DEPARTMENT OF ECONOMICS AND FINANCE

SCHOOL OF BUSINESS AND ECONOMICS

UNIVERSITY OF CANTERBURY

CHRISTCHURCH, NEW ZEALAND

Optimal Capital Allocation Between Earth and Space Insurance: A Standard Portfolio Theory Approach

Yuechen Dai
Richard Watt
Kuntal Das

WORKING PAPER

No. 2/2026

Department of Economics and Finance
UC Business School
University of Canterbury
Private Bag 4800, Christchurch
New Zealand

WORKING PAPER No. 2/2026

Optimal Capital Allocation Between Earth and Space Insurance: A Standard Portfolio Theory Approach

Yuechen Dai[1]
Richard Watt[1†]
Kuntal Das[1]

April 2026

Abstract: Over the past decade, the number of operational satellites in lower-earth orbit (LEO) has experienced exponential growth. This has led to levels of traffic congestion in the LEO environment that are making it increasingly risky. Satellites can collide with each other, and they can also be destroyed or damaged by orbital debris. While satellite insurance for such risks is not new, it does not seem to have grown at the rate one might expect given the increased number of insurable assets, and the increased risks they face. In fact, some insurers that once offered to underwrite orbital satellites now prefer to stay out of that particular market, citing excessive risk. This article uses standard financial economics modelling to explore whether this is indeed optimal. We find that in fact an expected utility maximizing insurer should always dedicate some of their underwriting capacity to orbital satellites. We also carry out a simulation to show that including orbital satellite insurance within an optimally structured insurance portfolio has the effect of enhancing long-run profitability.

Keywords: Satellite insurance; capital allocation; portfolio optimization; insurance underwriting; diversification; CRRA utility; efficient frontier; space risk.

JEL Classifications: G22, G11, D81, G32, C61

[1] Department of Economics and Finance University of Canterbury, NEW ZEALAND

† Corresponding author: Richard Watt. Email: richard.watt@canterbury.ac.nz

Optimal Capital Allocation Between Earth and Space Insurance: A Standard Portfolio Theory Approach

Yuechen Dai, Richard Watt and Kuntal Das

Abstract: Over the past decade, the number of operational satellites in lower-earth orbit (LEO) has experienced exponential growth. This has led to levels of traffic congestion in the LEO environment that are making it increasingly risky. Satellites can collide with each other, and they can also be destroyed or damaged by orbital debris. While satellite insurance for such risks is not new, it does not seem to have grown at the rate one might expect given the increased number of insurable assets, and the increased risks they face. In fact, some insurers that once offered to underwrite orbital satellites now prefer to stay out of that particular market, citing excessive risk. This article uses standard financial economics modelling to explore whether this is indeed optimal. We find that in fact an expected utility maximizing insurer should always dedicate some of their underwriting capacity to orbital satellites. We also carry out a simulation to show that including orbital satellite insurance within an optimally structured insurance portfolio has the effect of enhancing long-run profitability.

1 Introduction

Satellite insurance is underwritten by a heterogeneous set of market participants, including both large multi-line insurance groups and specialized insurers focused on aerospace lines. For large diversified insurers, satellite coverage typically represents only a small share of overall underwriting activity, which remains dominated by mass-market lines such as health, life, and property–casualty insurance. By contrast, for specialized market participants, satellite insurance can constitute a core line of business. Moreover, insurers' involvement in satellite insurance is not fixed but contingent on prevailing market conditions, with underwriting exposure adjusted over time in response to loss experience, pricing adequacy, and capital constraints. Total market capacity—defined as the aggregate amount of risk that insurers are willing to underwrite—has historically exhibited pronounced fluctuations in the space insurance market. Industry evidence suggests that capacity and premium levels tend to move in opposite directions (see Figure 1 below). In periods of abundant capacity, competition among insurers intensifies, enabling buyers to secure more favorable contract terms and lower rates. By contrast, following episodes of poor underwriting performance, insurers typically retain earnings to rebuild reserves, leading to a contraction in available capacity and an increase in premium rates (Aon, 2016). This pattern reflects the well-known sensitivity of the satellite insurance market to loss events, where a small number of large claims can materially affect market conditions.[1] As a result, satellite insurance remains a relatively small and volatile segment within global insurance portfolios (Manikowski & Weiss, 2013).

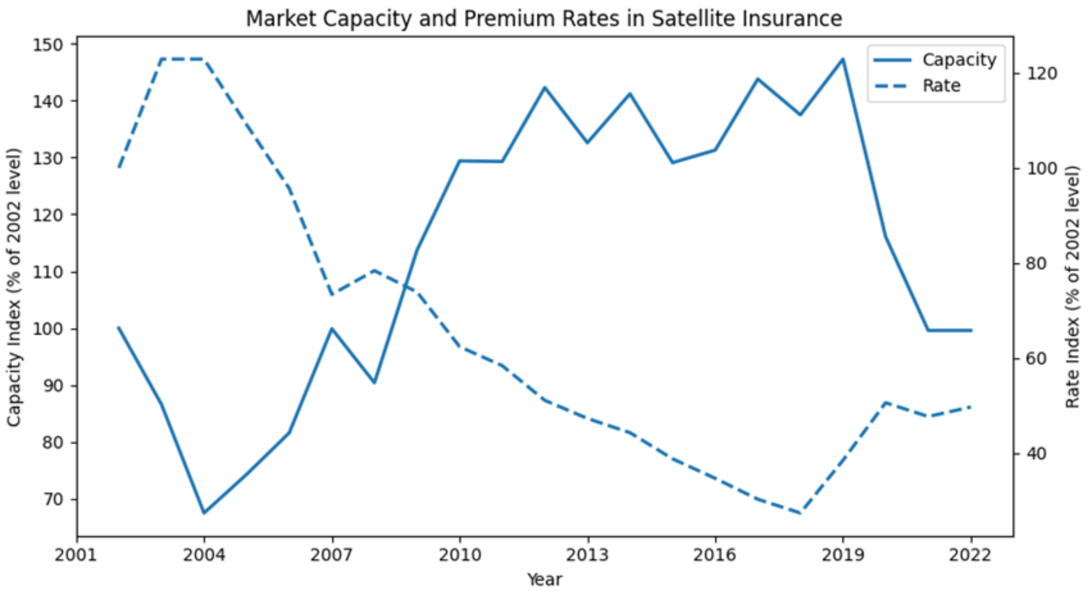

Figure 1 Market capacity and premium rates in satellite insurance constructed from proprietary AXA data and normalized to the 2002 level

Complementary qualitative evidence from industry sources further reinforces what we see in Figure 1. According to industry commentary by Chris Kunstadter, President of Triton Space LLC and former Global Head of Space at AXA XL, as competition intensified between 2012 and 2018, market premiums declined. Subsequent losses in 2018 and 2019 triggered a market hardening phase, during which several insurers exited the space insurance segment due to its inherent volatility. However, despite this hardening episode, aggregate annual premiums from 2018 through 2024 failed to recover to levels widely regarded as adequate. This reflected a combination of factors, including some clients' decisions to forego insurance, reductions in sums insured as satellites and launch services became less expensive, and the market's inability to sustain rate increases following 2019 (Arthur J. Gallagher, 2025).

[1] Manikowski and Weiss (2012, 2013) analyze satellite insurance underwriting, revealing substantial variation over 1968–2010 in both rates and market capacity, with these variables displaying recurrent cyclical patterns rather than stable long-run relationships.

Exceptionally large losses occurred in 2023. Two claims in particular—relating to the Viasat-3 and Inmarsat-6 F2 satellites, both of which reportedly experienced in-orbit failures, accounted for a substantial share of total market losses. The scale of these losses, estimated to be on the order of USD 800 million, had not been observed for more than two decades. As in earlier cycles, these unexpected claims prompted a sharp increase in premium rates, widely interpreted as signaling a market-wide revaluation of space insurance risk (Clyde & Co., 2024). The financial impact of the 2023 loss experience further accelerated an ongoing withdrawal of insurers from the sector (Arthur J. Gallagher, 2025). In late 2023, BRIT exited the direct space insurance market, adding to a sequence of departures that had already included Swiss Re, which withdrew from space underwriting in 2019. Other insurers and reinsurers—including AGCS, AIG, Allianz, Aspen Re, and Canopius—have similarly exited the market over recent time period (Brown, 2025).

Underlying the wave of insurer exits is a fundamental challenge inherent to satellite insurance. Unlike terrestrial lines of business—such as property or motor insurance—where claims can typically be investigated and fault can often be established, losses in space frequently occur in environments where attribution is highly uncertain or impossible. The inability to verify causes of failure or assign responsibility constrains insurers' capacity to manage moral hazard and assess risk, reinforcing the tendency for capital withdrawal following large loss events (Brown, 2025). However, recent developments suggest that this contraction may not be permanent. While several insurers have withdrawn from the space insurance market in recent years, there are indications that a number of these insurers are now returning, and that new market participants are also emerging, such as Helvetia. Nevertheless, profitability has remained uneven across market participants (AXA XL, 2025).

At the same time, the developments on the demand side have reinforced the recent contraction in underwriting capacity. A key driver is the divergence between rapidly declining satellite launch and manufacturing costs and persistently high insurance premiums. In spite of the increased orbital congestion and space debris risks, many satellite operators have opted to forego in-orbit insurance altogether. Instead, operators have increasingly adopted cost-reduction strategies, such as lowering satellite manufacturing costs and accelerating launch rates with more satellites (Brown, 2025). This shift away from insurance coverage is reflected in current market participation. Recent estimates indicate that over 10,000 active satellites were in orbit, yet only about 300 were insured against in-orbit losses, with coverage concentrated almost entirely among high-value geostationary satellites. In LEO, there are more than 9,000 satellites of which only around 50 are insured (Swinhoe, 2024). Regulatory frameworks contribute to this outcome. In the United States, for example, launch operators are required by law to obtain liability insurance for launch activities, but no equivalent requirement applies once a satellite is operational in orbit. Consequently, large constellation operators increasingly rely on self-insurance, SpaceX, for instance, seeks little to no third-party in-orbit insurance for its Starlink satellites (Brown, 2025). However, from a longer-term perspective, developments on the demand side do not necessarily point to a persistent decline in space insurance demand. Looking ahead, increased government procurement of commercial space products and services is expected to expand both civil and national security–related activities within the space economy, thereby creating additional sources of demand for insurance coverage. At the same time, the emergence of in-orbit servicing capabilities—such as satellite inspection, repair, and life-extension services—introduces new risk-management opportunities. These developments may not only generate new forms of insurable activity but also help mitigate operational risks, potentially reducing the frequency and severity of insurance claims and benefiting both insurers and satellite operators (Gallagher, 2025). Taken together, this suggests substantial long-term potential for expansion in satellite insurance, provided that pricing structures and risk management mechanisms can adapt to the evolving risk environment.

For multi-line insurers, the supply of satellite insurance coverage is determined within a broader portfolio allocation problem, in which capital is distributed across business lines with heterogeneous risk and return profiles. On the one hand, while the in-orbit satellite insurance business is not strongly correlated with other insurance risks, the ability to underwrite satellites will inevitably be affected by the rest of the insurer's business. In theory, the objective of insurance companies is to make profits. With relatively few insurers in the underwriting space, the loss of even a few satellite insurers can significantly reduce the underwriting capacity of the overall market, thereby exerting upward pressure on premium rates. On the other hand, some believe that a loss in one area of satellite insurance directly affects the ability of the insurance company to cover other risks (Manikowski, 2004), since the indemnities of satellite insurance policies are often large (Manikowski & Weiss, 2013). In fact, the satellite insurance market is very sensitive. Any single accident can distort it, triggering wild swings in premiums.

Conventional insurance lines benefit from a rich historical data set of frequent accidents and relatively low indemnities per claim (e.g., auto insurance and life insurance). These characteristics increase insurers’ confidence in risk estimates and in continuing to underwrite risks profitably. Specifically, as the number of independent risks in the insurance portfolio increases, the variance of aggregate claims decreases, enabling insurers to offer lower premiums. This dynamic drives up demand for insurance, increases coverage of losses, and ultimately enhances the profitability of insurers while promoting the stability of the industry. Frequent claims provide valuable feedback mechanisms that facilitate insurers' ability to adjust their premiums and coverage decisions, enabling them to continuously improve their strategies (Ermoliev et al., 2000). In contrast, in the satellite insurance market, accidents are relatively infrequent, and when accidents do occur, the claims can be very large. This poses a challenge to the application of the law of large numbers (Manikowski & Weiss, 2013) – satellite insurance becomes similar in essence to catastrophe insurance.[2] Therefore, underwriting satellite insurance becomes a challenging decision-making problem, requiring careful consideration of risk management strategies (Ermoliev et al., 2000).

From a portfolio-theoretic perspective, concentrating underwriting activity exclusively on earth-based insurance lines is not generally optimal, as it forgoes potential diversification benefits from space and limits the efficient trade-off between risk and return (space insurance typically offers higher average returns than earth-based insurance, but it is also more volatile). Within this broader context, insurers face an explicit capital allocation problem. By distributing capital across earth-based insurance lines, space-based insurance lines, and risk-free assets, insurers can trade off expected returns against risk in a manner consistent with utility maximization. This consideration motivates the present article, which examines how in-orbit satellite insurance can be optimally integrated as a regular component of an insurer's underwriting portfolio rather than treated as a specialized activity. We explicitly incorporate optimal capital allocation into insurers' decision-making under satellite insurance exposure. Specifically, we examine how insurers should optimally allocate available funds among earth-

[2] Typically, catastrophe insurance (as in, for example, cover for damage from earthquakes, tsunamis, volcanic eruptions, etc.) is challenging since one single event can give rise to many insurance claims, that is, the number and value of claims are positively correlated. Even though catastrophic events are few and far between, they lead to very large aggregate claims when they do occur. In satellite cover, while the claims are not generally correlated, they can be huge, and they are also few and far between. These characteristics set this sort of insurance apart from other instances of property and casualty losses. The unpredictable and infrequent nature of accidents makes them hard to underwrite effectively, and this line of insurance can imply a huge potential risk to insurers (Dong et al., 1996). Some insurers have concluded that the satellite insurance industry needs more events to accurately measure risk. However, due to the relatively small number of insured satellites at present, the satellite insurance industry may not have enough accidents to refer to and form a comprehensive judgment on satellite risks (Manikowski, 2004).

based insurance businesses, which offer relatively stable returns, space-based insurance businesses, which involve higher risk but potentially higher expected returns, and risk-free assets that provide guaranteed returns. This is achieved using a standard two-stage portfolio optimization framework, in which an optimal risky portfolio is first constructed and subsequently combined with a risk-free asset based on standard utility-maximizing portfolio theory.

2 The Standard Optimal Portfolio Model Applied to the Business of Insurance

The present article conceptualizes each line of insurance as an investable asset class, where the underwriting return serves as a proxy for expected return. Although this transformation is uncommon in the literature, it enables the application of mean-variance portfolio theory and CRRA-based utility optimization to the capital allocation problem. This abstraction provides a consistent and interpretable structure for analyzing solvency management and capital deployment under heterogeneous risk classes. Our model draws on modern portfolio theory to construct a two-stage capital allocation model that captures cross-sectional portfolio composition. The theoretical foundation of portfolio optimization can be traced back to the seminal work of Markowitz (1952) and Tobin (1958) which formalized the risk–return trade-off and laid the groundwork for modern portfolio theory. The upshot of this seminal literature is the result that all investors, regardless of risk preferences, would hold a combination of the risk-free asset and a single mutual fund of risky assets (the so-called “one-fund theorem”).

Merton (1975) embedded this structure within a continuous-time, utility-maximizing framework. Under CRRA preferences, Merton demonstrated that the investment decision can be decomposed into two separable components: (i) selecting the optimal risky portfolio based on market characteristics, and (ii) determining the total proportion of wealth to allocate between the risky portfolio and the risk-free asset based on investor preferences. This mutual fund separation result enables tractable modeling of dynamic portfolio decisions under uncertainty. Campbell and Viceira (2002) further developed this approach in the context of a long-term investment strategy, showing how a similar two-stage structure can guide strategic asset allocation for institutional investors facing uncertain returns and long horizons. Although our study focuses on capital allocation across insurance lines rather than financial assets, the two-stage structure is conceptually consistent with these portfolio models.

The present article builds on this perspective by embedding a two-stage, single-period capital allocation framework into the analysis of insurers' internal business decisions. In the first stage, capital is optimally allocated between earth-based and space-based insurance lines to construct a risk-efficient underwriting portfolio. In the second stage, the insurer chooses the overall risky investment level by maximizing expected utility under a CRRA framework with a risk-free asset included in the setup. While the modeling framework remains, micro-level focusing on an individual insurer's internal capital allocation, the underlying parameters are calibrated using macro-level industry data. Specifically, earth-based insurance returns are estimated from OECD insurance data across three countries (1998–2022), while space-based insurance returns rely on proprietary historical data from AXA.

3 Data Selection and Parameter Settings

In order to incorporate space-based insurance as a distinct risky asset within the portfolio optimization framework, it is necessary to calibrate its expected return and volatility. However, historical financial data on satellite in-orbit insurance is not only sparse but also highly confidential, as underwriting performance and loss experience in this niche market are typically proprietary to a few specialized insurers and not publicly disclosed. We contacted an executive in the Space Insurance division of AXA who supplied us with data from the end of the 19th century to the present (1998-2022) for in-orbit satellite insurance. Table 1 presents a list of the in-orbit insurance metrics utilized in this article along with their respective definitions.

Table 1. Variable definitions

Premiums	The total amount of premiums paid for in-orbit insurance for all insured satellites during the year
Claims	Claims paid for losses that occurred during in-orbit operations
Underwriting result	Premiums less claims
Loss ratio	Claims divided by premiums

These variables are used to construct an estimate of the underwriting return for the space-based insurance line, defined as the ratio of underwriting result to premium income, which is also equal to 1 minus loss ratio:

$$r_{s,i} = \frac{Premium\ Income_i - Claims_i}{Premium\ Income_i} = 1 - Loss\ ratio_i$$

This measure captures the gross underwriting margin before accounting for operating expenses, commissions, or taxes. Given that expense data for satellite insurance is not available, this simplified definition provides a consistent and transparent proxy for technical profitability.

To construct the earth-based insurance asset in the portfolio model, we use historical insurance data obtained from the OECD Statistics database. Specifically, we extracted annual series of "Premiums" and "Claims" for three OECD countries (France, Germany, and the United Kingdom) from 1998 to 2022. These data include gross written premiums earned and gross claims paid by insurers and are publicly available through the OECD platform. This line of business serves as a representative proxy for a traditional earth-based insurance asset with relatively stable risk-return characteristics. For each country-year observation, we compute the loss ratio as the ratio of claims to premiums and derive the underwriting return using the same formulation as in the space insurance line. The resulting time series of underwriting returns is then used to compute the sample mean and standard deviation, which are interpreted as the expected return and volatility of the earth-based insurance asset.

Table 2 presents the calibrated values used as the initial values for simulation. These are calculated from the raw data that is provided in Appendix 1 and 2. The expected return of earth-based insurance is 25.71%, with a standard deviation of 17.72%, reflecting moderate profitability and volatility. In contrast, the space-based insurance line exhibits a significantly higher expected return of 49.43%, but also a much larger standard deviation of 68.02%, consistent with the rare but high-impact nature of satellite insurance claims.

Table 2. Expected return and standard deviation of insurance lines

Insurance line	Expected return (μ_i)	Standard deviation (σ_i)
Earth insurance	0.2571	0.1772
Space insurance	0.4943	0.6802

As for the risk-free rate, whenever we bring it into consideration it is set to 5%, which is very close to the one-year U.S. Treasury bond rate.

4. Two-Stage Portfolio Allocation Model

Assume an insurer who has a set amount of capital, K, which it can use to develop and support the insurance business lines in which it participates, and to retain a risk-free reserve fund against which any business losses can be financed. The insurer must decide which fraction of K to invest in each of their business lines (essentially, the budget that is set aside for running each line of insurance underwriting), and how much to hold back in reserves invested in a risk-free asset. Each line of insurance business is treated analogously to a financial asset, where the underwriting profit margin, which is approximated by $1 - Loss\ ratio$, is used as the asset's return. Although this is not a market-traded financial return in the traditional sense, it represents the core profitability of the insurance line and captures its contribution to the insurer's capital over time. This abstraction enables the application of portfolio theory tools, such as mean-variance optimization and CRRA-based allocation, to evaluate diversification strategies across insurance lines.

In the first stage, the insurer constructs an optimal risky portfolio by allocating capital between two risky assets, namely earth-based insurance and space-based insurance. In the second stage, the insurer decides the overall investment allocation between the optimal risky portfolio and holding reserves in a risk-free asset, based on an expected utility maximization framework. The underlying idea of optimizing a risk-return trade-off is conceptually aligned with Markowitz (1952)'s original formulation.

Assume then that the insurer has total investible capital equal to K dollars with which to operate business. There are three options for allocating capital, namely (1) underwriting risks on earth, (2) underwriting risks in space, and (3) investing in a risk-free asset (i.e. holding a reserve fund). We refer to the first two options (the regular business of underwriting risks) as the "risky" part of the portfolio, and the third as the "non-risky" or "safe" part. The problem of the insurer is to select the fractions of total capital to dedicate to each part of the portfolio in order to maximise their expected utility.

In the first stage, the insurer focuses on constructing an optimal risky portfolio composed solely of investments in underwriting insurance in earth-based risks and in space-based risks. Denote by α the fraction of K that the insurer dedicates to these risky business lines. For risks on earth, the expected return per dollar invested is $E(r_e) \equiv \mu_e$, with standard deviation σ_e. The proportion invested in the space insurance business gives an expected return per-dollar invested of $E(r_s) \equiv \mu_s$, with standard deviation σ_s. Critically, the correlation between these two business lines is zero (i.e. co-variance is 0) – accidents in space cannot affect accidents on earth, and vice versa. The two risky business lines give us two points in mean-standard deviation space (one for the on-earth line and another for the in-space line), and the insurer shares $\alpha K \equiv K_R$ between them. The fraction of K_R that the insurance company dedicates to the space business line is denoted by α_s, and in the same way, the fraction dedicated to underwriting risks on earth is α_e. The overall expected rate of return of an insurance company's risky business portfolio is then:

$$E(r_R) = \alpha_e \mu_e + \alpha_s \mu_s \equiv \mu_R$$

Since there is zero co-variance between risks on earth and risks in space, the variance of this risky business portfolio is:

$$\sigma_R^2 = \alpha_e^2 \sigma_e^2 + \alpha_s^2 \sigma_s^2$$

Next, introduce a risk-free asset with return r_f. The insurer invests a fraction of total capital in the risk-free asset equal to $1 - \alpha$. Now the insurer optimizes a three-way capital allocation among earth risks, space risks, and the risk-free asset by allocating its total investable capital between three options:

1. Earth-based risks: Expected return μ_e, standard deviation σ_e.
2. Space-based risks: Expected return μ_s, standard deviation σ_s.
3. Risk-free asset: Fixed return (r_f), zero standard deviation.

This is a standard model in finance, with a standard solution in mean-standard deviation space. First, one constructs the investment opportunities frontier between the two risky options (here, underwriting risks on earth and in space). Second, one locates the capital markets line (CML), which is the straight line starting from the risk-free point on the vertical axis that is tangent to the risky investment opportunities line. This essentially involves maximizing the Sharpe ratio on the risky portfolio frontier. The point of tangency indicates the optimal shares of the two risky ventures within the risky part of the portfolio. Finally, the optimal shares of the total reserve fund between risk-free and risky elements are found by maximizing the expected utility of the insurer on the CML.

4.1 Analysis of an Indifference Curve in Mean–Standard Deviation Space

The optimal shares of the two risky assets in the risky part of the overall portfolio are well known to be (Tobin, 1958)[3]

$$\alpha_s^* = \frac{\varphi_s/\sigma_s^2}{\varphi_e/\sigma_e^2 + \varphi_s/\sigma_s^2}$$

where

$$\varphi_i \equiv \mu_i - r_f \quad for\ i = e, s.$$

Accordingly, the optimal weight for earth insurance in the risky portfolio is:

$$\alpha_e^* = 1 - \alpha_s^*.$$

Given the optimal structure of the risky component of the insurer's portfolio as characterized by Tobin's separation result, the remaining problem is to determine the optimal overall exposure to risk, that is, the fraction of total capital allocated to the risky portfolio as opposed to the risk-free asset. To address this second-stage decision, the insurer's preferences are characterized using a certainty-equivalent representation derived from a second-order Taylor approximation to expected utility.

The second order Taylor's approximation to the certainty equivalent wealth of a lottery with mean μ and standard deviation σ (Pratt, 1978) is

$$C = \mu - \frac{\sigma^2}{2} R^a(\mu)$$

Take this as preference function, since utility is a monotone increasing function of certainty equivalent. Under CRRA preferences with risk-aversion parameter ρ, the coefficient of absolute risk aversion is given by $R^a(\mu) = \frac{\rho}{\mu}$, so the expected utility function becomes

$$U(\mu, \sigma, r) = \mu - \frac{\sigma^2 \rho}{2\mu}$$

To find the equation of an indifference curve, this expression for μ needs to be solved, holding expected utility at $\bar{U}$. The solution is $\mu = \frac{1}{2}\left(\bar{U} + \sqrt{2\rho\sigma^2 + \bar{U}^2}\right)$ (Levy & Markowitz, 1979). The derivative of this with respect to σ (i.e. slope of indifference curve) is

[3] Details of all of the mathematical derivations of the model are given in Appendix 3.

$$\left.\frac{\partial\mu}{\partial\sigma}\right|_{dU=0} = \frac{\sigma\rho}{\sqrt{2\rho\sigma^2 + \bar{U}^2}}$$

Notice that

$$\lim_{\sigma\to\infty}\left(\frac{\sigma\rho}{\sqrt{2\rho\sigma^2 + \bar{U}^2}}\right) = \frac{\rho}{\sqrt{2\rho}} = \sqrt{\frac{\rho}{2}}$$

Therefore, for a finite tangency between an indifference curve and the CML to exist, the slope of the CML must be less than $\sqrt{\frac{\rho}{2}}$. Otherwise, no finite solution exists.

4.2 Analysis of the Risky Options

Take any two assets, defined by (μ_1, σ_1) and (μ_2, σ_2), with $\mu_2 > \mu_1$ and $\sigma_2 > \sigma_1$. The risky investment frontier between these two assets is given by the following piecewise expression,

$$\mu(\sigma) = \begin{cases} \mu_2 - \left(\frac{\sigma_2^2 - \sqrt{\sigma^2(\sigma_1^2 + \sigma_2^2) - \sigma_1^2\sigma_2^2}}{\sigma_1^2 + \sigma_2^2}\right) \times (\mu_2 - \mu_1), & if\ \mu(\sigma) > \mu_2 - \left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}\right) \times (\mu_2 - \mu_1) \\ \mu_2 - \left(\frac{\sigma_2^2 + \sqrt{\sigma^2(\sigma_1^2 + \sigma_2^2) - \sigma_1^2\sigma_2^2}}{\sigma_1^2 + \sigma_2^2}\right) \times (\mu_2 - \mu_1), & if\ \mu(\sigma) < \mu_2 - \left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}\right) \times (\mu_2 - \mu_1) \end{cases}$$

Notice that in order for there to be a real solution, we require $\sigma^2(\sigma_1^2 + \sigma_2^2) > \sigma_1^2\sigma_2^2$, therefore, the variance in the risky portfolio must satisfy

$$\sigma^2 > \frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2}$$

that is, the risky portfolio standard deviation must satisfy

$$\sigma > \sqrt{\frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2}}$$

So, the smallest possible standard deviation that can be obtained using these two assets is $\sqrt{\frac{\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_2^2}}$. This also determines the split in the piecewise function above. At exactly this minimum standard deviation, the upper section and the lower section are both equal to

$$\mu_2 - \left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}\right) \times (\mu_2 - \mu_1)$$

Using asset 1 to be insurance on earth (now indexed by e), and asset 2 to be insurance of orbiting satellites in space (now indexed by s), the raw data are the following

$$\mu_e = 0.2571,\ \sigma_e = 0.1772,\ \mu_s = 0.4943,\ \sigma_s = 0.6802.$$

To account for operational costs (i.e. costs that are not actuarial), for which we assume that a loading factor of around 0.3 has been included in the premium,[4] these numbers become (rounded to decimal places):

$$\mu_e = 0.18\ \ \sigma_e = 0.12,\ \mu_s = 0.35,\ \sigma_s = 0.48.$$

Using these numbers, the efficient frontier is illustrated in the figure below, where the two circles indicate the positions of the two assets.

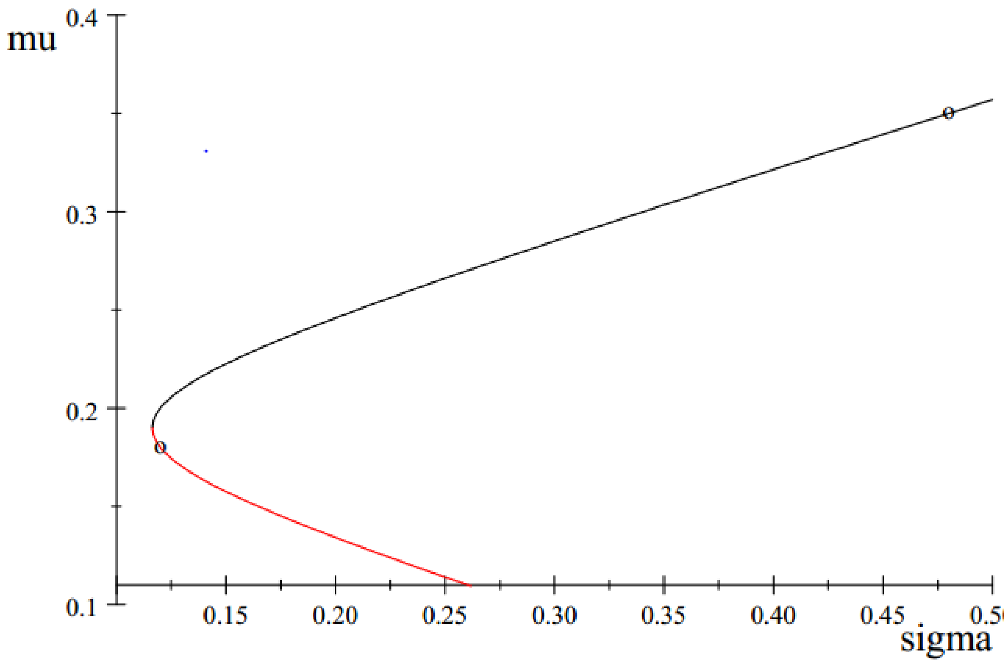

Figure 2. Efficient frontier constructed from earth-based and space-based insurance assets

Indifference curves can now be added to the figure using the expressions derived above. An optimal portfolio is identified at the point where an indifference curve is tangent to the investment frontier. Since all indifference curves have positive slope, any optimal allocation must lie on the upper section of the investment frontier. In the figure below, the blue indifference

[4] This is simply an estimate, but a quick AI search online indicates that for insurance companies, the "Operational Costs (Expense Ratio)" which covers the costs of running the business, but does not include actuarial costs of claims, for P&C insurance, usually falls between 20% and 40% of the total collected premiums.

curve corresponds to relative risk aversion of 2, while the brown curve corresponds to relative risk aversion of 1. A higher degree of risk aversion is associated with a larger fraction of investable wealth allocated to the less risky asset. In both cases, the majority of the investment is concentrated on the safer asset, although some exposure to the riskier asset remains.

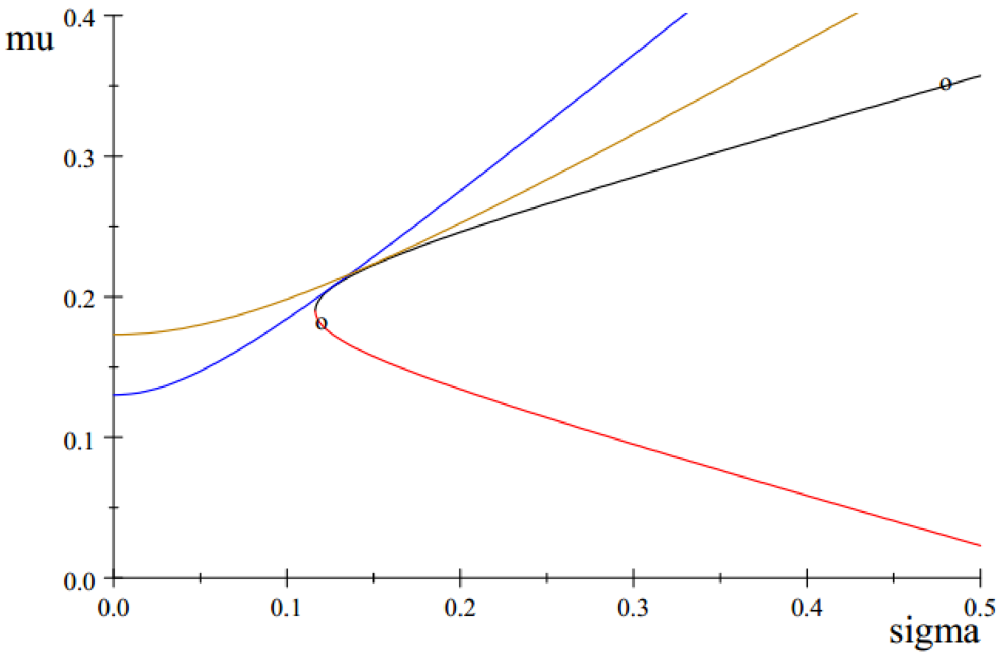

Figure 3. Indifference curves and optimal portfolio choice on the efficient frontier

4.3 Addition of a Risk-Free Asset

The risk-free asset is now introduced with return $r_f = 0.05$. The portfolio can therefore be decomposed into an investment in risk-free assets and an investment in risky assets. The fraction of wealth allocated to earth-based insurance within the risky portfolio is determined by maximizing the Sharpe ratio along the investment frontier. Based on the numerical example above, the fraction allocated to earth-based insurance in the risky portfolio is given by:

$$\omega_e = \frac{\left(\frac{0.18-0.05}{0.12^2}\right)}{\left(\frac{0.18-0.05}{0.12^2}\right) + \left(\frac{0.35-0.05}{0.48^2}\right)} \approx 0.87395$$

So, at the point at which the Sharpe ratio is maximized, the optimal risky portion of an insurance portfolio should include only about 87% invested in earth risks, and so 13% in space. The risky portfolio mean is

$$\mu_R = \alpha_e\mu_e + \alpha_s\mu_s = 0.87395 \times 0.18 + (1 - 0.87395) \times 0.35 \approx 0.20143$$

and the risky portfolio standard deviation is

$$\sigma_R = \sqrt{\alpha_e^2\sigma_e^2 + \alpha_s^2\sigma_s^2} = \sqrt{0.87395^2 \times 0.12^2 + (1 - 0.87395)^2 \times 0.48^2} \approx 0.12108$$

The slope of the capital market line is $\frac{0.20143-0.05}{0.12108} \approx 1.2507$. The following graph shows the CML and the portfolio frontier.

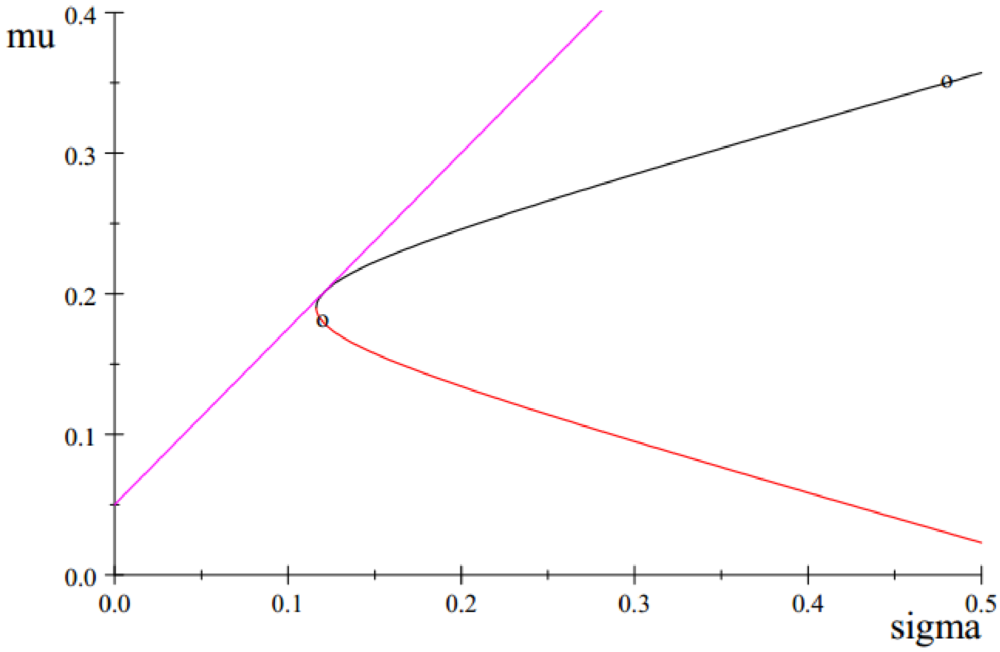

Figure 4. Capital market line and optimal risky portfolio

At the point of tangency between the portfolio frontier and the CML, expected utility is

$$U = \mu_R - \frac{{\sigma_R}^2\rho}{2\mu_R} = 0.20143 - \frac{0.12108^2 \times \rho}{2 \times 0.20143}$$

Recall, that a necessary condition for a tangency between the CML and an indifference curve is that the slope of the CML be less than $\sqrt{\frac{\rho}{2}}$. Accordingly, in the present example, this condition requires that

$$1.2507 < \sqrt{\frac{\rho}{2}}$$

Which means

$$\rho > 3.1285$$

So, there is no solution that involves the risk-free asset at 0.05 when the level of risk aversion is at or below 3.1285.[5] In contrast, a finite solution does exist for higher levels of risk aversion.

Take, for example, $\rho = 3.5$. At the point where the CML is tangent to the risky frontier, expected utility with this risk aversion level is

$$U = \mu_R - \frac{{\sigma_R}^2\rho}{2\mu_R} = 0.20143 - \frac{0.12108^2 \times 3.5}{2 \times 0.20143} \approx 0.074062$$

Notice that it is greater (i.e. better in an expected utility sense) than investing only in risk-free assets, which would yield expected utility equal to 0.05.

The slope of an indifference curve at the point where the CML is tangent to the risky frontier is

$$\left.\frac{\partial\mu}{\partial\sigma}\right|_{dU=0} = \frac{\sigma\rho}{\sqrt{2\rho\sigma^2 + \bar{U}^2}} = \frac{0.12108 \times 4}{\sqrt{2 \times 0.12108^2 \times 3.5 + 0.074062^2}} \approx 1.2889$$

Now, notice that the indifference curve is steeper than the CML at the risky portfolio (1.2889 > 1.2507), so in this case the solution involves a positive position in risk-free option.

To find the exact solution, the equation of the CML must be solved together with the tangency condition. Specifically, the CML is given by

$$\mu = 0.05 + 1.2494\sigma$$

and the tangency requirement is

$$\frac{\sigma\rho}{\sqrt{2\rho\sigma^2 + \bar{U}^2}} = 1.2494$$

where U is expected utility, $\mu - \frac{\sigma^2\rho}{2\mu}$.

Since U depends on both the mean and standard deviation, this gets rather complicated but can be done in another way. Expected utility is $U(\mu, \sigma, r) = \mu - \frac{\sigma^2\rho}{2\mu}$. Holding U constant, an expression for the slope of an indifference curve can be found by using the implicit function theorem,

[5] Essentially, with risk aversion numbers below 3.1285, the investor (here the insurer) would like to borrow infinite money at the risk-free rate and use it to expand their insurance business lines.

$$\frac{\partial\mu}{\partial\sigma} = -\frac{\left(-\frac{\sigma\rho}{\mu}\right)}{\left(1+\frac{\sigma^2\rho}{2\mu^2}\right)} = \frac{\frac{\sigma\rho}{\mu}}{\frac{2\mu^2+\sigma^2\rho}{2\mu^2}} = \left(\frac{2\mu^2}{\mu}\right)\left(\frac{\sigma\rho}{2\mu^2+\sigma^2\rho}\right) = \frac{2\mu\sigma\rho}{2\mu^2+\sigma^2\rho}$$

Just to check, using $\sigma_R = 0.12108$ and $\mu_R = 0.20143$, this formula with $\rho = 3.5$ gives the slope to be

$$\frac{2\times 0.20143\times 0.12108\times 3.5}{2\times 0.20143^2 + 0.12108^2\times 3.5} \approx 1.2889$$

so exactly as before. So, the two equations which need to be used to find the tangency point are

$$\mu = 0.05 + 1.2507\sigma$$

and

$$\frac{2\mu\sigma\rho}{2\mu^2+\sigma^2\rho} = 1.2507$$

Using $\rho = 3.5$, the following system must be solved,

$$\mu = 0.05 + 1.2507\sigma$$

$$\frac{7\mu\sigma}{2\mu^2+4\sigma^2} = 1.2507$$

The solution is $\sigma \approx 0.08273$ $\mu \approx 0.15347$.

The value of expected utility at the point is

$$U = \mu - \frac{\sigma^2\rho}{2\mu} = 0.15347 - \frac{0.08273^2\times 3.5}{2\times 0.15347} \approx 0.07526$$

This is now better than using only the risk-free asset, and also better than what the investor would get using only the point at which the Sharpe ratio is tangent to the investment frontier (0.07526 > 0.074062).

The optimal share of total wealth allocated to the risky portfolio can be obtained by solving the equation for the optimal portfolio mean. Let the share of total wealth dedicated to the risky portfolio be denoted by α, so that the share allocated to the risk-free option is $1-\alpha$. The risky portfolio mean is 0.20143, while the risk-free mean is 0.05. Finally, the optimal total portfolio mean is 0.15347. Therefore,

$$0.15347 = 0.05\times(1-\alpha) + 0.20143\times\alpha$$

It is simple to work out that the solution is $\alpha \approx 0.68329$. Therefore, a little over two-thirds of total wealth is dedicated to risky insurance lines, and a little less than one-third is dedicated to risk-free assets. Of the 68.3% of total investable wealth that is dedicated to insurance business lines, it was calculated above that 87.4 % should be spent on insurance on earth and the other 12.6% should be used to underwrite in-orbit satellites. Thus, of the total investable wealth, 31.7% would be held as cash reserves (the risk-free option), 59.7% (87.4% of 68.3%) should be dedicated to underwriting risks on earth, and 8.6% (12.6% of 68.3%) should be dedicated to underwriting risks involving in-orbit satellites.

The solution obtained naturally depends critically on the level of risk aversion assumed for the insurer. In the example above, a risk aversion parameter of $\rho = 3.5$ is assumed, so even if the statistical data for the insurance business lines were to remain the same, insurers with different levels of risk aversion would optimally allocate their investable wealth in different ways. Perhaps the most important implication of that is that as risk aversion rises, the optimal solution will move downward along the CML, implying a greater involvement of the risk-free assets, and a smaller involvement in insurance. And, as risk aversion goes down, the opposite would be the case. It is worthwhile exploring that second case in a little more detail.

4.4 Lower Levels of Risk Aversion

Above, an example is calculated in which relative risk aversion is assumed to be 3.5. In that solution, the insurer prefers to hold some of their investable wealth in the risk-free asset. Clearly, if risk aversion were to decline from 3.5, the resulting solutions would involve a greater amount of total wealth being dedicated to the risky part of the portfolio (insurance business lines) and therefore less to risk-free assets. But recall that there is a lower limit to risk aversion that admits a finite solution. With the current statistical characteristics, it is calculated above that the lower limit of risk aversion for a finite solution is $\rho > 3.1285$. So, if risk-aversion is actually less than 3.1285, the insurer would like to borrow infinitely at the risk-free rate and then invest infinitely into insurance. Such an outcome is absurd, so further assumptions must be imposed on the model to characterize reasonable solutions for low levels of risk aversion.

To begin with, it is interesting to identify the level of risk aversion that would lead to a solution at exactly the point on the CML at which no risk-free asset is demanded. To answer this question, the level of risk aversion is determined such that an indifference curve is tangent to the CML at the coordinates $(\mu, \sigma) = (0.20143, 0.12108)$. The corresponding tangency condition is given by

$$\frac{2\mu\sigma\rho}{2\mu^2 + \sigma^2\rho} = 1.2507$$

Substitute in the relevant values for mean and standard deviation:

$$\frac{2 \times 0.20143 \times 0.12108 \times \rho}{2 \times 0.20143^2 + \ 0.12108^2 \times \rho} = 1.2507$$

A little bit of algebra reveals that the solution is $\rho \approx 3.3339$. That is, for any risk aversion level that is above 3.3339, the insurer will always include the risk-free asset in the portfolio (as was the case in the example above). If risk aversion is below 3.3339 but still above 3.1285, then the company will want to borrow a finite amount of money at the risk-free rate and invest it (along with their initial wealth stake) in the market portfolio, and if risk aversion goes below 3.1285, there is no finite solution to the issue of borrowing at the risk-free rate. In that last case, and potentially in the case in which a finite amount of borrowing of the risk-free assets would be optimal, it is convenient to consider imposing a limit on the amount which can in fact be borrowed at the risk-free rate. Exactly where that limit should be set is a matter that is beyond the scope of the present article and would most relevantly be determined by regulatory authorities. In order to show how standard financial model would work with a limit on borrowing, we simply assume that no borrowing at all is feasible, that is, the smallest amount of investable wealth that can be held as risk-free funds is 0.

4.5 Small Risk Aversion and No Option for Lending

Consider now the case in which relative risk aversion is relatively low and there is no option for lending. Specifically, consider $\rho = 1$, which corresponds to logarithmic utility.[6] All remaining parameter values are held constant. The following graph illustrates the point at which the indifference curve (in blue) is tangent to the risky portfolio frontier, as well as the tangency of the Sharpe ratio line (in magenta).

[6] The choice of log utility preferences is well-known to correspond to the optimal investment path of a decision maker wanting to maximise the long-run growth of their capital fund, which we consider appropriate for an insurer with long term financial objectives. This is referred to in the literature as the "Kelly Capital Growth Criterion" (see, for example, MacLean et al. 2011).

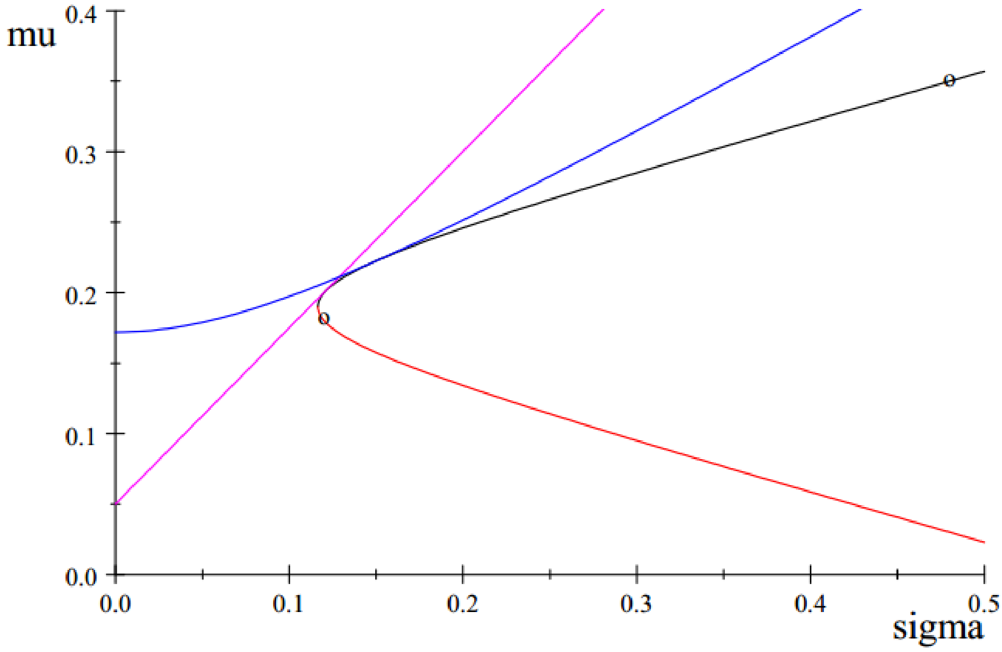

Figure 5. Optimal risky portfolio choice under low risk aversion without access to risk-free lending

Clearly, when borrowing at the risk-free rate is ruled out, the investor prefers to allocate a larger fraction of investable funds to the riskier asset than would be implied by the Sharpe-ratio–maximizing portfolio. To identify the relevant allocation, the tangency between the upper frontier of the investment portfolio curve and an indifference curve must be determined. This requires computing the slope of the upper frontier of the investment portfolio curve. The equation of that curve (using our initial data) is

$$\mu = \mu_s - \left(\frac{\sigma_s^2 - \sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}}{\sigma_e^2 + \sigma_s^2}\right) \times (\mu_s - \mu_e)$$

$$= 0.35 - \left(\frac{0.48^2 - \sqrt{\sigma^2(0.12^2 + 0.48^2)} - 0.12^2 \times 0.48^2}{0.12^2 + 0.48^2}\right) \times (0.35 - 0.18)$$

$$\approx 0.69444 \times \sqrt{0.2448\sigma^2 - 0.0033178} + 0.19$$

The slope of this is:

$$\frac{\partial \mu}{\partial \sigma} = \frac{0.17\sigma}{\sqrt{0.2448\sigma^2 - 0.0033178}}$$

On the other hand, as derived above, the derivative of utility with respect to σ (i.e., the slope of the indifference curve) is

$$\left.\frac{\partial\mu}{\partial\sigma}\right|_{dU=0} = \frac{2\mu\sigma\rho}{2\mu^2 + \sigma^2\rho}$$

So, the tangency condition is

$$\frac{0.17\sigma}{\sqrt{0.2448\sigma^2 - 0.0033178}} = \frac{2\mu\sigma\rho}{2\mu^2 + \sigma^2\rho}$$

$$0.17(2\mu^2 + \sigma^2\rho) = \sqrt{0.2448\sigma^2 - 0.0033} \times 2\mu\rho$$

Substitute in $\rho = 1$

$$0.17(2\mu^2 + \sigma^2) = 2\mu\sqrt{0.2448\sigma^2 - 0.0033178}$$

Collect the mean terms

$$0.34\mu^2 - 2\mu\sqrt{0.2448\sigma^2 - 0.0033178} + 0.17\sigma^2 = 0$$

This is a second-order equation in μ.

The other equation required to find the solution is the upper frontier of the investment curve

$$\mu = 0.69444 \times \sqrt{0.2448\sigma^2 - 0.0033178} + 0.19$$

The simultaneous solution to these two equations is $\sigma \approx 0.14937$, and $\mu \approx 0.22215$, which is in perfect consonance with the graph. Expected utility at this point is

$$U = \mu - \frac{\sigma^2\rho}{2\mu} = 0.22215 - \frac{0.14937^2}{2 \times 0.22215} \approx 0.17193$$

At the optimal position, the fraction of total investable wealth allocated to each of the two risky assets can be readily calculated. Denote by α_e the fraction of wealth that is invested in insurance on earth, and so $1 - \alpha_e$ is dedicated to insurance in space. It follows that

$$0.22215 = \alpha_e \times 0.18 + (1 - \alpha_e^*) \times 0.35$$

It is straight-forward to solve this out to get $\alpha_e \approx 0.75206$. So, a log utility insurer faced with the assumed parameter values on the two risky assets, and no ability to borrow at the risk-free rate of 0.05, will dedicate just over 75% of their investable fund to insurance on earth, and just under 25% to insurance in space. This compares to $\alpha_e \approx 0.87395$ at the point at which the CML is tangent to the risky investment frontier, that is, by maximizing expected utility as opposed to minimizing the Sharpe ratio, a log utility investor in this market would shift a little over 12% of their investable fund from insurance on earth to insurance in space.

5 Simulation of a Multi-Period Model with No Borrowing

Up to now in this article, we have restricted our attention to the case of an insurer who maximizes the expected utility of their capital allocation choice over possible insurance business lines. Using the standard Markowitzian portfolio model, in this setting the dual facts that (1) underwriting in-orbit satellite risks involves greater expected return and greater risk than does underwriting risks on earth, and (2) the two business lines are uncorrelated, any risk averse investor should always include some of the riskier asset in an expected utility maximizing portfolio. However, it is perhaps easy to see that real-world insurers are more concerned with the evolution of their retained profits over time than their expected utility. Therefore, in this section we consider exactly how profits might evolve over time when the allocation of insurer capital maximizes their expected utility. To do that, we need to employ simulation analysis.[7]

5.1 Theoretical Framework

In order to run a dynamic simulation that tracks the optimal capital allocations over time, the following are the steps required:

1. Choose a numerical value for risk aversion, ρ, and a value for the risk-free rate, r_f.
2. Calculate the numerical values of the means and standard deviations of the two insurance business lines for the given period. In period 1, these are just the values from the real-world data of 1998-2022.
3. Use the means and standard deviations from step 2 to calculate the point on the upper frontier of the investment curve that maximizes the Sharpe ratio.
4. Calculate the value of risk aversion for which exactly the point found in step 3 is optimal. Call this value δ.
5. Compare δ and ρ. If $\rho \geq \delta$, calculate the optimal capital allocation as the tangency point between an indifference curve and the CML, and calculate the three fractions for capital allocation – fraction for risk-free ($\alpha_{f,t}$), fraction for on earth insurance ($\alpha_{e,t}$), and fraction for in space insurance ($\alpha_{s,t}$). If $\rho < \delta$, then calculate the point of tangency between the indifference curve using ρ and the upper frontier of the investment curve.

[7] The Python code for the simulations that follow is available from the authors by request.

Then calculate the two fractions for optimal capital allocation – the fraction for on-earth insurance ($\alpha_{e,t}$) and fraction for in-space insurance ($\alpha_{s,t}$).

6. Take a random draw from each of the current distributions for returns on-earth and in-space insurance (using Normal distributions with the current means and standard deviations). Call the numbers that are drawn $y_{e,t}$ and $y_{s,t}$, respectively.
7. Let $K_t > 0$ denote the investable capital at the beginning of period t, and the portfolio weights

$$\alpha_{f,t}, \alpha_{e,t}, \alpha_{s,t} \geq 0, and\ \alpha_{f,t} + \alpha_{e,t} + \alpha_{s,t} = 1$$

Define the portfolio net return rate at period t as

$$r_t = \alpha_{f,t} \times r_f + \alpha_{e,t} \times y_{e,t} + \alpha_{s,t} \times y_{s,t}$$

The periodic profit is then defined as

$$\pi_t = r_t \times K_t$$

8. Let d denote the dividend payout ratio. Dividends paid at the end of period t are given by

$$D_t = \pi_t \times d$$

Capital at the beginning of period $t + 1$ evolves according to

$$K_{t+1} = K_t + (1 - d) \times \pi_t = K_t + \pi_t - D_t$$

9. The new values $y_{e,t}$ and $y_{s,t}$ are added to the total data set, and the insurer updates its beliefs by recomputing the sample mean and standard deviation for each insurance line. Let $\hat{\mu}_{j,t}$ and $\hat{\sigma}_{j,t}$ denote the insurer's estimated mean and standard deviation of returns on insurance line $j \in \{e, s\}$ after observing returns up to the end of period t, which are used for decision-making in period $t + 1$. Initial beliefs satisfy $\hat{\mu}_{j,0} = \mu_j$, $\hat{\sigma}_{j,0} = \sigma_j$.
10. Finally, the decision process returns to step 1.

5.2 Simulation Parameter Settings and Results

The simulation is conducted over $t = 50$ periods. The insurer starts with an initial capital level normalized to $K_0 = 1$. This normalization does not affect the qualitative results and allows capital dynamics to be interpreted in relative terms. A risk-free asset is available and yields a constant gross return of $r_f = 0.05$. The insurer follows a fixed dividend payout rule. In each period, a constant fraction $d = 0.2$ of operating profits is distributed as dividends to shareholders, while the remaining profits are retained and added to capital. The insurer's portfolio choice is governed by the mean–standard deviation preference representation consistent with CRRA utility. The coefficient of relative risk aversion is set to $\rho = 1$, which corresponds to log-utility. Based on industry-calibrated values, the data-generating processes for returns are assumed to follow normal distributions with parameters initially set at $\mu_e = 0.18$, $\sigma_e = 0.12$, $\mu_s = 0.35$, $\sigma_s = 0.48$. These values, which are those that correspond to the actual experience of the real-world insurance market over the past 25 years, reflect the higher expected return and substantially higher volatility of space insurance relative to traditional Earth-based insurance lines, and the assumption that 30% of premium revenue is used up as non-claims operational costs. The correlation between Earth-based and space-based insurance returns is set to 0. Those are, therefore, the initial beliefs that the insurer holds regarding the mean and volatility of each insurance line. These beliefs are interpreted as being based on a finite amount of prior information, represented by pseudo-sample sizes: $n_{e,0} = n_{s,0} = 25$ (since we collected 25 years of data).

5.2.1 Single-Path Simulation Results

First, we illustrate the outcomes from a single realization of the simulation, in which the model is run once over the assumed 50-year horizon and the time evolution of key variables is recorded along that single path. The evolution of the mean returns for both in-orbit and on-earth insurance in the data generated in the simulation, over our assumed 50-year period is shown in Figure 5, and the evolution of the volatility parameter is shown in Figure 6. The mean of returns in space insurance takes an initial upward trend and then remains relatively stable. It is persistently above the initial calibrated value, but this is only because of the initial good outcomes that the stochastic process has given in this simulation. The fact that there is greater movement in the mean of space returns than the mean of on-earth returns is of course due to the much higher volatility in space compared to on-earth.

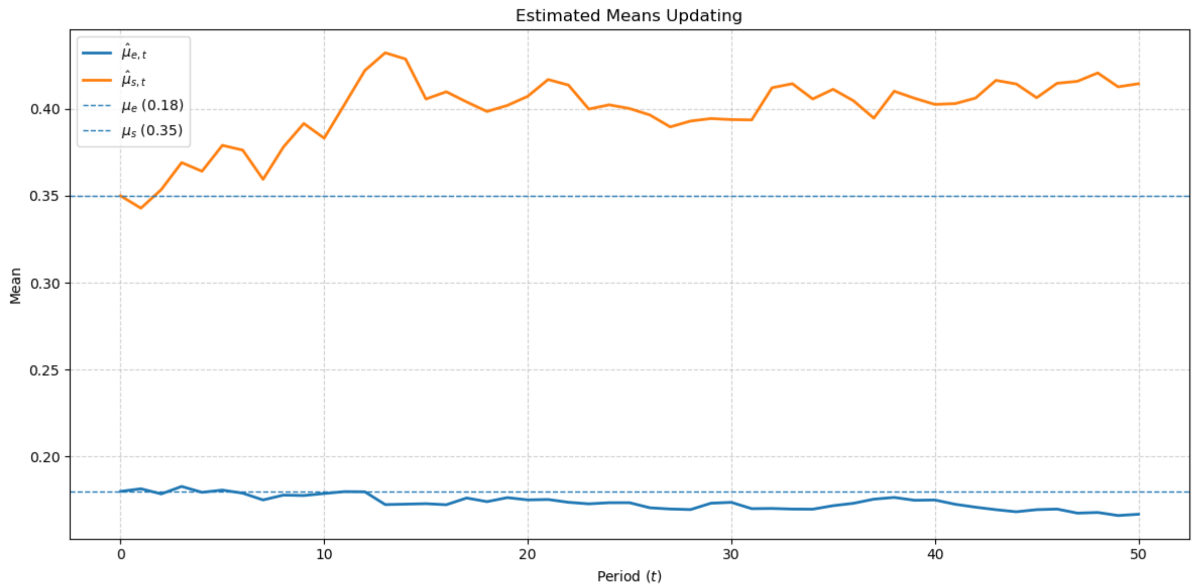

Figure 5: Evolution of total sample mean returns

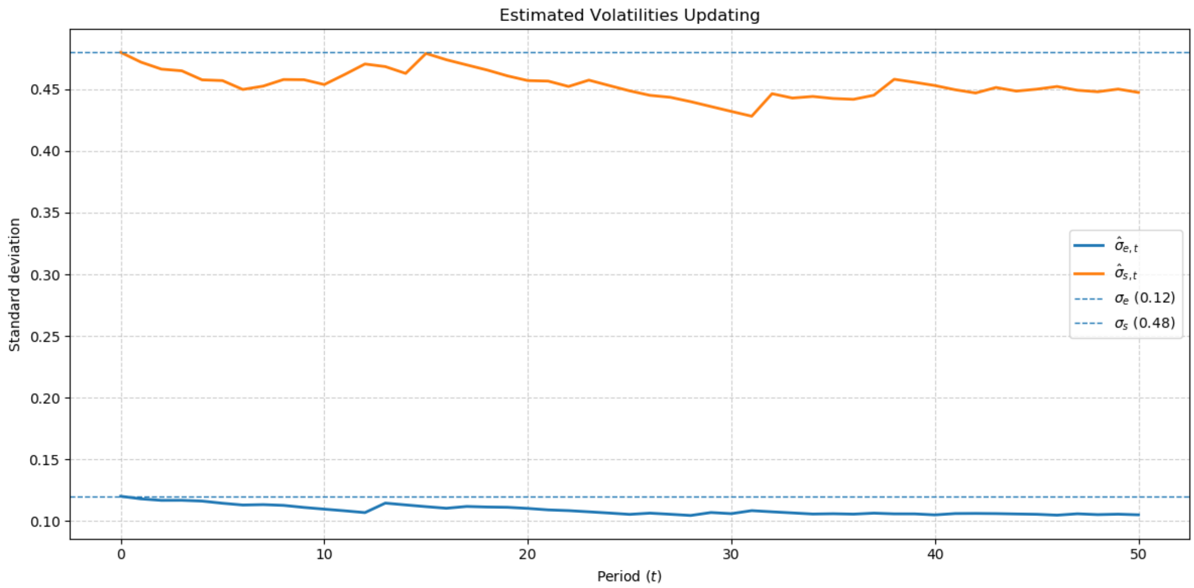

Figure 6: Evolution of total sample standard deviation of returns

Figure 7 shows how the optimal shares of the total portfolio change over time, given the evolution of the simulation sample.

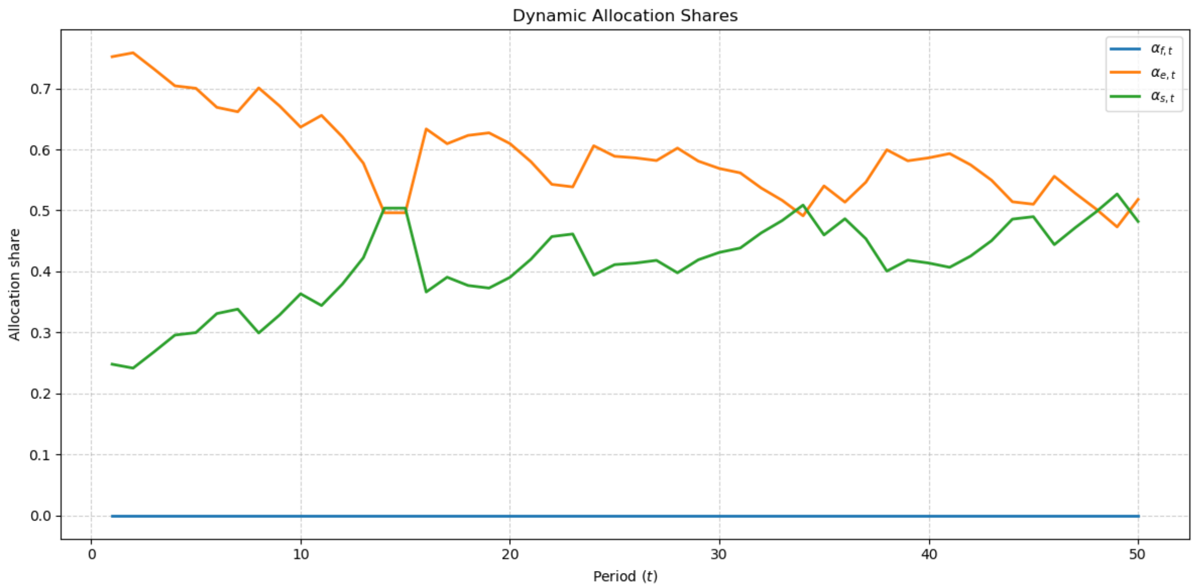

Figure 7: Optimal portfolio investment shares over time

Given the assumption of CRRA equal to 1, at no point in the simulation does the insurer ever allocate any capital to holding the risk-free asset. Of course, this would change if we were to use a higher risk aversion parameter. The portfolio weights are not constant over time. Instead, the insurer gradually shifts capital from on-earth insurance to space insurance. Specifically, investment in on-earth insurance declines from roughly three-quarters at the start to around one-half by the end of the horizon, while investment in space insurance rises correspondingly from about one-quarter to nearly one-half. Around period 14 and period 34 there are short-lived rebalancing episodes in which the two risky allocations briefly converge toward a near 1:1, after which the weights continue to fluctuate. This is, of course, a rather natural expectation given the stochastic nature of the problem. In any case, the insurer allocates a larger share of capital to earth-based insurance for most of the simulation horizon, while maintaining a persistent but limited exposure to space insurance. The absence of investment in the risk-free asset reflects that, under log utility, the optimal allocation lies strictly on the risky efficient frontier rather than on the capital market line. Importantly, dynamic learning does not eliminate space insurance from the optimal portfolio but instead leads to persistent interior allocations.

5.2.2 Monte Carlo Simulation Results

To complement the single-path illustration, we next consider results based on repeated simulations. Specifically, the model is run 100 independent times, each over the same 50-period horizon. For each simulation run, the optimal capital allocation shares are recorded period by period. The values shown in Figure 8 represent, for each period, the cross-simulation averages of the allocation shares across these 100 runs. Figure 8 therefore depicts the expected evolution of capital allocation over time, abstracting from idiosyncratic fluctuations specific to any single realization. Compared to the single-run results, the paths are substantially smoother. Under the three-asset specification, the insurer allocates a larger share of capital to earth-based insurance for most of the simulation horizon, while maintaining a persistent and non-negligible exposure to space insurance. The expected allocation to the risk-free asset remains effectively zero throughout, consistent with the model operating on the risky frontier under the chosen level of risk aversion.

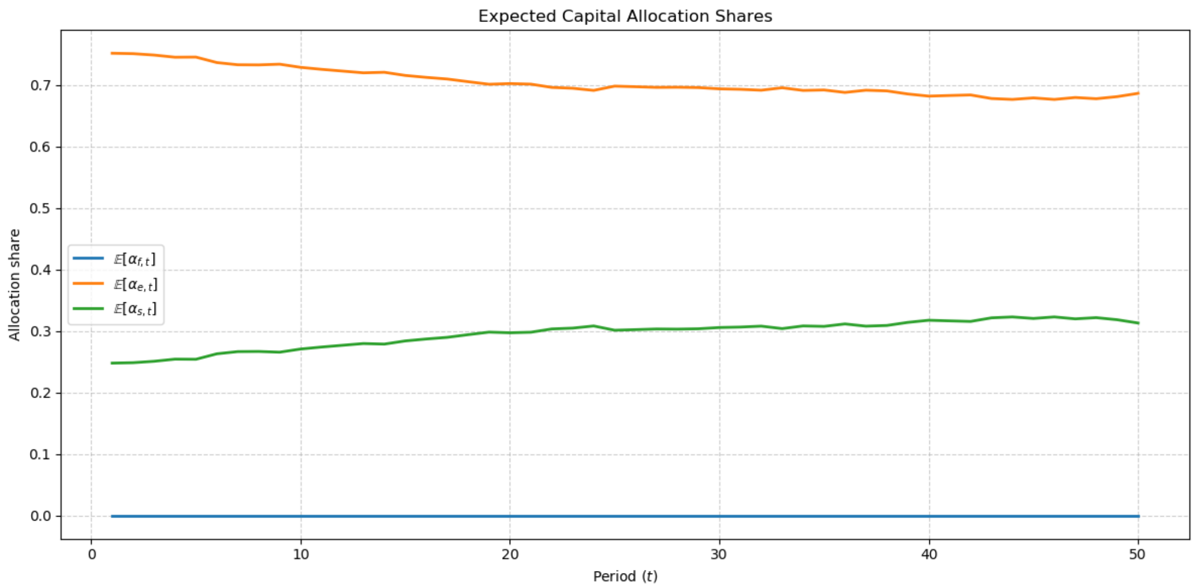

Figure 8: Optimal expected portfolio investment shares over time (three lines)

For comparison, we also conduct an analogous set of 100 simulations for an insurer that can allocate capital only between earth-based insurance and the risk-free asset. The corresponding average allocation paths are shown in Figure 9. In this case, the insurer's portfolio is overwhelmingly concentrated on earth-based insurance, with only a minimal expected allocation to the risk-free asset emerging over time. Together, these results highlight that allowing access to space insurance leads to persistent interior allocations across risky lines, whereas excluding space insurance results in a near-corner solution concentrated in on-earth insurance.

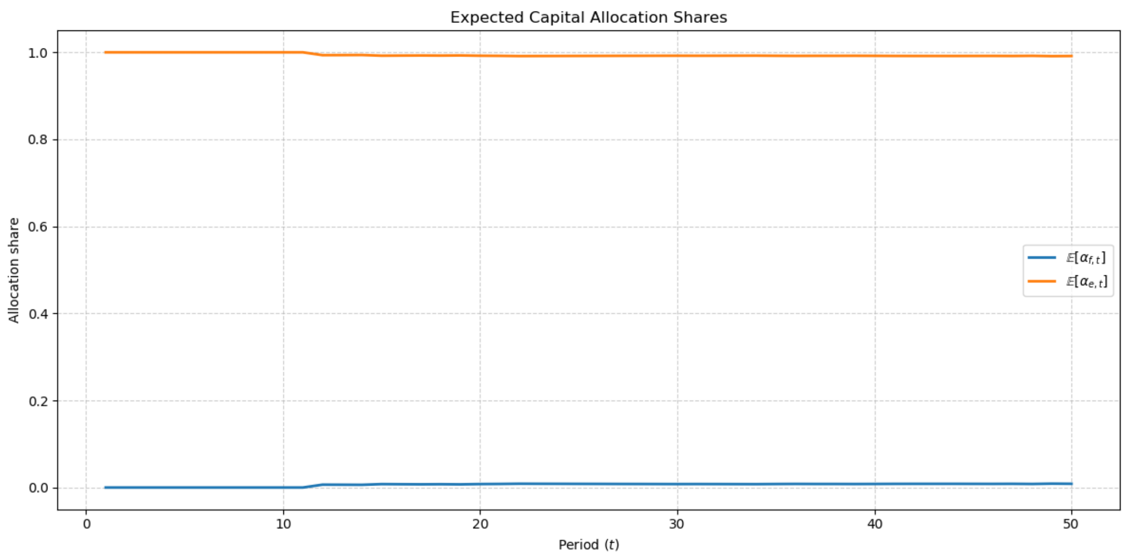

Figure 9: Optimal expected portfolio investment shares over time (two lines)

Finally, of interest is the time-series that is generated for the amount of capital stock at the end of each period. To examine the implications of portfolio composition for insurer capital dynamics, we compare capital evolution under repeated simulations for two representative types of insurance companies. The first type allocates capital between a risk-free asset and earth-based insurance only, while the second type allocates capital across a risk-free asset, earth-based insurance, and space insurance. For each type, the model is simulated 100 times, and capital paths are tracked over time. At each period, the reported capital level corresponds to the cross-simulation average of end-of-period capital across these 100 runs.

This is shown in Figure 10. Recalling that the simulation assigns at each period end a fraction of 20% of that period's return (profit) as dividends to shareholders; consequently, only the remaining 80% of profits are retained and carried forward as additions to capital in the subsequent period. Despite the 20% outflow of profits as dividends in each period, the simulation reveals an extremely healthy capital stock dynamic under both configurations. Importantly, the divergence between the two capital paths emerges only gradually. In the early and middle stages of the simulation, the expected capital levels of the two insurers remain relatively close. However, as time progresses, the insurer that includes space insurance in its portfolio experiences a markedly faster acceleration of capital growth. This widening gap reflects the compounding effects of higher expected returns from space-based underwriting once accumulated capital reaches a sufficient scale. This outcome is driven by disciplined portfolio allocation, learning-based belief updating, and partial profit retention, which together generate compounding effects over time.

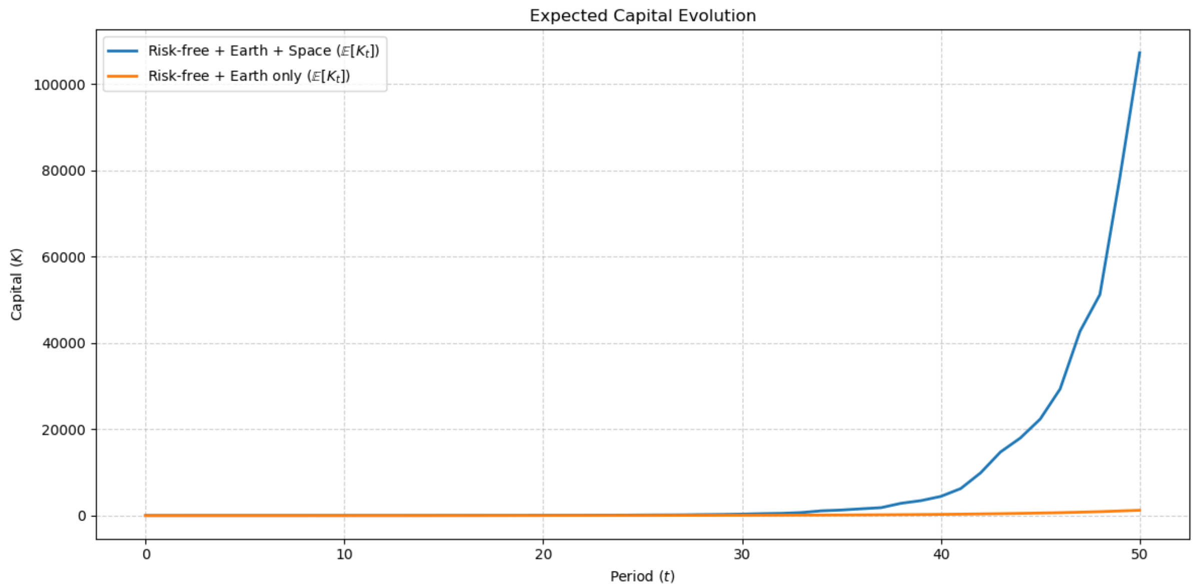

Figure 10: Evolution of expected period-end retained capital over time

6 Conclusion

This article has developed a two-stage portfolio-theoretic framework to analyze how a multi-line insurer can integrate satellite in-orbit insurance into its broader underwriting portfolio. The motivation is grounded in the empirical reality that satellite insurance, while potentially profitable, exhibits unusually pronounced volatility relative to mass-market insurance lines, and that there is constant movement in and out of the in-orbit satellite market by insurers. We find the fact that some insurers have preferred to leave this market due to the elevated risk that it implies somewhat at odds with standard portfolio theory, which predicts that whenever there exists two investable assets that lie on a positively sloped line in mean-standard deviation space, where there is zero correlation between them, any risk averse investor would create an optimal portfolio that includes participation in both assets.

To formalize this allocation problem, the article treats each line of underwriting activity as an investable pseudo-asset, with an underwriting return proxy derived from profitability. Within this setup, we are then able to employ the standard optimal portfolio theory as originally proposed by Markowitz (1952) and Tobin (1958). Assuming CRRA preferences and calibrating the model with a 25-year sample of real-world data, we show that indeed all insurers should invest a strictly positive fraction of their investable capital in underwriting in-orbit insurance. The only asset class that might disappear from an optimally constructed portfolio is the risk-free asset, assuming a limit on borrowing, and assuming sufficiently low risk aversion.

We also conducted a numerical simulation in order to see how the model would perform over time, with special interest in how the stock of investable capital changes. The simulation reveals that although participation in underwriting of in-orbit insurance is generally a minor element in an optimal portfolio, it contributes to a very healthy dynamic of capital stock, along with consistently high (and growing) dividend payouts to shareholders.

References

Aon. (2016). “Insuring space activities: A white paper.” https://www.aon.com/russia/files/Insuring_Space_Activities_whitepaper.pdf

Arthur J. Gallagher. (2025). “Space insurance market update: Q2 2025.” https://specialty.ajg.com/plane-talking/space-insurance-update-q2-2025

AXA XL. (2025). “Adapting to a new era: How the space insurance market is transforming.” https://axaxl.com/fast-fast-forward/articles/adapting-to-a-new-era_how-the-space-insurance-market-is-transforming

Brown, T. (2025). “Satellites keep breaking up in space. Insurance won’t cover them.” https://www.space.com/space-exploration/satellites/satellites-keep-breaking-up-in-space-insurance-wont-cover-them

Campbell, J. Y., & Viceira, L. M. (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economics. https://doi.org/10.1093/0198296940.001.0001

Clyde & Co. (2024). “One small leap: Clyde & Co.” *Space Bulletin* – January 2024. https://www.clydeco.com/en/insights/2024/01/space-bulletin-january-2024

Dong, W., Shah, H., & Wong, F. (1996). “A rational approach to pricing of catastrophe insurance.” *Journal of Risk and Uncertainty*, 12(2), 201-218. https://doi.org/10.1007/BF00055794

Ermoliev, Y. M., Ermolieva, T. Y., MacDonald, G. J., & Norkin, V. I. (2000). “Stochastic optimization of insurance portfolios for managing exposure to catastrophic risks.” *Annals of Operations Research*, 99, 207-225. https://doi.org/10.1023/A:1019244405392

Levy, H., & Markowitz, H. M. (1979). “Approximating expected utility by a function of mean and variance.” *World Scientific-Nobel Laureate Series*, 281.

MacLean, L.C., E.O. Thorp and W.T. Ziemba (eds.) (2011), *The Kelly Capital Growth Investment Criterion*, London: World Scientific Publishing Ltd.

Manikowski, P. (2004). "The satellite insurance market as a part of the insurance market." *Argumenta Oeconomica*, 16(2), 139-151.

Manikowski, P., & Weiss, M. A. (2012). "Cyclicality or volatility? The satellite insurance market." *Space Policy*, 28(3), 192-198. https://doi.org/10.1016/j.spacepol.2012.01.001

Manikowski, P., & Weiss, M. A. (2013). "The satellite insurance market and underwriting cycles." *The Geneva Risk and Insurance Review*, 38, 148-182. https://doi.org/10.1057/grir.2013.2

Markowitz, H. (1952). "Portfolio selection." *The Journal of Finance*, 7(1), 77-91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Merton, R. C. (1975). "Optimum consumption and portfolio rules in a continuous-time model." In *Stochastic optimization models in finance* (pp. 621-661). Academic Press. https://doi.org/10.1016/B978-0-12-780850-5.50052-6

Pratt, J. W. (1964). "Risk aversion in the small and in the large." *Econometrica*, 32, 122-136.

Swinhoe, D. (2024). "Paying the premium: Why 2023 was a bad year for space insurance." https://www.datacenterdynamics.com/en/analysis/paying-the-premium-why-2023-was-a-bad-year-for-space-insurance/

Tobin, J. (1958). "Liquidity preference as behavior towards risk." *The Review of Economic Studies*, 25(2), 65-86. https://doi.org/10.2307/2296205

Appendix 1. Annual Gross Premiums, Claims, Loss Ratios, and Returns (1 − Loss ratio) in Satellite In-Orbit Insurance (1998–2022)

Year	In orbit gross premium	In orbit claims	Loss ratios	Returns (1−loss ratio)
1998	255.379692	569.317122	2.229297	-1.229297
1999	266.736824	55.755573	0.209028	0.790972
2000	251.943160	412.251163	1.636286	-0.636286

2001	264.389009	142.944470	0.540660	0.459340
2002	211.182409	13.681113	0.064783	0.935217
2003	205.981990	211.772414	1.028111	-0.028111
2004	209.269690	0.000000	0.000000	1.000000
2005	189.027842	43.503602	0.230144	0.769856
2006	189.027842	32.772000	0.173371	0.826629
2007	207.204446	80.656053	0.389258	0.610742
2008	229.136928	88.667500	0.386963	0.613037
2009	245.680570	0.000000	0.000000	1.000000
2010	192.827370	6.528296	0.033856	0.966144
2011	181.785657	0.000000	0.000000	1.000000
2012	180.798766	32.418893	0.179309	0.820691
2013	144.380646	72.447988	0.501785	0.498215
2014	142.292887	22.317880	0.156845	0.843155
2015	142.129481	158.500000	1.115180	-0.115180
2016	129.373826	3.962588	0.030629	0.969371
2017	108.673527	9.637906	0.088687	0.911313
2018	97.235569	240.259604	2.470902	-1.470902
2019	156.102199	27.086957	0.173521	0.826479
2020	195.204622	38.149488	0.195433	0.804567
2021	180.362716	83.374483	0.462260	0.537740
2022	159.746844	55.184706	0.345451	0.654549

Notes: Premiums and claims values are reported in USD million. For presentation clarity, figures are rounded in the table.

Appendix 2. Annual Gross Premiums, Claims, Loss Ratios, and Returns (1 − Loss ratio) for Insurance Companies in France, Germany, and the United Kingdom (1998–2022)

Country	France		Germany		The United Kingdom	
Year	Gross written premiums	Gross claims payments	Gross written premiums	Gross claims payments	Gross written premiums	Gross claims payments
1998	123,354	78,456	178,521	121,552	191,180	90,374
1999	129,077	78,016	181,265	129,755	215,198	137,185
2000	128,410	83,718	164,773	117,618	256,363	155,876
2001	122,085	79,835	166,059	130,307	230,522	156,741
2002	131,873	89,127	192,476	142,179	252,689	170,063
2003	173,491	112,677	236,129	166,654	297,663	165,057
2004	211,501	132,712	291,502	179,434	336,327	39,923
2005	234,404	138,214	296,248	187,166	358,541	43,967
2006	292,646	180,148	264,402	178,929	525,858	376,780
2007	313,113	175,596	280,534	197,035	644,556	519,264
2008	279,450	205,978	303,291	220,315	454,666	437,125
2009	302,184	195,668	247,779	214,685	342,609	331,644
2010	297,324	193,166	241,052	206,871	325,141	324,674
2011	289,082	240,944	255,708	248,906	339,365	345,502
2012	258,641	230,848	302,739	216,007	368,525	349,852
2013	275,750	223,557	333,713	242,461	340,385	347,474
2014	294,033	231,967	338,764	195,733	365,688	358,024

2015	250,873	201,639	293,087	167,507	336,743	348,888
2016	314,251	252,764	294,661	173,432	403,794	286,215
2017	314,319	278,561	311,208	173,812	394,929	418,738
2018	338,708	279,259	335,669	187,607	471,031	479,672
2019	332,472	271,367	344,268	241,641	418,559	399,096
2020	302,729	274,227	365,897	251,310	380,954	358,215
2021	362,503	310,473	396,718	279,786	445,432	429,176
2022	349,319	291,761	351,091	261,981	446,337	387,775

Unit: USD million

Country	France		Germany		The United Kingdom	
Year	Loss ratios	Returns (1−loss ratio)	Loss ratios	Returns (1−loss ratio)	Loss ratios	Returns (1−loss ratio)
1998	0.636019	0.363981	0.680887	0.319113	0.472719	0.527281
1999	0.604416	0.395584	0.715829	0.284171	0.637484	0.362516
2000	0.651959	0.348041	0.713818	0.286182	0.608027	0.391973
2001	0.653926	0.346074	0.784704	0.215296	0.679941	0.320059
2002	0.675858	0.324142	0.738685	0.261315	0.673013	0.326987
2003	0.649470	0.350530	0.705776	0.294224	0.554512	0.445488
2004	0.627477	0.372523	0.615550	0.384450	0.118702	0.881298
2005	0.589639	0.410361	0.631788	0.368212	0.122629	0.877371
2006	0.615584	0.384416	0.676730	0.323270	0.716506	0.283494
2007	0.560808	0.439192	0.702358	0.297642	0.805615	0.194385
2008	0.737082	0.262918	0.726416	0.273584	0.961418	0.038582
2009	0.647512	0.352488	0.866437	0.133563	0.967995	0.032005

2010	0.649680	0.350320	0.858203	0.141797	0.998562	0.001438
2011	0.833480	0.166520	0.973399	0.026601	1.018084	-0.018084
2012	0.892543	0.107457	0.713509	0.286491	0.949332	0.050668
2013	0.810723	0.189277	0.726557	0.273443	1.020827	-0.020827
2014	0.788916	0.211084	0.577787	0.422213	0.979042	0.020958
2015	0.803747	0.196253	0.571529	0.428471	1.036066	-0.036066
2016	0.804339	0.195661	0.588582	0.411418	0.708814	0.291186
2017	0.886237	0.113763	0.558508	0.441492	1.060286	-0.060286
2018	0.824484	0.175516	0.558905	0.441095	1.018344	-0.018344
2019	0.816211	0.183789	0.701897	0.298103	0.953500	0.046500
2020	0.905849	0.094151	0.686834	0.313166	0.940312	0.059688
2021	0.856469	0.143531	0.705252	0.294748	0.963506	0.036494
2022	0.835227	0.164773	0.746192	0.253808	0.868795	0.131205

Notes: Premiums and claims values are reported in USD million. For presentation clarity, figures are rounded in the table.

Appendix 3: Derivation of Key Optimization Results for Two-Stage Portfolio Allocation

3.1 Maximizing the Sharpe Ratio

The Sharpe ratio of the risky portfolio is

$$S_p = \frac{E(r_p) - r_f}{\sigma_p}$$

The expected return of the risky portfolio is

$$E(r_p) = \alpha_e\, E(r_e) + \alpha_s E(r_s)$$

The portfolio variance is

$$\sigma_p^2 = {\alpha_e}^2{\sigma_e}^2 + {\alpha_s}^2{\sigma_s}^2 = (1-\alpha_s)^2{\sigma_e}^2 + {\alpha_s}^2{\sigma_s}^2$$

Substituting into the Sharpe ratio

$$S_p = \frac{(1-\alpha_s)E(r_e) + \alpha_s E(r_s) - r_f}{\sqrt{(1-\alpha_s)^2{\sigma_e}^2 + {\alpha_s}^2{\sigma_s}^2}}$$

Define excess returns

$$R_s = E(r_s) - r_f$$

$$R_e = E(r_e) - r_f$$

Then the Sharpe ratio simplifies to

$$S_p = \frac{(1-\alpha_s)R_e + \alpha_s R_s}{\sqrt{(1-\alpha_s)^2{\sigma_e}^2 + {\alpha_s}^2{\sigma_s}^2}}$$

Let

$$N(\alpha_s) = (1-\alpha_s)R_e + \alpha_s R_s$$

$$D(\alpha_s) = \sqrt{(1-\alpha_s)^2{\sigma_e}^2 + {\alpha_s}^2{\sigma_s}^2}$$

Take derivatives

$$\frac{dN}{d\alpha_s} = R_s - R_e$$

$$\frac{dD}{d\alpha_s} = \frac{1}{2D(\alpha_s)}[-2(1-\alpha_s){\sigma_e}^2 + 2\alpha_s{\sigma_s}^2]$$

Apply the quotient rule

$$\frac{dS_p}{d\alpha_s} = \frac{\frac{dN}{d\alpha_s}\cdot D(\alpha_s) - N(\alpha_s)\cdot\frac{dD}{d\alpha_s}}{D(\alpha_s)^2}$$

Set the derivative equal to zero

$$\frac{dS_p}{d\alpha_s} = 0$$

Multiply both sides by $D(\alpha_s)^2$ to eliminate the denominator

$$(R_s - R_e)D(\alpha_s) - N(\alpha_s) \cdot \frac{1}{2D(\alpha_s)}[-2(1-\alpha_s)\sigma_e{}^2 + 2\alpha_s\sigma_s{}^2] = 0$$

Multiply through by $2D(\alpha_s)$ to simplify

$$2(R_s - R_e)D(\alpha_s)^2 - N(\alpha_s)[-2(1-\alpha_s)\sigma_e{}^2 + 2\alpha_s\sigma_s{}^2] = 0$$

Substituting $N(\alpha_s)$ and $D(\alpha_s)$ into the formula

$$2(R_s - R_e)((1-\alpha_s)^2\sigma_e{}^2 + \alpha_s{}^2\sigma_s{}^2) - ((1-\alpha_s)R_e + \alpha_s R_s)[-2(1-\alpha_s)\sigma_e{}^2 + 2\alpha_s\sigma_s{}^2] = 0$$

Now rearrange the terms and solve for $\alpha_s{}^*$. After algebraic simplification (see standard mean-variance optimization derivation), obtain

$$\alpha_s{}^* = \frac{R_s/\sigma_s^2}{R_e/\sigma_e^2 + R_s/\sigma_s^2}$$

This result implies that the optimal weight is proportional to each asset's excess return per unit of variance (reward-to-risk ratio).

3.2 Finding the Equation of the Efficient Risky Portfolio Curve

Given a mean and standard deviation for earth insurance (μ_e, σ_e) and the same for in-orbit satellite insurance, (μ_s, σ_s), and given a fraction α_e of business dedicated to earth, the resulting portfolio mean is

$$\mu_{\alpha_e} = \alpha_e\mu_e + (1-\alpha_e)\mu_s = \mu_s - \alpha_e(\mu_s - \mu_e)$$

and, since there is zero correlation between the two assets, the portfolio variance is

$$\sigma_{\alpha_e}^2 = \alpha_e{}^2\sigma_e^2 + (1-\alpha_e)^2\sigma_s^2$$

It is necessary to solve the second equation for α_e and substitute that into the first equation to get the equation for the frontier of investment options, $\mu_{\alpha_e}(\sigma_{\alpha_e})$. So, write the variance as

$$\sigma_{\alpha_e}^2 = \alpha_e{}^2\sigma_e^2 + (1 + \alpha_e{}^2 - 2\alpha_e)\sigma_s^2$$

Collecting terms, this simplifies the following second-order equation

$$(\sigma_e^2 + \sigma_s^2){\alpha_e}^2 - 2\sigma_s^2\alpha_e + \left(\sigma_s^2 - \sigma_{\alpha_e}^2\right) = 0$$

Using the quadratic formula, then get

$$\alpha_e = \frac{\sigma_s^2 \pm \sqrt{\sigma_{\alpha_e}^2(\sigma_e^2 + \sigma_s^2) - \sigma_s^2\sigma_e^2}}{\sigma_e^2 + \sigma_s^2}$$

Notice that there is a real solution only if

$$\sigma_{\alpha_e}^2(\sigma_e^2 + \sigma_s^2) \geq \sigma_s^2\sigma_e^2$$

$$\sigma_{\alpha_e}^2 \geq \frac{\sigma_s^2\sigma_e^2}{\sigma_e^2 + \sigma_s^2}$$

This shows that the minimum variance that can be achieved with these two assets is precisely $\frac{\sigma_s^2\sigma_e^2}{\sigma_e^2+\sigma_s^2}$.

The efficient investment frontier is then given by

$$\mu_{\alpha_e}\left(\sigma_{\alpha_e}\right) = \mu_s - \left(\frac{\sigma_s^2 \pm \sqrt{\sigma_{\alpha_e}^2(\sigma_e^2 + \sigma_s^2) - \sigma_s^2\sigma_e^2}}{\sigma_e^2 + \sigma_s^2}\right)(\mu_s - \mu_e)$$

This is expressible more generally as a general piecewise function

$$\mu(\sigma) = \begin{cases} \mu_s - \left(\dfrac{\sigma_s^2 - \sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}}{\sigma_e^2 + \sigma_s^2}\right) \times (\mu_s - \mu_e), & if\ \mu(\sigma) > \mu_s - \left(\dfrac{\sigma_s^2}{\sigma_e^2 + \sigma_s^2}\right) \times (\mu_s - \mu_e) \\ \mu_s - \left(\dfrac{\sigma_s^2 + \sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}}{\sigma_e^2 + \sigma_s^2}\right) \times (\mu_s - \mu_e), & if\ \mu(\sigma) < \mu_s - \left(\dfrac{\sigma_s^2}{\sigma_e^2 + \sigma_s^2}\right) \times (\mu_s - \mu_e) \end{cases}$$

3.3 Finding an Optimal Risky Point with No Risk-Free Asset

To begin with, since the objective is to identify the highest point of tangency between the efficient portfolio curve and an indifference curve, the solution must lie on the upper section of the portfolio curve. The equation of this part of the curve is

$$\mu = \mu_s - \left(\frac{\sigma_s^2 - \sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}}{\sigma_e^2 + \sigma_s^2} \right) \times (\mu_s - \mu_e)$$

where of course $\mu_s > \mu_e$ and $\sigma_e^2 < \sigma_s^2$. A little bit of rearranging suffices to write this as

$$\frac{(\mu_s - \mu)(\sigma_e^2 + \sigma_s^2)}{\mu_s - \mu_e} = \sigma_s^2 - \sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}$$

This is one equation in the two unknowns μ and σ that is used to find the solution. The other equation is the tangency requirement. The slope of the portfolio curve (upper section only) is

$$\left.\frac{\partial\mu}{\partial\sigma}\right|_{pc} = \frac{2(\mu_s - \mu_e)}{\sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}}$$

and it has been shown above that the slope of an indifference curve at any given point is

$$\left.\frac{\partial\mu}{\partial\sigma}\right|_{dU=0} = \frac{2\mu\sigma\rho}{2\mu^2 + \sigma^2\rho}$$

Thus, the second equation in the two unknowns is

$$\frac{2(\mu_s - \mu_e)}{\sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}} = \frac{2\mu\sigma\rho}{2\mu^2 + \sigma^2\rho}$$

After dividing both sides by 2, the expression becomes

$$\frac{\mu_s - \mu_e}{\sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}} = \frac{\mu\sigma\rho}{2\mu^2 + \sigma^2\rho}$$

In short, the following two simultaneous equations in μ and σ must be solved as functions of the five parameters $(\mu_s, \mu_e, \sigma_s, \sigma_e, \rho)$.

$$\frac{\mu_s - \mu_e}{\sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}} = \frac{\mu\sigma\rho}{2\mu^2 + \sigma^2\rho}$$

$$\frac{(\mu_s - \mu)(\sigma_e^2 + \sigma_s^2)}{\mu_s - \mu_e} = \sigma_s^2 - \sqrt{\sigma^2(\sigma_e^2 + \sigma_s^2) - \sigma_e^2\sigma_s^2}$$

These can be written in lots of different ways, but however they are expressed, the calculation of the values of our two variables that satisfy these two equations is very difficult to work out as closed form expressions of the 5 parameters in general. However, as soon as the equations are populated with specific numerical parameter values, they are no longer difficult to solve. This was the strategy employed in the text above.