Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Risk curves in enterprise risk management: linking performance targets to risk responses

Nikolaos Vourdas

nvourdas@uoa.gr

National and Kapodistrian University of Athens

Vassilis Stathopoulos

National and Kapodistrian University of Athens

Research Article

Keywords: enterprise risk management, risk curves, resource exposure, risk appetite, performance targets, risk responses

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11266874/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Risk curves in enterprise risk management: linking performance targets to risk responses

Abstract

Enterprise risk management links performance commitments to decisions about resources and risk responses. This paper develops an exposure-based framework for selecting those decisions under a quantitative risk-appetite criterion. Coefficients remain fixed within each assessed regime. A general formulation separates a policy-specific risk profile from the least committed expenditure needed to meet the criterion. Analytical and hypothetical examples show when resource reallocation helps, how response costs change selection, and why local marginal costs cannot price every response switch. A multi-output activity model compares expenditure and expected-total-cost objectives, tests insurance pricing and programme effectiveness, and introduces binding resource constraints. Assessing a selected plan unchanged under alternative assumptions distinguishes its robustness from the options available after reassessment. A documented input procedure and lexicographic implementation make the calculations reproducible. The framework supports financial appraisal within ERM while retaining management responsibility for objectives, criteria and approval.

Keywords: enterprise risk management; risk curves; resource exposure; risk appetite; performance targets; risk responses.

1. Introduction

A business objective becomes an operating commitment when management assigns the resources needed to deliver it. Enterprise risk management (ERM) supports that commitment by bringing risk into strategy-setting, performance management and the pursuit of value. COSO treats these activities as part of the organisation's culture, capabilities and practices, rather than as a separate risk function (COSO, 2017a). A useful assessment therefore connects the objective with the operating plan, its resource requirements, its financial consequences and the decisions for which management is accountable.

COSO uses risk profiles to show the amount of risk associated with different levels of performance. Performance lies on the horizontal axis and the selected risk measure on the vertical axis. The profile supports discussion of business objectives, Risk Appetite, Performance Targets, Tolerance and Risk Responses. Its shape follows the business context and the risks assessed; COSO does not prescribe a universal curve or a particular mathematical formula (COSO, 2017a, pp. 15–16; 2017b, Appendix D).

This paper develops one quantitative construction for a production-related business objective. Its starting point is resource exposure. Producing more output generally requires a different quantity of equipment, labour or other resources. Expected loss can therefore change even when the expected loss per unit of each resource remains constant. The performance relationship enters through resource requirements and any additional, explicitly identified exposure terms.

The research question is: how can an exposure-based risk curve connect a performance target to a resource plan and a proportionate risk response within ERM? We first derive the risk profile of a cost-minimising operating policy. We then assess whether that policy meets a specified financial criterion, compare alternative allocations and responses, and examine whether unit-level plans can be implemented together. This separates three decisions that are easily conflated: minimising production expenditure, minimising expected loss and minimising expenditure subject to the selected risk-appetite criterion.

The framework links three objects: the risk profile of a stated operating policy, the least expenditure needed to meet a risk criterion, and the resource-and-response choice delivering that expenditure. A general formulation accommodates different technologies and vectors of outputs. Its marginal-cost analysis makes the financial consequences of changing the criterion explicit. Two contrasting production examples and a portfolio example show how the construction works; a documented input procedure specifies what an application must supply. The contribution is a decision protocol that makes three potential errors visible: selecting a response from curve height alone, treating the marginal cost within one response as the cost of switching responses, and treating a reoptimised scenario as evidence that the selected plan is robust. It specifies the additional cost, input and feasibility information needed to avoid those errors.

The study is analytical, with hypothetical numerical applications. Coefficients, prices and contract terms are inputs to the assessment, not estimates from company records. The activity-model inputs and the illustrative ledger are separately declared assumptions. The Supplement contains the derivations, complete comparisons, input records and numerical checks. The main article concentrates on the reasoning and its consequences for strategy, objectives and resource decisions.

2. Position within ERM research

Decision-oriented ERM already links risk to capital allocation and business planning. Nocco and Stulz (2006) distinguish entity-wide protection of access to resources from business-unit ownership of risk–return decisions. Their discussion of project appraisal connects local choices to their contribution to total company risk. Gleißner and Berger (2024) likewise favour value drivers, quantitative aggregation and integration into planning and budgeting, rather than minimising individual risks. The present study works within that decision-oriented view. It examines the narrower task of costing feasible responses to a specified operating commitment; it does not replace a valuation or capital-adequacy model.

Aven (2004, 2010) examines how assessments of risk and uncertainty support decisions and how acceptance criteria interact with other decision criteria. That distinction matters here: an expenditure minimum, an expected-total-cost minimum and management approval are different outcomes. Donnelly *et al* (2012) address risk registers and corporate risk communication. Crawford and Jabbour (2024) consider the interaction between ERM and managerial judgement. Marc *et al* (2023) find, in their sample, that the relationship between advanced risk-system design and effectiveness is mediated by interactive use. These contributions motivate reporting assumptions and alternatives with the numerical recommendation, rather than treating an optimisation result as an approval.

Reviews identify substantial variation in ERM definitions, implementation and evaluation (Bromiley *et al*, 2015; Anton and Afloarei Nucu, 2020). Empirical studies also examine different performance measures and organisational settings (Gordon *et al*, 2009; Callahan and Soileau, 2017). That literature establishes the importance and context of ERM, but does not calibrate the hypothetical production and loss inputs used below.

Table 1 separates the established foundations from the contribution developed here. Standard production theory supplies conditional resource demands (Varian, 2010). Activity analysis and constrained optimisation supply feasible-set modelling, resource selection and sensitivity analysis (Dantzig, 1963; Boyd and Vandenberghe, 2004). COSO supplies the objective-centred language and governance context. The paper combines these foundations into a traceable sequence: specify the policy behind a financial risk profile; price the feasible resource-and-response alternatives; identify the cost of changing the criterion; test the selected plan against consequential assumptions; and check joint enterprise commitments.

Table 1. Established foundations and the paper's applied contribution.

Foundation	**Role of established work**	**Connection developed here**
Decision-oriented ERM	Nocco and Stulz (2006): resource access and unit risk ownership. Gleißner and Berger (2024): value drivers, planning and aggregation.	Cost a specified performance commitment; keep its financial appraisal distinct from enterprise valuation.
Production and optimisation	Conditional demands, feasible sets and dual sensitivity (Varian, 2010; Dantzig, 1963; Boyd and Vandenberghe, 2004).	Connect the policy profile, priced response alternatives and cost of tightening a criterion; expose response switches.
COSO ERM	Objectives, performance, appetite, responses, portfolio review and governance (COSO, 2017a, 2017b).	Report target, policy, selected package, resource commitments and the quantitative criterion in one decision record.
Assessment and its use	Interpretation of uncertainty and decision criteria (Aven, 2004, 2010); interactive use (Marc et al, 2023).	Separate testing the unchanged plan from reoptimising alternatives; identify assumptions that can change approval.

The analytical distinctions are consequential even in small examples. A lower reference risk curve can accompany a smaller affordable target. Cost and risk minima coincide only under a particular alignment of input weights in the two-factor model. Discrete packages can create response switches that cannot be priced with one local multiplier. Finally, an alternative plan that becomes feasible after reassessment does not establish that the original commitment was protected.

3. Constructing the risk curve from resource exposure

3.1. A fixed assessment regime

Let q denote planned output during one assessment period. A regime s specifies the strategy, technology, resource prices, existing controls and any additional risk response. For each distinct loss category i, κ_{is} is expected retained financial loss per unit of exposure. In this work and for simplicity reasons, it is constant throughout the declared operating range of that regime. An operating policy π determines the resource plan at each output, giving exposure $e_{is}(q; \pi)$. The resulting curve is

$$R_s^{\pi}(q) = \sum_{i=1}^{m} \kappa_{is}\ e_{is}(q;\pi), \qquad \kappa_{is} \text{ constant.} \tag{1}$$

For the basic model, the exposures are equipment units K and labour hours L. Their coefficients have different denominators, but both products are expected monetary losses over the same period. Equipment quantities are continuous exposure units for analytical purposes; an operational application with indivisible machines would use a discrete feasible set.

Performance dependence is already represented in $e_{is}(q; \pi)$. With unchanged coefficients, a change in expected loss follows from the change in exposure. A different strategy or an implemented response may justify a different, separately documented coefficient set. Each set remains fixed when its own profile is plotted. The model, in this work, does not automatically increase or decrease a coefficient when output or the equipment–labour mix changes.

Financial consequences must be defined without overlap. Planned wages, equipment costs and insurance premiums enter expenditure. Uncertain repair, absence, rework or interruption expenses enter the relevant loss categories. Ordinal likelihood–impact scores are not converted into euros. Expected losses of distinct consequences can be added without an independence assumption; tail losses and liquidity effects require a joint loss model and separate criteria.

The control state also needs a name. A profile based on existing measures describes assessed residual risk. Inherent risk requires an explicitly specified absence of the relevant management actions. Target Residual Risk is management's preferred remaining risk, while Actual Residual Risk concerns the risk remaining after actions have been taken (COSO, 2017a, p. 77).

3.2. Selecting resources and responses together

Let y be the vector of required outputs and x the vector of operating decisions. For each admissible regime s, $X_s(y)$ records the technical requirements, available resources and other obligations. The set can describe a production function, a network of activities or a discrete equipment choice. It states explicitly whether output must equal the target or may exceed a minimum requirement. The set of regimes contains the response packages management is prepared to consider, including the existing arrangements.

Total expenditure $C_s(x)$ includes production expenditure and response expenditure. A response may have a fixed implementation charge, a premium proportional to an insured exposure, or a stated operating cost. Expenditure already included in a resource price is not added again. The expected-loss assessment $R_s(x,y)$ uses the same fixed-coefficient sum as Equation (1). The joint decision is

$$V_s(\mathbf{y}, A) = \min_{\substack{\mathbf{x}\in\mathcal{X}_s(\mathbf{y}) \\ R_s(\mathbf{x},\mathbf{y})\le A}} C_s(\mathbf{x}),$$
$$V(\mathbf{y}, A) = \min_{s\in\mathcal{S}} V_s(\mathbf{y}, A). \tag{2}$$

$V_s(y,A)$ is the least expenditure within one regime at the quantitative Risk Appetite criterion A; $V(y,A)$ is the least across regimes. An infeasible regime has value $+\infty$ and cannot be selected. Solving each regime and comparing its optimum is equivalent to selecting the regime and its resources jointly. It allows a stricter criterion to change both the resource mix and the selected risk response. The costs and effects of available responses are inputs; the quantities and selected response are decisions.

The principal objective is committed expenditure: the resource purchases, payroll allocation, programme charges and premiums that management must authorise for the period. This is appropriate for a stated budget-allocation mandate in which a separate expected-loss criterion has already been approved. It does not make retained losses economically irrelevant. A manager seeking the lowest expected financial cost should instead minimise expenditure plus expected retained loss over the same feasible alternatives. Section 5.2 compares these two mandates without changing targets, risk limits or the response menu. Neither mandate, by itself, establishes strategic value: benefits, liquidity, severe-loss exposure and service obligations require their own assessment.

Within a response regime, equal expenditure can correspond to different assessed losses. The implementation therefore minimises expenditure first, expected loss second on the expenditure-optimal set, and then uses a fixed coordinate order for any remaining tie. The same declared priorities apply between regime solutions. This makes the reported plan, and thus its risk profile, reproducible rather than dependent on an arbitrary solver choice. Supplement S18 specifies the algorithm and tests degenerate cases.

A reference profile answers a different question. It plots the loss of a specified policy, such as cost minimisation without the appetite constraint, and is derived independently of A. The profile of the jointly selected plan also depends on the criterion used in its selection. For multiple outputs, a two-dimensional curve requires an explicit path y(q); otherwise the result is a surface or a set of target-specific decisions. A reported curve must identify its path, policy and tie-breaking rule.

This is conditional, static optimisation with a general feasible-set description. Application to another organisation requires its own technology, exposures, costs and response evidence; it does not require a universal production equation. Supplement S11 gives conditions for existence, an implementation algorithm and the distinction between convex subproblems and discrete response selection. Review and Revision reruns the assessment with justified, versioned inputs; it is not an automatic dynamic-learning rule.

3.3. The cost-minimising reference policy

Consider the technology q = 10√(KL), with positive resource prices r and w. Minimising rK + wL at a specified output gives

$$K^C(q) = \frac{q}{10}\sqrt{\frac{w}{r}}, \qquad L^C(q) = \frac{q}{10}\sqrt{\frac{r}{w}}. \tag{3}$$

The numerical benchmark sets r = 5, w = 20, κK = 2 and κL = 7. The reference policy then uses K = 0.2q and L = 0.05q. Expenditure is 2q, equipment-related expected loss is 0.4q and labour-related expected loss is 0.35q. Total expected loss is 0.75q. Increasing output from 1,000 to 1,300 therefore changes expected loss from €750 to €975, while both risk coefficients remain unchanged.

Figure 1 places this construction within the ERM process. Management establishes the business objective and the proposed target. The operating unit identifies resource requirements. Finance and risk specialists examine the cost and exposure calculations, and the responsible management body selects the plan. The portfolio check and subsequent review connect the individual assessment with the wider organisation.

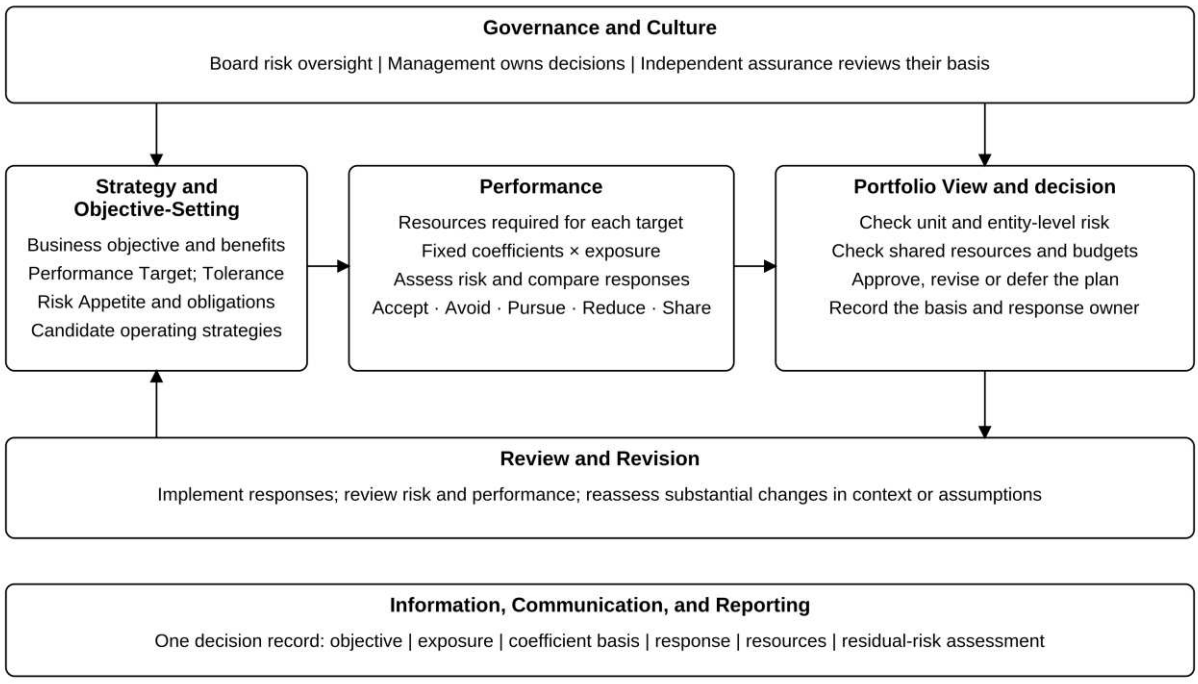


Figure 1. From a business objective to an ERM decision. The diagram links Strategy and Objective-Setting, Performance and Review and Revision, supported by Governance and Culture and Information, Communication, and Reporting. Portfolio View is part of Performance. The arrangement is an application of COSO (2017a; 2017b, Appendix C), not a prescribed organisation chart. Coefficients are fixed within the assessed regime.

3.4. Risk Appetite, Performance Target and Tolerance

Figure 2 overlays a €980 expected-loss criterion on the reference profile. At the target of 1,000 units, expected loss is €750. The reference policy reaches the criterion at approximately 1,306.67 units. That crossing is a boundary of the policy's financial assessment.

Risk Appetite describes the types and amount of risk the organisation is willing to accept in pursuing value. The monetary ceiling here is one operational criterion within that broader expression. Tolerance concerns

acceptable variation in performance around the target, whereas Risk Capacity concerns the risk the entity can absorb. A spending budget and a stock of equipment are resource constraints, not measurements of Risk Capacity (COSO, 2017a, Principle 7; 2017b, pp. 20–22; Martens and Rittenberg, 2020).

The numerical criterion leaves €230 of expected-loss headroom at the 1,000-unit target. This does not prove that management should increase output. The benefit of additional performance, service obligations, other risks and available resources remain part of the decision. Likewise, the lowest point of a risk curve need not be the most valuable performance target.

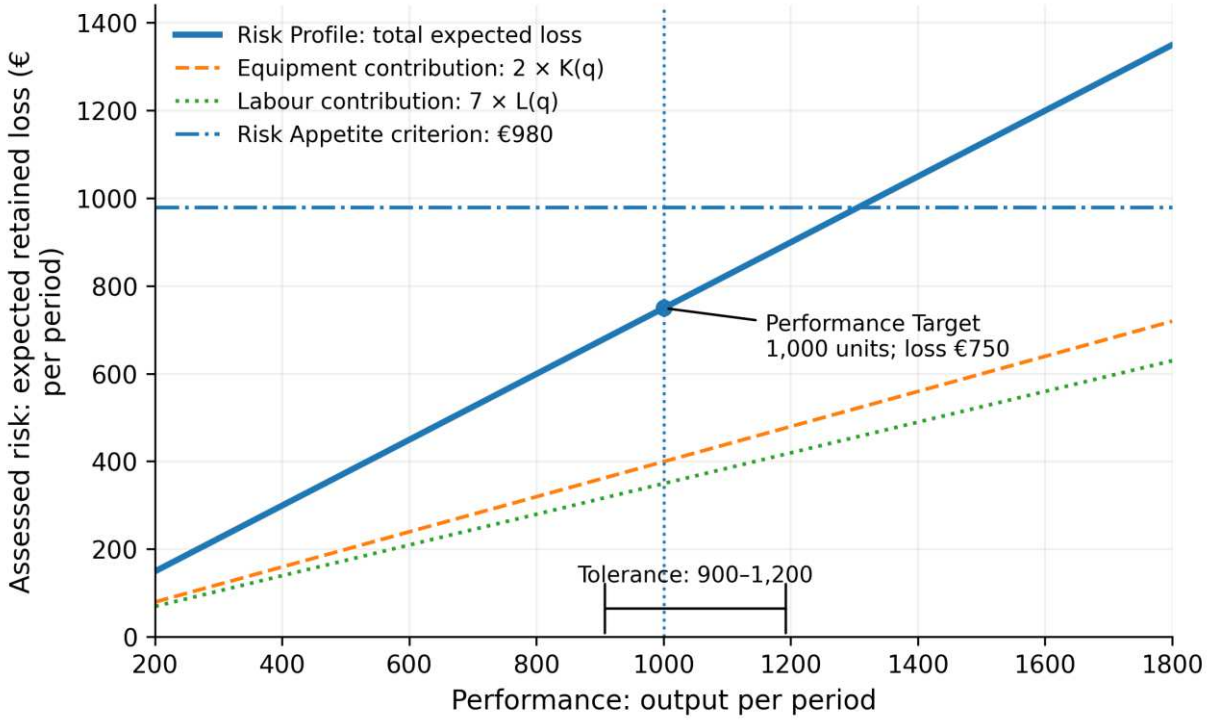


Figure 2. Business-objective Risk Profile under fixed coefficients. The total curve is R = 0.75q, comprising equipment loss 0.4q and labour loss 0.35q. The Risk Appetite criterion is €980 per period. The Performance Target of 1,000 and illustrative Tolerance of 900–1,200 are separately specified management choices. No Risk Capacity value is inferred from these means or from resource availability. All values are hypothetical.

4. Using the curve to select resources and responses

4.1. Reduce: meet the criterion at the least additional cost

At output 1,300, the reference plan uses K = 260 and L = 65, costs €2,600 and has expected loss of €975. For this separate decision, management sets a €973 expected-loss criterion. The question is whether a different equipment–labour mix can meet it. Coefficients remain 2 and 7 throughout the comparison. The resource decision is

$$\min_{K,L>0}\{rK + wL\} \quad \text{subject to} \quad 10\sqrt{KL} = q, \quad \kappa_K K + \kappa_L L \leq A. \tag{4}$$

Table 2 and Figure 3 show four distinct criteria applied to the same target. The least-cost allocation meeting €973 uses approximately 247.785 equipment units and 68.204 labour hours. Equipment-related expected loss falls from €520 to €495.57, while labour-related expected loss rises from €455 to €477.43. The net reduction is €2 and additional expenditure is €3.01.

Table 2. Four resource decisions at the same Performance Target of 1,300 units. Monetary amounts are euros per period.

Decision criterion	K	L	Expenditure	Expected loss
1. Minimum expenditure	260.000	65.000	2,600.00	975.00
2. Minimum expenditure plus expected loss	255.315	66.193	2,600.43	973.98
3. Least expenditure meeting criterion	247.785	68.204	2,603.01	973.00
4. Minimum expected loss	243.208	69.488	2,605.80	972.83

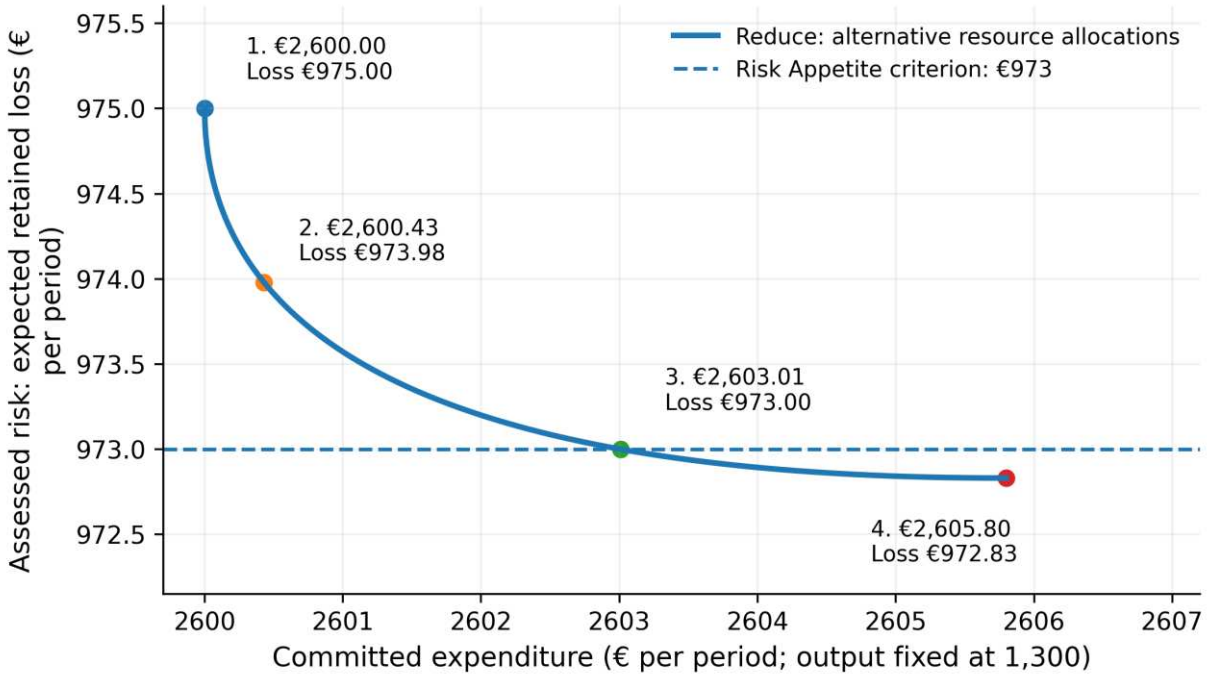


Figure 3. Comparing resource decisions at Performance Target 1,300. Points 1–4 correspond to Table 2: minimum expenditure; minimum expenditure plus expected loss; least expenditure meeting the €973 Risk Appetite criterion; and minimum expected loss. The proposed Reduce response at point 3 changes exposure only: κK = 2 and κL = 7 at every point. This is a cost–risk comparison at fixed performance, not a performance-axis Risk Profile. The close vertical scale makes the small numerical differences visible.

The risk-minimising allocation achieves expected loss of €972.83, but costs €2,605.80. It is therefore more expensive than necessary to meet the €973 criterion. Conversely, minimising expenditure plus expected loss yields €973.98 of expected loss and fails that separate criterion. The distinction is a management choice about the decision rule, not an inconsistency between the calculations.

The room for reallocation depends on the relative cost and loss weights, not just on the presence of a risk limit. Write a = κK and b = κL. For this technology,

$$\frac{R^{C}(q)}{R_{\min}(q)} = \frac{1}{2}\left(\sqrt{\frac{aw}{br}} + \sqrt{\frac{br}{aw}}\right) \geq 1. \qquad (5)$$

The ratio equals one exactly when aw = br. In that case the expenditure-minimising and loss-minimising allocations coincide, and changing the resource mix cannot reduce expected loss at the fixed output. The benchmark ratio is approximately 1.00223: at most about 0.22% of its reference expected loss can be removed by reallocation. This explains the small vertical range in Figure 3. It is a property of the illustrative

inputs, not an estimate of the value of ERM. Supplement S12 provides the derivation and contrasting weight combinations.

The same model also identifies a limit to reallocation. If management instead requires expected loss no higher than €970, no equipment–labour mix can satisfy it at output 1,300 without another response. The minimum remains €972.83. A mathematically infeasible reallocation calls for consideration of other Risk Responses or review of the objective; it does not itself constitute a COSO Avoid decision.

4.2. Share: compare retained loss and response expenditure

COSO includes insurance within Share responses (2017a, Principle 13). We examine two hypothetical contracts against the €970 criterion. Equipment insurance retains one quarter of equipment-related loss and adds €3 per equipment unit. Its effective coefficient is therefore 0.5 and its equipment price is 8. Labour insurance retains five sevenths of labour-related loss and adds €5 per hour, giving coefficient 5 and labour price 25. These are contract-specific constants, not functions of performance.

Resource prices change when premiums are included, so each contract has its own cost-minimising mix. Table 3 reports those plans at output 1,300. Both meet the criterion. Labour insurance has lower expenditure, while equipment insurance provides a larger reduction in retained expected loss. Premiums are already included in total expenditure and are not added twice.

Table 3. Share responses at output 1,300. Each row uses its own cost-minimising resource policy; the Risk Appetite criterion is €970.

Response	K	L	Expenditure €	Retained loss €
N	260.000	65.000	2,600.00	975.00
IK	205.548	82.219	3,288.77	678.31
IL	290.689	58.138	2,906.89	872.07

N: reference; IK: equipment insurance; IL: labour insurance. Premiums included in expenditure are €616.64 for IK and €290.69 for IL. N does not meet the €970 criterion.

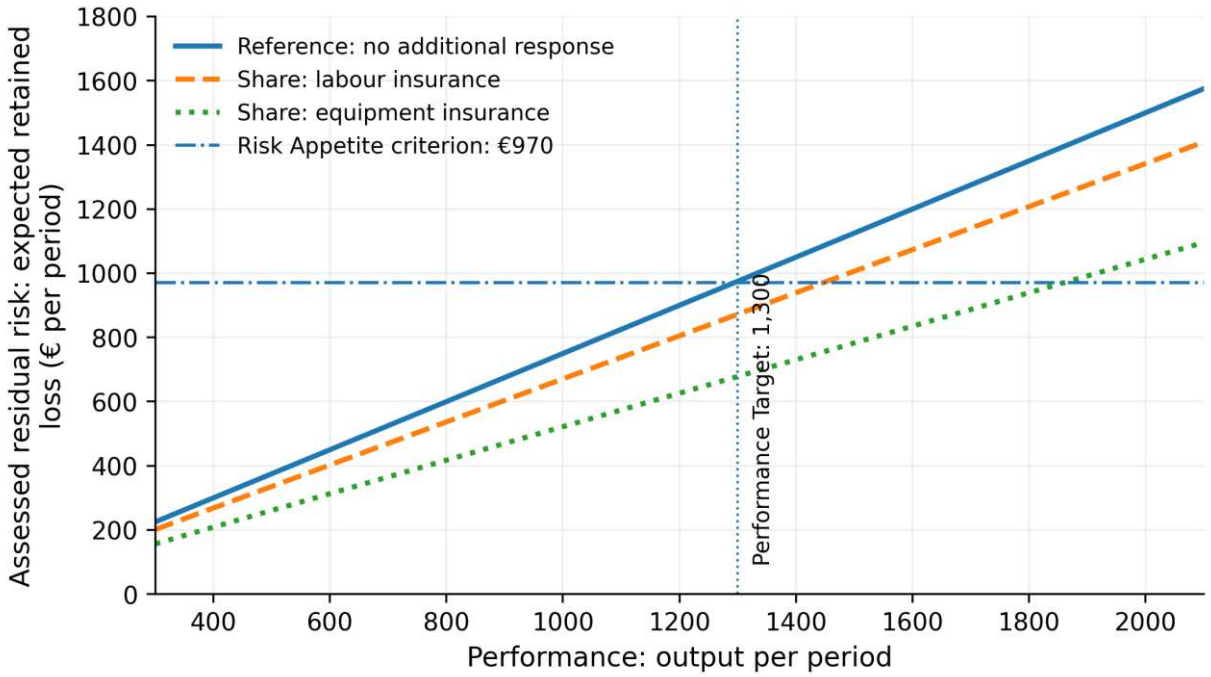


Figure 4. Risk Profiles under alternative Share responses. Each curve follows the cost-minimising resource policy for the stated contract. Coefficients are fixed within each policy: reference (2, 7), equipment insurance (0.5, 7), and labour insurance (2, 5). The line at €970 is the selected Risk Appetite criterion. Its intersections are policy-specific financial boundaries, not automatically approved Tolerance limits. Insurance changes retained consequences; it does not by itself secure service continuity.

Figure 4 shows why response selection also depends on the target. The cost-minimising labour-insurance policy reaches €970 at about 1,445.99 units, whereas the equipment-insurance policy reaches it at about 1,859.04. Further reallocation can extend a policy's feasible range, so these crossings are not the maximum attainable outputs across all input mixes. At output 1,600, even the risk-minimising labour-insurance allocation exceeds €970. Equipment insurance meets it at expenditure of €4,047.72.

An expenditure budget adds another decision. With a €3,500 budget and the same €970 criterion, the largest output supported by any of the three response regimes is about 1,523.81 units. It uses labour insurance and a mix of K = 270 and L = 86; expenditure and expected loss meet their limits exactly. Equipment insurance supports only 1,383.50 units within that budget despite its lower reference risk curve. Supplement S4 gives the full optimisation. A lower risk curve and a larger affordable target are different properties.

4.3. Alternative strategies and resource commitments

COSO calls for evaluation of the risk associated with alternative strategies (2017a, Principle 8). Consider an alternative operating arrangement with labour price 10 and labour coefficient 9, while equipment price and coefficient remain 5 and 2. These four constants define a new regime. At output 1,000, its cost-minimising plan costs €1,414.21 and has expected loss of €919.24, compared with €2,000 and €750 under the reference strategy. It reaches the €980 criterion at about 1,066.10 units. The lower production cost is accompanied by less headroom under that policy. Benefits and wider consequences are needed before choosing between the strategies.

The time horizon also changes the feasible resource set. If equipment is committed at K0, labour is $q^2/(100K0)$ and expected loss is $2K0 + 7q^2/(100K0)$. The equipment coefficient and equipment contribution

are fixed; the labour exposure changes with output. At output 1,000, K0 = 200 gives expected loss of €750, while K0 = 400 gives €975. More equipment reduces the labour requirement but increases the equipment exposure. It therefore does not automatically enlarge the output range allowed by a given risk criterion. Supplement S5 derives the boundaries and shows the full comparison.

4.4. What does a tighter risk criterion cost?

The slope of a Risk Profile and the marginal cost of tightening an appetite criterion answer different questions. The first measures the change in expected loss when performance changes under a stated policy. The second holds the target fixed and asks how much more expenditure is required when the expected-loss limit is reduced.

For a convex regime problem with strong duality and a differentiable value function, the multiplier λ_s on the loss constraint gives that marginal expenditure:

$$M_s(\mathbf{y}, A) = -\frac{\partial V_s(\mathbf{y}, A)}{\partial A} = \lambda_s \geq 0. \tag{6}$$

Its unit is euros of additional expenditure per euro reduction in the expected-loss limit, for the same period. This is the standard sensitivity interpretation of the constraint multiplier (Boyd and Vandenberghe, 2004, Section 5.6), applied here to risk-response appraisal. It is not a probability, an insurance price or the value of ERM.

At output 1,300 without additional insurance, the marginal expenditure is about €0.97 at A = 974 and €6.90 at A = 973. These are local rates: tightening from 974 to 973 costs €2.60 in total, rather than either endpoint rate. The rate rises sharply near the minimum achievable loss of €972.83. With only the three specified regimes N, IK and IL available, crossing below that minimum changes the selected response; the least expenditure jumps from about €2,605.80 to €2,906.89. A derivative within the uninsured regime cannot price that switch. Its size depends on the available packages; fractional or combined cover, if admitted and priced, would define a different decision problem. Supplement S12 reports the derivatives, finite changes and the one-sided comparison at the boundary.

For ERM, the implication is to report the cost of the proposed change in criterion together with the response it requires. Management can then compare a small allocation adjustment with a different response package. The numerical analysis informs the choice of appetite; it does not determine the organisation's willingness or ability to bear risk.

5. A different technology: three resources and two outputs

5.1. Joint choices under fixed resource recipes

The same decision structure can be used without a Cobb–Douglas production function. Consider a separate hypothetical operation producing 100 units of product A and 80 of product B. Each product can use an

automated or manual route. Four activity quantities are chosen, and a technical matrix D converts them into equipment, labour and energy requirements:

$$x_1 + x_2 = 100, \quad x_3 + x_4 = 80, \quad (K, L, E)^{\mathsf{T}} = D\mathbf{x}, \quad \mathbf{x} \geq 0. \tag{7}$$

The matrix and all units are given in Supplement S13. Automated routes use more equipment and energy but less labour. Resource prices are 5, 20 and 2; the corresponding reference loss coefficients are 4, 1 and 0.5. Resource availability is 100 equipment units, 80 labour hours and 140 energy units. These inputs are stipulated for the illustration, not estimated operating requirements.

Management compares no additional response with two packages. A Reduce programme costs €100 per period and is assumed to reduce the expected equipment-related consequence by 60%, lowering its coefficient from 4 to 1.6. A Share contract retains half of the equipment consequence and charges €0.80 per equipment unit. Labour and energy coefficients are unchanged. Each package has fixed coefficients over the operating range. The programme and contract are alternatives in this example; their combination is not an available option.

Figure 5(a) shows two levels of selection. For the same output targets and risk criterion, the production routes are optimised separately within each response regime, including its programme charge or premium. The feasible regime optima are then compared to identify the least-expenditure plan. These are parallel alternatives, not a sequence of responses; no additional programme still permits resource reallocation. The candidate plan remains subject to the assumption tests and enterprise review described in Sections 6 and 7.

Without the loss constraint, all output uses automated routes: K = 88, L = 26 and E = 114, with expenditure €1,188 and expected loss €435. With a €430 criterion, a small shift to the manual route for product B is sufficient. At €420, equipment insurance is cheaper than a larger reallocation. At €250, the Reduce programme is preferred. Table 4 reports the jointly selected plans.

Table 4. Joint choices for outputs (100,80) in the activity model. The criterion and monetary amounts are euros per period.

Criterion A	Selected plan	Expenditure	Expected loss	Labour hours
430	Exposure reallocation	1,213.93	430.00	27.85
420	Share: insurance	1,258.40	259.00	26.00
250	Reduce programme	1,288.00	223.80	26.00
220	Reduce programme	1,356.21	220.00	30.87

All prices and coefficients are fixed within each regime. Costs include the selected programme charge or premium. Resource limits: K ≤ 100, L ≤ 80, E ≤ 140. Detailed allocations and provenance are in Supplement S13.

At the €250 criterion, no unassisted reallocation is feasible within the resource limits; its minimum loss is €293.75. The Reduce programme supports the target at expenditure €1,288 and expected loss €223.80. The best Share plan also meets the criterion, at €1,367.71. If the programme charge were €200 rather than €100, Share would be preferred. Thus the response expenditure has a direct effect on the selected plan, rather than being reported only after selection.

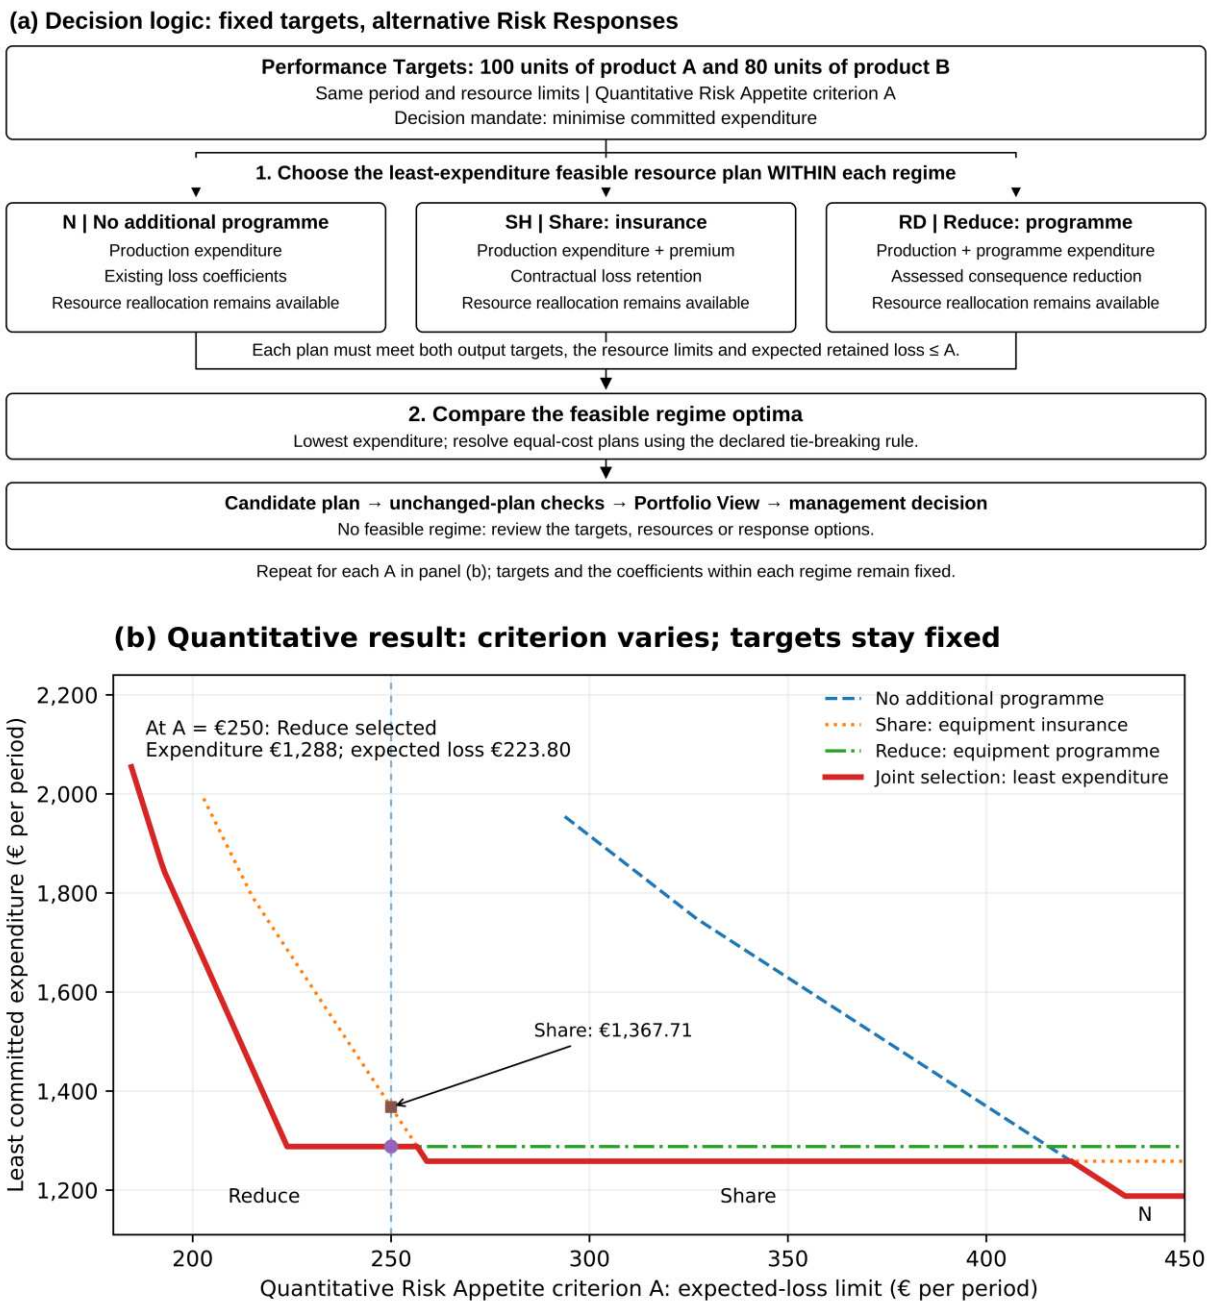


Figure 5. Joint selection of resources and Risk Responses at fixed Performance Targets. (a) Optimise resources within each regime, compare feasible minima and send the candidate plan for assumption testing, Portfolio View and management review. (b) Repeat the calculation as the expected-loss criterion A varies. Targets remain (100,80) and coefficients stay fixed within each regime. At A = €250, N is infeasible; Share costs €1,367.71 and Reduce costs €1,288 with expected loss €223.80. The lower envelope is the comparison result, not a fourth response. Panel (b) plots the criterion, not performance. All values are hypothetical.

Figure 5(b) repeats the comparison for each value of A while holding targets, technology, resource availability, prices and regime coefficients fixed. Moving left tightens the criterion. The lower envelope records the least committed expenditure among feasible regimes; it is not a fourth response or a path of increasing performance. The vertical guide at €250 links the process in panel (a) to the corresponding decision in Table 4.

Within each regime the problem is a linear programme. Its expenditure value is non-increasing and convex in A, but the minimum across discrete packages need not be convex. A tighter criterion can select a package with spare room below its loss limit and zero local marginal expenditure. At A = 250, Reduce has €26.20 of expected-loss headroom; at A = 220, further tightening within that regime costs about €17.95 per euro

locally. Supplement S13 reports the full allocations, analytical solutions, dual checks and fixed-input comparisons.

5.2. The decision mandate and insurance price matter

At the €430 criterion, the selected reallocation commits €1,213.93 and leaves expected loss of €430. The feasible Reduce programme commits €1,288 and leaves expected loss of €223.80. Their expenditure-plus-loss totals are €1,643.93 and €1,511.80. Minimising that total, rather than committed expenditure, selects Reduce. At €420 the expenditure mandate selects Share while the expected-total-cost mandate again selects Reduce. The objectives agree at €250 for the stated inputs. Supplement S15.1 solves both objectives under identical targets, resource limits and expected-loss criteria. Reporting both financial quantities prevents the phrase ‘least cost’ from concealing a different decision mandate.

The illustrative Share premium also needs context. Its €0.80 charge per equipment unit is below the €2 expected recovery implied by the organisation’s stipulated loss assessment: 4 minus 2 per exposure unit. At the reference allocation, premium and assessed expected recovery are €70.40 and €176. The price is an input chosen to illustrate the comparison, not a calibrated market quotation or a pricing model. We do not assume that the insurer shares the organisation’s assessment.

Supplement S15.2 therefore varies the premium to €1, €2, €2.40, €3 and €4 per unit, with unchanged coverage. At a premium of €2.40 and criterion €420, expenditure minimisation selects unassisted reallocation at €1,265.78 rather than the nominal Share plan. At a €300 criterion it selects Reduce at €1,288. At €250, Reduce remains preferred. The intermediate Share choice is sensitive to its favourable assumed price; the broader requirement to compare priced response packages remains. These alternatives are fixed contract scenarios, not performance-dependent coefficient adjustments.

Figure 6 complements the fixed-target expenditure comparison in Figure 5(b). It plots risk against the explicit performance path y(q) = (q,0.8q), using each regime’s cost-reference policy. The lower expected-loss line does not include the cost of obtaining it. Reading the profile together with the expenditure comparison links the Performance Target, financial risk criterion and selected Risk Response.

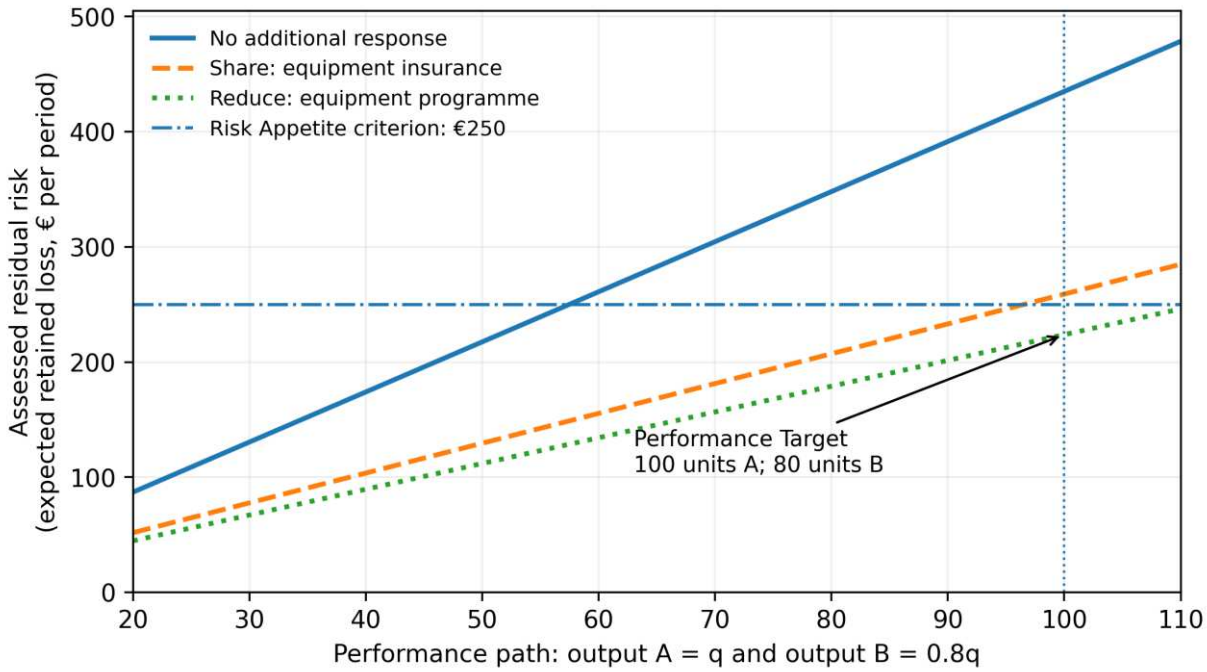


Figure 6. Business-objective Risk Profiles for the activity technology. Product A equals q and product B equals 0.8q; both quantities are retained in the optimisation. Each curve uses its regime's all-automated cost-reference policy on 20–110 and fixed coefficients. The €250 Risk Appetite criterion and target (100,80) are shown. The Reduce programme changes the assessed equipment consequence; Share changes retained financial loss. Programme cost and premium enter the separate expenditure comparison in Figure 5(b). Inputs are hypothetical.

6. Testing the selected plan and resource availability

6.1. A scenario optimum does not protect the unchanged plan

Before approval, the selected resources and response are evaluated unchanged under alternative fixed-input assumptions. At the €250 criterion the nominal choice is Reduce with all production automated, expected loss €223.80 and expenditure €1,288. Increasing all its loss coefficients by 20% gives expected loss €268.56 for that same plan. It fails the financial criterion. A separate optimisation within Reduce restores feasibility by moving 39.66 units of product B to the manual route, increasing period expenditure to €1,565.61. Table 5 distinguishes the unchanged-plan assessment from this alternative allocation.

Table 5. Test the selected plan before reoptimising: output (100,80), criterion €250.

Assessment assumption	Unchanged-plan loss €	Meets criterion?	Reallocated RD cost €	Reallocated RD loss €
Nominal	223.80	Yes	1,288.00	223.80
Coefficients +20%	268.56	No	1,565.61	250.00
Programme effect 50%	259.00	No	1,402.55	250.00

The unchanged plan is the all-automated Reduce (RD) plan costing €1,288. Reallocated columns show a separate full-period least-expenditure plan within RD; they do not assume costless post-implementation adjustment. Scenario inputs are hypothetical constants.

Response effectiveness is examined independently of a common coefficient change. A programme delivering 50% rather than 60% equipment-loss reduction gives expected loss €259 for the unchanged plan. Reallocation within that programme can meet the criterion at €1,402.55. A fresh appraisal across packages instead selects Share at €1,367.71. For the nominal prices and criterion, the expenditure preference switches

from Share to Reduce at approximately 50.82% effectiveness. This differs from the 52.56% reduction needed for the all-automated programme plan to meet the criterion without reallocation. The former is a response-choice threshold; the latter is a fixed-plan feasibility threshold. Supplement S16 derives and tests both.

These calculations describe pre-approval appraisal or a new full-period planning exercise. Reallocating an already implemented plan may incur switching costs, sunk programme charges or delay; those costs must be added before presenting the alternative as an available corrective action. Changing packages requires renewed approval and cannot be treated as a costless mid-period adjustment. No probability is assigned to the deterministic scenarios. Coefficients are constant within each scenario; uncertainty about their values does not turn them into functions of performance.

6.2. A binding resource can change the response

Resource availability is part of the decision, not merely a check after risk minimisation. Retain the activity model's targets, prices and loss coefficients, but reduce energy availability from 140 to 80 units. At the €250 criterion, the least-expenditure plan now uses Share, with 100 automated units of A and a split of 12 automated and 68 manual units of B. It uses K = 60.8, L = 60 and E = 80, commits €1,712.64 and leaves expected loss €221.60. Reduce is feasible with the same route mix at €1,764 and expected loss €197.28, but requires more committed expenditure. The energy constraint therefore changes both the operating mix and the selected response from their nominal values.

The role of another resource is visible when labour availability is also reduced to 58 hours. Saving 34 energy units from the all-automated plan requires at least 34 additional labour hours in the available routes. Total labour must therefore be at least 60 hours. None of the three response packages can meet the exact targets with both limits, because neither insurance nor the programme changes the technical recipes. Supplement S17 verifies this bound, compares all branches, and also reports an equipment-scarcity case. The result calls for another technical option, additional resources or objective review, not a claim that every enterprise with limited energy should purchase insurance.

7. Application in ERM, strategy and portfolio decisions

7.1. From the risk register to the operating proposal

The inputs are assembled before the resource-and-response comparison. First, the operating owner specifies the objective, output units, period and admissible plans. Second, engineering or process records support the technical constraints and exposure requirements. Third, Finance reconciles distinct retained consequences with incident and exposure records, separating planned expenditure, recoveries and losses. Fourth, the assessment states the coefficients, response effects, their source and the operating domain to which they apply.

With comparable and complete records, a coefficient can be estimated as total retained financial loss divided by total exposure. For a suitable frequency–severity model, the same quantity is estimated by event count

per exposure unit multiplied by mean retained severity. The denominator and loss boundary must match. NIST/SEMATECH (n.d.) provides the corresponding count-to-exposure logic for a constant repair-rate model; applying it outside that observation process requires its own justification.

Each input is identified as observed, estimated, elicited or stipulated. Zero recorded events do not establish zero risk. Changes in technology, incomplete loss reporting or unrepresented control conditions require further assessment rather than an automatic coefficient adjustment. A programme's effectiveness is not established by the baseline incident ledger. In the activity example, the assumed 60% equipment-loss reduction needs separate support; at the reference mix, at least 52.56% is needed to meet A = 250 without reallocation. This makes the evidence requirement specific to the proposed decision.

The final steps are to test the selected plan unchanged, compare justified alternatives where needed, and release a versioned record with an owner, Finance reconciliation, risk challenge and management approval. Supplement S14 gives a six-step procedure, a constructed ledger, a complete input register and sensitivity thresholds. The ledger illustrates the calculation; it is not company evidence. Costs, the coefficient basis and unsupported assumptions remain visible in the decision presented to management.

COSO's Accept, Avoid, Pursue, Reduce and Share categories describe management responses, not numerical labels automatically assigned by a model. An exposure reallocation can support Reduce; a contract can support Share. Accept remains a management decision taking account of the selected criteria and other obligations. A higher target may involve Pursue, but being mathematically feasible is insufficient to establish its strategic value. Review of the business objective or strategy remains available when the assessed options are unsuitable (COSO, 2017a, Principle 13).

Supplement S9 provides an application map for service objectives, preserving the sequence from objective selection and performance definition to risk identification, profile construction, governance criteria, triggers and portfolio review. For water and wastewater activities, coverage, continuity and restoration require their own service–resource relationships. A financial production example is not a substitute for assessing those objectives or their mandatory service conditions.

7.2. Develops Portfolio View: plans must work together

A local plan may meet its own criterion and still conflict with enterprise commitments. Consider two hypothetical units, each with output 1,300 and a €980 unit criterion. A distinct central consequence adds €200 to total expected loss. Management also specifies an entity criterion of €1,800 and 125 shared labour hours. These common constraints are illustrative inputs to a new finite calculation, not estimates from an actual organisation.

Each unit may choose five resource mixes and either no additional insurance, equipment insurance or labour insurance. All 225 pairs are evaluated with the same fixed coefficients and prices as the unit examples. Table 6 shows how the chosen pair changes as management incorporates the common conditions. The unit-only choices exceed entity appetite. Adding the entity criterion changes the insurance choice, but still requires

130 labour hours. The joint check selects a different pair that meets the local, entity and shared-resource conditions.

Table 6. Portfolio View changes the selected pair. Unit criteria: €980; entity criterion: €1,800; shared labour: 125 hours. All monetary values are euros per period.

Conditions in selection	Expenditure	Labour hours	Entity loss
Unit limits only	5,200.00	130.00	2,150.00
Unit and entity limits	5,980.00	130.00	1,760.00
Unit, entity and shared resources	6,338.65	113.15	1,727.69

The entity loss includes the distinct central consequence of €200. The final row alone meets every condition. This is a finite 225-pair comparison, not a continuous portfolio optimum.

The final pair uses equipment insurance at one unit and labour insurance with more equipment at the other. Their expected losses are €585.00 and €942.69; adding the distinct central €200 gives €1,727.69. Shared labour use is 113.15 hours. The extra expenditure is the cost of meeting the specified joint requirements within this finite choice set, not a measured return on investment in ERM.

This illustrates COSO's Portfolio View: unit decisions are reviewed against the enterprise's objectives, risk appetite and capabilities (2017a, Principle 14; 2017b, pp. 26–27). Expected losses are added only once for each financial consequence. The calculation does not infer a portfolio tail measure from their sum, and the central term is not a statistical dependence parameter. The full candidate set and the selected resource quantities are in Supplement S7.

7.3. Implementation, reporting and review

Management approves the target and the plan together, records the response owner and establishes review triggers. Operating units implement the actions. Finance checks expenditure and loss accounting; the risk function challenges the assessment and its scope. Independent assurance reviews the process and reports through the appropriate governance arrangements without taking over the operating decision (COSO, 2017b, Appendix C).

Review distinguishes a movement along the documented policy from a change in technology, control effectiveness, prices, loss experience or the scope of the business objective. A material change prompts reassessment of the regime and a newly documented profile. Any new coefficient set must be supported and documented for its own operating domain. This makes the basis of each reported profile stable and reviewable, while allowing management to respond to changes in the business context.

Reporting should retain the target, the chosen resources, expected loss, expenditure, applicable criteria and the assumptions that could change the decision. Regular discussion of these elements is consistent with the emphasis on interactive use in Marc *et al* (2023) and on risk communication in Donnelly *et al* (2012). It also keeps a financial test connected to judgement rather than allowing the calculation to become an approval in its own right.

8. Discussion and conclusions

The framework provides a traceable financial appraisal of a performance commitment. A documented operating policy determines exposure, fixed coefficients translate that exposure into expected retained consequences, and explicit response costs enter the selection of the resource plan. Its methodological contribution is the connection between the policy-specific Risk Profile, the least-expenditure response decision and the evidence needed to support that decision within ERM.

Three findings organise this connection. First, the curve and the decision objective must be read together. A lower risk curve does not identify the most affordable target, and the least committed expenditure need not minimise expenditure plus expected loss. The activity-model comparison demonstrates that different mandates can choose different responses from the same feasible alternatives. Management must state the mandate and review both expenditure and retained consequences.

Second, the cost of changing a criterion depends on the available adjustments. In the two-factor model, alignment between cost and loss weights determines the scope for reallocation. Local marginal expenditure increases near the minimum feasible loss, while crossing that boundary may require a different response package. The reported discontinuity belongs to the declared package set; it is not a universal property of risk taking. Insurance pricing, programme effectiveness and binding resources can each change the preferred response in the activity example.

Third, a scenario-optimal plan and a robustly supported commitment are different objects. Reoptimising under a less favourable coefficient set can identify a feasible alternative without protecting the original allocation. Evaluating the nominal plan unchanged exposes that distinction. Evidence collection can then focus on claims that matter to feasibility or response selection, such as the programme's effectiveness. This provides a concrete basis for the management discussion of assumptions and confidence envisaged by COSO and the decision-support literature.

The model remains conditional on a declared technical set, consequence accounting and operating domain. Its financial measure is an expectation, not a complete description of severe losses, liquidity stress, non-financial consequences or unknown risks. Nominal production is deterministic in the examples: technical feasibility does not guarantee realised delivery during disruptive events. Statutory and service requirements must be represented separately. Fixed-resource commitments and nonlinear exposure functions can yield nonlinear profiles with unchanged coefficients; Supplement S5–S6 develops these optional cases without proposing a universal curve shape. Dynamic learning, endogenous appetite preferences, Risk Capacity estimation and the total value of ERM are outside the present analysis.

For strategy, the implication is to consider the target together with its priced and technically feasible response requirements. For Risk Responses, it is to distinguish the selected economic objective, the sensitivity of the recommendation and the resources needed to implement alternatives. For Portfolio View, it is to check that locally acceptable plans can proceed together. For Review and Revision, it is to compare the

selected commitment with new evidence before treating a recalculated optimum as an available remedy. The resulting decision record supports management judgement; it does not replace it.

Declarations

Data and code availability

All numerical inputs are hypothetical. The computational supplement contains the fixed coefficient sets, calculation code, complete results, figure sources and numerical checks. The study uses analytical solutions, deterministic sensitivity comparisons, linear programming and finite enumeration. The computational material includes the lexicographic solver, objective and premium comparisons, unchanged-plan stress tests, effectiveness and resource-scarcity scenarios, a constructed input ledger and automated verification. Supplement S11–S18 supplies the general formulation and full technical details. No company observations or stochastic estimation experiment are reported.

Ethics and participation

No human participants, animals, personal data or confidential company records were used.

References

Anton, S.G. and Afloarei Nucu, A.E. (2020) Enterprise risk management: A literature review and agenda for future research. Journal of Risk and Financial Management 13(11): 281. doi:10.3390/jrfm13110281.

Aven, T. (2004) On how to approach risk and uncertainty to support decision-making. Risk Management 6(4): 27–39. doi:10.1057/palgrave.rm.8240196.

Aven, T. (2010) An integrated framework for decision support on risk and uncertainty. Risk Management 12(4): 285–300. doi:10.1057/rm.2010.6.

Boyd, S. and Vandenberghe, L. (2004) Convex Optimization. Cambridge: Cambridge University Press. doi:10.1017/CBO9780511804441.

Bromiley, P., McShane, M., Nair, A. and Rustambekov, E. (2015) Enterprise risk management: Review, critique and research directions. Long Range Planning 48(4): 265–276. doi:10.1016/j.lrp.2014.07.005.

Callahan, C. and Soileau, J. (2017) Does enterprise risk management enhance operating performance? Advances in Accounting 37: 122–139. doi:10.1016/j.adiac.2017.01.001.

COSO (2017a) Enterprise Risk Management—Integrating with Strategy and Performance. Volume I. Committee of Sponsoring Organizations of the Treadway Commission.

COSO (2017b) Enterprise Risk Management—Integrating with Strategy and Performance: Appendices. Volume II. Committee of Sponsoring Organizations of the Treadway Commission.

Crawford, J. and Jabbour, M. (2024) The relationship between enterprise risk management and managerial judgement in decision-making: A systematic literature review. International Journal of Management Reviews 26(1): 110–136. doi:10.1111/ijmr.12337.

Dantzig, G.B. (1963) Linear Programming and Extensions. Santa Monica, CA: RAND Corporation, R-366-PR. doi:10.7249/R366.

Donnelly, R., Clement, J., Le Heron, R. and St George, J. (2012) Redesigning risk frameworks and registers to support the assessment and communication of risk in the corporate context: Lessons from a corporate risk manager in action. Risk Management 14(3): 222–247. doi:10.1057/rm.2012.3.

Gleißner, W. and Berger, T.B. (2024) Enterprise risk management: Improving embedded risk management and risk governance. Risks 12(12): 196. doi:10.3390/risks12120196.

Gordon, L.A., Loeb, M.P. and Tseng, C.-Y. (2009) Enterprise risk management and firm performance: A contingency perspective. Journal of Accounting and Public Policy 28(4): 301–327. doi:10.1016/j.jaccpubpol.2009.06.006.

Marc, M., Arena, M. and Peljhan, D. (2023) The role of interactive style of use in improving risk management effectiveness. Risk Management 25: 9. doi:10.1057/s41283-023-00114-4.

Martens, F. and Rittenberg, L. (2020) Risk Appetite—Critical to Success. COSO.

NIST/SEMATECH (n.d.) e-Handbook of Statistical Methods. Section 8.4.5.1, Constant repair rate (HPP/exponential) model. National Institute of Standards and Technology. https://www.itl.nist.gov/div898/handbook/apr/section4/apr451.htm (accessed 4 October 2026).

Nocco, B.W. and Stulz, R.M. (2006) Enterprise risk management: Theory and practice. Journal of Applied Corporate Finance 18(4): 8–20. doi:10.1111/j.1745-6622.2006.00106.x.

Varian, H.R. (2010) Intermediate Microeconomics: A Modern Approach. 8th edn. New York: W. W. Norton.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ERMSupplementForsubmission.docx