Actuaries in Space: The Pricing Frontier Report

SOA Research Institute 2026 Student Research Case Study Challenge

VALCS Consulting
University of New South Wales

Alan Steny, Carol Zhang, Liya Ruan, Sarina Truong & Vihaan Jain
Advisor: Xiao Xu

Table of Contents

1. Executive Summary (All) *3*

2. Product Design *3*

2.1 Equipment Failure (EF) 3

2.2 Cargo Loss (CL) 4

2.3 Workers' Compensation (WC) 4

2.4 Business Interruption (BI) 5

3. Summary of Pricing and Capital Modelling *6*

3.1 Frequency Model 6

3.2. Severity Model 7

3.3. Aggregate Loss Distribution 7

3.4 Capital Modelling 8

3.5 Modelling Interplanetary Dependencies 10

4. Risk Assessment *10*

4.1. Risk identification by Solar System 10

4.2 Threat Table 10

.... 11

4.3 Risk Matrix 11

4.4 Scenario Testing 11

5. Key Assumptions *11*

A comprehensive list of assumptions can be found in Appendix 5.1. .... *12*

6. Data and Data Limitations *12*

6.1 Data Limitations 12

6.2 AI Usage 12

7. References *13*

Appendix *14*

Appendix 2.1 14

Appendix 2.2 14

Appendix 2.3 16

Appendix 2.4 Appendix 2.4A 18

Appendix 3.1 19

Appendix 3.2 22

Appendix 3.3 26

.... 27

Appendix 3.4 .... **27**
Appendix 4.2 .... **31**
Appendix 4.4.1 .... **31**
Appendix 4.4.2 .... **32**
Appendix 4.4.3 .... **32**
Appendix 5.1 .... **33**
Appendix 5.1A **33**
Appendix 6.1 .... **35**
Appendix 6.1A **35**
Appendix 6.1B Further Considerations **35**

1. Executive Summary

To pioneer success in interstellar commercial insurance, VALCS Consulting recommends the following targeted product suite for Cosmic Quarry Mining Corporation. Our rigorous capital modelling projects expected costs, revenues and tail risks by applying a cost-of-capital risk margin to the 99.5% tail of our simulated aggregate loss distributions. This strategy ensures our capital buffer withstands extreme volatility while generating an expected $9.2 billion aggregate return and a 132% return on opening capital over the next 10 years.

We recommend offering Equipment Failure, as its predictable severity provides a profitable avenue for market expansion. Workers' Compensation is also recommended, fulfilling a necessary social mandate while remaining commercially viable through protective stop-loss reinsurance. Business Interruption is viable because strictly defined shutdown triggers and aggregate limits effectively contain exposure. Conversely, we recommend excluding Cargo Loss, as the extreme value concentration of transported metals generates uninsurable tail risk and unsustainable capital requirements.

These products dynamically scale to accommodate Cosmic Quarry's ten-year growth by explicitly incorporating 15% and 25% expansion trajectories, projected interplanetary inflation, and annually renegotiated reinsurance attachments. Furthermore, pricing features specific tweaks for each solar system's unique risk profile. We apply tailored premium multipliers and targeted deductibles to address kinetic debris in the Helionis Cluster, severe ambient radiation in the Bayesia System and unpredictable gravitational anomalies in Oryn Delta.

Our framework actively supports Cosmic Quarry's ESG objectives. Workers' Compensation enforces strong social standards by mandating unlimited liability for injured personnel while excluding negligence claims to incentivise rigorous safety cultures (Shin et al., 2011). Business Interruption policies promote environmental stewardship by requiring strict adherence to interstellar guidelines to minimise hazardous debris. For enterprise risk management (ERM), integrating these policies transfers severe operational volatility to Galaxy General. Meanwhile, excluding cargo loss establishes clear risk retention boundaries, requiring internal committees to implement physical transport mitigations that transform unpredictable cosmic threats into manageable operating expenses.

2. Product Design

2.1 Equipment Failure (EF)

Equipment Failure is expected to be a profitable product with low capital requirements, providing an opportunity to expand Galaxy General's market share in the space mining industry. Compared to other product lines, equipment failure exhibits more predictable claim severity, as it is closely linked to the type of the equipment. This predictability enables more accurate modelling of expected losses and effective risk-based pricing even across different solar systems. Additionally, there is strong prospective growth for this line of coverage, driven by the expansion plans of the Cosmic Quarry Mining Corporation and broader space mining operations, creating a sizable and growing market. Finally, offering equipment failure coverage aligns closely with Galaxy General's expertise in mining technologies, allowing the company to leverage its knowledge of operational risks and engineering controls to manage and mitigate potential losses effectively.

Feature	Terms
Per-claim limit	Unlimited
Coverage	All equipment type present in Cosmic Quarry's inventory dataset
Premium basis	Per equipment type, GLM risk-rated by solar system, equipment age, usage intensity and maintenance interval

Coverage trigger: Equipment failure insurance is designed to cover the costs associated with repairing or replacing specific equipment following an unforeseen and sudden physical failure.

Benefit structure: We will pay for the cost to restore the machine to its original condition prior to breakdown, including the costs of dismantling, freight, replacement parts and re-installation. If the cost of repair exceeds the actual cash value of the equipment at the time of breakdown, we will pay the actual cash value of the equipment.

Exclusions: We will not cover any events outside the defined benefit structure, nor any equipment breakdowns resulting from causes that are not sudden and unforeseen. For a comprehensive list of exclusions, please refer to *Appendix 2.1A*.

2.2 Cargo Loss (CL)

While Cargo Loss represents a significant operational risk for Cosmic Quarry Mining Corporation, our analysis indicates that offering cargo loss insurance is not currently commercially viable for Galaxy General. The transport of extremely high-value metals such as gold and platinum creates severe loss potential, requiring substantial capital and resulting in premiums that would likely be unsustainable for the client. Providing coverage under these conditions would expose the insurer to disproportionate balance sheet volatility relative to the premium income generated. This will be further discussed in *Section 3.4*.

Accordingly, we do not recommend introducing a cargo loss product currently. Should the value concentration of transported metals decrease in the future, or if market conditions allow risks to become more diversifiable and predictable, Galaxy General may reconsider offering cargo loss coverage. The proposed design of a potential restricted product is provided in *Appendix 2.2* for management's reference.

2.3 Workers' Compensation (WC)

Workers Compensation is a necessary product for the Cosmic Quarry's Operations across the solar systems. Occupation injury losses are linked to identifiable events and with well supported historical data alongside rigorous modelling. The range in claims severity seen through a fitted σ^2 = 1.73, depicts the need for careful product design especially in the catastrophic tail end without compromising social insurance mandate that underpins workers' compensation.

Feature	Terms
Per-claim limit	Unlimited
Reinsurance structure	Portfolio stop-loss attaching at Đ6,650,395 (VaR 99%)
Reinsurance premium	Đ2,993 Year 1 (120% load on expected ceded loss)
Coverage	All occupations, full-time and contract workers equally
Premium basis	Per employee, GLM risk-rated by occupation, solar system and employment type

Cosmic Quarry is indemnified for the full cost of any covered claim with no per-claim deductible, sublimit, or benefit cap (*Appendix 2.3A and 2.3B*). No per-claim limit was imposed as doing so would leave seriously injured workers including those suffering permanent disability or fatal accidents partially unpaid, contradicting the core social insurance mandate of workers' compensation (ILO, 1964).

<u>*Coverage Trigger:*</u> Work-related injuries or illnesses sustained by Cosmic Quarry employees across all operating solar systems. (See *Appendix 2.3A* for comprehensive list).

<u>*Exclusions:*</u> Pre-existing conditions, off-duty incidents, intoxication, self-inflicted injuries, acts of war and claims within the 90-day seasoning period (See *Appendix 2.3B* for comprehensive list). Unlimited per-claim liability necessitates explicit tail risk management. Rather than restricting coverage, Galaxy General employs stop-loss reinsurance well established in actuarial literature as the optimal mechanism for relieving the primary insurer from unbounded tail exposure (Ignatov et al., 2022). When total portfolio losses exceed Đ6,650,395, the reinsurer assumes all additional losses unconditionally. Under the S5 Pandemic scenario, gross losses are projected at Đ26.4M, yet Galaxy General's retained exposure remains capped at the attachment point across all simulated paths. The reinsurance premium is transparently included in the gross premium rather than recouped via benefit caps (*Appendix 2.3C*).

The product maintains profitability across all non-pandemic scenarios throughout the ten-year projection horizon, with a combined ratio below 1.0 and positive net revenue in every year (*Appendix 2.3E*).

2.4 Business Interruption (BI)

Business Interruption is a strong candidate for Galaxy General as its losses are well-attributed to identifiable operational shutdown events and historical experience is sufficiently observable to support projections. With proactive risk-reduction features including waiting periods, deductibles, occurrence limits and exclusions, BI is well-positioned as a leading product in the universe.

It is recommended that BI insurance lines offer three streams of coverage (*Product A, B, C*) differentiated by their attachment points, with a full product comparison provided in *Appendix 2.4D*. Product B is endorsed as the optimal option for balancing premium adequacy, coverage breadth and capital efficiency. Its $1.0m deductible was selected to eliminate high-frequency, low-severity attritional claims that would otherwise inflate administrative costs and distort loss experience. The 14-day waiting period was calibrated to the observed distribution of shutdown durations in the historical data, ensuring only genuine prolonged interruptions trigger indemnification. The $15m per-occurrence limit reflects the 95th percentile of projected single-event losses across the portfolio, providing meaningful coverage breadth without exposing Galaxy General to extreme tail events at the per-claim level. An annual aggregate limit of $45m, set at three times the per-occurrence limit, caps total annual outflows across multiple events and protects capital in high-frequency loss years. A full coverage list can be found in *Appendix 2.4A*.

Under Product B, Galaxy General indemnifies the insured for proven revenue loss arising from an operational interruption once the waiting period has been exceeded. Gross loss is calculated as the number of interrupted days multiplied by the station's average daily revenue over the 30 days immediately preceding the event, with the $1m deductible retained by the insured and Galaxy General meeting the residual loss up to the $15m per-occurrence limit. A full solar system breakdown can be found in *Appendix 2.4C*.

Feature	Terms
Per-claim limit	Unlimited
Coverage	Business interruption cover triggered by equipment failure shutdown, partial output reduction exceeding 30%, supply chain interruption preventing commodity transport, mandatory evacuation due to catastrophic or radiative events, debris field closure exceeding regulatory safety thresholds, or crew safety shutdown where workforce capacity becomes operationally insufficient
Premium basis	Per equipment type, GLM risk-rated by solar system, equipment age, usage intensity and maintenance interval

Exclusions: Coverage excludes losses arising from scheduled maintenance, gradual deterioration, market price fluctuations, pre-existing faults, acts of war, double recovery under other policies or claims occurring within the first 90 days of a new policy. The policy strictly excludes any revenue losses, regulatory fines or legal penalties resulting from environmental disputes or compliance breaches. A comprehensive exclusions list can be found in *Appendix 2.4B.*

3. Summary of Pricing and Capital Modelling

3.1 Frequency Model

A Negative Binomial GLM was selected as the frequency model across all product lines. Claim counts are discrete, non-negative, and preliminary exploratory data analysis indicated variability that violates the Poisson assumption. Given the concentration of policies with zero claims, a Zero-Inflated Negative Binomial model was initially considered to explicitly account for the excess zeros (Frees 012, p. 354). However, the EDA revealed that the zero counts do not arise from a separate process, such as deductible thresholds. Consequently, the standard Negative Binomial was preferred, as it accommodates overdispersion while maintaining interpretability.

The frequency model incorporates covariates that are observable or inferable at the point of underwriting and are identified as key drivers of claim occurrence based on historical data. For a detailed summary of the model selection process and chosen parameters for each product line, refer to *Appendix 3.1.*

A key limitation of the frequency model is that it was constructed using historical claim data from different solar systems, which were used as proxies for new systems. While risk factors were manually applied to account for differences in expected risk, the accuracy of predictions remains variable. Therefore, it is important to update the model as future claim data becomes available. Additionally, the univariate Negative Binomial assumes independence between claim counts, which may not hold in scenarios that affect multiple product lines simultaneously. This limitation was partially mitigated through scenario testing, discussed further in *Section 4.3*.

Product	Model	E(X)	SD(X)
Business Interruption	NB ~ (0.2020,0.72)	0.20199	0.48087
Equipment Failure	NB ~ (0.1033, 0.43)	0.10329	0.35814

Workers Compensation	NB ~ (0.0147, 7.84)	0.01469	0.12330
Cargo Loss	NB ~ (0.442, 0.16)	0.422778	1.346285

3.2. Severity Model

A Log Normal GLM was selected as the severity model across all product lines. Claim costs are positive and right skewed and losses scale multiplicatively rather than by fixed shifts making the Log Normal suitable, implemented as a Gaussian GLM on log transformed claim amounts. This is supported by QQ plots of log-transformed claims aligning closely with the theoretical line, well-behaved residual plots, and marginal AIC comparisons ranking Log-Normal above Gamma, Weibull, and Pareto alternatives across all lines (*Appendix 3.2)*.

Severity model covariates were restricted to variables observable at the point of pricing. Where a statistically significant predictor such as exposure was not directly observable for future policyholders, it was simulated from a uniform distribution consistent with the empirical spread in historical data and used as a covariate within the GLM to generate forward looking severity predictions. This preserved predictive value without requiring future observability.

The primary limitation of the Log-Normal is that it underestimates tail risk relative to heavier tailed alternatives such as Pareto or Burr. This is partially mitigated through specialised stress models that incorporate adverse risk factors, and through a multiplicative bias correction applied during simulation to address the systematic underestimation that arises when converting log-space predictions back to dollar scale.

Product	Model	E(X) ($ Scale)	SD(X) ($ Scale)
Business Interruption	LN ~ (14.58, 1.2)	4,370,702.02	7,784,089.34
Equipment Failure	LN ~ (11.36, 0.570)	86,084.01	54,504.0
Workers Compensation	LN ~ (7.827, 1.32)	5,958.10	12,839.8
Cargo Loss (Gold)	LN ~ (17.19, 0.795)	42,686,682.5	42,114,906.1
Cargo Loss (Platinum)	LN ~ (15.78, 0.79)	1,0329,328.4	10,185,109.9
Cargo Loss (Other)	LN ~ (12.08, 0.80)	388,834.4	493,530.7

3.3. Aggregate Loss Distribution

The frequency and severity models were subsequently combined to construct a stochastic aggregate loss model. A total of 10,000 Monte Carlo simulations were performed for each product line to generate the aggregate loss distribution. Histograms of the simulated aggregate losses for each product line in the year 2175 are presented below (*see Appendix 3.3*).

From the Figure below we can see that the aggregate loss for equipment failure exhibits an approximately symmetric, near-normal distribution, reflecting its stable severity of claims as discuss in *Section 2.1* above. Similarly, business interruption also has a near normal aggregate loss distribution. In contrast, the remaining product lines display more pronounced right-skewness, indicating the potential for occasional large losses. For workers' compensation, the impact of this skewness is limited, as individual claim sizes are relatively small and therefore contribute only modestly to overall portfolio volatility. However, the right-skewness is far more pronounced for Cargo Loss, where the potential for very large claims leads to substantial volatility in aggregate losses, thus leading to our decision to exclude cargo loss from our product coverage.

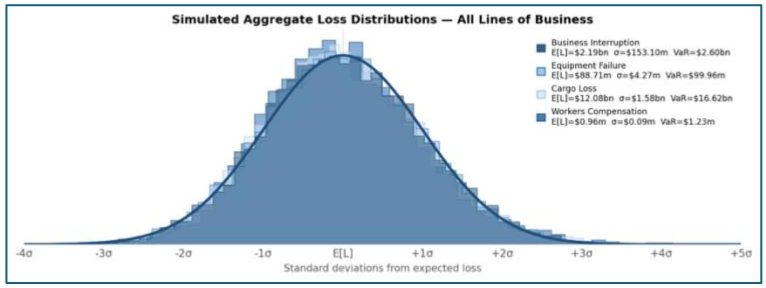

3.4 Capital Modelling

The capital model for Cosmic Quarry's portfolio is projected across a 2175 - 2185 window using 10,000 bootstrapped simulations for each product line. The simulations were then stress tested under 11 stress testing scenarios and are mapped across 3 capital requirements (100%, 150%, 200% SCR). In addition, the model incorporates a 40% quota share reinsurance structure with a 25% ceding commission returned to the insurer.

Aggregate Costs

Year	E[Cost]	Std Dev	P99 Loss	TVaR 99.5%	Range (E ± 2σ)
2176	$9.14B	$1.00B	$11.69B	$12.37B	$7.14B - $11.14B
2185	$14.48B	$1.59B	$18.50B	$19.45B	$11.29B - $17.66B

A short-term projection of net losses is estimated to be *$9.14B* and is expected to grow long-term to *$14.48B*, indicating a 58.4% growth rate. This is largely driven by the increase in expected policies that Cosmic Quarry will face in achieving their 15% and 25% 10-year growth rates. The distribution also exhibits a pronounced tail risk, reflecting the potential for low frequency but high severity catastrophic events which materially influence the upper loss quantiles. A full year-year breakdown can be found in *Appendix 3.4B/C*.

Premiums reflect expected losses plus a risk margin based on the 99.5% tail risk of the aggregate loss distribution to ensure Galaxy General remains solvent under extreme events. See *Appendix 3.4A* for further detail on derivation and capital modelling methodology.

Product Line	Expected Loss	Risk Margin	Gross Premium
Cargo Loss	$12,075.96M	$976.22M	$19,335.05M
Equipment Failure	$88.71M	$2.26M	$135.49M
Workers Compensation	$0.96M	$56,774.69	$1.51M
Business Interruption	$2,190.92M	$81.81M	$3,380.85M

The capital model produces a total gross premium requirement of $22.85B across the four proposed insurance products. This premium reflects expected losses plus a risk margin calibrated to the 99.5% tail of the aggregate loss distribution, ensuring Galaxy General remains solvent under extreme events. However, this level of premium would place a significant financial burden on Cosmic Quarry Mining Corporation. For context, Cosmic Quarry reported net income of $14.76B in FY2174, meaning the full insurance program would exceed the company's annual profit. The majority of this cost is driven by the cargo loss line, which alone requires $19.34B in premium due to the extreme severity risk associated with transporting high-value metals such as gold and platinum. As a result, we recommend excluding cargo loss coverage from the portfolio currently, while proceeding with equipment failure, workers' compensation, and business interruption, whose premiums appear financially sustainable relative to Cosmic Quarry's operating scale.

Aggregate Returns

Year	E[Total Return]	Std Dev	P5 (1-in-20 downside)	Solvency Ratio
2176	$1.01B	$1.12B	($0.93B)	2.52x
2185	$1.33B	$1.78B	($1.74B)	2.15x

The above table provides a short-term to long-term view of returns and is assumed under a 150% SCR capital requirement. The return includes both their underwriting returns (net written premium *less* net claims *less* operating expenses and claims handling costs) and their investment income earned on their capital base. The 40% quota share reinsurance results in 40% of the premium and claims flowing to the reinsurer and has a 25% ceding commission returned. The 2076 aggregate revenue is estimated at just over $1bn and projected to grow to be $1.3B in 2085. The 10-year aggregate total is $9.2Bn (discounted) and indicated a net profit margin of 5.25% and return on opening capital of 132%. For volatility, an individual yearly breakdown shows that in 2185 tail risk may exceed returns ($1.78B vs $1.33B), however, individual bad years are absorbed by the 150% capital buffer. Moreover, the model has a recapitalisation assumption such that if capital falls below zero on any simulation path, there is a single capital injection to 100% SCR. The probability of needing this injection across the 10 years is 0.16%, indicating that the capital buffer is sufficient in all but extreme catastrophic events. Overall, VALCS believes there is suitable returns over the 10-year period to not warrant concern for Galaxy General.

Net Revenue

Year	E[Net Revenue]	Std Dev	Loss Ratio	Net Premium	Net Claims
2176	$0.646B	$1.121B	53.99%	$16.925B	$9.138B
2185	$1.069B	$1.782B	53.85%	$26.884B	$14.477B

This table indicates the combined portfolio of recommended products is expected to remain financially viable over both the short (1Y) and long (10Y) term. In the short term, the portfolio generates an expected net revenue of $0.646B, supported by $16.925B in net premium against $9.138B in net claims, resulting in a loss ratio of approximately 54%. Over the long term, expected net revenue increases to $1.069B, with net premium rising to $26.884B and net claims to $14.477B, reflecting the assumed growth of Cosmic Quarry's operations across the three solar systems. Despite this expansion, the loss ratio is expected to remain stable, assuming claims costs and operational expenses scale efficiently with the portfolio. While the standard deviation of net revenue increases in absolute terms, indicating greater earnings volatility as the portfolio grows, the positive expected net revenue in both periods demonstrates that the recommended product mix remains financially sustainable across both short and long-term horizons.

Stress Test

We define a stress test as the event where an existing model parameter is pushed to an extreme value. The model structure is unchanged and only the magnitude of a known input is shocked. Stress testing indicates that the portfolio remains resilient under most individual shocks, including lower investment income, workers compensation deterioration and solar system

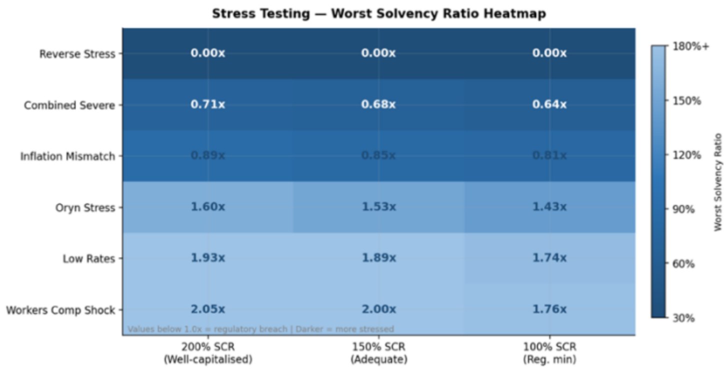

specific underperformance. However, solvency becomes vulnerable under macroeconomic inflation mismatches and severe multi-factor events. In the combined severe scenario, where claim frequency rises by 20%, inflation increases by 2%, investment income falls by 2% and workers compensation losses increase, the solvency ratio declines to 0.71× and ruin probability rises to 14.5%, highlighting correlated shocks as the primary threat to capital adequacy. Overall, Galaxy General's portfolio remains sufficiently capitalised to withstand most extreme scenarios. (see *Appendix 3.4D* for more details)

3.5 Modelling Interplanetary Dependencies

To incorporate shared nuances between systems a dependency structure was introduced in the form of a shared annual portfolio shock that effects all solar systems simultaneously. The result of this was immaterial to expected losses but it increases portfolio tail risk by allowing multi-system deterioration. As a result, the diversified required capital increased to approximately $3.6B from $3.2B, reducing the diversification benefit from roughly $1.4B to $1.0B. Overall, the dependency adjustment primarily affects extreme outcomes rather than average profitability, hence VALCS concludes it is not a significant concern.

4. Risk Assessment

4.1. Risk identification by Solar System

Solar system risk profiles are defined by unique stellar and orbital characteristics. In the Helionis Cluster solar system, irregular gravitational resonances cause frequent micro collisions and shifting debris clouds, necessitating the periodic repositioning of communication satellites. The Bayesia System solar system is instead defined by a binary star pair that generates sharp electromagnetic spikes, resulting in elevated ambient radiation and temperature extremes on the mining planet. Conversely, the Oryn Delta solar system presents a low visibility environment with the potential for unpredictable solar flares. Its primary hazard is an asymmetric asteroid ring that creates localised zones of rapid orbital shear and fluctuating gravitational gradients which threaten deep extraction infrastructure.

4.2 Threat Table

The following table presents the most unique and complex threats to the portfolio. For a more comprehensive list of threats, please refer to *Appendix 4.2*.

Threat	Description	Mitigation Strategy
Strong Environmental Risk	Extreme simultaneous total losses from systemic environmental hazards like the rapid orbital shear and fluctuating gravitational gradients in the Oryn Delta.	Exclude Cargo Loss entirely due to uninsurable tail risk. Apply stringent limits of $15M per occurrence and $45M annually for Business Interruption policies.

Technology Risk	Systemic failures in fleets or AI systems where a correlated cyber-attack could trigger simultaneous Equipment Failure and Business Interruption claims across systems.	Mandate redundant manual backup systems and strict engineering logs. Deploy orbital AI adjudicators to rapidly isolate and assess claims regarding systemic technological failures.
Regulatory & Compliance Risk	Operational halts or penalties from environmental disputes given the history Cosmic Quarry has with major settlements and fines for landscape alteration and orbital saturation.	Explicitly exclude regulatory fines and penalties from all policy wording. Require strict adherence to environmental guidelines to trigger Business Interruption coverage and prevent moral hazard.

4.3 Risk Matrix

The Risk Matrix visually maps the top identified threats to the portfolio based on their estimated likelihood and potential financial impact prior to intervention. It illustrates a diverse threat landscape that ranges from rare but catastrophic environmental tail events to highly probable and routine financial fluctuations. Crucially, comprehensive mitigation strategies such as stringent policy limits, reinsurance structures and strict exclusions have been proactively embedded to address every plotted threat and safeguard the solvency of Galaxy General.

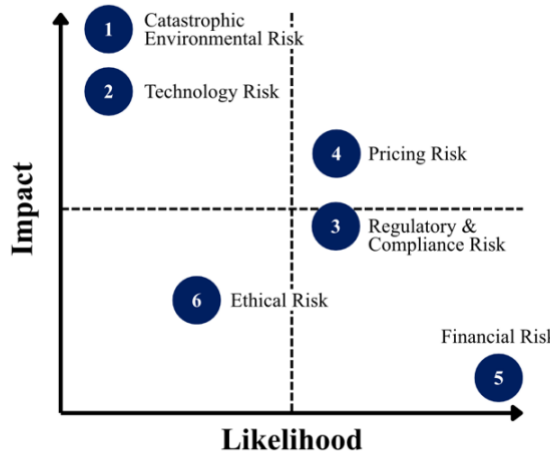

4.4 Scenario Testing

To evaluate portfolio resilience under extreme conditions we conducted rigorous scenario testing. We define a scenario test as the introduction of a new risk not captured in the base model such as a correlated multi system event a causal chain between product lines or an external market shock.

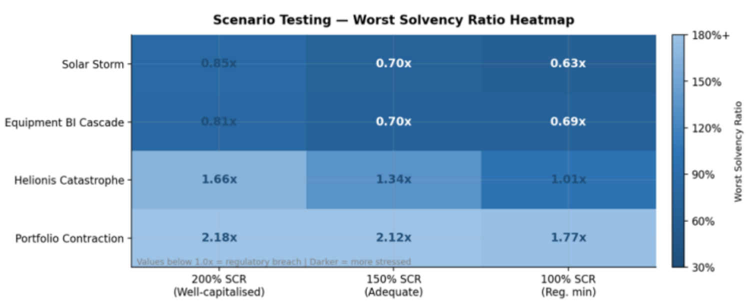

The accompanying heatmap illustrates the financial impact of unmodelled events on capital adequacy. A regulatory minimum of 100% SCR leaves the portfolio vulnerable to ruin during a Solar Storm or Equipment BI Cascade, whereas 150% and 200% buffers successfully absorb moderate shocks. However, because even the strongest capital positions fail against extreme correlated events, the data reinforces our recommendation to deploy portfolio stop loss reinsurance. A full case scenario analysis is provided in *Appendix 4.4 to Appendix 4.4E.*

5. Key Assumptions

Assumption	Value	Rationale	Sensitivity
Inflation & Discount Rate	2175: carry forward 2174 rate; 2176–77: 3yr avg spot; 2178–80: 5yr avg spot; 2181–85: 10yr avg spot	Term-matched rates per AASB 1023 (AASB, 2023)	Longer-duration rates amplify discounting impact in later years

New Business Growth	Helionis & Bayesia: +25% over 10yrs; Oryn Delta: +15% over 10yrs (linear)	Aligned to CQ's expected system-level expansion	Faster growth increases exposure and aggregate loss volatility
Expense Ratios & CoC	Operating: 25% GP; Claims: 12% net loss; CoC: 12%	Industry proxies for specialised commercial lines	A 5pp shift in operating ratio moves the combined ratio by ~5pp
Capital Management	Distribute 25% of capital above 150% SCR; recapitalise at 0% SCR to 100%; max 1 injection per path	Maintains ROE discipline; single injection cap reflects real-world distress constraints	Higher distribution rates erode the solvency buffer under stress
Diversification	Stand-alone required capital per system vs. diversified combined portfolio capital	Capital saving from non-simultaneous adverse outcomes across systems	May be overstated if cross-system dependencies were not explicitly modelled

A comprehensive list of assumptions can be found in *Appendix 5.1.*

6. Data and Data Limitations

The reports analysis relied solely on project provided data, key limitations, assumptions, and their impacts are summarised below. Detailed procedures and considerations taken can be found in *Appendix 6.1A & Appendix 6.1B*.

6.1 Data Limitations

Limitation	Issue	Impact
Data Quality	Typos, missing fields, out-of-range values	Cleaning preserved ~10% more observations via cross-dataset validation
Limited Historical Coverage	No claims data for new systems; cross-system correlations only provided for equipment failure	Proxy modelling required, introducing additional assumptions for all other lines
Missing Data	Some records lacked financial exposure information	Financial risk relationships could not be fully incorporated
Missing Demographics	Worker attributes absent in one dataset	Imputed via worker_id joins, introducing minor uncertainty
Implausible Claim Values	Raw amounts outside reasonable ranges	Hard bounds applied; may marginally compress the extreme severity tail
Missing Observations	Remaining missing values post-cleaning	Complete-case filtering used; may introduce selection bias if missingness is systematic

6.2 AI Usage

Our team used AI as a supportive tool to work more efficiently throughout the project. During the technical analysis we used AI to help debug our code and provide a general sense check of our modelling outputs. For the written report we developed all the original concepts and drafted the core content ourselves before using AI to refine our spelling and sentence structure. Most importantly human actuarial judgement guided every single decision we made.

7. References

Frees, E 2012, *Regression Modeling* with Actuarial and Financial Applications, Cambridge University Press

Australian Accounting Standards Board 2023, AASB 1023, *General Insurance Contracts*
Sura Engineering Pty Ltd 2022, *SURA Engineering Equipment Breakdown Policy Wording,* <https://sura.com.au/app/uploads/2022/02/SURA-Engineering-Equipment-Breakdown-Policy-Wording-11-2018-FINAL.pdf>.

International Labour Organization. (1964). Employment Injury Benefits Convention, 1964 (No. 121). ILO.

Ignatov, Z., Kaishev, V., & Krachunov, S. (2022). Effect of stop-loss reinsurance on the primary insurer solvency. *Risks, 10*(10), 193. https://doi.org/10.3390/risks10100193

Shin, I., Oh, J. B., & Yi, H. K. (2011). Workers' compensation insurance and occupational injuries. *Safety and health at work*, *2*(2), 148–157. https://doi.org/10.5491/SHAW.2011.2.2.148

Appendix

Appendix 2.1

Appendix 2.1A – Equipment Failure Coverage Exclusions

Exclusion	Rationale
Any causes that are not sudden and unforeseen, including but not limited to: • Natural wear and gradual deterioration (predictable operating parameter within routine maintenance and servicing) • External events leading to equipment damage: including but not limited to fire, smoke, soot, chemical explosion, wind, storm, flood, hail, snow, frost, ice, water, natural disasters, nuclear hazards and war • Intentional damage or negligence • Accidental leakage or spillage • Alterations, overhauls or maintenance • Experimenting or overloading	These are additional risks to be insured through separate property insurance policies.
Consequential revenue loss	This is covered in business interruption

Appendix 2.2

Appendix 2.2A Cargo Loss Potential Product Design

Cargo loss coverage is a critical component of insurance within the interstellar mining economy. As the transportation of metals and ores underpins the entire mining and exploration value chain, leaving a fundamental operational pillar uninsured would expose Cosmic Quarry to significant balance sheet volatility and limit sustainable expansion across solar systems. While cargo loss represents a high-risk portfolio due to high value shipments such as gold and platinum, our modelling has identified specific segments and contract conditions under which Galaxy General may be able to offer coverage while remaining both profitable and competitive in the market.

Coverage Triggers

Coverage is triggered by physical loss or damage to insured cargo during transit between the declared origin and destination. Covered events include total loss, partial loss, or damage arising from operational transport hazards such as debris impact, radiation-induced container compromise, vessel collision, handling incidents, or other route-specific transit events.

Coverage is limited to operational and industrial cargo types with moderate and predictable value profiles. The following cargo categories are eligible for coverage:

- Titanium
- Lithium
- Rare earths
- Cobalt
- Operational supplies

Cargo Loss Coverage Exclusions

Exclusion	Rationale
Gold Cargo	Extremely high insured values create severe loss tail risk, requiring capital levels that make premiums commercially unviable.
Platinum Cargo	High-value concentration risk leads to extreme claim severity and capital requirements beyond the insurer's sustainable risk appetite.
Wear and Deterioration	Excluded as gradual deterioration reflects normal use rather than an insurable event.
Pre-Existing Damage	Excluded to ensure coverage applies only to losses arising during the policy period.
Packaging Defects	Excluded to maintain insured responsibility for appropriate packaging standards.
Misconduct or Fraud	Excluded to mitigate moral hazard and restrict coverage to accidental losses.
Delay Without Physical Damage	Excluded as pure delay represents an operational risk rather than property damage.
Regulatory or Legal Seizure	Excluded as such losses arise from legal or regulatory actions rather than accidental damage.
Unauthorized Transit Conditions	Excluded to enforce compliance with approved routes and security protocols assumed in underwriting.
Undeclared Cargo or Route Deviations	Excluded as undeclared changes materially alter the insured risk profile.
Systemic Catastrophic Events	Excluded or separately reinsured due to their potential for large accumulation losses.

Benefit Structure

Product Line	Policy Limit	Deductible	Settlement Terms
Standard Product	$3M	2% of Insured Value	Market Value

The cargo loss product is designed to provide targeted protection for operational shipments while maintaining a sustainable risk profile for the insurer. Coverage is limited to standard cargo types such as industrial materials and operational supplies, which typically have lower and more predictable values than precious metal shipments.

A policy limit of $3M per shipment is applied to ensure meaningful protection against material cargo losses while containing the insurer's exposure to extreme severity events. This limit aligns with the value profile of most operational cargo and helps maintain capital efficiency within the insurer's portfolio.

Claims are settled on a market value basis, reflecting the prevailing value of the cargo at the time of loss. This approach is appropriate for standard cargo as these goods generally have stable or depreciating market values, allowing claims to be settled based on the actual financial loss incurred.

A deductible of 2% of the insured value is applied to ensure proportional risk sharing between the insurer and the insured. This structure scales with shipment value, discourages small attritional claims and reducing claims administration costs while preserving coverage for material losses. Cargo types with extremely high insured values, such as gold and platinum, are excluded from coverage due to their disproportionate loss severity and the significant capital requirements needed to insure them sustainably. By focusing coverage on standard cargo, the product remains commercially viable while still addressing a meaningful operational risk faced by Cosmic Quarry's interstellar transport network.

Appendix 2.3

Appendix 2.3A – WC Coverage Inclusions

Trigger	Description	Evidence Required
Occupational injury	Physical injury during operational duties on any CQ site or vessel	Incident report and medical assessment
Gravity related trauma	Musculoskeletal, impact, or cardiovascular injuries from non-standard gravity	Medical report citing gravity exposure and operational log
Life support failure	Medical costs and disability from life-support system malfunction	Engineering failure log and medical records
Psychological injury	Stress, trauma, or PTSD arising from occupational events	Certified psychological assessment
Radiation exposure	Health impacts from documented radiation events above operational thresholds	Radiation monitoring logs and medical assessment
Fatal injury	Lump sum and ongoing dependant payments for occupational fatalities	Coroner's report and employment records

Appendix 2.3B – WC Coverage Exclusions

Exclusion	Rationale
Pre-existing conditions	Outside insurable interest, must be disclosed at underwriting
Off-duty incidents	Outside operational scope of the policy
Intoxication-related injury	Gross misconduct exclusion
Self-inflicted injury	Not arising from an occupational hazard
Acts of war or deliberate sabotage	Uninsurable systemic risk
90-day policy seasoning period	Adverse selection and moral hazard protection at inception

Appendix 2.3C - WC Gross vs Net Reinsurance Plot

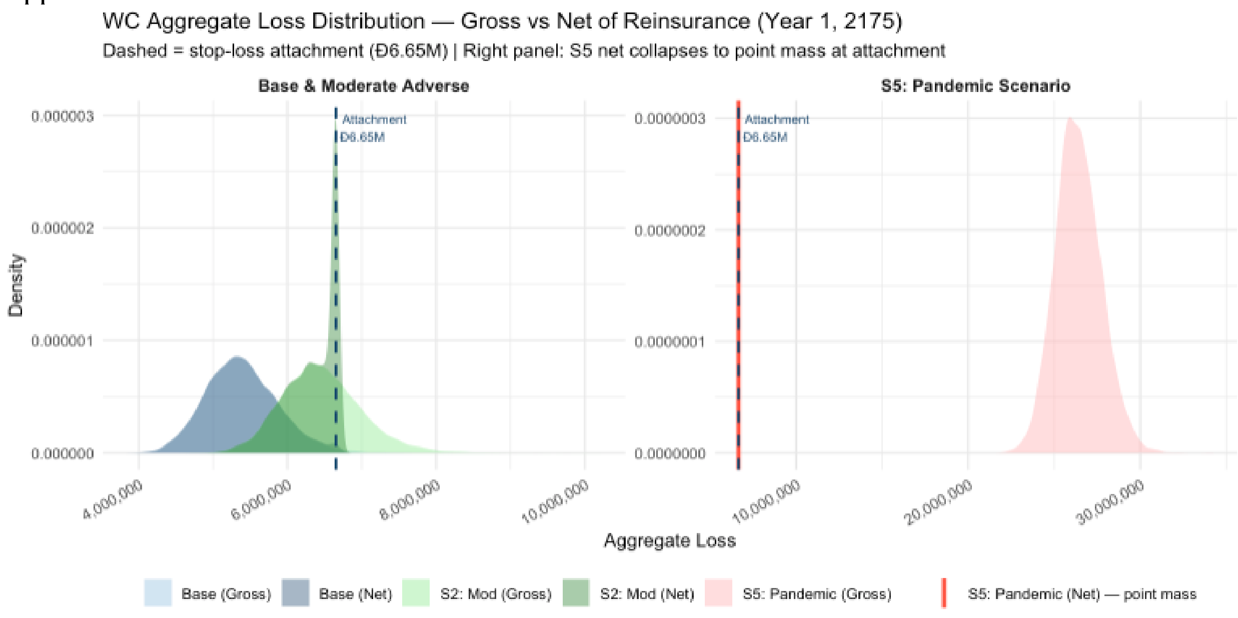

Appendix 2.3D - WC Projected Aggregate Loss

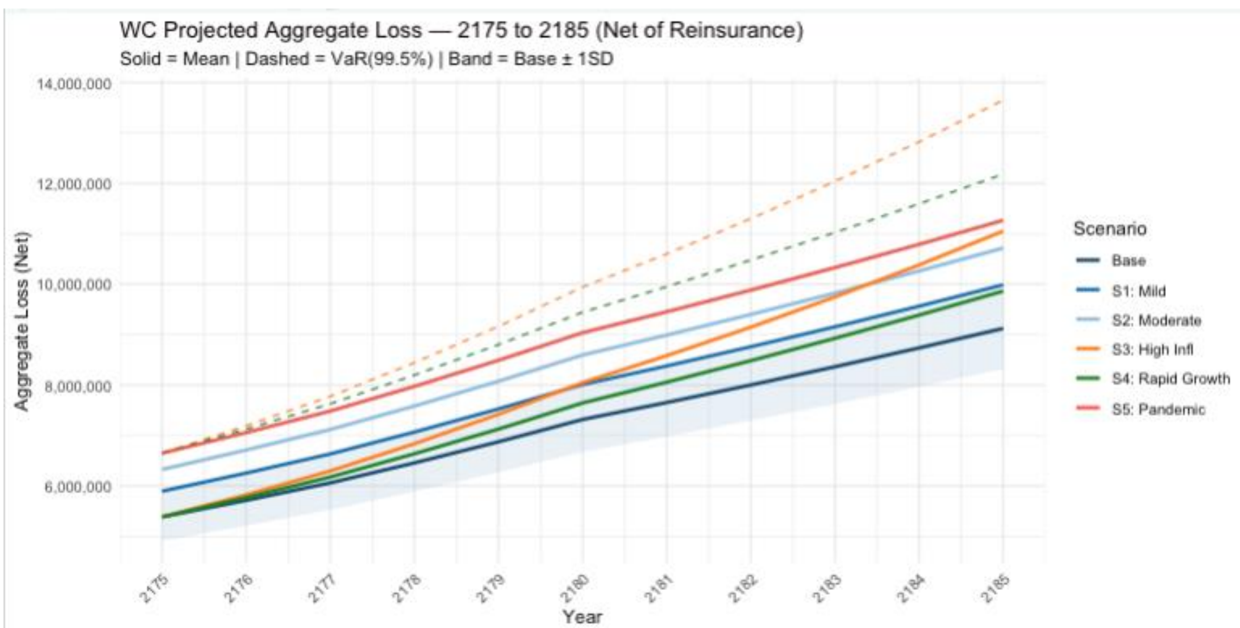

Appendix 2.3E - WC Net Revenue for Reinsurance

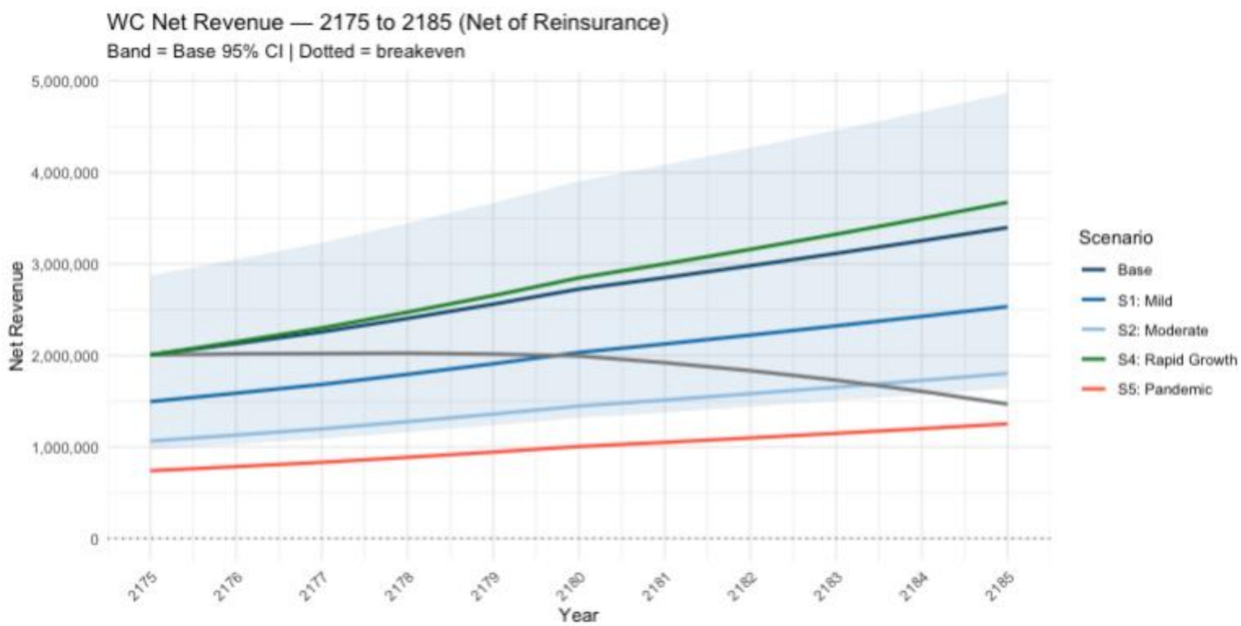

Appendix 2.3F - WC Methodology
Across 10,000 Monte Carlo simulations, the Base scenario produces an expected annual aggregate loss of Đ5.39M and VaR(99%) of Đ6.65M, with $\sigma^2 = 1.73$ on the fitted log-normal severity distribution reflecting wide dispersion in individual claim costs.

Appendix 2.4

Appendix 2.4A
Coverage inclusions

Trigger Category	Event Description	Evidence Required
Equipment Failure Shutdown	The stoppage of extraction operations caused by failed by transport equipment	Operational output log proving zero output during the interruption period
Partial Output Reduction	Equipment malfunction results in over 30% loss in normal output	Operational output log proving a minimum 30% reduction
Supply Chain Interruption	Extracted commodities are unbaled to be transported to designated delivery points	Transport system logs
Mandatory Evacuation	Catastrophic or radiative event requiring a station evacuation	Official evacuation order or safety authority notification
Debris Field Closure	Debris density exceeds government mandated safety thresholds	Closure notice and debris logs
Crew Safety Shutdown	Critical workforce capacity no longer operationally sufficient	Workers' compensation records

Appendix 2.4B
Coverage Exclusions

Exclusion Category	What is Excluded

Scheduled Maintenance	Pre-determined operational shutdowns including repair and machine testing.
Gradual Deterioration	Progressive standard wear or degradation of machinery.
Market Price Fluctuation	Commodity price volatility resulting in revenue loss.
Pre-existing Faults	Known equipment defects and risks known prior to policy inception.
Acts of War or Sabotage	Losses caused by hostile action piracy or armed conflict.
Double Recovery	Loss already recoverable under Equipment, Cargo, or Workers' Compensation policies.
New Policy Seasoning	Claims occurring within the first 90 days of a new policy.

Appendix 2.4C (Pure Premium by Solar System and Product)

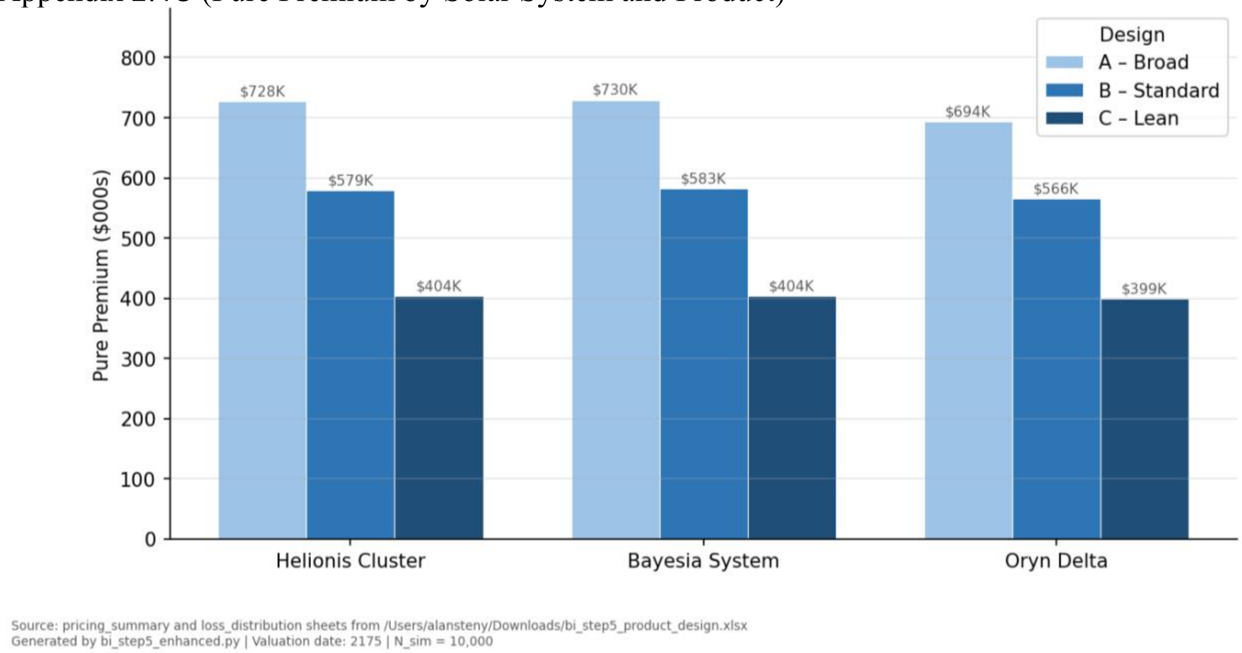

Appendix 2.4D

Product	Deductible	Per-Occurrence Limit	Waiting Period	Annual Aggregate Limit
A: Broad Coverage	$0.5m	$25m	7 days	$75m
B: Standard Coverage	$1.0m	$15m	14 days	$45m
C: Lean / Catastrophe Layer	$2.0m	$10m	30 days	$30m

Appendix 3.1

Appendix 3.1A – Equipment Failure

To select the frequency model, the dataset was split into an 80% training set and a 20% testing set to evaluate both in-sample fit and out-of-sample predictive performance. The standard Negative Binomial (NB) was compared against a full Zero-Inflated Negative Binomial (ZINB) model using AIC and BIC for goodness-of-fit, along with out-of-sample RMSE and prediction accuracy on the testing dataset. The NB model was selected based on the best overall balance between predictive performance, goodness-of-fit, and model interpretability.

The explanatory variables included in the final NB model were selected based on both their statistical significance and patterns observed during exploratory data analysis (EDA). The final model incorporates equipment age, maintenance interval, usage intensity, equipment type, solar system, and exposure. All included variables were highly statistically significant ($p < 0.001$), and EDA indicated that they meaningfully influence claim frequency. Equipment type, solar system, and equipment age were identified as the most influential predictors, as illustrated in the figures below.

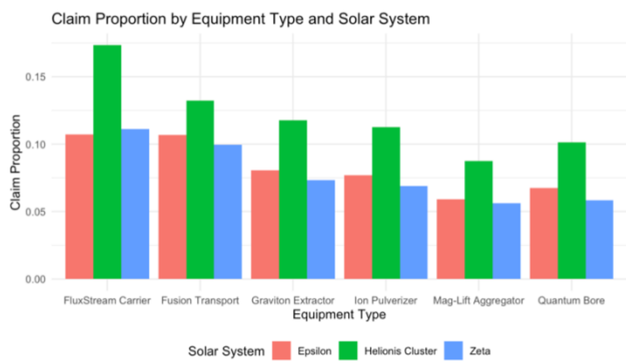

Appendix 3.1B – Cargo Loss

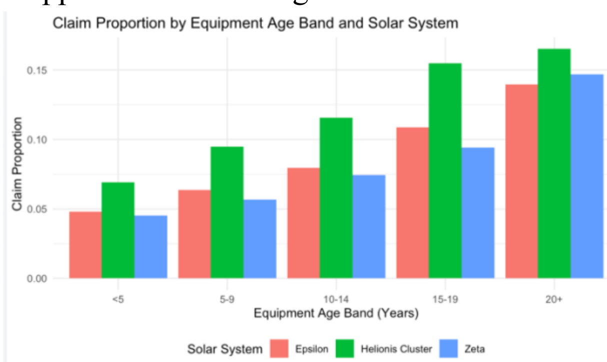

The dataset was split into training and testing sets (80/20). A Poisson GLM was initially fitted as the baseline frequency model; however, substantial overdispersion was observed (Pearson dispersion = 4.64), indicating that the variance of claim counts exceeded the mean. A Negative Binomial GLM was therefore adopted to accommodate this additional variability. A reduced specification was then developed to improve the balance between interpretability and accuracy, retaining route risk, container type, debris density, and pilot experience, which were the most statistically significant predictors of claim frequency and were also supported by the EDA.

Model	AIC	BIC
Poisson GLM	159288.5	159506.6
Negative Binomial Full	140612.0	140839.5
Negative Binomial Reduced	140606.5	140710.8

Model selection criteria such as AIC and BIC were used to compare the full and reduced models. The reduced Negative Binomial model achieved the lowest AIC and BIC, indicating that the removed variables did not materially contribute to explaining claim frequency.

Model	RMSE	MAE
Poisson GLM	0.6405548	0.4310131
Negative Binomial Full	0.7013468	0.5532492
Negative Binomial Reduced	0.7011412	0.5534879

To assess predictive performance, the candidate models were evaluated on the test dataset. While the Poisson model produced marginally lower RMSE and MAE values, it violates the Poisson mean variance assumption. The Negative Binomial model provides a more appropriate representation of the claim frequency process by allowing the variance to exceed the mean. The reduced Negative Binomial model achieved nearly identical predictive performance to the full model, indicating that the removed variables did not materially affect predictive accuracy. Therefore, the reduced Negative Binomial model was selected as the final frequency model due to its improved simplicity and appropriate treatment of overdispersion.

Appendix 3.1C – Workers Compensation

The dataset was split into training and testing sets using an 80/20 random partition, yielding 105,205 training and 26,302 test records for the frequency model, and 1,520 training and 381 test records for the severity model.

A Poisson GLM was initially considered as the baseline frequency model. All Negative Binomial specifications produced θ estimates ranging from 7.84 (M4) to 13.52 (M1), confirming genuine overdispersion the variance of claim counts materially exceeds the mean and validating the adoption of the Negative Binomial family across all candidate models.

Five candidate models were evaluated. Models M1 and M2 were excluded from pricing use as they include predictors not available in the CQ portfolio (e.g., psych_stress_index, protective_gear_quality, supervision_level). M6 was also excluded as the addition of safety_training_index over M4 produced only marginal improvement. A Zero-Inflated Negative Binomial (M7) was evaluated against M4 on the same CQ predictor set; the AIC difference was −19.5 in favour of M4, indicating ZINB did not improve fit and added unnecessary complexity.

Table: Frequency Model Information Criteria

Model	Predictors	AIC	CQ-Deployable
M1: Full NB	Occupation, employment type, solar system, accident history, psych stress, hours/week, safety training, protective gear, experience, supervision, gravity, salary	15,559.00	No
M2: EDA NB (stress test)	Occupation, accident history, psych stress, safety training, experience, supervision	15,559.30	No
M6: CQ+Safety NB	Occupation, solar system, experience, safety training	15,575.60	No
M4: CQ NB ★	Occupation, employment type, solar system, experience	15,593.60	Yes
M7: CQ ZINB	Occupation, employment type, solar system, experience	15,613.20	Yes

Table : Frequency Model Predictive Performance (Test Set)

Model	RMSE	MAE	Test Deviance
M1: Full NB	0.119668	0.0283	3,055.60
M2: EDA NB (stress test)	0.119668	0.0283	3,051.30

M6: CQ+Safety NB	0.119668	0.0283	3,051.10
M4: CQ NB ★	0.119668	0.0283	3,053.40
M7: CQ ZINB	0.119668	0.0283	3,053.40

All five models produced virtually identical RMSE and MAE on the test set, indicating that the additional predictors in M1 and M2 do not improve out-of-sample predictive accuracy. The likelihood ratio tests confirmed each successive addition of predictors produced statistically significant improvements in in-sample fit (M4→M6: $\chi^2(3) = 23.26$, $p < 0.001$; M6→M2: $\chi^2(4) = 25.80$, $p < 0.001$; M2→M1: $\chi^2(13) = 25.61$, $p = 0.019$), though this did not translate to improved test-set performance, suggesting the additional variables capture within-sample noise rather than generalizable signal. M4 (CQ-Deployable Negative Binomial) was therefore selected as the final pricing model, with M2 retained as the stress testing frequency model due to its richer predictor set.

Appendix 3.2

This choice is motivated across each product line with WC, injury costs span a wide range from minor strains to catastrophic life support failures. With Equipment Failure, the cost of repairs increases multiplicatively with the age and kind of equipment. The structural variations in loss ratio distributions among cargo categories are captured by stratified Log-Normal GLMs by cargo tier (Gold, Platinum, Other) for cargo loss. Downtime expenses for business interruptions increase multiplicatively with the length of the interruption and the size of the activity.

Appendix 3.2A – Equipment Failure

Exploratory data analysis indicated that claim severities follow a log-normal distribution, which was confirmed using a Q-Q plot of the logarithm of claim amounts (see Figure below). Consequently, a Generalised Linear Model (GLM) with a Gaussian distribution and log link was used to model claim severity.

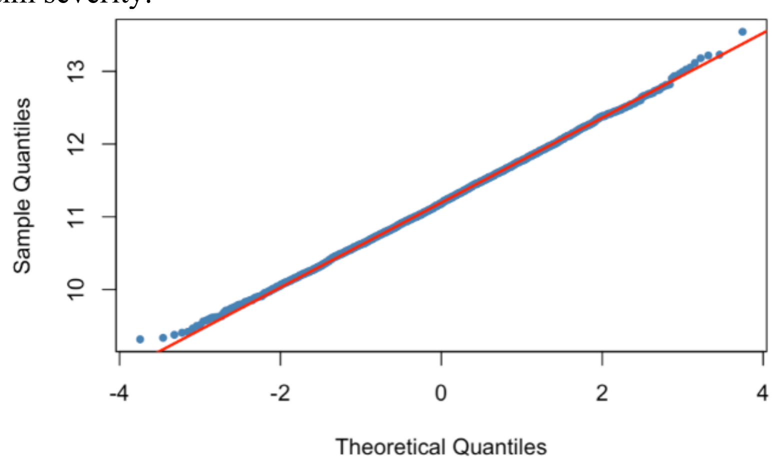

A full model incorporating all potential explanatory variables was initially fitted. A two-way stepwise model selection procedure was then applied to remove variables that did not significantly contribute to model performance. The resulting reduced model retained equipment type, usage intensity, and solar system as predictors of claim severity.

The relationships identified in the final severity model are consistent with the patterns observed during the exploratory data analysis. As illustrated in the figure below, claim amounts vary systematically across both equipment type and solar system. Equipment such as Quantum Bore exhibit noticeably higher median claim amounts and wider dispersion compared to other equipment types, indicating a greater potential repair or replacement cost when failures occur. Additionally, differences across solar systems are evident, with claims in Zeta generally displaying higher medians and variability relative to Helionis Cluster and Epsilon.

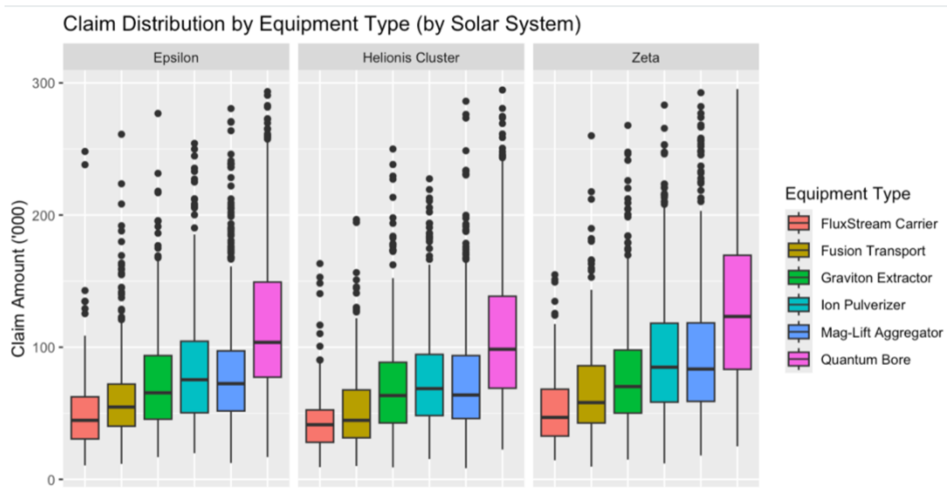

Appendix 3.2B – Cargo Loss

During exploratory data analysis (EDA), the severity histogram exhibited a bimodal shape (see Figure), indicating the presence of two distinct peaks in the loss distribution. Further investigation revealed that claim severities differed substantially by cargo type. When analysed separately, the severity distributions for gold, platinum, and other cargo types each followed approximately lognormal distributions, with gold exhibiting the highest mean severity, platinum the second highest, and the remaining cargo types showing relatively similar mean values.

A lognormal GLM was initially fitted to the full dataset, however the resulting fit underestimated the tail behaviour and did not adequately capture the observed loss distribution. This suggested that a single severity model was not appropriate across all cargo types. Consequently, the dataset was segmented into gold, platinum, and a combined "other" category, and separate lognormal models were fitted to each subset with their own parameters. The resulting fits showed a much stronger alignment with the

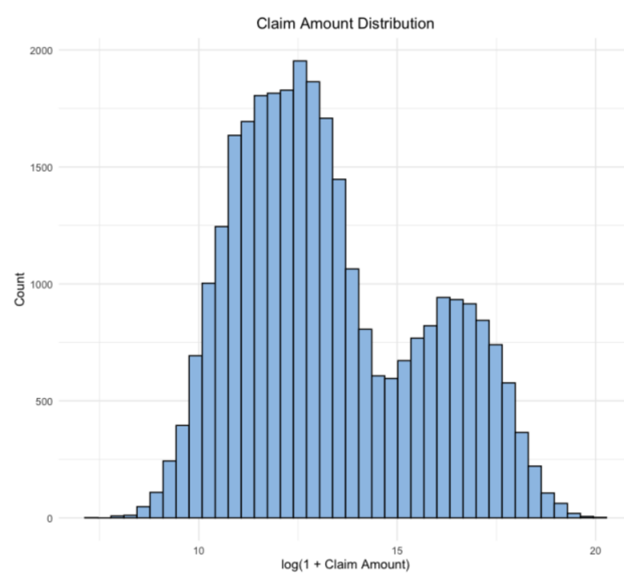

Figure: Log claim amount distribution

observed data, which is supported by the QQ plots presented in Figure where all three models demonstrate a good fit across the distribution.

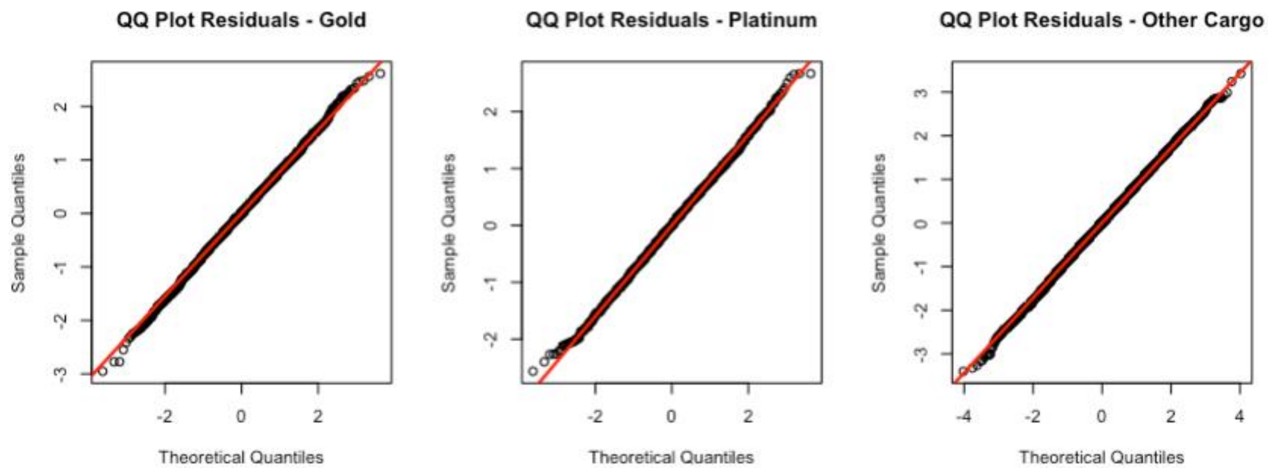

Figure: QQ plots for lognormal models by cargo type

Appendix 3.2C – Workers Compensation Severity Model

The marginal distribution of claim amounts was assessed prior to GLM fitting. The severity distribution exhibited strong right-skew (skewness = 4.379, kurtosis = 25.319), with the mean ($8,114) considerably exceeding the median ($2,043) and the P99 reaching 12.8× the mean. The mean excess plot showed a persistently upward trend, consistent with a heavy-tailed distribution. Parametric MLE fits confirmed the log-normal as the best-fitting marginal family.

Table: Marginal Distribution Fit (claims ÷ 1,000)

Distribution	AIC	BIC
Log-Normal	8,001	8,012
Pareto	8,113	8,123
Weibull	8,582	8,593
Gamma	8,891	8,902

Five GLM specifications were then fitted. The Gamma family (S1, S3) and Inverse Gaussian (S6) were evaluated as alternatives to the log-normal. The Tweedie family was excluded due to package conflicts and unreliable convergence on this dataset. S1 and S4 include injury type and cause, which are not available in the CQ portfolio and were therefore restricted to stress testing use.

Model	Distribution	AIC	RMSE	MAE	MAPE (%)	R^2	CQ-Deployable
S5: Lognormal CQ ★	Log-Normal	5,178.30	16,783	7,730	369	0.0285 (Adj-R^2)	Yes
S4: LogNormal	Log-Normal	5,047.60	16,641	7,638	315.2	0.119 (Adj-R^2)	No

EDA (stress test)							
S6: InvGaussian CQ	Inverse Gaussian	28,805.20	17,031	9,318	514.2	0.0203 (Pseudo-R^2)	Yes
S3: Gamma CQ	Gamma	29,892.40	16,985	9,218	507	0.0406 (Pseudo-R^2)	Yes
S1: Gamma Full	Gamma	29,762.00	17,175	9,003	450.9	0.127 (Pseudo-R^2)	No

Among CQ-deployable models, S5 achieved the lowest AIC by a considerable margin (5,178 vs. 28,805 for S6), confirming the log-normal family as substantially more appropriate for this severity distribution than Gamma or Inverse Gaussian. S5 also produced the lowest MAE of all CQ-deployable models.

A portfolio bias check was conducted by comparing mean predicted claim amounts against the test set mean of $6,815.

Table: Portfolio Bias Check (Test Set Mean = $6,815)

Model	**Mean Predicted ($)**	**Bias Ratio**	**Bias (%)**
S6: InvGaussian CQ	8,104	1.19	0.189
S1: Gamma Full	8,049	1.181	0.181
S3: Gamma CQ	8,008	1.175	0.175
S4: LogNormal EDA (stress)	6,341	0.93	−6.96%
S5: LogNormal CQ ★	6,163	0.904	−9.57%

The log-normal retransformation correction applies the bias-correction multiplier when converting log-scale predictions back to dollar scale, where $\sigma^2 = 1.7306$ yields a multiplier of 2.3757. Despite this correction, S5 returned a bias ratio of 0.904, indicating a −9.6% systematic under-prediction at the portfolio level. While this exceeds the ±5% internal threshold, the Gamma and Inverse Gaussian models over-predicted by approximately +18%, which is the more dangerous direction for pricing adequacy. S5 (LogNormal CQ) was therefore selected as the final pricing severity model, with the known downward bias noted as a conservative pricing feature. S4 was retained as the stress testing severity model, benefiting from the additional signal from injury type and injury cause.

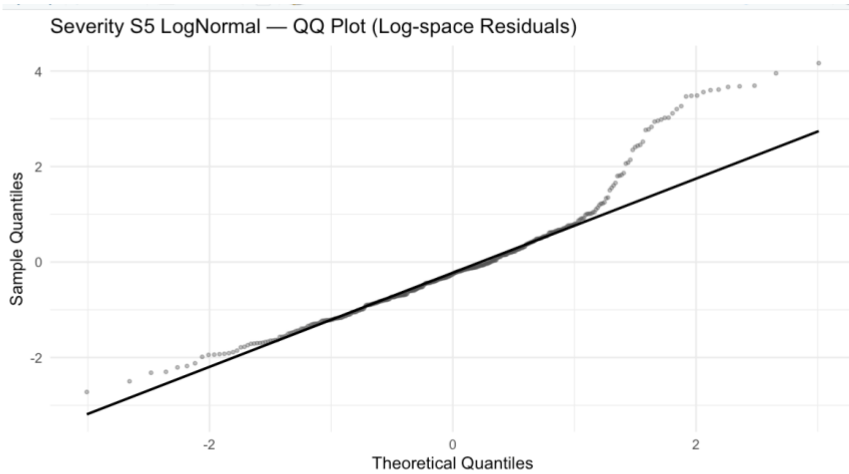

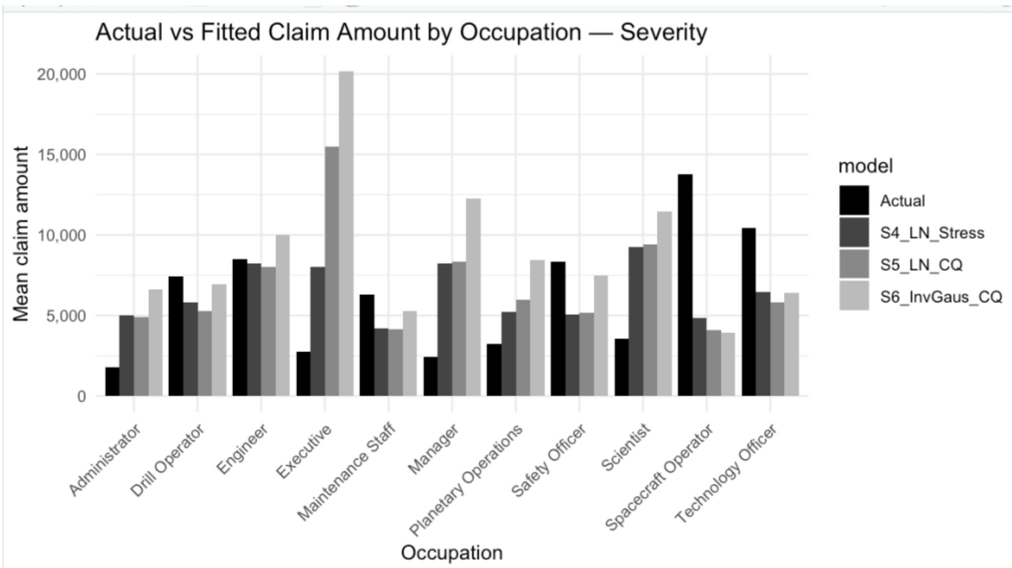

Appendix 3.2D – Business Interruption

Appendix 3.3

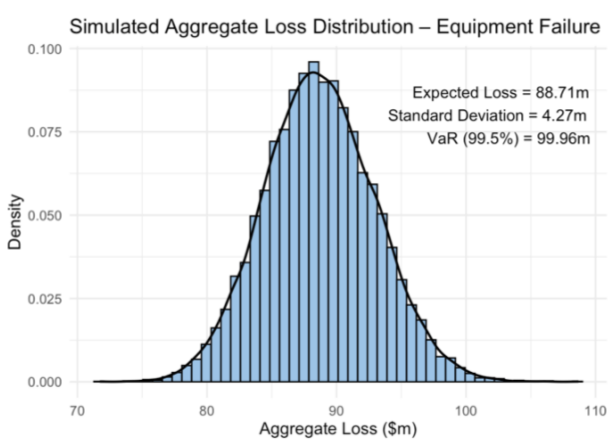

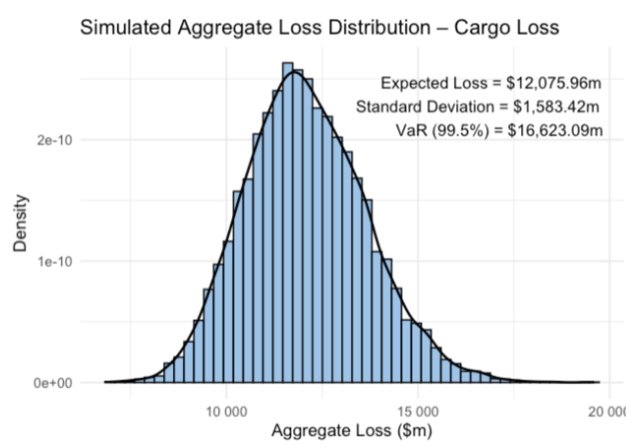

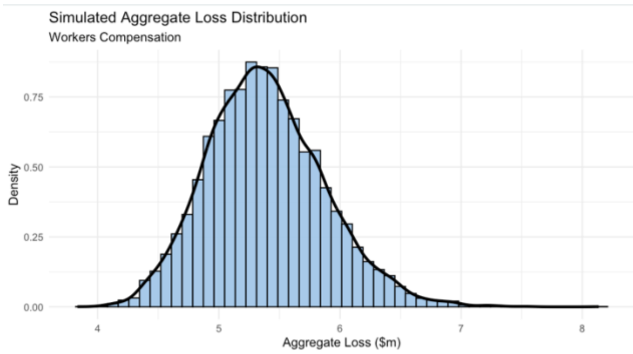

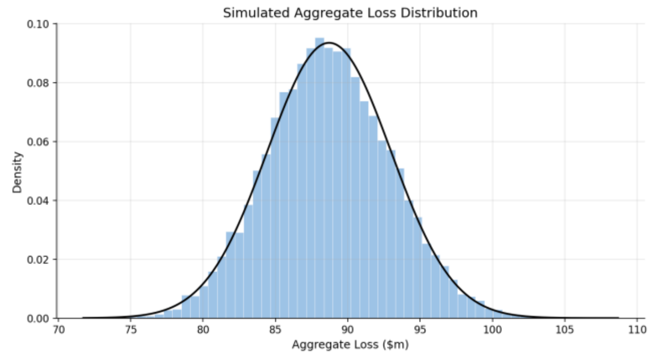

Appendix 3.4

Appendix 3.4A – Capital Modelling Methodology

This appendix sets out the methodology used to project portfolio capital requirements, profitability and solvency for Galaxy General Insurance Company's proposed Cosmic Quarry portfolio. The model was designed to meet the case study requirement to provide short- and long-term ranges for aggregate costs, returns and net revenue, including expected values, variance and tail behaviour, with additional allowance for stress testing and multi-system dependency.

Modelling Objective

A stochastic capital model was developed to assess the solvency, profitability, and capital adequacy of Galaxy General's proposed insurance portfolio for Cosmic Quarry Mining Corporation. The model projects portfolio performance over a 10-year horizon (2176–2185) and evaluates expected returns, variability, tail risk, recapitalisation requirements, and probability of ruin.

The framework supports the case study requirement to analyse aggregate loss distributions, profitability, and capital adequacy under both baseline and stressed environments.

Portfolio Structure

The portfolio consists of four insurance lines across three operating systems:

- Cargo Loss
- Equipment Failure
- Workers' Compensation
- Business Interruption

For each solar system $s \in$ {Helionis, Bayesia, Oryn}, aggregate losses are defined as

$$L_s = L_{Cargo,s} + L_{Equip,s} + L_{WC,s} + L_{BI,s}$$

Workers' compensation losses are retained separately to enable targeted stress testing.
Aggregate portfolio losses are obtained by summing system-level losses

Premium Calibration

Premiums are determined using a cost-of-capital pricing framework based on the simulated aggregate loss distributions.

Key quantities calculated for each line include:

- Expected loss $E[L]$
- Value-at-Risk at the 99.5th percentile

- Tail Value-at-Risk (TVaR)

Required capital is defined as

$$SCR = VaR_{99.5\%}(L) - E[L]$$

A risk margin equal to the cost of capital applied to required capital is added to the expected loss. The technical premium is therefore

$$Premium = \frac{E[L](1 + \text{Claims Expense}) + (\text{Cost of Capital} \times SCR)}{1 - \text{Operating Expense}}$$

System-level premiums are obtained by aggregating premiums across all product lines

<u>Projection Framework</u>

The capital model simulates portfolio outcomes over 10 years using Monte Carlo simulation. For each simulation path and year, the model:

1. Bootstraps aggregate losses from historical distributions
2. Scales losses for exposure growth and inflation
3. Applies stress and dependency adjustments
4. Calculates reinsurance recoveries
5. Computes underwriting and investment results
6. Updates capital and solvency metrics

This structure preserves the empirical shape of the aggregate loss distributions while allowing macroeconomic and operational assumptions to evolve through time.

<u>Reinsurance Structure</u>

The portfolio includes a proportional quota-share treaty with the following parameters:

- Quota share cession = 40%
- Ceding commission = 25%

Net premium is therefore

$$Premium_{net} = Premium_{gross} - Premium_{ceded} + Commission$$

Net losses are

$$Loss_{net} = Loss_{gross} - Recovery$$

This structure reduces retained volatility while maintaining underwriting revenue.

<u>Cross-System Dependency</u>

Losses across solar systems are partially correlated through a shared annual shock. A lognormal common factor X_t is simulated for each year:

$$X_t \sim Lognormal(\mu, \sigma)$$

with parameters chosen such that

$$E[X_t] = 1$$

The dependency multiplier applied to all systems is

$$M_t = (1 - w) + wX_t$$

where:

- Dependency weight w = 35%
- Shock volatility σ = 15%

This introduces correlated loss deterioration across systems while preserving individual system variability.

Workers' Compensation Stress

Workers' compensation is modelled separately to allow targeted deterioration scenarios. Under stress conditions, WC losses increase by

$$WC_{stress} = 1.40 \times WC$$

from year 3 onward.

This reflects the potential impact of injury severity inflation, regulatory changes, or operational hazards

Capital Dynamics

At each projection step, underwriting profit is calculated as

$$UW = Premium_{net} - Loss_{net} - Expenses$$

where

$$Expenses = 25\% \times Premium_{gross}$$

Investment income is

$$Investment = Capital \times r$$

Total annual result is

$$Result = UW + Investment$$

Capital evolves as

$$Capital_{t+1} = Capital_t + Result$$

Capital Targets

Three initial capital levels are tested:

1. 200% SCR – Strong Capital Buffer
2. 150% SCR – Adequate Solvency
3. 100% SCR – Regulatory minimum

Initial capital is defined as

$$Capital_0 = Multiplier \times SCR_{Year1}$$

Surplus Distribution and Recapitalisation

If capital exceeds 150% of required capital, 25% of excess capital is distributed. If capital falls below zero, the model allows a single recapitalisation event, restoring capital to the SCR level. This ensures recapitalisation is treated as a severe management intervention rather than a routine adjustment

Solvency Metrics

Key solvency indicators produced by the model include:

- Solvency ratio
- Required capital
- Recapitalisation probability
- Probability of ruin

The solvency ratio is defined as

$$Solvency = \frac{Capital}{SCR}$$

Probability of ruin represents the proportion of simulation paths where capital becomes negative

Model LimitationsSeveral limitations should be acknowledged:

- Dependency is represented by a single common shock rather than a full copula structure
- Loss distributions are bootstrapped from historical simulations
- Reinsurance is simplified as proportional quota share
- Premiums are assumed constant apart from inflation adjustments

Despite these simplifications, the framework provides a transparent and robust basis for evaluating capital adequacy and portfolio resilience.

Appendix 3.4B - A Year-by-Year Net Claims Trajectory — 150% SCR

Year	E[Net Claims]	Std Dev	P99	TVaR 99.5%	E ± 2σ Range
2176	$9.138B	$1.001B	$11.685B	$12.37B	$7.14B – $11.14B
2177	$9.648B	$1.057B	$12.323B	$13.08B	$7.53B – $11.76B
2178	$10.241B	$1.115B	$13.072B	$13.89B	$8.01B – $12.47B
2179	$10.905B	$1.180B	$13.964B	$14.78B	$8.54B – $13.26B
2180	$11.599B	$1.286B	$14.866B	$15.76B	$9.03B – $14.17B
2181	$12.131B	$1.320B	$15.442B	$16.36B	$9.49B – $14.77B
2182	$12.666B	$1.379B	$16.201B	$17.15B	$9.91B – $15.42B
2183	$13.240B	$1.471B	$16.880B	$18.14B	$10.30B – $16.18B
2184	$13.844B	$1.497B	$17.684B	$18.75B	$10.85B – $16.84B
2185	$14.477B	$1.598B	$18.504B	$19.45B	$11.28B – $17.67B

Appendix 3.4C – Year-1 Claims Breakdown by Coverage Line (Gross, pre-reinsurance)

Line	System	E[Gross Loss]	VaR 99.5%	Gross Premium	Loss Ratio
Cargo	Helionis Cluster	$4.69B	$7.66B	$7.48B	62.72%
Cargo	Bayesia System	$3.75B	$6.09B	$5.97B	62.77%
Cargo	Oryn Delta	$3.63B	$6.46B	$5.88B	61.80%
Equipment	Helionis Cluster	$0.06B	$0.07B	$0.09B	65.82%
Equipment	Bayesia System	$0.02B	$0.03B	$0.03B	65.15%
Equipment	Oryn Delta	$0.01B	$0.02B	$0.02B	64.32%
Workers Compensation	Helionis Cluster	$0.00B	$0.00B	$0.00B	64.32%
Workers Compensation	Bayesia System	$0.00B	$0.00B	$0.00B	62.70%
Workers Compensation	Oryn Delta	$0.00B	$0.00B	$0.00B	62.31%
Business Interruption	Helionis Cluster	$1.24B	$1.55B	$1.90B	65.26%
Business Interruption	Bayesia System	$0.59B	$0.80B	$0.92B	64.48%
Business Interruption	Oryn Delta	$0.36B	$0.52B	$0.56B	63.78%

Appendix 3.4D – Stress Tests

Stress test	Description	Worst solvency ratio	Ruin prob (yr 10)	Solvency breach?
Inflation Mismatch	Claims inflation +200bp, premium inflation +50bp	0.89x	1.5%	Yes
Low Rates	Investment return −200bp	1.93x	0.0%	No
Workers Comp Shock	WC losses ×1.40 from year 3	2.05x	0.0%	No
Oryn Stress	Oryn Delta losses ×1.20 + exposure growth −2%	1.60x	0.1%	No
Combined Severe	Frequency +20%, inflation +200bp, rates −200bp, WC shock	0.71x	14.5%	Yes
Reverse Stress	Losses ×1.40 all systems + WC ×1.25	0.00x	100.0%	Yes

Appendix 4.2

Threat	Description	Mitigation Strategy
Pricing Risk	The risk that modelled premiums fail to capture actual future loss experiences, particularly in the tail. For instance, the Log-Normal severity model may underestimate the financial impact of unpredictable solar flares from the dim dwarf star in the Oryn Delta system.	Implement a multiplicative bias correction during simulations. Continually true-up Log-Normal assumptions against emerging claims data and utilise the Q-RISK Engine to stress test heavier tailed distributions.
Financial Risk	The risk that macroeconomic pricing assumptions deviate significantly from expectations. Systemic economic shocks would alter the yield rate curves used for discounting expected losses and the interplanetary inflation rates, severely impacting the projected 10-year profitability.	Base premiums are already modelled to incorporate expected annual interplanetary inflation. To mitigate extreme inflation shocks, Galaxy General will continuously monitor macroeconomic indicators and adjust premium rates during annual reviews if actual inflation outpaces our projections.
Ethical Risk	The risk of compromising the social insurance mandate that underpins WC by failing to adequately protect severely injured personnel or discriminating against specific worker classifications.	Maintain the unlimited per-claim liability for WC to ensure seriously injured workers are made whole. Protect Galaxy General's capital from this unlimited exposure via the proposed portfolio stop-loss reinsurance attaching at the 99% VaR.

Appendix 4.4

Appendix 4.4A

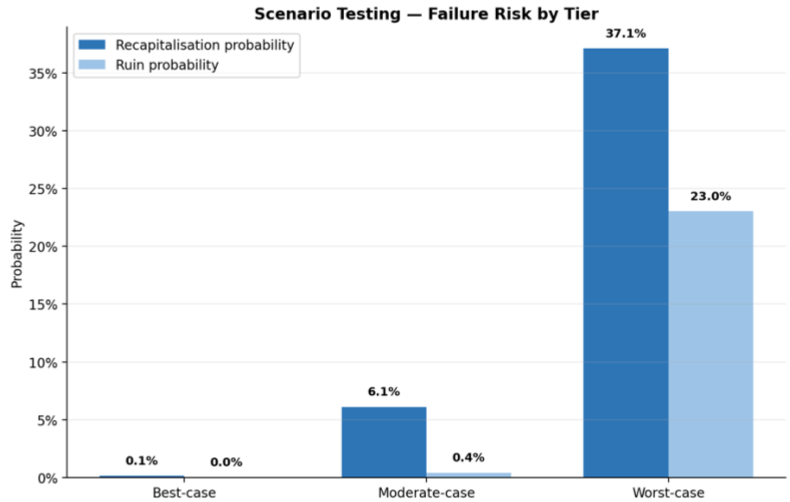

.

Appendix 4.4B

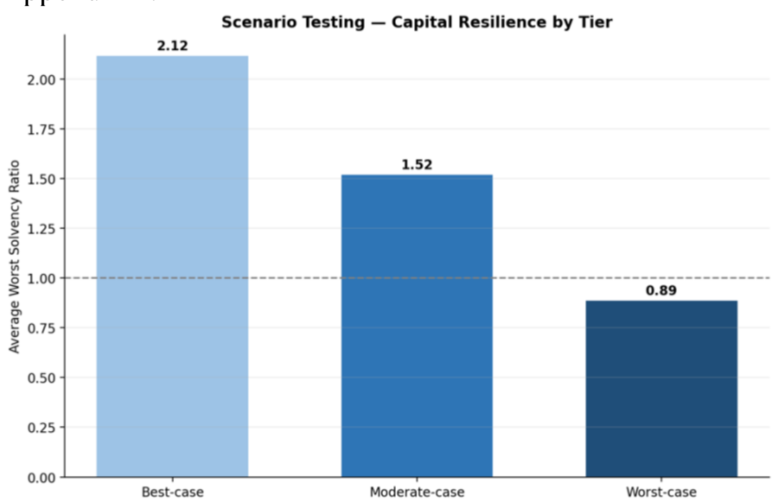

Appendix 4.4C

Scenario	Description	Worst solvency ratio	vs Base	Recap prob (yr 10)	Ruin prob (yr 10)
Helionis Catastrophe	An isolated 60% loss spike in year three falling completely outside standard attritional frequency projections.	1.66x	-0.40x	0.8%	0.0%
Solar Storm	A massive external event impacting all three solar systems simultaneously introducing severe cross system dependencies uncaptured by the base model.	0.85x	-1.20x	10.3%	0.2%
Equipment BI Cascade	An initial equipment failure triggering compounding business interruption loss over	0.81x	-1.24x	19.1%	0.5%

	two consecutive years to test unmodelled causal chains.				
Portfolio Contraction	A sudden business volume loss from external market disruptions like regulatory shutdowns or client insolvency falling outside standard growth projections.	2.18x	+0.13x	0.1%	0.0%

Appendix 4.4D

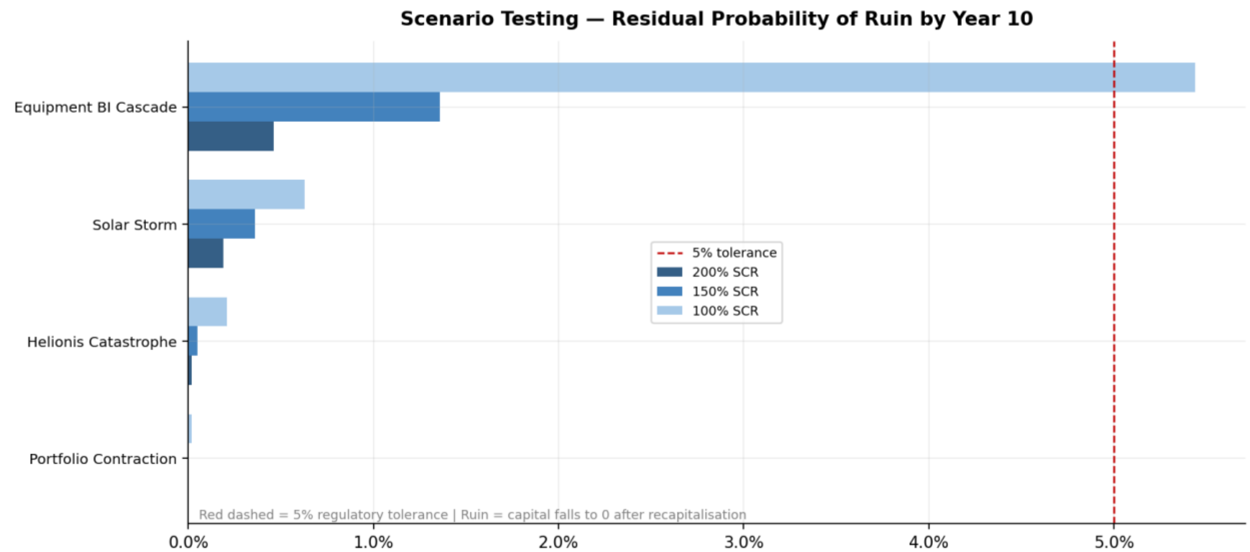

Appendix 4.4E

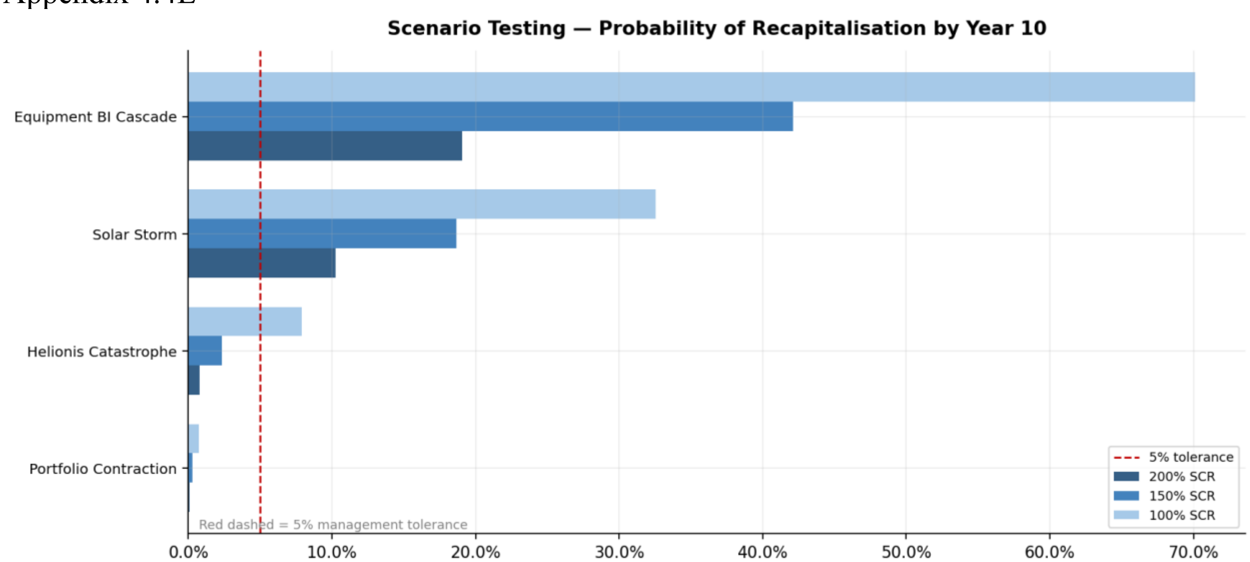

Appendix 5.1

Appendix 5.1A

Metric	Assumed value	Reasoning behind assumption	The impact it could have

Inflation & Discount Rate	2175: same rate as 2174 2176 - 2177: 3 years' average (1 yr spot rate for discount) 2178 - 2180: 5 years' average (1 yr spot rate for discount) 2181 - 2185: 10 years' average (10 yr spot rate for discount)	We adopted a simplified approach to AASB1023 guidance, where the discount rates are selected to reflect the different nature and term of the obligations across each future period (AASB, 2023).	
New Business Development Rate	Helionis Cluster & Bayesia System: linear development to achieve 25% increase over the next 10 years Oryn Delta: linear development to achieve 15% increase over the next 10 years	This aligns with Cosmic Quarry's expected expansion in each respective solar system over the next 10 years.	
Expense ratios & cost of capital	Operating expense: 25% of gross premium. Claims expense: 12% of net loss. CoC rate: 12%	Proxies based on industry norms for specialised commercial lines; no internal data to calibrate against for a novel line	A 5pp move in the operating ratio shifts the combined ratio by ~5pp and alters technical premium
Capital management (distribution & recapitalisation)	Surplus distributed at 25% of capital above 150% SCR. Recapitalisation triggered at 0% SCR, restored to 100% SCR; max 1 injection per simulation path	Distributes excess capital to maintain ROE discipline; single-injection cap reflects realistic constraints on raising capital under distress	Higher distribution rates reduce the solvency buffer available under stress.
Diversification	Before projecting capital, the model compares: • sum of stand-alone required capitals by solar system vs • diversified required capital of the combined portfolio	This tells you how much capital is saved because not all bad outcomes happen simultaneously.	Diversification is captured by aggregating simulated system losses, but may be overstated if cross-system dependencies were not explicitly modelled.
Reinsurance (WC)	Portfolio stop loss set to 6650395 (VaR 99%).	Attachment set at VaR 99% to ensure that it only applies for the tail risk	It aggregates grow due to inflation and

	Reinsurance Premium load on expected cred loss.	events under base conditions.	headcount growth. The reinsurance section triggers more in later years. In practice this attachment should be negotiated annually.

Appendix 6.1

Appendix 6.1A

Steps taken to improve incomplete, incorrect claims data:

- Missing data fields
  Each column within the historical claims datasets contains several blank fields. We reconstructed the dataset by cross-referencing between severity and frequency sheet. This is achieved through unique identifiers created by combining variables such as policy_id and other ID variable (i.e. equipment_id for equipment failure) which was then matched between the datasets.

- Inconsistent values or typos
  Character columns contained "_???" suffixes. These suffixes were removed to retain valid prefix, allowing data standardisation and ensured consistent modelling inputs.

- Negative data and data outside sensible range
  Variables contained values outside physically or contractually plausible ranges (i.e. negative values in age, decimal claim counts). This was corrected by matching valid values from paired dataset, unmatched OOB values were set to NA and excluded via complete case filter. The excluded NA observations were less than 3% of total observations.

Appendix 6.1B Further Considerations

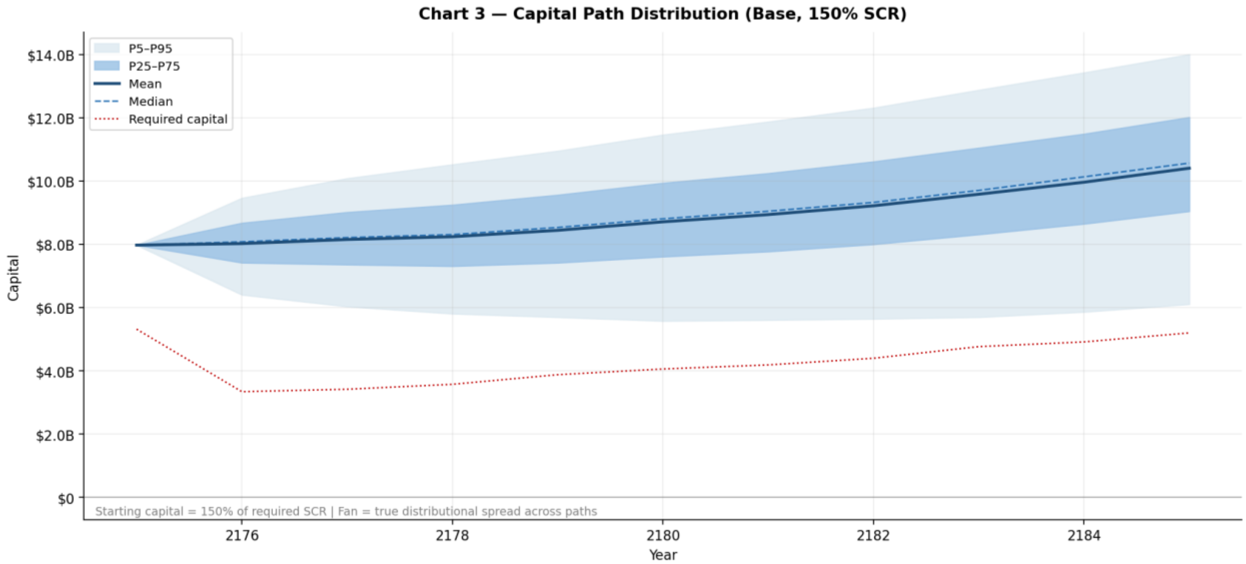
Chart 3 — Capital Path Distribution (Base, 150% SCR)
P5–P95
P25–P75
Mean
Median
Required capital
$14.0B
$12.0B
$10.0B
$8.0B
$6.0B
$4.0B
$2.0B
$0
Capital
Starting capital = 150% of required SCR | Fan = true distributional spread across paths
2176
2178
2180
2182
2184
Year

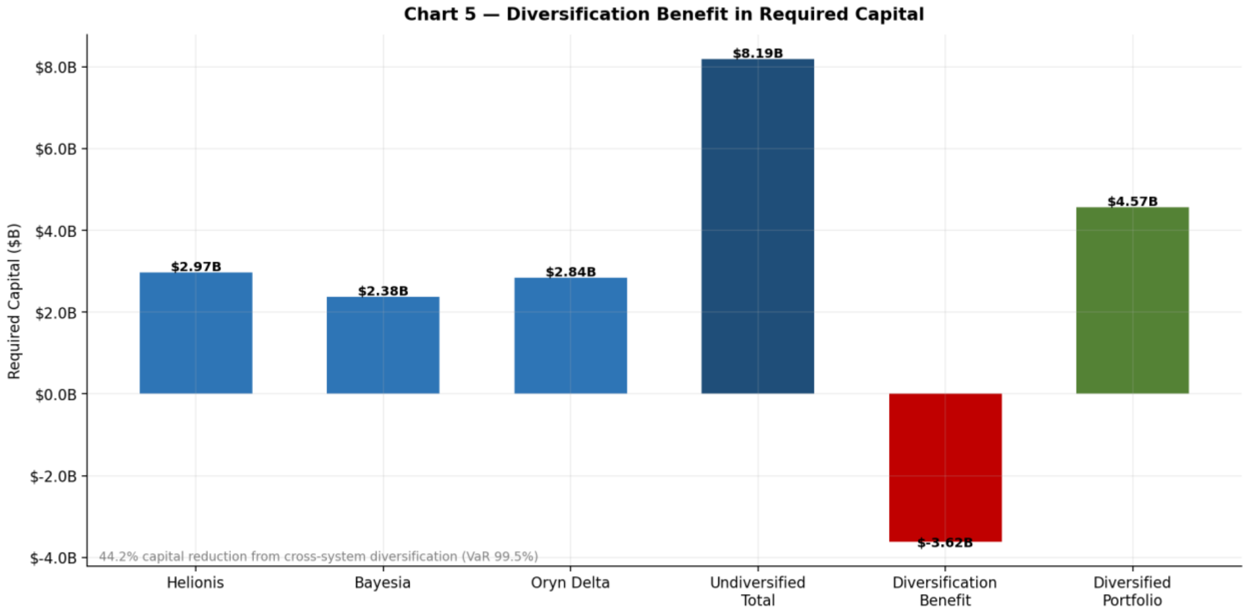
Chart 5 — Diversification Benefit in Required Capital
$8.0B
$6.0B
$4.0B
$2.0B
$0.0B
$-2.0B
$-4.0B
Required Capital ($B)
$2.97B
$2.38B
$2.84B
$8.19B
$-3.62B
$4.57B
44.2% capital reduction from cross-system diversification (VaR 99.5%)
Helionis
Bayesia
Oryn Delta
Undiversified Total
Diversification Benefit
Diversified Portfolio

Chart 2 — Probability of Recapitalisation by Year 10

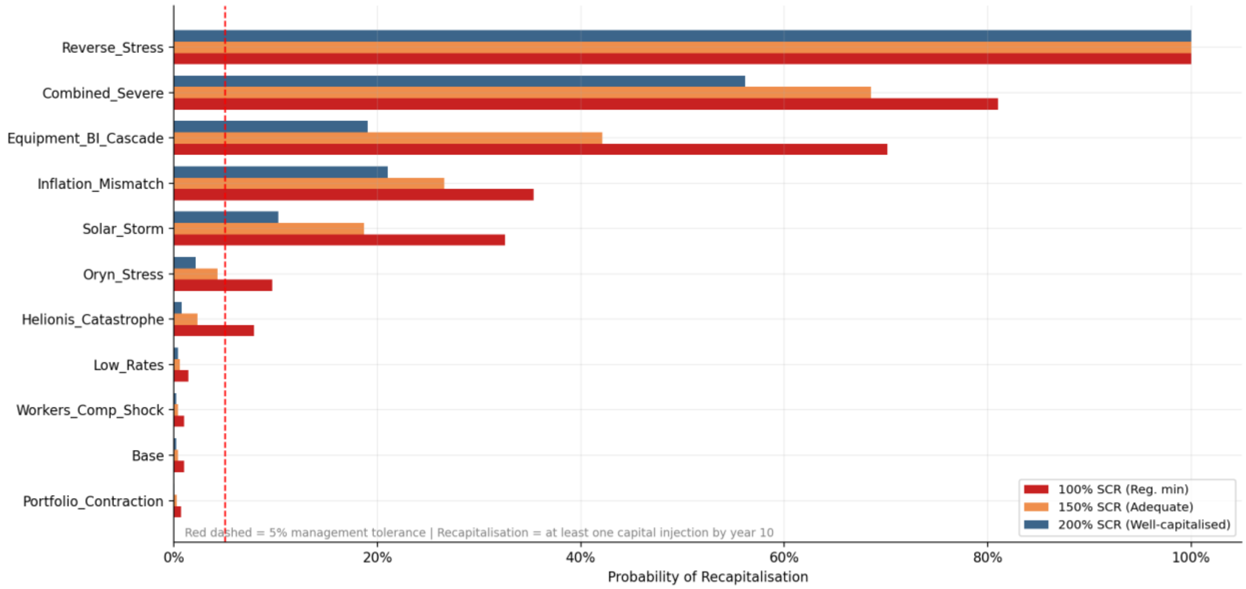

Chart 2B — Residual Probability of Ruin by Year 10

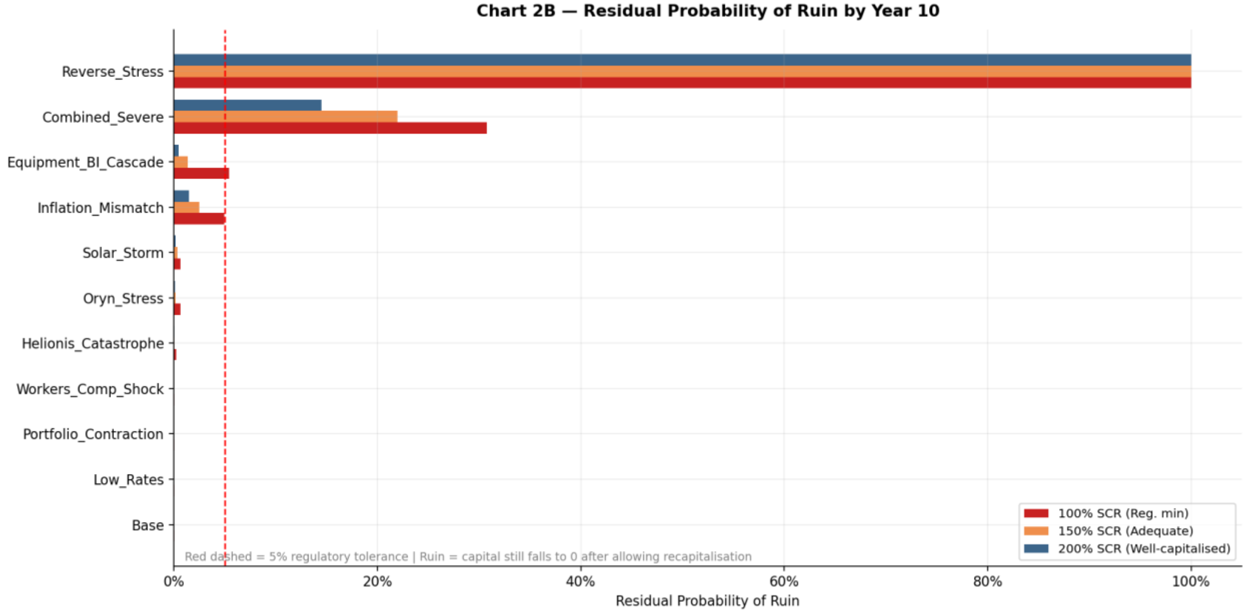

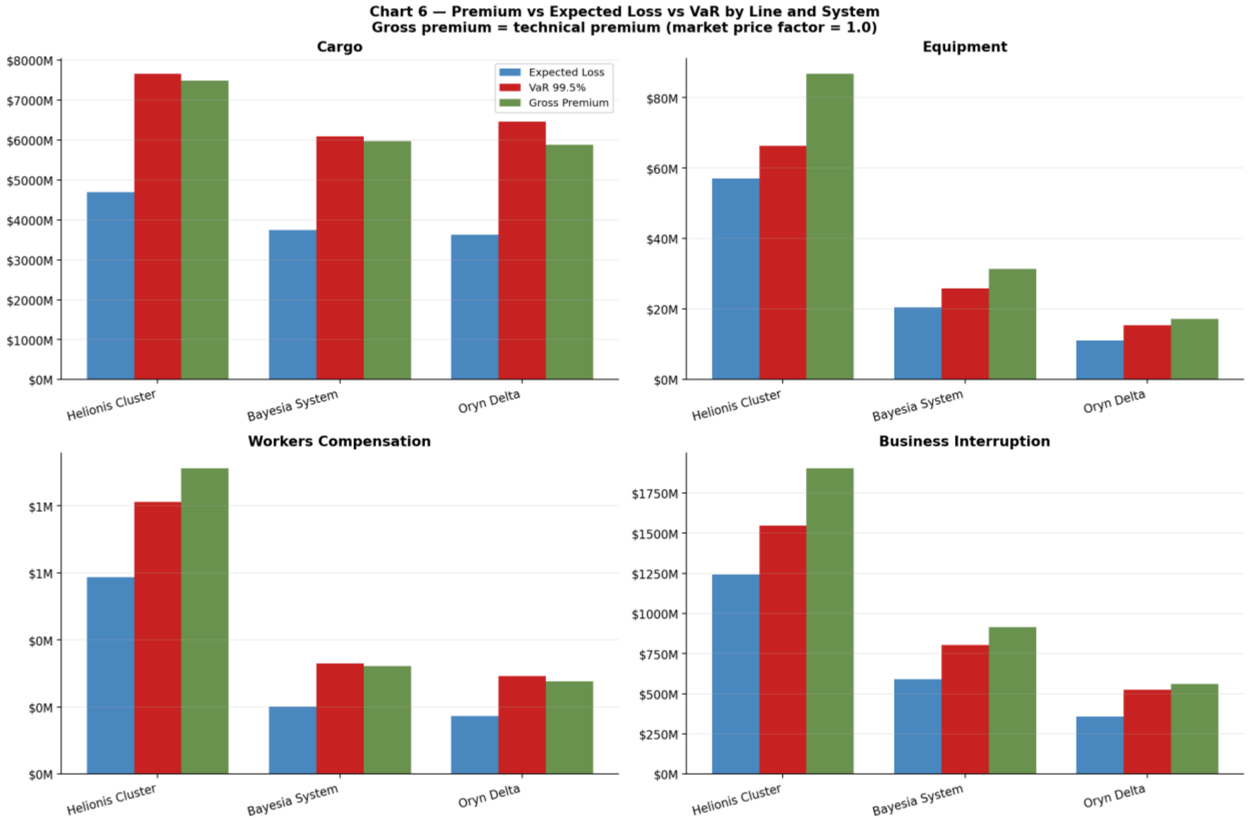
Chart 6 — Premium vs Expected Loss vs VaR by Line and System
Gross premium = technical premium (market price factor = 1.0)
Cargo
Equipment
Workers Compensation
Business Interruption
Expected Loss
VaR 99.5%
Gross Premium
Helionis Cluster
Bayesia System
Oryn Delta
$8000M
$7000M
$6000M
$5000M
$4000M
$3000M
$2000M
$1000M
$0M
$80M
$60M
$40M
$20M
$1M
$1750M
$1500M
$1250M
$1000M
$750M
$500M
$250M