 states the corresponding identifying assumptions.

3 Theoretical Framework and Hypothesis Development

This chapter develops the theoretical framework from which the hypotheses of the analysis are derived. Its starting point is the research question that guides the thesis:
"To what extent did the geopolitical shock generated by Russia's invasion of Ukraine causally affect the stock-market performance of listed defence firms in Europe and the United States relative to firms operating in sectors without a direct conflict-specific military-demand channel?"

The thesis targets the total effect of the geopolitical shock on the relative stock-market performance of defence firms. This effect includes the political and military response through which the shock was transmitted to financial markets. Budgetary commitments, procurement programmes and military assistance generated by the invasion are therefore included within the treatment mechanism. The empirical design does not identify the direct effect of the invasion separately from this response. The framework connects the arrival of geopolitical information to firm valuation and identifies the channels through which the shock may affect defence producers differently from the rest of the market.

3.1 Geopolitical Shocks, Expectations and Firm Valuation

The connection between a geopolitical shock and firm valuation operates through the information conveyed to financial markets. The geopolitical shock generated by Russia's invasion of Ukraine changed the expected political and economic environment in which firms would operate, with effects that varied across industries. The shock could worsen the general outlook for economic activity and increase the compensation required for bearing risk. It could also improve the expected demand faced by firms whose products formed part of the political and military response. The resulting stock-market response therefore depends on expected cash flows, required returns and firms' positions within the economic transmission of the shock.

3.1.1 Stock Prices as Forward-Looking Valuations

Stock prices are forward-looking because they reflect the information available to investors about the future economic position of a firm. In an informationally efficient market, newly available public information is incorporated into prices as market participants revise their expectations [18]. This adjustment does not require an immediate change in the current revenues or profits of the firm. Information that alters the probability of future contracts, the expected profitability of future operations or the risks associated with those operations can affect the current valuation even when the contemporaneous accounting position remains unchanged.

The present value of a firm depends on two elements: the cash flows investors expect it to generate and the rate at which those cash flows are discounted. Campbell decomposes unexpected stock returns into revisions concerning future cash flows and revisions concerning future expected returns [19]. In economic terms, a price increase may therefore follow from an improvement in the expected stream of cash flows, from a reduction in the return required by investors or from a combination of the two. Conversely, expectations of lower cash flows or a higher required return reduce the present value of the firm.

The distinction is particularly relevant when the information concerns a geopolitical event. A conflict may alter the revenues, costs and investment opportunities that investors expect for a firm. At the same time, it may increase uncertainty and the risk premium applied to those expected outcomes. These components can move in the same direction or offset each other, depending on the position of the firm within the economic transmission of the shock. An observed stock return captures their combined valuation effect, while price data alone do not separate the part attributable to cash-flow news from the part attributable to changes in the discount rate.

Stock-market performance in this thesis therefore measures the combined valuation revision reflected in returns. Realised accounting outcomes remain outside the dependent variable.

3.1.2 Geopolitical Risk, Expected Cash Flows and Discount Rates

Geopolitical risk concerns the threat, realisation and escalation of adverse events associated with wars, terrorism and tensions between states. An increase in this risk changes

both the expected economic environment and the uncertainty surrounding it. The index developed by Caldara and Iacoviello associates heightened geopolitical risk with adverse effects on economic activity, including lower investment and employment and a greater probability of severe negative outcomes [1]. These consequences provide a first channel through which a geopolitical shock can reduce the expected cash flows of firms.

A conflict may affect expected cash flows through several related mechanisms. Disruption to production and trade can raise input costs or restrict access to established markets. Greater uncertainty can lead firms to postpone investment and households to revise expenditure. A deterioration in expected economic activity can reduce future demand, while the political response to the conflict may impose additional constraints on particular transactions and operations. The precise importance of these mechanisms varies across countries and industries, but their combined effect can worsen the expected profitability of a broad part of the market.

Geopolitical risk also affects the rate at which future cash flows are valued. When uncertainty about future states of the world increases, investors may require greater compensation for holding risky assets. The corresponding increase in the required return reduces the present value of a given stream of expected cash flows. Evidence from the period preceding the Iraq War shows that changes in perceived war risk affected equity prices and a broader set of financial variables, demonstrating that the probability and expected consequences of a conflict can be priced before its final resolution [2].

At the aggregate level, geopolitical risk can adversely affect valuations through revisions to expected cash flows and required returns. The incidence of these channels depends on firms' economic exposure, allowing an adverse aggregate effect to coexist with different sectoral responses.

3.1.3 The Economic Specificity of the Defence Industry

The defence industry occupies a distinctive economic position because demand for its principal products originates predominantly from governments. National authorities define security requirements, allocate military budgets and determine which capabilities are acquired. They also influence the supply side of the market through regulation, export controls, the financing of research and development and the design of procurement procedures. The relationship between geopolitical risk and expected demand is therefore

more direct for defence producers than for firms whose products are purchased mainly by private consumers [7].

Defence procurement also differs from ordinary commercial exchange. Major military systems require substantial initial investment, specialised knowledge and long development periods. The resulting technological and institutional barriers limit entry and contribute to a market composed of a relatively small number of buyers and qualified producers. Procurement contracts must address uncertainty over costs, performance and technological development, creating incentive problems that influence the allocation of risks and returns between governments and contractors [22]. The market can consequently approach a bilateral structure in which a dominant public purchaser interacts with a limited number of specialised suppliers [7].

These characteristics make the expected demand of defence producers particularly sensitive to changes in the security assessments of governments. A deterioration in the geopolitical environment may lead public authorities to raise military expenditure, replenish existing stocks or acquire capabilities that had previously received limited priority. Such decisions enlarge the set of potential programmes and contracts available to the industry. Where technological knowledge and productive capacity require time to develop, a revision of expected demand may also increase the economic value of the capabilities already held by incumbent producers.

A larger military budget does not translate automatically or proportionally into firm value. The transmission depends on the allocation of expenditure across programmes, on the products supplied by the firm, on its eligibility to compete for the relevant contracts, on productive capacity and on the margins allowed by procurement arrangements. An announcement of aggregate spending therefore changes the set of economic opportunities available to producers. The amount of additional profit generated by those opportunities remains uncertain. Its valuation effect depends on the probability that those opportunities will generate cash flows for the individual producer.

The distinctive feature of the defence industry is consequently the presence of a demand channel connected directly to the political response to geopolitical risk. Firms across the market may face worsening economic conditions and a higher required return, while defence producers may simultaneously face an increase in expected public demand.

The net effect on their valuation depends on the relative strength of these channels and on the market's assessment of how future expenditure will be transmitted to the industry.

3.1.4 Expectations Before Accounting Realisation

The economic effects of defence expenditure are realised through a sequence that can extend over several years. A political commitment to higher military spending must be translated into budgetary allocations and procurement programmes. These programmes are followed by tenders or negotiations, contract awards, production and delivery. Revenue is recognised as the firm performs the contractual obligations under the applicable accounting rules. The initial political information and its final appearance in the income statement are therefore separated by the institutional and productive stages of the procurement process.

Once a contract has been awarded, firms may disclose its uncompleted portion as order backlog. The backlog represents contracted work that remains to be performed and has not yet been fully recognised as revenue. It provides information about future activity, but it appears only after a general expectation of military demand has been converted into identifiable orders. Before that point, investors may already assign probabilities to the programmes likely to be approved, the firms likely to obtain them and the cash flows those contracts may generate.

The forward-looking nature of stock prices allows these expectations to be incorporated before accounting realisation. Evidence on listed defence firms shows that their market value responds to earnings information, procurement bids, contract awards, public military expenditure and geopolitical events [3]. These forms of information affect the expected distribution of future cash flows even though their consequences may not yet be observable in current revenues.

The same principle has been used to extract information about government expenditure from the returns of military contractors. Fisher and Peters construct a measure based on innovations in the accumulated excess returns of large United States defence contractors to identify changes in expectations about future military spending [23]. Their approach rests on the proposition that firms with a direct economic interest in military procurement are repriced when the expected path of that expenditure changes.

The timing of market valuation and accounting realisation explains why stock returns are the relevant outcome for the analysis. The geopolitical shock can modify expectations about future public demand when information arrives, while contracts, production and revenues are realised through subsequent institutional and productive stages.

Taken together, the framework identifies a broad geopolitical-risk channel and a defence-specific public-demand channel. Their interaction provides the theoretical basis for examining the average relative response of defence producers and its variation across geographical markets, firms and periods.

3.2 Sectoral Exposure and the Average Post-Invasion Differential

The framework developed in Section 3.1 implies that the geopolitical shock generated by Russia's invasion of Ukraine affected firm valuations through broad market channels and a defence-specific demand channel. The broad component operates through the macroeconomic conditions and required returns relevant to firms across the market. The sector-specific component arises from revisions in the military demand expected by defence producers. H1 concerns the balance between these components for defence firms relative to firms without a direct conflict-specific military-demand channel.

3.2.1 Broad Market Channels of the Geopolitical Shock

The geopolitical shock generated by Russia's invasion of Ukraine created adverse conditions for the wider economy. It increased uncertainty concerning the duration and possible escalation of the conflict, disrupted established commercial relations and affected expectations about future economic activity. Investors also had to evaluate sanctions, interruptions to production and trade, changes in commodity markets and the broader policy response. These developments could reduce expected firm cash flows and increase the compensation required for holding risky assets.

These developments represent the application of the broad geopolitical-risk channel defined in Section 3.1.2. Their intensity varied across industries according to exposure to energy costs, international trade, financial conditions and consumer demand. The

channels were not generated by the production of defence goods and were therefore relevant to defence producers and firms operating elsewhere in the economy.

Firms in the comparison group are consequently not assumed to be unaffected by the geopolitical shock. They remain exposed to uncertainty, risk aversion and its macroeconomic and financial consequences. Their defining characteristic is the absence of a direct demand channel arising from military procurement. Their performance therefore provides a comparison for the components of the shock extending beyond the defence industry.

3.2.2 The Defence-Specific Demand Channel

Defence producers were exposed to the broad adverse channels of the geopolitical shock, but the invasion also changed the security problem that their products are intended to address. Within Europe, the immediate response identified the need to replenish stockpiles, restore readiness, expand existing forces and reinforce military capabilities [11]. Military expenditure in Europe increased by 13 per cent in real terms during 2022, showing that the reassessment had begun to affect realised public spending [14]. Together, these developments created a positive expected-demand channel specific to the defence industry.

Within the framework developed in Section 3.1.4, the invasion revised the probabilities assigned to future military expenditure, procurement programmes and contract awards before the corresponding revenues were realised. Subsequent information about the scale and composition of the political and military response allowed investors to update the expected cash flows of firms capable of supplying the required systems. The defence-specific demand channel could therefore offset part of the adverse broad-market response affecting the same firms.

Evidence concerning Russia's invasion provides direct support for the expected sectoral response. Martins, Correia and Gouveia document positive short-term stock-market returns among major listed defence firms around the event [4]. Their analysis concerns the immediate reaction and does not determine whether the differential continued over a longer horizon.

The valuation effect depends on the probability that aggregate commitments will be converted into procurement programmes and on the expected position of each firm

within those programmes. Investors consequently value the expected share and profitability of future opportunities available to each producer.

3.2.3 The Expected Sectoral Differential

The theoretically relevant comparison is between the valuation revision experienced by defence producers and that experienced by firms exposed to the general consequences of the invasion without benefiting directly from military demand. A comparison of defence firms before and after February 2022 in isolation would combine the defence-specific channel with changes affecting the market as a whole. The sectoral differential instead expresses the additional response associated with the economic position of the defence industry.

A positive differential can emerge even if the general consequences of the geopolitical shock reduce the value of defence firms in absolute terms. It requires the revision in expected military demand to produce a more favourable average response among defence producers than among firms without the same direct channel. The prediction concerns the average difference between the two groups and does not require every defence producer to outperform every control firm.

The first hypothesis is therefore:

“H1. Following the geopolitical shock generated by Russia’s invasion of Ukraine, defence firms exhibit a positive average post-invasion return differential relative to firms operating in sectors without a direct conflict-specific military-demand channel.”

3.3 Geographical Exposure and Transatlantic Heterogeneity

The sectoral differential developed in the preceding section may vary across geographical markets. Its magnitude depends on how the geopolitical shock changes regional security conditions, expectations concerning public expenditure and the procurement opportunities available to producers. Europe and the United States differed in their proximity to the conflict, their pre-invasion defence equilibria and the extent to which the shock represented a departure from previous expectations. These dimensions provide the theoretical basis for examining the geographical heterogeneity of the post-invasion differential.

3.3.1 Geographical Proximity and Perceived Security Risk

The geographical location of a conflict affects the intensity with which its economic and security consequences are perceived. Countries closer to the conflict face greater exposure to disruptions in trade, energy supply and cross-border production. Proximity can also increase the perceived probability that the conflict will escalate or produce further security consequences. Investors may therefore assign greater economic and political relevance to the same geopolitical event when it occurs closer to the market in which a firm operates.

Federle et al. document a proximity penalty in the stock-market response to Russia's invasion of Ukraine: markets and firms located closer to the conflict experienced more negative returns following the invasion [13]. A substantial part of the relationship is associated with differences in trade exposure, while the remaining geographical component is consistent with the pricing of military disaster risk. Their evidence establishes that physical distance from the conflict was relevant to its financial transmission.

For the European market, the invasion consequently represented both an economic disruption and a direct deterioration of the regional security environment. The possibility of further escalation, the vulnerability of neighbouring states and the need to reinforce collective security acquired greater immediacy. The United States was politically and militarily involved in the response but remained geographically more distant from the conflict and less exposed to its direct regional consequences.

Geographical proximity produces two potential valuation channels. Greater exposure to the conflict can worsen the general economic outlook and increase the required returns applied across the European market. A more immediate security threat can also increase the expected scale of the political and military response, creating a sector-specific revision in military demand. The interaction between these broad and sector-specific channels provides a basis for a larger defence-control differential in Europe.

3.3.2 Different Pre-Invasion Defence Equilibria

The valuation effect of new information depends on its difference from the expectations already incorporated into prices. An increase in military expenditure produces a smaller revision when a high level of expenditure is already expected and a larger revision when previous expectations reflect a lower and relatively stable level. The transatlantic

comparison therefore concerns the change relative to the respective pre-invasion equilibria as well as the absolute size of the two defence markets.

European defence policy developed for several decades within the conditions associated with the post-Cold War peace dividend. The reduction in the perceived probability of a major interstate conflict allowed European governments to maintain relatively contained military expenditure and to rely partly on the capabilities provided through the Atlantic alliance. Evidence covering the period from 1991 to 2020 identifies persistent free-riding incentives within NATO and a limited adjustment of allied expenditure to the growth of Russian military spending before the invasion [6]. These conditions contributed to an equilibrium in which a large and rapid expansion of European defence demand was not the central expectation underlying the valuation of the industry.

The United States provides a relevant comparison because the conflict was geographically more distant from its domestic market and produced a less immediate change in its regional security environment. United States producers retained potential demand gains. The European security environment may therefore have experienced a sharper discontinuity relative to its pre-invasion position.

The geopolitical shock could consequently generate different valuation revisions on the two sides of the Atlantic. In Europe, the invasion challenged an established expectation of comparatively limited military requirements and created the prospect of rebuilding capabilities that had received less priority during the preceding period. In the United States, it reinforced a demand environment that was already more extensive and persistent. Since stock prices respond to changes relative to prior expectations, the greater discontinuity of the European equilibrium can produce a stronger repricing even if the absolute size of the United States defence market remains larger.

3.3.3 The Geography of the Procurement Response

A revision in military expenditure affects the value of a defence firm only when the additional resources are expected to generate procurement opportunities from which that firm can benefit. The geographical distribution of the valuation response therefore depends on both the location of the increase in expenditure and the allocation of the resulting contracts. The connection between national security priorities, defence budgets and industrial policy makes the geographical destination of procurement economically relevant.

The response to the invasion was not uniform across countries. The largest increases in military expenditure were concentrated among states facing the most immediate change in their security environment, particularly those located closer to Russia and Ukraine [14]. Military assistance, the replacement of transferred equipment and the reinforcement of national capabilities further increased the demand associated with the European response. These developments made the revision of public demand particularly important for firms positioned within the European defence market.

Additional military expenditure may be allocated differently from producers' existing market shares. Defence procurement reflects national security considerations, industrial capabilities and the objective of retaining strategically important production within the relevant jurisdiction. Cross-border consolidation can reduce national procurement bias and increase the acquisition of European systems [8]. Procurement and defence-industrial policies consequently shape the relationship between military budgets and the firms that ultimately receive the corresponding demand [7].

Defence firms obtain contracts across regional boundaries. United States firms can participate in European procurement. European producers may supply the United States and other international customers. Headquarters identify the institutional and industrial environment in which a firm is primarily positioned and provide a partial measure of its commercial exposure. Domestic and regional procurement preferences make this position relevant to the probability that a firm will benefit from the policy response.

3.3.4 The Expected Transatlantic Differential

The expected geographical difference results from the interaction of proximity, prior expectations and procurement. Europe experienced a more immediate deterioration in its security environment and a larger departure from its pre-invasion defence equilibrium. These conditions increased both the negative general exposure of the European market and the expected military-demand response directed towards the defence sector. In the United States, the conflict was geographically more distant, and the existing defence market already incorporated a higher and more stable level of public demand.

The prediction concerns the difference between defence firms and control firms within each area. The European differential compares European defence producers with European firms operating in sectors without a direct military-demand channel. The United

States differential makes the corresponding comparison within the United States market. This structure separates the expected sectoral asymmetry from a direct comparison of the general performance of European and United States equities.

The empirical analysis tests the combined implication of these mechanisms. If the revision in military-demand expectations was greater relative to the previous European equilibrium, the difference between defence producers and their respective control firms should also be greater in Europe. The second hypothesis is therefore:
"H2. The post-invasion return differential between defence firms and their respective control firms is larger in Europe than in the United States".

3.4 Firm-Level Exposure to the Defence Business

The preceding hypotheses classify firms according to their membership in the defence sector. Within that sector, the economic importance of military activities varies considerably. Some producers obtain most of their revenues from arms sales, while others combine military production with substantial civil operations. This variation provides a further dimension for examining the transmission of the geopolitical shock generated by Russia's invasion of Ukraine. H3 concerns whether firms with greater pre-invasion dependence on the defence business experience a larger stock-market response.

3.4.1 Sector Membership and Economic Exposure

A sectoral classification identifies firms whose activities include the production of military goods and services. The economic importance of these activities varies across producers. The defence industry includes highly specialised firms and diversified groups operating in civil aerospace, engineering, electronics and other industrial markets. Firms classified within the same sector may therefore differ considerably in the proportion of their revenues connected to military demand.

This distinction is relevant because the invasion does not necessarily change expectations concerning every segment of a diversified firm. The expected increase in military expenditure is directed towards the defence component of its activity, while its civil operations remain exposed to their respective demand conditions and to the general consequences of geopolitical risk. The effect on the valuation of the entire firm should

consequently depend, at least in part, on the economic weight of the component to which the positive demand revision applies.

The mechanism can be understood through a sum-of-parts interpretation of firm value. The expected cash flows of a diversified company originate from several business segments. If the geopolitical shock raises the expected cash flows associated with military production while leaving the other segments unchanged, a firm whose activity is concentrated in defence receives a larger revision relative to its total business than a firm for which defence is a secondary activity. If the civil segments are adversely affected by the wider economic consequences of the conflict, diversification may further reduce the net revision associated with military demand.

A continuous measure of exposure therefore contains information that a binary sectoral indicator cannot represent. Sector membership identifies whether the defence-demand channel is present, while the share of revenues derived from arms sales indicates its importance within the existing operations of the firm.

3.4.2 Pre-Invasion Arms-Sales Share as an Exposure Proxy

The use of arms sales as a share of total revenues follows from the expected-cash-flow mechanism. Consider two defence firms facing a comparable upward revision in the expected demand for their military products. For the firm that derives a larger part of its activity from arms sales, the affected business represents a greater proportion of its overall revenue base. The corresponding revision should therefore have a greater influence on the expected cash flows and valuation of the firm as a whole.

The measure must refer to the period preceding the invasion. A contemporaneous or post-invasion revenue share could incorporate contracts, acquisitions or strategic adjustments produced by the geopolitical shock and would therefore be partly determined by the treatment. Pre-invasion arms-sales dependence describes the economic position of the firm before the arrival of the new geopolitical information and provides an ex ante indicator of its sensitivity to a revision in military demand.

Existing evidence supports the broader proposition that differences in exposure to government demand can produce heterogeneous stock-market responses. Belo, Gala and Li show that differences in industry exposure to government spending are reflected in

cash flows and in the cross-section of stock returns [24]. This establishes a connection between the economic importance of public demand and the valuation of individual firms.

Evidence concerning military producers supports the same proposition. Martins, Correia and Gouveia construct portfolios weighted by the proportion of revenues generated by defence activities and find a stronger short-term response among portfolios with greater defence-sales exposure following Russia's invasion [4]. Amodeo and Briganti use the reliance of United States contractors on military revenues to estimate their cross-sectional response to events that altered the expected path of defence expenditure [25]. Their working paper applies revenue dependence directly as an indicator of exposure to military-spending news.

The arms-sales share measures firms' existing economic dependence on military activity. Its coverage of additional post-invasion demand is incomplete because it does not observe the composition of future programmes, productive capacity, expected contract profitability or the geographical origin of customers. Its theoretical relevance derives from the proposition that a larger pre-existing military component should increase the sensitivity of the firm's overall valuation to information concerning military demand.

3.4.3 The Expected Firm-Level Gradient

The firm-level prediction follows from the difference between exposure to the defence sector and dependence on it. The first identifies whether the military-demand channel is present. The second indicates how extensively a revision of that demand can affect the total activity of the firm. If investors relate the expected benefits of higher military expenditure to the existing economic composition of producers, the stock-market response should increase as the importance of arms sales increases.

The relevant comparison is conducted among defence firms, all of which are exposed to the military-demand channel but differ in their pre-invasion dependence on that business. Within the measurement scope defined in Section 3.4.2, H3 tests whether the arms-sales share provides a systematic ex ante predictor of the post-invasion response. The third hypothesis is therefore:

"H3. Among defence firms, the post-invasion change in stock-market returns is positively related to the pre-invasion share of revenues derived from arms sales".

3.5 Information Arrival and Persistence

The previous sections develop the direction and cross-sectional distribution of the post-invasion differential. Section 3.5 examines its temporal dimension. Financial markets may respond when information about the geopolitical shock first becomes available and revise valuations as subsequent decisions clarify its economic consequences. H4 concerns whether the return differential remains observable beyond the immediate market response.

3.5.1 The Immediate Incorporation of Geopolitical Information

Under the valuation framework developed in Section 3.1.1, prices adjust when new public information changes investors' expectations. The invasion of 24 February 2022 provided information about the security environment, the probability of a prolonged conflict and the likelihood of a political and military response. If investors interpreted the event as implying a durable increase in military requirements, the initial revision in the expected cash flows of defence producers could be incorporated immediately into stock prices.

The date of an event does not necessarily coincide with the first arrival of information concerning it. Before the invasion, the concentration of Russian forces and the public discussion of a possible attack could already have changed the probability investors assigned to the event. The subsequent invasion resolved part of that uncertainty and provided additional information about the scale and nature of the conflict. The initial repricing may therefore include both an anticipatory component and a further adjustment when the event occurred.

Evidence concerning Russia's invasion confirms the relevance of the immediate event window. Martins, Correia and Gouveia document positive abnormal returns among major listed defence firms around the invasion [4]. Their findings identify the initial market response but do not determine whether the return differential remained observable over a longer horizon.

The initial return reflects the information available when the invasion occurred. The event increased the expected probability of additional military expenditure, while the size, composition and duration of the policy response remained uncertain. The subsequent resolution of these uncertainties created the possibility of further valuation adjustments.

3.5.2 Subsequent Policy Information and Continued Updating

The political, military and industrial consequences of the invasion were not fully determined on 24 February 2022. Governments subsequently announced budget revisions, procurement initiatives and longer-term capability commitments [11, 15]. At firm level, procurement bids, contract awards and related corporate disclosures could provide additional information about the allocation of expected demand [3]. These developments progressively clarified how the geopolitical shock and its associated policy response could affect the expected cash flows of defence producers.

The timing of expenditure news is distinct from the timing of expenditure itself. Ramey shows that changes in government spending are frequently anticipated and that their economic interpretation depends on identifying when information about future expenditure becomes available [26]. Ben Zeev and Pappa apply this principle specifically to anticipated defence spending shocks, identifying news about future military expenditure separately from contemporaneous spending and documenting an immediate response in the excess returns of defence contractors [27]. The revision in expectations constitutes the relevant economic event, while accounting realisation occurs at a later stage.

Successive revisions in the expected path of government demand can generate successive adjustments in the valuation of defence producers. This application follows the framework developed in Section 3.1.4, where contractor returns provide information about expected military expenditure before its accounting realisation.

Continued repricing is consistent with market efficiency when each adjustment reflects information that was not fully available during the preceding period. The expected passage of time between a budget announcement, contract award and revenue recognition does not independently generate additional returns. A further return arises when new information changes the probability, magnitude, timing or profitability of previously anticipated cash flows. Policy and procurement developments can therefore confirm, expand or reduce the initial valuation revision.

3.5.3 The Expected Persistence of the Return Differential

The persistence examined in this thesis concerns the return differential. A positive differential produces further divergence between the relative price paths of defence and control firms. An existing price-level gap can remain after their returns become equal. H4

examines whether additional relative returns remain observable at annual frequency beyond the initial adjustment. The hypothesis permits monthly variation and non-monotonic adjustment within the post-treatment period.

Previous evidence indicates that the relationship between geopolitical information and defence-sector returns can differ across investment horizons. Zhang et al. find that the co-movement between geopolitical risk and the returns and volatility of global defence firms varies across time-frequency scales during the period surrounding the invasion [21]. Their analysis concerns associations across investment horizons within a short daily sample and does not establish persistence across subsequent post-treatment years. Gurdgiev, Henrichsen and Mulhair find that military-budget announcements can have dynamic effects on defence-stock performance. They also document subsequent moderation in the response to direct conflict-related information [20]. Together, these findings establish the relevance of examining the temporal dimension of the response without determining its duration in advance.

In the setting considered here, the geopolitical shock was followed by information concerning military budgets, procurement programmes and longer-term commitments. These developments provided a basis for further revisions in expected military demand after the initial market adjustment. The fourth hypothesis is therefore:

“H4. The positive post-invasion return differential of defence firms relative to firms without a direct conflict-specific military-demand channel remains observable at annual frequency beyond 2022.”

4 Research Design

This chapter presents the research design of the thesis. It defines the empirical approach adopted to test the hypotheses developed in the previous chapter.

4.1 Overview and Empirical Approach

This thesis adopts a Difference-in-Differences design to examine the stock-market response to the geopolitical shock generated by Russia's invasion of Ukraine. The treatment group comprises listed defence producers, while the control group comprises firms without a direct conflict-specific military-demand channel. The control firms may nevertheless be exposed to the broader macroeconomic and financial consequences of the conflict.

The causal estimand targeted by the analysis is the total effect of the geopolitical shock, including the political and military response through which it was transmitted to expectations concerning defence demand. The empirical design does not identify the direct effect of the invasion separately from the effects of the policy measures, military assistance and revisions in expected military expenditure that followed it.

Four specifications address the hypotheses developed in Chapter 3. The baseline specification estimates the average post-treatment return differential between defence and control firms. A geographical interaction examines whether this differential varies between Europe and the United States. A continuous measure of pre-invasion defence exposure examines its firm-level distribution within the treatment group. A dynamic specification describes the evolution of the return differential over the post-treatment period.

The treatment indicator switches on in February 2022 for all defence producers. This common temporal boundary reflects the timing of the invasion and does not imply identical exposure across firms or exclude the possibility of anticipatory adjustment. The credibility of the estimates is examined through alternative specifications, placebo dates, alternative observation windows and the construction of a synthetic counterfactual. The remainder of the chapter presents the data, the econometric specifications and the synthetic-control design.

4.2 Data

This section describes the data on which the analysis is built. It presents the sources from which the data are drawn, the procedure through which the sample is constructed, the composition of the treatment and the control group and the definition of the variables used in the estimation.

4.2.1 Data Sources

The analysis relies on stock prices, on an industry classification and on an external ranking of arms producers. The identification of the defence firms and the measurement of their exposure to the defence business are based on the SIPRI Arms Industry Database, which ranks the largest arms producing companies in the world and reports for each of them the share of revenues obtained from arms sales. The initial universes from which the control firms are drawn are the constituent lists of the STOXX Europe 600 and of the S&P 500. Adjusted closing prices for all firms in the sample were retrieved from Yahoo Finance and cover the period from January 2010 to June 2026.

The prices are adjusted for dividends and for stock splits and the outcome variable therefore measures a total return. Without this adjustment, a two-for-one split would appear as a return of minus fifty per cent and firms with a high dividend payout would appear to underperform firms that distribute little. Since the control group contains many established dividend payers, a systematic difference in payout policy between the two groups would translate directly into a distortion of the estimated effect.

The monthly observation corresponds to the closing price at the end of the month. February 2022 therefore already incorporates the market reaction to the invasion of 24 February, and the month is assigned to the post-treatment period. The alternative definition, which excludes February from the post-treatment window, is reported among the robustness checks.

The industry classification adopted is GICS, applied uniformly to the two markets. For the United States firms the classification is available directly from the index constituent data. For the European firms it is not, and it was therefore reconstructed firm by firm from the industry information reported by Yahoo Finance, from which the correspondence with GICS is direct for the categories employed here.

Yahoo Finance is a freely available source and not a professional data provider. The choice is nevertheless adequate to the purpose of this work, which requires monthly closing prices over a long horizon and no intraday information. The retrieved series contain no missing observations within the listing period of each firm and all of them extend to the end of the observation window, while a small number of firms were admitted to listing after January 2010 and their series begin later.

4.2.2 Sample Construction

The treatment group and the control group are built through different procedures and are therefore described separately. The firms in the defence sector are identified from an external classification, while the control firms are drawn randomly from a pool of eligible companies. Both groups are identified initially by common institutional and geographical criteria. The perimeter of the analysis comprises the countries that took part in the Western response to the invasion, which are the member states of the European Union, the United Kingdom, Norway and the United States. Participation is understood as adherence to the sanctions regime and to the instruments of cooperation and industrial financing of defence that followed the invasion.

The treated firms are selected from the SIPRI Arms Industry Database, restricted to the editions covering the years from 2010 to 2021. The population is closed before the shock by construction, so the membership of the treatment group cannot depend on what happened after February 2022. The following conditions are then applied in sequence. The firm must be headquartered within the geographical perimeter of the analysis, it must be listed on a regulated market, and a price series must be available over the pre-treatment period and must extend to the end of the observation window. It must appear in at least four of the twelve pre-treatment years, since the measure of exposure to the defence business is computed as a median over the period. Its arms sales must exceed ten per cent of total sales, both on average and at the median, so that firms with a marginal exposure are not classified as defence firms. Finally, its corporate history must be free of the discontinuities that would compromise the price series. The requirement concerning the price series is less restrictive than the one applied to the control firms, since a moderate reduction of the pre-treatment series does not compromise the estimation, whereas the exclusion of the firms concerned would impair the coverage of the sector. A control firm that

fails to meet the condition is replaced by another name drawn from the pool, whereas a treated firm has no substitute.

The conditions are cumulative and the sample is their intersection, so the order in which they are applied does not affect the outcome. The procedure reduces the population from 170 harmonised company names to 36 firms. The geographical restriction leaves 114 companies, the threshold on defence intensity 97, and the requirement of a minimum presence in the ranking 69. Of these, 30 are not listed and 3 do not have a usable price series.

The control firms are drawn from the constituents of the S&P 500 and the STOXX Europe 600. The treated group is defined through an external classification, while the eligibility of a control firm depends on its location within the geographical perimeter and on the absence of a direct conflict-specific exposure channel. This condition is assessed at both sector and firm level. Sectors affected through macroeconomic channels shared with the wider market, including inflation, interest rates and general risk aversion, are retained. The fixed effects address the average components of these developments that are common within the relevant market comparison. A direct sector-specific channel could instead generate a differential movement in the control group and contaminate the estimated defence-control comparison.

Energy producers are excluded. The shock reached this sector directly. An event study on a global sample of energy firms reports positive cumulative abnormal returns during the Ukraine shock of 24 February 2022 and an outperformance of the general market [28]. Controls of this kind would react in the same direction as the treated firms and would compress the estimated effect.

Utilities are removed by the same logic. In the European wholesale electricity market, the price is set by the cost of the last plant called to meet demand, according to the principle of uniform marginal pricing [29].

This mechanism operates in Europe, where firms did not retain the advantage in energy costs that American firms preserved after the invasion [30]. Utilities are nevertheless removed from the S&P 500 as well, because the second hypothesis compares the effect estimated in the two areas, and a difference in the sectoral composition of the two control groups would enter that comparison as a confounding factor.

Financial firms are removed because the sanctions imposed on Russia generated a direct sector-specific response. Bank stocks fell more than the broader market on the day of the invasion, with the largest decline recorded in Europe [31].

Airlines are removed because of the closure of the airspace. Russian and Ukrainian airspace became unavailable to Western carriers, which lost part of their network towards Asia and had to divert the remaining traffic at higher costs, and the price reaction of the segment around the invasion was negative and stronger for carriers based in Europe [32].

Aerospace and defence firms are removed from both indices. These firms are the object of the analysis, and their presence among the controls would bias the estimate towards zero.

The exclusions operate at two levels of the classification. Energy, utilities and financials are removed at the GICS-sector level. Within industrials, the exclusions are limited to airlines and aerospace and defence.

The exclusions leave a pool of 738 eligible firms, 360 from the S&P 500 and 378 from the STOXX Europe 600. The target was set at fifty firms from each market in order to preserve geographical balance, provide a sufficiently large donor pool for the synthetic-control exercise and keep the firm-level screening feasible. The names are drawn from this pool in random order and are examined one at a time in the order in which they are drawn. The random order is generated with a fixed seed recorded in the extraction script, so that the sequence in which the names are examined can be reproduced. Sector eligibility is not sufficient, since it does not capture the history of the single company, and each name is therefore verified individually before being admitted.

The verification concerns the continuity of the price series and the exposure of the individual firm. The firm must have been listed without interruption from January 2010 to the end of the observation window. It must not have carried out a spin-off or distributed shares to its own shareholders within that period, and the legal entity must be continuous, with no merger, no change of listing venue and no delisting. Multiple share classes of the same company are counted once. Firm-level screening was also used to identify direct exposure channels not captured by the sector classification. Exclusion was based on documented operational exposure and was not determined by post-invasion stock returns or changes in corporate revenues. A control firm was excluded when clear documented evidence showed that the invasion or the measures that followed it had forced the closure,

suspension or disposal of a material operation in Russia. Firms were not excluded merely because they operated in Russia or because their stock price declined after February 2022. Firms selling commodities whose Russian supply was directly disrupted were also excluded. Since comparable firm-level data on Russian trade and operational exposure were not available, the screening was qualitative and restricted to cases for which the evidence was sufficiently clear.

A firm that does not satisfy one of these conditions is discarded and replaced by the following name in the same random order. Eligibility is a property of the firm and does not depend on the position in which the name is drawn, so examining the names in sequence until the required number is reached is equivalent in distribution to a simple random sample drawn from the eligible subpopulation. In total 180 names were examined to cover 100 positions, 81 for the 50 United States firms and 99 for the 50 European ones. Eighty names were discarded and are reported in Appendix A.3: forty-nine on the European side and thirty-one on the United States side. The final sample contains 136 firms, 36 treated and 100 controls, observed at monthly frequency from January 2010 to June 2026. The following section describes how the two groups are composed.

4.2.3 Treatment and Control Groups

The treatment group contains 36 firms, 21 headquartered in the United States and 15 in Europe. The composition of the group reflects the structure of the listed defence industry, in which public ownership and unlisted arrangements are more frequent on the European side. The control group is balanced by construction, with 50 firms drawn from each of the two indices.

Within the treatment group the exposure to the defence business ranges from 0.111, the value of Parker-Hannifin, to 0.952, the value of BAE Systems, with a median of 0.462. The lower bound is a direct consequence of the ten per cent threshold applied in the selection, which truncates the support of the variable on the left. The third hypothesis is therefore tested on this range, and the estimate cannot be extended to the part of the distribution that the sample does not cover.

The exposure measure is distributed almost identically across the two geographical areas. Its firm-level mean is 0.510 among United States producers and 0.512 among European producers. The difference is 0.0016 and the correlation between the European

indicator and the exposure measure is 0.003. The geographical classification therefore explains virtually none of the cross-sectional variation in the selected measure within the treatment group.

The eligibility criteria do not act uniformly on the two markets. They penalise firms with complex corporate histories and firms that disposed of foreign activities during the window, and both features are more frequent in Europe than in the United States. Of the eighty names discarded during the extraction, forty-nine were European and thirty-one were headquartered in the United States. The European control group that results contains no firms classified under materials, although these account for approximately one tenth of the pool from which it is drawn, and consumer staples are under-represented. The absence of materials firms follows from the firm-level conditions, which remove the producers of the commodities affected by the disruption of Russian supply, and the discards recorded in Appendix A.3 show that firms of this kind were drawn and set aside. The sectoral composition of the control group therefore differs from that of the pool.

A second consequence follows from the requirement of a continuous price series. Since the condition extends beyond the treatment date to the end of the observation window, firms that were delisted, that entered insolvency or that were acquired after February 2022 are excluded from both groups. The sample is thus composed exclusively of firms that remained listed until the end of the period, and the estimates refer to this subset. Applying the condition to both groups provides a common sample definition, but it does not eliminate the possibility of survivorship selection if the probability of remaining listed differs between treated and control firms. The estimates therefore refer to firms that remained listed until the end of the observation window and cannot be extended to firms that left the market after the shock.

A third consideration concerns the conditions related directly to the war, which are applied at firm level to the control group alone. Including firms affected through these channels would contaminate the control group in the direction opposite to the expected response of the defence producers. A direct negative shock to the controls would enlarge the estimated difference between the two groups and would be attributed by the estimator to the treatment. Excluding the clearly documented cases therefore reduces this source of upward bias. The screening does not claim to identify every form of commercial exposure

to Russia, and the estimates remain conditional on the control group constructed through these criteria.

4.2.4 Outcome Variables

The dependent variable of the analysis is the monthly logarithmic return of each firm, computed on the adjusted closing prices described in the preceding sections. The adoption of a return in place of a price removes the differences of scale between firms, which would otherwise dominate any comparison within a sample that combines companies of markedly different size and denominated in different currencies. The logarithmic form is used because it is additive over time, so that the log return over a longer interval equals the sum of the monthly log returns it comprises. Log changes of equal absolute magnitude and opposite sign therefore cancel exactly. The computation entails the loss of the first observation of each firm, as the return of any given month is defined with respect to the price of the preceding one.

The outcome is winsorised at the first and ninety-ninth percentiles calculated over the complete estimation sample of 26,663 non-missing firm-month observations. The corresponding thresholds are approximately −0.2103 and 0.2181 in monthly logarithmic returns. Observations below and above these thresholds are replaced by the respective percentile values. The procedure differs from the deletion of the extreme observations, since no observation is removed from the panel and its structure is preserved. The adjustment is required by the sensitivity of the least squares estimator to extreme values. The distribution of monthly returns contains observations of considerable magnitude, which originate in circumstances specific to the individual firm, and without the adjustment a limited number of companies would determine an estimate that is intended to describe the response of the market as a whole. The unwinsorised series is retained, and the estimates obtained on it are reported among the robustness checks.

The specifications are constructed on three indicators and one continuous variable. The indicator $treat$ equals one for firms in the treatment group, while post equals one for months from February 2022 onwards. Their interaction, denoted did, carries the coefficient that estimates the average post-treatment return differential and constitutes the parameter of interest for the first hypothesis. The indicator Europe equals one for firms headquartered in Europe, and its interaction with did allows the estimated differential to

vary between the two areas, as examined by the second hypothesis. The fourth hypothesis introduces no additional variable and is examined by decomposing did into separate coefficients for periods relative to February 2022, thereby describing the evolution of the return differential over time.

The intensity of exposure to the defence business is measured by the continuous variable $DefShare_{\text{i}}$ defined as the median of the ratio between arms sales and total sales reported by SIPRI for each firm over the years from 2010 to 2021, and set to zero for the control firms, which do not appear in the ranking. The median is preferred to the arithmetic mean because the mean is sensitive to extreme values, whereas the median is determined by the central observation of the distribution and is not displaced by an isolated year. The variable is constructed exclusively on years preceding the treatment, so that it cannot incorporate the reaction of the firms to the shock. Its interaction with post, estimated within the treatment group alone, examines the third hypothesis.

The resulting panel comprises 26,799 observations on 136 firms, covering the period from January 2010 to June 2026. The panel is not perfectly balanced, since a small number of firms were admitted to listing after January 2010 and their series begin later. Once the first month of each firm is excluded, 26,663 observations enter the estimation. The observation window terminates in June 2026, the last month for which a complete series is available. The descriptive statistics of the sample and of the variables defined above are reported in Appendix A.4.

4.3 Empirical Strategy

This section presents the econometric strategy used to estimate the post-treatment return differentials, the assumptions required for their causal interpretation and the tests used to assess the credibility of those assumptions.

4.3.1 Baseline Difference-in-Differences Specification

The average post-treatment return differential is estimated by least squares on the panel described in the preceding section, with fixed effects for the firm and for the month:

$$r(i,t) = \alpha(i) + \delta(t) + \beta \cdot [Treat(i) \times Post(t)] + \varepsilon(i,t) \qquad (4.1)$$

The index i runs over the 136 firms of the sample and t over the months from January 2010 to June 2026. The variable $r(i,t)$ is the monthly logarithmic return of firm i in month, $\alpha(i)$ is a firm fixed effect, $\delta(t)$ is a month fixed effect, $Treat(i)$ takes value one for the firms of the defence group and $Post(t)$ takes value one from February 2022 onwards.

The two indicators do not enter the specification separately. $Treat(i)$ does not vary over time and is therefore collinear with the set of firm fixed effects, while $Post(t)$ takes the same value for every firm within a given month and is collinear with the set of month fixed effects. Neither main effect is separately identified, and the parameter that remains estimable is the coefficient β on their product. Once both sets of fixed effects have been removed, β is identified by the variation that is internal to the firm and internal to the month, so that no comparison is drawn between firms observed in different months.

The coefficient is a difference between two differences. It contrasts the variation in the average monthly return of the defence firms between the two periods with the variation recorded by the control firms over the same time interval, net of the permanent differences between firms and of the movements common to the entire market within each month. Whether the resulting quantity may be interpreted as the effect of the shock rests on the assumption examined in Section 4.3.4.

Since the outcome is a monthly logarithmic return, β is measured in monthly log points and represents a differential in average monthly logarithmic returns. For coefficients of limited magnitude, 100β provides an approximate percentage-point interpretation. The transformation $100[exp(\beta) - 1]$ provides the corresponding exact percentage change.

Standard errors are clustered at firm level, which allows the residuals of a given firm to be arbitrarily correlated over time. Treating them as independent would understate the uncertainty of a difference-in-differences estimated on a long panel [33]. The sample contains 136 clusters, a number sufficient for the approximation on which the correction rests.

4.3.2 Specification for the Heterogeneity across Markets

The baseline specification estimates a single coefficient for the whole treatment group, averaged across firms and months. The second and third hypotheses examine

geographical and firm-level exposure heterogeneity in the estimated average post-treatment return differential. H2 examines the institutional and geographical division between the two markets. H3 examines the intensity of firms' engagement in the defence business, as presented in Section 4.3.3. Each specification augments the interaction carrying the average effect with an interaction involving a time-invariant firm characteristic. Relative to the baseline, the H2 specification changes the fixed-effect structure, and the H3 specification changes the estimation sample.

The second hypothesis is tested with the specification:

$$r(i,t) = \alpha(i) + \delta(a(i),t) + \beta_1 \cdot [Treat(i) \times Post(t)] + \beta_2 \cdot [Treat(i) \times Post(t) \times Europe(i)] + \varepsilon(i,t) \quad (4.2)$$

where $a(i)$ denotes the area in which firm i is headquartered, taking the two values Europe and the United States, and $Europe(i)$ is the indicator that equals one when $a(i)$ is Europe. The term $\delta(a(i),t)$ is a fixed effect for each combination of area and month, and it replaces the month fixed effect of the baseline. Where the latter assigns a single parameter to each month, common to every firm of the sample, the former assigns two, one for the European firms and one for the United States firms. Each parameter absorbs the average return recorded by its own market in that month, so the coefficients are identified from within-area, within-month variation.

The indicator $Europe(i)$ does not enter separately, being constant over time and therefore collinear with the firm fixed effects. Of the two terms that remain, the first takes value one for every firm of the treatment group in the months from February 2022 onwards, while the second takes value one only for the European firms among them. The effect estimated on the United States firms is therefore β_1, the effect estimated on the European firms is the sum of β_1 and β_2, and β_2 is the difference between the two.

The hypothesis holds that β_2 is positive, and the null hypothesis against which it is tested sets β_2 to zero, which corresponds to an effect of the same magnitude in the two areas. The quantity on which the hypothesis bears is therefore estimated as a coefficient of the specification, with a standard error of its own, and is not reconstructed from the difference between two separate estimates.

The fixed effects by area and month are required by the structure of the specification. A triple difference requires the presence of every interaction of lower order, among which the interaction between area and period. A month fixed effect common to the two areas does not contain that term, since it removes only what is common to all the firms of the sample within a given month. The interaction between area and period would therefore be omitted from the model, and being correlated with the triple interaction its omission would be transmitted to the coefficient β_2.

The terms omitted in this way are not negligible in this sample. It combines two currency areas, two monetary authorities and two energy markets, so that conditions bearing on all the firms of one market and on none of the other are a regular feature of the period under examination. Under a month fixed effect common to both, a part of what concerned the European firms as a whole would be attributed to the European firms of the treatment group. The fixed effects by area and month remove that component in advance, and the differential that β_2 measures is therefore one between defence firms and their own controls.

4.3.3 Specification for the Heterogeneity across Firms

The third hypothesis is tested with the specification:

$$r_{i,t} = \alpha_i + \delta_t + \gamma(Post_t \times DefShare_i) + \varepsilon_{i,t} \quad (4.3)$$

Equation (4.3) is estimated on the firms of the treatment group alone, using firm and month fixed effects. $DefShare_i$ is the pre-treatment exposure measure defined in Section 4.2.4. Its time-invariant main effect is absorbed by the firm fixed effects, while the month fixed effects absorb movements common to the treated firms within each month. The coefficient γ estimates whether the change in returns following February 2022 varies systematically with firms' pre-invasion dependence on arms sales. A causal interpretation would require that, in the absence of the geopolitical shock, changes in returns would not have been systematically related to $DefShare_i$. The coefficient is therefore interpreted as the estimated firm-level exposure gradient in the observed post-treatment response.

Since DefShare(i) is a share, γ is expressed as a differential of monthly return per unit of that share and is therefore not comparable in magnitude with the coefficient of the

baseline. The differential implied between two producers is obtained by multiplying γ by the distance between their exposures.

The restriction of the sample is required by the meaning of the coefficient. On the full sample the variable takes value zero for every control firm, so that its interaction with the post-treatment period would separate the treated firms from the controls and would identify once again the average effect estimated by the baseline, rather than the gradient internal to the treatment group. Estimated on the treated firms alone, the coefficient compares producers that differ in the intensity of their exposure, which is the comparison that the hypothesis concerns.

The exposure measure is distributed almost identically across the two geographical areas. Its mean is 0.510 among United States producers and 0.512 among European producers, while its correlation with the European indicator is 0.003. The selected exposure measure is therefore effectively orthogonal to the geographical classification within the treatment group. The specification consequently retains common month fixed effects and does not introduce a separate area-by-month structure.

This specification is estimated on the 36 firms of the treatment group. The estimator and the treatment of the standard errors are those of the baseline, but the clusters are considerably fewer, and the approximation on which the correction rests is correspondingly weaker. The inference on this hypothesis is accordingly less reliable than that on the other two, and this is taken into account when the result is discussed.

The second hypothesis differentiates the effect along the geographical and institutional division, the third along the intensity of exposure at firm level, while the fourth, presented in Section 4.3.5, describes its evolution over time.

4.3.4 Identification Assumptions and Parallel Trends

The interpretation of the baseline coefficient β as the effect of the shock rests on an assumption about what would have occurred in its absence. Writing $r^0(i,t)$ for the return that firm i would have recorded in month t had the invasion not taken place, the assumption requires that:

$$E[r^0(i,t)|Treat(i)=1] - E[r^0(i,t)|Treat(i)=0] = Constant \qquad (4.4)$$

be constant across the months of the observation window. The condition concerns the difference between the two groups and not the two returns separately, so it does not require the two groups to perform alike, and it allows the value of one to grow persistently faster than that of the other. Since the outcome is a logarithmic return, a constant difference in returns corresponds to a constant rate of divergence between the two price paths, and what the condition excludes is a change in that rate during the post-treatment period arising from causes other than the shock.

The identifying condition in equation (4.4) concerns the treated-control comparison underlying H1. The geographical specification for H2 requires the corresponding condition to hold within each regional market. The treated-only specification for H3 requires that counterfactual changes in returns would not have varied systematically with pre-treatment defence exposure. The dynamic specification for H4 inherits the treated-control condition for the periods represented by its coefficients. The four specifications therefore rely on related but distinct counterfactual assumptions.

The interpretation rests on a second condition: the treatment must have no causal effect before its implementation [34]. In this application, that condition is not automatic. By January 2022, the concentration of Russian forces and the public discussion of a possible attack were already observable [35]. Equity prices incorporate expectations about future cash flows and anticipated policy developments. Part of the repricing may therefore have preceded the invasion.

The condition cannot be verified, because it refers to a period in which the treatment did occur. What can be examined is the implication it carries for the years preceding February 2022. If the difference between the two groups was already changing before the shock, the proposition that it would have remained constant afterwards loses credibility. The dynamic specification presented in the following section provides the means, since it estimates a separate coefficient for each pre-treatment period, and those coefficients are tested jointly against the hypothesis that they are all equal to zero.

The coefficients on which the examination rests are not absolute quantities. Each of them measures the difference between the two groups in a given period against the same difference in the period taken as reference, so that its sign as well as its magnitude depends on which period is chosen. A reference that is itself unusual for one of the two groups displaces the entire profile. For this reason, the dynamic specification is estimated

against more than one reference, and the examination is conducted on the set of estimates rather than on any single one of them.

The observation window contains the two years of the pandemic, in which a contraction and a recovery of exceptional size reached the firms of the sample through channels unrelated to the shock under examination. Whether they reached the two groups to the same degree bears directly on the assumption, since a difference confined to those two years would displace the pre-treatment profile without implying anything about the years that follow.

To strengthen the credibility of the design, Section 4.4 introduces a second strategy that constructs the counterfactual from the pre-treatment path of the treated unit. Its identifying credibility depends on the quality of the resulting constructed counterfactual.

4.3.5 Dynamic Specification of the Post-Treatment Effect

The fourth hypothesis concerns the evolution of the return differential over the post-treatment period, which the baseline specification cannot describe. The baseline coefficient averages all post-treatment months and cannot distinguish a return adjustment concentrated around the invasion from a sequence of continuing positive differentials. An initial repricing followed by no further differential may leave a persistent difference in price levels, while continuing positive return differentials imply that the price-level divergence continues to widen. The post-treatment period is therefore decomposed in two ways, by elapsed time and by event.

The first decomposition replaces the single interaction with a set of them, one for each period measured relative to the invasion:

$$r(i,t) = \alpha(i) + \delta(t) + \Sigma(k \neq k_0)\beta(k) \cdot [Treat(i) \times D(k,t)] + \varepsilon(i,t) \qquad (4.5)$$

where $D(k,t)$ takes value one when month t belongs to the period lying k periods from the treatment, and k indexes the periods of the specification with the exception of k_0. The complete set of the terms $Treat(i) \times D(k,t)$ sums to $Treat(i)$, which the firm fixed effect already absorbs, so that a specification containing all of them could not be estimated. Each coefficient is accordingly interpreted as a difference with respect to the period that has been excluded. Equation (4.5) describes both versions of the specification,

which differ in the unit of k, in the range over which k runs and in the sample on which each version is estimated.

The coefficients serve two distinct purposes. The pre-treatment coefficients are used to assess the identifying implication discussed in Section 4.3.4 and should be jointly indistinguishable from zero under parallel trends. The post-treatment coefficients describe the evolution of the estimated return differential following the geopolitical shock generated by Russia's invasion of Ukraine. The resulting profile is presented graphically, and H4 is evaluated from its overall temporal pattern rather than from any single coefficient.

Equation (4.5) is estimated at two frequencies, which serve different purposes and cover different spans. The annual version is estimated over the whole observation window and describes the profile of the effect across the years. The coefficient of 2026 is estimated but not reported, since the year is incomplete. The monthly version is restricted to the interval surrounding the invasion, from one year before it to two years after it, and resolves into its component months a reaction that the annual version represents by a single coefficient.

The bounds of that interval follow from its purpose. Monthly resolution is warranted where the movement is concentrated and where an annual coefficient would conceal its shape. Over a longer horizon the annual coefficients describe the profile adequately. The two versions differ also in the sample on which they are estimated. The annual version employs the whole panel, whereas the monthly version is estimated on the firm-months falling within the interval, and its fixed effects are therefore estimated on that interval as well.

The choice of the reference period affects every coefficient of the profile. The annual version uses the year preceding the invasion as its primary reference and is accompanied by two alternatives. The first uses 2019, which precedes the pandemic and the period discussed in Section 4.3.4. The second uses the average of the four years preceding the invasion, distributing the influence of the pandemic across a wider reference period and assigning the reference weight across multiple years. Section 4.3.6 reports both alternatives as robustness checks. Features that vary across reference periods indicate sensitivity to the selected normalisation.

The monthly version adopts November 2021 in place of the month immediately preceding the invasion. The concentration of forces at the border was public information during the closing months of 2021, and a reference placed in January 2022 would therefore risk falling within the reaction rather than outside it. The months lying between the reference and the invasion remain among the coefficients estimated on the pre-treatment period, where they may reflect anticipation of the event rather than a departure from the assumption examined in Section 4.3.4.

The second decomposition organises the post-treatment period into event-based phases:

$$r_{i,t} = \alpha_i + \delta_t + \theta_1(Treat_i \times Post1_t) + \theta_2(Treat_i \times Post2_t) + \theta_3(Treat_i \times Post3_t) + \varepsilon_{i,t} \quad (4.6)$$

where $Post_1(t)$ takes value one from February 2022, $Post_2(t)$ from July 2023 and $Post_3(t)$ from June 2025. The second and the third date are those of the NATO summits at which the commitments on military spending were revised, held at Vilnius and at The Hague. The month of each summit is assigned to the phase that it opens, in accordance with the convention adopted for February 2022 in Section 4.2.1.

The three indicators are cumulative, and all of them therefore take value one during the final phase. Each coefficient is an increment with respect to the phase that precedes it. The differential prevailing within a given phase is obtained by summing the coefficients up to that point, with a standard error computed on the linear combination.

Equations (4.5) and (4.6) describe the same profile in different ways. Equation (4.5) exhibits the shape of the profile without imposing any structure on it. Equation (4.6) is estimated over the whole post-treatment period and partitions it into blocks delimited by the dates on which the commitments on public spending were revised.

The event-based partition supports an interpretation that the design cannot establish directly. The specification identifies a change in the prevailing return differential at a given date without identifying the event responsible for that change. The phases delimited by the summits contain additional political, military and geopolitical developments whose contributions are not separately estimated. The final phase is also the shortest of the three, and its coefficient is estimated over a correspondingly smaller number of months.

Standard errors are clustered at firm level in every specification of this section.

4.3.6 Robustness Design and Placebo Tests

The verifications described in this section vary one element of the design at a time and leave the remainder unchanged, and any change in the estimate is therefore attributable to that element alone. The estimates are reported in Chapter 5.

Five verifications bear on the baseline specification.

R1.1, unwinsorised outcome. The estimation is repeated on the raw series in place of the one winsorised at the first and the ninety-ninth percentile. The thresholds adopted in Section 4.2.4 are conventional and are not implied by the data. An estimate that moved appreciably once they are removed would rest on the treatment of a small number of observations rather than on the behaviour of the sample as a whole.

R1.2, fixed effects by area and month. The month fixed effect is replaced by a fixed effect for each combination of area and month. Under the month fixed effect of the baseline any movement confined to one of the two markets remains in the residual, and the two groups do not contain the same proportion of European firms. The area-month structure removes that movement in advance and leaves only the variation internal to each market. It is the structure required by the specification for the second hypothesis, for the reason given in Section 4.3.2, and its use on the baseline establishes whether the average effect is sensitive to it.

R1.3, lagged outcome. The return of the preceding month enters among the regressors. The baseline is static, and the estimate it produces would be open to question if the persistence of returns were strong enough to matter. The length of the panel allows a specification of this kind to be estimated without recourse to instrumental variables, for the reason given in Section 4.3.7.

R1.4, short observation window. The observation window is closed at the beginning of 2023. This restriction concentrates the estimate on the earlier phase of the geopolitical shock and reduces the influence of later armed conflicts and independent geopolitical escalations. It does not identify the direct effect of the invasion separately from the political and military response, sanctions, military assistance and other developments observed during 2022. Equation (4.6) addresses the temporal composition of the post-treatment period by separating later phases rather than excluding them.

R1.5, post-treatment period from March 2022. The beginning of the post-treatment period is moved forward by one month. February contains the days that precede the

invasion, and its assignment to the post-treatment period rests on the argument given in Section 4.2.1 rather than on the calendar. Excluding that month establishes whether the assignment governs the result.

R2.1, common month fixed effects. The specification for the second hypothesis is estimated with a month fixed effect common to the two areas, a version which does not contain the interaction between area and period required by the triple difference. It is reported in order to measure what the omission costs, since Section 4.3.2 establishes that the omission is not admissible.

R4.1 and R4.2, alternative reference periods. The dynamic specification is estimated against 2019 and against the average of the four years preceding the invasion, as described in Section 4.3.5. Its coefficients are differences with respect to the period taken as reference, and a reference that is itself unusual displaces the entire profile. The two alternatives establish which features of the profile belong to the data and which to the reference.

R4.3, pre-treatment coefficients by area. The examination described in Section 4.3.4 is repeated for the firms in each market. This verification assesses whether the departure is confined to one market or shared across both. A departure confined to one market would indicate that the pooled result arises from combining the two markets. A departure shared across both would concern the design as a whole.

R4.4, exclusion of years affected by other geopolitical and policy shocks. To examine whether the joint rejection of constant pre-treatment differentials is driven by years containing other shocks relevant to defence-sector valuations, the annual pre-treatment diagnostic is repeated after excluding 2011, 2014 and 2016. These years contain, respectively, NATO's military intervention in Libya [40], Russia's annexation of Crimea and the associated reassessment of NATO defence investment [9], and the unexpected United States presidential election, which generated heterogeneous firm-level equity-market reactions [41]. The verification is conducted for the full, European and United States samples.

P1 to P8, placebo dates. The treatment is assigned to dates on which no shock occurred. The dates are eight, the month of February of each year from 2012 to 2019, the first of them chosen so that every fictitious treatment is preceded by at least two years of observations. For each of them the sample is truncated before February 2022, and the

period following the fictitious treatment therefore contains no month of the real post-treatment period. The years 2020 and 2021 are not used, since the period discussed in Section 4.3.4 falls within them. What the exercise establishes is whether a date chosen arbitrarily can produce an effect comparable to the one estimated on the real date, in sign as well as in magnitude.

The assumption on which the identification rests cannot be settled by these verifications.

4.3.7 Econometric Considerations

The baseline parameter estimated in Section 4.3.1 is the average post-treatment return differential between defence and control firms. It is estimated by least squares conditional on firm and month fixed effects. Its interpretation as a causal effect remains conditional on the counterfactual assumptions described in Section 4.3.4.

The treatment date is the same for every treated firm, since the invasion reached the whole sector on 24 February 2022. Every comparison is therefore drawn between firms exposed to the shock and firms never exposed to it. The two-way fixed effects estimator has been shown to be a weighted average of two-by-two comparisons, and where units are treated at different moments some of those comparisons use already treated units as controls, which can attach negative weights to the average [36]. A design with a single treatment date and a never-treated comparison group does not contain comparisons in which already treated units serve as controls. The distortion associated with staggered treatment timing therefore does not arise. This property does not independently establish causal identification, which continues to depend on the counterfactual assumptions described in Section 4.3.4.

The baseline is static, and the bias that affects dynamic panels with fixed effects does not arise in it. That bias concerns the version reported as R1.3 among the robustness checks, where the lagged outcome enters among the regressors. It is of the order of the inverse of the number of periods [37], and with close to two hundred periods it is negligible, which is why that version requires no instrumental variables.

Two forms of dependence may arise in a panel of returns. Observations belonging to the same firm may be correlated over time, and clustering at firm level accommodates this serial dependence without imposing a specific structure. Returns of different firms

may also be correlated within the same month. Month fixed effects remove average movements common to the firms in the relevant comparison, although they do not necessarily eliminate every form of residual cross-sectional dependence. The reported inference is therefore conditional on the firm-level clustering structure adopted in the analysis [38].

The number of clusters varies across specifications. The baseline and the specification for the second hypothesis rest on the 136 firms of the sample, the third hypothesis on the 36 firms of the treatment group, and the examination by area on 65 and on 71 firms. Where the number of clusters is smaller, the finite-sample approximation underlying the correction is weaker. Cluster-robust standard errors may be downward biased in such settings, which can produce tests that reject too frequently.

Fixed effects form part of the identification strategy. Firm fixed effects absorb time invariant differences across companies, while month fixed effects absorb shocks common to all firms in a given month. In H2, fixed effects defined for each combination of area and month perform this function separately within the European and United States markets. A random effects specification would not reproduce this structure and is therefore not appropriate for the analysis.

The four hypotheses use related fixed effects specifications, but rest on different identifying assumptions, as discussed in Section 4.3.4. In the specifications that compare treated and control firms, the control group is selected according to eligibility based on conflict exposure. Section 4.4 follows a different approach. It constructs the counterfactual by matching return trajectories observed before treatment, so its validity depends on the quality of the match and on a different set of assumptions.

4.4 A Complementary Identification Strategy: The Synthetic Control Method

This section presents a second identification strategy, based on a synthetic counterfactual, which complements the design set out above. It describes the method and the outcome variable on which it is applied, the units and the donors on which the exercise is conducted, and the procedure by which the results are assessed.

4.4.1 Motivation and Method

The Difference-in-Differences design supports a causal interpretation only under the counterfactual assumptions examined in Section 4.3.4. A second strategy is therefore adopted whose identifying credibility depends on the quality of a constructed pre-treatment counterfactual. The method produces a separate counterfactual for each treated unit examined. These units comprise aggregates of European and United States defence producers and three illustrative European firm-level applications.

The synthetic control method constructs a counterfactual from a weighted combination of donor units [39]. The weights are non-negative and sum to one and are selected so that the synthetic combination reproduces the pre-treatment characteristics and valuation trajectory of the treated unit as closely as possible. In this application, the treated units are defence producers or an aggregate of defence producers. The donor firms do not possess the direct conflict-specific military-demand channel defining the treatment group, although they may remain exposed to the broader macroeconomic and financial consequences of the geopolitical shock. The post-treatment distance between the treated and synthetic series measures the divergence relative to this constructed counterfactual.

The outcome variable employed in this exercise is not the monthly return used in the preceding sections. It is a price index, obtained by cumulating the monthly returns of each unit and rebasing every series to one hundred in January 2010. The unwinsorised series is used, since the index has to follow the price actually observed and truncating the largest movements would let it drift away from it.

The reason for adopting a different outcome variable lies in what the method makes the donors reproduce. The weights are chosen so that the donors follow the series of the treated unit before February 2022, and what they end up following depends on which series is used. A sequence of monthly returns consists for the most part of movements that belong to a single month and are not repeated, so that the donors selected would be the firms that happened to move in the same direction as the treated unit in particular months, which carries no information about their behaviour thereafter. A price index instead traces the valuation path followed by the unit over the complete period and allows the donor combination to reproduce its pre-treatment dynamics. A close correspondence between the two paths provides evidence of empirical similarity before the treatment and

does not establish that the treated and donor firms were driven by identical underlying economic mechanisms.

The rebasing places all series on a common initial scale. Since the weights are non-negative and sum to one, the synthetic unit is a weighted average of the donor indices at every date. Share prices differ substantially in their original levels and are denominated in different currencies, which would make an exercise conducted on untransformed prices dependent on arbitrary differences in scale. Setting each index equal to one hundred in January 2010 removes this initial scale difference. Rebasing does not by itself guarantee that the complete predictor vector of the treated unit belongs to the convex hull of the donor units. The feasibility of the synthetic counterfactual is therefore assessed through predictor balance and the pre-treatment fit obtained from the data.

4.4.2 Treated Units and Donor Pool

Five treated units are examined. Two are equal-weighted aggregates of the European and United States defence firms with complete price histories. The other three are Rheinmetall, Thales and Leonardo, which provide illustrative firm-level applications of the method. These individual cases do not constitute a representative sample or a formal test of the distribution of the effect across the European defence industry.

Each regional aggregate is built by taking the mean of the unwinsorised monthly returns of the eligible defence firms in each month and cumulating those means into an index.

The method requires a complete series for every unit, since a weighted combination cannot be formed for a month in which one of the series is absent. Fincantieri is therefore left out of this part of the analysis, having been admitted to listing in 2014. The average unit is accordingly computed on fourteen of the fifteen European firms of the treatment group. The corresponding United States aggregate contains seventeen defence firms with complete price histories.

The European aggregate and the three individual European firms use the fifty European control firms as their donor pool. The United States aggregate uses the fifty United States control firms. In each case, the donor firms lack the direct military-demand channel defining the treatment group, although they may remain exposed to the broader macroeconomic and financial consequences of the geopolitical shock.

The principal specifications use regional donor pools because the synthetic-control method does not include time fixed effects capable of absorbing divergences between European and United States markets. Each treated unit is therefore paired with control firms from the same regional market. A robustness exercise subsequently opens the European donor pool to United States control firms.

The predictor set comprises seven pre-treatment values of the price index: December 2010, December 2012, December 2014, December 2016, December 2018, December 2020 and November 2021. These values are distributed across the complete pre-treatment period. The estimation uses January 2010 to January 2022 as the pre-treatment MSPE window, and the quality of the resulting fit is evaluated over all 145 monthly observations within that interval. Firm characteristics are not included as additional predictors because the available data for this exercise consist of stock-price information. The donor combination is therefore selected on the basis of the pre-treatment valuation trajectory rather than similarity in firms' industrial characteristics.

4.4.3 Inference and Stability

Conventional standard errors are not available for this design because each exercise examines one treated unit at a time. Inference evaluates how unusual the observed divergence is relative to placebo units drawn from firms without the direct conflict-specific military-demand channel. These firms may still have been exposed to the broader macroeconomic and financial consequences of the geopolitical shock.

The test statistic follows the post-treatment to pre-treatment prediction-error normalisation proposed by Abadie, Diamond and Hainmueller [39], but this application reports root mean squared prediction errors. The post/pre RMSPE ratio is the square root of the post/pre MSPE ratio used in the original study; this monotonic transformation leaves the ranking of the treated and placebo units unchanged. For each unit, the pre-treatment RMSPE measures the root mean squared distance between the observed and synthetic series before February 2022, while the post-treatment RMSPE measures the corresponding distance afterwards. Their ratio indicates how much the prediction error increases after treatment relative to the quality of the initial fit. This normalisation prevents a large post-treatment distance generated by a poor pre-treatment match from being interpreted as evidence of an unusual divergence.

The same exercise is repeated for each donor firm. Every donor is treated in turn as though it had been exposed to the treatment in February 2022, receives a synthetic replica constructed from the remaining donor firms and produces its own RMSPE ratio. For the individual treated firms, this procedure compares firm-level treated units with firm-level placebo units. For both regional aggregates, the benchmark distribution remains composed of individual donor firms from the corresponding market. Their ranks therefore indicate how unusual each aggregate divergence is relative to the regional donor firms and do not constitute exact randomisation tests.

The rank-based p-value is determined by the position of the treated unit within the placebo distribution. It is calculated as the number of placebo ratios at least as large as the observed ratio, plus the treated unit itself, divided by the total number of placebo units plus one. With fifty donors, the smallest attainable value is therefore one in fifty-one. This value represents the floor imposed by the size of the donor pool and should be interpreted as a placebo-based inferential benchmark.

A secondary restricted distribution excludes placebo runs whose pre-treatment RMSPE exceeds five times the pre-treatment RMSPE of the treated unit. These units are removed because the method does not reproduce their trajectories with a sufficiently comparable degree of accuracy. Results are reported for both the complete placebo distribution and the restricted distribution, so that the inferential conclusion does not depend solely on the treatment of poorly fitted placebo units.

The stability of the result is assessed by removing the donors one at a time. The estimation is repeated as many times as there are donors carrying a substantial weight, each time without one of them, and the set of paths that results shows whether the counterfactual depends on a single company. A synthetic unit resting on one donor is fragile, since the divergence it produces may reflect the fortunes of that company rather than those of the treated unit.

The synthetic-control analysis complements the Difference-in-Differences design. The Difference-in-Differences specification estimates an average monthly return differential across the firm-level sample. The synthetic-control analysis evaluates the price-index trajectories of equal-weighted regional aggregates against constructed counterfactuals. Each regional aggregate is evaluated with its own donor pool and placebo distribution, so the resulting causal interpretation applies to that aggregate within its regional

counterfactual. The causal transatlantic difference and the division of the total response between the invasion and the ensuing political and military response lie outside the parameters identified by these designs.

5 Results

This chapter reports the estimates obtained from the specifications presented in Chapter 4 and evaluates them in relation to the four hypotheses. Section 5.1 presents the average post-treatment return differential, Section 5.2 examines its geographical distribution, Section 5.3 assesses its relationship with pre-invasion defence exposure and Section 5.4 describes its evolution over the post-treatment period. Section 5.5 evaluates the identifying assumptions underlying the Difference-in-Differences estimates, while Section 5.6 presents the synthetic counterfactuals for the European and United States treated units.

All tables in this chapter report estimates obtained from the monthly firm-level panel described in Section 4.2. The dependent variable is the monthly logarithmic return, winsorised at the first and ninety-ninth percentiles, and the sample covers January 2010 to June 2026. Unless a table indicates otherwise, every specification includes firm and month fixed effects. Standard errors are clustered at firm level and reported in parentheses below the coefficients. Deviations from this common structure are stated in the notes to the individual table.

5.1 Average Post-Treatment Return Differential

This section evaluates H1 by estimating the average post-treatment return differential between defence and control firms. The baseline specification provides the central firm-level estimate, and the robustness exercises assess its stability across alternative outcomes, fixed-effect structures, dynamic specifications and treatment windows. This benchmark establishes the aggregate pattern examined across markets, firm-level exposure and time in the subsequent sections.

5.1.1 The Baseline Estimate

The baseline post-treatment return differential is estimated with equation (4.1), which relates the monthly logarithmic return of each firm to firm fixed effects, month fixed effects and the interaction between the treatment and post-treatment indicators.

The interaction coefficient compares the change in the average monthly return of defence firms after February 2022 with the corresponding change among the control firms, after accounting for permanent firm characteristics and movements common to the sample within each month. It is the empirical quantity corresponding to the average treatment effect on the treated targeted by H1. Its interpretation as a causal effect nevertheless requires the counterfactual assumptions examined in Section 5.5. The coefficient is an average monthly logarithmic-return differential and not a cumulative return over the complete post-treatment period.

Table 5.1 - Average post-treatment return differential, estimate of equation (4.1)

	(1) M1
β: Treat(i) x Post(t)	0.0166***
	(0.0031)
Observations	26,663
Firms (clusters)	136
Within R-squared	0.2514
RMSE	0.0673

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

The estimate of β is 0.0166, with a standard error of 0.0031. This corresponds approximately to 1.66 percentage points of additional average monthly return for the defence firms relative to the control group during the post-treatment period, conditional on firm and month fixed effects. The coefficient is 5.43 times its standard error, and the confidence interval runs from 0.0106 to 0.0226.

The residual standard deviation reported in Table 5.1 is 0.0673, which represents the scale of the monthly variation left unexplained by the two sets of fixed effects. The estimated differential is limited relative to the variation observed for an individual firm in a single month. Its precision reflects the repeated observations available across the panel and the aggregation of information across 136 firm clusters. Since inference is clustered at firm level, the 26,663 firm-month observations should not be interpreted as statistically independent observations.

If a monthly logarithmic-return differential of 0.0166 were sustained for twelve months, it would correspond to approximately 22.0 per cent in compounded terms. This calculation converts the monthly coefficient into annual units under a constant-

differential assumption. It does not constitute an estimate of a separate annual causal effect. The temporal stability of the differential is examined directly in Section 5.4.

Relative to the dispersion of the outcome, whose standard deviation is 0.0775 in the estimation sample, the coefficient corresponds to 0.21 standard deviations. The cumulative divergence between the defence firms and a constructed counterfactual is estimated in Section 5.6, where its statistical relevance is assessed through placebo-based rank inference.

H1 is supported. Defence firms recorded a positive and statistically significant post-treatment return differential relative to firms without the direct conflict-specific military-demand channel characterising the treatment group. The Difference-in-Differences specification establishes this differential within the observed sample. Its causal attribution is not secured by this design alone because of the evidence presented in Section 5.5 and is evaluated for the regional aggregate trajectories through the synthetic-control analysis of Section 5.6.

The outcome of equation (4.1) is a stock-market return. Prices can respond to information concerning future military expenditure and procurement before the corresponding contracts, production, revenues or profits become observable. The estimated coefficient directly measures a relative market repricing. Its timing is consistent with the expectation channel developed in Section 3.2, although the analysis does not directly observe investors' expectations, expected cash flows or realised procurement benefits.

A rise in price may follow either from an increase in expected cash flows or from a fall in the rate at which those flows are discounted, and a design built on prices alone does not separate the two. The estimate is also an average along two dimensions, the firms and the months, and neither of them is uniform, as Sections 5.2 and 5.4 show.

5.1.2 Robustness of the Estimate

The verifications described in Section 4.3.6 vary one element of the design at a time while leaving the remainder unchanged. Support for H1 requires the sign, statistical relevance and order of magnitude of the estimated differential to remain reasonably stable. Its exact magnitude need not remain unchanged because R1.4 and R1.5 alter the months included in the post-treatment comparison, and Section 5.4 shows that the differential varies over time.

Table 5.2 - Robustness of the average post-treatment return differential

Panel A: Outcome, fixed effects and dynamics

	(1) R1.1 No winso-risation	(2) R1.2 Area-month FE	(3) R1.3 Dynamic panel
Lagged monthly return			-0.0098
			(0.0069)
Treat(i) x Post(t)	0.0171***	0.0165***	0.0168***
	(0.0034)	(0.0031)	(0.0031)
Observations	26,663	26,663	26,527
Firms (clusters)	136	136	136
Within R-squared	0.2451	0.2815	0.2524
RMSE	0.0726	0.0662	0.0673

Panel B: Estimation window

	(4) R1.4 Short window	(5) R1.5 Post from Mar 2022
Treat(i) x Post(t)	0.0274***	
	(0.0032)	
Treat(i) x Post(t), from March 2022		0.0134***
		(0.0030)
Observations	21,087	26,663
Firms (clusters)	136	136
Within R-squared	0.2859	0.2507
RMSE	0.0644	0.0674

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Estimated on the unwinsorised series, R1.1 gives 0.0171 against 0.0166 in the baseline, a difference of 3 per cent. The thresholds adopted in Section 4.2.4 are conventional choices. Their removal changes the estimate by only 3 per cent, indicating that the result reflects the behaviour of the sample as a whole. The winsorised specification adopted as the baseline is also the more conservative of the two and has the smaller standard error, 0.0031 against 0.0034.

Under fixed effects for each combination of area and month, R1.2 gives 0.0165. The verification answers the objection that the average effect might reflect a movement common to one of the two markets in the months following the invasion. Absorbing the

monthly movement of each market separately leaves the coefficient unchanged at the third decimal. The same fixed-effect structure is not optional for the second hypothesis, for the reason given in Section 4.3.2, and it is worth establishing in advance that the average effect does not depend on it.

With the return of the preceding month among the regressors, R1.3 gives 0.0168. The coefficient on the lagged outcome is −0.0098 with a standard error of 0.0069 and is not distinguishable from zero. Monthly returns display no persistence in this sample, and the estimate of interest does not move. The bias that affects dynamic panels with fixed effects is negligible over close to two hundred periods, for the reason given in Section 4.3.7, and the specification requires no instrumental variables.

Closing the observation window at the beginning of 2023, R1.4 produces an estimate of 0.0274, which exceeds the baseline by 65 per cent. The restriction concentrates the estimate on the earlier phase of the geopolitical shock and excludes the later phases in which the differential becomes smaller. It also reduces the influence of subsequent armed conflicts and independent geopolitical escalations. The restricted window still includes the invasion and the political and military response generated by it, including sanctions, military assistance and revisions in expected military expenditure. It therefore does not identify the direct effect of the invasion separately from the broader treatment mechanism.

Moving the beginning of the post-treatment period to March 2022, R1.5 gives 0.0134, which falls short of the baseline by 19 per cent. February 2022 contains the days that precede and start the invasion, and its assignment to the post-treatment period rests on the argument given in Section 4.2.1. The resulting reduction demonstrates the sensitivity of the average post-treatment differential to the inclusion of February 2022. It does not identify the separate causal contribution of that month.

The coefficient remains positive and significant at the one per cent level in all five verifications, with a standard error between 0.0030 and 0.0034 throughout. The two displacements of any size, in R1.4 and in R1.5, both concern the period over which the differential is estimated and are accounted for by the profile of Section 5.4. No verification alters the conclusion drawn on the first hypothesis.

5.2 Geographical Heterogeneity of the Return Differential

This section evaluates H2 by examining the distribution of the post-treatment return differential across the European and United States market groups. The preferred specification estimates each defence-control comparison within its respective regional market environment and identifies the European premium directly within the model. The resulting comparison characterises the geographical dimension of the response and provides the regional basis for the synthetic-control analysis presented later in the chapter.

5.2.1 The Difference between the Two Markets

The second hypothesis is tested with equation (4.2), which differs from the baseline in the structure of the fixed effects and in the interaction terms. A fixed effect is estimated for each combination of area and month, while the interaction carrying the baseline post-treatment differential is interacted with the indicator for a European headquarters.

Three quantities follow from the specification. The coefficient β_1 is the estimated post-treatment differential for the United States firms. The coefficient β_2 measures the additional European premium relative to that differential. The estimated differential for the European firms is therefore $\beta_1 + \beta_2$. H2 concerns β_2, which is estimated directly with its own standard error.

Table 5.3 - Post-treatment return differential by area, estimate of equation (4.2)

	(1) M2 - H2 EU vs US
β_1: Treat(i) x Post(t)	0.0043
	(0.0031)
β_2: Treat(i) x Post(t) x Europe(i)	0.0281***
	(0.0052)
Observations	26,663
Firms (clusters)	136
Within R-squared	0.2828
RMSE	0.0662

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

The estimated differential for the United States firms is 0.0043, with a standard error of 0.0031. The additional European premium is 0.0281, with a standard error of 0.0052, and differs from zero at the one per cent level. The estimated differential for the European

firms, obtained as the sum of the two coefficients, is 0.0323 with a standard error of 0.0041. It is 7.85 times its standard error, and its confidence interval ranges from 0.0242 to 0.0405.

Relative to the dispersion of the outcome, whose standard deviation is 0.0775 in the estimation sample, the estimated differential for the European firms corresponds to 0.42 standard deviations and the estimated differential for the United States firms to 0.055. The two markets are separated by a factor of eight, and the European differential is twice the average differential estimated in Section 5.1.

The estimated differential for the United States firms is not statistically distinguishable from zero. Its confidence interval ranges from −0.0020 to 0.0105 and therefore excludes both the European estimate of 0.0323 and the average estimate of 0.0166 obtained from the full sample. The evidence remains compatible with small positive or negative differentials and excludes a differential of the magnitude recorded in Europe. The design therefore provides insufficient statistical precision to determine the sign and exact magnitude of the underlying United States post-treatment differential.

The composition of the treatment group bears on how the average of Section 5.1 should be read. Of the 1,908 treated firm-months in the post-treatment period, 1,113 belong to United States firms and 795 to European ones, in a proportion of 58 to 42 per cent. The average effect is therefore drawn towards the market in which the effect is small, and it understates the differential recorded by the European firms by a factor of about two.

Figure 5.1 - Post-treatment return differential by area, with 95 per cent confidence intervals

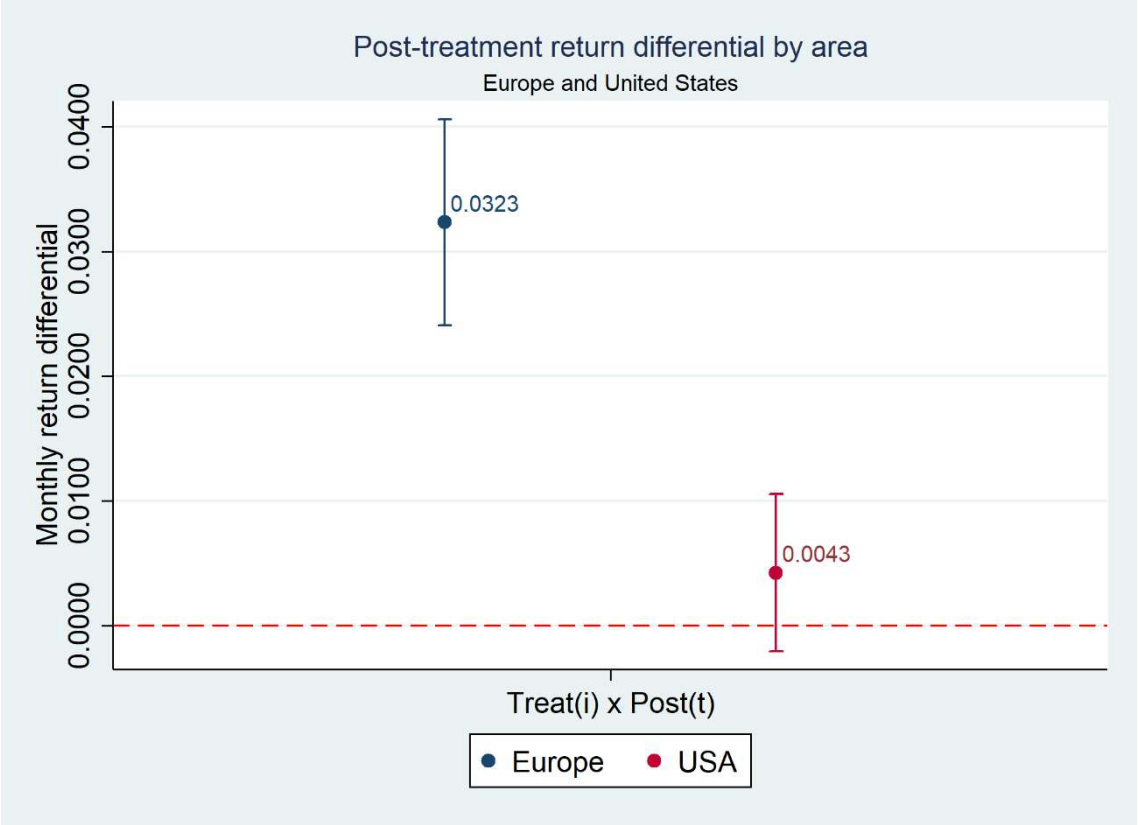


Figure 5.1 reports the two estimated differentials of Table 5.3 in graphical form, each with its 95 per cent confidence interval. What it shows is the distance between the two markets set against the precision with which each of them is measured.

H2 is supported as a proposition concerning geographical heterogeneity in the estimated post-treatment differential. The difference between the European and United States estimates is positive, statistically significant and larger than the average differential obtained from the full sample. Since this comparison is estimated through the Difference-in-Differences specification, its causal interpretation inherits the parallel-trends limitation examined in Section 5.5. The synthetic-control evidence presented in Section 5.6 supports a causal interpretation of the European aggregate response, while no statistically distinguishable causal effect is detected for the United States aggregate. The two regional exercises do not identify the transatlantic difference as a separate causal parameter.

The geographical pattern is consistent with the geopolitical shock and the associated political and military response producing a larger revision of defence-spending and security expectations in Europe. European countries entered the shock from a lower level of military expenditure and from the institutional position developed after the end of the

Cold War, while the United States had a larger and more continuously funded defence market. This difference provides a plausible interpretation of the stronger European repricing. The analysis does not directly observe investors' expectations or identify the relative contribution of the underlying mechanisms.

Pre-invasion valuation levels provide a second possible explanation. European defence producers had accumulated a lower valuation than their own market over the decade preceding the invasion, and the same information could produce a larger proportional movement in a security whose price already incorporates modest expectations. Initial valuation differences may therefore contribute to the observed transatlantic estimate. Estimating the relative contributions of initial valuations and information revisions would require direct measures of expectations, which lie outside the available price data. Chapter 6 includes this limitation.

The estimated post-treatment differential is not uniform across the two markets. Whether it is uniform within the treatment group, across firms that differ in the intensity of their engagement in the defence business, is the object of Section 5.3.

5.2.2 Robustness of the Estimate

The verification described in Section 4.3.6 varies the structure of the fixed effects and leaves the remainder of the specification unchanged. Section 4.3.2 establishes that a triple difference requires the interaction between area and period, and that a month fixed effect common to the two areas does not contain it. R2.1 estimates equation (4.2) without that term, and it is the only verification reported for this hypothesis, since the remaining elements of the design are those of the baseline and are examined in Section 5.1.2.

Table 5.4 - Area-specific return differentials under common month fixed effects, verification R2.1

	(1) R2.1 H2 common month FE
β_1: Treat(i) x Post(t)	0.0081***
	(0.0029)
β_2: Treat(i) x Post(t) x Europe(i)	0.0203***
	(0.0041)
Observations	26,663
Firms (clusters)	136
Within R-squared	0.2522
RMSE	0.0673

*Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1*

Under a month fixed effect common to the two areas, the European premium falls from 0.0281 to 0.0203, a reduction of 28 per cent, while the estimated differential for the United States firms rises from 0.0043 to 0.0081 and becomes statistically significant at the one per cent level. The sign and statistical significance of the European premium are preserved. Its magnitude and the interpretation of the United States coefficient are nevertheless sensitive to the fixed-effect structure.

The displacement of the coefficients has a single source. Under a month fixed effect common to the two areas the treated firms of each market are compared with a control baseline drawn from both, which lies below the United States controls and above the European ones over a period in which the European market performed worse than the United States one. Each of the two estimated differentials is displaced by about 0.0039, upwards for the United States firms and downwards for the European ones. The premium is the difference between the two estimated differentials and absorbs both displacements, which is why it moves by 0.0078, the reduction of 28 per cent reported above. The same omission leaves unchanged the single average effect of R1.2, reported in Section 5.1.2, where the two displacements fall on one coefficient and offset one another instead of adding. Table 5.4 quantifies the importance of the omitted area-by-period component. Its omission reduces the European premium by 0.0078, corresponding to 28 per cent of the preferred estimate, and increases the United States differential by 0.0039, an amount close to the coefficient itself. The smaller standard errors obtained from the restricted specification reflect the estimation of fewer parameters and do not make it the preferred model.

The area-by-month fixed effects remain necessary for the within-market comparison required by H2.

5.3 The Intensity of Exposure

The third hypothesis is tested with equation (4.3), estimated on the 36 firms of the treatment group with firm and month fixed effects. The exposure measure does not vary over time, and its main effect is absorbed by the firm fixed effects. The interaction between exposure and the post-treatment period estimates whether changes in returns after February 2022 vary systematically with firms' pre-invasion dependence on arms sales. The coefficient γ therefore represents the estimated firm-level exposure gradient. A causal interpretation would require that counterfactual changes in returns were unrelated to the selected exposure measure.

Table 5.5 - Post-treatment return differential by exposure intensity, estimate of equation (4.3)

VARIABLES	(1) M3 - H3 intensity
1.post#c.def_share	0.0024
	(0.0094)
Observations	6,963
Number of firm_id	36
Within R-squared	0.3506
RMSE	0.0641

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Since the measure of exposure is a share, γ is expressed as a differential of monthly return per unit of that share, and the quantity that can be read directly is the differential implied between two producers, obtained by multiplying γ by the distance between their exposures, as Section 4.3.3 sets out.

The evidence does not support the third hypothesis. Over a variation of one standard deviation in exposure, which is 0.257 among the treated firms, the implied differential lies between −0.43 and +0.55 percentage points per month. Between the least and the most exposed firms in the sample, whose exposures are 0.111 and 0.952, the interval widens to between −1.40 and +1.81 percentage points per month. The estimates therefore provide no evidence of a systematic gradient measurable through the pre-treatment share

of revenues derived from arms sales. Their precision is nevertheless insufficient to exclude economically substantial differences between firms at the extremes of the observed exposure distribution.

The exposure is distributed almost identically in the two markets, with a mean of 0.512 among the European firms and of 0.510 among the United States ones.

The support of the measure is truncated on the left at ten per cent by the selection rule described in Section 4.2.3. The comparison is therefore limited to firms that are already substantially engaged in the defence business and establishes nothing about producers below that threshold. The specification contains 36 firm clusters, making its cluster-robust inference less reliable than that of the full-sample specifications. The standard errors may be downward biased in a sample of this size. Greater finite-sample uncertainty would not convert the non-rejection into support for H3, although it would widen the confidence interval and further limit the relationships that the analysis can exclude.

Two mechanisms are compatible with the estimate. First, investors may have revised expectations for the sector-wide demand environment, producing a broadly shared valuation response across defence producers. Second, firm-level exposure may depend on characteristics omitted from the pre-invasion arms-sales share. The ability to capture additional orders depends on productive capacity, product composition and customer geography. The estimated coefficient therefore applies specifically to the selected arms-sales-share measure.

5.4 The Evolution of the Return Differential over Time

The baseline specification estimates a single coefficient over the complete post-treatment period and cannot distinguish an initial repricing followed by no further return differential from a sequence of continuing positive differentials. The first pattern can leave a persistent difference in price levels, while the second implies that the price-level divergence continues to widen. The two decompositions presented in Section 4.3.5 examine this distinction by elapsed time and by event phase.

5.4.1 The Profile of the Return Differential

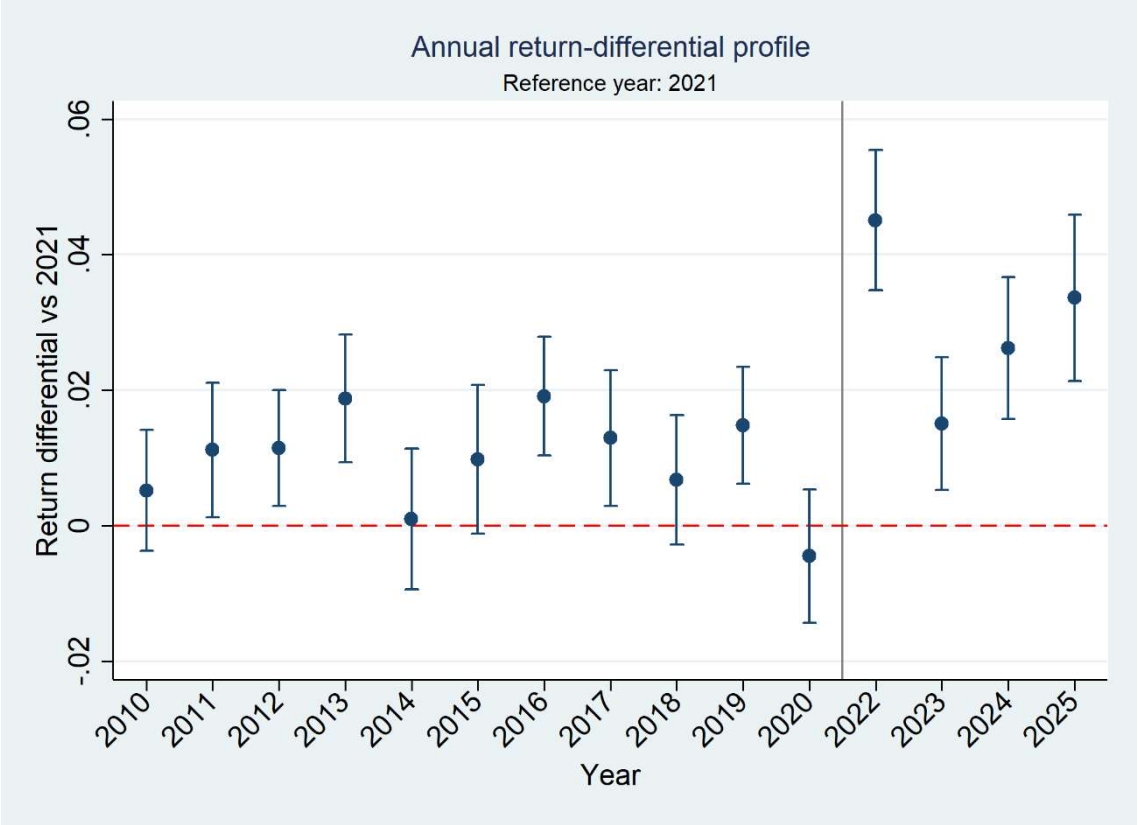
The dynamic specification is estimated at two frequencies. The annual version covers the whole observation window and takes 2021 as its reference year. The monthly version covers the interval from February 2021 to February 2024 and uses November 2021 as its reference month. This reference avoids using January 2022, when anticipatory adjustment may already have occurred, while retaining December 2021 and January 2022 among the estimated pre-treatment coefficients. The coefficient of 2026 is estimated but not reported, since the year is incomplete. The coefficients of the years preceding the invasion serve a different purpose and are examined in Section 5.5. The complete set of annual coefficients, including the years preceding the invasion and the estimates obtained against alternative reference periods, is reported in Table B.1, and the complete set of monthly coefficients in Table B.2.

Table 5.6 - Annual return differentials in the post-treatment period, estimate of equation (4.5)

VARIABLES	(1) M4 - Annual return differential profile post-treatment years
Treat(i) x Year 2022	0.0451***
	(0.0052)
Treat(i) x Year 2023	0.0151***
	(0.0050)
Treat(i) x Year 2024	0.0262***
	(0.0053)
Treat(i) x Year 2025	0.0336***
	(0.0062)
Observations	26,663
Number of firm_id	136
Within R-squared	0.2546
RMSE	0.0672

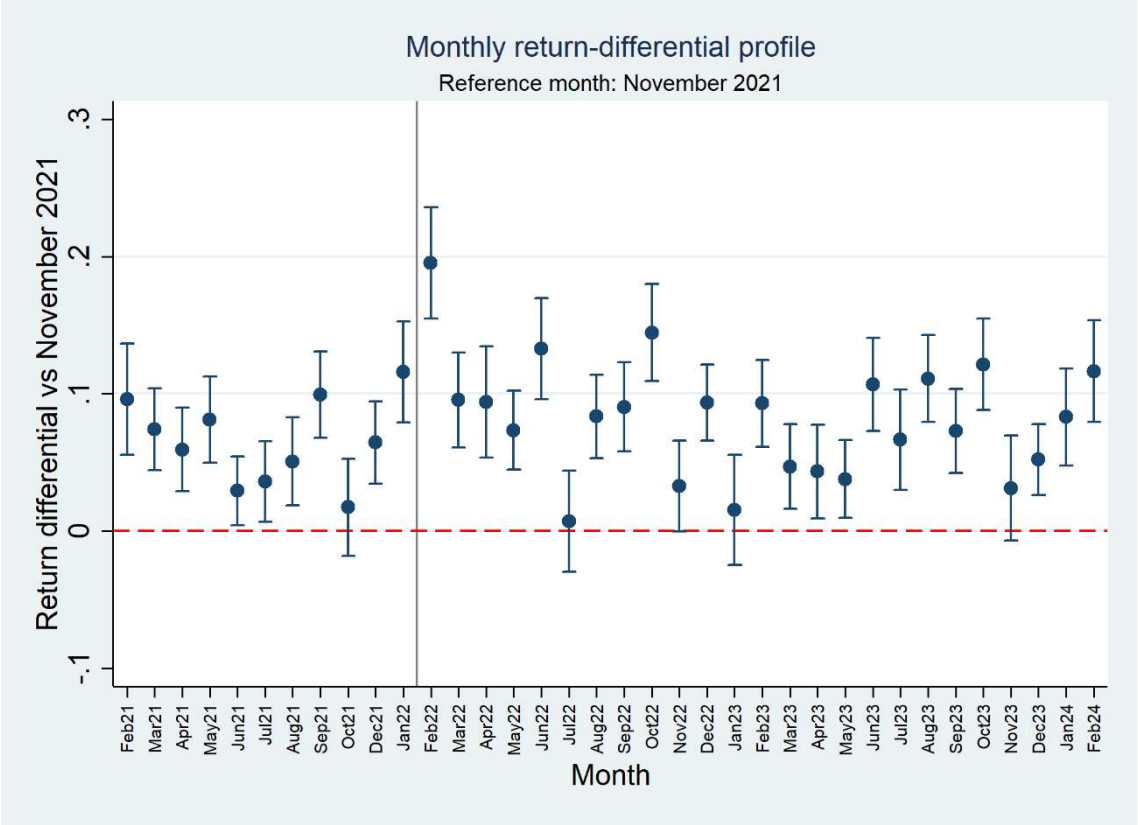
Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Figure 5.2 - Annual return-differential profile, reference year 2021



The annual post-treatment differential is largest in 2022 and remains positive relative to the 2021 reference year throughout the subsequent complete years. Each of the four post-treatment coefficients differs from zero at the one per cent level. The profile records a large differential in 2022, a reduction in 2023 and positive differentials in 2024 and 2025. These coefficients concern monthly returns relative to the reference year and do not measure the persistence of a price-level difference.

Figure 5.3 - Monthly return-differential profile, reference month November 2021



The monthly resolution shows what the annual coefficient of 2022 compresses into a single number. The differential is 0.1954 in February 2022, the month of the invasion, and it is the largest coefficient of the whole window. It is 0.0955 in March, 0.0940 in April, 0.0735 in May and 0.1328 in June, all four differing from zero at the one per cent level. In July 2022 it is 0.0070 with a standard error of 0.0186 and is no longer distinguishable from the reference month. The reaction is therefore concentrated in the month of the invasion and in the four that follow it, after which the monthly differential returns to values comparable with those recorded before the shock, while the annual profile shows that a smaller positive differential persists in the following years.

Two cautions govern how this profile is read, and both concern the reference month. The coefficients estimated on the months preceding the invasion are all positive and of considerable size, from 0.0962 in February 2021 to 0.0174 in October 2021, which places the level of the whole profile above zero irrespective of the shock. The level is fixed by a single month, and November 2021 was itself a month in which the treated group performed poorly relative to the controls. The monthly version is therefore informative about

the shape of the reaction and not about its level, and it is the same feature that produces the joint test reported in Section 5.5.

Within the profile, the comparison between January 2022 and the preceding months does not depend on the selected reference. The coefficients from February to October 2021 average 0.0605, December 2021 stands at 0.0645 and January 2022 rises to 0.1161. This pattern is consistent with anticipatory adjustment before the invasion. Its causal origin remains unidentified at monthly frequency. The earlier coefficients of 2021 form part of the pre-treatment discrepancy examined in Section 5.5.

5.4.2 The Decomposition by Event

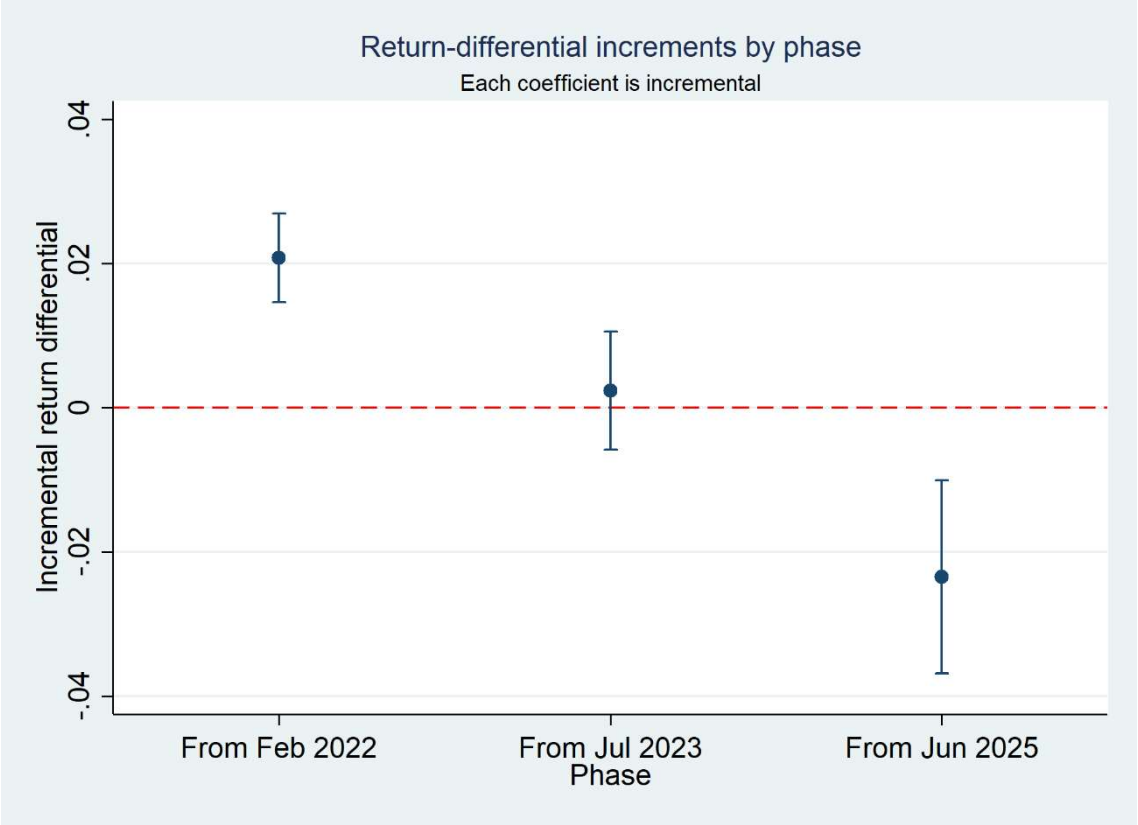
The three indicators of equation (4.6) are cumulative, and each coefficient is an increment with respect to the phase that precedes it. The differential prevailing within a phase is the sum of the coefficients up to that point, and its standard error is computed on the linear combination.

Table 5.7 - Incremental return differentials by phase, estimate of equation (4.6)

VARIABLES	(1) M5 - Phase de-composition
Treat(i) x Phase 1 (Feb 2022 to Jun 2023)	0.0208***
	(0.0031)
Treat(i) x Phase 2 increment (from Jul 2023)	0.0023
	(0.0041)
Treat(i) x Phase 3 increment (from Jun 2025)	-0.0235***
	(0.0068)
Observations	26,663
Number of firm_id	136
Within R-squared	0.2522
RMSE	0.0673

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Figure 5.4 - Incremental return differentials by phase, with 95 per cent confidence intervals



The coefficient for the phase beginning in February 2022 is 0.0208, with a standard error of 0.0031. The additional increment beginning in July 2023 is 0.0023, with a standard error of 0.0041, and is not statistically distinguishable from zero. The increment beginning in June 2025 is −0.0235, with a standard error of 0.0068, and differs from zero at the one per cent level. These dates delimit the phases of the specification and do not identify the separate causal effect of the invasion or of either NATO summit.

The differentials prevailing in the three phases follow from those increments. Between February 2022 and June 2023 the differential is 0.0208, with a confidence interval from 0.0146 to 0.0269. Between July 2023 and May 2025 it is 0.0231, with an interval from 0.0139 to 0.0323. From June 2025 to the end of the window it is −0.0004, with a standard error of 0.0055 and a p-value of 0.944, and its interval runs from −1.13 to +1.05 percentage points per month.

Two features limit the interpretation of the final phase. It contains thirteen months and is the shortest of the three phases. Its interval is compatible with both a modest positive differential and a modest negative differential. The specification identifies a change in the prevailing return differential at the selected boundary without identifying the event

responsible for that change. Each phase also contains additional political, military and geopolitical developments whose contributions are not separately estimated.

The temporal pattern is consistent with successive revisions of expectations after the geopolitical shock generated by Russia's invasion of Ukraine. The design does not establish this mechanism directly. From June 2025 onwards, the estimated differential is −0.0004 per month and is not statistically distinguishable from zero. This result is consistent with the exhaustion of the return differential and provides no evidence that the price-level divergence accumulated during the preceding period was reversed.

The persistence of positive differentials during the earlier phases is consistent with information reaching the market over time. Concurrent geopolitical developments, subsequent policy measures and changes in the returns required by investors provide additional compatible mechanisms. The available price data identify the combined valuation response to these mechanisms. Their individual contributions lie outside the identifying scope of the design.

H4 is therefore supported. The annual return differential remains positive and statistically significant in 2023, 2024 and 2025, satisfying the prediction that the post-invasion differential remains observable beyond 2022. The monthly and event-based decompositions locate this persistence within the post-treatment period: the initial adjustment is concentrated between February and June 2022, positive differentials remain observable through May 2025, and the estimate becomes statistically indistinguishable from zero from June 2025 onwards. The causal interpretation of this temporal pattern remains conditional on the identifying evidence examined in Section 5.5.

5.5 The Identifying Assumption

The Difference-in-Differences estimates reported in Sections 5.1 to 5.4 rely on the counterfactual conditions presented in Section 4.3.4. In the absence of the geopolitical shock generated by Russia's invasion of Ukraine, including the political and military response through which it affected financial markets, the return differential between defence and control firms must have remained constant. The design also requires that the treatment mechanism exerted no influence before February 2022. Neither counterfactual condition

can be observed directly, but their implications can be examined using the pre-treatment coefficients.

5.5.1 The Departure from Parallel Trends

The pre-treatment coefficients do not support the implication of parallel trends under any of the reference periods considered. The departure is visible in the full sample and in the two markets examined separately, although the evidence is weaker for the European firms. What remains to be established is its magnitude and its position in time. Statistical rejection and economic magnitude are distinct questions: a joint test asks whether the pre-treatment coefficients are collectively equal to zero, whereas a long panel may identify a departure that is statistically detectable but economically modest. The profile and magnitude of the departure are therefore examined first, followed by the joint tests.

Table 5.8 - Pre-treatment coefficients against two reference periods, extract of Table B.1

VARIABLES	(1) M4 - Annual event study base 2021	(3) R4.1 Event study base 2019
yr2010_t	0.0052	-0.0096**
	(0.0045)	(0.0047)
yr2011_t	0.0112**	-0.0036
	(0.0050)	(0.0052)
yr2012_t	0.0115***	-0.0033
	(0.0043)	(0.0044)
yr2013_t	0.0188***	0.0039
	(0.0048)	(0.0049)
yr2014_t	0.0010	-0.0138**
	(0.0053)	(0.0056)
yr2015_t	0.0098*	-0.0050
	(0.0055)	(0.0051)
yr2016_t	0.0191***	0.0043
	(0.0044)	(0.0040)
yr2017_t	0.0129**	-0.0019
	(0.0051)	(0.0054)
yr2018_t	0.0068	-0.0081*
	(0.0048)	(0.0048)
yr2019_t	0.0148***	
	(0.0044)	
yr2020_t	-0.0045	-0.0193***
	(0.0050)	(0.0046)
Yr2021_t		-0.0148***
		(0.0044)
yr2022_t	0.0451***	0.0303***
	(0.0052)	(0.0055)

yr2023_t	0.0151***	0.0002
	(0.0050)	(0.0054)
yr2024_t	0.0262***	0.0114*
	(0.0053)	(0.0060)
yr2025_t	0.0336***	0.0188***
	(0.0062)	(0.0068)
Observations	26,663	26,663
Number of firm_id	136	136
Within R-squared	0.2546	0.2546
RMSE	0.0672	0.0672

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Column 1 is estimated with 2021 as the reference year and column 2 with 2019, and the cell corresponding to the reference year of each column is empty by construction. The post-treatment coefficients are reported in Table 5.6. The complete monthly profile is reported in Table B.2, and the estimates obtained against the mean of the four years preceding the invasion are reported in Table B.1. Standard errors clustered at firm level in parentheses.

The sign of those coefficients is a property of the reference period, as Section 4.3.4 anticipates. Against 2021 ten of the eleven are positive, and against 2019 nine of the eleven are negative, since the two specifications differ by a constant equal to the coefficient of the year that has been displaced. Only the distances between the coefficients are invariant, and the discussion that follows rests on those distances.

The number of coefficients that differ from zero individually also depends on the selected reference year. Against 2021, six of the eleven coefficients differ from zero at the five per cent level and one further coefficient at the ten per cent level. Against the pre-pandemic reference, four differ from zero, including 2020 and 2021. The profile indicates a discrepancy of modest size distributed across the full pre-treatment period. The statistical significance assigned to an individual year depends on the year used as its reference.

Section 4.3.4 identifies the two years of the pandemic as the period in which the two groups may have been affected unequally. Against the pre-pandemic reference the nine years from 2010 to 2018 average 0.0059 in absolute value, whereas 2020 stands at −0.0193 and 2021 at −0.0148, between two and a half and three times as large. Both are negative, which is to say that the defence firms lost ground on the controls over those two

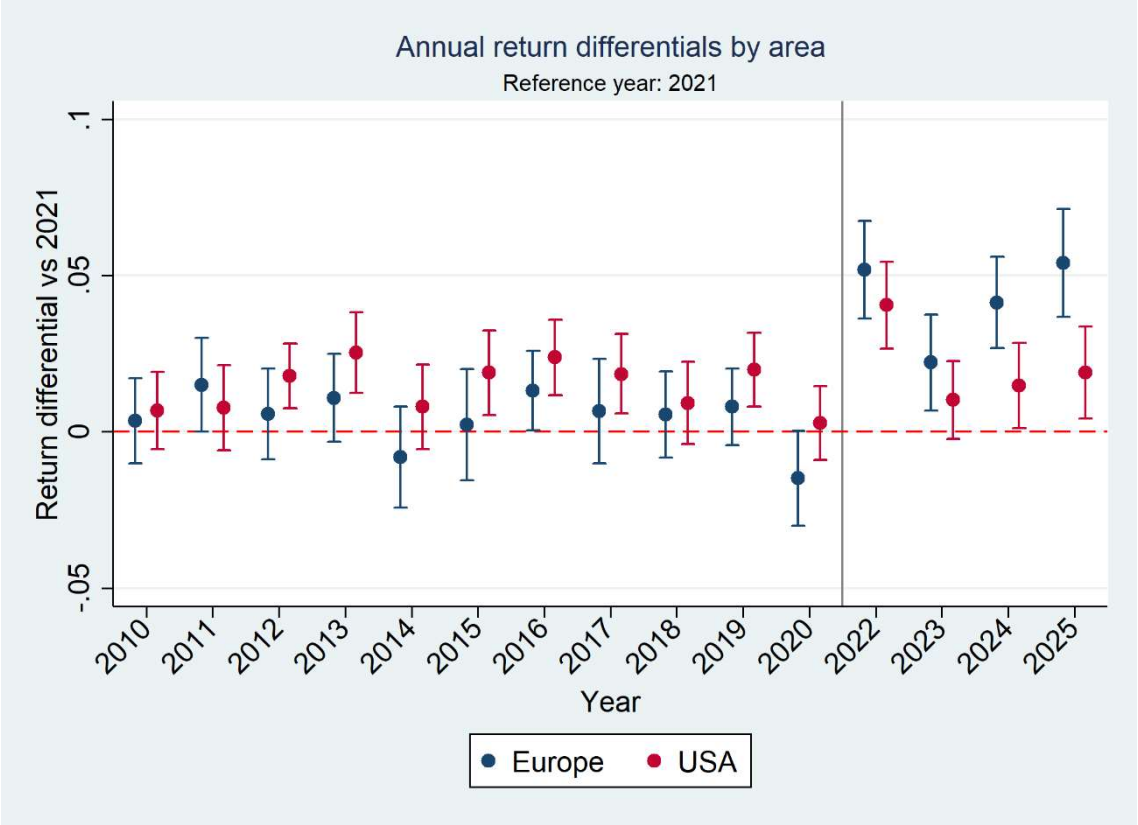
years. The discrepancy is nevertheless not confined to those two years, which is the case Section 4.3.4 treats as the less damaging of the two. It is present over the whole of the observation window and is composed, over the decade, of coefficients that are individually small and share a direction.

The magnitude of the pre-treatment discrepancy can be compared descriptively with the post-treatment coefficients. Before the invasion, the coefficients average 0.0105 in absolute value, compared with 0.0300 afterwards, producing a ratio of 2.9. Against the 2019 reference, the corresponding values are 0.0059 and 0.0152, producing a ratio of 2.6. These comparisons have no associated standard errors and depend on absolute coefficients whose signs vary with the reference year. They indicate that the pre-treatment discrepancy is economically relevant but do not measure the portion of the post-treatment differential attributable to the geopolitical shock.

The joint tests confirm that reading and locate the discrepancy in time. Tested as one set against 2021, the eleven pre-treatment coefficients give $F(11,135) = 4.10$, with a p-value below 0.0001. Tested against 2019 they give the same statistic. The restriction is in both cases that the differential was constant, and a Wald statistic does not respond to a change of normalisation. Dropping 2021 from the sample leaves the decade unchanged as well, at $F(9,135) = 2.95$. No test therefore owes its rejection to the choice of reference year. What the two tests against 2019 separate is the source of the rejection. Over the decade the statistic is $F(9,135) = 2.95$, with $p = 0.0031$, and seven of those nine years do not differ from zero individually. There the rejection is produced by the direction the coefficients share. Over the two pandemic years the statistic is $F(2,135) = 10.66$ on two degrees of freedom, and there it is produced by their size.

Two tests require additional qualification. The monthly test includes January 2022, whose coefficient is consistent with anticipatory adjustment, so part of its rejection may concern information related to the treatment under examination. The tests conducted separately by area also rely on fewer clusters than the full-sample specification. For the European firms, $F(11,64) = 1.99$ with $p = 0.0440$, while for the United States firms $F(11,70) = 3.51$ with $p = 0.0006$. Inference based on 65 and 71 clusters is less secure than full-sample inference, but the analysis cannot establish that finite-sample bias explains either result. The evidence of a pre-treatment discrepancy is therefore weaker in Europe and considerably stronger in the United States.

Figure 5.5 - Annual return differentials by area, reference year 2021



The verification designated R4.3 in Section 4.3.6 was the one most likely to leave the design intact. A discrepancy confined to one market would indicate that the pooled result arises from the combination of the two. Figure 5.5 shows it in both, and Section 4.3.6 assigns that case to the design itself. The discrepancy is also better established in the market where Section 5.2 finds no statistically detectable post-treatment differential, which is the United States.

The additional sensitivity check designated R4.4 in Section 4.3.6 repeats the annual pre-treatment test after excluding 2011, 2014 and 2016. The hypothesis of constant pre-treatment differentials continues to be rejected in the full sample, with $F(8,135) = 4.26$ and $p = 0.0001$, and in the United States sample, with $F(8,70) = 3.77$ and $p = 0.0010$. For Europe, the statistic is $F(8,64) = 1.93$ with $p = 0.0707$, compared with $p = 0.0440$ in the complete pre-treatment sample. The European test therefore no longer rejects at the five per cent level, although it remains borderline and rejects at ten per cent. This verification weakens the evidence against parallel trends for Europe but does not validate the pooled or United States comparisons.

5.5.2 Anticipation

The monthly evidence is inconsistent with a strict absence of anticipatory adjustment. January 2022 records a coefficient of 0.1161, while the nine preceding months of 2021 average 0.0605. The concentration of Russian forces and the public discussion of a possible invasion provided information that could have affected prices before February 2022. December 2021 remains comparable with the earlier months, while the increase is concentrated in January. This pattern is consistent with anticipation without establishing that the geopolitical shock was the only cause of the January movement.

If the January coefficient reflects anticipatory repricing caused by the same geopolitical shock, its inclusion in the pre-treatment period raises the average return of the defence firms before the formal treatment date and attenuates the baseline coefficient. Under this interpretation, the estimated post-treatment differential understates the total market response. The presence of pre-treatment discrepancies and the possibility of other influences on January returns prevent the coefficient from being treated as a formal lower bound.

5.5.3 The Placebo Dates

The treatment is assigned to eight dates on which no shock occurred, the month of February of each year from 2012 to 2019. In each case the sample is truncated before February 2022, and 19,455 observations enter the estimation. The design admits no more than eight dates. Before 2012 the fictitious treatment would not be preceded by two years of observations, and 2020 and 2021 are excluded for the reason given in Section 4.3.6. The exercise examines whether an arbitrarily selected placebo date can produce a post-treatment differential comparable with the one estimated from February 2022.

Table 5.9 - Return differentials estimated at eight placebo dates, verifications P1 to P8

VARIABLES	(1) Feb 2012	(2) Feb 2013	(3) Feb 2014	(4) Feb 2015
Treat(i) x Post(t), placebo February 2019	0.0006	-0.0000	-0.0042**	0.0004
	(0.0028)	(0.0025)	(0.0021)	(0.0021)
Observations	19,455	19,455	19,455	19,455
Number of firm_id	136	136	136	136
Within R-squared	0.2685	0.2685	0.2686	0.2685
RMSE	0.0635	0.0635	0.0635	0.0635

Table 5.9 (continued)

VARIABLES	(5) Feb 2016	(6) Feb 2017	(7) Feb 2018	(8) Feb 2019
Treat(i) x Post(t), placebo February	-0.0002	-0.0042**	-0.0065***	-0.0067**
	(0.0023)	(0.0019)	(0.0023)	(0.0026)
Observations	19,455	19,455	19,455	19,455
Number of firm_id	136	136	136	136
Within R-squared	0.2685	0.2686	0.2688	0.2688
RMSE	0.0635	0.0635	0.0635	0.0635

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Four of the eight placebo coefficients differ from zero at the five per cent level, and all four are negative. None produces a positive differential comparable with the estimate of 0.0166 obtained from February 2022. The high rejection rate strengthens the evidence of pre-treatment instability. The later placebo windows assign increasing weight to the pandemic years, while the February 2014 window begins in a year in which the treated firms already underperformed the controls. The placebo estimates therefore reproduce the discrepancy documented by the dynamic coefficients and provide no independent validation of the causal interpretation.

The findings of Section 5.5 apply to the Difference-in-Differences estimates reported in Sections 5.1 to 5.4. These coefficients establish post-treatment differentials between the two groups but do not independently identify the causal effect of the geopolitical shock. The pre-treatment discrepancy is economically relevant relative to the baseline estimate. Its implications depend on whether the earlier difference would have persisted

or whether the pandemic years represented a temporary contraction followed by recovery. The available Difference-in-Differences evidence does not determine the direction of this potential bias. Section 5.6 therefore evaluates the European and United States aggregate responses through separate counterfactuals based on their respective pre-treatment trajectories.

The second reading requires an estimate of the path the treated firms would have followed without the shock. Section 4.4 constructs this counterfactual from observed pre-treatment donor trajectories. Section 5.6 reports the resulting estimates and the conditions supporting their interpretation.

5.6 The Synthetic Counterfactual

Section 5.5 rejects the identifying implication of parallel trends, leaving the coefficients of Sections 5.1 to 5.4 as estimates of post-treatment differentials. The synthetic-control analysis evaluates whether the aggregate trajectories observed in Europe and the United States diverge from counterfactuals constructed from their respective pre-treatment paths. The resulting evidence remains conditional on the donor pools, predictor set and inferential procedure described in Section 4.4.

The regional synthetic-control exercises complement H1 and H2. Each treated unit receives a separate counterfactual and is evaluated within its corresponding placebo distribution. The regional estimates therefore do not constitute a direct causal test of the difference between Europe and the United States. The method does not identify the firm-level exposure gradient examined by H3 or the temporal decomposition examined by H4.

Five treated units are examined: the European and United States aggregates and three individual European producers.

5.6.1 The Average European Producer

The treated unit is the average of the fourteen European defence firms whose price series covers the whole window, as defined in Section 4.4.2. Before February 2022 the donors reproduce its index with a root mean squared error of 10.24 index points, which is 6.6 per cent of the average level of the series. No individual firm in this chapter is reproduced to better than 12.2 per cent. The difference follows from the construction of the unit.

Averaging fourteen series removes the movements specific to each firm, and the component that remains is the one a combination of donors can track.

That error acquires a scale from the distance it is used to measure. At January 2022, the last month before the treatment, the treated unit stands 4.96 index points below its replica. At the end of the observation window the two series stand 373.6 points apart. The residual of the construction is therefore one seventy-fifth of the divergence reported below. The requirement stated in Section 4.4.3 is met for this unit.

The method matches the two series on seven quantities, and none of them is a characteristic of the firms. The data available for this work are prices, as Section 4.4.2 states, so the predictors are seven values of the index itself at seven pre-treatment dates. Those dates are December of 2010, 2012, 2014, 2016, 2018 and 2020, and November 2021. They are distributed across the whole pre-treatment period. The donors are selected to reproduce the treated index at those seven points, and the fit reported above is the outcome over all 145 pre-treatment months.

Two of those seven dates fall inside the pandemic, which is the period Section 5.5 holds responsible for most of the pre-treatment discrepancy. December 2020 and November 2021 are explicit matching targets, and the outcome can be read on the realised fit. Over 2020 and 2021 the error is 5.1 per cent of the level of those two years, against 7.0 per cent over the decade before them, and the average gap across the twenty-four months is 0.81 index points. Section 5.5 leaves two readings of that period open: a temporary contraction from which the treated group would have rebounded, and a relative decline that would have continued. The replica followed the contraction, and it followed the recovery. The movement is inside this counterfactual.

Seven donors compose the replica. Telefónica carries 55 per cent of the weight and Sodexo 22 per cent, Prysmian 10 per cent and Tele2 6.5 per cent, with Adidas, Ubisoft and SES dividing the remaining 6 per cent. The weights are concentrated on one donor, and the estimate is submitted on that account to the examination reported in Section 5.6.4.

Table 5.10 - Pre-treatment fit and placebo-based rank inference for the four European treated units.

Panel A: Counterfactual fit and divergence				
Treated unit	Pre-treatment RMSPE	Pre-treatment RMSPE (% of mean index)	Post-treatment RMSPE	Post/pre RMSPE ratio
Rheinmetall	26.54	15.86	2,643.19	99.59
European defence aggregate	10.24	6.61	277.65	27.12
Thales	27.71	12.18	386.06	13.93
Leonardo	21.12	23.54	274.91	13.02
Panel B: Placebo-based rank inference				
Treated unit	Full-sample rank	Full-sample p-value	Filtered rank	Filtered p-value
Rheinmetall	1/51	0.020	1/47	0.021
European defence aggregate	1/51	0.020	1/34	0.029
Thales	3/51	0.059	3/47	0.064
Leonardo	4/51	0.078	4/44	0.091

Notes: The outcome is the price index defined in Section 4.4.1, rebased to one hundred in January 2010. Pre-treatment fit is measured over January 2010 to January 2022.

After the treatment date the post-treatment error of the replica is 27.12 times the pre-treatment error. That ratio is the statistic on which the inference of Section 4.4.3 rests. The separation begins in February 2022. The date is established on the change in the difference with respect to January 2022, which is the quantity the treatment moves, and from that month onwards the change exceeds twice the pre-treatment error without interruption.

Figure 5.6 - Observed and synthetic price indices for the European defence aggregate

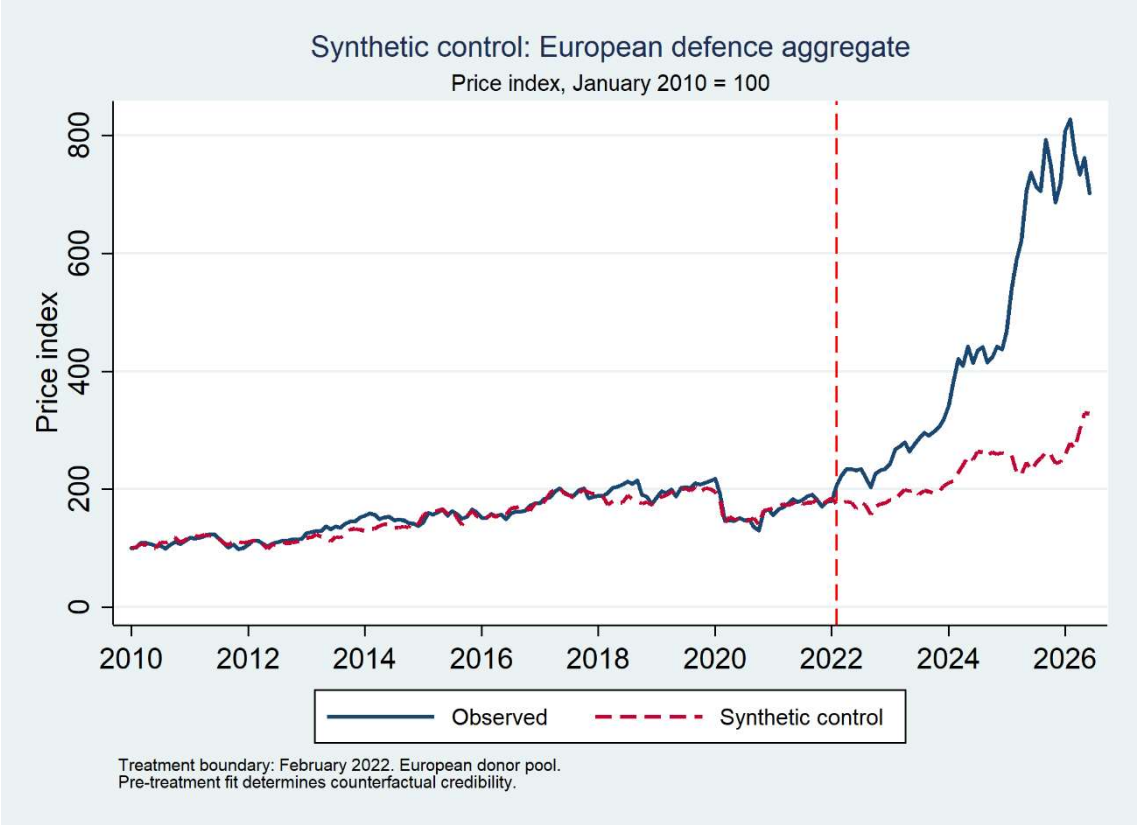


Figure 5.7 - Synthetic-control gap for the European defence aggregate

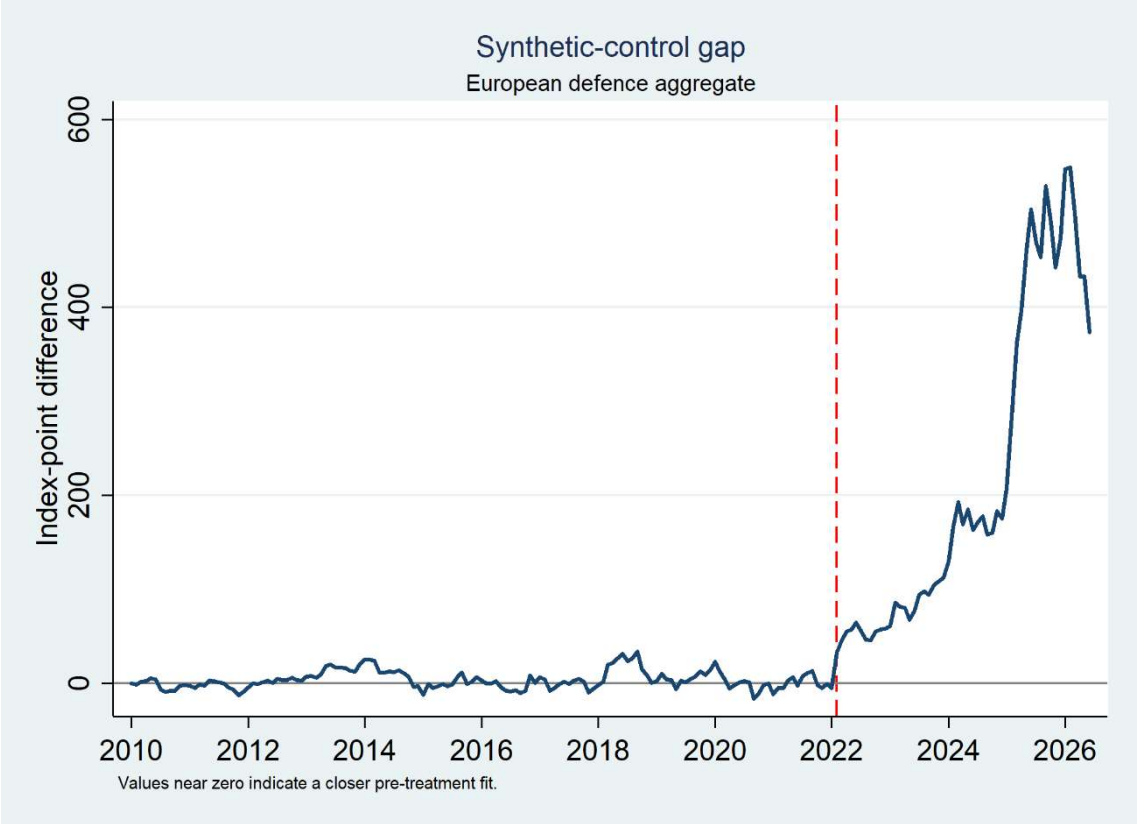
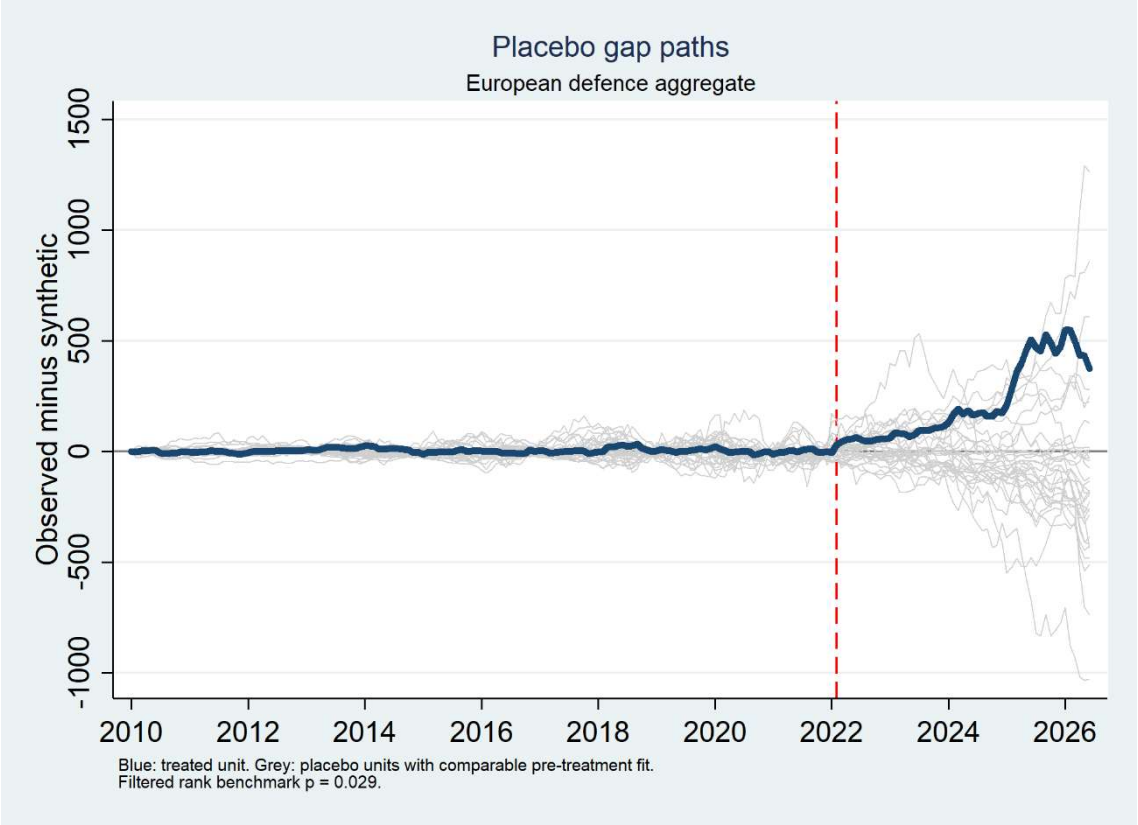


Figure 5.6 reports the two series, indistinguishable at the scale of the figure until the end of 2021 and separate from the beginning of 2022. Figure 5.7 reports the difference between them month by month. In each of the fifty-three post-treatment months that difference exceeds 20.5 index points, which is twice the pre-treatment error. The smallest, 32.4 points, is recorded in February 2022, and the largest, 549.0 points, in February 2026. Over the twelve preceding years the difference exceeded the same threshold in eleven months out of one hundred and forty-five, and it never exceeded 35 points. What distinguishes the two periods is the persistence of the excess and the level it reaches.

Each of the fifty European control firms is treated in turn as though it had entered the treatment in February 2022, receives a synthetic replica and produces its own RMSPE ratio. For the individual treated firms, this procedure compares firm-level treated units with firm-level placebo units. For the aggregated European treated unit, the benchmark distribution remains composed of individual donor firms. Its rank therefore measures how unusual the aggregate divergence is relative to those firms and should not be interpreted as an exact randomisation test. The aggregated European unit ranks first among the fifty-one units, producing a placebo-based benchmark of 0.0196. After excluding placebo units whose pre-treatment RMSPE exceeds five times that of the treated unit, it remains first among the thirty-four units retained, with a benchmark value of 0.0294.

Figure 5.8 - Placebo gap paths for the European defence aggregate and placebo units with comparable pre-treatment fit



The size of the divergence is measured from January 2022, the last pre-treatment month, and the interval therefore includes the month of the invasion. Over the fifty-three months that follow, the treated unit grows by a factor of 3.90 and its replica by a factor of 1.78. The ratio between the two growth factors is 2.20, a divergence of 119.5 per cent. Both factors are computed from the January 2022 values of the two series, which differ by 2.7 per cent, and that mismatch enters both bases of the measure.

The synthetic-control evidence supports a causal interpretation for the aggregated European treated unit within the European donor pool employed by the design. The synthetic series represents the estimated trajectory that the aggregate would have followed had it continued to move with the weighted donor combination reproducing its pre-treatment path. The treated and synthetic series separate in February 2022, and the divergence persists throughout the post-treatment period. This interpretation remains conditional on the constructed counterfactual and concerns the total geopolitical shock, including the political and military response generated by the invasion. It does not identify the direct effect of the invasion separately from that response.

5.6.2 The Three Individual Producers

The aggregate analysis does not identify the distribution of the response across individual firms. Rheinmetall, Thales and Leonardo are therefore examined as illustrative firm-level applications of the method. They do not constitute a representative sample or a formal test of heterogeneity across the European defence industry. All three separate from their replicas within two months of the invasion, although the quality of their pre-treatment fit and the magnitude of their subsequent divergence differ substantially.

Rheinmetall is reproduced to 15.9 per cent of its level before the treatment, and over the two pandemic years the error falls to 6.6 per cent. Its ratio of post-treatment to pre-treatment error is 99.59, the highest of the four European units. The firm stands first among the fifty-one units of the placebo distribution, and first among the forty-seven that survive the exclusion of poorly reproduced donors. From January 2022 the divergence amounts to 857 per cent, the largest of the four. One qualification attaches to that figure. At January 2022 the firm already stood 17 per cent above its replica, and the fit error of that month enters the measure.

Thales is reproduced to 12.2 per cent of its level, the closest fit among the three individual firms. Its ratio is 13.93 and places it third among the fifty-one units, with a p-value of 0.0588. From January 2022 the divergence amounts to 116 per cent, the smallest of the four European units, computed from a month in which the two series differ by 5.0 per cent. The error of its replica rises over the pandemic, from 11.3 per cent of the level over the decade to 13.9 per cent over 2020 and 2021. The argument made for the average unit in Section 5.6.1 therefore does not extend to this firm.

Leonardo is reproduced with an error of 21.12 index points, corresponding to 23.5 per cent of its average pre-treatment level and representing the least precise fit among the four European treated units. Its post-treatment to pre-treatment RMSPE ratio is 13.02 and ranks fourth among the fifty-one units, with a p-value of 0.0784. This provides only weak placebo-based evidence at the ten per cent level. In January 2022 the synthetic replica stood 34.7 per cent above the observed series, so part of the subsequent divergence closes an existing pre-treatment mismatch. The estimated divergence of 290 per cent is therefore the least reliable of the four European units and cannot be used to draw conclusions about the pre-treatment trajectory of the European sector as a whole.

The four European units agree on the direction of the divergence and on its approximate timing. Their estimated magnitudes vary from 116 per cent for Thales to 857 per cent for Rheinmetall, while the January 2022 mismatches range from 2.7 to 34.7 per cent. These differences are descriptive and partly reflect variation in pre-treatment fit and in the measurement base. The three individual applications do not establish sector-wide heterogeneity, and Section 5.3 provides no evidence of a systematic gradient based on pre-invasion arms-sales exposure.

5.6.3 The United States Aggregate

The United States aggregate comprises the seventeen defence firms with complete price histories and is matched to the fifty United States control firms. Table 5.11 reports the counterfactual fit and placebo-based rank inference for this separate regional synthetic-control exercise, constructed as in Section 5.6.1.

Table 5.11 - Synthetic-control results for the United States defence aggregate

Panel A: Counterfactual fit and divergence				
Treated unit	Pre-treatment RMSPE	Pre-treatment RMSPE (% of mean index)	Post-treatment RMSPE	Post/pre RMSPE ratio
United States defence aggregate	20.08	7.43	339.14	16.89
Panel B: Placebo-based rank inference				
Treated unit	Full-sample rank	Full-sample p-value	Filtered rank	Filtered p-value
United States defence aggregate	16/51	0.314	12/41	0.293

Notes: The outcome is the price index defined in Section 4.4.1, rebased to one hundred in January 2010. Pre-treatment fit is measured over January 2010 to January 2022.

The pre-treatment fit is acceptable. In both the full and filtered placebo distributions, the post-treatment divergence falls within the variation observed among untreated firms. The procedure therefore provides insufficient evidence to determine the sign or magnitude of a causal effect for the United States aggregate.

5.6.4 Robustness, Limitations and Verdict

The weights of the four European replicas are concentrated on one donor. Telefónica carries 55 per cent of the average European producer and 69 per cent of Rheinmetall, Ipsen

carries 55 per cent of Thales, and Nokia carries 80 per cent of Leonardo. Each estimation is therefore repeated with each substantial donor removed in turn. Excluding Telefónica brings the ratio of the average European producer from 27.12 to 21.43 and the ratio of Rheinmetall from 99.59 to 84.96. Re-running the inference on the forty-nine remaining donors leaves both units first in their distributions. Across every exclusion the ratio moves by at most 14.7 per cent for Rheinmetall, 15.7 per cent for Leonardo, 22.6 per cent for the average European producer and 26.1 per cent for Thales. No unit owes its divergence to a single donor. The post-treatment to pre-treatment error ratios obtained after excluding each substantial donor are reported in Table B.4.

The restriction of the donor pool to European firms is the second element varied, and Section 4.4.2 announces the exercise while motivating the restriction in advance. Opening the pool to the fifty United States control firms moves the ratios by at most 10 per cent, and the ranks are what changes. The average European producer falls from first among fifty-one to eighth among a hundred and one, with a p-value of 0.0792, Thales falls to twenty-eighth and Leonardo to thirty-sixth. Rheinmetall remains first, and its p-value of 0.0099 is one in a hundred and one, the floor of the enlarged test. Seven control firms now depart from their own replicas by more than the average European producer, five of them United States securities. On the index construction of Section 4.4.1 the United States control firms grew by a factor of 1.67 from January 2022, against 0.90 for the European ones. The two pools are therefore not interchangeable, and the significance of the average European producer at the five per cent level holds under the European pool alone.

Table 5.12 - Post/pre RMSPE ratios and placebo-based inference under the European and open donor pools

	European donor pool			Open donor pool		
Treated unit	Ratio	Rank	p-value	Ratio	Rank	p-value
Rheinmetall	99.59	1/51	0.020	96.24	1/101	0.010
European defence aggregate	27.12	1/51	0.020	29.83	8/101	0.079
Thales	13.93	3/51	0.059	14.30	28/101	0.277
Leonardo	13.02	4/51	0.078	12.21	36/101	0.356

Four limitations bound the conclusions of this section. The donor weights are selected exclusively from the pre-treatment trajectory of the price index and do not incorporate firms' industrial characteristics. Each regional donor pool imposes a minimum attainable rank-based value of one in fifty-one. For both regional aggregates, the placebo

distributions are composed of individual firms. Their ranks provide conditional inferential benchmarks based on comparisons between an aggregate treated unit and individual placebo firms. Finally, the regional aggregates have separate donor pools and placebo distributions and therefore do not identify the transatlantic difference as a single causal parameter.

The two strategies agree on the direction and statistical relevance of the European post-treatment divergence within their respective inferential procedures. Section 5.2 estimates an average monthly return differential of 0.0323 for European firms relative to the European control group, conditional on area-by-month and firm fixed effects. The synthetic-control divergence corresponds to approximately 0.0148 per month over the same fifty-three months and is measured against a weighted replica selected for its pre-treatment fit. The magnitudes reflect the outcome and counterfactual specific to each design. Their agreement therefore concerns direction, timing and statistical relevance within their respective identifying scopes.

Taken together, the findings support a causal interpretation of the favourable relative repricing observed for the aggregated European defence producer following the geopolitical shock generated by Russia's invasion of Ukraine. The synthetic counterfactual reproduces the pre-treatment trajectory of the aggregate, including the pandemic period, and separates from the treated series in February 2022. The aggregate ranks first relative to the European donor firms and remains robust to the removal of each principal donor. The difference between an aggregated treated unit and individual placebo units, together with the sensitivity of the rank to the geographical composition of the donor pool, limits the strength of this conclusion. The evidence therefore applies to the aggregated European treated unit within the European counterfactual employed by the design. It does not establish a causal effect for every European producer, a causal transatlantic difference or the direct effect of the invasion separately from the political and military response it generated.

6 Conclusions

This chapter brings together the evidence developed in the preceding analysis and relates it to the research question and hypotheses formulated in Chapter 3. The central issue is whether and to what extent the geopolitical shock generated by Russia's invasion of Ukraine affected the stock-market performance of listed defence producers relative to firms operating in sectors without direct exposure to military demand.

The conclusions draw on two complementary empirical strategies. The Difference-in-Differences design estimates the post-invasion differential across the firm-level panel and examines its geographical, firm-level and temporal distribution. The synthetic-control analysis constructs separate counterfactuals from the pre-invasion trajectories of regional defence aggregates and selected European firms. The two strategies use different outcomes, units of analysis and counterfactuals. Their findings are therefore interpreted within their respective identifying scopes before being combined into an overall answer to the research question.

6.1 Research Question and Analytical Scope

6.1.1 Research Question and Analytical Dimensions

This thesis addresses the following research question:

"To what extent did the geopolitical shock generated by Russia's invasion of Ukraine causally affect the stock-market performance of listed defence firms in Europe and the United States relative to firms operating in sectors without a direct conflict-specific military-demand channel?"

The primary firm-level estimand is the change, from February 2022 onwards, in the average monthly logarithmic-return differential between defence and control firms relative to the corresponding pre-invasion difference. A positive estimate indicates that defence producers achieved higher average returns than the comparison group during the post-invasion period.

The study targets the total effect of the geopolitical shock on the relative performance of defence firms. This effect includes the political and military response through

which the shock was transmitted to financial markets. The empirical design does not identify the direct effect of the invasion separately from this response. Defence producers occupy a distinct position within this transmission mechanism because military expenditure, procurement programmes and demand for productive capacity are central components of the policy response. The control firms provide a comparison for contemporaneous market conditions affecting companies in the same geographical areas, including the broader macroeconomic and financial consequences of the conflict.

The research question is examined through four analytical dimensions. H1 concerns the existence of an average post-invasion differential. H2 examines its distribution between European and United States producers. H3 tests its relationship with firms' pre-invasion dependence on arms sales. H4 examines its temporal persistence. These dimensions structure the presentation of the findings in Section 6.2, while their causal interpretation remains bounded by the counterfactual evidence available within each empirical design.

6.1.2 Empirical and Geographical Scope

The analysis covers listed defence producers and control firms headquartered in the United States and within the European perimeter of the study. This perimeter comprises the member states of the European Union, the United Kingdom and Norway. It reflects the institutional context in which sanctions, military cooperation, defence financing and procurement commitments developed after February 2022.

The firm-level panel contains 36 defence producers identified through the SIPRI Arms Industry Database and 100 control firms drawn from the S&P 500 and the STOXX Europe 600. The defence group comprises 21 United States firms and 15 European firms, while the control group contains 50 firms from each geographical area. The observation window extends from January 2010 to June 2026 and provides a long pre-invasion period for evaluating the trajectories preceding the geopolitical shock.

The Difference-in-Differences analysis measures stock-market performance through monthly logarithmic returns calculated from adjusted closing prices. The synthetic-control analysis uses a cumulative price index constructed from unwinsorised monthly returns and rebased to one hundred in January 2010. The two outcomes capture

complementary dimensions of the relative market response to the geopolitical shock and are interpreted within the identifying scope of their respective empirical designs.

6.2 Findings and Answer to the Research Question

The empirical analysis identifies a positive post-invasion return differential whose magnitude varies across geographical markets, firms and periods. Since the two empirical strategies rely on different outcomes and counterfactuals, their findings are first assessed in relation to the four hypotheses and then combined into an answer to the research question.

6.2.1 H1: The Average Post-Invasion Differential

The first hypothesis predicts that defence firms experienced a positive average post-invasion return differential relative to firms without a direct conflict-specific exposure channel. The baseline Difference-in-Differences specification estimates a differential of 0.0166 monthly log points, with a standard error of 0.0031. This corresponds to approximately 1.67 percentage points in simple-return terms. The estimate is positive and statistically significant at the one per cent level after accounting for permanent firm characteristics and movements common to the sample within each month.

The result remains positive and statistically significant across the five principal robustness checks. Its sign is preserved when the outcome is left unwinsorised, area-by-month fixed effects are introduced, lagged returns are included, the observation window is shortened, and February 2022 is excluded. The lower estimates obtained under the two temporal restrictions are consistent with the concentration of part of the response around the invasion.

The pre-treatment dynamic coefficients are inconsistent with the parallel-trends assumption in the full sample. The Difference-in-Differences estimate therefore documents a positive post-invasion differential but does not independently establish its causal attribution to the geopolitical shock. The synthetic-control analysis provides an additional counterfactual for the average of the fourteen European defence producers with complete price histories. The synthetic replica reproduces the twelve-year pre-invasion trajectory

with a pre-treatment error equal to 6.6 per cent of the average index level, including the contraction and recovery recorded during the pandemic.

The treated series separates from its synthetic replica in February 2022. Its ratio of post-treatment to pre-treatment prediction error is 27.12 and ranks first among the treated unit and the fifty European placebo units, producing a rank-based p-value of 0.0196. The result remains robust to the restricted placebo distribution and to the removal of the principal donors. Over the fifty-three post-treatment months, the treated series records a cumulative divergence of 119.5 per cent relative to the synthetic counterfactual.

H1 is therefore supported as a proposition concerning the existence of a positive average post-invasion differential. The firm-level panel establishes this differential in the transatlantic sample. The synthetic-control analysis supports a causal interpretation for the aggregated European defence producer within the counterfactual and donor pool employed by the design.

6.2.2 H2: Transatlantic Heterogeneity

The second hypothesis predicts a larger post-invasion return differential in Europe than in the United States. The preferred specification estimates each defence-control differential within its respective geographical market through area-by-month fixed effects. The European estimate is 0.0323 per month, with a standard error of 0.0041. The United States estimate is 0.0043, with a standard error of 0.0031, and is statistically indistinguishable from zero. The difference between the two estimates is 0.0281 and is statistically significant at the one per cent level.

The alternative specification with a common month fixed effect preserves the positive European differential but also produces a positive and statistically significant United States estimate. This change confirms the importance of controlling for the broader divergence between European and United States markets after 2022. The area-by-month fixed effects provide the within-market comparisons required to assess the geographical distribution of the defence-sector response.

The concentration of the estimated return differential in Europe is consistent with a larger revision of security and military-demand expectations in that market. Europe entered the geopolitical shock from a position of comparatively lower defence expenditure, while the United States had a larger and more continuously funded defence market.

Differences in pre-invasion valuations may also have contributed to the magnitude of the European adjustment. The price data do not permit the relative importance of these mechanisms to be determined.

H2 is therefore supported as a proposition concerning geographical heterogeneity in the estimated post-invasion differential. The Difference-in-Differences specification documents a marked transatlantic difference. The synthetic-control evidence supports a causal interpretation of the European aggregate response, while no statistically distinguishable causal effect is detected for the United States aggregate. Since the two regional analyses use separate donor pools and placebo distributions, they do not identify the transatlantic difference as a single causal parameter.

6.2.3 H3: Pre-Invasion Defence Revenue Exposure

The third hypothesis predicts that the post-invasion return differential increases with firms' pre-existing dependence on military activity. This relationship is estimated within the treatment group using the median share of arms sales in total revenues from 2010 to 2021.

The interaction between the post-treatment period and the arms-sales share is statistically indistinguishable from zero. Its confidence interval includes positive and negative relationships and remains compatible with economically meaningful differences between firms at the extremes of the observed distribution. The analysis therefore provides no evidence of a systematic gradient measurable through the selected variable.

H3 is not supported. The pre-invasion arms-sales share does not provide a sufficient basis for explaining the firm-level distribution of the post-invasion return differential. Relevant heterogeneity may depend on firms' product portfolios, procurement awards, productive capacity, expected contract margins and geographical exposure. These characteristics are not captured by the selected revenue measure.

6.2.4 H4: Persistence of the Return Differential

The fourth hypothesis predicts that the positive post-invasion return differential extended beyond the year of the invasion. The annual dynamic specification records positive and statistically significant coefficients in every complete post-treatment year from 2022 to

2025. The profile shows a large differential in 2022, a partial reduction in 2023 and smaller positive differentials in the subsequent years.

The monthly specification identifies the timing of the initial adjustment more precisely. The largest coefficient occurs in February 2022 and is followed by positive and statistically significant estimates through June. The estimates subsequently return to values comparable with the reference month. The monthly profile also identifies an anticipatory movement in January 2022, when information concerning the concentration of Russian forces was already available to investors. Part of the market adjustment may therefore have preceded the formal treatment date.

The event-phase specification identifies a differential of approximately two percentage points per month from February 2022 to June 2023. A similar differential remains observable from July 2023 to May 2025. From June 2025 to the end of the observation window, the estimate falls to −0.0004 per month, with a p-value of 0.944. No additional return differential is statistically detectable during this final phase, and the analysis identifies no statistically significant negative reversal.

The temporal pattern is consistent with successive market adjustments as governments specified military budgets, procurement programmes and longer-term commitments. Each period also contains additional policy and geopolitical developments. The coefficients therefore describe the evolution of the post-invasion differential without identifying the separate causal contribution of individual announcements or events.

H4 is supported as a proposition concerning the persistence of the estimated return differential. The differential remained observable beyond 2022 and during the principal subsequent post-treatment periods. From June 2025 onwards, it is no longer statistically detectable within the final event phase.

6.2.5 Research Question Answer

The findings provide a geographically delimited answer to the research question. The Difference-in-Differences analysis documents a positive average post-invasion return differential in the transatlantic sample. This differential is concentrated among European defence producers, while the preferred specification produces a small and statistically undetectable estimate for United States producers. The synthetic-control analysis supports a causal interpretation of the European divergence by constructing a counterfactual

that reproduces the pre-invasion trajectory of the aggregated European defence producer and separates from it in February 2022. A comparable causal effect is not established for the United States aggregate.

The return differential became pronounced around the invasion, with evidence of anticipatory adjustment in January 2022. It remained observable during the subsequent post-treatment periods and became statistically undetectable from June 2025 onwards. Its firm-level distribution is not systematically explained by the pre-invasion share of revenues derived from arms sales. The evidence therefore supports H1, H2 and H4 as propositions concerning the estimated post-invasion differential, while H3 is not supported.

Within the geographical, sample and counterfactual perimeter of the analysis, the synthetic-control evidence supports a causal interpretation of the favourable relative repricing observed for the aggregated European defence producer. This result concerns the total effect of the geopolitical shock generated by Russia's invasion of Ukraine, including the political and military response through which it was transmitted to financial markets.

6.3 Limitations of the Analysis

The findings of this thesis are subject to limitations concerning sample representativeness, causal identification, measurement and the attribution of post-treatment dynamics.

6.3.1 Sample and External Validity

The treatment group comprises major listed defence producers identified through the SIPRI Arms Industry Database and headquartered in the United States or within the European perimeter of the study. Privately held producers, wholly state-owned producers without listed equity, smaller firms outside the SIPRI selection, companies without usable stock-price histories and producers headquartered in other geographical areas are not observed. The findings therefore cannot be generalised automatically to the global defence industry or to firms operating under different ownership and procurement arrangements.

The requirement that price histories extend to June 2026 excludes firms that were acquired, delisted or entered insolvency before the end of the observation window. Applying the same requirement to treatment and control firms provides a consistent sample definition but does not eliminate survivorship selection if continued listing is related to

treatment status or post-invasion performance. The use of the 2026 constituent lists of the S&P 500 and the STOXX Europe 600 also excludes firms that had previously left these indices.

The comparison group is conditional on its construction procedure. Fifty control firms were selected from each market through a reproducible random ordering and then subjected to qualitative firm-level screening. The screening excludes firms with identified direct exposure to the conflict but cannot guarantee that every indirect or firm-specific exposure was detected. The analysis also uses one realised control sample without repeating the estimates across alternative random selections. The findings are therefore conditional on the sectoral composition and individual firms included in the selected comparison group.

The estimates do not represent capitalisation-weighted defence portfolios. The firm-level regressions assign observations to individual firms without weighting their contribution by market value. The regional treated units used in the synthetic-control analysis are averages of producers with complete price histories. The estimates consequently represent an average listed producer within the relevant sample and do not measure the aggregate market value or investable performance of the European or United States defence industry.

6.3.2 Limits to Causal Identification

The principal identification limitation concerns the Difference-in-Differences analysis. The pre-treatment coefficients are inconsistent with the parallel-trends assumption in the full sample and in both geographical markets, although the departure is weaker among European firms. When 2011, 2014 and 2016 are excluded from the annual pre-treatment diagnostic, the p-value of the European joint test rises from 0.0440 to 0.0707, weakening the evidence against parallel trends for Europe, while the hypothesis of constant pre-treatment differentials continues to be rejected in the full and United States samples. The coefficients estimated for H1 to H4 document post-invasion differentials between defence and control firms but do not independently establish that the geopolitical shock caused them. Firm and time fixed effects remove permanent firm characteristics and common market movements. They cannot eliminate unobserved factors that affected the two groups differently over the same period.

The synthetic-control analysis strengthens causal identification for the regional aggregate trajectories examined in relation to H1. For the European treated unit, the pre-treatment fit, post-invasion separation, placebo distribution and leave-one-out exercises support a causal interpretation of the observed divergence within the European donor pool. This interpretation remains conditional on the selected counterfactual. Expanding the donor pool to include United States controls leaves the estimated magnitude broadly comparable but weakens its rank-based statistical evidence. The strength of the European result therefore depends partly on the institutional and geographical comparability of the donors.

The synthetic-control analysis for the aggregated United States producer achieves an acceptable pre-treatment fit, but its post-treatment divergence is not statistically unusual within the placebo distribution. The analysis therefore does not detect a statistically distinguishable causal effect for the United States aggregate. This result does not establish that the true effect is zero or that the geopolitical shock produced a negative response. It establishes that the observed divergence cannot be credibly distinguished from the variation recorded among untreated placebo units under the inference procedure employed.

The two regional synthetic controls do not provide a formal causal test of the difference between the European and United States effects. Each aggregate has its own counterfactual, donor pool and placebo distribution. The transatlantic contrast estimated for H2 remains a result of the firm-level Difference-in-Differences specification and inherits its parallel-trends limitation. The synthetic evidence supports a causal interpretation of the European aggregate response but does not identify the geographical difference as a single causal parameter. It also does not convert the firm-level interaction estimated for H3 or the temporal coefficients estimated for H4 into separate causal effects.

6.3.3 Measurement of Outcomes and Economic Exposure

The dependent variables measure stock-market valuations. Stock prices incorporate revisions in expected future cash flows and in the returns required by investors. They can adjust before additional military expenditure becomes observable in contracts, production, revenues or profits. Price data do not permit the separate estimation of expected cash-flow revisions and changes in discount rates. The estimated return differential therefore identifies a relative market repricing associated with the information generated by

the geopolitical shock. Realised procurement benefits, production outcomes and corporate profitability remain outside the measurement scope of the analysis.

Geographical exposure is measured through the location of firms' headquarters. This classification does not capture the geographical distribution of customers, production facilities, procurement programmes or contracts. European firms may benefit from United States or non-European procurement, while United States producers may receive contracts generated by European military expenditure. Returns are also measured in the respective listing currencies. The area-by-month fixed effects absorb movements common to firms within each regional market, but firm-specific exposure to exchange-rate fluctuations, cross-border activities and foreign procurement remains only partially observed.

The pre-invasion share of revenues derived from arms sales is an incomplete measure of economic exposure. The variable is observed for 36 treated firms, and its support is truncated by the ten per cent threshold applied during sample construction. It measures existing dependence on military revenues but does not capture the allocation of subsequent procurement programmes, eligibility for new contracts, product portfolios, productive capacity, order backlogs, expected margins or the geographical origin of future demand. The estimated coefficient provides no evidence that the selected revenue measure systematically explains relative repricing within the observed sample. This result does not establish that firm-level exposure or productive characteristics are economically irrelevant.

6.3.4 Concurrent Events and Post-Treatment Attribution

The treatment date assigns February 2022 to the geopolitical shock generated by Russia's invasion of Ukraine. Relevant information was already available before the invasion, as the concentration of Russian forces and the public discussion of a possible attack allowed investors to revise their expectations during the preceding months. The February closing price incorporates market information following 24 February, but the monthly frequency does not identify the exact timing of the initial adjustment. The anticipatory movement observed in January 2022 indicates that part of the repricing may have preceded the formal treatment date.

The political and military response to the invasion forms part of the treatment mechanism examined in this thesis. Military-spending announcements, procurement decisions, military assistance and NATO commitments transmitted the geopolitical shock to expectations concerning defence demand. Under the estimand adopted in the analysis, the post-treatment differential represents the combined market response to the invasion and to the policy measures generated by it. The empirical design does not identify the direct effect of the invasion separately from this endogenous response.

The post-treatment period also contains potentially confounding events that did not arise as part of this mechanism. Other armed conflicts and independent geopolitical escalations introduced additional information capable of affecting defence-sector valuations. Fixed effects absorb the components of these events that produce common movements within the relevant market comparison. They cannot remove events or exposures that affect defence producers differently from the control firms. The annual and event-phase coefficients may therefore combine the treatment mechanism with the influence of concurrent events. They do not identify the separate causal contribution of an individual announcement, conflict or geopolitical development.

The event-based specification uses the NATO summits in Vilnius and The Hague to delimit economically relevant phases of the post-treatment period. The use of these dates as phase boundaries does not establish that the summits caused subsequent changes in the estimated differential. Each phase contains additional policy, military and geopolitical developments. The final phase is also shorter than the preceding phases and consequently provides a less precise estimate.

No additional relative return differential is statistically detectable during the final phase. This result does not identify the exact date on which the valuation adjustment ended or establish whether the relative price-level divergence accumulated earlier was subsequently reversed. The dynamic specifications therefore support conclusions about the evolution and persistence of the estimated post-treatment differential within the temporal resolution of the analysis.

6.4 Financing and Industrial Implications

The findings identify a possible connection between the relative repricing of European defence producers and the financing of their industrial growth. The synthetic-control analysis supports a causal interpretation of the favourable repricing observed for the aggregated European producer within the counterfactual employed by the study. This result does not apply uniformly to every company in the sample, but it indicates that the geopolitical shock and the political and military response associated with it materially changed the relative market valuation of the European defence aggregate. For listed producers, such a change may affect the terms on which new equity can be raised and therefore become relevant to decisions concerning the financing of future investment.

A higher market valuation does not itself provide the firm with additional cash. It becomes financially relevant if the company issues new shares. For a given amount of gross proceeds, a higher issue price requires the placement of fewer shares and, other conditions being equal, causes less dilution of existing shareholders' percentage ownership. The firm may therefore consider an equity issue alongside retained earnings and debt when choosing how to finance a growth plan. Equity issuance is more frequent when firms have relatively high market valuations [42]. The repricing documented in this thesis may consequently have increased the relevance of equity financing for the European listed producers represented in the analysis. This is a potential financing channel, since the empirical analysis does not examine subsequent share issues or changes in firms' capital structures.

Capital raised through an equity issue could support different forms of industrial growth. One option is internal investment in productive assets, technological development and specialised skills. This use is particularly relevant in the defence industry, where major programmes require substantial initial expenditure, specialised knowledge and long development periods [7]. Equity financing may allow a firm to undertake such investments without relying entirely on internally generated funds or increasing its debt by the full amount required. A second option is the acquisition of businesses or assets that provide complementary technologies and expertise. The choice between developing a capability internally and acquiring it from another firm depends on the resources already available to the company, the time and cost required for internal development and the

characteristics of the capability being sought. The findings identify a potential channel through which a favourable market valuation may enter subsequent financing decisions.

The temporal findings qualify this implication. The estimated return differential remained positive beyond 2022 and was detectable during the principal post-invasion periods through May 2025. From June 2025 to the end of the observation window, no additional differential was statistically detectable in the complete firm-level sample. This final result concerns the rate at which the relative valuation changed during that phase. It does not show that the price divergence accumulated earlier was reversed.

Taken together, the findings suggest that the post-invasion repricing may have consequences extending beyond the returns earned by existing shareholders. For the listed European producers covered by the study, it may also affect the feasibility and ownership implications of raising equity to finance industrial growth. The capital obtained could be allocated to internal development or to the acquisition of complementary capabilities, while the value of either strategy would depend on its expected costs and benefits. The analysis does not establish that firms used this channel, but it identifies a concrete route through which the financial response to a geopolitical shock may influence later industrial decisions.

6.5 Future Research

The findings and limitations of this thesis identify three main directions for future research. These concern the causal decomposition of the geopolitical shock and the isolation of concurrent events, the external validity of the findings across markets and conflicts, and the relationship between financial valuations, industrial outcomes and firms' economic exposure to military demand.

The first direction concerns the causal decomposition of the geopolitical shock generated by Russia's invasion of Ukraine. This thesis targets its total effect, including the political and military response through which it was transmitted to financial markets. Future studies could distinguish the direct effect of the invasion from the effects of specific policy announcements, military budgets, procurement programmes and military assistance. They could also isolate the influence of other conflicts and independent geopolitical escalations. Higher-frequency data and narrower event windows could improve the

analysis of anticipation and market timing, subject to explicit assumptions concerning simultaneous and previously anticipated events.

The second direction concerns the generalisability of the findings. A comparable empirical framework could be applied to other wars and geopolitical shocks. Suitable cases would require an identifiable arrival of information, an adequate pre-treatment period and a credible comparison group. Comparative studies could examine whether the sectoral response, geographical asymmetry and persistence documented in this thesis also emerge in other institutional and security contexts. Further research could also estimate the difference between European and United States producers as a causal parameter using institutionally and geographically comparable counterfactuals.

The third direction concerns the relationship between stock-market repricing and realised industrial outcomes. Contract awards, order backlogs, production, revenues, realised margins and investment in productive capacity could be used to assess whether market expectations were subsequently reflected in firm-level activity. Future studies could also combine the pre-invasion arms-sales share with information on product portfolios, procurement awards, productive capacity and the geographical distribution of customers, facilities and contracts. Integrating financial and operating data would provide a more complete account of the transmission from the geopolitical shock to the defence industry and could help explain the heterogeneity observed among producers.

These extensions would strengthen the identification and interpretation of the relationship between geopolitical shocks and defence-industry performance. They would clarify how changes in expected military demand are incorporated into financial valuations and subsequently reflected in realised industrial activity.

6.6 Final Remarks

Taken together, the findings indicate that the geopolitical shock generated by Russia's invasion of Ukraine was followed by a favourable relative repricing of listed defence producers, concentrated among European firms. The Difference-in-Differences estimates document this differential in the firm-level sample, while the synthetic-control analysis supports a causal interpretation of the divergence observed for the aggregated European defence producer within the counterfactual and donor pool employed. The geographical

and temporal patterns are consistent with a substantial revision of expected military demand following the shock and the associated political and military response. Within the limitations of the empirical design, this thesis provides evidence on the financial transmission of a major geopolitical shock to the market valuation of firms operating in a strategically important industry.

References

[1] Caldara, D. and Iacoviello, M. (2022) "*Measuring Geopolitical Risk*". Available at: https://doi.org/10.1257/aer.20191823

[2] Rigobon, R. and Sack, B. (2005) "*The Effects of War Risk on US Financial Markets*". Available at: https://doi.org/10.1016/j.jbankfin.2004.06.040

[3] Capelle-Blancard, G. and Couderc, N. (2008) "*What Drives the Market Value of Firms in the Defense Industry?*". Available at: https://doi.org/10.1016/j.rfe.2007.02.001

[4] Martins, A.M., Correia, P.A.P. and Gouveia, R. (2025) "*War! Good News for Defense Firms? Analysis of the Impact of Russia-Ukraine Conflict*". Available at: https://doi.org/10.1108/JES-11-2023-0667

[5] Knight, M., Loayza, N. and Villanueva, D. (1996) "*The Peace Dividend: Military Spending Cuts and Economic Growth*". Available at: https://doi.org/10.2307/3867351

[6] George, J. and Sandler, T. (2022) "*NATO Defense Demand, Free Riding, and the Russo-Ukrainian War in 2022*". Available at: https://doi.org/10.1007/s40812-022-00228-y

[7] Hartley, K. (2007) "*The Arms Industry, Procurement and Industrial Policies*", in Hartley, K. and Sandler, T. (eds) Handbook of Defense Economics. Vol. 2. Amsterdam: Elsevier. Available at: https://doi.org/10.1016/S1574-0013(06)02033-3

[8] Kluth, M. (2017) "*European Defence Industry Consolidation and Domestic Procurement Bias*". Available at: https://doi.org/10.1080/14751798.2017.1302576

[9] North Atlantic Treaty Organization (2014) "*Wales Summit Declaration*", 5 September. Available at: https://www.nato.int/cps/en/natohq/official_texts_112964.htm

[10] North Atlantic Treaty Organization (2022) *The Secretary General's Annual Report 2021*. Brussels: NATO. Available at: https://www.nato.int/nato_static_fl2014/assets/pdf/2022/3/pdf/sgar21-en.pdf

[11] European Commission and High Representative of the Union for Foreign Affairs and Security Policy (2022) *Joint Communication on the Defence Investment Gaps Analysis and Way Forward*. JOIN(2022) 24 final, 18 May. Brussels: European Commission. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52022JC0024

[12] Granát, M.P., Lehmann, K., Nagy, O. and Neszveda, G. (2023) "*Expect the Unexpected: Did the Equity Markets Anticipate the Russo-Ukrainian War?*". Available at: https://doi.org/10.1016/j.frl.2023.104301

[13] Federle, J., Meier, A., Müller, G.J. and Sehn, V. (2026) "*Proximity to War: The Stock Market Response to the Russian Invasion of Ukraine*". Available at: https://doi.org/10.1111/jmcb.13226

[14] Tian, N., Lopes da Silva, D., Liang, X., Scarazzato, L., Béraud-Sudreau, L. and Assis, A. (2023) *Trends in World Military Expenditure, 2022*. SIPRI Fact Sheet. Stockholm: Stockholm International Peace Research Institute. Available at: https://doi.org/10.55163/PNVP2622

[15] North Atlantic Treaty Organization (2023) "*Vilnius Summit Communiqué*", 11 July. Available at: https://www.nato.int/cps/en/natohq/official_texts_217320.htm

[16] North Atlantic Treaty Organization (2025) "*The Hague Summit Declaration*", 25 June. Available at: https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2025/06/25/the-hague-summit-declaration

[17] Stockholm International Peace Research Institute (2026) "*SIPRI Arms Industry Database*". Available at: https://www.sipri.org/databases/armsindustry

[18] Fama, E.F. (1970) "*Efficient Capital Markets: A Review of Theory and Empirical Work*". Available at: https://doi.org/10.2307/2325486

[19] Campbell, J.Y. (1991) "*A Variance Decomposition for Stock Returns*". Available at: https://doi.org/10.2307/2233809

[20] Gurdgiev, C., Henrichsen, A. and Mulhair, A. (2022) "*The Budgets of Wars: Analysis of the U.S. Defense Stocks in the Post-Cold War Era*". Available at: https://doi.org/10.1016/j.iref.2022.06.023

[21] Zhang, Z., Bouri, E., Klein, T. and Jalkh, N. (2022) "*Geopolitical Risk and the Returns and Volatility of Global Defense Companies: A New Race to Arms?*". Available at: https://doi.org/10.1016/j.irfa.2022.102327

[22] Rogerson, W.P. (1994) "*Economic Incentives and the Defense Procurement Process*". Available at: https://doi.org/10.1257/jep.8.4.65

[23] Fisher, J.D.M. and Peters, R. (2010) "*Using Stock Returns to Identify Government Spending Shocks*". Available at: https://doi.org/10.1111/j.1468-0297.2010.02355.x

[24] Belo, F., Gala, V.D. and Li, J. (2013) "*Government Spending, Political Cycles, and the Cross Section of Stock Returns*". Available at: https://doi.org/10.1016/j.jfineco.2012.08.016

[25] Amodeo, F. and Briganti, E. (2025) *High-Frequency Cross-Sectional Identification of Military News Shocks*. Staff Working Paper 2025-27. Ottawa: Bank of Canada. Available at: https://doi.org/10.34989/swp-2025-27

[26] Ramey, V.A. (2011) "*Identifying Government Spending Shocks: It's All in the Timing*". Available at: https://doi.org/10.1093/qje/qjq008

[27] Ben Zeev, N. and Pappa, E. (2017) "*Chronicle of a War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks*". Available at: https://doi.org/10.1111/ecoj.12349

[28] Nerlinger, M. and Utz, S. (2022) "*The Impact of the Russia-Ukraine Conflict on Energy Firms: A Capital Market Perspective*". Available at: https://doi.org/10.1016/j.frl.2022.103243

[29] European Union Agency for the Cooperation of Energy Regulators (2022) *ACER's Final Assessment of the EU Wholesale Electricity Market Design*. Ljubljana: ACER. Available at: https://www.acer.europa.eu/sites/default/files/documents/Publications/Final_Assessment_EU_Wholesale_Electricity_Market_Design.pdf

[30] Ferriani, F. and Gazzani, A. (2023) "*The Invasion of Ukraine and the Energy Crisis: Comparative Advantages in Equity Valuations*". Available at: https://doi.org/10.1016/j.frl.2023.104604

[31] Boubaker, S., Nguyen, N., Trinh, V.Q. and Vu, T. (2023) "*Market Reaction to the Russian Ukrainian War: A Global Analysis of the Banking Industry*". Available at: https://doi.org/10.1108/RAF-10-2022-0294

[32] Martins, A.M. and Cró, S. (2024) "*Mayday! Mayday! The Airlines Stock Returns Are Failing: Analysis of the Impact of Russia-Ukraine War*". Available at: https://doi.org/10.1108/JES-08-2023-0390

[33] Bertrand, M., Duflo, E. and Mullainathan, S. (2004) "*How Much Should We Trust Differences-in-Differences Estimates?*". Available at: https://doi.org/10.1162/003355304772839588

[34] Roth, J., Sant'Anna, P.H.C., Bilinski, A. and Poe, J. (2023) "*What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature*". Available at: https://doi.org/10.1016/j.jeconom.2023.03.008

[35] North Atlantic Treaty Organization (2022) "*NATO Allies Send More Ships, Jets to Enhance Deterrence and Defence in Eastern Europe*", NATO News, 24 January. Available at: https://www.nato.int/en/news-and-events/articles/news/2022/01/24/nato-allies-send-more-ships-jets-to-enhance-deterrence-and-defence-in-eastern-europe

[36] Goodman-Bacon, A. (2021) "*Difference-in-Differences with Variation in Treatment Timing*". Available at: https://doi.org/10.1016/j.jeconom.2021.03.014

[37] Nickell, S. (1981) "*Biases in Dynamic Models with Fixed Effects*". Available at: https://doi.org/10.2307/1911408

[38] Petersen, M.A. (2009) "*Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches*". Available at: https://doi.org/10.1093/rfs/hhn053

[39] Abadie, A., Diamond, A. and Hainmueller, J. (2010) "*Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program*". Available at: https://doi.org/10.1198/jasa.2009.ap08746

[40] North Atlantic Treaty Organization (2012) The Secretary General's Annual Report 2011. Brussels: NATO. Available at: https://www.nato.int/content/dam/nato/legacy-wcm/media_pdf/pdf_publications/20120125_Annual_Report_2011_en.pdf

[41] Wagner, A.F., Zeckhauser, R.J. and Ziegler, A. (2018) "*Company Stock Price Reactions to the 2016 Election Shock: Trump, Taxes, and Trade*". Available at: https://doi.org/10.1016/j.jfineco.2018.06.013

[42] Baker, M. and Wurgler, J. (2002) "*Market Timing and Capital Structure*". Available at: https://doi.org/10.1111/1540-6261.00414

APPENDIX A

Table A.1 Treatment Group: 36 Listed Defence Firms

#	Ticker	Company	Area	Defence share
1	AIR.PA	Airbus	Europe	0.180
2	BA.L	BAE Systems	Europe	0.952
3	BAB.L	Babcock International Group	Europe	0.535
4	CHG.L	Chemring Group	Europe	0.948
5	AM.PA	Dassault Aviation Group	Europe	0.366
6	FCT.MI	Fincantieri	Europe	0.326
7	KOG.OL	Kongsberg Gruppen	Europe	0.414
8	LDO.MI	Leonardo	Europe	0.643
9	QQ.L	QinetiQ	Europe	0.752
10	RHM.DE	Rheinmetall	Europe	0.509
11	RR.L	Rolls-Royce	Europe	0.261
12	SAAB-B.ST	Saab	Europe	0.834
13	SAF.PA	Safran	Europe	0.153
14	SRP.L	Serco Group	Europe	0.297
15	HO.PA	Thales	Europe	0.504
16	ACM	AECOM	United States	0.130
17	BWXT	BWX Technologies	United States	0.779
18	BA	Boeing	United States	0.364
19	BAH	Booz Allen Hamilton	United States	0.694
20	CACI	CACI International	United States	0.699
21	CW	Curtiss-Wright Corp.	United States	0.415
22	GD	General Dynamics Corp.	United States	0.672
23	HII	HII	United States	0.950
24	HON	Honeywell International	United States	0.133
25	KBR	KBR	United States	0.420
26	LHX	L3Harris Technologies	United States	0.814
27	LDOS	Leidos	United States	0.597
28	LMT	Lockheed Martin Corp.	United States	0.830
29	MOG-A	Moog	United States	0.346
30	NOC	Northrop Grumman Corp.	United States	0.832
31	OSK	Oshkosh Corp.	United States	0.298
32	PH	Parker-Hannifin Corp.	United States	0.111
33	SAIC	SAIC	United States	0.640
34	TDY	Teledyne Technologies	United States	0.339
35	TXT	Textron	United States	0.297
36	TDG	TransDigm Group	United States	0.349

Notes. The treatment group contains the 36 listed defence producers included in the firm-level panel. Defence share is the median ratio of arms sales to total revenues over the period from 2010 to 2021. Area identifies the geographical classification used in the empirical analysis.

Table A.2 Control Group: 100 Listed Firms

#	Ticker	Company	Country	GICS Sector
1	STX	Seagate Technology	United States	Information Technology
2	WSM	Williams-Sonoma	United States	Consumer Discretionary
3	BSX	Boston Scientific	United States	Health Care
4	NVR	NVR	United States	Consumer Discretionary
5	LVS	Las Vegas Sands	United States	Consumer Discretionary
6	HSY	Hershey	United States	Consumer Staples
7	ODFL	Old Dominion Freight	United States	Industrials
8	EL	Estee Lauder	United States	Consumer Staples
9	ARE	Alexandria Real Estate	United States	Real Estate

#	Ticker	Company	Country	GICS Sector
10	JBHT	J.B. Hunt	United States	Industrials
11	WELL	Welltower	United States	Real Estate
12	KIM	Kimco Realty	United States	Real Estate
13	AVGO	Broadcom	United States	Information Technology
14	HST	Host Hotels & Resorts	United States	Real Estate
15	TTWO	Take-Two Interactive	United States	Communication Services
16	CMG	Chipotle	United States	Consumer Discretionary
17	FRT	Federal Realty	United States	Real Estate
18	CDNS	Cadence Design	United States	Information Technology
19	PWR	Quanta Services	United States	Industrials
20	TYL	Tyler Technologies	United States	Information Technology
21	GLW	Corning	United States	Information Technology
22	EME	EMCOR Group	United States	Industrials
23	MRVL	Marvell Technology	United States	Information Technology
24	AMAT	Applied Materials	United States	Information Technology
25	LRCX	Lam Research	United States	Information Technology
26	SMCI	Super Micro Computer	United States	Information Technology
27	CBRE	CBRE Group	United States	Real Estate
28	UHS	Universal Health Services	United States	Health Care
29	ORCL	Oracle	United States	Information Technology
30	ACN	Accenture	United States	Information Technology
31	CHD	Church & Dwight	United States	Consumer Staples
32	COR	Cencora	United States	Health Care
33	MSFT	Microsoft	United States	Information Technology
34	HUM	Humana	United States	Health Care
35	DHI	D.R. Horton	United States	Consumer Discretionary
36	WST	West Pharmaceutical	United States	Health Care
37	CLX	Clorox	United States	Consumer Staples
38	AVB	AvalonBay	United States	Real Estate
39	CHRW	C.H. Robinson	United States	Industrials
40	PM	Philip Morris Intl	United States	Consumer Staples
41	SJM	J.M. Smucker	United States	Consumer Staples
42	CVS	CVS Health	United States	Health Care
43	AMD	Advanced Micro Devices	United States	Information Technology
44	LOW	Lowe's	United States	Consumer Discretionary
45	COO	Cooper Companies	United States	Health Care
46	ROP	Roper Technologies	United States	Information Technology
47	HD	Home Depot	United States	Consumer Discretionary
48	RSG	Republic Services	United States	Industrials
49	ECL	Ecolab	United States	Materials
50	DGX	Quest Diagnostics	United States	Health Care
51	DSY.PA	Dassault Systemes	France	Information Technology
52	DEMANT.CO	Demant	Denmark	Health Care
53	ABF.L	Associated British Foods	Great Britain	Consumer Staples
54	SW.PA	Sodexo	France	Consumer Discretionary
55	DLN.L	Derwent London	Great Britain	Real Estate
56	AMBU-B.CO	Ambu B	Denmark	Health Care
57	INCH.L	Inchcape	Great Britain	Consumer Discretionary
58	IPN.PA	Ipsen	France	Health Care
59	TELIA.ST	Telia Company	Sweden	Communication Services
60	BIM.PA	bioMerieux	France	Health Care
61	AFX.DE	Carl Zeiss Meditec	Germany	Health Care
62	HUSQ-B.ST	Husqvarna B	Sweden	Consumer Discretionary
63	INDT.ST	Indutrade	Sweden	Industrials
64	REY.MI	Reply	Italy	Information Technology
65	GRG.L	Greggs	Great Britain	Consumer Discretionary
66	NEM.DE	Nemetschek	Germany	Information Technology
67	SESG.PA	SES	Luxembourg	Information Technology
68	BNZL.L	Bunzl	Great Britain	Industrials
69	DPLM.L	Diploma	Great Britain	Industrials

#	Ticker	Company	Country	GICS Sector
70	ACS.MC	ACS	Spain	Industrials
71	DHL.DE	Deutsche Post DHL	Germany	Industrials
72	UBI.PA	Ubisoft Entertainment	France	Consumer Discretionary
73	QIA.DE	Qiagen	Germany	Health Care
74	BMW.DE	BMW	Germany	Consumer Discretionary
75	TEP.PA	Teleperformance	France	Industrials
76	COV.PA	Covivio	France	Real Estate
77	RAND.AS	Randstad	Netherlands	Industrials
78	MC.PA	LVMH Moet Hennessy	France	Consumer Discretionary
79	TEL2-B.ST	Tele2 B	Sweden	Communication Services
80	NOKIA.HE	Nokia	Finland	Information Technology
81	CPG.L	Compass Group	Great Britain	Consumer Discretionary
82	EN.PA	Bouygues	France	Industrials
83	BWY.L	Bellway	Great Britain	Consumer Discretionary
84	SAFE.L	Safestore Holdings	Great Britain	Real Estate
85	TEF.MC	Telefonica	Spain	Communication Services
86	REC.MI	Recordati	Italy	Health Care
87	TEL.OL	Telenor	Norway	Communication Services
88	REL.L	RELX	Great Britain	Industrials
89	GN.CO	GN Store Nord	Denmark	Consumer Discretionary
90	ADS.DE	Adidas	Germany	Consumer Discretionary
91	CAST.ST	Castellum	Sweden	Real Estate
92	PRY.MI	Prysmian	Italy	Industrials
93	FABG.ST	Fabege	Sweden	Real Estate
94	IFX.DE	Infineon Technologies	Germany	Information Technology
95	AALB.AS	Aalberts	Netherlands	Industrials
96	OR.PA	L'Oreal	France	Consumer Staples
97	STMPA.PA	STMicroelectronics	Italy	Information Technology
98	BT-A.L	BT Group	Great Britain	Communication Services
99	HIK.L	Hikma Pharmaceuticals	Great Britain	Health Care
100	BC8.DE	Bechtle	Germany	Information Technology

Notes. The control group contains 50 firms from the United States and 50 firms from the European perimeter. Firms were drawn from the eligible S&P 500 and STOXX Europe 600 pools through a reproducible random ordering and subsequently examined against the sample-selection criteria.

Table A.3 Selection Record: 80 Discarded Draws

Area	Draw	Company	Criterion	GICS Sector
United States	1	Biogen	shareholder distribution	Health Care
United States	8	Comcast	shareholder distribution	Communication Services
United States	11	Pfizer	shareholder distribution	Health Care
United States	12	Mosaic	commodity affected by the shock	Materials
United States	21	3M	shareholder distribution	Industrials
United States	23	Airbnb	listing after 2010	Consumer Discretionary
United States	29	Smurfit Westrock	newly created entity	Materials
United States	30	Allegion	listing after 2010	Industrials
United States	33	Alphabet Cl. C	dual share class	Communication Services
United States	34	Ford Motor	Russian operations divested	Consumer Discretionary
United States	35	Warner Bros Discovery	newly created entity	Communication Services
United States	36	GE Vernova	listing after 2010	Industrials
United States	37	Union Pacific	merger	Industrials
United States	39	Revvity	corporate discontinuity	Health Care
United States	41	CF Industries	commodity affected by the shock	Materials
United States	46	CRH	listing after 2010	Materials
United States	47	IQVIA	listing after 2010	Health Care
United States	48	T-Mobile US	newly created entity	Communication Services
United States	49	Kraft Heinz	newly created entity	Consumer Staples
United States	54	Medtronic	shareholder distribution	Health Care
United States	57	Merck & Co	shareholder distribution	Health Care

Area	Draw	Company	Criterion	GICS Sector
United States	59	Western Digital	shareholder distribution	Information Technology
United States	60	Aptiv	listing after 2010	Consumer Discretionary
United States	61	Dover	shareholder distribution	Industrials
United States	63	Bunge Global	commodity affected by the shock	Consumer Staples
United States	65	Norfolk Southern	merger	Industrials
United States	66	Leidos	overlap with treatment group	Industrials
United States	71	ServiceNow	listing after 2010	Information Technology
United States	72	DuPont de Nemours	corporate discontinuity	Materials
United States	74	Caterpillar	Russian operations divested	Industrials
United States	77	Arista Networks	listing after 2010	Information Technology
Europe	1	Nexans	Russian operations divested	Industrials
Europe	2	Boliden	commodity affected by the shock	Materials
Europe	3	Flutter Entertainment	listing after 2010	Consumer Discretionary
Europe	4	Frontline	commodity affected by the shock	Energy
Europe	11	Daimler Truck	listing after 2010	Industrials
Europe	14	Holmen	commodity affected by the shock	Materials
Europe	15	Volkswagen Pref.	Russian operations divested	Consumer Discretionary
Europe	17	Wienerberger	Russian operations divested	Materials
Europe	20	Unilever	shareholder distribution	Consumer Staples
Europe	22	Universal Music	listing after 2010	Communication Services
Europe	23	Danone	Russian operations divested	Consumer Staples
Europe	25	B&M European Value	listing after 2010	Consumer Discretionary
Europe	27	Spectris	acquired or delisted	Information Technology
Europe	29	Intl Distributions Serv.	listing after 2010	Industrials
Europe	30	Groupe SEB	not reached in time	Consumer Discretionary
Europe	32	SimCorp	acquired or delisted	Information Technology
Europe	33	Heineken	Russian operations divested	Consumer Staples
Europe	34	Melrose Industries	shareholder distribution	Industrials
Europe	35	Rio Tinto	commodity affected by the shock	Materials
Europe	37	Siemens	Russian operations divested	Industrials
Europe	42	DSM Firmenich	newly created entity	Materials
Europe	43	GSK	shareholder distribution	Health Care
Europe	45	OCI	commodity affected by the shock	Materials
Europe	47	Yara	commodity affected by the shock	Materials
Europe	48	Merlin Properties	listing after 2010	Real Estate
Europe	51	Huhtamaki	Russian operations divested	Materials
Europe	53	Addtech	shareholder distribution	Industrials
Europe	54	Francaise des Jeux	listing after 2010	Consumer Discretionary
Europe	56	Voestalpine	commodity affected by the shock	Materials
Europe	58	Just Eat Takeaway	acquired or delisted	Consumer Discretionary
Europe	59	Electrolux	Russian operations divested	Consumer Discretionary
Europe	61	Valeo	Russian operations divested	Consumer Discretionary
Europe	64	Ferrari	listing after 2010	Consumer Discretionary
Europe	67	TietoEVRY	corporate discontinuity	Information Technology
Europe	69	Ackermans & van Haaren	sector inconsistency	Financials
Europe	70	JDE Peet's	listing after 2010	Consumer Staples
Europe	71	Kering	shareholder distribution	Consumer Discretionary
Europe	72	ArcelorMittal	commodity affected by the shock	Materials
Europe	74	Svenska Cellulosa	shareholder distribution	Materials
Europe	75	Air Liquide	Russian operations divested	Materials
Europe	76	Ferrovial	corporate discontinuity	Industrials
Europe	81	Fresenius	corporate discontinuity	Health Care
Europe	83	Reckitt Benckiser	corporate discontinuity	Consumer Staples
Europe	85	Antofagasta	commodity affected by the shock	Materials
Europe	86	DS Smith	acquired or delisted	Materials
Europe	87	Fuchs Petrolub Pref.	dual share class	Materials
Europe	94	Sartorius Pref.	dual share class	Health Care
Europe	97	LPP	Russian operations divested	Consumer Discretionary
Europe	98	UPM-Kymmene	Russian operations divested	Materials

Notes. The table records the firms examined and excluded during the construction of the control group, together with the applicable exclusion criterion. A total of 180 firms were examined to obtain the final 100 control firms: 81 firms for the 50 United States positions and 99 firms for the 50 European positions.

Table A.4 Descriptive Statistics for the Estimation Sample

Sample	Variable	Obs.	Mean	SD	Min.	Max.
Firm-month	Monthly logarithmic return, winsorised	26,663	0.0107	0.0775	−0.2103	0.2181
Firm-month	Monthly logarithmic return, unwinsorised	26,663	0.0105	0.0833	−0.6632	0.7534
Firm-month	Defence-firm indicator	26,663	0.2611	0.4393	0.0000	1.0000
Firm-month	Post-treatment indicator	26,663	0.2703	0.4441	0.0000	1.0000
Firm-month	Defence × Post interaction	26,663	0.0716	0.2578	0.0000	1.0000
Firm-month	European-firm indicator	26,663	0.4782	0.4995	0.0000	1.0000
Firm level	Median pre-invasion arms-sales share	36	0.5106	0.2573	0.1110	0.9520

Notes. Firm-month statistics are calculated on the 26,663 observations entering the Difference-in-Differences estimations. The arms-sales share is reported once for each of the 36 treated firms. The winsorised return is bounded at the first and ninety-ninth percentiles of the estimation sample.

APPENDIX B

Table B.1. Complete Annual Return-Differential Estimates under Alternative Reference Periods

Year	Coef. (ref. 2021)	SE	Coef. (ref. 2019)	SE	Coef. (ref. 2018-2021)	SE
2010	0.0052	0.0045	−0.0096	0.0047	0.0009	0.0037
2011	0.0112	0.0050	−0.0036	0.0052	0.0069	0.0042
2012	0.0115	0.0043	−0.0033	0.0044	0.0072	0.0036
2013	0.0188	0.0048	0.0039	0.0049	0.0145	0.0039
2014	0.0010	0.0053	−0.0138	0.0056	−0.0033	0.0045
2015	0.0098	0.0055	−0.0050	0.0051	0.0055	0.0045
2016	0.0191	0.0044	0.0043	0.0040	0.0148	0.0034
2017	0.0129	0.0051	−0.0019	0.0054	0.0087	0.0044
2018	0.0068	0.0048	−0.0081	0.0048		
2019	0.0148	0.0044				
2020	−0.0045	0.0050	−0.0193	0.0046		
2021			−0.0148	0.0044		
2022	0.0451	0.0052	0.0303	0.0055	0.0408	0.0046
2023	0.0151	0.0050	0.0002	0.0054	0.0108	0.0044
2024	0.0262	0.0053	0.0114	0.0060	0.0219	0.0050
2025	0.0336	0.0062	0.0188	0.0068	0.0294	0.0063

Notes: Each coefficient reports the treated-control monthly logarithmic-return differential for the indicated year relative to the stated reference period. All specifications include firm and month fixed effects. Standard errors are clustered by firm. Blank cells identify the omitted reference year or reference period. The partial observations for 2026 are not reported.

Table B.2. Complete Monthly Return-Differential Estimates, Reference Month November 2021

Month	Return differential	Standard error
Feb 2021	0.0962	0.0205
Mar 2021	0.0742	0.0150
Apr 2021	0.0594	0.0154
May 2021	0.0812	0.0159
Jun 2021	0.0294	0.0127
Jul 2021	0.0362	0.0148
Aug 2021	0.0507	0.0162
Sep 2021	0.0994	0.0158
Oct 2021	0.0174	0.0178
Dec 2021	0.0645	0.0152
Jan 2022	0.1161	0.0186
Feb 2022	0.1954	0.0204
Mar 2022	0.0955	0.0174
Apr 2022	0.0940	0.0204
May 2022	0.0735	0.0145
Jun 2022	0.1328	0.0186
Jul 2022	0.0070	0.0186
Aug 2022	0.0835	0.0154
Sep 2022	0.0903	0.0164
Oct 2022	0.1446	0.0178
Nov 2022	0.0329	0.0167
Dec 2022	0.0934	0.0140
Jan 2023	0.0153	0.0203
Feb 2023	0.0930	0.0160

Month	Return differential	Standard error
Mar 2023	0.0471	0.0155
Apr 2023	0.0434	0.0172
May 2023	0.0380	0.0144
Jun 2023	0.1066	0.0172
Jul 2023	0.0665	0.0185
Aug 2023	0.1110	0.0160
Sep 2023	0.0729	0.0156
Oct 2023	0.1215	0.0168
Nov 2023	0.0313	0.0193
Dec 2023	0.0521	0.0130
Jan 2024	0.0830	0.0179
Feb 2024	0.1164	0.0187

Notes. The table reports the monthly treated-control return differential from February 2021 to February 2024. November 2021 is the omitted reference month and is therefore not reported. The specification includes firm and month fixed effects. Standard errors are clustered by firm.

Table B.3. Joint Tests of the Pre-Treatment Coefficients

Test	F-statistic	Numerator degrees of freedom	Residual degrees of freedom	p-value
Annual, base 2021	4.1023	11	135	< 0.001
Monthly, base Nov 2021	8.5596	11	135	< 0.001
Decade 2010-2018, base 2019	2.9534	9	135	0.0031
Years 2020-2021, base 2019	10.6570	2	135	< 0.001
Full pre-treatment set, base 2019	4.1023	11	135	< 0.001
Base mean 2018-2021	4.2283	8	135	< 0.001
European firms only, base 2021	1.9897	11	64	0.0440
United States firms only, base 2021	3.5135	11	70	< 0.001
European firms only, excluding 2011, 2014, 2016. Base 2021	1.9288	8	64	0.0707

Notes. The reported Wald tests evaluate the null hypothesis that the selected pre-treatment interaction coefficients are jointly equal to zero. Numerator and residual degrees of freedom are reported for each test. Values below 0.001 are reported as $p < 0.001$.

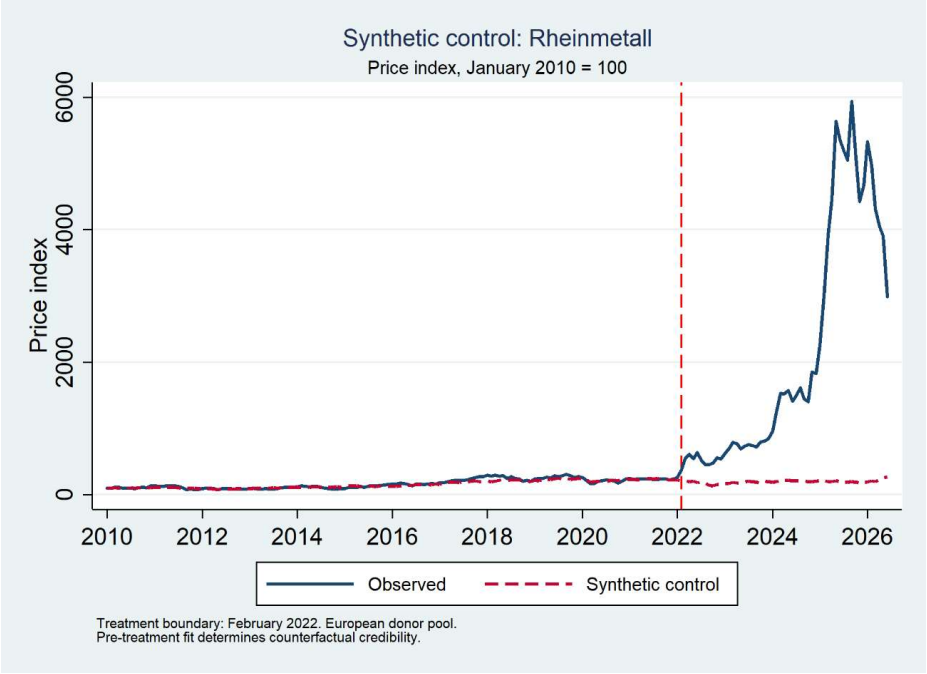
Table B.4. Leave-One-Out Donor Sensitivity of the Synthetic-Control Estimates

Region	Treated unit	Excluded donor	Weight	RMSPE pre, %	RMSPE ratio	Final gap
Europe	Defence aggregate	Baseline		6.61	27.12	373.6
Europe	Defence aggregate	TEF.MC	0.550	7.09	21.43	240.5
Europe	Defence aggregate	SW.PA	0.224	6.50	24.51	274.4
Europe	Defence aggregate	PRY.MI	0.099	6.37	33.05	515.4

Region	Treated unit	Excluded donor	Weight	RMSPE pre, %	RMSPE ratio	Final gap
Europe	Defence aggregate	TEL2-B.ST	0.065	6.34	33.25	513.6
Europe	Leonardo	Baseline		23.54	13.02	346.9
Europe	Leonardo	NOKIA.HE	0.805	27.04	11.19	405.1
Europe	Leonardo	SW.PA	0.133	25.68	10.97	276.5
Europe	Leonardo	TEF.MC	0.050	23.58	13.06	348.5
Europe	Rheinmetall	Baseline		15.86	99.59	2723.1
Europe	Rheinmetall	TEF.MC	0.691	16.99	84.96	2048.9
Europe	Rheinmetall	ADS.DE	0.130	17.91	88.18	2641.7
Europe	Rheinmetall	ACS.MC	0.076	15.94	99.96	2789.9
Europe	Rheinmetall	UBI.PA	0.063	17.43	87.84	2501.6
Europe	Thales	Baseline		12.18	13.93	496.5
Europe	Thales	IPN.PA	0.547	12.83	16.23	646.1
Europe	Thales	SW.PA	0.343	12.11	10.29	245.9
Europe	Thales	UBI.PA	0.089	12.12	11.09	324.5
United States	Defence aggregate	Baseline		7.43	16.89	−509.5
United States	Defence aggregate	SJM	0.243	7.22	24.91	−816.7
United States	Defence aggregate	BSX	0.220	8.82	11.78	−1003.0
United States	Defence aggregate	EME	0.123	8.93	12.68	−327.2
United States	Defence aggregate	SMCI	0.083	7.82	10.84	−723.2
United States	Defence aggregate	GLW	0.055	7.45	18.41	−410.3

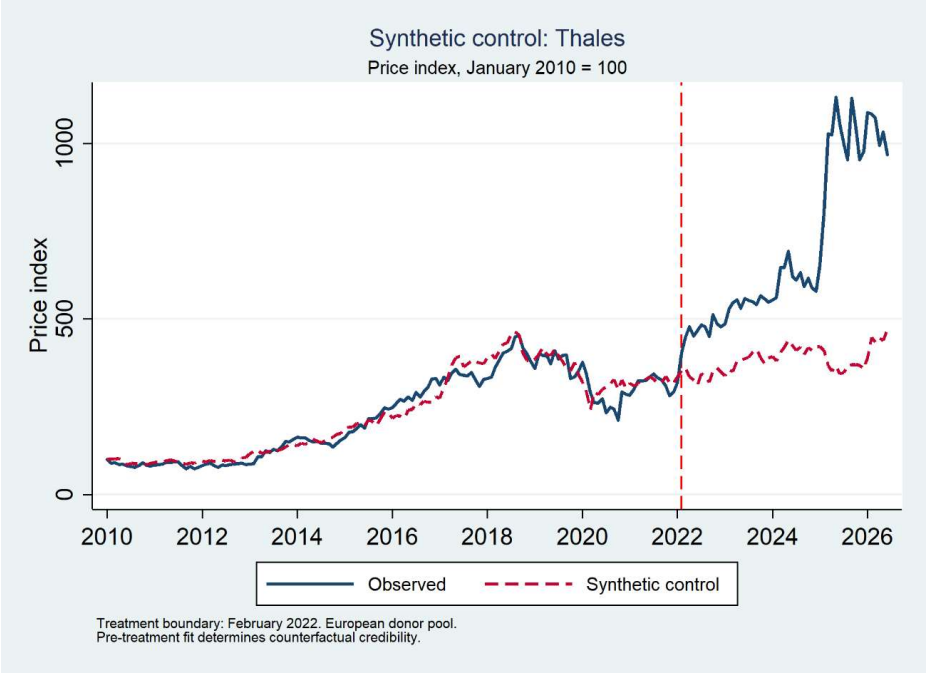
Notes. Baseline rows report the original synthetic-control estimates. The remaining rows report estimates obtained after separately excluding each principal donor. Donor weight refers to the weight assigned to the excluded firm in the baseline synthetic control. Pre-treatment RMSPE is expressed as a percentage of the average treated index. The final gap is the difference between the treated and synthetic price indices in June 2026.

Figure B.1 Observed Synthetic Price Indices for Rheinmetall



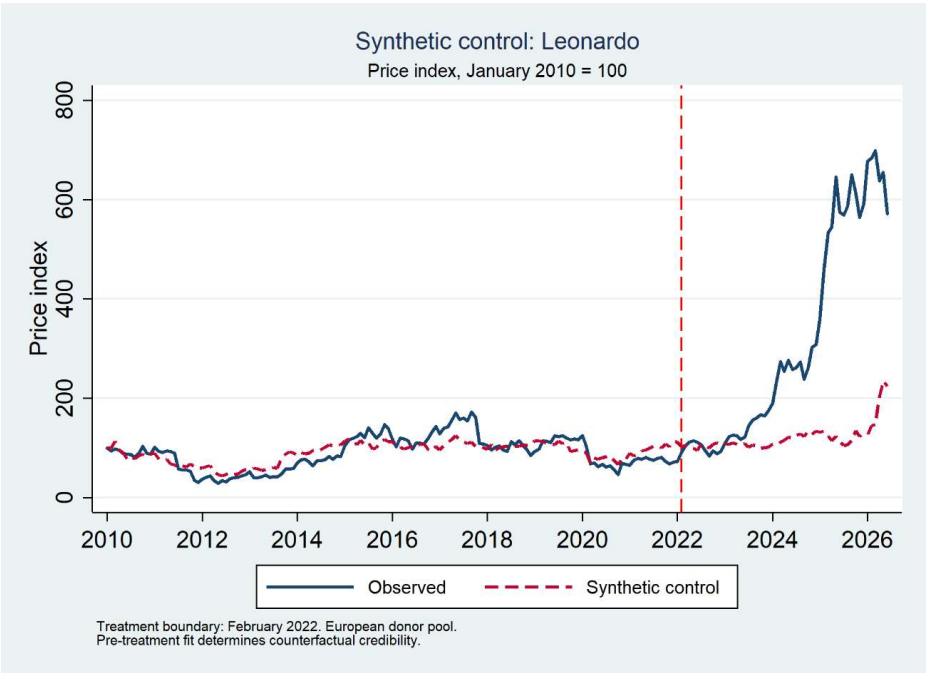
Notes. Price indices are rebased to 100 in January 2010. The vertical line identifies February 2022. The synthetic counterfactual is constructed from eligible European control firms.

Figure B.2 Observed and Synthetic Price Indices for Thales



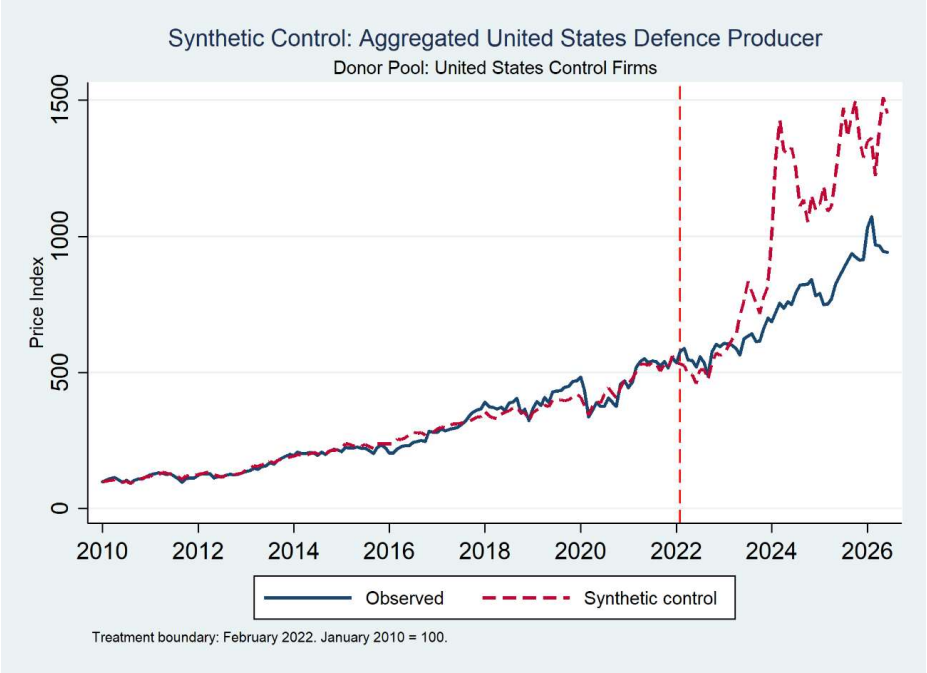
Notes. Price indices are rebased to 100 in January 2010. The vertical line identifies February 2022. The synthetic counterfactual is constructed from eligible European control firms.

Figure B.3. Observed and Synthetic Price Indices for Leonardo



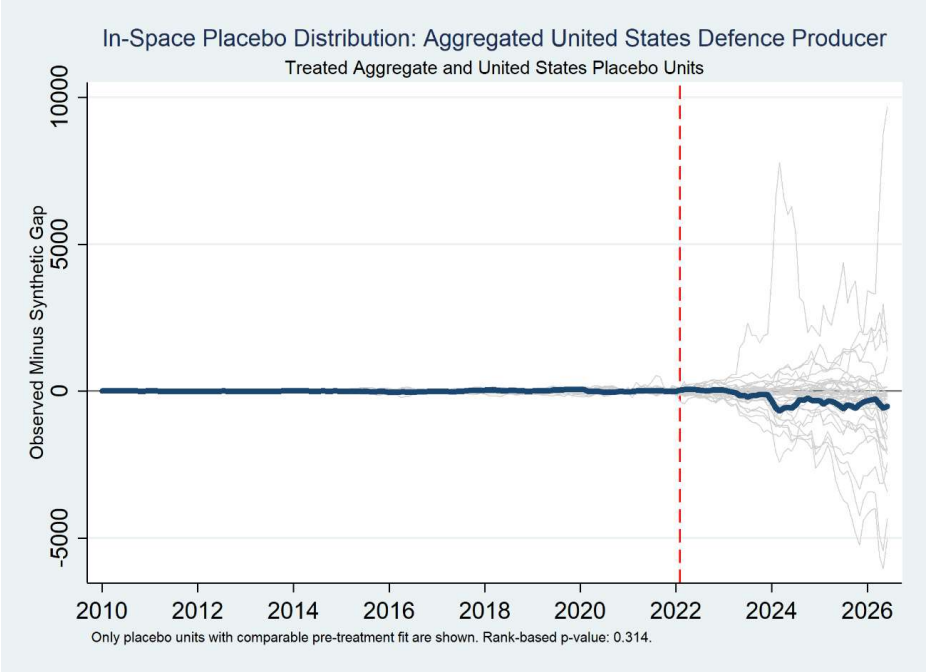
Notes. Price indices are rebased to 100 in January 2010. The vertical line identifies February 2022. The synthetic counterfactual is constructed from eligible European control firms.

Figure B.4. United States Defence Aggregate and Synthetic Control



Notes. Price indices are rebased to 100 in January 2010. The vertical line identifies February 2022. The synthetic counterfactual is constructed from eligible United States control firms.

Figure B.5. In-Space Placebo Distribution for the Aggregated United States Defence



Notes. The solid line shows the treated aggregate gap and the grey lines placebo gaps passing the fit restriction. The vertical line marks February 2022; rank-based p = 0.314.