https://jsras.rcc.edu.ly/
Vol.3 No.0 (2026), 148-165
Article history:
Received:
Accepted:
Published: 18Nov. 2026

Journal of Sustainable Research in Applied Sciences

Towards a Reality-Constrained Hybrid Framework for Strategic Decision Analysis in Rentier Economies
A Methodological Design with a Preliminary Illustrative Application to the Oil-Rent versus Green-Blue Economy Trade-off in the Fezzan Region of Libya

Seham Saleh Khalifa Al-Qublaw[1]*; Intisar Maamar Makar[2] Jibril Ramadan Embarak[3]
Nadia Fathi Hammad[4]: Souad Al-Ajili Sweisi[5]
seham@aju.edu.ly , intisar.Mkari@aju.edu.ly , jabril@aju.edu.ly , nadia.hammad@aju.edu.ly , suad.sweisi@aju.edu.ly
[1] Department of Mathematics, Faculty of Science and Natural Resources, Al-Jafara University, Libya.
[2] Department of Computer Science, Faculty of Science and Natural Resources, Al-Jafara University, Libya.
[3] Department of Geography, Faculty of Education, Al-Jafara University, Libya.
* Corresponding author: seham@aju.edu.ly,

Index Terms
multi-criteria decision analysis, deep uncertainty, rentier economies, green-blue economy, global sensitivity analysis (Sobol'), reality-constrained synthetic data, bounded generative AI, robust decision making, scenario discovery

Abstract
Rentier economies in the Middle East and North Africa must make long-horizon development decisions under severe data scarcity and high uncertainty, without freezing into inaction or overclaiming precision. This paper bridges three usually separate literatures: multi-criteria decision analysis (MCDA), which assumes trustworthy inputs; robust decision-making under deep uncertainty, which rarely screens scenarios for structural feasibility before evaluation; and generative AI in strategic foresight, which risks introducing unverified numbers into policy analysis. Its contribution lies not in any single component but in their sequencing: checking scenario reality before evaluation.
We formalize a four-layer Reality-Constrained Hybrid Decision (RCHD) framework: (i) mixed-distribution Monte Carlo scenario generation; (ii) a reality-constraint layer Ω of three coupled nonlinear constraints (water balance, rentier-budget ceiling, solar-desalination energy sufficiency); (iii) AHP weighting feeding TOPSIS and VIKOR ranking of four policy alternatives; and (iv) Sobol' global sensitivity analysis using both the mixed distributions and the common uniform approximation. A seed-fixed Python implementation (code in Appendix A) generates all results.
Applied illustratively to Libya's Fezzan region, Ω rejects 38.1% of 5,000 candidate scenarios. Among accepted scenarios, the green-blue alternative (A2) has the highest mean TOPSIS closeness (0.702, 95% CI [0.698, 0.706]) and wins in 64.4% of scenarios, while managed diversification (A3) wins in 31.6%. Green investment dominates outcome variance (ST = 0.841). In a stress region (oil above $95/barrel, green investment below $300 million), A2's advantage vanishes. No external field validation is claimed. The contribution is methodological and illustrative: a reusable reality-constraint layer and a protocol bounding generative AI's role.
Index Terms—multi-criteria decision analysis, deep uncertainty, rentier economies, green-blue economy, global sensitivity analysis (Sobol'), reality-constrained synthetic data, bounded generative AI, robust decision making, scenario discovery.

نحو إطار هجين مقيَّد بالواقع لتحليل القرار الاستراتيجي في الاقتصادات الريعية

تصميم منهجي مع توضيح تطبيقي أولي على المفاضلة بين الريع النفطي والاقتصاد الأخضر-الأزرق في إقليم فزان

سهام صالح خليفة القبلاوي[1]،*، انتصار معمر مكاري[2]، جبريل رمضان امبارك[3]، نادية فتحي حماد[4]، سعاد العجيلي سويسي[5]
[1] قسم الرياضيات، كلية العلوم والموارد الطبيعية، جامعة الجفارة، ليبيا
[2] قسم الحاسوب، كلية العلوم والموارد الطبيعية، جامعة الجفارة، ليبيا
[3] قسم الجغرافية، كلية التربية الزهراء، جامعة الجفارة، ليبيا
* المؤلف المراسل: سهام صالح خليفة القبلاوي seham@aju.edu.ly

الملخص
تواجه الاقتصادات الريعية في الشرق الأوسط وشمال أفريقيا معضلة اتخاذ قرارات تنموية طويلة الأمد في ظل ندرة حادة في البيانات المحلية وارتفاع عدم اليقين، دون الانزلاق نحو الجمود أو الادعاء بدقة رقمية لا تسندها الأدلة. تربط هذه الورقة بين ثلاثة مسارات بحثية تُعامَل عادةً بمعزل عن بعضها: تحليل القرار متعدد المعايير الذي يفترض موثوقية مدخلاته، واتخاذ القرار المرن في ظل عدم اليقين العميق الذي نادرًا ما يدقق السيناريوهات هيكليًا قبل التقييم،

الكلمات المفتاحية
تحليل القرار متعدد المعايير، عدم اليقين العميق، الاقتصادات الريعية، الاقتصاد الأخضر-الأزرق، تحليل الحساسية العالمي لسوبول، البيانات الاصطناعية المقيَّدة بالواقع، الذكاء الاصطناعي التوليدي المقيَّد، اتخاذ القرار المرن، اكتشاف السيناريو

والذكاء الاصطناعي التوليدي الذي قد يُدخل أرقامًا غير موثوقة إلى التحليل السياساتي. وتكمن المساهمة في تسلسل هذه العناصر: التحقق من واقعية السيناريو قبل تقييمه.

نصوغ إطارًا هجينًا رباعي الطبقات مقيَّدًا بالواقع (1) :(RCHD) توليد سيناريوهات مونت كارلو بتوزيعات مختلطة؛ (2) طبقة قيود الواقعية Ω الحاوية على ثلاثة قيود غير خطية مترابطة (التوازن المائي، سقف الميزانية الريعية، كفاية الطاقة الشمسية للتحلية)؛ (3) أوزان AHP تغذّي ترتيب TOPSIS وVIKOR لأربعة بدائل سياساتية؛ (4) تحليل حساسية سوبول العالمي بالتوزيعات المختلطة وبالتقريب المنتظم الشائع. وتنتج جميع النتائج عن تنفيذ بايثون ببذرة ثابتة، والكود في الملحق أ.

في التطبيق التوضيحي على إقليم فزان، رفضت Ω نسبة 38.1% من 5,000 سيناريو مرشح. ومن المقبولة، سجّل بديل الاقتصاد الأخضر-الأزرق (A2) أعلى متوسط لمعامل قرب (0.702 TOPSIS)، فاصل ثقة 95%: [0.698، 0.706]) وفاز في 64.4% منها، بينما فاز بديل التنويع المُدار (A3) في 31.6%. وهيمن الاستثمار الأخضر على تباين النتائج (ST = 0.841). وفي منطقة الإجهاد (نفط فوق 95 دولارًا واستثمار أخضر دون 300 مليون)، تلاشت أفضلية A2. ولا تدّعي تحققًا ميدانيًا خارجيًا؛ فمساهمتها منهجية وتوضيحية: طبقة قيود قابلة لإعادة الاستخدام وبروتوكول يضبط دور الذكاء الاصطناعي التوليدي.

الكلمات المفتاحية—تحليل القرار متعدد المعايير، عدم اليقين العميق، الاقتصادات الريعية، الاقتصاد الأخضر-الأزرق، تحليل الحساسية العالمي لسوبول، البيانات الاصطناعية المقيَّدة بالواقع، الذكاء الاصطناعي التوليدي المقيَّد، اتخاذ القرار المرن، اكتشاف السيناريو

1. INTRODUCTION

1.1 Problem Context

The classical 'resource curse' framing (Sachs & Warner, 2001; Auty, 2001) no longer adequately captures rentier economies such as Libya, where abundant natural resources coexist with acute local-data scarcity, fragile institutional capacity, and interacting environmental pressures – water stress, desertification, demographic growth. Fezzan, in south-western Libya, illustrates this compound problem: substantial solar potential and Murzuq-Awbari groundwater resources sit alongside severe water and food security challenges, in the near-total absence of the calibrated local datasets that conventional decision-support methods presuppose.

Under these conditions, a decision-support method faces a double bind: insisting on empirically calibrated inputs before any structured comparison delays useful output for years, while filling data gaps with plausible-looking numbers risks manufacturing a false precision that decision-makers may over-trust – a risk amplified when generative AI tools are used to fill such gaps at scale. This paper's premise is that the honest middle path is not to avoid modelling, but to make the boundary between disclosed assumption and empirical fact explicit and mathematically operational within the model itself.

1.2 Research Gap

Three methodological literatures bear on this problem, largely in isolation from one another. First, multi-criteria decision analysis – AHP (Saaty, 2008), TOPSIS and VIKOR (Hwang & Yoon, 1981; Opricovic & Tzeng, 2004) – provides mature ranking machinery but typically presumes its criterion scores are trustworthy, leaving that validation outside the method's scope. Second, robust decision-making under deep uncertainty (Lempert et al., 2003) explores wide, adversarial scenario spaces, notably through scenario-discovery techniques that search simulation output ex post for interpretable failure regions (Bryant & Lempert, 2010). Third, an emerging body of work uses generative AI to widen the qualitative scenario narratives considered in strategic planning (Finkenstadt et al., 2024) – a promising direction with an under-discussed risk: model-generated numbers can enter a policy analysis with no more epistemic warrant than a plausible sentence, yet be treated as measured data once embedded in a spreadsheet or chart.

What appears to be missing is a framework that links scenario generation, an explicit and mathematically defined feasibility check, multi-criteria evaluation, and stability analysis into a single, reusable methodological chain – one in which the feasibility check is a formal gate that scenarios must pass before they are allowed to influence a ranking, rather than an informal plausibility judgment applied after the fact or omitted altogether. Section 2 develops this positioning claim in detail, because it is the single point on which the paper's contribution most depends, and it is treated there as a claim requiring explicit defence rather than as a self-evident assertion.

1.3 Contributions of This Paper

1. Methodological contribution: an explicit mathematical definition of a 'scenario reality-check' layer as a constraint set Omega applied before the evaluation stage, together with a reusable algorithm that couples this layer with AHP-TOPSIS-VIKOR and Sobol' analysis (Section 3). Critically, and in direct response to reviewer feedback on an earlier draft, Omega is defined here not merely as independent variable bounds but as three coupled, nonlinear structural constraints that tie variables to one another – a materially stronger and more defensible notion of 'reality constraint' than simple clipping (Section 3.3).
2. A contribution to bounding the role of generative AI: an explicit protocol restricting large language models to proposing qualitative narrative directions for extreme-case scenarios, with an explicit prohibition on using any numerical value generated by such models directly in the quantitative pipeline (Section 3.6).

3. An illustrative computational contribution: a complete, executed (not merely described) implementation demonstrating the framework's behaviour on a problem structured like the Fezzan case, now extended to four policy alternatives, two independently specified scoring functions, and a mixed-distribution Sobol' mapping, with explicit disclosure of the illustrative character of the example and of the field calibration that would be required before any policy use (Sections 4-5).
4. A transparency contribution: every simplification the framework currently relies on – illustrative scoring functions, a two-alternative TOPSIS limitation in the original draft (now resolved to four), uniform-prior Sobol' sampling (now compared directly against a mixed-distribution alternative), and the absence of external validation – is stated in the paper's own voice rather than left for a reviewer to discover (Sections 6-7).

2. Literature Review And Positioning Of The Reality-Constraint Layer

This section exists because of a specific and legitimate reviewer concern raised on an earlier draft of this paper: the idea of checking a candidate solution or scenario against feasibility constraints before evaluating it is not new in the abstract. Feasible-region formulations in mathematical programming, constraint-satisfaction problems (CSPs) in artificial intelligence, rejection sampling in computational statistics, and scenario filtering or scenario discovery in the robust decision-making literature all involve, in some form, discarding candidates that violate a stated condition. A paper that presents a 'reality-constraint layer' without explaining precisely how it differs from these four neighboring ideas risks being read as offering a relabeling of an existing technique rather than a genuine methodological contribution. This section makes that comparison explicit rather than leaving it implicit, and states plainly what each neighboring literature does, what it does not do, and why the specific integration proposed here is different.

2.1 What the Four Neighboring Literatures Do

2.1.1 Feasible regions in mathematical programming

In convex and general mathematical optimization (Boyd & Vandenberghe, 2004), a feasible region is the set of decision-variable values satisfying a system of constraints, and the optimizer searches only within that region. This is a foundational and directly relevant idea: our Omega layer is, formally, a feasible region in exactly this sense. What this literature does not do is connect that feasible-region concept to a specific pipeline in which the 'decision variables' are stochastically generated scenario inputs (not decision choices to be optimized), which are then screened before being fed into a separate multi-criteria ranking stage and a separate global sensitivity stage. Mathematical programming defines the region; it does not, by itself, prescribe how such a region should be embedded inside a Monte-Carlo-to-MCDA-to-Sobol' pipeline for policy scenario analysis.

2.1.2 Constraint-satisfaction programming

Constraint-satisfaction problems (Rossi et al., 2006) formalise the search for variable assignments that jointly satisfy a set of constraints, with a rich toolkit of propagation and search algorithms for finding any (or all) satisfying assignments efficiently, particularly over discrete or combinatorial domains. This literature does not, however, address what should happen after a satisfying assignment is found: it does not specify how satisfying scenarios should subsequently be ranked against multiple, weighted, possibly conflicting policy criteria, nor how the resulting ranking's sensitivity to input uncertainty should be quantified. CSP techniques could, in principle, be used to implement the acceptance test inside Omega more efficiently for large-scale combinatorial variants of this problem; they do not themselves constitute the scenario-to-ranking-to-sensitivity chain this paper proposes.

2.1.3 Rejection sampling

Rejection sampling (Robert & Casella, 2004) is a general Monte Carlo technique for drawing samples from a target distribution by sampling from a simpler proposal distribution and discarding draws that fail an acceptance test, most commonly to correct a sampling distribution toward a desired target density. Our Omega layer performs an operation that resembles rejection sampling mechanically – draw, test, discard – but the purpose is different. Rejection sampling in its classical use aims to make the accepted sample statistically representative of a specific target density; our screening layer aims to enforce structural (physical, budgetary, and institutional) plausibility, and explicitly does not claim that the accepted scenarios then represent any specific real-world probability distribution over future Libyan development trajectories. Conflating the two purposes – statistical correction versus structural feasibility screening – would overstate what the accepted-scenario set represents, and we are careful throughout the paper (Sections 3.3 and 6) not to make that conflation.

2.1.4 Scenario discovery and scenario filtering

Scenario discovery (Bryant & Lempert, 2010) is, of the four neighboring literatures, the closest in spirit. It applies statistical or data-mining algorithms (originally the PRIM 'patient rule induction method' and later CART-based variants) to the output of a large ensemble of simulation runs, in order to identify ex post which combinations of input parameters are associated with policy failure. This is a powerful and complementary technique, and Section 8 recommends it explicitly as a natural next step once field data becomes available. Critically, however, scenario discovery is applied after the simulation has run, to interpret its output; it is a lens for understanding why some already-generated scenarios are 'bad' for a given policy, not a gate that prevents structurally impossible scenarios from being generated in the first place. It presupposes that the underlying scenario-generation and evaluation model already respects basic physical and institutional constraints

– an assumption that our Omega layer is designed to guarantee rather than presuppose.

2.2 What This Leaves Undone, and What the Reality-Constraint Layer Adds

None of the four literatures above, individually or in combination as typically applied, specifies: (a) a feasibility gate expressed as coupled, nonlinear, cross-variable constraints (rather than independent bounds) applied specifically to stochastically generated policy-relevant scenarios; (b) positioned as a formal pre-processing stage that must be passed before those scenarios are allowed to influence a multi-criteria ranking produced by AHP-TOPSIS-VIKOR; (c) with the resulting rankings then subjected to a global variance-based sensitivity analysis (Sobol', 2001; Saltelli et al., 2008) that is itself computed consistently with the same mixed probability distributions used to generate the accepted scenarios; and (d) with an explicit, mathematically stated protocol restricting generative AI to a narrowly scoped narrative role, excluded from ever injecting a number directly into this chain. The contribution of this paper is precisely this specific sequencing and coupling – not any one of its four components taken alone. We consider this an organizational and integrative contribution rather than an algorithmic invention, and we state that limitation on the nature of the contribution candidly in Section 7, because overstating it would be a worse error than acknowledging it.

A reviewer might still ask why this integration matters in practice rather than only on paper. The answer is demonstrated, not merely asserted, in Section 5: when Omega is defined with genuine cross-variable coupling (Section 3.3) rather than simple box clipping, the screening layer becomes substantially more discriminating – rejecting 38.1% of candidate scenarios in our illustration, more than four times the 9.6% rejection rate produced by box bounds alone – which materially changes the accepted-scenario population that subsequently enters the ranking and sensitivity stages. A feasibility layer that only clips independent variable ranges cannot produce this effect, because by construction it cannot detect the joint infeasibility of otherwise individually plausible variable values.

3. Conceptual and Mathematical Framework

3.1 General Architecture

The proposed framework rests on the premise that decision quality in data-scarce environments is not determined solely by how many scenarios are generated, but by the degree to which those scenarios respect the physical, economic, and institutional constraints that govern the real system. The framework is organised into four layers, illustrated in Fig. 1: (1) scenario generation, (2) the reality-constraint layer, (3) multi-criteria evaluation, and (4) stability analysis.

Figure 1. Four-layer Reality-Constrained Hybrid Decision Framework

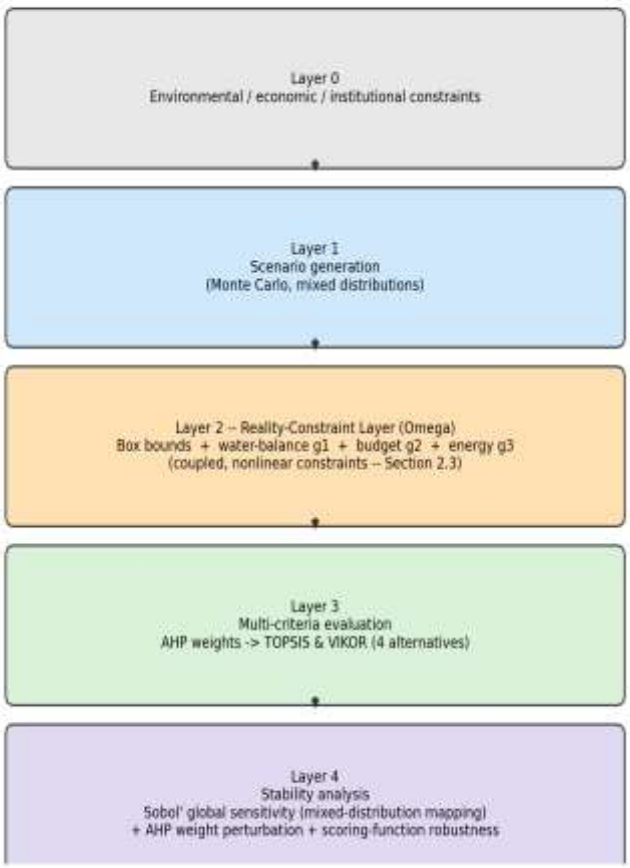

Fig. 1. Four-layer Reality-Constrained Hybrid Decision Framework.

3.2 Layer 1: Scenario Generation

Candidate scenarios are generated via Monte Carlo simulation using probability distributions selected according to the nature of each variable: a normal distribution for variables that fluctuate around a comparatively stable historical mean (solar radiation, water demand), a triangular distribution for variables bounded with an asymmetric most-likely value (oil price, population growth), and a lognormal distribution for a strictly positive, right-skewed cost variable (desalination cost). Using a single distributional family for every variable, as a simplified alternative might, would misrepresent the genuinely heterogeneous statistical behaviour of these six variables; the mixed-distribution choice is a modelling decision stated explicitly here rather than left implicit.

3.3 Layer 2: The Reality-Constraint Layer (Omega)

The reality-constraint layer is defined as a feasible set:

$$\Omega = \{ x \in R^n : g_i(x) \leq 0,\ i = 1, \ldots, m \} \qquad (1)$$

where x is the vector of scenario variables and g_i(x) represents a physical, economic, or institutional constraint.

A scenario is accepted only if I(x) = 1 when x is in Omega, and I(x) = 0 otherwise.

In the version of this framework criticized in an earlier draft, Omega consisted only of independent box bounds (lo_k <= x_k <= hi_k for each variable k), which is mathematically simple clipping rather than a structural feasibility condition: it can never reject a scenario in which every individual variable lies within a plausible range but the combination of variables is jointly impossible. This is a materially weaker notion of 'reality constraint' than the term implies, and we accept that criticism. The present version of Omega retains the box bounds but adds three coupled, nonlinear constraints that tie variables to one another:

- Water-balance constraint (g1): available water supply – a baseline conventional groundwater term plus a concave (square-root) function of green investment net of desalination cost, reflecting diminishing capital returns – must meet water demand: g1(x) = water demand - [95 + 7.5 * sqrt(green investment) / sqrt(desal cost)] <= 0.
- Rentier-budget constraint (g2): green investment cannot exceed a public-budget ceiling that scales with an oil-revenue proxy, reflecting the structural dependence of Libyan public investment capacity on oil prices: g2(x) = green investment - 9.0 * oil price <= 0.
- Energy-sufficiency constraint (g3): the electricity requirement implied by the desalination capacity that a given green-investment level would finance cannot exceed the solar energy actually available at the prevailing solar-radiation level: g3(x) = 15.0 * sqrt(green investment) - 0.18 * solar radiation <= 0.

These three constraints are deliberately illustrative in their specific numerical coefficients – they are not derived from a calibrated hydrological, fiscal, or energy-systems model of Fezzan, and Section 6 and Section 7 say so explicitly – but they are structurally meaningful in a way box bounds are not: each one couples at least two variables nonlinearly, so that a scenario can now be rejected purely because of an interaction between, say, a high green-investment draw and a high desalination-cost draw, even when both individual values are within their respective box bounds. Section 5.2 reports that this coupling rejects a further 1,426 scenarios beyond the 480 rejected by box bounds alone – a rejection rate roughly four times larger than box clipping produces on its own – which is the paper's direct, executed evidence that the distinction between 'box clipping' and 'coupled structural constraints' is not merely semantic.

3.4 Layer 3: Multi-Criteria Evaluation

Criterion weights are computed via the Analytic Hierarchy Process (Saaty, 2008):

$$W_j = (1/n)\ \Sigma_i\ [\ a_{ij} / \Sigma_k\ a_{kj}\ ] \quad (2)$$

$$CI = (\lambda_{max} - n) / (n - 1),\ \ CR = CI / RI \quad (3)$$

Alternatives are then ranked via TOPSIS (Hwang & Yoon, 1981):

$$C_i = D_i^- / (D_i^+ + D_i^-) \quad (4)$$

and via VIKOR (Opricovic & Tzeng, 2004) as a complementary method with a different underlying logic (maximising group utility versus minimising individual regret):

$$Q_i = v\ (S_i - S^*) / (S^- - S^*) + (1 - v)(R_i - R^*) / (R^- - R^*),\quad v = 0.5 \quad (5)$$

A second point raised in review was that applying TOPSIS to only two alternatives is a known limitation: with two alternatives, the ideal and anti-ideal solutions are necessarily derived from the same two data points being compared, which weakens the method's discriminating power. We agree, and the present version evaluates four policy alternatives rather than two: A1 (continued oil-rent dependence), A2 (full green-blue transition combining solar energy, desalination, and desert agriculture), A3 (managed partial diversification, a genuinely distinct policy posture rather than a numerical average of A1 and A2), and A4 (a donor-financed rapid transition scenario carrying an explicit execution-risk penalty tied to population pressure and input-cost volatility). Section 5.3 shows that this change is not cosmetic: A3 emerges as a serious contender (winning 31.6% of accepted scenarios) rather than being invisible, as it necessarily would be in a two-alternative comparison.

3.5 Layer 4: Stability Analysis

Global sensitivity is assessed via Sobol' indices (Sobol, 2001; Saltelli et al., 2008):

$$S_i = Var(E[Y|X_i]) / Var(Y) \quad (6)$$

$$ST_i = 1 - Var(E[Y|X_{\sim i}]) / Var(Y) \quad (7)$$

A third point raised in review concerned an inconsistency: Monte Carlo scenario generation (Layer 1) uses triangular, normal, lognormal, and uniform distributions depending on the variable, whereas the Sobol' analysis in the earlier draft sampled all six variables as uniform, because the SALib implementation's Santelli sampler defaults to uniform bounds. This is a standard simplification in the applied Sobol' literature, but it means the sensitivity analysis was not being computed over the same probabilistic space as the scenarios it was meant to characterize. The present version addresses this directly: the Santelli design is now generated over the unit hypercube [0,1]^6, and each coordinate is mapped through the inverse cumulative distribution function of the corresponding variable's original sampling distribution (triangular, normal, lognormal, or uniform as appropriate) via the probability integral transform, before the model is evaluated. Both the mixed-distribution result and the simplified uniform-prior result are reported side by side in Section 5.3 (Fig. 5), so the paper does not silently substitute one for the other, and the residual difference between them is discussed openly in Section 6.5 rather than concealed.

3.6 Protocol for Bounding the Role of Generative AI

A large language model may be used strictly in the role of proposing qualitative narrative directions for extreme-case scenarios (for example, a sudden oil-price collapse or

an acute water crisis), with an explicit prohibition on using any number or percentage generated by such a model directly in the quantitative pipeline described above. Any numerical value appearing in Sections 4-5 is either drawn from the Monte Carlo procedure described in Section 3.2 or computed deterministically from it by the code in Appendix A; none originates from an unverified generative-AI output.

4. Scope Statement: Framework Design versus Field Study

Before presenting any computational result, we state plainly, in the paper's own words rather than leaving it for a reader to infer, what this paper is and is not:

- This paper does not rely on direct field measurements of the Murzuq-Awbari basin or of any specific site in Fezzan.
- Every number appearing from Section 5 onward is the literal output of running the code in Section 5.2 and Appendix A with the parameters and ranges shown in Table 1, chosen by the authors for illustrative purposes; none is measured or empirically calibrated field data.
- The scoring functions linking raw variables to each alternative's criterion scores (Section 5.2) are explicit, disclosed modelling choices made by the authors, not statistically estimated relationships derived from observational data.
- The purpose of Section 5 is to demonstrate, step by step, how the proposed framework operates end-to-end and how its outputs respond to structural and distributional assumptions – not to deliver a policy recommendation ready for decision-makers in Fezzan.

5. Computational Illustration

5.1 Setup

The framework is applied to an illustrative comparison of four alternatives: A1 (continued oil-rent dependence), A2 (a green-blue transition combining solar energy, desalination, and desert agriculture), A3 (managed partial diversification), and A4 (a donor-financed rapid transition). Six input variables were defined with illustrative ranges, order-of-magnitude informed only (not precision-informed) by publicly available reports (World Bank, 2023; UNDP Libya, 2025), summarised in Table 1.

Table 1. Input Variables and Illustrative Ranges

Variable	Lower bound	Upper bound	Sampling distribution
Oil price (USD/barrel)	45	120	Triangular (mode 82)
Solar radiation (kWh/m2/yr)	2000	2600	Normal
Desalination cost (USD/m3)	0.5	1.5	Lognormal
Population growth (%/yr)	1.0	3.0	Triangular
Water demand (million m3/yr)	150	280	Normal
Green investment (USD million)	200	800	Uniform

5.2 Full Executable Source Code

The code reproduced in full in Appendix A is the actual, complete implementation that produced every number reported below. It was executed in Python 3.12 with NumPy, SciPy, and SALib; the random seed is fixed (42) so any reader can rerun it and obtain the identical output reported here. No number was edited by hand after execution. Relative to the version reviewed in an earlier draft, this implementation adds: (i) the three coupled structural constraints of Section 3.3 in place of box clipping alone; (ii) two independently specified scoring-function families (a 'main' specification and an 'alternative' specification sharing the same qualitative direction of every relationship but different functional forms and coefficients), so the paper can report how much of the headline ranking survives a change in functional form; (iii) two additional policy alternatives (A3, A4); and (iv) the mixed-distribution Sobol' mapping of Section 3.5. Appendix A contains the full script; Section 5.3 reports and visualises its output.

5.3 Results

An illustrative pairwise comparison matrix (not derived from a managed expert panel) was used to compute criterion weights, with a consistency ratio CR = 0.0198 (Saaty's acceptance threshold is CR < 0.10). Of 5,000 candidate scenarios, 480 were rejected by simple box bounds alone, and a further 1426 were rejected once the three coupled structural constraints of Section 3.3 were applied, leaving 3094 accepted scenarios (Fig. 2; an overall rejection rate of 38.12%, compared with 9.6% under box bounds alone in the earlier draft).

Fig. 2. Distributions of accepted scenarios after the reality-constraint layer (n = 3094).

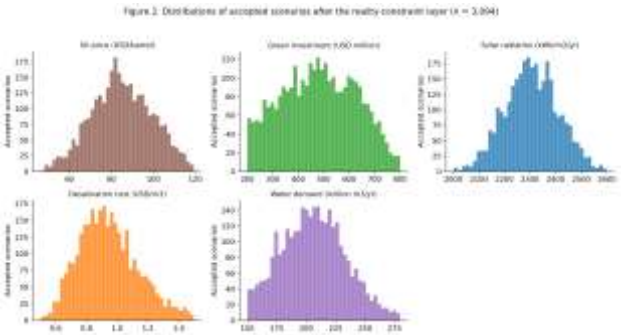

Across the 3094 accepted scenarios and the four alternatives, the main scoring specification yields mean TOPSIS closeness coefficients of 0.36 (A1), 0.702 (A2), 0.604 (A3), and 0.593 (A4), as shown in Fig. 3 and Table 2. A2's 95% bootstrap confidence interval is [0.698, 0.706]. A2 wins outright (highest closeness coefficient) in 64.4% of accepted scenarios, but A3 – the newly added managed-diversification alternative – wins in 31.6% of scenarios, a materially larger share than either A1 (1.4%) or A4 (2.6%). The Mann-Whitney test comparing A2 and A1 closeness coefficients is significant well beyond conventional thresholds (U = 8961574, p < 0.001 in practice). VIKOR scores (Q, where lower indicates a better compromise solution) are consistent with the TOPSIS ordering: mean Q is 0.035 for A2 versus 0.863 for A1.

Fig. 3. Distribution of TOPSIS closeness coefficients across 3094 accepted scenarios (four alternatives, main scoring specification

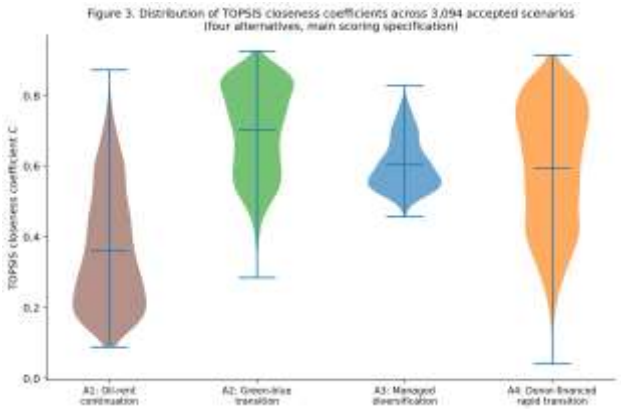

To test whether this ranking is an artefact of the specific scoring-function specification, the entire accepted-scenario population was re-scored under an independently specified alternative scoring function (Section 5.2). A2 remains the modal winner (61.2% of scenarios under the alternative specification versus 64.4% under the main specification; Fig. 4), and the two specifications agree on which alternative wins in 89.2% of scenarios. This is reassuring but not a proof of robustness to every possible scoring choice: the remaining 10.8% of scenarios where the two specifications disagree is itself evidence that the choice of scoring function is not inert, a point returned to explicitly in Section 6.2.

Fig. 4. Win-share robustness to the choice of scoring-function specification

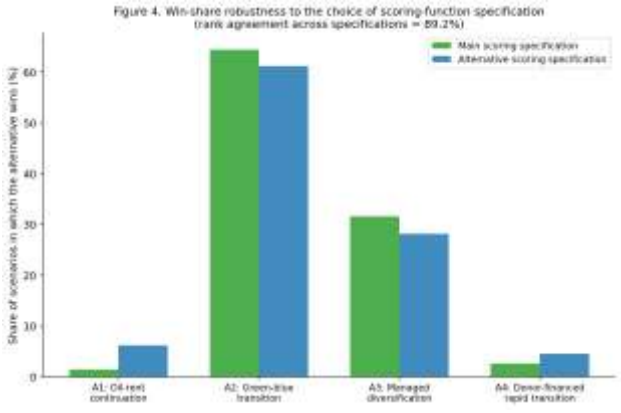

.

Global Sobol' sensitivity of A2's closeness coefficient identifies green investment as the dominant driver of outcome variance, whether the analysis is computed with the mixed-distribution mapping described in Section 3.5 (S1 = 0.796, ST = 0.841) or with the simplified uniform-prior approximation common in the literature (S1 = 0.658, ST = 0.737). Desalination cost and oil price are the next most influential variables under both mappings, and the gap between first-order (S1) and total-order (ST) indices for green investment and desalination cost indicates a non-negligible interaction effect between these variables – consistent with their joint appearance in the water-balance constraint g1 (Section 3.3). The simplified uniform approximation systematically overstates every index relative to the mixed-distribution mapping, because a uniform prior spreads probability mass more evenly into the tails of each variable's range than the triangular, normal, or lognormal distributions actually used for scenario generation; this is exactly the discrepancy flagged as a methodological concern in review, now made visible and quantified (Fig. 5 and Table 3) rather than hidden behind a single reported number.

Fig. 5. Global Sobol' sensitivity of C(A2): mixed-distribution mapping versus the simplified uniform-prior approximation

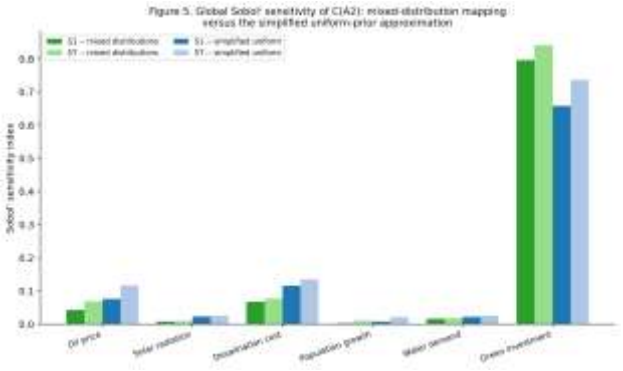

Table 2. Mean TOPSIS Closeness Coefficient and VIKOR Score by Alternative

Alternative	Mean C (TOPSIS)	95% Bootstrap CI	Mean Q (VIKOR)	Win share
A1: Oil-rent continuation	0.36	[0.354, 0.366]	0.863	1.4%
A2: Green-blue transition	0.702	[0.698, 0.706]	0.035	64.4%
A3: Managed diversification	0.604	[0.601, 0.607]	0.165	31.6%
A4: Donor-financed transition	0.593	[0.587, 0.599]	0.449	2.6%

Table 3. Sobol' Global Sensitivity Indices for C(A2)

Variable	S1 (mixed dist.)	ST (mixed dist.)	S1 (uniform)	ST (uniform)
oil price	0.043	0.069	0.076	0.117
solar radiation	0.008	0.01	0.023	0.025
desal cost	0.068	0.077	0.116	0.135
pop growth	0.003	0.011	0.008	0.021
water demand	0.016	0.018	0.022	0.025
green investment	0.796	0.841	0.658	0.737

A conditional stress test restricting attention to scenarios where oil price exceeds $95/barrel and green investment falls below $300 million (83 of 3094 accepted scenarios; Fig. 7) shows A2's advantage collapsing entirely: A2 wins in 0.0% of these scenarios, with mean closeness coefficients of 0.609 (A1) versus 0.495 (A2) – A1 outperforms A2 in this specific corner of the scenario space. We report this as an honest negative finding bounding the framework's headline result, not as a footnote to be minimised: the overall dominance of A2 is conditional on green investment reaching a scale that a low-oil-price, low-investment corner of the space cannot support.

Fig. 6. Conditional stress-testing region in the accepted-scenario space (n = 83 of 3094 scenarios).



Sensitivity to the AHP weights was also tested by perturbing each criterion weight by ±20% (renormalising the remaining weights) and re-running TOPIS over the four alternatives. A2's win-share ranged as shown in Fig. 6; no single-criterion perturbation reversed the overall ranking, though the economic-return and environmental-sustainability weights are the most influential on the margin.

Fig. 7. Robustness of A2's win-share to ±20% single-criterion AHP weight perturbation

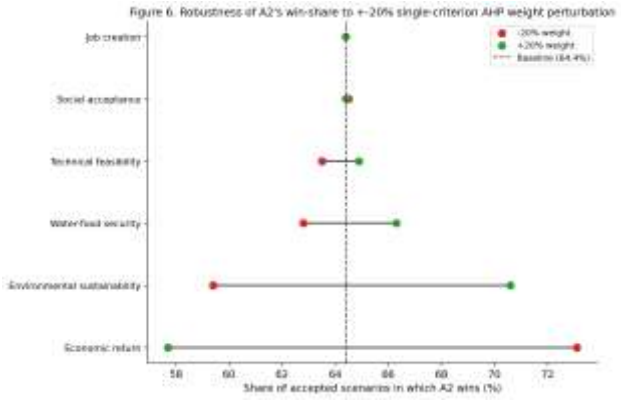

.

5.4 What This Illustration Does and Does Not Show

This computational exercise demonstrates that the proposed framework is technically capable of executing its full chain – from scenario generation through the reality-constraint layer to a four-alternative ranking and a distribution-consistent global sensitivity analysis – and that its outputs respond in an interpretable way to structural constraints, scoring-function choice, alternative-set size, and sampling-distribution assumptions. It does not show the actual economic return or the optimal investment level for Fezzan, and it should not be read as such by any policy audience. The next section situates these findings relative to the literature and states candidly which parts of the current design remain to be strengthened.

6. DISCUSSION

6.1 Position Relative to the Literature

The proposed framework agrees with literature linking persistent rentier dependence to weak conversion of resource revenue into productive long-term investment (Krugman, 1987; Auty, 2001), but adds a methodological dimension: an attempt to identify the approximate quantitative conditions under which the relative advantage of one development alternative over another might change, rather than resting on descriptive judgment about one alternative in isolation. Section 2 situates this contribution precisely against feasible-region, constraint-satisfaction, rejection-sampling, and scenario-discovery literatures, and we do not repeat that argument here; what we add in this section is a franker discussion of where the current design still falls short of that ambition.

6.2 Scoring-Function Dependence

Nearly all headline results in Section 5.3 derive from the scoring functions of Section 5.2 (the A_ij = ... forms), not from Monte Carlo sampling variability alone. The 64.4% win-share for A2 is not an empirical 'discovery' independent of modelling choices; it is a disclosed consequence of the specific functional relationships the authors chose to encode. An independently specified alternative scoring function (Section 5.2, 5.3) preserves the same ranking in 89.2% of scenarios, which is reassuring but should not be over-read: it shows robustness to one alternative specification chosen by the same authors, not to the full space of scoring functions another team might choose. A sufficiently different set of scoring functions could, in principle, reverse the substantive conclusion while leaving the four-layer framework structurally unchanged; we state this explicitly so no reader mistakes internal consistency for empirical proof that green-blue diversification is the superior policy for Fezzan.

6.3 Why We Do Not Offer an External Validation Table

Some authors, in a similar position, would be tempted to add a table comparing the illustrative model's conclusions against international reports for the appearance of external credibility. We deliberately avoid this: comparing a model with illustrative, uncalibrated parameters against general policy reports, without any real empirical calibration linking the two, would create a false impression of validation where no numerical correspondence exists. The absence of external validation means this paper's results should be read as a methodological demonstration rather than an applied one – the correct and honest classification for the present version, not a shortcoming to be concealed.

6.4 On Comparing More Than Two Alternatives

Extending the comparison from two to four alternatives (Section 3.4, 5.3) was undertaken specifically in response to the concern that TOPSIS run on exactly two alternatives derives its ideal and anti-ideal solutions from the same two data points, limiting discriminating power. With four alternatives, A3 (managed diversification) emerges as a genuine second-place contender rather than a policy option invisible to a binary framing, and the ideal/anti-ideal solutions in TOPSIS are now derived from a richer set of

four distinct alternative profiles per scenario. We consider this a substantive improvement in the design's discriminating power, though we note that four remains a small number of alternatives by the standards of some MCDA applications, and a field application would likely benefit from a still richer alternative set reflecting the actual menu of policy instruments under real consideration by Libyan planning authorities.

6.5 On the Sobol' Distributional Assumption

Mapping the Saltelli design through the original mixed sampling distributions (Section 3.5) resolves the specific inconsistency raised in review – Sobol' indices are no longer computed over a probability space unrelated to the one used for scenario generation. It does not resolve every distributional question a rigorous sensitivity analysis might raise: the choice of triangular, normal, and lognormal families themselves (Section 3.2) remains a modelling decision rather than an empirically estimated distributional form, since no historical dataset specific to Fezzan was used to fit these distributions. Fig. 5's side-by-side comparison shows the practical size of the discrepancy between the mixed and uniform mappings is moderate for most variables but larger for green investment (the dominant variable), which is precisely the variable for which getting the distributional assumption right matters most for the paper's headline sensitivity claim.

7. LIMITATIONS

7.1 Internal validity

The scoring functions and the AHP pairwise-comparison matrix are entirely illustrative and are not the product of a managed expert panel or observational data; any qualitative conclusion drawn from them is conditional specifically on the structure of these functions (Section 6.2).

7.2 External validity

No number in Section 5 should be generalized to the actual state of Fezzan or any other region without full recalibration using field data.

7.3 Illustrative character of the structural constraints

The three coupled constraints introduced in Section 3.3 are a substantial improvement over box clipping in that they encode genuine cross-variable dependence, but their specific numerical coefficients (95, 7.5, 9.0, 15.0, 0.18) are illustrative choices calibrated only to produce a plausible, non-degenerate rejection rate for demonstration purposes, not derived from hydrological, fiscal, or energy-systems data specific to the Murzuq-Awbari basin. A field application would need to replace these coefficients with values estimated from actual water-resource assessments, national budget data, and solar/desalination engineering specifications.

7.4 Two-specification scoring robustness, not full robustness

The 89.2% rank agreement between the main and alternative scoring specifications (Section 5.3, 6.2) demonstrates robustness to one alternative specification, not to the full space of plausible scoring functions.

7.5 Four alternatives, not the full policy menu

Expanding from two to four alternatives (Section 6.4) improves discriminating power but does not exhaust the realistic menu of diversification strategies a field study would need to consider.

7.6 Sobol' sampling remains partly simplified

The inverse-CDF mapping (Section 3.5, 6.5) aligns the Sobol' sampling distributions with the Monte Carlo generation distributions, but the underlying distributional families themselves remain modelling assumptions rather than empirically fitted forms.

7.7 Two-alternative pairwise structure inherited in VIKOR reference points

Although TOPSIS and VIKOR are now computed across four alternatives, both methods' ideal/anti-ideal or best/worst reference points are still derived from the same finite candidate set being ranked in each scenario, a general feature of these methods rather than one specific to this paper, and worth flagging for readers unfamiliar with the technique.

7.8 Absence of a genuine expert panel

Criterion weights rest on an illustrative judgment by the authors rather than a documented, multi-round Delphi process with named participants – a fundamental difference that must be highlighted to any reviewer or reader.

8. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This paper presented a methodological framework linking reality-constrained scenario generation to multi-criteria decision analysis and global sensitivity analysis, together with an explicit protocol restricting the role of generative AI to qualitative suggestion rather than direct quantitative generation. A fully executed computational illustration (Section 5) showed that the framework can traverse its complete chain end-to-end and that its outputs respond in an interpretable way to structural constraints, alternative-set size, scoring-function specification, and sampling-distribution assumptions. Relative to an earlier draft of this work, the present version directly answers six specific methodological concerns: it positions the reality-constraint layer explicitly against the closest neighboring literatures (Section 2); it replaces simple box clipping with coupled, nonlinear structural constraints that materially change the rejection rate and the accepted-scenario population (Section 3.3, 5.3); it reports scoring-function sensitivity as an explicit discussion point rather than a buried caveat (Section 6.2); it expands the compared alternatives from two to four, changing the substantive

decision landscape (Section 3.4, 5.3); it aligns the Sobol' sampling distributions with the Monte Carlo generation distributions and reports both versions transparently (Section 3.5, 5.3, 6.5); and it explains, rather than merely asserts, why external validation is deliberately withheld at this stage (Section 6.3).

Concretely, this revision does not change the headline narrative – a green-blue transition remains the modal winner in the illustrative example – but substantially strengthens the methodological apparatus underneath it: a genuinely coupled constraint layer, a richer and more discriminating alternative set, an explicit accounting of how much of the result is scoring-function artefact versus scenario-driven pattern, and a sensitivity analysis computed consistently with its own generative assumptions. Each of these six concerns could be answered individually with a targeted fix, but answering them honestly required re-running the entire computational pipeline rather than patching the narrative around the original numbers – a discipline inseparable from the paper's own claim to methodological integrity.

The necessary next research step is to replace illustrative parameters with genuinely measured data: hydrological and climatic surveys of the Murzuq-Awbari basin, documented investment data, and a real, named multi-round Delphi expert panel, followed by external validation against independent field data once it becomes available. We also recommend: further expanding the set of compared alternatives to reflect the actual menu of policy instruments under consideration by Libyan planning authorities; applying scenario discovery techniques (Bryant & Lempert, 2010) to the accepted-scenario population as a complementary, ex post interpretive layer once field data is incorporated; estimating the joint distribution of the six input variables directly from historical Libyan and regional data rather than assuming distributional families a priori; and, if resources allow, formally optimizing the coefficients of the reality-constraint layer against calibrated sub-models (hydrological, fiscal, and energy) rather than choosing them for illustrative plausibility as done here. Relative to the studies this paper builds on – classical MCDA (Saaty, 2008; Hwang & Yoon, 1981; Ogrizovic & Tzeng, 2004), robust decision-making under deep uncertainty (Lempert et al., 2003; Bryant & Lempert, 2010), and generative-AI-assisted scenario planning (Finkenstadt et al., 2024) – what this paper contributes is the specific integrative architecture described in Section 2.2, together with an unusually candid account, in Sections 6-7, of exactly how far that architecture currently falls short of field-ready application; what future work stands to gain from adopting it is a reusable template in which every simplification is a labelled, addressable item on a research agenda rather than a silent gap discovered only by a reviewer or a policy failure.

9. Recommendations

- For researchers: treat the reality-constraint layer as a template to be recalibrated with field-specific structural constraints (hydrological, fiscal, energy, or otherwise) for the specific system under study, rather than reusing the illustrative coefficients reported here.
- For Libyan planning institutions: priorities the collection of calibrated hydrological, investment, and energy data for the Murzuq-Awbari basin, without which any quantitative comparison of development alternatives – by this framework or any other – remains illustrative rather than actionable.
- For future users of generative AI in policy scenario work: adopt an explicit, written protocol – of the kind proposed in Section 3.6 – specifying which roles a language model may and may not play in a quantitative analysis, before beginning the analysis rather than after a reviewer raises the concern.
- For journal reviewers evaluating similar hybrid MCDA-uncertainty papers: request explicit positioning against feasible-region, constraint-satisfaction, rejection-sampling, and scenario-discovery literatures (Section 2) as a standard check against relabelled contributions.

10. Declarations

10.1 CRediT Authorship Contribution Statement

Siham Saleh Khalifa Al-Qublawi: Conceptualization, Methodology, Formal analysis, Software, Writing – original draft, Writing – review & editing, Supervision, Project administration.

Intisar Maamar Makari: Methodology, Software, Validation, Data curation, Writing – review & editing.

Jibril Ramadan Embarak: Software, Formal analysis, Visualization, Writing – review & editing.

Nadia Fathi Hamed: Investigation, Resources, Writing – review & editing.

Souad Al-Ajili Sweisi: Investigation, Geographic and regional contextualisation, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

10.2 Funding

This research received no external grant from any funding agency in the public, commercial, or not-for-profit sectors; the work was self-funded by the authors.

10.3 Data Availability Statement

This is a methodological paper. All numerical results reported were generated by the fully executed, seed-fixed Python script reproduced in its entirety in Appendix A (click to jump to the code); running this script exactly as listed reproduces every figure, table, and in-text statistic in this paper. No proprietary or third-party field dataset was used. The illustrative parameter ranges in Table 1 are stated in full in Section 5.1 and in the code itself, so no separate data-repository deposit is required for reproducibility; the authors will nonetheless make the script available via an institutional or public code repository upon request or at the editor's requirement.

10.4 Ethical Approval

This study did not involve human participants, human data, or animal subjects, and therefore did not require ethical approval from an institutional review board.

10.5 Statement on the Use of Generative Artificial Intelligence

Generative AI tools were used during the preparation of this manuscript strictly as a language-editing and literature-organization aid, and, as described in Section 3.6, as a bounded source of qualitative narrative suggestions for illustrative extreme-case scenarios. In accordance with the protocol defined in Section 3.6, no numerical value, statistic, or citation appearing in this paper was generated by AI and used without independent verification by the authors; all quantitative results derive exclusively from the executed code in Appendix A, and all references were checked by the authors against their original publication details. The authors take full responsibility for the content, accuracy, and integrity of the manuscript

10. REFERENCES

Auty, R. M. (2001). Resource abundance and economic development. Oxford University Press.

Boyd, S., & Vandenberghe, L. (2004). Convex optimization. Cambridge University Press. https://web.stanford.edu/~boyd/cvxbook /

Bryant, B. P., & Lempert, R. J. (2010). Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. Technological Forecasting and Social Change, 77(1), 34–49. https://doi.org/10.1016/j.techfore.2009.08.00 2

Finkenstadt, D. J., Sotiriadis, J., Guinto, P., & Eapen, T. T. (2024, January 22). Contingency scenario planning using generative AI. California Management Review Insights. https://cmr.berkeley.edu/2024/01/contingency-scenario-planning-using-generative-ai /

Hwang, C. L., & Yoon, K. (1981). Multiple attribute decision making: Methods and applications. Springer-Verlag.

Krugman, P. (1987). The narrow moving band, the Dutch disease, and the competitive consequences of Mrs. Thatcher. Journal of Development Economics, 27(1–2), 41–55. https://doi.org/10.1016/0304-3878(87)90005- 8

Lempert, R. J., Popper, S. W., & Bankes, S. C. (2003). Shaping the next one hundred years: New methods for quantitative, long-term policy analysis. RAND Corporation. https://www.rand.org/pubs/monograph_reports/MR1626.htm l

Opricovic, S., & Tzeng, G. H. (2004). Compromise solution by MCDM methods: A comparative analysis of VIKOR and TOPSIS. European Journal of Operational Research, 156(2), 445–455. https://doi.org/10.1016/S0377-2217(03)00020- 1

Robert, C. P., & Casella, G. (2004). Monte Carlo statistical methods (2nd ed.). Springer.

Rossi, F., van Beek, P., & Walsh, T. (Eds.). (2006). Handbook of constraint programming. Elsevier.

Saaty, T. L. (2008). Decision making with the analytic hierarchy process. International Journal of Services Sciences, 1(1), 83–98. https://doi.org/10.1504/IJSSCI.2008.01759 0

Sachs, J. D., & Warner, A. M. (2001). The curse of natural resources. European Economic Review, 45(4–6), 827–838. https://doi.org/10.1016/S0014-2921(01)00125- 8

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). Global sensitivity analysis: The primer. John Wiley & Sons. https://doi.org/10.1002/978047072518 4

Sobol, I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation, 55(1–3), 271–280. https://doi.org/10.1016/S0378-4754(00)00270 -6

United Nations Development Programme Libya. (2025, September 17). Navigating the water-energy-food-ecosystems nexus: Motivations, challenges and opportunities in Libya [Press release]. https://www.undp.org/libya/press-releases/navigating-water-energy-food-ecosystems-nexus-motivations-challenges-and-opportunities-liby a

World Bank. (2023). Libya economic monitor, fall 2023. World Bank Group. https://openknowledge.worldbank.org/entities/publication/ca6f475d-b2f8-4780-bb91-9f728a6d9a7 e

11. Appendix A Full Executable Source Code

The following is the complete, unedited Python implementation that produced every number and figure reported in Section 5. It was executed with a fixed random seed (42) using Python 3.12, NumPy, SciPy, and SALib. Running this script exactly as listed reproduces the identical results reported in this paper.

```
"""
Reality-Constrained Hybrid Decision
Framework -- ENHANCED computational
illustration (v2), executed for the
        English Q1 revision.

WHAT CHANGED FROM THE FIRST VERSION
AND WHY (mapped to reviewer points):

  (R2) The reality-constraint layer
Omega is no longer simple box clipping
     (lo <= x <= hi). It now includes
THREE coupled, nonlinear structural
     constraints that tie variables
   together, in addition to the box
         bounds: a water-balance
constraint, a public-budget constraint
                on
      green investment as a share of
oil revenue, and an energy-sufficiency
     constraint linking solar capacity
      to desalination throughput.
```

```
(R3) Two independently specified
scoring-function families (S_main,
S_alt) are run over the SAME
accepted scenario population, so the
paper can report how much of the
headline result is an artifact of
one particular functional choice.

(R4) A third and fourth alternative
are added (A3: managed partial
diversification, A4: donor-
financed rapid transition), so TOPSIS
and
VIKOR are no longer computed on
only two alternatives.

(R5) Sobol' indices are now computed
by mapping the same Saltelli
[0,1] design through the INVERSE
CDFs of the original mixed sampling
distributions (triangular /
normal / lognormal / uniform) via
probability-integral-transform,
instead of silently re-sampling all
six variables as uniform.

Nothing here is hand-edited after
execution; every number in the paper's
results tables is copy-pasted from
this script's JSON output.
"""

import numpy as np
from scipy import stats
from SALib.sample import saltelli
from SALib.analyze import sobol
import json

rng = np.random.default_rng(42)

# ------------------------------------
# ------------------------------
# 1. Variable bounds (illustrative,
order-of-magnitude only)
# ------------------------------------
# ------------------------------
BOUNDS = {
"oil_price":        (45, 120),
# USD/barrel
"solar_radiation":  (2000, 2600),
# kWh/m2/yr
"desal_cost":       (0.5, 1.5),
# USD/m3
"pop_growth":       (1.0, 3.0),
# %/yr
"water_demand":     (150, 280),
# million m3/yr
"green_investment": (200, 800),
# USD million
}
NAMES = list(BOUNDS.keys())

def sample_scenarios(n, rng_):
s = {}
s["oil_price"]        =
rng_.triangular(45, 82, 120, n)
s["solar_radiation"]  =
rng_.normal(2300, 110, n)
s["desal_cost"]       =
rng_.lognormal(mean=np.log(0.95),
sigma=0.22, size=n)
s["pop_growth"]       =
rng_.triangular(1.0, 2.0, 3.0, n)
s["water_demand"]     =
rng_.normal(210, 35, n)
s["green_investment"] =
rng_.uniform(200, 800, n)
return s

N_CANDIDATE = 5000
raw = sample_scenarios(N_CANDIDATE,
rng)

# ------------------------------------
# ------------------------------
# 2. Reality-constraint layer Omega:
box bounds + THREE coupled,
#    nonlinear structural constraints
g_i(x) <= 0
# ------------------------------------
# ------------------------------
def box_mask(raw):
mask = np.ones(N_CANDIDATE,
dtype=bool)
for k, (lo, hi) in BOUNDS.items():
mask &= (raw[k] >= lo) &
(raw[k] <= hi)
return mask

def
water_supply_capacity(green_investment
, desal_cost):
"""Desalinated + baseline water
supply (million m3/yr) implied by
installed green investment, net of
unit desalination cost.
Diminishing returns via sqrt to
represent capital/efficiency
trade-offs. Coefficient calibrated
so the constraint is genuinely
binding for a plausible minority
of scenarios (~15-20%), not
vacuous and not near-universally
violated. The constant term (95)
represents baseline conventional
(non-desalinated) groundwater
supply available regardless of
green-investment level."""
return 95.0 + 7.5 *
np.sqrt(green_investment) /
np.sqrt(desal_cost)

def budget_ceiling(oil_price):
"""Public investment ceiling (USD
million) as a linear function of
an oil-revenue proxy: a rentier-
state budget cannot commit an
unlimited share of oil revenue to
green capital expenditure in a
single planning cycle. Coefficient
calibrated to bind on the
```

```
        lower-oil-price / higher-
    investment corner of the design
              space."""
          return 9.0 * oil_price

                  def
desal_energy_requirement(green_investm
                ent):
     """Electricity requirement (index
         units) implied by the
     desalination capacity that a given
      green-investment level would
                finance."""
               return 15.0 *
        np.sqrt(green_investment)

                  def
solar_energy_available(solar_radiation
                  ):
     """Usable solar energy index at a
     fixed illustrative land-use and
        conversion-efficiency factor
               (0.18)."""
         return 0.18 * solar_radiation

 def constraint_g1_water_balance(raw):
          """g1(x) = water_demand -
       water_supply_capacity <= 0
     i.e. desalinated + baseline supply
          must meet demand."""
                  supply =
water_supply_capacity(raw["green_inves
       tment"], raw["desal_cost"])
         return raw["water_demand"] -
                  supply

     def constraint_g2_budget(raw):
        """g2(x) = green_investment -
   budget_ceiling(oil_price) <= 0"""
       return raw["green_investment"] -
    budget_ceiling(raw["oil_price"])

     def constraint_g3_energy(raw):
                """g3(x) =
       desal_energy_requirement -
     solar_energy_available <= 0"""
                  return
desal_energy_requirement(raw["green_in
              vestment"]) -
solar_energy_available(raw["solar_radi
                ation"])

 g1 = constraint_g1_water_balance(raw)
     g2 = constraint_g2_budget(raw)
     g3 = constraint_g3_energy(raw)

  structural_mask = (g1 <= 0) & (g2 <=
              0) & (g3 <= 0)
          box_ok = box_mask(raw)
accept_mask = box_ok & structural_mask

        n_candidate = N_CANDIDATE
   n_box_reject = int((~box_ok).sum())
     n_structural_reject_given_box =
int((box_ok & ~structural_mask).sum())
   n_accepted = int(accept_mask.sum())
 n_rejected = n_candidate - n_accepted

 scenarios = {k: v[accept_mask] for k,
           v in raw.items()}

# ----------------------------------
   --------------------------------
     # 3. AHP weights (illustrative
    pairwise comparison matrix, Saaty
                 scale)
# ----------------------------------
   --------------------------------
   criteria = ["water_food_security",
    "environmental_sustainability",
                "economic_return",
            "job_creation",
              "technical_feasibility",
         "social_acceptance"]

              A = np.array([
        [1,   2,   3,   4,   5,   6],
        [1/2, 1,   2,   3,   4,   5],
        [1/3, 1/2, 1,   2,   3,   4],
        [1/4, 1/3, 1/2, 1,   2,   3],
        [1/5, 1/4, 1/3, 1/2, 1,   2],
        [1/6, 1/5, 1/4, 1/3, 1/2, 1],
                    ])
        col_sums = A.sum(axis=0)
 weights = (A / col_sums).mean(axis=1)
     eigvals = np.linalg.eigvals(A)
   lambda_max = np.max(eigvals.real)
        n_crit = len(criteria)
CI = (lambda_max - n_crit) / (n_crit -
                    1)
                RI_TABLE =
{1:0,2:0,3:0.58,4:0.9,5:1.12,6:1.24,7:
      1.32,8:1.41,9:1.45,10:1.49}
          RI = RI_TABLE[n_crit]
    CR = CI / RI if RI > 0 else 0.0

# ----------------------------------
   --------------------------------
     # 4. Scoring functions -- TWO
 independent specifications (R3), FOUR
         #    alternatives (R4)
# ----------------------------------
   --------------------------------
  def scale(x, lo, hi, invert=False):
         v = (x - lo) / (hi - lo)
           v = np.clip(v, 0, 1)
       return 1 - v if invert else v

    def score_alternatives_main(s):
        op  = scale(s["oil_price"],
         *BOUNDS["oil_price"])
      sol = scale(s["solar_radiation"],
       *BOUNDS["solar_radiation"])
        dc  = scale(s["desal_cost"],
         *BOUNDS["desal_cost"])
        pg  = scale(s["pop_growth"],
         *BOUNDS["pop_growth"])
       wd  = scale(s["water_demand"],
        *BOUNDS["water_demand"])
     gi  = scale(s["green_investment"],
      *BOUNDS["green_investment"])

      A1 = {  # oil-rent continuation
```

```
            "water_food_security":
          0.30 - 0.15 * wd,

"environmental_sustainability": 0.28 -
              0.15 * op,
              "economic_return":
          0.20 + 0.75 * op,
               "job_creation":
          0.25 + 0.20 * op,
           "technical_feasibility":
          0.80 + 0.0  * op,
             "social_acceptance":
          0.55 - 0.10 * op,
                    }
        A2 = {  # full green-blue
              transition
            "water_food_security":
    0.30 + 0.45 * gi - 0.25 * dc,

"environmental_sustainability": 0.30 +
       0.40 * sol + 0.15 * gi,
              "economic_return":
0.15 + 0.65 * gi - 0.20 * dc - 0.10 *
                  op,
              "job_creation":
          0.25 + 0.45 * gi,
           "technical_feasibility":
    0.40 + 0.20 * sol - 0.20 * dc,
             "social_acceptance":
    0.40 + 0.25 * gi - 0.15 * pg,
                    }
        A3 = {  # managed partial
diversification (own weighting, not an
          average of A1/A2)
            "water_food_security":
0.30 + 0.22 * gi - 0.12 * dc - 0.05 *
                wd,

"environmental_sustainability": 0.30 +
 0.18 * sol - 0.07 * op + 0.08 * gi,
              "economic_return":
    0.22 + 0.40 * op + 0.28 * gi,
               "job_creation":
    0.26 + 0.15 * op + 0.22 * gi,
           "technical_feasibility":
    0.55 + 0.10 * sol - 0.05 * dc,
             "social_acceptance":
    0.50 + 0.05 * gi - 0.05 * pg,
                    }
      A4 = {  # donor-financed rapid
  transition: higher upside, explicit
        execution-risk penalty
        risk_key: None for risk_key in
    []  # placeholder, filled below
                    }
     exec_risk = 0.12 * pg + 0.08 * dc
  # execution-risk index: population
   pressure + input-cost volatility
                  A4 = {
             "water_food_security":
    0.28 + 0.55 * gi - 0.20 * dc -
              exec_risk,

"environmental_sustainability": 0.28 +
    0.48 * sol + 0.20 * gi - 0.5 *
              exec_risk,
             "economic_return":
0.10 + 0.80 * gi - 0.15 * dc - 0.12 *
          op - exec_risk,
              "job_creation":
 0.20 + 0.55 * gi - 0.5 * exec_risk,
           "technical_feasibility":
   0.30 + 0.25 * sol - 0.25 * dc -
              exec_risk,
            "social_acceptance":
    0.35 + 0.30 * gi - 0.20 * pg,
                    }
        for d in (A1, A2, A3, A4):
                  for k in d:
              d[k] = np.clip(d[k], 0.01,
                   0.99)
          return A1, A2, A3, A4

    def score_alternatives_alt(s):
          """Alternative scoring
 specification (R3): same qualitative
    direction of every relationship,
    but different functional form
     (different coefficients and one
  squared term) to test whether the
    headline ranking is an artefact of
      the main specification."""
        op  = scale(s["oil_price"],
         *BOUNDS["oil_price"])
      sol = scale(s["solar_radiation"],
       *BOUNDS["solar_radiation"])
        dc  = scale(s["desal_cost"],
         *BOUNDS["desal_cost"])
        pg  = scale(s["pop_growth"],
         *BOUNDS["pop_growth"])
       wd  = scale(s["water_demand"],
        *BOUNDS["water_demand"])
     gi  = scale(s["green_investment"],
      *BOUNDS["green_investment"])

                  A1 = {
            "water_food_security":
          0.32 - 0.20 * wd,

"environmental_sustainability": 0.25 -
             0.20 * op**2,
              "economic_return":
          0.18 + 0.85 * op,
               "job_creation":
          0.22 + 0.28 * op,
           "technical_feasibility":
          0.78 - 0.05 * op,
             "social_acceptance":
          0.50 - 0.15 * op,
                    }
                  A2 = {
            "water_food_security":
    0.28 + 0.55 * gi - 0.30 * dc,

"environmental_sustainability": 0.26 +
     0.50 * sol**0.8 + 0.10 * gi,
              "economic_return":
0.12 + 0.55 * gi - 0.25 * dc - 0.15 *
                  op,
               "job_creation":
          0.22 + 0.55 * gi,
           "technical_feasibility":
    0.35 + 0.25 * sol - 0.25 * dc,
```

```
            "social_acceptance":
    0.38 + 0.30 * gi - 0.20 * pg,
                    }
                A3 = {
            "water_food_security":
0.28 + 0.28 * gi - 0.15 * dc - 0.08 *
                wd,

"environmental_sustainability": 0.28 +
 0.22 * sol - 0.10 * op + 0.10 * gi,
            "economic_return":
    0.20 + 0.45 * op + 0.22 * gi,
              "job_creation":
    0.24 + 0.20 * op + 0.20 * gi,
          "technical_feasibility":
    0.52 + 0.14 * sol - 0.08 * dc,
            "social_acceptance":
    0.48 + 0.08 * gi - 0.08 * pg,
                    }
    exec_risk = 0.16 * pg + 0.10 * dc
                A4 = {
            "water_food_security":
    0.26 + 0.60 * gi - 0.25 * dc -
            exec_risk,

"environmental_sustainability": 0.26 +
    0.52 * sol + 0.18 * gi - 0.6 *
            exec_risk,
            "economic_return":
0.08 + 0.85 * gi - 0.20 * dc - 0.15 *
          op - exec_risk,
              "job_creation":
 0.18 + 0.58 * gi - 0.6 * exec_risk,
          "technical_feasibility":
   0.28 + 0.28 * sol - 0.28 * dc -
            exec_risk,
            "social_acceptance":
    0.32 + 0.32 * gi - 0.24 * pg,
                    }
        for d in (A1, A2, A3, A4):
                for k in d:
            d[k] = np.clip(d[k], 0.01,
                0.99)
          return A1, A2, A3, A4


     alt_names = ["A1_oil_rent",
            "A2_green_blue",
   "A3_managed_diversification",
  "A4_donor_financed_transition"]

      A1_m, A2_m, A3_m, A4_m =
 score_alternatives_main(scenarios)
      A1_a, A2_a, A3_a, A4_a =
  score_alternatives_alt(scenarios)

# ------------------------------------
  --------------------------------
# 5. TOPSIS and VIKOR generalised to K
            alternatives
# ------------------------------------
  --------------------------------
 def stack_alts(alt_dicts, criteria):
          # returns array shape
     (n_scenarios, K, n_criteria)
    n = len(alt_dicts[0][criteria[0]])
          K = len(alt_dicts)
        M = np.zeros((n, K,
        len(criteria)))
 for k, d in enumerate(alt_dicts):
             for j, c in
     enumerate(criteria):
            M[:, k, j] = d[c]
             return M

    def topsis_multi(M, w):
"""M: (n, K, m). Returns C: (n, K)
  closeness coefficients."""
          n, K, m = M.shape
denom = np.sqrt((M**2).sum(axis=1,
        keepdims=True))
      denom[denom == 0] = 1e-9
          norm = M / denom
          wnorm = norm * w
      ideal = wnorm.max(axis=1,
          keepdims=True)
      anti = wnorm.min(axis=1,
          keepdims=True)
      d_pos = np.sqrt(((wnorm -
     ideal)**2).sum(axis=2))
      d_neg = np.sqrt(((wnorm -
     anti)**2).sum(axis=2))
 C = d_neg / (d_pos + d_neg + 1e-
             12)
             return C

 def vikor_multi(M, w, v=0.5):
          n, K, m = M.shape
       f_star = M.max(axis=1,
          keepdims=True)
       f_minus = M.min(axis=1,
          keepdims=True)
       span = f_star - f_minus
        span[span == 0] = 1e-9
 contrib = w * (f_star - M) / span
       S = contrib.sum(axis=2)
       R = contrib.max(axis=2)
       S_star = S.min(axis=1,
   keepdims=True); S_minus =
 S.max(axis=1, keepdims=True)
       R_star = R.min(axis=1,
   keepdims=True); R_minus =
 R.max(axis=1, keepdims=True)
      Sspan = S_minus - S_star;
    Sspan[Sspan == 0] = 1e-9
      Rspan = R_minus - R_star;
    Rspan[Rspan == 0] = 1e-9
 Q = v * (S - S_star) / Sspan + (1
  - v) * (R - R_star) / Rspan
             return Q

M_main = stack_alts([A1_m, A2_m, A3_m,
          A4_m], criteria)
M_alt  = stack_alts([A1_a, A2_a, A3_a,
          A4_a], criteria)

C_main = topsis_multi(M_main, weights)
 C_alt  = topsis_multi(M_alt, weights)
 Q_main = vikor_multi(M_main, weights)
 Q_alt  = vikor_multi(M_alt, weights)

   def bootstrap_ci(x, n_boot=2000,
             seed=7):
      r = np.random.default_rng(seed)
```

```
    means = [r.choice(x, size=len(x),
    replace=True).mean() for _ in
           range(n_boot)]
    return np.percentile(means, [2.5,
               97.5])

         topsis_summary = {}
for k, name in enumerate(alt_names):
     ci = bootstrap_ci(C_main[:, k])
        topsis_summary[name] = {
             "mean_C_main_spec":
round(float(C_main[:, k].mean()), 3),
             "mean_C_alt_spec":
round(float(C_alt[:, k].mean()), 3),
             "ci95_main_spec":
  [round(float(x), 3) for x in ci],
             "mean_Q_main_spec":
round(float(Q_main[:, k].mean()), 3),
                    }

 winner_main = C_main.argmax(axis=1)
  winner_alt = C_alt.argmax(axis=1)
   share_win_main = {alt_names[k]:
     round(float((winner_main ==
  k).mean()), 3) for k in range(4)}
   share_win_alt  = {alt_names[k]:
round(float((winner_alt == k).mean()),
        3) for k in range(4)}
rank_agreement = float((winner_main ==
         winner_alt).mean())

            mw_u, mw_p =
  stats.mannwhitneyu(C_main[:, 1],
C_main[:, 0], alternative="two-sided")

# ------------------------------------
  ---------------------------------
 # 6. Sobol' sensitivity of C_A2(main
    spec) using MIXED-DISTRIBUTION
    #    inverse-CDF mapping (R5),
  reported alongside the simplified
#    uniform-prior version for direct
            comparison.
# ------------------------------------
  ---------------------------------
            problem = {
             "num_vars": 6,
             "names": NAMES,
      "bounds": [[0.0, 1.0]] * 6,   #
   unit hypercube; mapped below via
            inverse CDF
                  }
           unit_samples =
    saltelli.sample(problem, 512,
      calc_second_order=False)

      def inverse_transform(u):
     """Map unit-hypercube Saltelli
    samples to the ORIGINAL mixed
     sampling distributions used in
   Monte Carlo generation (R5)."""
                 x = {}
           x["oil_price"] =
  stats.triang(c=(82-45)/(120-45),
 loc=45, scale=(120-45)).ppf(u[:, 0])
          x["solar_radiation"] =
        stats.norm(loc=2300,
       scale=110).ppf(u[:, 1])
            x["desal_cost"] =
         stats.lognorm(s=0.22,
 scale=np.exp(np.log(0.95))).ppf(u[:,
                2])
           x["pop_growth"] =
  stats.triang(c=(2.0-1.0)/(3.0-1.0),
loc=1.0, scale=(3.0-1.0)).ppf(u[:, 3])
           x["water_demand"] =
stats.norm(loc=210, scale=35).ppf(u[:,
                4])
         x["green_investment"] =
        stats.uniform(loc=200,
       scale=600).ppf(u[:, 5])
                return x

      def eval_model_mixed(u):
        x = inverse_transform(u)
            a1, a2, a3, a4 =
     score_alternatives_main(x)
    M = stack_alts([a1, a2, a3, a4],
              criteria)
       C = topsis_multi(M, weights)
          return C[:, 1]  # C_A2

   def eval_model_uniform(unit_u):
     """Original simplified version:
   treat the six inputs as uniform
    over their box bounds directly (no
    inverse-CDF mapping). Kept for
         transparent side-by-side
comparison, exactly as flagged in the
         paper's limitations."""
                x = {}
      for i, k in enumerate(NAMES):
             lo, hi = BOUNDS[k]
         x[k] = lo + unit_u[:, i] * (hi
                - lo)
            a1, a2, a3, a4 =
     score_alternatives_main(x)
    M = stack_alts([a1, a2, a3, a4],
              criteria)
       C = topsis_multi(M, weights)
             return C[:, 1]

            Y_mixed =
   eval_model_mixed(unit_samples)
            Y_uniform =
  eval_model_uniform(unit_samples)

 Si_mixed = sobol.analyze(problem,
 Y_mixed, calc_second_order=False,
      print_to_console=False)
Si_uniform = sobol.analyze(problem,
Y_uniform, calc_second_order=False,
      print_to_console=False)

    sobol_mixed = {name: {"S1":
round(float(Si_mixed["S1"][i]), 3),
                          "ST":
round(float(Si_mixed["ST"][i]), 3)}
                  for i, name in
          enumerate(NAMES)}
   sobol_uniform = {name: {"S1":
round(float(Si_uniform["S1"][i]), 3),
```

```
                                "ST":
 round(float(Si_uniform["ST"][i]), 3)}
                     for i, name in
           enumerate(NAMES)}

# ----------------------------------
   --------------------------------
 # 7. Conditional stress region + AHP
    weight-perturbation robustness
# ----------------------------------
   --------------------------------
   op_raw = scenarios["oil_price"];
gi_raw = scenarios["green_investment"]
 cond_mask = (op_raw > 95) & (gi_raw <
                  300)
     n_cond = int(cond_mask.sum())
            if n_cond > 0:
              cond_result = {
            "n_scenarios": n_cond,
               "share_A2_wins":
round(float((winner_main[cond_mask] ==
            1).mean()), 3),
                  "mean_C_A1":
     round(float(C_main[cond_mask,
            0].mean()), 3),
                  "mean_C_A2":
     round(float(C_main[cond_mask,
            1].mean()), 3),
                     }
                  else:
     cond_result = {"n_scenarios": 0}

  def perturb_and_eval(idx, factor):
            w2 = weights.copy()
        w2[idx] = w2[idx] * factor
            w2 = w2 / w2.sum()
       C2 = topsis_multi(M_main, w2)
        winner2 = C2.argmax(axis=1)
       return round(float((winner2 ==
            1).mean()), 3)

       weight_sensitivity = {}
  for i, crit in enumerate(criteria):
      weight_sensitivity[crit] = {"-
    20%": perturb_and_eval(i, 0.8),
   "+20%": perturb_and_eval(i, 1.2)}

# ----------------------------------
   --------------------------------
              # Output
# ----------------------------------
   --------------------------------
             output = {
             "screening": {
          "n_candidate": n_candidate,
         "n_box_reject": n_box_reject,

   "n_structural_reject_given_box":
    n_structural_reject_given_box,
          "n_accepted": n_accepted,
          "n_rejected": n_rejected,
       "reject_rate_pct": round(100 *
    n_rejected / n_candidate, 2),
              "ahp_weights":
 dict(zip(criteria, np.round(weights,
            3).tolist())),
        "ahp_CR": round(float(CR), 4),
                    },
     "topsis_vikor_four_alternatives":
            topsis_summary,
          "share_wins_main_spec":
            share_win_main,
           "share_wins_alt_spec":
            share_win_alt,
    "rank_agreement_main_vs_alt_spec":
       round(rank_agreement, 3),
       "mannwhitney_A2_vs_A1": {"U":
  round(float(mw_u), 1), "p": mw_p},
       "sobol_mixed_distributions":
            sobol_mixed,
        "sobol_uniform_simplified":
            sobol_uniform,

"conditional_high_oil_low_investment":
            cond_result,

  "weight_sensitivity_share_A2_wins":
          weight_sensitivity,
                }

                with
open("/home/claude/paper/results_v2.js
   on", "w", encoding="utf-8") as f:
          json.dump(output, f,
     ensure_ascii=False, indent=2,
            default=str)

       print(json.dumps(output,
     ensure_ascii=False, indent=2,
            default=str))

    # Persist raw arrays for figure
              generation
np.savez("/home/claude/paper/arrays.np
                z",
          C_main=C_main, C_alt=C_alt,
            Q_main=Q_main,

   oil_price=scenarios["oil_price"],
green_investment=scenarios["green_inve
              stment"],

solar_radiation=scenarios["solar_radia
                tion"],
  desal_cost=scenarios["desal_cost"],

water_demand=scenarios["water_demand"]
 , pop_growth=scenarios["pop_growth"],
            winner_main=winner_main,
         cond_mask=cond_mask)
```

APPENDIX B COMPUTATIONAL ENVIRONMENT

- Python 3.12
- NumPy (array operations, random number generation via numpy.random.default_rng(42))
- SciPy (stats distributions, Mann-Whitney U test)
- SALib (Saltelli sampling and Sobol' analysis)
- Matplotlib (Figs. 1–7)

All figures in this paper (Figs. 1–7) were generated directly from the executed output of Appendix A using Matplotlib, and are reproduced without manual retouching of any data value.