

**GEOPOLITICAL RISK, OIL PRICE EXPOSURE, AND
FIRM'S TAIL RISK:
EVIDENCE FROM EXPECTED SHORTFALL ANALYSIS
OF KOMPAS 100 CONSTITUENTS**



THESIS

By:

CHERYL JUSUF

22111004

Academic Supervisor:

Dr. Ardo Ryan Dwitanto, SE, MSM, CFP

**BACHELOR OF BUSINESS ADMINISTRATION PROGRAM
INSTITUT IPMI JAKARTA 2026**

**GEOPOLITICAL RISK, OIL PRICE EXPOSURE, AND
FIRM'S TAIL RISK:
EVIDENCE FROM EXPECTED SHORTFALL ANALYSIS OF
KOMPAS 100 CONSTITUENTS**

Prepared By:
Cheryl Jusuf (22111004)

A THESIS

Submitted in partial fulfillment of the requirements for the degree of Bachelor of
Business Administration

CERTIFICATE OF APPROVAL

Name & Student ID : Cheryl Jusuf (22111004)

Topic : Geopolitical Risk, Oil Price Exposure, and Firm's Tail
Risk: Evidence from Expected Shortfall Analysis of Kompas 100 Constituents

We hereby declare that this Thesis is from student's own work, has been read and presented to Institut IPMI Board of Examiners, and has been accepted as part of the requirements needed to obtain a Bachelor of Business Administration Degree and has been found to be satisfactory.

Jakarta, August 27th 2026

Examined by,



Dr. Ardo Ryan Dwitanto,
SE, MSM, CFP



Dety Nurfadilah, B.If.,
MBA



Dr. Firdaus Fanny Putera
Perdana

Acknowledged by,



Prof. Dr. Christina Whidya Utami MM., CLC., CPM (A)

Vice Rector of Academic and Student Affairs

NON-PLAGIARISM DECLARATION FORM

This Thesis is a presentation of our original research work. Wherever contribution of others are involved, every effort is made to indicate this clearly, with due reference to the literature, and acknowledgement of collaborative research and discussions.

Also, this work is being submitted in partial fulfillment of the requirements for the Bachelor of Business Administration degree and has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Jakarta, 14 August 2026

Materai 10.000



Cheryl Jusuf

TABLE OF CONTENTS

LIST OF FIGURES.....	9
LIST OF TABLES.....	10
ACKNOWLEDGEMENT	11
ABSTRACT	12
CHAPTER 1.....	13
INTRODUCTION.....	13
1.1. Background of the Study.....	13
1.2. Problem Identification.....	22
1.3. Research Questions	23
1.4. Research Objectives	24
1.5. Scope of Study	24
1.6. Limitations	25
1.7. Significance of the Study	26
1.7.1. Theoretical Contribution	26
1.7.2. Practical Contribution	27
1.7.3. Directions for Future Research	27
1.8. Systematic of the Thesis.....	28
CHAPTER 2.....	29
LITERATURE REVIEW	29
2.1. Previous Research	29
2.1.1. Geopolitical Risk and Equity Markets	31
2.1.2. Oil Price Shocks and Emerging Market Equities.....	32

2.1.3. Cost Structure, Operating Leverage, and Equity Risk	33
2.1.4. Profitability and Firm Resilience	34
2.1.5. Expected Shortfall as a Firm-Level Outcome Variable.....	34
2.1.6. Research Gap and Value Added to Previous Studies	35
2.2. Theoretical Framework	37
2.2.1. Theory of Risk.....	38
2.2.2. From Volatility to Tail Risk.....	41
2.2.3. The Oil Cost Channel and Asymmetry of Oil Shocks	42
2.2.4. Cost Intensity and Operating Leverage: What EXPO Measures	44
2.2.5. The Sentiment Channel	46
2.2.6. Profitability as a Risk Buffer.....	46
2.3. Hypothesis Development	47
2.4. Research Framework.....	49
CHAPTER 3.....	51
RESEARCH METHODOLOGY	51
3.1. Research Design.....	51
3.1.1. Choice of Annual Frequency.....	52
3.2. Population and Sample.....	53
3.3. Variable Operationalization.....	54
3.3.1. Expected Shortfall (ES5), Dependent Variable	54
3.3.2. GPR Index, Independent Variable	57
3.3.3. GEAR (Geopolitical Energy-Adjusted Risk), Independent Variable	58
3.3.3.1. Construct Validity of GEAR	62
3.3.3.2. Exclusion of Negative Oil Price Movements.....	64
3.3.3.3. Uniform GEAR Across Firms	65

3.3.4. Firm Exposure (EXPO), Independent Variable	66
3.3.4.1. Construct Validity of EXPO	67
3.3.5. ROA (Return on Assets), Control Variable	68
3.3.6. Interaction Term ($\text{GEAR0} \times \text{EXPO}$)	69
3.3.7. Regression Specification	70
3.4. Analytical Procedure	72
3.4.1. Descriptive Statistics	73
3.4.2. Correlation Matrix and Multicollinearity	74
3.4.3. Classical Assumption Tests	77
3.4.4. Panel Model Selection	79
3.4.5. Significance Testing	84
3.4.6. Robustness Checks	86
3.5. Data Collection, Software, and Tools	87
3.5.1. Data Verification	88
CHAPTER 4.....	90
RESULTS AND DISCUSSION	90
4.1. Findings and Analysis	90
4.1.1. Descriptive Statistics	90
4.1.2. Panel Data Model Selection	92
4.1.3. Correlation Matrix	94
4.1.4. Hypothesis Testing	96
4.2. Robustness Checks	101
4.3. Discussion	104
4.3.1. Geopolitical Risk and Tail Risk (H1)	104
4.3.2. Geopolitical Energy-Adjusted Risk and Tail Risk (H2).....	105

4.3.3. Firm Cost Exposure and Tail Risk (H3).....	106
4.3.4. The Moderating Role of Cost Structure (H4).....	107
4.3.5. Profitability as a Control Variable.....	108
4.3.6. Summary of Findings.....	109
CHAPTER 5.....	110
CONCLUSION AND RECOMMENDATIONS	110
5.1. Conclusion.....	110
5.2. Limitations	112
5.3. Implications of the Study	115
5.3.1. Theoretical Implications.....	115
5.3.2. Managerial and Practical Implications.....	117
5.4. Recommendations	118
5.4.1. Recommendations for Investors and Portfolio Managers.....	118
5.4.2. Recommendations for Regulators and Policymakers.....	118
5.4.3. Recommendations for Future Research	119
REFERENCES.....	123

LIST OF FIGURES

Figure 1.1. Geopolitical Risk (GPR) Index, January 2014 – April 2026	14
Figure 1.2. Global Crude Oil Prices (Brent, WTI, Murban), January 2014 – May 2026.....	17
Figure 2.1. Research Framework	50
Figure 3.1. Research Process.....	73

LIST OF TABLES

Table 2.1. Summary of Previous Related Research	29
Table 2.2. Research Hypotheses.....	48
Table 3.1. Simulated Daily Returns (Sorted Ascending)	56
Table 3.2. Variable Operationalization Summary	70
Table 3.3. Interpretation of Correlation Magnitude	75
Table 3.4. Panel Model Selection Tests.....	81
Table 4.1. Descriptive Statistics	90
Table 4.2. Chow Test Result.....	92
Table 4.3. Hausman Test Result	93
Table 4.4. Breusch-Pagan Lagrange Multiplier Test Result.....	93
Table 4.5. Summary of Panel Model Selection Tests.....	93
Table 4.6. Pearson Correlation Matrix	94
Table 4.7. Variance Inflation Factors.....	96
Table 4.8. Heteroskedasticity Test Results	96
Table 4.9. Regression Results Without Moderating Variable.....	96
Table 4.10. Regression Results With the Interaction Term	97
Table 4.11. Comparison of Panel Estimation Models	98
Table 4.12. Summary of Hypothesis Testing	100
Table 4.13. Robustness Check Results.....	101
Table 4.14. Leave-One-Year-Out Estimation	103

ACKNOWLEDGEMENT

First and foremost, I would like to express my deepest gratitude to Allah SWT for His blessings, guidance, and strength throughout the process of completing this thesis. It has given me the perseverance and determination needed to overcome the challenges encountered throughout this research journey.

I would like to express my sincere appreciation to my mentor, Dr. Ardo Ryan Dwitanto, SE, MSM, CFP, for his invaluable guidance, constructive feedback, encouragement, and continuous support throughout the preparation and completion of this thesis. His insights and recommendations have been instrumental in improving the quality and direction of this research.

My deepest gratitude also goes to my family and friends, who have continuously supported, encouraged, and pushed me to complete this thesis. Their patience, prayers, motivation, and constant reminders have given me the strength to keep moving forward, especially during the more challenging stages of this journey.

I would also like to extend my special appreciation to Bahana Sekuritas for kindly providing access to Bloomberg Terminal facilities, which greatly assisted me in obtaining the necessary data for this research. Their support was invaluable to the data collection and research process.

I am also sincerely grateful to Bu Dety Nurfadilah, B.IfB., MBA, Dr. Firdaus Fanny Putera Perdana and Prof. Ir. H.M. Roy Sembel, MBA, Ph.D, who served as examiners and provided insightful feedback, constructive comments, and valuable perspectives that helped improve and strengthen this thesis.

Finally, I would like to express my appreciation to everyone who has contributed, directly or indirectly, to the completion of this thesis. Every form of support, encouragement, and assistance has played an important role in helping me reach this achievement. For that, I am truly grateful.

ABSTRACT

This study investigates the influence of geopolitical risk, energy exposure, and firm cost structure on the equity tail risk of non-financial companies listed on the Indonesia Stock Exchange during 2015–2025. Tail risk is measured through Expected Shortfall at the fifth percentile (ES5), computed from daily returns as the mean of the worst five percent of trading days. Two variables are constructed for this study: Geopolitical Energy-Adjusted Risk (GEAR), which weights the annual change in a three-benchmark composite crude oil price by Indonesia's Middle East share of petroleum imports, and Firm Exposure (EXPO), measured as cost of goods sold to net sales, which also serves as the moderating variable. The Geopolitical Risk Index enters as a second geopolitical measure and return on assets as a control. Estimation is on 54 firms and 594 firm-year observations. Geopolitical risk, energy exposure and cost intensity each come out negative against tail risk, the energy estimate consistently so across specifications and the other two with less security, which means that in years of higher geopolitical tension the extreme losses recorded across the sample were milder rather than deeper. Two features of the setting offer a plausible reading of that pattern, though neither is tested directly. The Indonesian listed market is composed largely of commodity-linked issuers, and Indonesia both produces and imports oil, so a price increase arrives as cost for some firms and as revenue for others. On cost structure, the negative coefficient is compatible with a cost-flexibility reading of Firm Exposure (EXPO) rather than the reading under which a high cost ratio simply means a thin operating margin, though the evidence does not settle the question. Cost structure was also tested as a moderator of the geopolitical energy channel and no moderation was detected. Whether an energy shock helps or harms a firm may turn on which side of the income statement it lands on, and that remains to be tested at the industry level.

Keywords: Expected Shortfall, Geopolitical Risk, Geopolitical Energy-Adjusted Risk, Firm Cost Exposure, Operating Leverage, Panel Data Regression, Kompas 100, Indonesia Stock Exchange

CHAPTER 1

INTRODUCTION

1.1. Background of the Study

Empirical finance has measured equity risk with variance, beta and standard deviation for most of its history. These describe how widely returns are spread. About the far left of the distribution, where investors actually lose money and where positions get closed under pressure and risk limits are tightened, they say comparatively little. Consider two stocks with identical standard deviations. One delivers frequent losses, none of them fatal. The other is quiet for years and then falls very hard across a handful of days. An average measure of variation cannot tell those cases apart, and the failure matters more in emerging markets, where extreme moves carry real consequences and have attracted less study than ordinary volatility.


Expected Shortfall, also known as Conditional Value at Risk, was built for exactly this region of the distribution. Value at Risk marks a loss threshold at a chosen confidence level and stops there. Expected shortfall looks past the threshold and averages the losses that occur beyond it, so it reports both where the tail begins and how bad matters become once returns fall inside. That property is what recommends it for thinking about severe shocks.

Making expected shortfall the thing to be explained raises the question of what moves it, and the candidates divide. Some are external and macroeconomic, reaching the firm from outside the country altogether, and geopolitical conflict is the most disruptive of these. Others are internal, principally the cost structure a company operates on and the profitability it can draw against. This study takes up both, and asks whether firm characteristics change how hard the external shock lands.

Geopolitical risk is awkward to work with because it is abstract. Tension between states cannot be read off a screen the way an interest rate can, and anyone proposing to treat it as a variable has to decide what to count.

Caldara and Iacoviello (2022) solve this by counting words. Their Geopolitical Risk (GPR) Index measures the share of articles across ten major English-language newspapers that discuss geopolitical threats and military tensions. The reasoning is that how much the press covers something is a reasonable proxy for how much uncertainty economic agents are actually absorbing. The index is published daily and is aggregated to annual averages for this study.

Figure 1.1. Geopolitical Risk (GPR) Index, January 2014 – April 2026



Source: Caldara and Iacoviello (2022); processed by Author (2026)

Figure 1.1 plots the index with the period's principal escalations marked. The baseline rises and grows more volatile. From an annual average of 99.6 in 2014 the index eased to 78.0 in 2020, climbed to 153.0 in 2022 alongside the invasion of Ukraine, and reached 211.2 in 2026 during the confrontation between the United States, Israel and Iran. Within the 2015–2025 observation window the record includes the territorial expansion of ISIS, the Saudi-led intervention in Yemen from March 2015, the American withdrawal from the JCPOA in May 2018 and renewed sanctions on Iranian crude, the killing of General Qasem Soleimani in January 2020, the Russia-Ukraine war from February 2022, the outbreak of the Israel-Hamas conflict in October 2023, and Houthi attacks on Red Sea shipping thereafter. Each

coincided with movements in crude prices, safe-haven assets and emerging market valuations, though direction and persistence varied considerably between episodes.

What stands out is how often the index returns to one region. That is no accident, and the reason matters, because the argument developed here rests on the Middle East being structurally unstable rather than episodically unlucky.

The region's borders were drawn largely from outside. The Sykes-Picot understanding of 1916 and the mandate system that followed the Ottoman collapse produced frontiers reflecting imperial convenience rather than tribal, sectarian or ethnic geography. States were created whose populations shared little and whose ruling families often held no historical claim to the territory granted them. Saudi Arabia is a partial exception, consolidated by the Al Saud themselves, but much of the Peninsula and the Levant was assembled externally. The result is a set of neighbours with overlapping claims, unresolved grievances, and little incentive to treat the map as settled.

Layered over this is the Sunni-Shia division, originating in the seventh-century dispute over succession to the Prophet Muhammad and still politically live. Iran is the demographic and institutional centre of Shia Islam; most Arab states, Saudi Arabia among them, are Sunni-majority and governed as such. The rivalry is not confined to theology. It has been pursued through proxies, financing and armed intervention across Iraq, Syria, Lebanon and Yemen, and it shapes alignment on nearly every regional question.


Iran shows how fast the political character of a regional state can change. Before 1979 it was a monarchy under Shah Mohammad Reza Pahlavi, aligned with Washington and embedded in Western security arrangements in the Gulf. The Islamic Revolution replaced the monarchy with a clerical republic, the Shah went into exile, and foreign policy reversed almost overnight. War with Saddam Hussein's Iraq followed from 1980 to 1988 at a cost of several hundred thousand lives. The American-led invasion of 2003 removed Saddam and transferred power to Iraq's Shia majority; Iranian influence subsequently expanded through political parties, armed groups and economic ties, deepening the sectarian fault line rather

than closing it. ISIS grew directly out of that unsettled aftermath. Israel adds a further axis, its establishment in 1948 and the wars that followed leaving disputes that reopen periodically, as events after October 2023 made clear.

Even the name of the waterway at the centre of the region's oil trade is contested: Iran and most international bodies say Persian Gulf, Arab states insist on Arabian Gulf, and the argument has outlasted several generations of diplomacy. The implication is larger than the example. Where parties cannot agree what to call a body of water, a durable settlement of borders, sectarian authority and energy rents is optimistic. For a financial economist the practical consequence is that Middle East tension is better modelled as a persistent condition of varying intensity than as a sequence of independent surprises.

The episode that made this impossible for Indonesian market participants to ignore was the Strait of Hormuz crisis of 2026. On 28 February the United States and Israel began coordinated airstrikes on Iran and Supreme Leader Ali Khamenei was killed. Iran retaliated against Israeli cities and American installations across the Gulf. On 4 March the Islamic Revolutionary Guard Corps declared the Strait closed and threatened any vessel attempting transit; ship-tracking data showed traffic falling roughly seventy percent within hours. The International Energy Agency (2026) called the interruption the largest supply disruption in the history of the oil market, affecting approximately twenty million barrels per day of transit. Brent moved from about USD 73 in late February to a peak of USD 126, the largest monthly increase on record, and gold passed USD 3,000 per ounce. Asian equities fell, with the Nikkei 225 down more than six percent and the KOSPI down 6.3 percent. In Indonesia the Jakarta Composite Index fell 4.57 percent on 4 March with 734 stocks declining, and by mid-March had recorded the weakest year-to-date performance among thirty-seven monitored markets. Energy was the only sector in positive territory.

Figure 1.2. Global Crude Oil Prices (Brent, WTI, Murban), January 2014 – May 2026



Source: Bloomberg Terminal (CO1, CL1, MUC1 Comdty); processed by Author (2026)

The 2026 crisis falls outside the estimation window, which runs from 2015 to 2025. It is used here as motivation, not as evidence. Its role is conceptual. The episode made a mechanism visible in unusually sharp form. Tension in a region supplying a large share of traded crude disrupted the physical movement of that crude. Prices rose. Equities across oil-importing Asia came under pressure while producers gained. Once a mechanism is stated it can be tested, but not on a single year. Testing it needs a panel long enough to contain real variation in both geopolitical intensity and oil prices, and full-year 2026 financial statements were not available when this was written. The 2015 to 2025 window contains the ISIS expansion, Yemen, the sanctions period, Soleimani, Ukraine and Israel-Hamas. These are repeated instances of the same mechanism at lower amplitude. The empirical claims in this thesis are about that window. The 2026 episode returns in Chapter 5 as an out-of-sample reference point.

Indonesia's position in the transmission chain follows from its energy balance. The country has been a net crude importer since roughly 2004. Crude oil and condensate production fell to 582,000 barrels per day in 2024, a four percent decline on the year as mature fields deplete, while consumption of petroleum and other liquids reached 1.7 million barrels per day, its highest level since 2018 (U.S. Energy Information Administration, 2025). The gap has widened rather than closed. Between roughly 30 and 58 percent of Indonesia's petroleum imports, crude and refined together, originated in the Middle East across the observation window, and those cargoes transit Hormuz. That share matters here in a specific, mechanical way: it is the weight converting a global price movement into an Indonesia-specific exposure, and it forms the second component of the principal explanatory variable. Broader energy policy questions, including refining capacity, reserve adequacy and subsidy design, shape how a shock is ultimately distributed between the state budget, consumers and firms. None enters the model, and they are set aside so that the background and the empirical scope describe the same study.

The two threads must be joined into something measurable. One route models the relationship in two stages: geopolitical risk raises oil prices, which raise equity tail risk. That indirect specification is defensible but treats oil as a generic commodity variable and loses the geopolitical content of the movement, since prices also move for reasons unrelated to conflict. This study takes the alternative route, combining both elements into a single variable termed Geopolitical Energy-Adjusted Risk, abbreviated GEAR.

GEAR has two components. The first is the annual change in the world crude price, measured as the average of Brent, West Texas Intermediate and Murban, which together span Atlantic basin, North American and Gulf pricing. The second is Indonesia's Middle East share of petroleum imports in the same year. GEAR is the product of the two, so a global price increase is scaled by how much of Indonesia's crude actually comes from the region where tension originates. Read in words: the more a country depends on imports from a geopolitically unstable region, the larger the share of any energy shock originating there that it must absorb, and the more risk is transferred onto firms operating within that economy.

One feature of the construction requires comment. GEAR retains only positive oil price changes, defined with a $\text{Max}(0, \cdot)$ operator so that years of falling prices are recorded as zero rather than negative exposure. The variable represents the cost of energy insecurity, and a decline driven by supply glut or weak demand is not a geopolitical shock in reverse. Rising and falling prices carry asymmetric information about regional tension. The choice discards information, and in this sample the censoring binds in six of eleven years, a cost weighed against the alternative in Chapter 3.

GPR and GEAR are retained separately because they answer different questions. GPR is global and agnostic about transmission: it registers tension in Eastern Europe, East Asia or the Middle East alike and assumes nothing about how that tension reaches Indonesian firms. GEAR is deliberately narrow, covering one region, one channel, and a weight specific to Indonesia's import pattern. Including both allows a question neither answers alone, namely whether Indonesian tail risk responds more to worldwide geopolitical tension or specifically to the Middle East energy channel.

Macroeconomic shocks do not land evenly across firms. Higher imported energy costs pass through the income statement of a petrochemical producer very differently from that of a tower operator. Any study working at firm level needs a variable that captures the difference. The measure used here is the ratio of cost of goods sold to net sales, called Firm Exposure and abbreviated EXPO. Its simplicity is deliberate.

This can be calculated for each firm in each year from audited statements, which is not true of more targeted measures. EXPO has two readings with different signs. Under the margin reading, for example, a firm with 85 percent of revenue consumed by direct costs has only a 15 percent cushion, which an input price increase eats into. However, the same increase at a firm reporting 50 percent is absorbed more comfortably, so high EXPO means vulnerability and a deeper tail. Under the flexibility reading, cost of goods sold is mostly variable and selling and administrative expense is mostly fixed, so a high ratio describes a cost base that

scales with output. Such a company has low operating leverage, and can reduce production and cost together when demand weakens. An airline provides a useful contrast because expenses such as aircraft leases can remain in place even when fewer seats are sold. From this perspective, a higher EXPO could indicate greater cost flexibility and therefore less sensitivity to a downturn, or in other words to tail risk. These interpretations lead to different expectations about the sign of the relationship between EXPO and tail risk. The variable is therefore not assigned a predetermined direction in the analysis. Whether a higher EXPO is associated with greater or lower downside risk is left to the results.

A caveat belongs here.

EXPO reports the size of a firm's direct cost base and nothing about its contents, so two firms showing the same ratio can meet the same oil price move from opposite directions. A cleaner variable isolating the fuel component would settle it and cannot be built from Indonesian disclosure; Section 3.3.4.1 works through the cases and what they mean for reading the coefficient.

The composition of the sample matters as much as the construction of the variable. Resource and commodity-linked issuers make up roughly a third of the sample by number and by capitalisation, and Indonesia sells gas, coal, nickel and palm oil as well as buying imported crude, so a geopolitical event that lifts commodity prices reaches a substantial part of this sample as revenue. They do not dominate the index, but their weight is large enough that when a shock lifts their revenue the pooled coefficient can take their sign. The sign of the GEAR relationship is therefore open before anything is estimated, a structural feature of the market rather than an artefact of sampling, and the point on which the setting departs from Japan or South Korea. Section 4.3.1 takes it up against the results.

Profitability is the second firm-level characteristic. Earnings retained in good years give a firm something to draw on when input costs rise, and that cushion should show up as a shallower left tail (Novy-Marx, 2013). It enters the model as a control, and Section 3.3.5 sets out what that does and does not imply.

Moderation is studied for one pair, GEAR and EXPO, and their product completes the model. It earns its place because the two readings of EXPO disagree about what sign it should take. On the first, a large direct cost base is a thin cushion, so a year of elevated energy prices should hurt cost-intensive firms more and the product should come out positive. The second reading treats that same base as something management can act on, since production comes down when demand softens and raw material, freight and contracted labour come down with it, which would leave cost-intensive firms no worse off and possibly better. The main effects cannot tell those apart. The product can.

GPR and crude prices are published daily. Annual averages are used anyway, because of what the study is trying to measure: not how the market reacted to any one incident, but how severe a firm's worst trading days were across a year of sustained tension. Caldara and Iacoviello built their index for exactly that reading, treating geopolitical risk as a standing condition present between the headlines, shaping hedging and investment decisions taken over months. Nothing can be done about frequency at the firm level in any case. EXPO and ROA come off audited annual statements and exist nowhere else, so a quarterly design would mean giving up the variables the study is built on. Little is lost by that. Expected shortfall is computed from daily returns inside each firm-year, so the tail days survive in the dependent variable even when the regressors do not.

Each piece of this exists in the literature. The combination does not. Geopolitical risk has been linked to emerging market equity returns at index level (Bouras, Christou, Gupta and Suleman, 2019; Zaremba, Cakici, Demir and Long, 2022), and cost structure to equity risk since Lev (1974) and Mandelker and Rhee (1984). The oil side is the most developed of the three, and Kilian and Park (2009) supply the finding this thesis leans on hardest: what an oil price increase does to equities depends on what caused it, since a supply disruption pushes returns down where a demand-driven increase does not. An oil variable carrying no information about origin is therefore measuring the wrong thing. Basher and Sadorsky (2006) establish oil as a priced factor across emerging markets, and for Indonesia the effects have been traced sector by sector (Sukarno and Setiawati, 2020; Wahyudi

and Rahmawati, 2025). Expected shortfall has a methodological literature of its own (Artzner et al., 1999; Acerbi and Tasche, 2002) and has been estimated on Indonesian firm-level panel data by Dwitanto, Manurung, Machdar and Hidayat (2026), though their determinants are governance and asset strategy, not macroeconomic conditions. No study the author has found puts expected shortfall on the left-hand side together with a geopolitical measure weighted by a country's own import dependence, a firm-level cost ratio, and an interaction between them.

The panel is balanced: 54 non-financial Kompas 100 constituents, 2015 to 2025, 594 firm-year observations, estimated by panel generalised least squares. Chow, Lagrange Multiplier and Hausman did not agree on a specification, so fixed effects with cross-section weighting is reported as primary with random effects alongside. Section 4.1.2 sets out why. The design is observational. Nothing in the specification isolates exogenous variation in geopolitical risk, so every result is a conditional association, and the wording used throughout the thesis keeps to that.

1.2. Problem Identification

The background above gives rise to five research problems.

1. Tensions in the Middle East are a chronic condition that reaches Indonesia mainly through the price of imported crude, but existing measures capture only one side of the relationship. The GPR Index indicates tension but says nothing about whether a country is exposed to it. Oil price changes capture the energy channel but do not tell us about geopolitical origin or sourcing. No measure combining the size of the price move with dependence on Middle East supply has been constructed or tested.
2. The exposure is not uniform across firms, and the sign of the difference is uncertain. Firms with a high cost ratio may be more vulnerable because thinner margins leave less cushion, or less vulnerable because a variable cost base can be reduced. The interaction between GEAR and EXPO is supposed to separate these, but its sign and significance are untested in Indonesia.

3. Most of this literature makes use of average returns or volatility, neither of which measures the severity of extreme losses. A single-day drop of 4.57 percent in the IHSG is a tail event, and a variance-based measure treats it as one observation among many. Expected shortfall is considered theoretically superior and regulator approved, but has not been applied to Indonesian firm-level panel data in a geopolitical context.
4. The relative importance of global and region-specific geopolitical risk for Indonesian equities is an unresolved question. It has not been estimated jointly at firm level whether firms respond more to worldwide tension or to the narrower Middle East energy channel.
5. Indonesia's commodity-heavy market complicates the standard oil-importer narrative. The country imports crude but exports coal, nickel, palm oil and gas, and much of the non-financial Kompas 100 derives revenue from commodity sales. Whether the net relationship is amplifying or dampening is an empirical question requiring firm-level data rather than index aggregates. Together these point to a single framework combining macro geopolitical risk, country-weighted energy exposure, firm cost structure, a profitability control and the interaction, estimated on firm-level panel data with expected shortfall as the dependent variable.

1.3. Research Questions

Based on the problem identification, the research questions are:

1. Does geopolitical risk (GPR) significantly affect the expected shortfall of Kompas 100 stocks?
2. Does geopolitical energy-adjusted risk (GEAR) significantly affect the expected shortfall of Kompas 100 stocks?
3. Does firm cost exposure (EXPO), measured as COGS to sales, significantly affect the expected shortfall of Kompas 100 stocks?
4. Does the interaction between GEAR and EXPO significantly affect expected shortfall, indicating a moderating role of cost structure?

The questions are stated non-directionally, using “significantly affect” rather than “increase” or “decrease”, for two reasons. Theory delivers no unambiguous prediction of sign for EXPO or the interaction, and imposing a restriction would risk overlooking a relationship running the other way. The commodity-weighted composition of the market likewise leaves the sign of the GEAR relationship open. The word “affect” is used in its statistical sense of a non-zero coefficient estimated on observational data, not as a claim of causal identification.

1.4. Research Objectives

The objectives of this study are:

1. To examine the relationship between geopolitical risk and the expected shortfall of Kompas 100 stocks.
2. To analyse the relationship between geopolitical energy-adjusted risk and expected shortfall.
3. To examine the relationship between firm cost exposure and expected shortfall.
4. To test whether the interaction between GEAR and firm cost exposure exhibits a significant moderating relationship with expected shortfall.

The study further aims to describe the behaviour of GPR, GEAR, EXPO and expected shortfall across the sample period, and to assess robustness across alternative specifications including different tail thresholds, exclusion of the pandemic year, and leave-one-year-out estimation.

1.5. Scope of Study

The research covers constituents of the Kompas 100 index on the Indonesia Stock Exchange over 2015 to 2025, giving eleven years of annual panel data. Banking and financial institutions, including BBKA, BMRI, BBRI, BBNI and BBTN, are excluded because their cost structure rests on cost of funds rather than cost of goods sold, which renders the COGS to sales ratio economically meaningless for them, and because bank costs are driven by interest rates rather than energy prices, so the

GEAR channel does not apply. A second criterion is the balanced panel requirement: each firm must have complete data for every variable across all eleven years, which removes firms listing after 2015, delisted during the period, or lacking continuous reporting. The final sample is 54 firms and 594 firm-year observations.

The dependent variable is expected shortfall at the fifth percentile in absolute value. Independent variables are GPR, GEAR under the $\text{Max}(0, \cdot)$ specification, and EXPO. Cost structure enters the interaction as the moderating variable, and ROA enters as a control. ES1 and ES10 are used as robustness checks. Every variable is annual. Because expected shortfall is computed inside each firm-year from daily returns, the information in individual trading days survives in the dependent variable. On the macro side the study measures tension sustained over a year instead of the reaction to a single incident, and on the firm side EXPO and ROA come from audited annual statements and cannot be observed more often on any consistent basis.

The 2026 Hormuz crisis sits outside the empirical work. It motivates the study and returns in Chapter 5 as an out-of-sample reference, but no 2026 observation enters the regression sample.

The IDX Index Committee reviews Kompas 100 composition semi-annually with the Kompas editorial team, looking at transaction frequency and value, market capitalisation on full and free-float bases, trading days against market totals, and fundamental considerations. The universe is therefore liquid and investable, which lowers the risk that thin trading distorts expected shortfall estimates built from daily returns. A balanced panel does carry a mild survivorship characteristic, since firms that failed or lost liquidity cannot appear in it. Large-capitalisation indices limit that problem, as delisting is uncommon and index changes usually substitute comparable firms.

1.6. Limitations

Several constraints bound what this study can establish. They are given here in outline so that the design is read alongside them, and Section 5.2 states each in full

with what it costs the interpretation. The evidence is observational, so the results are conditional associations and not causal effects. The two macro variables take one value per year, which means their coefficients rest on eleven annual observations and not on the full 594. GEAR is built from a national import share and from oil price changes floored at zero, making it a measure of national exposure in years of rising prices instead of firm-specific geopolitical intensity. EXPO measures the size of a firm's direct cost base and not its energy content. The observation window includes the COVID-19 pandemic, which produced extreme tail events in 2020 with no geopolitical origin. The model is linear. Financial institutions are excluded, and the balanced panel requirement leaves only firms listed throughout the period. Expected shortfall is computed from simple returns rather than logarithmic returns, following the convention agreed with the academic supervisor.

1.7. Significance of the Study

1.7.1. Theoretical Contribution

The contribution is primarily one of measurement. Two variables are constructed here that have not previously been applied to Indonesian firm-level tail risk, and the claim rests on demonstrating that they carry explanatory content rather than on novelty in estimation technique.

GEAR is a country-weighted measure of geopolitical energy exposure. It advances no new theory of geopolitical risk; it combines two established series, a threebenchmark crude average and Indonesia's Middle East import share, so that a global price movement is expressed as a national exposure specific to the region generating the tension. Existing work uses either a raw oil series, common to every country and carrying no geopolitical content, or a risk index carrying no energy content and no national weighting. Neither speaks to how much of a shock a particular economy is positioned to absorb. Estimated jointly with GPR, GEAR can be tested directly for explanatory power beyond the general measure.

EXPO is applied as a determinant of expected shortfall. Cost structure has an established relationship with systematic risk (Lev, 1974; Mandelker and Rhee, 1984; Novy-Marx, 2011) and more recently with crash risk (Zhu and Zang, 2025), but the author has not located a study relating a direct cost intensity ratio to a coherent tail measure at firm level, and none in Indonesia. The variable also carries an unresolved ambiguity of sign, which this setting is positioned to address.

A third, more qualified contribution concerns the dependent variable. Expected shortfall is established in risk measurement and banking regulation and has already been applied to Indonesian firm-level data by Dwitanto et al. (2026). What remains undone is modelling firm-level expected shortfall against geopolitical and energy determinants.

The interaction is retained as a test of whether cost structure conditions transmission, connecting the geopolitical literature to the operating leverage literature. It is a component of the model rather than a principal contribution; the value of the study does not depend on the moderation reaching significance.

1.7.2. Practical Contribution

For risk managers, the results indicate which firm characteristics accompany deeper or shallower tails during geopolitical stress, which is directly usable in stress testing and position sizing. For investors, the study offers a screening logic grounded in observable financial statement data rather than sector labels, an approach whose relevance was illustrated by the divergence between energy and non-energy stocks on the IDX during March 2026. For policymakers at Bank Indonesia and the Otoritas Jasa Keuangan, firm-level evidence on where tail exposure to imported energy shocks concentrates supports more targeted macroprudential monitoring than aggregate index analysis permits.

1.7.3. Directions for Future Research

The GEAR methodology can be adapted to other oil-importing emerging markets, including Thailand, the Philippines, India and Turkey, by substituting their respective import shares. The expected shortfall framework can be applied at

higher frequency where firm accounting variables are held fixed. The interaction structure extends naturally to financial leverage, firm size, foreign revenue share and hedging activity. Finally, the 2026 Hormuz crisis offers a natural out-ofsample test of whether firms the model identifies as carrying high predicted tail risk were in fact among the worst performers that March.

1.8. Systematic of the Thesis

Chapter 1 establishes expected shortfall as the risk measure of interest, explains why Middle East tension is treated as a persistent structural condition, introduces GEAR and EXPO, and states the problem identification, questions, objectives, scope, limitations and significance. The 2026 Hormuz crisis is presented as motivation rather than evidence.

Chapter 2 reviews prior empirical work on geopolitical risk, oil shocks, cost structure and tail risk measurement, sets out the theoretical framework with arbitrage pricing theory central, consolidates the research gap, and develops the hypotheses.

Chapter 3 describes the research design, population and sample, variable operationalization with numbered equations and construct-validity discussion for GEAR and EXPO, the panel specification, model selection and classical assumption tests, and robustness checks.

Chapter 4 presents descriptive statistics, correlation and multicollinearity diagnostics, model selection results, regression output and discussion, including interpretation of coefficient signs running counter to initial expectations.

Chapter 5 summarises the findings, sets out theoretical and practical implications, revisits the 2026 Hormuz crisis as an out-of-sample reference, and identifies limitations and directions for future research.

CHAPTER 2

LITERATURE REVIEW

2.1. Previous Research

This section reviews the empirical work on which the study builds, organised into five areas: geopolitical risk and equity markets, oil price shocks and emerging market equities, cost structure and equity risk, profitability and firm resilience, and expected shortfall as a firm-level outcome. Each subsection reports what the studies examined and what they found. Synthesis, identification of what remains unaddressed, and the statement of what this thesis adds are reserved for the research gap subsection, so that the account of existing work stays separate from the defence of the model proposed here.

Table 2.1. Summary of Previous Related Research

No.	Author(s)	Year	Journal	Country / Sample	Method	Key Findings
1	Caldara & Iacoviello	2022	American Economic Review	Global, 120 years	VAR	GPR shocks foreshadow lower investment and stock prices; higher disaster probability and larger downside risks; capital shifts from emerging to advanced economies
2	Bouras, Christou, Gupta & Suleman	2019	Emerging Markets Finance & Trade	18 EMs incl. Indonesia	Panel GARCH	Global GPR dominates country-specific GPR in driving EM stock volatility
3	Zaremba, Cakici, Demir & Long	2022	Journal of Financial Stability	48 countries	Fama-MacBeth	Stocks with lowest GPR beta earn 3.9% higher annualised returns; attributed to overreaction via the availability heuristic
4	Ndako, Salisu & Ogunsiji	2021	Asian Economics Letters	Indonesia, Malaysia	GARCHMIDAS	Global and country GPR raise volatility of Indonesian Islamic stocks; effect larger for Indonesia
5	Hoque & Zaidi	2020	Borsa Istanbul Review	Fragile Five incl. Indonesia	Markovswitching	GPR effects are nonlinear and regime-dependent; Indonesia shows lagged negative effects in high-volatility regimes

6	Kilian & Park	2009	International Economic Review	United States	Structural VAR	Supply and demand oil shocks have opposite equity effects; Middle East supply disruptions depress returns
7	Mork	1989	Journal of Political Economy	United States	Asymmetric OLS	Oil price increases are significantly related to output while decreases are not; symmetric specifications are misspecified
8	Hamilton	2003	Journal of Econometrics	United States	Nonlinear NOPI	The oil-macroeconomy relation is nonlinear and asymmetric; a net oil price increase measure outperforms the linear change specification
9	Basher & Sadorsky	2006	Global Finance Journal	21 emerging markets	Multi-factor model	Oil price risk is a significant pricing factor in EM equity returns
10	Sukarno & Setiawati	2020	Int. J. Energy Econ. & Policy	Indonesia, 10 sectors	VAR, IRF	Mining responds positively to oil price increases; all other IDX sectors respond negatively
11	Wahyudi & Rahmawati	2025	Cogent Economics & Finance	Indonesia, 11 sectors	VAR, T-GARCH	Heterogeneous sectoral responses to oil volatility; calls for extension to geopolitical tension
12	Lev	1974	J. Financial & Quant. Analysis	United States	Crosssectional	Operating leverage is positively associated with earnings variability and systematic risk
13	Mandelker & Rhee	1984	J. Financial & Quant. Analysis	United States	Crosssectional	Degree of operating and financial leverage are both positively associated with systematic risk of common stocks
14	Novy-Marx	2011	Review of Finance	United States	Panel, FF factors	Operating leverage is a priced risk factor in the cross-section of returns
15	Zhu & Zang	2025	Economic Modelling	China	Panel	Firms with high operating leverage exhibit significantly higher stock price crash risk
16	Novy-Marx	2013	J. Financial Economics	United States	Panel, FF factors	Gross profitability is among the strongest fundamental predictors of returns
17	Fama & French	2015	J. Financial Economics	United States	Five-factor model	Profitability (RMW) is a distinct pricing factor alongside size, value and investment

18	Artzner, Delbaen, Eber & Heath	1999	Mathematical Finance	Theoretical	Axiomatic	Defines coherent risk measures; ES is subadditive while VaR is not
19	Acerbi & Tasche	2002	Economic Notes	Theoretical	Analytical	Establishes ES as a coherent alternative to VaR and resolves estimation issues for discrete distributions
20	Dwitanto, Manurung, Machdar & Hidayat	2026	Architectural Image Studies	Indonesia, 19 property firms, 2013-2023, 209 obs	Panel RE (GLS), mediation	Controlling ownership is negatively associated with ES5 and with landbanking; landbanking is negatively associated with ES5, giving partial mediation. Leverage negative, exchange rate positive; liquidity and ROA not significant
21	Baur & Lucey	2010	Financial Review	US, UK, Germany	GARCH	Gold acts as a safe haven for equities only during extreme downturns, and temporarily

Source: Author (2026)

Research on geopolitical risk and financial markets expanded rapidly after Caldara and Iacoviello (2022) published the GPR Index, which gave researchers a standardised high-frequency measure suited to quantitative models. Earlier work relied on event dummies for wars or terrorist attacks, agency political risk ratings, or defence spending. Each carried a serious drawback: dummies are binary and cannot register intensity, agency ratings update only monthly or quarterly and blend several risk dimensions into one score, and defence spending is a policy outcome rather than a measure of risk. The GPR Index supplies a continuous daily series from a transparent, replicable text-analysis procedure. Developed by economists at the Federal Reserve Board, it has since been used by the European Central Bank in assessing downside risks to euro area growth and by the International Monetary Fund in the World Economic Outlook, and it has become the standard measure across work on equity markets, commodities, sovereign spreads, foreign direct investment and corporate investment.

2.1.1. Geopolitical Risk and Equity Markets

Caldara and Iacoviello (2022) applied their index in a vector autoregression covering roughly 120 years. They report three results. Higher geopolitical risk

comes before lower investment and lower stock prices. It goes with a higher probability of disaster and larger downside risks. And during these periods capital moves away from emerging economies toward advanced ones.

Bouras, Christou, Gupta and Suleman (2019) estimated a panel GARCH model across 18 emerging markets, Indonesia among them. Global geopolitical risk turned out to matter more than the country-specific component. Once the global measure is in the model, country-specific risk has only a limited direct effect.

Zaremba, Cakici, Demir and Long (2022) ran Fama-MacBeth regressions across 48 countries. Stocks with the lowest GPR beta earned 3.9 percent higher annualised returns than those with the highest. They explain this as overreaction working through the availability heuristic: vivid headlines lead investors to put too much weight on bad outcomes.

Two studies look at Indonesia directly. Ndako, Salisu and Ogunsiji (2021) used GARCH-MIDAS and found that both global and country-level geopolitical risk raise the return volatility of Indonesian Islamic stocks. The effect was larger for Indonesia than for Malaysia. Hoque and Zaidi (2020) applied a Markov-switching model to the Fragile Five economies. They found the effects nonlinear and dependent on regime, with Indonesia showing lagged negative effects mainly in high-volatility states.

2.1.2. Oil Price Shocks and Emerging Market Equities

Kilian and Park (2009) used a structural VAR to split oil price movements into three components: supply-driven, aggregate-demand-driven and precautionary demand-driven. Their main finding is that the source of a price change matters more than its size. Increases caused by supply disruption, including Middle East conflict, push equity returns down. Increases caused by stronger global demand are neutral or even favourable. A twenty-dollar rise from Chinese demand growth and a twenty-dollar rise from a Gulf supply interruption are two different events that happen to produce the same headline number.

The asymmetry of oil price effects was documented earlier and separately. Mork (1989) split oil price changes into positive and negative components and estimated a coefficient on each. Increases were significantly related to output. Decreases were not. He concluded that specifications imposing symmetry are misspecified. Hamilton (1996, 2003) took this further with the net oil price increase measure. It equals zero unless the current price exceeds its maximum over a preceding window, and otherwise records only the excess. This censored version outperforms the linear percentage change when forecasting real activity.

Basher and Sadorsky (2006) estimated multi-factor models across 21 emerging markets and confirmed oil price risk as a statistically significant pricing factor, distinct from conventional market and size factors.

For Indonesia, Sukarno and Setiawati (2020) used a VAR with impulse response analysis across ten IDX sectors. Mining responds positively to oil price increases. Every other sector responds negatively. Wahyudi and Rahmawati (2025) reached similar conclusions using a VAR and threshold GARCH across eleven sectors. They document responses that differ by sector and close with an explicit call to extend the analysis to geopolitical tension.

2.1.3. Cost Structure, Operating Leverage, and Equity Risk

Lev (1974) established the foundational result. Operating leverage is positively associated with earnings variability and with beta. The mechanism is simple. A firm carrying a large fixed cost base cannot cut costs proportionately when sales fall, so a percentage change in revenue turns into a bigger percentage change in operating income. Mandelker and Rhee (1984) put operating and financial leverage into the same test. Both were positively and significantly associated with systematic risk, and between them they accounted for a large share of the crosssectional variation in beta. Novy-Marx (2011) took the question back into the Fama-French framework and found operating leverage priced in the cross-section rather than merely correlated with beta, with higher-leverage firms earning higher average returns. Recent work has moved the outcome variable off beta and onto the left tail. Zhu and Zang (2025), working on China, define crash risk by the probability and

severity of extreme negative returns and find high operating leverage firms carrying significantly more of it. The Lev tradition thus arrives at an asymmetric outcome rather than a symmetric second-moment one.

2.1.4. Profitability and Firm Resilience

Novy-Marx (2013) measured gross profitability as revenue less cost of goods sold over total assets and found it among the strongest fundamental predictors of returns, performing at least as well as book-to-market. He reads it as capturing a quality dimension in productive assets that conventional value measures miss.

Fama and French (2015) brought profitability into their five-factor model as RMW and showed it carries explanatory power of its own for average returns. Financial slack is the common mechanism. Firms earning strong returns on assets generate more internal cash flow, carry stronger balance sheets and fund more of their own activity. A very profitable firm can take several periods of high input costs and still make money, while a firm running near breakeven can slip into loss, breach covenants and end up looking for distress financing.

2.1.5. Expected Shortfall as a Firm-Level Outcome Variable

Artzner, Delbaen, Eber and Heath (1999) set out what a coherent risk measure has to satisfy, namely monotonicity, translation invariance, positive homogeneity and subadditivity, and showed that expected shortfall meets all four where value at risk fails subadditivity. The practical cost of that failure is that VaR can make a diversified portfolio look riskier than the sum of its positions. Acerbi and Tasche (2002) then established expected shortfall as the natural coherent alternative and cleared up the estimation ambiguities that arise when the return distribution is discrete, which is the case that matters for a measure computed from a finite sample of daily returns. The Basel Committee took the same view in January 2019.

Alongside this, a separate literature established that downside risk is priced apart from total risk. The findings of Ang, Chen and Xing (2006) and of Bali, Demirtas and Levy (2009) are set out in Section 2.2.1; what matters for the present purpose

is that firms with high downside beta record larger expected shortfall values, so a factor priced in downside beta should be visible in expected shortfall.

The study most directly comparable in design and setting is Dwitanto, Manurung, Machdar and Hidayat (2026). Working with 19 property developers listed on the IDX over 2013 to 2023, a balanced panel of 209 firm-year observations, they examined whether ownership concentration and landbanking intensity relate to extreme downside risk. Their dependent variable is expected shortfall at the five percent tail, constructed by taking the fifth percentile of the annual distribution of daily returns as the VaR threshold and averaging the returns below it, and estimation is by panel generalised least squares with a random effects specification.

Controlling ownership is negatively and significantly associated with expected shortfall and negatively associated with landbanking intensity, while landbanking is itself negatively associated with tail risk. Mediation analysis indicates partial mediation, so controlling shareholders appear linked to lower tail risk both directly and through more disciplined land accumulation. Leverage is negatively and significantly associated with expected shortfall, the exchange rate positively so, and neither liquidity nor return on assets reaches significance in the tail-risk model. Two features of that design bear on this one. They treat landbanking as ambiguous in sign, since a real options reading makes land reserves a form of flexibility while corporate finance and agency readings make the same reserves capital rigidity, and the net direction is left as an empirical question. And they combine firm-level characteristics with a macroeconomic variable in a single firm-level tail-risk regression, which is the structural precedent this study extends by making the macroeconomic side the object of interest rather than a control.

2.1.6. Research Gap and Value Added to Previous Studies

Each of these five strands is well developed on its own. Together they do not connect. The geopolitical literature shows that GPR moves equity outcomes, but it works at index or country level, so its emerging market evidence describes aggregate volatility and not particular firms. The oil literature shows that the source and direction of a price change matter, and it offers a mature way of handling

asymmetry, but in Indonesia it has been applied at sector level. The cost structure literature links operating leverage to beta and to crash risk, but it does so unconditionally, without identifying the shock. The profitability literature explains why some firms absorb shocks better, though nobody has joined it to a geopolitical setting. Expected shortfall has now been used on Indonesian firm-level data by Dwitanto et al. (2026), but with governance determinants and not macroeconomic ones.

Four gaps follow.

1. Indonesian evidence on geopolitical risk is aggregate. Ndako et al. (2021) and Hoque and Zaidi (2020) both work with index or portfolio volatility, so neither can say which firms carry the greatest exposure or why. Zaremba et al. (2022) do show that firms differ in their GPR loading, but they compare across countries, not within one. They also do not link that loading to any firm characteristic an investor could observe.
2. No existing measure weights a global oil price movement by a country's own import dependence on the region where the tension originates. Studies use one of two things. Either a raw oil price series, which is the same for every country and carries no geopolitical content, or a raw risk index, which carries no energy content and no national weighting.
3. Firm-level expected shortfall has been modelled in Indonesia, but not against geopolitical or energy determinants. Dwitanto et al. (2026) show that ES5 works as a dependent variable on IDX panel data, and that a macroeconomic regressor can sit alongside firm characteristics. Their explanatory variables, however, are governance and asset-strategy measures.
4. Nobody has estimated the cost structure relationship conditionally, under an identified energy shock. Lev (1974), Mandelker and Rhee (1984) and NovyMarx (2011) all relate cost structure to risk unconditionally. Zhu and Zang (2025) relate it to crash risk without specifying what produces the crash.

What this study adds is best described as a contribution to measurement. It is worth saying plainly what is not being claimed. Expected shortfall is not a new risk

measure, and it is not new to Indonesian firm-level research. Geopolitical risk is not an unstudied determinant of equity outcomes. Panel generalised least squares is standard practice. The claim rests on two variables.

GEAR is offered as a measure of geopolitical energy exposure, weighted by a country's own import dependence on the region where the tension originates. Its two components are established. The combination is not. Whether it is useful is an empirical question, and estimating it jointly with GPR is designed to answer that. If GEAR explains something the general index does not, then the weighting is doing work that neither an unweighted oil series nor an unweighted risk index can do.

EXPO is offered as a firm-level determinant of tail risk. The link between cost structure and systematic risk has been established for decades. In that tradition, though, the outcome variable is beta or crash risk, and the measure is usually the degree of operating leverage, not a direct cost ratio. Relating cost of goods sold to net sales against a coherent tail measure at firm level appears not to have been attempted, and certainly not in Indonesia. The variable also carries an unresolved question about the sign it should take, which this design can address.

The interaction stays in the model as a test of whether cost structure conditions how the shock is transmitted. It is not a principal contribution. The claim does not depend on the moderation reaching significance, and treating it as central would tie the value of the whole study to one coefficient. What is offered is the pair of constructed variables, tested on 594 firm-year observations of non-financial Kompas 100 constituents.

2.2. Theoretical Framework

The framework below is deliberately narrow, developing only those concepts that do work in specifying the model or generating the hypotheses. The first subsection follows risk theory from Knight through Markowitz, Sharpe and Jorion to expected shortfall, ending with arbitrage pricing theory, which is the framework behind the empirical model and which grew directly out of the limitations of the single-factor

approach before it. Each remaining subsection handles one component of the model: the step from volatility evidence to a tail outcome, the cost channel behind GEAR and its asymmetry, the construct behind EXPO, the sentiment channel behind GPR, and the case for including ROA as a control.

2.2.1. Theory of Risk

Over the past century risk has moved from a loose concept toward something practical and quantifiable, particularly as applied in corporate financial management and investment decision-making. The word itself comes from the Italian *risicare*, to dare, which already implies that risk is not something that simply befalls people but something bound up with deliberate choice under uncertainty. Choice, expectation and consequence have been attached to the idea from the start.

Knight (1921) drew the line the field still works with. Where probabilities can be estimated, expected values computed and decisions optimised against them, there is risk. Where there is no data, no distribution and nothing to estimate from, there is uncertainty, and it yields to judgement rather than to calculation. Geopolitical conflict falls awkwardly across that line, which is why the distinction is worth keeping. Any single escalation is close to Knightian uncertainty. The level of tension sustained across a year is not, and measuring it is what the GPR Index sets out to do.

Markowitz (1952) made risk calculable by defining it as the variability of returns and measuring it with variance and standard deviation. In the mean-variance model the investor minimises risk for a given return or maximises return for a given risk, and assets whose returns correlate weakly lower the risk of the portfolio holding them. Asset allocation theory, multi-factor models and risk-adjusted performance evaluation all rest on that definition.

The objection that matters here is the one Markowitz raised against his own measure. Variance counts a return five percent above the mean exactly as it counts one five percent below, and no investor holding the position treats those two days as equivalent. His later answer was semi-variance, which counts only the deviations falling below the mean or below a chosen target (Markowitz, 1991).

That is the first version of downside risk, and it is the ancestor of everything this thesis measures. Expected shortfall is the same instinct carried further and made coherent, and the distance between the two ideas is most of the reason a separate literature on tail measurement exists at all.

Sharpe (1964) went in a different direction, toward the portion of risk that diversification cannot remove, captured in beta. CAPM matters here for one reason: it is the first account of how a shock reaching the whole market reaches an individual firm, which is the mechanism this study is asking about when it puts a national energy variable against firm-level tail risk. Its assumptions do not survive contact with a non-normal left tail, and that problem is taken up below.

Later work took beta apart instead. Ang, Chen and Xing (2006) separated its upside and downside components and found that stocks loading heavily on the downside earn returns conventional beta fails to predict, which is direct evidence that the two directions are priced differently. Bali, Demirtas and Levy (2009) went further and tested downside beta against expected shortfall itself, finding the two closely linked in the cross-section, with high downside beta firms recording the larger expected shortfall values. Symmetric measures throw away information that investors act on.

Jorion (2007) brought the quantile-based approach together through value at risk and expected shortfall. Risk here is defined by the probability and size of losses in the far left of the return distribution. Volatility on its own cannot capture rare events with large consequences, so measurement has to focus on the tail itself. Value at risk reports a loss threshold that will not be exceeded at a stated confidence level. Expected shortfall goes past that threshold and reports the average of what happens beyond it. The difference is not presentational. Because expected shortfall satisfies subadditivity, combining two positions can never make measured risk larger than the sum of the parts, while value at risk can, and a diversified book can be made to look more dangerous than its own components. That defect is what led the Basel Committee to replace value at risk with expected shortfall for market risk capital

under the Fundamental Review of the Trading Book (Bank for International Settlements, 2019).

Expected shortfall at the fifth percentile is accordingly the dependent variable throughout this thesis, and the choice is the end point of the evolution traced above, from total variance through semi-variance and downside risk to the coherent measures international regulators now endorse. It reports the depth of the tail and not only its edge. It satisfies the coherence axioms of Artzner, Delbaen, Eber and Heath (1999), which is what makes it internally consistent when aggregated across firms and portfolios. And it is the measure banking supervisors actually use, which gives any result obtained here a practical reference point outside the academic literature.

Deciding what to measure does not decide what explains it. Sharpe's model offers one answer, and it does not survive contact with the present problem. CAPM works with a single factor, the return on the market portfolio, and it reaches that result by assuming either quadratic utility or normally distributed returns. Neither assumption sits comfortably in a study whose entire subject is the non-normal left tail. Ross (1976) supplies the alternative. Under Arbitrage Pricing Theory, returns are generated by a linear combination of exposures to several systematic factors plus an idiosyncratic term, and absent arbitrage the expected return is a linear function of those exposures. What makes it the right framework here is not what it specifies but what it leaves open: it names no factors at all, requiring only that a factor be systematic, in the sense of touching many assets rather than one, and that some economic reasoning stand behind it.

Geopolitical risk and energy exposure clear both requirements. They can therefore enter as factors in their own right, rather than being folded into an undifferentiated market return where whatever they do to Indonesian firms disappears into the index. The second permission matters more. Because APT allows two assets to load differently on the same factor, a single national energy shock can reach a diesel-burning mining contractor and a tower operator drawing grid electricity with quite different force, and nothing in the theory objects, which is the formal licence for

the entire firm-level half of this study, since without it the only defensible model would treat every firm in the panel as equally exposed to a shock that plainly does not reach them equally. Macro factors above, firm characteristics below. Both halves sit inside the same framework.

Mapping it onto the estimating equation is direct. GPR and GEAR are the macrolevel factors, the first for broad geopolitical uncertainty and the second for energy cost pressure weighted by national import dependence, while EXPO sets the loading: a firm's ratio of direct costs to revenue describes how much of its income statement stands exposed to input price movements. The product $\text{GEAR} \times \text{EXPO}$ is a conditional loading. It lets the sensitivity of tail risk to the energy factor vary with cost structure instead of holding it fixed across the panel, which is the whole content of the fourth hypothesis. ROA is in the equation to absorb profitability differences that would otherwise land in the error term.

One departure from APT as Ross wrote it should not pass unnoticed. His theory concerns expected returns; the dependent variable here measures tail severity.

What is estimated is more accurately described as a multi-factor exposure model in the APT tradition, one whose coefficients map systematic factors into the severity of firms' downside outcomes and not into their mean returns. Ang, Chen and Xing (2006) and Bali, Demirtas and Levy (2009) do the same thing, applying factor logic to downside risk measures, and Bali et al. establish the empirical link between downside factor loadings and expected shortfall directly. All of these studies assume that markets price systematic exposures efficiently enough that cross-sectional differences in risk reflect real differences in exposure and not persistent mispricing. This study makes the same assumption and does not test it.

2.2.2. From Volatility to Tail Risk

A conceptual step needs stating, because most of the geopolitical evidence reviewed above concerns volatility whereas the outcome here is a tail measure. Bouras et al. (2019) estimate panel GARCH, Ndako et al. (2021) GARCHMIDAS, and Hoque and Zaidi (2020) a Markov-switching volatility model. All three establish that geopolitical risk associates with movements in the conditional second

moment. None establishes association with the severity of the left tail, and the second does not follow automatically from the first.

Variance is a whole-distribution statistic and expected shortfall is not. An increase in conditional variance is consistent with several different changes in shape: both tails may widen, the left tail alone may fatten, or the right tail may fatten while the left is unchanged, which would raise variance without raising expected shortfall at all. Expected shortfall responds only to realisations below the fifth percentile and is insensitive to everything above it. A documented volatility relationship therefore constrains but does not determine the corresponding tail relationship, and certainly does not fix its sign.

Three considerations support the expectation that the link holds here. Caldara and Iacoviello (2022) report not merely a variance result but that geopolitical risk associates with a higher probability of disaster and with larger downside risks specifically, which is a statement about the left tail. Bali et al. (2009) demonstrate an empirical link between downside factor loadings and expected shortfall, so a factor priced in downside beta should be visible in expected shortfall. And Zhu and Zang (2025) find operating leverage predicts crash risk, an explicitly left-tail outcome, suggesting the firm-characteristic channel also carries through.

Two considerations point the other way and explain why the hypotheses are nondirectional. Geopolitical escalation raises commodity prices, and for a commodity producer that is a favourable realisation; in an index with substantial commodity weight, tension could widen the right tail more than the left, raising volatility while leaving expected shortfall unchanged or reduced. Further, the macro variables here are annual averages measuring sustained tension rather than event timing, and there is no reason a sustained-tension relationship with the tail must resemble an eventwindow relationship with volatility.

2.2.3. The Oil Cost Channel and Asymmetry of Oil Shocks

For a net oil importer, a price increase is an adverse terms-of-trade movement. The country pays more for the same physical volume, the import bill rises, the current account deteriorates, and the currency typically weakens, raising the domestic cost

of all imported inputs rather than energy alone. At firm level this arrives through cost of goods sold: fuel, freight, petrochemical feedstock, packaging, and electricity where generation is fossil-fuelled. Firms whose direct costs absorb a large share of revenue take more of the increase onto their income statement, which is the mechanism connecting the macro shock to the cross-section (Kilian and Park, 2009).

One qualification matters for the Indonesian setting. Commodity and trading firms typically carry very high cost ratios, but their revenues also move with commodity prices, so the same shock arrives on both sides of the income statement. That natural hedge is unavailable to firms with fixed output prices, and it is one reason the direction of the cost structure relationship is treated as open.

GEAR retains only positive oil movements, using a $\text{Max}(0, \cdot)$ operator so falling years enter as zero. This is substantive and requires defence beyond the observation that increases are the interesting case. Four arguments support it.

1. The generating processes differ. In the Kilian and Park (2009) decomposition, large increases and large declines are typically produced by different disturbances. Within the observation period the sharp increases trace to supply and precautionary-demand shocks such as renewed sanctions on Iranian exports and disruption of Red Sea shipping, whereas the large declines trace to supply glut and collapsing demand. A symmetric variable would record the 2020 demand collapse as a large negative geopolitical energy shock, inverting the construct rather than extending it.
2. The econometric precedent is established. Mork (1989) split oil changes into increases and decreases and found only increases carried significant explanatory power for output. Hamilton (1996, 2003) formalised this into the net oil price increase measure, zero unless the price exceeds its recent maximum and otherwise recording only the excess. The $\text{Max}(0, \cdot)$ operator applies the same censoring logic at annual frequency, weighted by import share.

3. Cost pass-through is itself asymmetric. Work on retail fuel and downstream input pricing documents the pattern labelled rockets and feathers, in which costs pass through faster and more completely when the underlying commodity rises than when it falls (Bacon, 1991; Borenstein, Cameron and Gilbert, 1997). If transmission from crude to firm costs is asymmetric, a regressor imposing symmetry is not merely wasteful but structurally wrong.
4. The dependent variable is asymmetric by construction. Expected shortfall measures adverse outcomes only. Pairing a one-sided outcome with a twosided regressor would ask the model to treat energy relief and energy stress as points on a single linear scale mapping onto a variable registering only stress.

The cost is real. Censoring discards the variation in declining years, reducing the effective range of the regressor and preventing any test of whether declines carry a symmetric relationship of opposite sign. A symmetric alternative built on the absolute change is estimated alongside for this reason.

2.2.4. Cost Intensity and Operating Leverage: What EXPO Measures

The literature invoked for the cost structure variable is the operating leverage literature, but the variable constructed is the ratio of cost of goods sold to net sales. The two are discussed together often enough to be treated as the same thing, and the difference decides what sign this study should expect.

Operating leverage, in the sense Lev (1974) and Mandelker and Rhee (1984) give it, is a statement about how costs behave when volume changes, captured by the fixed-to-variable ratio or by the elasticity of operating income with respect to sales. A firm carrying heavy fixed costs cannot shed them when output falls, so a given percentage decline in revenue produces a larger percentage decline in operating income. Earnings swing harder than sales. Systematic risk rises with them, and the left tail deepens.

Cost of goods sold to net sales is a statement about magnitude. It reports how much of each unit of revenue direct production cost consumes, and therefore how thin

the gross margin is, and says nothing at all about whether those costs would fall if output did. What connects the two is an accounting approximation: cost of goods sold is predominantly variable, selling and administrative expense predominantly fixed. Under that mapping a high ratio describes a cost base weighted toward variable cost, which is to say a firm with low operating leverage, so Lev and Mandelker-Rhee, read through the accounting rather than through the ratio's surface, predict a negative relationship with tail risk. That is the opposite of what the ratio looks like on first inspection, and it is why Section 2.3 states H3 without imposing a direction.

The approximation holds unevenly across industries. Cost of goods sold routinely carries fixed manufacturing overhead inside it, depreciation on production assets, plant supervision and allocated facility costs among them, and how much varies with how capital-intensive the production process is. A cement producer running kilns books a substantial fixed component within its cost of sales. A fuel distributor books almost none. Both can report similar ratios while sitting at opposite ends of the operating leverage distribution the ratio is standing in for, so the measurement error in EXPO is not random noise but error correlated with industry, and Section 3.3.4.1 returns to what that costs the interpretation.

Because the two readings are opposed rather than differing in emphasis, H3 is stated without directional restriction. Dwitanto et al. (2026) treat landbanking intensity the same way and on the same reasoning.

A direct measure of the degree of operating leverage was considered and rejected on three grounds. Estimating it from a time-series regression of operating income on sales requires a long stable series per firm, which eleven annual observations do not provide and which the 2020 structural break would distort. Constructing it from the accounting fixed-variable split is impossible because Indonesian issuers need not disclose that split. And any assumed allocation would import an assumption at least as strong as the one it replaced.

The interaction offers partial resolution. Suppose EXPO measures margin thinness. Then a firm consuming most of its revenue in direct costs enters a year of elevated

energy prices with less room to absorb them, its earnings deteriorate faster than a competitor's, and the interaction should come out positive. On the cost-flexibility reading the same firm is no worse off and possibly better, so the interaction should sit at zero or turn negative. The product therefore does two jobs at once. It tests moderation in the ordinary sense, and it tells us which of the two readings the data are operating under.

2.2.5. The Sentiment Channel

GPR reaches firms by a route GEAR does not, which is why the model carries both. Tension raises uncertainty about future cash flows and about the rates at which they are discounted, and investors respond by reducing exposure to assets they read as risky. Baur and Lucey (2010) document the behaviour in gold, which acts as a safe haven for equities during extreme downturns and not the rest of the time; Caldara and Iacoviello (2022) find the country-level version, with capital moving out of emerging markets toward advanced ones. Portfolio reallocation is not input cost pass-through. It reaches a tower operator and a petrochemical producer alike, whatever sits in either firm's cost base.

Indonesia is more exposed to this than most. Hoque and Zaidi (2020) count it among the Fragile Five, a group defined by current account deficits and dependence on foreign portfolio inflows, and therefore by vulnerability to sudden reversals when global risk appetite turns. Their finding that the effects are regimedependent, showing up mainly in high-volatility states, fits a channel that switches on under stress rather than running continuously. Two channels reaching firms by different routes is the reason GPR and GEAR enter as separate regressors rather than one.

2.2.6. Profitability as a Risk Buffer

Return on assets stands in for the financial slack mechanism of Novy-Marx (2013) and Fama and French (2015). A firm earning strong returns on its asset base holds more internal cash flow, funds more of its own activity and can carry elevated input costs for longer before profitability turns and financing tightens, so higher

profitability should come with a shallower left tail. It enters the model as a control, and Section 3.3.5 sets out what that does and does not imply.

2.3. Hypothesis Development

Four hypotheses follow, corresponding one to one with the research questions in Chapter 1. All are stated non-directionally, since for each variable the discussion above identifies at least two channels operating with opposite signs and offers no basis for asserting which dominates.

The hypotheses are phrased in terms of significant association rather than significant effect. The design is observational and contains no mechanism for isolating exogenous variation, so it can establish whether variables move together conditional on the other regressors but cannot establish that one produces the other.

Table 2.2. Research Hypotheses

Hypothesis	Path	Statement	Theoretical Basis
H1	GPR → ES5	Geopolitical risk is significantly associated with expected shortfall	Ross (1976); Caldara & Iacoviello (2022); Bouras et al. (2019); Baur & Lucey (2010); sentiment and flight-to-quality channel
H2	GEAR → ES5	Geopolitical energy-adjusted risk is significantly associated with expected shortfall	Ross (1976); Kilian & Park (2009); Mork (1989); Hamilton (1996, 2003); Basher & Sadorsky (2006); oil cost channel
H3	EXPO → ES5	Firm cost exposure, measured as COGS to net sales, is significantly associated with expected shortfall	Lev (1974); Mandelker & Rhee (1984); Novy-Marx (2011); Zhu & Zang (2025); cost intensity versus cost flexibility
H4	GEAR × EXPO → ES5	The interaction between GEAR and firm cost exposure is significantly associated with expected shortfall	Ross (1976) conditional factor loading; cost amplification versus commodity pass-through
Control	ROA → ES5	Return on assets is included as a control; no separate hypothesis is formulated	Novy-Marx (2013); Fama & French (2015); Dwitanto et al. (2026)

Note: the dependent variable in every hypothesis is expected shortfall, whose selection rests on the risk measurement literature developed above and is therefore not repeated in the column. Source: Author (2026)

H1: Geopolitical risk is significantly associated with expected shortfall.

Elevated tension raises uncertainty and prompts global investors to reduce emerging market exposure, deepening the left tail. Against this, Zaremba et al. (2022) find evidence of overreaction to geopolitical headlines, which at annual

frequency could be partly reversed within the same observation year. Indonesia's position as a commodity producer provides a further offsetting mechanism, since escalation raises the prices of what much of the sample sells. The net direction is left open.

H2: Geopolitical energy-adjusted risk is significantly associated with expected shortfall. An increase in world oil prices weighted by Middle East import dependence transmits cost pressure through cost of goods sold and through a weaker currency (Kilian and Park, 2009). The offsetting consideration is that much of the sample consists of commodity producers and energy-linked firms for whom a price increase is revenue rather than cost. Which channel dominates at sample level is what the coefficient is being asked to reveal.

H3: Firm cost exposure is significantly associated with expected shortfall. Under the margin reading, cost-intensive firms carry thinner cushions and should exhibit deeper tails (Lev, 1974). Under the flexibility reading, the same firms hold cost bases that scale with output and should exhibit shallower tails. The two predict opposite signs, which is why no direction is imposed.

H4: Firm cost exposure moderates the effect of Geopolitical energy-adjusted risk on expected shortfall. A positive coefficient would indicate cost-intensive firms are disproportionately exposed in years of elevated energy cost pressure, supporting the margin reading. A negative coefficient would indicate partial shielding, supporting the flexibility reading or commodity pass-through. The interaction therefore carries construct-validity information alongside testing moderation.

Stating all hypotheses non-directionally rests on three grounds. Theoretically, each variable has at least two channels running in opposite directions, and imposing a direction would assert a dominance the theory does not deliver. Methodologically, a one-sided restriction risks a Type II error in which a genuine relationship running counter to expectation is recorded as no relationship. And there is precedent in closely comparable work, since Dwitanto et al. (2026) treat landbanking intensity the same way in an Indonesian firm-level expected shortfall setting.


A final word on why both GPR and GEAR appear. GPR registers tension wherever it occurs and through whatever route it reaches asset prices, including routes unrelated to energy; a crisis that spikes uncertainty without moving oil prices raises GPR while leaving GEAR largely unaffected. GEAR registers energy cost pressure weighted by import dependence and is silent about developments that do not reach the oil market. The two are correlated, since many events do both, but well short of unity. Including both allows the sentiment and fundamental cost channels to be estimated separately rather than confounded in a single regressor.

2.4. Research Framework

The research framework comprises three independent variables, GPR, GEAR and EXPO, together with one interaction term and one control variable, each examined for its association with expected shortfall at the fifth percentile. GPR and GEAR enter as macro-level factors common to all firms within a given year, while EXPO is a firm-level characteristic that both relates directly to tail risk and moderates the path running from GEAR to the dependent variable. The full estimating equation, with each term defined and operationalised, is presented in Chapter 3.

Figure 2.1 presents the framework diagrammatically. GPR and GEAR enter as macro-level factors common to all firms within a year, and both are hypothesised to relate directly to expected shortfall. EXPO occupies the centre of the diagram in two capacities: it relates directly to expected shortfall in its own right, and it moderates the path running from GEAR to the dependent variable, which is the conditional factor loading described above. ROA does not appear in the diagram. It enters the estimating equation as a control variable, absorbing cross-firm variation in profitability, but no hypothesis is attached to it and it forms no part of the conceptual model.

Figure 2.1. Research Framework



Source: Author (2026)

CHAPTER 3

RESEARCH METHODOLOGY

3.1. Research Design

This study applies a quantitative, explanatory design to secondary panel data (Sugiyono, 2017). Estimation is by Panel Estimated Generalized Least Squares with cross-section fixed effects and cross-section weights, with White diagonal standard errors correcting residual heteroskedasticity (White, 1980). The specification follows from the sequential selection procedure described later in this chapter; because that procedure did not resolve cleanly, random effects estimates are reported alongside the fixed effects results rather than suppressed. All estimation is carried out in EViews 13.

What the design can and cannot establish should be stated at the outset. Panel regression on observational data can determine whether variables move together once other measured factors are held constant, and can quantify the strength and direction of that co-movement. It cannot isolate causal effects. There is no randomisation, instrument, discontinuity or natural experiment in the specification, so nothing rules out an unmeasured factor driving both the geopolitical variables and equity tail risk. Every result is a conditional association, and the language throughout reflects this. Where the words effect or impact appear in the names of standard tests, such as the Fixed Effect Model, they carry conventional econometric meaning rather than causal claims.

The design was chosen over the alternatives for three reasons. The research questions concern the strength, direction and significance of relationships between specified variables, which calls for regression rather than descriptive or exploratory analysis. The data come from Bloomberg Terminal, audited financial statements and an established academic index, all objective and replicable, avoiding the response bias of surveys or interviews. And the panel structure exploits variation in two dimensions: 54 firms over 11 years give 594 observations for six parameters plus two variance components, a ratio comfortably exceeding the thresholds

recommended by Hair et al. (2014), though the qualification below regarding the macro coefficients should be read alongside it.

Panel EGLS is preferred to ordinary least squares because OLS ignores the panel structure and treats all 594 observations as independent, which is untenable when eleven observations belong to one firm and share unobserved characteristics; the resulting standard errors are understated. Both fixed and random effects address this, the first by giving each firm its own intercept and the second by modelling the firm component as a random draw. The EGLS estimator additionally folds an estimated variance structure into the weighting matrix. Here that structure is crosssectional: each firm is weighted by the inverse of its estimated residual variance, so inherently volatile firms do not dominate the fit. This is the standard remedy for cross-sectional heteroskedasticity in EViews, and it explains why the reported coefficients differ from unweighted counterparts, a sensitivity examined among the robustness checks.

3.1.1. Choice of Annual Frequency

GPR and crude prices are published daily, so the obvious objection is that annual aggregation discards the short-horizon shock information a study of geopolitical risk should want. Three responses apply.

First, on what is being measured. The study does not estimate the market reaction to a discrete event but the relationship between sustained tension over a period and the severity of firms' worst trading days in that same period. Geopolitical risk in the Caldara and Iacoviello sense is a standing condition rather than a sequence of surprises: tension persists between headlines, conditioning hedging behaviour, inventory policy, supplier contracting and portfolio allocation across the year. An event-window design would answer a different question and would require abandoning the firm-level accounting variables that constitute the contribution.

Second, and easily overlooked, the dependent variable is itself built from daily data. For each firm-year, expected shortfall is computed from roughly 245 daily returns, of which the worst twelve are retained and averaged. Daily granularity survives

inside the outcome variable, precisely in the region of interest. What is aggregated to annual frequency is the regressor, not the outcome.

Third, on feasibility. EXPO and ROA derive from audited annual statements. Indonesian issuers publish interim reports, but quarterly disclosure of cost of goods sold is neither universal nor consistent across the sample, and reconstructing a quarterly series would introduce more measurement error than it removed. The Middle East import share is published annually. A higher-frequency design would therefore require dropping the firm-level variables or fabricating them by interpolation.

The cost is real and should be stated rather than argued away. The design cannot date the response, cannot measure how quickly a development is absorbed into prices, and cannot distinguish a January shock from a November one within the same year. It also reduces the number of distinct values the macro regressors take, with consequences for power discussed below.

3.2. Population and Sample

The population comprises all stocks listed on the Indonesia Stock Exchange, approximately 900 securities at the end of the study period. The sampling frame is the Kompas 100 index, the hundred most liquid and largest-capitalisation stocks. The observation period runs from 2015 to 2025.

Banking and financial institutions are excluded, covering the major banks including BBCA, BMRI, BBRI, BBNI, BBTN, ARTO and BRIS, together with other financial institutions such as PNLF. A bank's cost structure is not comparable with a non-financial firm's. Banks report no conventional cost of goods sold line; their principal cost is the cost of funds, the interest paid on deposits and borrowings to finance lending. A COGS to sales ratio is therefore meaningless for a bank and cannot sit on the same scale as one for a manufacturer or distributor. A second reason reinforces this: the cost of funds is driven by the Bank Indonesia policy rate, deposit competition and monetary conditions, not by crude prices. GEAR operates

through input costs, and that channel does not apply to an institution whose principal input is money.

The reduction from 100 constituents to 54 eligible firms follows from the balanced panel requirement, which demands complete data for every variable in every year. Many current constituents listed after 2015 and have neither daily return data nor audited statements for the earlier years: GOTO listed in April 2022, BUKA in August 2021, ARTO converted from Bank Artos in 2020, MBMA listed in April 2023 and ADMR in December 2021. Because the balanced structure is what permits Panel EGLS estimation with consistent dimensions, all firms without complete eleven-year coverage are excluded. The final sample is 54 cross-sections and 594 firm-year observations.

One consequence deserves acknowledgement. Restricting the sample to continuously listed firms introduces a mild survivorship characteristic, since firms that failed, delisted or lost liquidity are absent by construction. The bias is limited by the nature of the index, where delisting is rare, but the results describe firms that survived the period rather than the full population that entered it.

3.3. Variable Operationalization

The model contains one dependent variable, three independent variables, one control and one interaction term. Each is defined below with a numbered equation, its components and sources, and the reasoning behind its form. Firm-level data are in rupiah and oil prices in dollars; no conversion is applied, since every variable is a ratio or percentage change and therefore dimensionless once constructed.

3.3.1. Expected Shortfall (ES5), Dependent Variable

Expected shortfall at the fifth percentile is computed separately for each firm in each year from that firm's daily returns, in four stages.

Stage 1: Daily returns

$$r_d = \frac{P_d - P_{d-1}}{P_{d-1}} \dots\dots\dots (3.1)$$

Where:

r_e : simple daily return on trading day d

P_e : closing price on day d (Bloomberg, PX_LAST)

P_{e-1} : closing price on the previous trading day

The numerator is the price change between consecutive trading days; dividing by the opening reference price converts it into a proportion of the capital at risk. A move from IDR 5,000 to IDR 4,750 gives -0.05 , a five percent loss. Scaling by the starting price is what makes returns comparable across stocks, since a fifty-rupiah move means something different at IDR 200 than at IDR 20,000.

Returns are computed on consecutive trading days rather than calendar days, so a Monday return is measured against the previous Friday's close. Simple rather than logarithmic returns are used, following the convention established with the academic supervisor. The two are close for small movements, since $\ln(1 + r)$ approximates r , but they diverge for the large negative moves populating the left tail, which is the region of interest. Simple returns are bounded below at -100 percent, matching the economic reality that a shareholder cannot lose more than the position. Because the first return of 2015 requires the closing price of the last trading day of 2014, the price series was downloaded from 2014. **Stage 2: Tail size**

$$n_\alpha = \lfloor \alpha \times N \rfloor \dots\dots\dots (3.2)$$

Where:

$ES(\alpha)$: Expected Shortfall at significance level α , with $\alpha = 0.05$ in the primary specification

n_α : number of worst trading days in the tail, computed as the floor of $\alpha \times N$

N : total trading days for that firm in that year, typically 240 to 250 on the IDX

r_e : daily returns sorted ascending, so the most negative occupies position 1

The floor function rounds down, since a fraction of a trading day has no meaning. With 245 trading days, $\text{floor}(0.05 \times 245) = 12$, so the tail comprises the twelve worst days. Rounding down keeps the tail strictly within the worst five percent.

Sorting ascending places the largest losses first, so the tail is simply the first n_a entries; sorting is therefore a necessary step rather than a presentational one. A minimum of 200 trading days per firm-year is required, below which the tail would hold fewer than ten observations and the average would be unstable.

Stage 3: Averaging the tail

$$ES(\alpha)_{i,t} = \frac{1}{n_\alpha} \sum_{j=1}^{n_\alpha} r_j \dots\dots\dots (3.3)$$

Expected shortfall is the arithmetic mean of the n_a worst returns. Each tail day receives equal weight; days outside the tail receive none and have no influence whatever. This is the property distinguishing expected shortfall from value at risk. VaR would report only the return at the tail boundary and say nothing about whether the days beyond it were marginally or catastrophically worse. Two firms could share a VaR while one lost a further two percent on its worst day and the other twenty. In the terminology of Artzner et al. (1999), VaR cuts the tail off; expected shortfall measures what is inside it.

Stage 4: Absolute value

$$|ES5|_{i,t} = |ES(0.05)_{i,t}| \dots\dots\dots (3.4)$$

The Stage 3 average is negative, since it averages losses. Taking the absolute value makes it positive, so a larger ES5 means more severe tail losses. This is a sign convention adopted so coefficient signs read intuitively: positive means a deeper tail, negative a shallower one. Applied consistently, it lets the reader take a positive coefficient at face value. Left in negative form, every coefficient sign would reverse.

Worked illustration

Suppose a stock trades on 20 days in a year with returns already sorted ascending, as in Table 3.1.

Table 3.1. Simulated Daily Returns (Sorted Ascending)

Day #	1	2	3	4	5	6	7	8	9
Return (%)	-3.21	-2.85	-2.14	-1.92	-1.55	-1.20	-0.87	-0.52	-0.18

Day #	10	11	12	13	14	15	16	17	18
Return (%)	+0.15	+0.34	+0.58	+0.89	+1.12	+1.45	+1.78	+2.05	+2.43

Note: days 19 and 20 return +2.87% and +3.50%. Hypothetical data for illustration.

1. Tail size: $n_a = \text{floor}(0.05 \times 20) = 1$. The worst five percent of twenty days is one day.
2. Tail identification: sorted ascending, the worst day is Day 1 at -3.21 percent.
3. Averaging: with one observation, $ES = -3.21\%$.
4. Absolute value: $|ES5| = 3.21\%$.

In the actual dataset each firm has 240 to 250 trading days, so $n_a = 12$ and ES5 averages the twelve worst returns. A firm whose twelve worst days ranged from -5.2 to -2.1 percent and averaged -3.15 would record ES5 of 3.15, read as: on its worst five percent of trading days the stock lost an average of 3.15 percent per day. Higher values indicate worse bad days; lower values indicate that even the worst days were mild. Variation in ES5 across firms and years is what the model seeks to explain.

The fifth percentile is the standard threshold for conditional value at risk and is used throughout as the primary dependent variable. For completeness, expected shortfall is also computed at the first and tenth percentiles by the identical fourstage procedure with a different value of alpha, and these serve as robustness checks in Chapter 4 to establish that the results do not depend on where the tail is cut.

3.3.2. GPR Index, Independent Variable

The daily Geopolitical Risk Index (Caldara and Iacoviello, 2022) is collapsed into one figure per year by arithmetic mean.

$$GPR_t = \frac{1}{D} \sum_{d=1}^D GPR_{daily,d} \dots\dots\dots (3.5)$$

Where:

GPR_t : annual average Geopolitical Risk Index for year t, an index number with no unit

$GPR_{daily,d}$: daily GPR reading on day d (matteoiacoviello.com)

D : number of daily observations in year t

The underlying index counts the share of articles across ten major Englishlanguage newspapers discussing geopolitical threats, scaled so the long-run average sits near 100. A reading of 150 indicates coverage roughly fifty percent above the historical norm. Because it is a share rather than a count, it is not distorted by long-run growth in news volume.

Averaging strips out the day-to-day volatility of a text-based measure, which spikes when a story breaks and decays over following weeks, while retaining year-to-year movement reflecting genuine shifts in tension. The annual average of 78.0 in 2020, when pandemic coverage crowded out geopolitical coverage, and 153.0 in 2022 illustrates the range within the sample.

One structural feature recurs throughout this chapter. GPR_t carries a t subscript and no i subscript: every firm receives the identical value in a given year. The variable has no cross-sectional variation, and its coefficient can only be identified from movement across the eleven years.

3.3.3. GEAR (Geopolitical Energy-Adjusted Risk), Independent Variable

GEAR represents Indonesia's exposure to energy cost pressure originating in the Middle East, and is built in five steps.

Step 1: Composite annual oil price

$$OilPrice_t = \frac{Brent_t + WTI_t + Murban_t}{3} \dots\dots\dots (3.6)$$

Where:

Brent : ICE Brent Crude annual average, USD per barrel (Bloomberg: CO1 Comdty)

WTI : NYMEX West Texas Intermediate annual average (Bloomberg: CL1 Comdty)

Murban : ICE Murban Crude annual average (Bloomberg: MUC1 Comdty)

Each benchmark is averaged from daily front-month settlements across the calendar year, then the three are averaged equally, there being no defensible basis for weighting one above another. Three benchmarks rather than one reflects Indonesia's sourcing. Brent references Atlantic basin and European crude and prices roughly two-thirds of internationally traded oil. WTI is the North American benchmark, relevant since the American export ban was lifted in December 2015. Murban is the Abu Dhabi grade and the reference most directly applicable to Gulf cargoes, the sourcing region the variable concerns. Averaging also guards against contract-specific dislocations: in April 2020 the expiring WTI contract settled negative because Cushing storage was exhausted, an event with no counterpart in Brent or Murban and no bearing on crude delivered to Indonesia. One qualification applies to the composite. The ICE Murban contract began trading only in 2021, so the average rests on Brent and WTI alone for 2015 through 2020 and on all three benchmarks from 2021 onward. Recomputing 2021, the year carrying the largest GEAR0 value, on the two-benchmark basis returns 23.54 against the reported 23.80, so the change in composition does not drive the result.

Step 2: Year-over-year change

$$\% \Delta OilPrice_t = \frac{OilPrice_t - OilPrice_{t-1}}{OilPrice_{t-1}} \times 100 \dots\dots\dots (3.7a)$$

Where:

$\% \Delta OilPrice_t$: year-over-year percentage change in the composite annual price

$OilPrice_t$: composite annual oil price for year t

$OilPrice_{t-1}$: composite annual oil price for the preceding year

The numerator is the change between consecutive years and the denominator the prior level, giving a proportional rather than absolute change. This matters because a ten-dollar move carries different weight depending on where the price starts: from USD 30 it is a thirty-three percent increase, from USD 100 it is ten. Expressing the change proportionally puts every year on a comparable scale. Working in changes rather than levels is equally deliberate, since the level is highly persistent and trending, and regressing an annual risk measure on a trending level risks a spurious relationship. The change is closer to stationary and represents the concept of interest, namely the emergence of new cost pressure rather than the standing cost of energy.

Step 3: Floor at zero

$$\% \Delta OilPrice_t^+ = \text{Max}(0, \% \Delta OilPrice_t) \dots\dots\dots (3.7b)$$

The $\text{Max}(0, \cdot)$ operator compares the change with zero and returns the larger. A rise passes through unaltered; a fall returns zero. The variable is censored from below, taking positive values or zero but never negative. The censoring binds more often than might be expected: across the eleven years GEAR0 is zero in six, leaving five distinct non-zero values on which the coefficient is identified. This substantially reduces effective variation and is treated as a limitation rather than a design feature, which is why the symmetric alternative described below is estimated alongside.

Step 4: Middle East import share

$$ME_ImportShare_t = \frac{Import_{crude,t}^{ME} + Import_{fuel,t}^{ME}}{Import_{crude,t}^{World} + Import_{fuel,t}^{World}} \dots\dots\dots (3.7c)$$

Where:

Import(ME, crude) : crude petroleum imports from Middle East countries (HS 270900)

Import(ME, fuel) : refined petroleum imports from Middle East countries (HS 2710)

Import(World, crude) : total crude petroleum imports from all origins (HS 270900)

Import(World, fuel) : total refined petroleum imports from all origins (HS 2710)

The numerator sums crude and refined petroleum sourced from the Middle East; the denominator sums the same categories across all origins. The quotient is the proportion of total petroleum imports from the region, a dimensionless figure between zero and one. Both crude and refined products are included because both are physically exposed to a Hormuz disruption, and Indonesia imports substantial volumes of gasoline and diesel alongside crude. Import values rather than volumes are used, which is what TradeMap reports; the ratio form means numerator and denominator are valued on the same basis, so no price distortion is introduced.

The Middle East is defined as Saudi Arabia, consistently the largest source, together with the United Arab Emirates, Iraq, Kuwait, Oman, Bahrain, Qatar, and Iran prior to full sanctions. Across the sample the share has moved between roughly 30 and 58 percent. This component is the strongest part of the construct: an officially reported, directly observable measure of physical dependence requiring no assumption about how the number was generated.

Step 5: Multiplication

$$GEAR0_t = \% \Delta OilPrice_t^+ \times ME_ImportShare_t \dots\dots\dots (3.7d)$$

Where:

GEAR0_t : Geopolitical Energy-Adjusted Risk for year t

%ΔOilPrice_t⁺ : floored year-over-year oil price change

ME_ImportShare_t : Middle East share of Indonesian petroleum imports

The multiplicative form encodes a specific proposition: exposure requires both a shock and a channel through which it arrives. A country importing nothing from the Middle East is untouched by a Hormuz disruption however far prices rise, and a year of falling prices delivers relief regardless of sourcing. Multiplying yields a variable that is zero whenever either condition is absent, which is the intended

behaviour; an additive construction would report positive exposure in a falling year merely because the import share was high. The form also scales correctly, since the same proportional increase produces a larger value when the import share is higher. This mirrors risk measurement generally, where exposure is the product of a hazard and the degree of contact with it.

The numerical scale depends on how the components are expressed, and the convention determines how coefficients are read. If the price change is in percentage points and the import share a decimal, a forty percent increase with a 45 percent share gives 18.0; if both are decimals, 0.18. The two differ by a factor of one hundred and the estimated coefficient differs by the same factor inversely, leaving fitted values identical. The scaling used is stated with the descriptive statistics in Chapter 4. The annual frequency is fixed by the components, since the TradeMap import share exists only annually.

3.3.3.1. Construct Validity of GEAR

A variable not previously used carries the burden of showing it measures what it claims. GEAR is offered as a measure of the intensity of energy cost pressure reaching Indonesia from Middle East conditions. Whether it delivers that depends on the oil price component picking up geopolitically driven movement, on the import share measuring exposure to it, and on the product capturing the two acting together.

The last of those is a matter of arithmetic and the second is a matter of official statistics. Import share is a directly observable measure of physical dependence, requiring no assumption about how the figure was generated, and the multiplicative form behaves as intended by demanding that both conditions be present before the variable registers anything at all.

The first is where the construct is weak. Oil prices move for many reasons and only some of them are geopolitical: demand cycles, OPEC and OPEC+ quota decisions, American shale supply growth, inventory dynamics, speculative positioning. GEAR cannot decompose the composite price by source, so a rise driven by Chinese industrial demand enters the variable with precisely the weight of one

driven by a Gulf supply interruption. The variable can be said with reasonable confidence to measure energy cost pressure weighted by regional import dependence. The geopolitical origin of that pressure it measures far less confidently. This is a limitation of discriminant validity, and calling GEAR a pure measure of geopolitical energy risk would claim more than the construction supports.

The design does not solve the problem, though it does contain some mitigation. GPR sits in the same regression, so the GEAR coefficient is estimated conditional on the general level of tension and the part of cost pressure that coincides with broad escalation is partly absorbed there; because the two series are correlated and neither is a clean measure the separation is imperfect, but it improves on estimating GEAR alone. Weighting by import share also pulls the variable toward the region where supply disruption is most likely to originate. A price movement caused by North American conditions still enters, but it enters scaled by Indonesia's dependence on the Gulf.

A stronger construction exists in principle. Kilian's structural decomposition splits oil price changes into supply, aggregate demand and precautionary demand components, and isolating the supply-disruption element before applying the import weight would produce something much closer to a pure geopolitical measure. It is not attempted here. The decomposition requires a structural VAR estimated on monthly global production, activity and inventory data, which is a project in its own right and outside the intended scope, and the monthly series it produced would then have to be pushed back to annual frequency to sit alongside the accounting variables, surrendering most of the precision it had bought. Chapter 5 lists it as future work.

Some of the claim can be checked against data already assembled, and Chapter 4 reports the result. If GEAR measures what it claims, it should be positively associated with the level of the composite oil price, since the two share a driver, and positively but well short of perfectly associated with GPR, since tension and energy cost pressure overlap without being the same thing. A correlation with GPR

near unity would indicate redundancy and one near zero would indicate no geopolitical content at all; the observed value falls between.

3.3.3.2. Exclusion of Negative Oil Price Movements

The floor discards the information in falling years, and the intuitive justification that falling prices bring relief, while correct, is not academically sufficient. Four arguments support the censoring and place it within an established tradition.

1. The generating processes differ. In the Kilian and Park (2009) decomposition, large increases and declines typically arise from different disturbances rather than one operating in reverse. Within the window the sharp increases trace to supply and precautionary-demand shocks; the large declines trace to the supply glut of 2014 to 2016 and the demand collapse of 2020. A symmetric variable would record the pandemic collapse as a large negative geopolitical energy shock, inverting the construct.
2. The econometric precedent is established. Mork (1989) separated increases from decreases and found only increases carried significant explanatory power for output, concluding symmetric specifications are misspecified. Hamilton (1996, 2003) developed this into the net oil price increase measure, zero unless the price exceeds its recent maximum, and showed it outperforms the linear change. The $\text{Max}(0, \cdot)$ operator applies the same logic at annual frequency.
3. Cost pass-through is itself asymmetric. Downstream pricing research documents the rockets and feathers pattern, in which increases pass through faster and more completely than decreases (Bacon, 1991; Borenstein, Cameron and Gilbert, 1997). If transmission is genuinely asymmetric, a symmetric regressor is structurally incorrect rather than merely inefficient.
4. The dependent variable is one-sided. Expected shortfall measures only adverse severity and is insensitive to the favourable side. Pairing it with a two-sided regressor would place energy relief and energy stress on a single linear scale mapping onto a variable registering only stress.

The cost is recorded alongside. Censoring removes variation, narrowing the range over which the coefficient is identified, and prevents any test of whether declines carry a symmetric relationship of opposite sign. A symmetric alternative built on the absolute change, denoted GEAR Absolute, was constructed and is retained as a robustness specification so the consequence can be observed rather than asserted.

3.3.3.3. Uniform GEAR Across Firms

GEAR0 carries a time subscript only: every firm receives the same value in a given year, and the variable contains no cross-sectional variation. This follows from building it from national trade statistics and global prices, neither observable at firm level, and it has three consequences.

The cost is heterogeneity. Firms differ in energy dependence in ways the variable cannot see: a mining contractor burning diesel in heavy equipment, a cement producer running kilns, a petrochemical processor buying naphtha as feedstock, and a tower operator drawing grid electricity all meet a crude shock differently, and GEAR treats them alike. Two firms reporting identical EXPO but with quite different fuel intensity receive the same value. What the coefficient describes, then, is the association between national energy cost pressure and the average tail risk across the sample, not the effect on any firm within it, and the discussion in Chapter 4 keeps to that reading.

Two further consequences, for identification and for inference, follow from the same construction. Both are set out with the year-effects question at Section 3.4.4.4 rather than anticipated here.

One element of the design does put cross-sectional variation back into the energy channel. EXPO varies across firms, so the product $\text{GEAR0} \times \text{EXPO}$ differs between firms within a year even where GEAR0 does not. The interaction is therefore identified from comparison within years, which is a considerably stronger basis than the eleven-point time series carrying the main effects, and it is the reason the moderation test survives the addition of year fixed effects when the main effects cannot.

3.3.4. Firm Exposure (EXPO), Independent Variable

Firm Exposure is the ratio of cost of goods sold to net sales.

$$EXPO_{i,t} = \frac{COGS_{i,t}}{NetSales_{i,t}} \dots\dots\dots (3.8)$$

Where:

EXPO_{i,t} : Firm Exposure for firm i in year t, a dimensionless ratio

COGS_{i,t} : Cost of Goods Sold for the fiscal year (Bloomberg: COGS_AND_OPER_EXP)

NetSales_{i,t} : Net Sales or Revenue for the fiscal year (Bloomberg: IS_REVENUE)

The calculation divides two figures from the same income statement for the same year. Both are in rupiah, so units cancel. A firm reporting revenue of IDR 100 billion against costs of IDR 82 billion has EXPO of 0.82: 82 percent of every rupiah of revenue is consumed by direct costs before any other expense, leaving an 18 percent gross margin. Taking the ratio rather than the level makes the measure comparable across firms, since a conglomerate with costs in the tens of trillions and a small manufacturer in the tens of billions cannot be compared on absolute figures but can be compared on proportion. Scaling by revenue also removes the effects of firm size, inflation and currency movement across the eleven-year window.

Data are drawn from Bloomberg and cross-checked against audited statements; where they disagree, the audited statement is authoritative. One definitional matter requires confirmation because it bears on interpretation: the Bloomberg field must be verified to return cost of goods sold alone rather than cost of goods sold combined with operating expenses. The distinction is material. A ratio of direct cost to sales measures production cost intensity and, because cost of goods sold is predominantly variable, carries an implication about flexibility. A ratio of total operating cost to sales includes the largely fixed selling and administrative component, which reverses that implication.

3.3.4.1. Construct Validity of EXPO

The proposition is that a high ratio of direct costs to revenue indicates greater vulnerability to oil-related shocks. It contains two links of unequal security.

The first, from ratio to cost intensity, is close to definitional. A firm consuming 85 percent of revenue in direct costs operates on a thinner cushion than one at 50 percent, and a given input price increase consumes a larger share of what remains.

The second, from cost intensity to oil vulnerability specifically, is weaker, because the ratio reports the size of a cost base and says nothing about what is inside it.

Darma Henwa (DEWA IJ Equity) and AKR Corporindo (AKRA IJ Equity) both sit near 90 percent, and they are almost the cleanest illustration the sample offers. Darma Henwa is a mining contractor. Diesel dominates its cost base, its output is priced under contract, and a higher crude price is simply an increase in the cost of doing business. AKR Corporindo distributes fuel and chemicals. The same crude price sits on its revenue line as well as its cost line, so an increase passes through both sides of the income statement and can leave the firm better off. EXPO returns two nearly identical numbers for two firms whose exposure to the same shock points in opposite directions.

Other failure modes run in each direction. A firm buying mostly finished goods, agricultural inputs or labour registers a high ratio while remaining comparatively insensitive to crude, and a firm with a modest ratio can be materially exposed where the energy component is large relative to its margin, in which case EXPO understates it rather than overstating it.

Isolating the fuel component of cost of goods sold would settle all three, and it cannot be done from Indonesian disclosure. Issuers are not required to disaggregate by input category on the face of the income statement. Some provide a breakdown in the notes, but the granularity varies by sector, by firm and by reporting year, so no consistent energy-cost ratio can be assembled across 54 firms and 11 periods. A partially populated variable would be worse than none, since the firms that disclose fuel costs are

disproportionately the firms for which fuel costs are material, which is selection running along exactly the dimension being measured.

What the design does instead is limit how far the ratio is asked to carry. The interaction supplies the one indirect test available within this specification: if EXPO genuinely captures vulnerability rather than cost intensity in general, its relationship with tail risk should be stronger in years when energy costs rose than in years when they did not. Two further checks would bear more directly on the concern, namely reporting the sector composition of the sample and re-estimating with energy sector firms excluded. Neither is carried out here, and both are identified in Section 5.4.3 as extensions this design leaves open.

An interpretive rule follows. If the EXPO coefficient is significant while the interaction is not, the evidence supports the conclusion that cost intensity is associated with tail risk in general, and does not support the narrower conclusion that firms are differentially vulnerable to oil shocks according to cost structure. Only a significant interaction supports the oil vulnerability reading. The variable is therefore described throughout as a measure of cost intensity, with inference to oil vulnerability made conditional on the interaction.

3.3.5. ROA (Return on Assets), Control Variable

Return on assets controls for profitability differences across firms.

$$ROA_{i,t} = \frac{NetIncome_{i,t}}{TotalAssets_{i,t}} \dots\dots\dots (3.9)$$

Where:

$ROA_{i,t}$: Return on Assets for firm i in year t, in percent (Bloomberg: RETURN_ON_ASSET)

$NetIncome_{i,t}$: net income for the fiscal year

$TotalAssets_{i,t}$: average of beginning and ending total assets

The numerator is a flow, profit generated over twelve months; the denominator a stock, the asset base generating it. Because a flow divides a stock, the denominator is the average of opening and closing balance sheets rather than the closing figure.

A firm doubling its assets in December would otherwise appear to have earned a full year's profit on a base held for one month. The ratio answers a specific question: for every rupiah of assets controlled, how many rupiah of profit were generated? A ROA of 8 percent means eight rupiah per hundred.

Return on assets was preferred to return on equity because it is unaffected by capital structure. Return on equity divides by shareholders' equity, which can become very small or negative for a firm that has bought back stock or accumulated losses, producing extreme or meaningless values. This is not hypothetical here: Bloomberg reports return on equity exceeding 200 percent for Unilever Indonesia owing to negative equity from accumulated buybacks. Total assets remain positive regardless of financing. The Bloomberg field applies a standardised calculation using trailing twelve-month net income over average total assets, ensuring identical treatment across all 54 firms.

ROA carries no separate hypothesis. It enters as a control on the established finding that profitability buffers adverse shocks (Novy-Marx, 2013; Fama and French, 2015), absorbing cross-firm variation that would otherwise remain in the error term and could confound the cost structure relationship, since cost-intensive firms tend also to be less profitable. It is not interacted with GEAR or GPR, and the model makes no claim that profitability alters the slope of the geopolitical relationship.

3.3.6. Interaction Term (GEAR0 × EXPO)

The interaction is the product of energy exposure and firm cost structure.

$$INTERACTION_{i,t} = GEAR0_t \times EXPO_{i,t} \dots\dots\dots (3.10)$$

Where:

INTERACTION_{i,t} : product of national energy cost pressure and firm cost intensity

GEAR0_t : Geopolitical Energy-Adjusted Risk in year t, common to all firms

EXPO_{i,t} : Firm Exposure for firm i in year t

For each firm-year the GEAR value applying to that year is multiplied by that firm's EXPO. Because GEAR is constant within a year and EXPO varies across firms,

the product varies in both dimensions, and that two-dimensional variation identifies the moderation relationship. An interaction is required because a main effects model imposes a restriction: it assumes the GEAR relationship is identical for every firm regardless of cost structure. Adding the product relaxes it, so the association between GEAR and ES5 becomes $\beta_2 + \beta_4 \times \text{EXPO}_{i,t}$ rather than a single number applying to all firms.

The coefficient β_4 measures how much the GEAR slope changes per one-unit increase in EXPO, and its magnitude cannot be interpreted without fixing the scaling of both components. If EXPO enters as a decimal, a firm at 80 percent enters as 0.80 and the slope change is $0.80 \times \beta_4$; in percentage points the same firm enters as 80, a difference of two orders of magnitude. Chapter 4 states the scaling and expresses magnitudes accordingly.

The sign carries the substantive content. A positive significant coefficient indicates that cost-intensive firms experience a larger increase in tail risk when energy pressure rises, supporting the margin reading. A negative significant coefficient indicates partial shielding, consistent either with cost flexibility, since a variable-weighted base can be reduced, or with commodity pass-through, since firms whose revenues move with commodity prices absorb the shock on both sides.

A caution follows from the two construct-validity discussions above. The credibility of the moderation result is bounded by that of its components. GEAR is an imperfect measure of geopolitically driven cost pressure and EXPO an imperfect measure of oil vulnerability; measurement error in either propagates into the product, and classical measurement error attenuates a coefficient toward zero. A significant interaction is therefore the informative outcome, since error works against finding one. A non-significant interaction admits two readings the data cannot separate, that no moderation exists or that it is obscured by noise, and should not be reported as evidence against moderation.

3.3.7. Regression Specification

The variables above enter a single panel regression estimated on the balanced panel of 54 firms over 11 years.

$$ES5_{i,t} = \beta_0 + \beta_1 GPR_t + \beta_2 GEAR0_t + \beta_3 EXPO_{i,t} + \beta_4 (GEAR0_t \times EXPO_{i,t}) + \beta_5 ROA_{i,t} + u_i + \varepsilon_{i,t} \quad \dots\dots\dots (3.11)$$

Where:

$ES5_{i,t}$: Expected Shortfall at the 5th percentile for firm i in year t , absolute value
 β_0 : constant term
 β_1 : coefficient on the annual average Geopolitical Risk Index
 β_2 : coefficient on GEAR0, evaluated at EXPO equal to zero
 β_3 : coefficient on Firm Exposure
 β_4 : coefficient on the interaction, the change in the GEAR slope per unit of EXPO
 β_5 : coefficient on the control variable Return on Assets
 u_i : firm-specific effect
 $\varepsilon_{i,t}$: idiosyncratic error term

The error separates into two components, which is what distinguishes panel estimation from pooled OLS. The term u_i captures whatever is persistently distinctive about firm i and does not change across the eleven years: its industry, business model, permanent risk profile. The term $\varepsilon_{i,t}$ captures what varies year to year and is unexplained by the regressors. Splitting the error acknowledges that eleven observations from one firm are not independent draws.

The regressors vary in different dimensions, and this determines what the model can identify. GPR and GEAR0 vary across time only. EXPO and ROA vary across both firms and time. The interaction varies across both because one component does.

Table 3.2. Variable Operationalization Summary

No.	Variable	Operationalization	Symbol	Role	Source
1	Expected Shortfall	Mean of the worst 5% of daily returns in the firm-year, absolute value	ES5	Dependent	Bloomberg (PX_LAST); simple returns
2	GPR Index	Caldara-Iacoviello daily GPR, annual arithmetic mean	GPR	Independent	matteoiacoviello.com
3	GEAR	Max(0, % change in 3-benchmark composite oil price) $\times$ ME import share	GEAR0	Independent	Bloomberg (CO1, CL1, MUC1); TradeMap.org (HS 270900, 2710)
4	Firm Exposure	COGS $\div$ net sales	EXPO	Independent	Bloomberg (BDH/BDP); audited financial statements


5	ROA	Net income ÷ average total assets	ROA	Control	Bloomberg (RETURN_ON_ASSET)
6	Interaction	GEAR0 × EXPO	GEAR0 × EXPO	Interaction	Computed from variables 3 and 4

Source: Author (2026)

3.4. Analytical Procedure

Analysis proceeds in a fixed sequence, each stage informing the next. Descriptive statistics establish distributional properties bearing on the choice of estimator. Correlation and variance inflation factors identify collinearity that would destabilise coefficient estimates. Model selection tests determine the panel specification. Classical assumption tests establish whether standard errors require correction. The regression is then estimated and the hypotheses examined. Robustness checks finally establish whether the conclusions survive the specification choices made along the way.

Figure 3.1. Research Process



Source: Author (2026)

3.4.1. Descriptive Statistics

Descriptive statistics summarise the sample before inferential analysis, allowing distributional characteristics to be understood on their own terms. Mishra et al. (2019) note that such measures are essential for detecting outliers, assessing normality and evaluating central tendency. The mean, standard deviation, skewness, kurtosis, minimum and maximum are reported for every variable.

3.4.1.1. Mean

$$\bar{X} = \frac{X_1 + X_2 + X_3 + \cdots + X_n}{n} = \frac{\sum_{i=1}^n X_i}{n} \dots\dots\dots (3.12)$$

Where:

$\bar{X}$: the arithmetic mean

$X_1, X_2, \dots, X_n$: the individual observations n

: the number of observations

The calculation sums every observation and divides by their count. Conceptually the mean is the balance point of the distribution, the value at which positive deviations above exactly offset negative deviations below. Because it weights every observation equally, a small number of extreme values can pull it away from the bulk of the data, which is why it is reported alongside the median, minimum, maximum and shape statistics.

3.4.1.2. Standard deviation

$$s = \sqrt{\frac{\sum_{i=1}^n (X_i - \bar{X})^2}{n - 1}} \dots\dots\dots (3.13)$$

Where:

s : the sample standard deviation

X_i : individual observation i $\bar{X}$

: the sample mean n : the

number of observations

The formula builds in four steps. Each observation's deviation from the mean is computed. Each is squared, which removes the sign so deviations above and below

both contribute positively rather than cancelling, and which gives disproportionate weight to large deviations. The squared deviations are summed and divided by $n - 1$ to give the variance. The square root returns the measure to the original units, so a standard deviation of ES5 is in percentage points rather than squared percentage points. The divisor is $n - 1$ because the mean was itself estimated from the same data, consuming one degree of freedom; dividing by n would bias the estimate downward.

For this study the dispersion of ES5 carries particular weight. Large cross-sectional dispersion indicates that some firms experienced far more severe tail losses than others in the same market and years, which is the variation the firm-level regressors are asked to explain.

3.4.2. Correlation Matrix and Multicollinearity

3.4.2.1. Pearson correlation coefficient

$$r_{xy} = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2 \cdot \sum_{i=1}^n (Y_i - \bar{Y})^2}} \dots\dots\dots (3.14)$$

Where:

r_{xy} : Pearson correlation coefficient between X and Y

X_i, Y_i : paired observations of the two

variables $\bar{X}, \bar{Y}$: the sample means n : the

number of paired observations

The numerator sums, across all observations, the product of each variable's deviation from its own mean. This is the covariance before scaling, and its sign carries the information. When an observation lies above the mean on both variables both deviations are positive and the product is positive; when below on both, both are negative and the product is again positive. Only when an observation lies above on one and below on the other is the product negative. A positive sum therefore indicates the variables move together.

The denominator scales by the product of the two dispersions. Without it the covariance would depend on the units of measurement and could take any

magnitude. Dividing bounds the result between -1 and $+1$, where $+1$ is a perfect positive linear relationship, -1 a perfect negative one and 0 no linear relationship. Two limitations apply: the coefficient captures linear association only, so a strong curved relationship can produce a value near zero, and correlation says nothing about which variable produces the other.

Table 3.3. Interpretation of Correlation Magnitude

Size of r	Interpretation
0.90 – 1.00	Very high correlation
0.70 – 0.89	High correlation
0.50 – 0.69	Moderate correlation
0.30 – 0.49	Low correlation
0.00 – 0.29	Little or no correlation

Source: Asuero et al. (2006)

The matrix serves three purposes here. It screens for pairwise collinearity among regressors, with 0.80 taken as the threshold of concern. It provides the convergent validity check on GEAR, since the correlation with GPR indicates whether the two capture overlapping but distinct constructs or are effectively redundant. And it establishes the arithmetic dependence between the interaction term and its components, which bears directly on how the main effects are read.

The third purpose is easily overlooked. A correlation matrix computed over the dependent variable and the four substantive regressors cannot reveal a collinearity problem residing in the product term, because the product does not appear in it. The matrix reported in Chapter 4 therefore includes $\text{GEAR0} \times \text{EXPO}$ as a row and column in its own right.

3.4.2.2. Variance Inflation Factor

$$VIF_j = \frac{1}{1 - R_j^2} \dots\dots\dots (3.15)$$

Where:

VIF_j : variance inflation factor for independent variable j

R_j^2 : R-squared from regressing variable j on all other independent variables

The calculation runs in two stages. Variable j is regressed on every other independent variable and the R-squared recorded, measuring how much of j 's variation the others already account for. The VIF is one divided by one minus that figure. The intuition is clearest at the extremes: if the others explain none of j , $R_j^2 = 0$ and $VIF = 1$, so the coefficient's variance is not inflated. If they explain 90 percent, $R_j^2 = 0.90$ and $VIF = 10$, meaning the variance is ten times what it would have been under independence and the standard error is inflated by $\sqrt{10}$, roughly 3.2.

Multicollinearity does not bias coefficients; it makes them imprecise. When two regressors move together closely the data contain little information about which is associated with the outcome, so estimates become unstable, standard errors widen, and a variable that genuinely matters may fail to reach significance. Conventional thresholds treat VIF below 5 as acceptable, 5 to 10 as warranting attention, and above 10 as severe (Hair, Black, Babin and Anderson, 2014).

One result shaped the diagnostic and should be anticipated. An interaction term is arithmetically bound to the variables from which it is constructed, so a product term registers an elevated inflation factor by construction rather than as evidence of redundancy in the data. Here the elevation is not marginal: GEAR0 and $GEAR0 \times EXPO$ correlate at 0.96, with an inflation factor close to 19. Including the product in the diagnostic would flag a condition that is certain in advance of observing any data and that carries no diagnostic content, so the factors reported in Chapter 4 are computed on the specification excluding it. A related interpretive point follows. Where a main effect is significant without the interaction and insignificant once the product is added, the change may reflect that arithmetic dependence rather than economics, which is one reason the specification without the interaction is retained as the primary model.

3.4.3. Classical Assumption Tests

3.4.3.1. Normality test (Jarque-Bera)

$$JB = \frac{n}{6} \left(S^2 + \frac{(K - 3)^2}{4} \right) \dots\dots\dots (3.16)$$

Where:

JB : the Jarque-Bera test statistic

n : the number of observations

S : sample skewness of the residuals

K : sample kurtosis of the residuals, equal to 3 under normality

The statistic combines two departures from normality. Skewness measures asymmetry: a normal distribution is symmetric with skewness zero, while a long left tail gives negative skewness. Kurtosis measures tail weight: a normal distribution has kurtosis exactly 3, and fatter tails give more. The formula squares the skewness and the deviation of kurtosis from 3, so departures in either direction contribute positively, then scales by the number of observations. Both components are zero under normality, so a perfectly normal sample gives $JB = 0$. The statistic follows a chi-squared distribution with two degrees of freedom, one per component.

The null is that residuals are normally distributed, rejected at p below 0.05. For financial return data rejection is expected: returns exhibit negative skewness and excess kurtosis, and that characteristic is part of this study's motivation rather than an inconvenience, being precisely why variance-based measures are inadequate. With 594 observations the central limit theorem also ensures coefficient estimates remain asymptotically normal even when residuals are not, so non-normality does not invalidate inference.

3.4.3.2. Multicollinearity test

The test examines whether independent variables are so closely related that their individual contributions cannot be separated. Where collinearity is severe, coefficients become unstable, standard errors inflate and t-statistics lose reliability.

Variance inflation factors are computed for GPR, GEAR0, EXPO, the interaction and ROA, supported by the correlation matrix with 0.80 as the pairwise threshold.

Null hypothesis (H_0): there is no multicollinearity among the independent variables ($VIF < 10$).

Alternative hypothesis (H_1): multicollinearity exists among the independent variables ($VIF \geq 10$).

$$H_1 : VIF \geq 10 \rightarrow \text{Multicollinearity present}$$

$$H_0 : VIF < 10 \rightarrow \text{No multicollinearity}$$

Where severe collinearity is detected, the remedies are removal of a redundant variable or estimation of the model without the affected term, which is the course taken for the interaction as described above.

3.4.3.3. Heteroskedasticity

Heteroskedasticity arises when error variance is not constant across observations. In financial panel data it is close to inevitable: larger firms typically exhibit lower return volatility, so dispersion differs systematically across the cross-section, and market conditions differ across years. The consequence is not bias but invalid standard errors, which makes every t-statistic and p-value unreliable.

Null hypothesis (H_0): the residuals are homoskedastic, with constant variance.

Alternative hypothesis (H_1): the residuals are heteroskedastic.

$$H_0 : \text{Var}(\varepsilon_i) = \sigma^2 \text{ (Homoscedasticity)}$$

$$H_1 : \text{Var}(\varepsilon_i) \neq \sigma^2 \text{ (Heteroscedasticity)}$$

The decision rule follows the p-value:

- $p \leq 0.10$: reject H_0 at the 10 percent level (*)
- $p \leq 0.05$: reject H_0 at the 5 percent level (**)
- $p \leq 0.01$: reject H_0 at the 1 percent level (***)
- $p > 0.10$: fail to reject H_0 , indicating homoskedasticity

Where heteroskedasticity is present, White diagonal robust standard errors are applied. The correction leaves coefficients unchanged and adjusts only their standard errors, using squared residuals to estimate each observation's variance individually rather than imposing a common one. It addresses heteroskedasticity but not the within-year correlation induced by regressors constant across firms in a given year.

3.4.4. Panel Model Selection

Three tests determine which of three candidate specifications applies: the Common Effect Model, pooling all observations under a single intercept; the Fixed Effect Model, giving each firm its own; and the Random Effect Model, treating the firm component as a random draw.

3.4.4.1. Chow test

The Chow test, known here as the redundant fixed effects test, compares the Common and Fixed Effect Models by examining whether firm-specific intercepts are jointly significant. It is an F-test comparing a restricted model, in which all firms share one intercept, against an unrestricted model in which each has its own, built from the improvement in fit obtained by relaxing the restriction scaled by the degrees of freedom consumed. A large improvement relative to the additional parameters rejects the restriction.

Null hypothesis (H_0): probability > 0.05 , the Common Effect Model is valid.

Alternative hypothesis (H_1): probability < 0.05 , the Fixed Effect Model is valid.

3.4.4.2. Hausman test

The Hausman test determines whether the Random or Fixed Effect Model applies by examining whether firm-specific effects are correlated with the explanatory variables, the assumption on which random effects depends. If correlated, random effects is inconsistent and fixed effects must be used; if not, random effects is consistent and more efficient.

The test compares the two coefficient sets. Fixed effects is consistent whether or not the correlation exists, because it removes the firm effects entirely; random effects is efficient but consistent only when the correlation is absent. Under the null both converge to the same values, so any difference should be small and attributable to sampling variation. A large systematic difference indicates the null is false. The statistic follows a chi-squared distribution with degrees of freedom equal to the number of coefficients compared.

Null hypothesis (H_0): probability > 0.05 , the Random Effect Model is valid.

Alternative hypothesis (H_1): probability < 0.05 , the Fixed Effect Model is valid.

A significance level of 0.10 is commonly applied in social science research, on the grounds that the test has low power in modest samples.

3.4.4.3. Lagrange Multiplier test

The Lagrange Multiplier test, in Breusch-Pagan form, compares the Random and Common Effect Models by examining whether the variance of the firm-specific error component differs significantly from zero. If that variance is zero there are no panel-level effects and pooled estimation suffices. The test proceeds from the pooled model's residuals: if firm effects exist, residuals belonging to the same firm share a common component and are correlated across years, producing a detectable pattern. The statistic follows a chi-squared distribution with one degree of freedom.

Null hypothesis (H_0): probability > 0.05 , the Common Effect Model is valid.

Alternative hypothesis (H_1): probability < 0.05 , the Random Effect Model is valid.

Table 3.4. Panel Model Selection Tests

Test	Comparison	H ₀ (Null)	H ₁ (Alternative)	Decision rule
Chow (Redundant FE)	CEM vs FEM	CEM is appropriate	FEM is appropriate	$p < 0.05$: reject H ₀ , use FEM
Hausman	FEM vs REM	REM is appropriate	FEM is appropriate	$p < 0.10$: reject H ₀ , use FEM
Lagrange Multiplier	CEM vs REM	CEM is appropriate	REM is appropriate	$p < 0.05$: reject H ₀ , use REM

Source: Author (2026)

The three are independent pairwise comparisons rather than a dependent sequence. Chow contrasts Common against Fixed, Hausman contrasts Fixed against Random, and the Lagrange Multiplier test contrasts Common against Random. Because none conditions on the outcome of another, the statistics are unaffected by the order in which they are computed; in practice that order is dictated by the effects specification each requires in the estimation software. They are presented here in the conventional order, moving outward from the simplest specification.

3.4.4.4. Time Effects and Identification of the Macro Coefficients

A further question arises that the three tests do not address: whether year effects should be included alongside firm effects. Because the principal explanatory variables vary across time rather than firms, omitted time-specific influences are a genuine concern.

The answer is definite and technical. GPR_t and $GEAR0_t$ take the same value for every firm within a year, so each is a function of t alone, and a complete set of year dummies spans the space of all functions of t . With year dummies in the equation both macro variables become perfectly collinear with them, their coefficients cannot be identified, and the software will either drop the variables or return a singular matrix. Year effects and macro regressors of this type cannot sit in one specification. That is algebra, not preference.

Since year effects cannot be included, the macro coefficients are identified from variation across eleven annual observations and nothing else. Any national factor that moves over time and is correlated with both geopolitical tension and equity tail risk is a confounder this design cannot remove, and the Bank Indonesia policy

rate, domestic output growth, the exchange rate, global risk appetite and the pandemic shock of 2020 all qualify. The coefficients on GPR and GEAR should be read as the association between those variables and the sample's average tail risk in a given year, conditional on the firm-level regressors. They are not evidence that geopolitical conditions, as opposed to some other national condition present in the same year, account for the pattern.

Power is affected by the same constraint. The panel has 594 rows, but the macro regressors take eleven distinct values, so the effective sample behind their coefficients is nearer eleven than 594. Standard errors computed as though the observations were independent will understate the true uncertainty, which is Moulton's (1990) point. Clustering by year would deal with it directly, though eleven clusters is too few for cluster-robust inference to be dependable.

Three coefficients sit in a better position. EXPO, ROA and the interaction all vary across firms inside a year. None of them is collinear with year dummies, and all three survive when year effects are added. That specification can therefore be estimated. The macro main effects drop out of it, while the three firm-level coefficients remain and are estimated free of confounding from anything common to all firms in a given year, whether geopolitical, monetary, fiscal or pandemic in origin. That specification matters most for the moderation hypothesis. Were the interaction to keep its sign and significance once year effects had absorbed every common annual influence, the moderation result would stop depending on the eleven-point macro series and become considerably more credible than the main effects. It is not estimated for this thesis, and Section 5.4.3 identifies it as the most direct extension available to the present design.

3.4.4.5. Resolution of the Model Selection Procedure

The model-selection tests do not point to a single specification. The Lagrange Multiplier test rejects the Common Effect Model in favour of the Random Effect Model. The Chow test rejects the Common Effect Model in favour of the Fixed Effect Model. The Hausman test is normally used to choose between fixed and random effects, but the cross-section test could not be performed because the

reported variance matrix was invalid. This happens when the difference between the coefficient covariance matrices does not satisfy the positive-definiteness condition that is required for the Hausman statistic. In this case the statistic is not defined and cannot be used to determine the preferred specification.

The test results still allow for both fixed and random effects. The study therefore focuses on the fixed effects specification with cross-section weights as the main model. The Chow test offers direct evidence that firm-specific intercepts matter. Fixed effects also does not require the firm-specific effects to be independent of the explanatory variables. This is an important consideration here because that relationship cannot be established from the Hausman test. The diagnostic tests are indicative of heteroskedasticity and therefore cross-section weighting is applied.

The random effects estimates are reported alongside the main results to show how the choice of specification affects the findings. The coefficient signs are broadly consistent across the two models, and the GPR coefficient remains statistically significant under both specifications, although the significance of the other variables differs. The greater number of significant coefficients under fixed effects is therefore not treated as evidence in favour of that model. The selection is based on the specification tests and the statistical properties of the estimator, not on which model produces the more favourable results.

The choice of weighting is handled separately in the robustness analysis. Removing the cross-section weights leaves the GPR and GEAR0 estimates significant but takes significance away from EXPO and ROA, so the weighting procedure has a substantive effect on two of the four direct estimates. The weighted model is kept as the main specification because the weights respond to the heteroskedasticity found in the diagnostics. To preserve transparency regarding the sensitivity of the results, the equivalent unweighted results are also reported.

There is also an important difference in the type of variation used by the two estimators. Fixed effects assigns a firm-specific intercept, and therefore absorbs time-invariant characteristics such as industry membership, business model or persistent cost structure differences. The estimates are therefore identified from

variation over time within the same firm. Random effects draws on two sources of variation, between firms and within firms. This distinction matters most for firm cost exposure, which is highly persistent and varies mainly between firms rather than within them. Geopolitical risk and energy exposure vary only over time and are unaffected by the choice.

3.4.5. Significance Testing

3.4.5.1. *t*-test

$$t = \frac{\hat{\beta}_j}{SE(\hat{\beta}_j)} \dots\dots\dots (3.17)$$

Where:

$\hat{\beta}_j$: the estimated regression coefficient for variable j
 $SE(\hat{\beta}_j)$: the standard error of that coefficient

The statistic divides the estimated coefficient by its own standard error, expressing the estimate in units of its own uncertainty. A coefficient of 0.5 with a standard error of 0.1 gives $t = 5$, five standard errors from zero and unlikely to have arisen by chance were the true value zero. The same coefficient with a standard error of 0.4 gives $t = 1.25$, entirely consistent with a true value of zero. Magnitude alone means nothing; it becomes interpretable only relative to precision. This is also where the standard error corrections discussed above become consequential, since any understatement of the denominator inflates the statistic and produces spurious significance.

Null hypothesis (H_0): $\beta_j = 0$, the variable has no significant relationship with ES5.

Alternative hypothesis (H_1): $\beta_j \neq 0$, the variable has a significant relationship with ES5.

The alternative is two-sided, consistent with the non-directional hypotheses developed in Chapter 2. The decision rule follows the p-value:

- $p \leq 0.10$: reject H_0 , significant at the 10 percent level (*)

- $p \leq 0.05$: reject H_0 , significant at the 5 percent level (**)
- $p \leq 0.01$: reject H_0 , significant at the 1 percent level (***)
- $p > 0.10$: fail to reject H_0 , no significant relationship

The p-value is the probability of observing a t-statistic at least as extreme as the one obtained, assuming the null is true. Failing to reject is not the same as establishing that the coefficient is zero; it indicates only that the data cannot distinguish it from zero at the chosen level. The level at which a coefficient is significant is therefore a property of its probability, not of its magnitude. A small coefficient estimated precisely can be significant at the one percent level while a larger one estimated imprecisely fails to clear ten percent, and the asterisks attached to a coefficient in the tables that follow record the probability bracket rather than the size of the estimate.

3.4.5.2. *F-test*

The F-test examines whether the independent variables jointly explain variation in ES5, testing the model as a whole (Ghozali, 2021). Every individual t-statistic may fall short of significance while the regressors jointly explain a meaningful share, which occurs when regressors are correlated and the data cannot attribute the explanation to any one. The statistic compares variation explained against variation unexplained, each scaled by its degrees of freedom.

Null hypothesis (H_0):

$$\beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = 0 \quad \dots\dots\dots (3.18a)$$

All independent variables are jointly unrelated to ES5.

Alternative hypothesis (H_1):

$$\text{At least one } \beta_i \neq 0 \quad \dots\dots\dots (3.18b)$$

At least one independent variable is related to ES5.

The decision rule follows the p-value of the F-statistic, with the same thresholds as the t-test.

3.4.5.3. *Coefficient of Determination*

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}} = 1 - \frac{\sum_{i=1}^n (Y_i - \hat{Y}_i)^2}{\sum_{i=1}^n (Y_i - \bar{Y})^2} \dots\dots\dots (3.19)$$

Where:

R^2 : the coefficient of determination

SS_{res} : residual sum of squares, the variation the model fails to explain

SS_{tot} : total sum of squares, the total variation in ES5 around its mean

Y_i : observed value of ES5 for observation i

$\hat{Y}_i$: value predicted by the regression

$\bar{Y}$: mean of the observed ES5 values

The construction compares two benchmarks. The total sum of squares measures how much ES5 varies around its own mean, the variation remaining if the mean were the only available prediction. The residual sum of squares measures what remains after fitting. Their ratio is the proportion unexplained, so subtracting from one gives the proportion explained. The measure is bounded between 0 and 1, where 0 means the model performs no better than predicting the sample mean and 1 means it reproduces every observation exactly. Adding any regressor can only raise R^2 or leave it unchanged, which is why adjusted R^2 is also reported: it penalises each additional parameter and can fall when a regressor adds nothing.

Interpretation must be calibrated to the setting. For panel models using macro-level predictors to explain firm-level tail risk, values of five to twenty percent are typical and acceptable, because much of any individual firm's tail risk is driven by idiosyncratic factors lying outside the model, including firm-specific news, ownership changes, index rebalancing and liquidity events. A low R^2 here indicates that firm-level tail risk is mostly idiosyncratic, which is a substantive finding about the market rather than a defect in the specification, and it does not undermine the individual coefficient estimates.

3.4.6. Robustness Checks

Several specification decisions could reasonably have been made otherwise, so each is tested by re-estimating under the alternative. The purpose is not to search for a favourable specification but to establish whether conclusions depend on

discretionary choices. Those carried out are reported in Chapter 4 whether or not they support the primary results.

- Alternative tail thresholds. ES1 and ES10 replace ES5, testing whether results depend on where the tail was cut.
- Symmetric GEAR. A variable built on the absolute change in the composite oil price replaces the censored form, showing directly what the $\text{Max}(0, \cdot)$ decision costs or contributes.
- Specification without the interaction, allowing the change in the main effect coefficients on adding it to be observed.
- Unweighted estimation, showing how far reported significance depends on the cross-section weighting.
- Random effects estimation, reported in parallel throughout because the Hausman test could not adjudicate.
- Leave-one-year-out estimation. Each year is dropped in turn, establishing whether any single year drives a reported relationship. Given that the macro coefficients rest on eleven annual observations, this is the most direct available test of the identification concern above.
- Exclusion of 2020, which combines an atypical tail-risk outcome with a low geopolitical risk reading and is the observation most likely to distort the GPR coefficient.

3.5. Data Collection, Software, and Tools

Estimation is carried out in EViews 13, which provides modules for panel estimation including Panel EGLS with cross-section fixed and random effects, cross-section weighting, and White diagonal standard errors. Expected shortfall is computed separately in Microsoft Excel from daily price data, with simple returns sorted ascending and the worst five percent averaged for each firm-year before the series is imported.

Daily stock prices were downloaded through the Bloomberg BDH function, retrieving PX_LAST for all 54 stocks across 2014 to 2025; the series begins in

2014 because the first daily return of 2015 requires the final 2014 close. Prices were cross-checked against IDX daily trading summaries for randomly selected dates.

Financial statement data were downloaded through the Bloomberg BDP function, retrieving Revenue, Cost of Goods Sold, Operating Income, Net Income, Total Assets and Return on Assets at annual frequency. Each was cross-checked against the audited statements published in company annual reports on the IDX website, with the audited figure treated as authoritative in cases of disagreement.

The Geopolitical Risk Index was downloaded from matteoiacoviello.com with the file timestamp recorded so the version can be identified; the series contains 4,473 daily observations from 1 January 2014 to 1 April 2026, from which calendar-year averages were computed. Oil price data were downloaded from Bloomberg at daily frequency for CO1, CL1 and MUC1 Comdty and averaged across the calendar year before combination. Middle East import share data were sourced from TradeMap.org using HS codes 270900 and 2710 disaggregated by country of origin.

The complete dataset was assembled in a single Excel workbook organised as a balanced panel, one row per firm-year and one column per variable, before import into EViews. Retaining it in this form allows every computed figure to be traced to its source series.

3.5.1. Data Verification

Three procedures were applied to the assembled panel before estimation.

Panel balance was verified by confirming that every firm holds exactly eleven annual observations and every year exactly 54 firms, giving 594 rows without gaps. No missing values were present in any series, so no imputation or listwise deletion was required and the estimation sample equals the full panel.

Internal consistency of the macro variables was verified next. GPR and GEAR0 are national series and must by construction take an identical value for every firm within a year, so each was checked by confirming that the number of distinct values

within each year equals one. An earlier version of the dataset failed this check for a single cross-section whose GEAR0 series was displaced by one period relative to the remaining firms; the misalignment was corrected before the results reported in Chapter 4 were produced. This class of error is silent in ordinary estimation output, since the regression runs and reports plausible coefficients regardless, and a within-year uniformity check is the only reliable way to detect it.

Range plausibility was assessed last. The minimum, maximum, mean and median of every series were inspected against economic expectation, and the composition of cost of goods sold reported under the Bloomberg field was examined for firmyears in which the ratio approached or exceeded unity, since a firm whose direct costs exhaust revenue is unusual and may indicate the field returns total operating cost rather than cost of goods sold alone. The outcome is reported with the descriptive statistics in Chapter 4.

CHAPTER 4

RESULTS AND DISCUSSION

4.1. Findings and Analysis

This chapter reports the empirical results obtained from the balanced panel of 54 non-financial Kompas 100 constituents observed over 2015 to 2025, a total of 594 firm-year observations. The presentation moves from descriptive statistics through the panel model selection procedure, the correlation and multicollinearity diagnostics, and finally to the regression estimates and hypothesis tests. Estimation was carried out in EViews 13 and every figure reported below derives from that output.

Two conventions apply throughout. Expected shortfall is expressed in absolute value, so a larger figure indicates more severe tail losses and a positive coefficient indicates that a variable accompanies a deeper left tail. And because the design is observational, the coefficients describe conditional associations rather than causal effects; the discussion is worded accordingly.

4.1.1. Descriptive Statistics

Table 4.1 summarises the distributional properties of every variable entering the model.

Table 4.1. Descriptive Statistics

Variable	Mean	Median	Maximum	Minimum	Std. Dev.
ES5	5.4269	5.2134	14.6673	0.3205	1.6895
GPR	109.1956	103.2213	152.9696	77.9835	23.7321
GEAR0	6.0846	0.0000	23.7965	0.0000	8.5146
EXPO	67.4545	69.6856	103.5035	12.7859	16.3438
ROA	7.0219	5.2645	55.7341	-31.8812	8.2627

Note: N = 594 firm-year observations (54 firms $\times$ 11 years). Source: EViews 13 output, processed by Author (2026)

The dependent variable averages 5.43, meaning that on its worst five percent of trading days the typical firm in the sample lost roughly 5.43 percent per day. The spread is considerable: ES5 runs from 0.32 to 14.67, so some firms endured tail days more than forty times as severe as others within the same market and the same eleven years. That dispersion is what the model attempts to explain, and had it been narrow there would have been little for the firm-level regressors to account for.

GPR ranges from 77.98 to 152.97 with a mean near 109. The low reading belongs to 2020, when pandemic coverage displaced geopolitical reporting, and the high to 2022 and the invasion of Ukraine.

GEAR0 requires comment because its median is zero. The $\text{Max}(0, \cdot)$ operator censors years in which the composite oil price fell, and within this sample that censoring binds in six of the eleven years. The variable therefore takes only five distinct non-zero values across the entire panel, and its standard deviation of 8.51 against a mean of 6.08 reflects that concentration at zero rather than continuous variation. This has consequences for statistical power that are taken up below.

EXPO averages 67.45 percent with a median slightly higher at 69.69, indicating mild left skew and a distribution weighted toward cost-intensive firms. The maximum of 103.50 percent, where direct costs exceeded revenue, occurs in lossmaking firm-years and is retained rather than trimmed, since such observations are economically meaningful rather than data errors. ROA spans -31.88 to 55.73 percent with a mean of 7.02 against a median of 5.26, the gap indicating that a minority of highly profitable firms lifts the average above the typical firm.

Every variable departs from normality, with Jarque-Bera probabilities of 0.0000 throughout. The pattern is expected rather than troubling. Financial return data are well documented to exhibit negative skewness and excess kurtosis, and that characteristic is precisely why a coherent tail measure was preferred to a variancebased one in the first place. With 594 observations the central limit theorem preserves asymptotic normality of the coefficient estimates, so inference remains valid despite non-normal residuals.

4.1.2. Panel Data Model Selection

Three tests were applied to choose among the Common Effect Model, the Fixed Effect Model and the Random Effect Model. Each is an independent pairwise comparison rather than a stage in a dependent sequence: Chow contrasts Common against Fixed, Hausman contrasts Fixed against Random, and the Lagrange Multiplier test contrasts Common against Random. Since none conditions on the outcome of another, the statistics are unaffected by the order of computation, and they are presented here in the conventional order.

4.1.2.1. Chow Test (Redundant Fixed Effects Test)

Table 4.2. Chow Test Result

Effects Test	Statistic	d.f.
Cross-section F	4.7083***	(53, 535)

*Note: *** significant at the 1 percent level. Source: EViews 13 output, processed by Author (2026)*

The null hypothesis holds that the Common Effect Model is appropriate; the alternative favours the Fixed Effect Model. With a probability of 0.0000 the null is rejected at the one percent level, indicating that firm-specific intercepts are jointly significant and that pooling all 594 observations under a single intercept is untenable. Firms differ systematically in their baseline level of tail risk.

4.1.2.2. Hausman Test (Correlated Random Effects Test)

Table 4.3. Hausman Test Result

Test Summary	Chi-Sq. Statistic	Chi-Sq. d.f.
Cross-section random	Not computable	5

Note: the statistic could not be validly computed and no value is reported; see below. Source: EViews 13 output, processed by Author (2026)

EViews returned a warning that the cross-section test variance is invalid, and the reason is visible in the coefficient comparison. The Hausman statistic requires the difference between the fixed effects and random effects coefficient covariance matrices to be positive definite. In this sample three of the five eigenvalues of that

difference are negative, which makes the statistic undefined rather than merely imprecise. No value or probability can therefore be interpreted, and the test is unable to adjudicate between the two specifications.

4.1.2.3. Lagrange Multiplier (LM) Test

Table 4.4. Breusch-Pagan Lagrange Multiplier Test Result

Test	Statistic	Chi-Sq. d.f.
Breusch-Pagan (cross-section)	147.0719***	1

*Note: *** significant at the 1 percent level. Source: EViews 13 output, processed by Author (2026)*

Here the null holds that the Common Effect Model is appropriate against an alternative favouring the Random Effect Model. The statistic of 147.0719 rejects the null decisively, confirming that the variance of the firm-specific error component differs significantly from zero and that panel-level effects exist.

4.1.2.4. Model Selection Outcome

Table 4.5. Summary of Panel Model Selection Tests

Test	Prob.	Points To	Conclusion
Chow (Redundant FE)	0.0000	Fixed Effect	Reject Common Effect Model
Lagrange Multiplier	0.0000	Random Effect	Reject Common Effect Model
Hausman	Invalid	Undetermined	Cannot adjudicate FEM vs REM

Source: Author (2026)

The three tests agree on one point and diverge on another. Both the Chow and Lagrange Multiplier tests reject the Common Effect Model, so pooled estimation is ruled out. What they cannot settle is the choice between fixed and random effects, and the Hausman test, which exists precisely to break that tie, was unavailable.

Both remaining specifications are therefore admissible on the evidence, and the choice rests on judgement. This study adopts the Fixed Effect Model with crosssection weights as its primary specification for three reasons. The Chow result is unambiguous in establishing that firm-specific intercepts belong in the model. The fixed effects estimator makes no assumption about correlation between firm effects and regressors, so it remains consistent whether or not such correlation exists, whereas random effects is consistent only under an assumption the Hausman

test could not verify. And cross-section weighting addresses the heteroskedasticity identified below.

Because the selection was not clean, random effects estimates are reported alongside the fixed effects results rather than set aside, so that the consequence of the choice remains visible to the reader. It should also be recorded that the fixed effects specification produces more significant coefficients than random effects; that outcome is not the basis for the choice, which rests on the Chow result and the consistency property described above, and the same specification would be reported had fewer coefficients attained significance.

4.1.3. Correlation Matrix

Table 4.6. Pearson Correlation Matrix

	EXPO	GPR	GEAR0	ROA	GEAR0*EXPO
EXPO	1.0000	-0.0004	-0.0262	-0.2575	0.1375
GPR	-0.0004	1.0000	-0.4913	0.0643	-0.4728
GEAR0	-0.0262	-0.4913	1.0000	0.0022	0.9597
ROA	-0.2575	0.0643	0.0022	1.0000	-0.0478
GEAR0*EXPO	0.1375	-0.4728	0.9597	-0.0478	1.0000

Source: EViews 13 output, processed by Author (2026)

Applying the conventional threshold of 0.80, no pair among the independent and control variables approaches it, the largest in absolute terms being -0.4913 between GPR and GEAR0. The single value above the threshold is the 0.9597 between GEAR0 and the interaction term, which reflects the arithmetic dependence of a product on the variable from which it is formed rather than redundancy in the data. For that reason the interaction is excluded from the variance inflation diagnostic reported below, and its consequence for the precision of the GEAR0 estimate is taken up with the regression results.

The correlation of -0.4913 between GPR and GEAR0 carries substantive interest. It is moderate in magnitude, which supports retaining both variables: they overlap without duplicating one another, and a coefficient approaching unity would have indicated redundancy. The negative sign, however, deserves attention. Examining

the annual series shows that the three highest-tension years in the sample carry a mean GEAR0 of 1.04, while the three lowest-tension years carry a mean of 13.23. GEAR0 registers zero in both 2024 and 2025, which rank third and second respectively for geopolitical tension across the whole window.

The explanation lies in the censoring. Geopolitical tension climbed through the latter part of the observation period while composite oil prices declined, so those years floor to zero under the $\text{Max}(0, \cdot)$ specification. This is not an error in construction but a consequence of oil prices and geopolitical tension decoupling in the back half of the window, and it reinforces the discriminant-validity qualification noted in the methodology: GEAR0 measures energy cost pressure weighted by import dependence, and its geopolitical content is narrower than the variable name alone suggests. The symmetric GEAR Absolute specification, which retains non-zero values in all eleven years, is examined among the robustness checks.

The remaining correlations are unremarkable. ROA and EXPO correlate at -0.2575 , consistent with cost-intensive firms being less profitable, which is precisely the confounding that motivated including ROA as a control. No other pair exceeds 0.14 in absolute value.

4.1.3.1. Multicollinearity Test

Variance inflation factors were computed for the regressors entering the estimated equation.

Table 4.7. Variance Inflation Factors

Variable	R ² (auxiliary)	VIF	Assessment
EXPO	0.0670	1.0718	No multicollinearity
GPR	0.2456	1.3256	No multicollinearity
GEAR0	0.2428	1.3207	No multicollinearity
ROA	0.0714	1.0769	No multicollinearity

Note: computed on the specification excluding the interaction term. Source: EViews 13 output, processed by Author (2026)

Variance inflation factors were computed for the three independent variables and the control variable. All four fall between 1.07 and 1.33, far below the threshold of 10 above which multicollinearity is conventionally treated as severe (Hair, Black,

Babin and Anderson, 2014). The null of no multicollinearity is therefore not rejected, and the auxiliary R-squared values, ranging from 0.067 to 0.246, confirm that no regressor is substantially explained by the others.

The interaction term is excluded from this diagnostic. A product term is arithmetically bound to the variables from which it is constructed, so it registers elevated inflation by construction rather than as evidence of redundancy in the data; including it would flag a condition that is certain in advance of observing any data and that carries no diagnostic content. The consequence of that dependence for the precision of the GEAR0 estimate is addressed directly in the regression results below.

4.1.3.2. Heteroskedasticity Test

Table 4.8. Heteroskedasticity Test Results

Test	Obs*R ²	d.f.	Prob. Chi-Square
Breusch-Pagan	17.4587***	5	0.0037
White (squared terms)	35.4863***	10	0.0001

*Note: *** significant at the 1 percent level, indicating rejection of the null of homoskedasticity.*

Source: EViews 13 output, processed by Author (2026)

Both tests reject the null of constant error variance, which is the expected outcome for a financial panel in which larger firms typically display lower return volatility than smaller ones and market conditions vary across years. Heteroskedasticity leaves the coefficient estimates unbiased but renders conventional standard errors invalid, so every reported specification applies White diagonal robust standard errors, and the primary model additionally applies cross-section weighting, which weights each firm by the inverse of its estimated residual variance.

4.1.4. Hypothesis Testing

4.1.4.1. Regression Model Results

Following the model selection reported above, the fixed-effects specification was estimated with cross-section weights. Cost exposure enters both as an independent variable and, in the interaction reported subsequently, as the moderating variable; following standard practice for a continuous moderator, the product term rather

than a second standalone entry carries the moderation, so the direct specification below contains one entry per variable. Two specifications are reported and both carry equal weight in the discussion that follows. The direct specification in Table 4.9 estimates the three main associations on their own. The full specification in Table 4.10 adds the product term and is the primary model, since it is the only one in which all four hypotheses are tested together. The direct specification is shown in Table 4.9.

Table 4.9. Regression Results Without Moderating Variable

Variable	Coefficient	Std. Error	t-Statistic	Prob.
EXPO	-0.016571*	0.009550	-1.735300	0.0833
GPR	-0.014439***	0.002278	-6.338200	0.0000
GEAR0	-0.049600***	0.005115	-9.697700	0.0000
ROA	-0.020214*	0.011664	-1.733000	0.0837

*Notes: Panel EGLS, cross-section fixed effects with cross-section weights, White diagonal standard errors. $R^2 = 0.4797$, Adjusted $R^2 = 0.4244$, F-statistic = 8.6708 (Prob. 0.0000). N = 594. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Source: EViews 13 output, processed by Author (2026)*

All four variables reach significance and all carry negative coefficients, but they do not reach it at the same level, and the difference matters for how they are read. GPR and GEAR0 return probabilities of 0.0000 and are significant at the one percent level. EXPO and ROA return probabilities of 0.0833 and 0.0837, which sit above 0.05 and below 0.10, so under the decision rule of Section 3.4.5.1 they are significant at the ten percent level only and carry a single asterisk. Cost-intensive firms, firms operating in years of heightened geopolitical tension, firms facing greater Middle East-weighted energy cost pressure, and more profitable firms each exhibit shallower left tails than their counterparts, with the evidence for the first and last of these appreciably weaker than for the two macro variables. The level of significance attaches to the probability, not to the size of the coefficient. The EXPO estimate of -0.0166 is not weakly significant because it is small; it is weakly significant because its standard error of 0.0096 leaves the t-statistic at -1.74, whereas GEAR0 at -0.0496 with a standard error of 0.0051 stands nearly ten standard errors from zero.

4.1.4.2. Regression Model With Moderating Variable

To test the fourth hypothesis, the model was re-estimated with the GEAR0-byEXPO interaction term added, so that cost exposure could act as a moderator of the energy-adjusted risk channel. Table 4.10 reports this specification.

Table 4.10. Regression Results With the Interaction Term

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	8.585834	0.777597	11.04149	0.0000
EXPO	-0.016904*	0.009626	-1.756198	0.0796
GPR	-0.014470***	0.002281	-6.344918	0.0000
GEAR0	-0.058369***	0.016670	-3.501505	0.0005
ROA	-0.019768*	0.011681	-1.692276	0.0912
GEAR0*EXPO	0.000136	0.000248	0.548891	0.5833

Notes: Panel EGLS, cross-section fixed effects with cross-section weights, White diagonal standard errors. $R^2 = 0.4792$, Adjusted $R^2 = 0.4227$, F-statistic = 8.4872 (Prob. 0.0000). $N = 594$. GEAR0 here is evaluated at EXPO equal to zero and its coefficient is therefore not directly comparable with Table 4.9. Source: EViews 13 output, processed by Author (2026)

Adding the interaction leaves the direct associations substantially intact: EXPO, GPR and ROA retain their signs, and the interaction itself is positive but does not attain significance (coefficient 0.000136, $p = 0.5833$). Table 4.10 is the primary model, since it is the specification in which all four hypotheses are tested together. The direct specification of Table 4.9 is reported alongside it and carries equal weight, because it estimates the three main associations without the arithmetic dependence that the product term introduces. The two agree on every conclusion. Each coefficient keeps the same sign and the same significance level in both, GPR and GEAR0 at one percent ($p = 0.0000$ and 0.0005 in Table 4.10) and EXPO and ROA at ten percent ($p = 0.0796$ and 0.0912), so no hypothesis outcome depends on which specification is read first.

One feature of the interacted specification warrants comment. When GEAR0 is entered alongside its own product with EXPO, its standard error rises from 0.005115 in Table 4.9 to 0.016670 in Table 4.10, roughly a threefold increase, while the coefficient itself moves only modestly. This reflects the arithmetic dependence between the product term and the variable from which it is formed rather than any change in the underlying relationship, and it is the reason both specifications are

reported rather than one. In the full specification the GEAR0 coefficient is evaluated at EXPO equal to zero, a point lying outside the observed range of 12.79 to 103.50, which is what inflates its variance and why its magnitude is not directly comparable with Table 4.9. In the direct specification, where GEAR0 is not carried in a product term, it is estimated as -0.0496 with a standard error of 0.0051. Both estimates are negative and significant at the one percent level, so the conclusion drawn for the second hypothesis is the same under either reading.

4.1.4.3. Panel Regression Model Summary Result

Table 4.11. Comparison of Panel Estimation Models

Variable	CEM	FEM	REM	FEM + Weights (Primary)
<i>Panel A: Direct specification (Table 4.9)</i>				
EXPO	0.0074*	-0.0022	0.0036	-0.0166*
GPR	-0.0086***	-0.0100***	-0.0094***	-0.0144***
GEAR0	-0.0444***	-0.0469***	-0.0458***	-0.0496***
ROA	-0.0356***	0.0067	-0.0101	-0.0202*
R ²	0.0837	0.3750	0.0597	0.4797
Observations	594	594	594	594
<i>Panel B: Full specification with product term (Table 4.10)</i>				
C	6.2784***	6.8208***	6.4552***	8.5858***
EXPO	0.0090*	-0.0009	0.0052	-0.0169*
GPR	-0.0086***	-0.0100***	-0.0094***	-0.0145***
GEAR0	-0.0255	-0.0275	-0.0268	-0.0584***
ROA	-0.0358***	0.0061	-0.0107	-0.0198*
GEAR0*EXPO	-0.0003	-0.0003	-0.0003	0.0001
R ²	0.0842	0.3755	0.0604	0.4792
Adjusted R ²	0.0764	0.3078	0.0524	0.4227
Observations	594	594	594	594

*Note: *** significant at 1%, ** at 5%, * at 10%. CEM, FEM and REM estimated with White robust standard errors and no weighting. Panel A reports the direct specification of Table 4.9; Panel B adds the product term and corresponds to Table 4.10. The final column of each panel applies cross-section weights. Source: EViews 13 output, processed by Author (2026)*

Table 4.11 places the four estimations side by side. Two results hold whatever the estimator and whether or not the weights are applied. GPR is negative and

significant at the one percent level in every column of both panels. GEAR0 is likewise negative and significant at the one percent level throughout Panel A, at -0.0444 under common effects, -0.0469 under fixed effects and -0.0458 under random effects against -0.0496 in the weighted model, so the energy estimate does not owe its standing to the weighting. EXPO and ROA behave differently. EXPO is insignificant under unweighted fixed and random effects and turns positive at the ten percent level under common effects; ROA is significant only under common effects and under the weighted model. Panel B shows the same pattern with one addition. Once the product term enters the unweighted estimations, the GEAR0 main effect falls to between -0.0255 and -0.0275 and loses significance, while the product term itself is insignificant in every column. This is the arithmetic dependence noted under Table 4.10, where entering GEAR0 beside its own product with EXPO roughly trebled its standard error, and it should not be read as a change in the underlying relationship: the direct estimates in Panel A and the unweighted leave-one-year-out results in Table 4.14 keep GEAR0 at the one percent level in every case.

The dependence on the weighting scheme should be stated plainly. Cross-section weighting responds to the heteroskedasticity diagnosed above and is not a discretionary preference, but two of the four direct estimates, those on EXPO and ROA, would not attain significance without it. The two macro estimates do not share that dependence. The unweighted results are therefore presented in full, and the conclusions drawn in the following chapter distinguish the estimates that rest on the weighting from those that do not.

It is also worth noting that the R-squared of approximately 0.48, which both specifications return, is elevated by the inclusion of 54 firm-specific intercepts rather than by the explanatory power of the regressors themselves. The Common Effect and Random Effect specifications, which do not absorb firm heterogeneity in the same way, return values below 0.09. The correct reading is that firm-level tail risk in this market is largely idiosyncratic, and that the macro and firm-level regressors explain a modest share of it, which is entirely usual for models of this type.

Table 4.12. Summary of Hypothesis Testing

Hyp.	Relationship	Coefficient	Prob.	Direction	Result
H1	GPR → ES5	-0.0145	0.0000	Negative	Supported
H2	GEAR0 → ES5	-0.0584	0.0005	Negative	Supported
H3	EXPO → ES5	-0.0169	0.0796	Negative	Supported (10% level)
H4	GEAR0 × EXPO → ES5	0.0001	0.5833	Positive	Not supported
—	ROA → ES5 (control)	-0.0198	0.0912	Negative	Control variable

Note: coefficients from the full specification of Table 4.10. The direct specification of Table 4.9 returns -0.0144 for GPR, -0.0496 for GEAR0, -0.0166 for EXPO and -0.0202 for ROA, with identical signs and identical significance levels throughout. Source: Author (2026)

Table 4.12 consolidates the hypothesis test outcomes. The estimated model may be written as follows:

$$ES5 = 8.5858 - 0.0169 EXPO - 0.0145 GPR - 0.0584 GEAR0 + 0.0001 (GEAR0 \times EXPO) - 0.0198 ROA$$

4.2. Robustness Checks

Several specification decisions could reasonably have been taken differently, and each was tested by re-estimating the model under the alternative. The battery is run on the direct specification, so that the effect of each alternative on the three main associations can be read without the arithmetic dependence that the product term introduces. Results are reported whether or not they support the main findings.

Table 4.13. Robustness Check Results

Variable	ES5	ES1	ES10	Excl. 2020
EXPO	-0.0166*	-0.0092	-0.0068	-0.0143
GPR	-0.0144***	-0.0107***	-0.0130***	-0.0148***
GEAR0	-0.0496***	-0.0631***	-0.0394***	-0.0462***
ROA	-0.0202*	-0.0168	-0.0113	-0.0193
Observations	594	594	594	540

*Note: *** significant at 1%, ** at 5%, * at 10%. Columns report alternative tail thresholds and exclusion of the pandemic year, estimated on the direct specification of Table 4.9. Panel EGLS, cross-section fixed effects with cross-section weights, White diagonal standard errors. Source: Author (2026)*

Two findings emerge. GEAR0 retains its sign and its significance at the one percent level in every specification tested. Its magnitude behaves sensibly as the tail is

varied, strengthening from -0.0394 at the ten percent threshold to -0.0496 at five percent and -0.0631 at one percent, which is what would be expected if the relationship genuinely concerns extreme outcomes rather than general downside. Each of those three estimates carries a probability below 0.01, so the change in magnitude across the thresholds is a change in the size of the association and not in the level at which it is significant. It also survives exclusion of the pandemic year, at -0.0462. This is the most robust result in the study.

The GPR result is more fragile, and the two exercises that bear on it need to be read as a sequence rather than as a single comparison, because they differ in two respects at once. Table 4.13 removes 2020 but keeps the cross-section weights; Table 4.14 removes 2020 and drops the weights as well. Four estimates already reported allow the two conditions to be separated. With the weights in place, GPR is -0.0144 on the full sample (Table 4.9) and -0.0148 with 2020 excluded (Table 4.13), significant at the one percent level in both. With the weights removed and all eleven years retained, the fixed effects estimate in Panel A of Table 4.11 is -0.0100, again at the one percent level. Only when the weights are removed and 2020 is dropped together does the coefficient fall to -0.0003 and lose significance (Table 4.14). Neither the weighting choice nor the pandemic year is sufficient on its own; the collapse requires both. A leave-one-year-out sequence run under the weighted specification for every year, rather than for 2020 alone, would complete the comparison and was not estimated here. On the evidence available, the GPR result is treated as sensitive to the joint choice of estimation weights and sample window, and as suggestive rather than established throughout. Table 4.14 examines the year dimension directly by dropping each year in turn.

Table 4.14. Leave-One-Year-Out Estimation

Year Dropped	GPR Coef.	GPR Prob.	GEAR0 Coef.	GEAR0 Prob.	Note
2015	-0.0091	0.0018	-0.0434	0.0000	
2016	-0.0100	0.0003	-0.0505	0.0000	
2017	-0.0106	0.0001	-0.0478	0.0000	
2018	-0.0120	0.0001	-0.0544	0.0000	

2019	-0.0104	0.0003	-0.0433	0.0000	
2020	-0.0003	0.9280	-0.0251	0.0015	GPR loses significance
2021	-0.0097	0.0007	-0.0526	0.0000	
2022	-0.0067	0.0991	-0.0441	0.0000	
2023	-0.0097	0.0005	-0.0533	0.0000	
2024	-0.0099	0.0006	-0.0472	0.0000	
2025	-0.0183	0.0000	-0.0474	0.0000	

*Note: direct specification of Table 4.9, fixed effects, unweighted, White robust standard errors.
Source: Author (2026)*

Removing any year other than 2020 leaves the GPR coefficient between -0.0067 and -0.0183 and significant throughout, while GEAR0 remains negative and significant at the one percent level in every case, including when 2020 is dropped. Removing 2020 alone collapses the GPR coefficient to -0.0003. The reason is visible in the annual series: 2020 pairs the lowest GPR reading in the sample, 77.98, with the highest mean ES5, and that combination is a pandemic effect rather than a geopolitical one. Because both macro variables are identified from only eleven annual observations, a single unusual year exerts considerable leverage, and this is exactly the identification constraint anticipated in the methodology. The GPR finding should accordingly be treated as suggestive rather than firmly established, whereas the GEAR0 finding is not sensitive to any single year.

A further specification announced in Chapter 3 was also estimated. Replacing GEAR0 with GEAR Absolute, which retains non-zero values in all eleven years rather than censoring falling-price years to zero, returns a coefficient of -0.0135 with a probability of 0.1864. The energy variable does not attain significance under the symmetric construction, while GPR remains significant at the five percent level. This indicates that the association is driven specifically by years of rising import-weighted energy cost rather than by the underlying oil price series in both directions, which is consistent with the asymmetric transmission mechanism motivating the $\text{Max}(0, \cdot)$ construction. The point is noted as a boundary on interpretation rather than a weakness, since the asymmetric form is the theoretically intended measure.

4.3. Discussion

4.3.1. Geopolitical Risk and Tail Risk (H1)

Geopolitical risk carries a significant negative association with expected shortfall, so periods of elevated global tension coincided with shallower rather than deeper left tails among Indonesian equities. The sign runs contrary to the intuitive expectation drawn from Caldara and Iacoviello (2022), who document capital moving away from emerging economies during such periods, and from Bouras et al. (2019), who find geopolitical risk raising emerging market volatility.

Two features of the Indonesian setting bear on the divergence. Neither is tested by the regression, which carries no variable for trade routing or for the commodity composition of revenue, so both are offered as readings of the setting that remain open to further work rather than as findings.

The first is geographical. The escalations that lifted GPR in the later sample years, the invasion of Ukraine in 2022, the Gaza conflict from late 2023 and the Houthi attacks on Red Sea shipping that followed, closed or degraded the Black Sea and the Asia to Europe corridor through Bab el-Mandeb and Suez. Transits through the Suez Canal fell by more than two fifths from their 2023 peak within months of the attacks beginning, and container services were diverted around the Cape of Good Hope at a cost of ten days or more per voyage (UNCTAD, 2024; Notteboom, Haralambides and Cullinane, 2024). Indonesia's export trade runs elsewhere. Across the sample period its principal buyers were China, the United States, India, Japan and the ASEAN members; in 2023 China alone took a quarter of total exports while the European Union took well under a tenth (Badan Pusat Statistik, 2024), and for the mineral fuels that dominate the commodity issuers in the Kompas 100, China and India together accounted for more than half of shipments in 2024 (International Trade Centre, 2026). Those destinations are reached through the South China Sea, the archipelagic straits and the Indian and Pacific oceans, none of which passes the chokepoints where the disruptions occurred. Freight rates on Pacific routes rose in sympathy (Notteboom et al., 2024), but the physical movement of Indonesian exports to their main markets was not interrupted. A rise

in global tension of the kind GPR records therefore reached the listed sector mainly through prices rather than through trade disruption. The Middle East, which supplies part of Indonesia's petroleum imports, is the exception, and isolating that single channel is what GEAR0 was built to do. Whether this routing position contributes to the negative coefficient cannot be settled from the present estimates; it is recorded here as a feature of the setting consistent with the sign observed.

The second is compositional, and it turns on a point the standard framing of Indonesia as an oil importer tends to obscure. The country is also a producer, and crude is not the whole of its exposure in either direction. Coal miners, metal producers, energy services companies and integrated commodity groups make up much of the non-financial Kompas 100, and for those issuers a geopolitically driven rise in commodity prices lands on the revenue line. Gas, coal, nickel and palm oil move with the same tension that moves crude, so a net importer of oil can still be, in the aggregate of its listed corporate sector, a beneficiary of energy stress. Japan and South Korea produce very little of what they consume, which is one reading of why the March 2026 declines there ran in one direction only while the Indonesian market's did not. The reading is plausible, but the pooled coefficient does not distinguish commodity producers from the rest of the sample, and the separate estimation that would test it is left to future work.

Zaremba et al. (2022) supply a third possibility that sits alongside these rather than replacing them. Vivid geopolitical headlines lead investors to over-weight bad outcomes, and at annual frequency an initial overreaction followed by correction inside the same observation year would leave little net effect on measured tail severity, or leave a firm looking more resilient across the year as a whole.

The qualification established above must nonetheless be carried forward. This coefficient rests on eleven annual observations and does not survive the exclusion of 2020 under the unweighted estimation.

4.3.2. Geopolitical Energy-Adjusted Risk and Tail Risk (H2)

GEAR0 is negative and significant at the one percent level, so the years carrying the heaviest Middle East-weighted energy cost pressure were the years of

shallowest tails. That is the opposite of what the literature predicts. Kilian and Park (2009) show supply-driven oil price increases pushing equity returns down in oilimporting economies, and Indonesia has been a net crude importer since 2004, so a positive coefficient was the expected result.

The compositional argument of Section 4.3.1 offers the more natural reading here, and it would apply with more force to GEAR0 than to GPR, because GEAR0 isolates the energy channel where GPR does not. It remains an interpretation. Sukarno and Setiawati (2020) found the same division at sector level on IDX data, with mining responding positively to oil price increases and every other sector negatively, and a pooled coefficient weights those two groups by their presence in the index; in the non-financial Kompas 100 the first group is large. Whether that composition accounts for the sign observed here is a question the present specification does not put, since no subgroup was estimated separately, and Section 5.4.3 places that test first among the extensions.

One caveat travels with the result. GEAR0 is zero in six of eleven years, and two of those zeros fall among the three highest-tension years in the window. The variable is measuring energy cost pressure weighted by import dependence, not geopolitical intensity, and H2 should be read as a statement about the first.

4.3.3. Firm Cost Exposure and Tail Risk (H3)

EXPO is negative and significant at the ten percent level in the primary specification, with a probability of 0.0796 in the full specification and 0.0833 in the direct one. Under the decision rule set out in Section 3.4.5.1 a probability between 0.05 and 0.10 clears only the loosest of the three thresholds: an estimate this far from zero would arise by chance roughly eight times in a hundred if the true coefficient were zero, which is why the result carries a single asterisk in Tables 4.9 and 4.10 and is treated here as tentative. Chapter 2 left the sign open because the ratio admits two readings that predict opposite things, one treating a large direct cost base as a thin cushion and the other as expense that can be reduced when trading conditions worsen. The negative sign is compatible with the second reading. It does not settle the matter, for reasons set out below.

The mechanism is the one Lev (1974) and Mandelker and Rhee (1984) described, arriving at an unfamiliar destination because the accounting has to be translated first. Their result concerns operating leverage: a firm carrying heavy fixed costs cannot shed expense when revenue falls, so earnings swing harder than sales and systematic risk rises. Cost of goods sold is predominantly variable while selling and administrative expense is predominantly fixed. A high COGS to sales ratio therefore describes a firm with relatively low operating leverage, which is the reverse of what the ratio looks like on first reading. Take a mining contractor whose costs are diesel, wages and equipment hire. When volumes fall, most of that falls with them. The firm defends its margin, its earnings hold up better than a competitor carrying the same revenue against depreciation and lease obligations, and its returns stay off the days that fill the left tail. Read that way, a negative coefficient is what the operating leverage literature would lead one to expect, and the estimate is consistent with it.

Two qualifications apply, and together they limit the reading to compatibility rather than confirmation. Significance holds only at the ten percent level, and Table 4.11 shows the estimate does not survive removal of the cross-section weights: it is insignificant under fixed and random effects and turns positive under common effects. Nor is it stable across the alternative tail thresholds of Table 4.13, where it loses significance at the one and ten percent tails and when 2020 is excluded. The fourth hypothesis also failed, so what the evidence supports is a statement about cost intensity in general and not about firms being differentially vulnerable to oil shocks. The cost flexibility reading is consistent with the sign; the margin thinness reading is not ruled out by evidence of this strength, and a fuel-specific cost variable would be needed to separate the two. The interpretive rule set out in the methodology applies here directly.

4.3.4. The Moderating Role of Cost Structure (H4)

The GEAR0 by EXPO interaction is positive but insignificant (0.000136, $p = 0.5833$), so H4 is not supported. Cost structure does not measurably change the association between energy-adjusted risk and tail risk in this sample, and no

moderation is claimed for it. The sign is still worth reporting, since it points where theory said it would even though the estimate cannot carry the claim. GEAR0 is associated with shallower tails on average, and a positive interaction, had it been precisely estimated, would mean that association weakens as a firm's direct cost intensity rises. Where would that be expected? In firms whose variable costs are heavily fuel and energy: a mining contractor burning diesel in heavy equipment, a logistics operator buying road and marine fuel, a cement producer running kilns. For those firms the cost side of an oil price increase is real. It works against the revenue benefit the commodity producers enjoy, so a larger cost base could offset part of the tail-risk compression that accompanies GEAR0, and the sign sits where that reasoning says it should. The estimate is far too imprecise to confirm that any such offset operates.

The imprecision has a source worth naming, because it changes what the null is evidence of. GEAR0 measures geopolitically driven cost pressure imperfectly, EXPO measures oil vulnerability imperfectly, and the product inherits the error in both. Classical measurement error attenuates coefficients toward zero, which means the test was working against a finding of moderation before any data were seen. Two possibilities remain open and this design cannot separate them: cost structure does not condition the energy relationship, or it does and the noise in the components has buried the estimate. Treating the result as evidence against moderation would claim the first when only the pair has been shown.

A variable isolating the fuel component of cost of goods sold would put the question properly. Section 3.3.4.1 explains why it cannot be built consistently across 54 firms and eleven reporting periods from Indonesian disclosure, and Section 5.4.3 keeps it on the list.

4.3.5. Profitability as a Control Variable

ROA is negative and significant at the ten percent level, with a probability of 0.0912 in the full specification and 0.0837 in the direct one, which is what the financial slack mechanism of Novy-Marx (2013) and Fama and French (2015) predicts: firms earning more on their asset base hold more internal cash flow and have more

room before bad conditions force financing decisions. No hypothesis was attached to it, since it sits in the model to absorb cross-firm variation in profitability that would otherwise contaminate the cost structure estimate, a concern the -0.2575 correlation between ROA and EXPO shows was justified.

The result echoes Dwitanto et al. (2026), who included ROA in the same capacity in an Indonesian firm-level expected shortfall model and did not find it significant. Here it attains marginal significance under the weighted specification but not under the unweighted one, which places it in a similar position.

4.3.6. Summary of Findings

Three of the four hypotheses receive support, though of uneven strength, and all three estimates carry negative signs; the moderation hypothesis is not supported. The picture that emerges is of an Indonesian equity market whose tail risk moved inversely with both geopolitical tension and Middle East-weighted energy cost pressure across 2015 to 2025. The commodity orientation of the listed universe, and Indonesia's position as a producer as well as an importer of energy, offer the most plausible reading of that pattern, but neither is tested by the estimates reported here, and both are carried into Chapter 5 as interpretations rather than as results.

Of the two macro findings, GEAR0 is the more secure. It holds its sign and its significance at the one percent level across every alternative tail threshold, across exclusion of the pandemic year, and across the leave-one-year-out estimation, whereas the GPR result depends on the single year 2020 and should be treated as provisional. This is a fortunate outcome, since GEAR0 is the variable constructed for this study and the demonstration that it carries explanatory content beyond the established geopolitical index is the principal empirical contribution. The weaker GPR result correspondingly limits what can be said about the general relationship between geopolitical tension and Indonesian equity tail risk. Both are reported as found.

CHAPTER 5

CONCLUSION AND RECOMMENDATIONS

5.1. Conclusion

This study asked whether geopolitical risk, energy exposure and firm cost structure are associated with the equity tail risk of Indonesian non-financial firms. Two separate literatures pointed toward the question but neither answered it. Studies of geopolitical risk show that rising tension moves emerging market equities, but they work at index level, so they cannot say which firms in a market are most exposed. Studies of cost structure link it to systematic risk, but they measure risk with beta or crash risk, not with a coherent tail measure. Neither literature offered a way to express how much of a global energy shock a particular economy has to absorb.

The study used a balanced panel of 54 non-financial Kompas 100 constituents from 2015 to 2025, giving 594 firm-year observations. Estimation was by Panel EGLS with cross-section fixed effects and cross-section weights. The four research questions are answered below.

1. Geopolitical risk is negatively associated with expected shortfall in the primary specification. The coefficient is -0.0145 in the full specification and -0.0144 in the direct specification, significant at the one percent level in both, so in years of higher global tension Indonesian firms recorded shallower left tails rather than deeper ones. The result is sensitive to the pandemic year and to the estimation choice taken together. It survives the exclusion of 2020 under the weighted specification and survives the removal of the weights on the full sample, but when both are changed at once the coefficient falls to -0.0003 and loses significance. The negative association is therefore reported as suggestive rather than settled. Indonesia's production of oil alongside its imports, and the commodity linkage of much of the listed universe, offer a plausible reading of the sign, since a geopolitically driven price increase reaches those firms as revenue; that reading was not tested directly and is left open.

2. Geopolitical energy-adjusted risk is negatively associated with expected shortfall, and this is the most consistent result in the study. The coefficient is -0.0584 in the full specification and -0.0496 in the direct specification, significant at the one percent level in both. Within the asymmetric specification used here, the result holds at both alternative tail thresholds, survives exclusion of the pandemic year, remains significant when any single year is dropped, and keeps its sign and significance under common, fixed and random effects without cross-section weights. The magnitude also grows as the tail narrows, from -0.0394 at the ten percent threshold to -0.0631 at the one percent threshold, which is what would be expected if the relationship concerns extreme outcomes. Two boundaries apply. The finding is a statement about the $\text{Max}(0, \cdot)$ construction, since the symmetric GEAR Absolute variant does not attain significance, and it says nothing about which firms meet energy cost pressure as revenue and which as cost, because no subgroup was estimated. Whether the revenue channel outweighs the cost channel in this market is a question the present design leaves to industry-level work.

3. Firm cost exposure is negatively associated with expected shortfall, with a coefficient of -0.0169 in the full specification and -0.0166 in the direct specification, at the ten percent level in both ($p = 0.0796$ and 0.0833). This result is tentative. The probability sits above 0.05 and below 0.10, so the estimate clears only the loosest of the three thresholds set in Chapter 3, and it loses significance when the cross-section weights are removed and at the alternative tail thresholds, so its standing depends on the primary specification. Theory alone could not settle the direction of the relationship, since two readings of the variable predicted opposite signs, and the negative sign is compatible with the cost flexibility reading: operating activity bears on variable costs more directly than on fixed ones, so when demand falls production can be cut back and raw materials, freight and contracted labour fall with it, while leases and depreciation do not. Evidence of this strength does not rule the margin thinness reading out, and the question remains open.

4. The interaction between geopolitical energy-adjusted risk and firm cost exposure is positive at 0.000136 with a p-value of 0.5833, so the fourth hypothesis is not supported and no moderation was detected in this sample. The positive sign is in the direction the proposed mechanism would predict, since for a firm whose variable cost base is heavily fuel the benefit that comes with energy cost pressure ought to be smaller than for a commodity producer meeting the same shock on its revenue line, but the estimate is nowhere near precise enough to establish that any such pattern exists. Section 4.3.4 sets out why the test was working against a finding from the outset, given the measurement error carried into the product by both of its components. What the null shows is that this design cannot identify moderation. It does not show that none is there, and it does not show that any is.

Overall, the empirical results indicate that Indonesian equity tail risk was negatively associated with both geopolitical risk and Middle East-weighted energy cost pressure during the sample period, the second association being the more secure of the two. One possible explanation is the commodity orientation of the Indonesian listed market. Indonesia's role as both an energy producer and an importer may also matter, since higher energy prices can create additional revenue for some firms while increasing costs for others. Both explanations are consistent with the estimates and neither is established by them; the industry-level comparison that would test them is set out in Section 5.4.3.

5.2. Limitations

The findings should be considered in light of several limitations.

1. The study uses an observational design. Panel regression can identify associations between variables after controlling for the other factors included in the model, but it does not by itself establish causality. There is no random assignment, instrumental-variable strategy, or discontinuity that would allow unobserved sources of bias to be ruled out. The estimated relationships

should therefore be interpreted as conditional associations rather than causal effects.

2. Both macro variables take one value per year. Although the panel contains 594 firm-year observations, the coefficients on the two annual macro variables are identified from only eleven annual observations. Year fixed effects cannot be included alongside these variables because they would be perfectly collinear with the annual observations. As a result, unobserved national factors that change over time and are related to both geopolitical conditions and equity tail risk may remain in the estimates. The leave-one-year-out analysis provides an indication of how sensitive the geopolitical risk coefficient is to this limited time variation.
3. GEAR is built from an aggregate national figure rather than from firm-level data. The measure uses Indonesia's national import share rather than firm-level information on sourcing or energy procurement. In this way, firms are assigned the same GEAR value in any given year, despite their supply chains, production processes, and dependence on energy inputs potentially differing widely. The measure therefore reflects the country's exposure to the external energy shock, rather than the precise exposure of each individual firm.
4. GEAR is designed specifically around geopolitical developments affecting the Middle East. Global oil prices can also respond to events in other regions. The Russian invasion of Ukraine in 2022, for example, contributed to a sharp movement in oil prices, and Russia is itself a major oil producer. Events of this kind are reflected in the broader Geopolitical Risk Index but are not directly incorporated into the Middle East-weighted component of GEAR. The results should therefore be understood as evidence on the Middle East channel rather than on geopolitical energy shocks originating anywhere in the world. A broader measure would need to account for these other sources of geopolitical disruption.
5. The $\text{Max}(0, \cdot)$ specification keeps only positive oil price changes. In this sample the censoring binds in six of the eleven years, leaving just five

distinct non-zero values. GEAR also reads zero in two of the three highesttension years, because tension rose while oil prices fell late in the observation period. What the variable measures with confidence is therefore energy cost pressure weighted by import dependence during periods of rising prices, rather than geopolitical intensity in the abstract; its geopolitical content is narrower than the name suggests. Under a symmetric specification retaining negative oil price changes the coefficient loses significance, which confirms that the association is specific to years of rising import-weighted energy cost. This is a boundary on interpretation rather than a defect, since the asymmetric form is the theoretically intended measure, but it should be kept in view when generalising the result.

6. Firm cost exposure is a broad proxy. The ratio of cost of goods sold to net sales measures how large a firm's direct cost base is. It says nothing about what that cost base contains, so two firms with the same ratio can face opposite exposures to oil. A fuel-specific ratio would be better. It cannot be built consistently across 54 firms and eleven reporting periods given current Indonesian disclosure practice.
7. Two of the four direct estimates depend on the cross-section weighting. Table 4.11 shows that EXPO and ROA lose significance once the weights are removed, while GPR and GEAR0 remain significant at the one percent level under common, fixed and random effects alike. When the product term is added, the unweighted GEAR0 main effect also loses significance, but that reflects the collinearity between GEAR0 and its own product rather than the weighting, as the direct estimates show. The weighting is a response to a diagnosed heteroskedasticity problem, not a matter of preference. Even so, the cost structure and profitability results should be read as conditional on it, and the GPR result carries the separate joint sensitivity to weighting and the pandemic year described in Section 4.2.
8. The model selection procedure did not resolve cleanly. Chow and the Lagrange Multiplier test both rejected the pooled specification. The Hausman test, which exists to settle the choice between fixed and random

effects, returned an invalid cross-section variance and could not be used. Fixed effects was adopted on grounds set out in Chapter 3, and the random effects estimates are reported next to it. The choice still rests on judgement, not on a decisive test.

9. The sample excludes banks and other financial institutions, because their cost structure is built on cost of funds and not on cost of goods sold. The exclusion was necessary. It does mean the findings say nothing about a sector that represents a large share of market capitalisation.
10. The balanced panel requirement excludes firms that listed after 2015 or were delisted during the period, which introduces a mild survivorship characteristic. The observation window also contains the COVID-19 pandemic. That produced extreme tail events in 2020 with no geopolitical cause at all.
11. The model captures linear relationships only. The interaction assumes the energy slope changes by a constant amount for each unit of cost exposure. Moderation could well be nonlinear, with threshold effects appearing at particular levels of cost intensity.
12. Expected shortfall is computed from simple rather than logarithmic returns, following the convention agreed with the academic supervisor. The two are close for small movements but diverge for the large negative returns that populate the left tail, so the choice has some influence on the level of the dependent variable.

5.3. Implications of the Study

5.3.1. Theoretical Implications

Two variables were built for this study and neither had been applied to Indonesian firm-level tail risk before. Both earned their place, though not equally and not in the way the design anticipated.

Geopolitical energy-adjusted risk holds its sign and significance at the one percent level across every alternative tail threshold, across exclusion of the pandemic year, and across the leave-one-year-out estimation. That is a stronger showing than GPR manages, and it is the result the construction was built to produce: weighting a global oil price movement by a country's own import dependence on the region generating the tension yields something the established index does not already contain. The two variables sit in the same regression, both attain significance, and their correlation of -0.4913 is moderate enough to rule out redundancy while high enough to confirm they are measuring related things. Neither an unweighted crude series, common to every country, nor an unweighted risk index, common to every sector, can express a global shock as a national exposure. This one can, and the method transfers directly to any oil-importing economy with a published import breakdown.

Firm cost exposure did something the design did not anticipate. The operating leverage literature has connected cost structure to beta and, more recently, to crash risk, but a direct cost intensity ratio set against a coherent tail measure at firm level appears not to have been tried, and certainly not in this market. Chapter 2 left the sign open because two defensible readings predicted opposite things. The estimate leans toward cost flexibility rather than margin thinness, a reading that only becomes legible once the accounting behind the ratio is made explicit, and one that, if it held under stronger evidence, would give pause to any stress-testing framework treating thin gross margins as a proxy for fragility. The evidence here is not strong enough to carry that conclusion on its own, given the ten percent significance and the sensitivity to weighting recorded in Chapter 4.

The dependent variable carries a quieter implication. Dwitanto, Manurung, Machdar and Hidayat (2026) established expected shortfall as a workable firmlevel outcome on Indonesian panel data using governance and asset-strategy determinants. Extending it to macroeconomic and geopolitical ones was not obviously going to work, since a variable taking eleven distinct values across 594 rows might reasonably have found nothing. It found something, which suggests the

framework accommodates both classes of regressor rather than only the firm-level class.

One assumption in the wider literature sits uneasily with this sample. Geopolitical risk is generally taken to harm emerging market equities, and here it accompanies shallower tails rather than deeper ones. What that implies is not that the received view is wrong but that its sign may depend on productive structure, a possibility this single-country design cannot confirm, and that cross-country studies pooling commodity importers with commodity exporters may be averaging across two effects running in opposite directions.

5.3.2. Managerial and Practical Implications

For risk managers, the results point to financial statement characteristics that can be observed directly and that may travel with tail behaviour in periods of geopolitical stress. Cost structure comes out of published accounts and needs no proprietary data, so it can be carried into stress testing and position sizing once its role has been tested more firmly. Cost-intensive firms here showed somewhat shallower tails in the primary specification, which, if it held up, would cut against the intuition that thin margins imply fragility, and stress tests built on that intuition would be worth re-examining.

For investors, the case is for testing financial statement data as a screen alongside sector labels. March 2026 on the Indonesia Stock Exchange, when energy and non-energy stocks moved apart, showed how coarse a guide sector classification can be to geopolitical exposure. Cost structure and commodity linkage are candidates for a finer one, though their value as screens has yet to be tested against sector labels directly.

For policymakers at Bank Indonesia and the Otoritas Jasa Keuangan, the inverse movement between Indonesian equity tail risk and Middle East energy cost pressure is consistent with the listed corporate sector being, in aggregate, a beneficiary of commodity-price shocks originating in that region. That complicates the usual framing of Indonesian vulnerability as an oil importer's problem. It also raises the question of whether macro-prudential monitoring of the equity market

should treat the commodity-linked segment separately from the rest. An aggregate estimate cannot answer that question, and the subgroup comparison it calls for is set out in Section 5.4.3.

One caution applies here. These results describe 2015 to 2025, a period in which oil prices and geopolitical tension came apart in the later years. The March 2026 episode, when the Indonesian market fell sharply alongside the rest of Asia, is not explained by the estimated relationships and lies outside the sample. Whether the historical pattern holds under a disruption on that scale is open, and nothing here should be treated as a forecast.

5.4. Recommendations

5.4.1. Recommendations for Investors and Portfolio Managers

Commodity linkage deserves consideration as an exposure dimension in its own right in portfolio construction, alongside the sector label, and the evidence here is a reason to test it rather than a demonstration that it works. Roughly a third of the index is made up of commodity-linked firms whose revenue rises when a geopolitical event pushes oil prices up, which is a plausible reason for aggregate energy pressure to arrive with shallower tails here; but the study did not estimate those firms separately, so how far a portfolio's sensitivity depends on its weighting toward them remains to be measured. Where a mandate requires tail risk to be contained during geopolitical stress, the results give grounds for examining cost structure and commodity linkage alongside the usual sector allocation, and for testing them in screening before either is relied upon. Stress-testing frameworks for Indonesian equities are also better calibrated on local evidence than imported from oil-importing markets, since a framework assuming that higher oil prices deepen equity tail risk would have had the sign of the aggregate relationship wrong for this market across the observation period.

5.4.2. Recommendations for Regulators and Policymakers

When the vulnerability of the equity market to external shocks is assessed, commodity exposure and cost structure deserve further testing as dimensions of

supervisory analysis. The combined relationship estimated here mixes firms whose exposure may run through different channels, and an aggregate estimate can hide the risk carried by firms that do not benefit when commodity prices rise; but the study did not estimate commodity-linked and other firms separately, so a recommendation to separate them in monitoring would run ahead of the evidence. What the findings do support is holding both sides of Indonesia's position in view, as an energy importer and a commodity exporter at the same time. A country's oil import balance does not by itself describe how an external energy shock reaches the corporate sector, since some companies may meet it as higher cost and others as higher revenue. Monitoring that allows for both channels, once the subgroup comparison has been made, would rest on a more complete basis.

5.4.3. Recommendations for Future Research

Several extensions follow from the limitations set out above. The first three respond directly to the questions the aggregate design leaves open and are placed first for that reason.

1. Testing commodity-linked and non-commodity firms separately. The reading offered for the negative signs on GPR and GEAR0 rests on the commodity orientation of the listed universe, and the most direct test of it is to estimate the two groups apart. Re-estimating the model with energy and commodity producers excluded, and then on that subgroup alone, would show whether the results depend on the firms whose revenues move with commodity prices, which bears on the construct validity of EXPO as much as on the macro estimates. A further step is to disaggregate by industry, distinguishing manufacturers whose production consumes fuel directly from producers whose revenues rise with it, so that the point at which the energy channel turns from revenue into cost can be located within the listed sector. Industry-level analysis of this kind would resolve what the aggregate estimates cannot, and until it is done the compositional explanation offered in Chapters 4 and 5 stands as a hypothesis.

Two related checks announced in Chapter 3 were not carried out and belong with this extension. Reporting the sector composition of the sample would show

which industries sit behind the estimated relationships, and adding year fixed effects would identify EXPO, ROA and the interaction free of any influence common to all firms in a given year, which is the most direct available strengthening of the moderation test, since the macro main effects rest on eleven annual observations while the interaction does not. Neither requires data beyond what has already been assembled.

2. Widening the sample beyond the Kompas 100. The index was chosen for liquidity and data availability, but it limits the number of firms in any one industry and so limits the subgroup comparison above. A sample drawn from the full population of non-financial issuers on the Indonesia Stock Exchange, or built from the IDX sector indices so as to maximise the contrast between commodity producers and energy-consuming industries, would give each group enough firms for separate estimation. Banks and other financial institutions, excluded here because the cost of goods sold ratio is not meaningful for them, could be brought in with cost-to-income or provision expense ratios in its place; they face geopolitical energy shocks through credit risk from oil-exposed borrowers, through policy rate responses to imported inflation and through currency exposure, none of which the present model captures.
3. Constructing an oil-related cost ratio. Where disclosure permits, isolating the fuel and energy component of cost of goods sold would give a direct measure of oil-related cost exposure and a sharper test of both the cost hypothesis and the moderation hypothesis than the broad cost intensity ratio used here. Section 3.3.4.1 explains why no such ratio can be assembled consistently across 54 firms and eleven reporting periods from published statements. Hand-collection from the notes for a smaller sample would make it feasible at the cost of the balanced panel, trading breadth for precision, and on this question that may be the more informative design. It would also separate the cost flexibility and margin thinness readings of EXPO, which the present evidence cannot.

4. Decomposing oil price changes by source. An important extension of the GEAR approach would be to identify the sources of oil price movements before the import-share weighting is applied. Kilian's structural decomposition may be used to distinguish supply shocks from aggregate demand and precautionary demand shocks, and the resulting measure would be closer to the geopolitical energy component that GEAR is intended to capture. This would also help with the discriminant validity issue noted above, since movements in oil prices unrelated to geopolitical supply disruptions could be treated separately. The refinement requires a structural VAR with monthly data on global production, economic activity and inventories, aggregated to the annual frequency used in the firm-level analysis.
5. Extending to higher frequency. Since expected shortfall is computed from daily returns while the regressors are annual, a design using daily or monthly geopolitical and oil price data with firm characteristics held fixed would allow the speed of transmission to be measured, which the present design cannot address.
6. Applying the framework to other markets. The GEAR methodology adapts directly to other oil-importing emerging economies, including Thailand, the Philippines, India and Turkey, by substituting their respective Middle East import shares. Comparing an importer without a large commodity export sector against Indonesia would test whether the negative sign observed here is attributable to productive structure, which the present single-country design cannot establish.
7. Extending the interaction structure. Financial leverage, firm size, foreign revenue share and hedging activity are all plausible moderators of the energy relationship and are observable from published data.
8. Testing the 2026 crisis out of sample. Once full-year 2026 financial statements become available, the estimated model can be used to generate predicted tail risk for each firm and compared against realised performance during the March 2026 disruption. Because that episode is far larger in

magnitude than anything within the estimation window, it offers a genuine test of whether the historical relationships hold under extreme conditions, and it is the most direct available check on the external validity of these findings.

A final observation concerns the period itself. The decoupling of oil prices from geopolitical tension in the later years of the sample, which produced zero values for the energy variable in two of the three highest-tension years, may prove characteristic of a period in which supply growth outside the Middle East reduced the sensitivity of prices to regional conflict. Whether that decoupling persists, and what it implies for the transmission of geopolitical risk to emerging market equities, is a question that further observation will answer more convincingly than additional analysis of the present sample.

REFERENCES

- Acerbi, C., & Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7), 1487–1503.
[https://doi.org/10.1016/S03784266\(02\)00283-2](https://doi.org/10.1016/S03784266(02)00283-2)
- Ang, A., Chen, J., & Xing, Y. (2006). Downside risk. *The Review of Financial Studies*, 19(4), 1191–1239. <https://doi.org/10.1093/rfs/hhj035>
- Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.
<https://doi.org/10.1111/14679965.00068>
- Asuero, A. G., Sayago, A., & González, A. G. (2006). The correlation coefficient: An overview. *Critical Reviews in Analytical Chemistry*, 36(1), 41–59. <https://doi.org/10.1080/10408340500526766>
- Bacon, R. W. (1991). Rockets and feathers: The asymmetric speed of adjustment of UK retail gasoline prices to cost changes. *Energy Economics*, 13(3), 211–218. [https://doi.org/10.1016/0140-9883\(91\)90022-R](https://doi.org/10.1016/0140-9883(91)90022-R)
- Badan Pusat Statistik. (2024). Statistik perdagangan luar negeri Indonesia: Ekspor 2023. BPS-Statistics Indonesia. <https://www.bps.go.id>
- Bali, T. G., Demirtas, K. O., & Levy, H. (2009). Is there an intertemporal relation between downside risk and expected returns? *Journal of Financial and Quantitative Analysis*, 44(4), 883–909.
<https://doi.org/10.1017/S0022109009990159>
- Bank for International Settlements. (2019). Minimum capital requirements for market risk (BCBS d457). Basel Committee on Banking Supervision.
<https://www.bis.org/bcbs/publ/d457.htm>
- Basher, S. A., & Sadorsky, P. (2006). Oil price risk and emerging stock markets. *Global Finance Journal*, 17(2), 224–251.
<https://doi.org/10.1016/j.gfj.2006.04.001>
- Baur, D. G., & Lucey, B. M. (2010). Is gold a hedge or a safe haven? An analysis of stocks, bonds and

- gold. *The Financial Review*, 45(2), 217–229.
<https://doi.org/10.1111/j.1540-6288.2010.00244.x>
- Borenstein, S., Cameron, A. C., & Gilbert, R. (1997). Do gasoline prices respond asymmetrically to crude oil price changes? *The Quarterly Journal of Economics*, 112(1), 305–339. <https://doi.org/10.1162/003355397555118>
- Bouras, C., Christou, C., Gupta, R., & Suleman, T. (2019). Geopolitical risks, returns, and volatility in emerging stock markets: Evidence from a panel GARCH model. *Emerging Markets Finance and Trade*, 55(8), 1841–1856. <https://doi.org/10.1080/1540496X.2018.1507906>
- Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225.
<https://doi.org/10.1257/aer.20191823>
- Dwitanto, A. R., Manurung, A. H., Machdar, N. M., & Hidayat, R. S. (2026). Controlling ownership, landbanking and extreme downside risk: Evidence from Indonesian listed property developers. *Architectural Image Studies*, 7(1), 2999–3014.
- Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22.
<https://doi.org/10.1016/j.jfineco.2014.10.010>
- Ghozali, I. (2021). *Aplikasi analisis multivariate dengan program IBM SPSS 26* (10th ed.). Badan Penerbit Universitas Diponegoro.
- Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2014). *Multivariate data analysis* (7th ed.). Pearson Education.
- Hamilton, J. D. (1996). This is what happened to the oil price–macroeconomy relationship. *Journal of Monetary Economics*, 38(2), 215–220.
[https://doi.org/10.1016/S0304-3932\(96\)01282-2](https://doi.org/10.1016/S0304-3932(96)01282-2)
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2), 363–398. [https://doi.org/10.1016/S0304-4076\(02\)00207-5](https://doi.org/10.1016/S0304-4076(02)00207-5)

- Hoque, M. E., & Zaidi, M. A. S. (2020). Global and country-specific geopolitical risk uncertainty and stock return of fragile emerging economies. *Borsa Istanbul Review*, 20(3), 197–213. <https://doi.org/10.1016/j.bir.2020.05.001>
- Iacoviello, M. (2026). Geopolitical risk (GPR) index, 2015–2025 [Data set]. Retrieved April 8, 2026, from <https://www.matteoiacoviello.com/gpr.htm>
- International Energy Agency. (2026). Oil market report. International Energy Agency. <https://www.iea.org/topics/oil-market-report>
- International Trade Centre. (2026). Trade Map: Trade statistics for international business development, 2015–2025 [Data set]. Retrieved April 8, 2026, from <https://www.trademap.org>
- Jorion, P. (2007). *Value at risk: The new benchmark for managing financial risk* (3rd ed.). McGraw-Hill.
- Kilian, L., & Park, C. (2009). The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50(4), 1267–1287. <https://doi.org/10.1111/j.1468-2354.2009.00568.x>
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.
- Lev, B. (1974). On the association between operating leverage and risk. *Journal of Financial and Quantitative Analysis*, 9(4), 627–641. <https://doi.org/10.2307/2329764>
- Mandelker, G. N., & Rhee, S. G. (1984). The impact of the degrees of operating and financial leverage on systematic risk of common stock. *Journal of Financial and Quantitative Analysis*, 19(1), 45–57. <https://doi.org/10.2307/2331000>
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91. <https://doi.org/10.1111/j.1540-6261.1952.tb01525.x>
- Markowitz, H. (1991). *Portfolio selection: Efficient diversification of investments* (2nd ed.). Blackwell.

- Mishra, P., Pandey, C. M., Singh, U., Gupta, A., Sahu, C., & Keshri, A. (2019). Descriptive statistics and normality tests for statistical data. *Annals of Cardiac Anaesthesia*, 22(1), 67–72.
https://doi.org/10.4103/aca.ACA_157_18
- Mork, K. A. (1989). Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *Journal of Political Economy*, 97(3), 740–744. <https://doi.org/10.1086/261625>
- Moulton, B. R. (1990). An illustration of a pitfall in estimating the effects of aggregate variables on micro units. *The Review of Economics and Statistics*, 72(2), 334–338. <https://doi.org/10.2307/2109724>
- Ndako, U. B., Salisu, A. A., & Ogunsiji, L. O. (2021). Geopolitical risk and the return volatility of Islamic stocks in Indonesia and Malaysia: A GARCHMIDAS approach. *Asian Economics Letters*, 2(3).
<https://doi.org/10.46557/001c.24843>
- Notteboom, T., Haralambides, H., & Cullinane, K. (2024). The Red Sea Crisis: Ramifications for vessel operations, shipping networks, and maritime supply chains. *Maritime Economics & Logistics*, 26(1), 1–20.
<https://doi.org/10.1057/s41278-024-00287-z>
- Novy-Marx, R. (2011). Operating leverage. *Review of Finance*, 15(1), 103–134.
<https://doi.org/10.1093/rof/rfq019>
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1), 1–28.
<https://doi.org/10.1016/j.jfineco.2013.01.003>
- Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3), 341–360. [https://doi.org/10.1016/0022-0531\(76\)90046-6](https://doi.org/10.1016/0022-0531(76)90046-6)
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3), 425–442.
<https://doi.org/10.1111/j.1540-6261.1964.tb02865.x>

- Sugiyono. (2017). Metode penelitian kuantitatif, kualitatif, dan R&D. Alfabeta.
- Sukarno, A. R., & Setiawati, R. (2020). The impact of world oil price shocks on the Indonesian sectoral stock market. *International Journal of Energy Economics and Policy*, 10(6), 12–19. <https://doi.org/10.32479/ijep.10121>
- UNCTAD. (2024). Review of maritime transport 2024: Navigating maritime chokepoints. United Nations Conference on Trade and Development. <https://unctad.org/publication/review-maritime-transport-2024>
- U.S. Energy Information Administration. (2025). Country analysis brief: Indonesia. U.S. Department of Energy. Retrieved from <https://www.eia.gov/international/analysis/country/IDN>
- Wahyudi, S. T., & Rahmawati, A. (2025). Examining the link between world oil price volatility and sectoral stock performance in Indonesia. *Cogent Economics & Finance*, 13(1), 2535484. <https://doi.org/10.1080/23322039.2025.2535484>
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48(4), 817–838. <https://doi.org/10.2307/1912934>
- Zarembo, A., Cakici, N., Demir, E., & Long, H. (2022). When bad news is good news: Geopolitical risk and the cross-section of emerging market stock returns. *Journal of Financial Stability*, 58, 100964. <https://doi.org/10.1016/j.jfs.2021.100964>
- Zhu, L., & Zang, W. (2025). Effect of operating leverage on stock price crash risk: Evidence from China. *Economic Modelling*, 153, 107320. <https://doi.org/10.1016/j.econmod.2025.107320>