UCLouvain

École polytechnique de Louvain

Identifying robust transition pathways under unexpected events using Modeling to Generate Alternatives and Multi-Armed Bandits

Author: **Mathieu Jourquin**
Supervisors: **Francesco Contino, Mahdi Kchaou**
Readers: **Francesco Contino, Diederik Coppitters, Mahdi Kchaou, Quentin Lété**
Academic year 2025–2026
Master [120] in Mathematical Engineering

Abstract

The urgency of climate change has placed the energy transition at the forefront of policy agendas worldwide, driving the development of optimisation-based energy system models as decision-support tools.

Yet cost-optimality alone offers a fragile basis for long-term planning. Low-probability but high-impact events, ranging from import disruptions and geopolitical shocks to technological setbacks and shifts in public acceptance, are systematically underrepresented in conventional energy planning, despite their demonstrated capacity to derail transition pathways. A planning framework that accounts for such contingencies is therefore necessary to yield strategies that are not merely optimal, but genuinely resilient.

This thesis addresses this gap by proposing an integrated multi-horizon framework combining two methodological pillars. The Modelling to Generate Alternatives approach is applied sequentially across transition phases to map the near-optimal solution space and identify structurally distinct trajectories, which are subsequently organised into a decision tree via k-means clustering. The robustness of these candidate pathways is then assessed through a Multi-Armed Bandit algorithm with an Upper Confidence Bound policy, enabling an efficient balance between the exploration of untested strategies and the exploitation of high-performing ones under a wide range of simulated unexpected events. The objective is to identify transition pathways that remain both economically acceptable and robust when disruptive events occur.

Applied to the Belgian energy system, the framework identifies a portfolio of robust transition pathways rather than a single optimal trajectory. The results suggest that robustness is favoured by diversified low-carbon supply structures, where different production and flexibility levers can compensate for one another under stress. In particular, pathways combining substantial solar and nuclear deployment tend to perform well because they rely on structurally different sources of low-carbon capacity. Conversely, pathways depending too strongly on a single externally constrained supply option are more exposed to disruption.

Overall, the analysis shows that several pathways within the near-optimal cost envelope can display markedly different robustness under unexpected events. Energy planning should therefore compare portfolios of near-optimal alternatives not only on cost, but also on their ability to remain feasible and effective under disruptive futures. This framework could be extended to additional decision horizons, broader sets of unexpected events, and more detailed operational representations of the energy system.

Acknowledgements

This thesis is the result of long months of exploration, trial and error, and learning. It would not have taken the shape it has without the support and guidance of several people.

I would like to express my sincere gratitude to Professor Francesco Contino for his critical and expert perspective on this work. His rigorous questioning and insightful recommendations were invaluable in sharpening both the methodology and the argumentation throughout this research.

I am also grateful to Dr. Diederik Coppitters for his early trust and encouragement, for helping me identify the subject of this thesis, and for his continued support in steering the direction of this work alongside my mentor.

I would especially like to thank my thesis mentor, Mahdi Kchaou, whose guidance played a central role in this work. The many discussions we shared were essential in shaping the ideas and the overall framework presented here. His ability to clarify my thinking, bring structure to my reasoning, and maintain a clear direction throughout the year was truly invaluable. I am deeply grateful for his consistent availability and his investment in this project. This thesis bears, in no small part, the mark of our collaboration.

Finally, I would like to thank my family and close friends for their steadfast support and encouragement throughout this year.

Contents

List of Figures

List of Abbreviations

AMPL	A Mathematical Programming Language
BEV	Battery Electric Vehicle
CCGT	Combined Cycle Gas Turbine
CCS	Carbon Capture and Storage
DAC	Direct Air Capture
ESP	EnergyScope Pathway
GHG	Greenhouse Gas
GWP	Global Warming Potential
H1/H2	Horizon 1 / Horizon 2 (decision years 2035 / 2040)
IPCC	Intergovernmental Panel on Climate Change
LP	Linear Programming
MAA	Modeling All Alternatives
MAB	Multi-Armed Bandit
MGA	Modeling to Generate Alternatives
PV	Photovoltaic
SMR	Small Modular Reactor
UCB	Upper Confidence Bound

Chapter 1

Introduction

1.1 Context & Motivation

The global climate crisis represents one of the most pressing challenges of our time. As documented by the IPCC, the accelerating impacts of climate change, on economic systems, food security, and human health, leave no room for complacency [1]. The scientific consensus is unambiguous: greenhouse gas emissions are the primary driver of these disruptions, and drastic reductions are necessary to limit warming to 1.5°C above pre-industrial levels. As illustrated in Figure 1.1, current implemented policies are projected to maintain global CO_2 emissions well above net zero through 2100, in stark contrast with the sharp reductions required to meet the 1.5°C or 2°C targets [1]. In response, the Paris Agreement and subsequent national commitments have made carbon neutrality around mid-century a central benchmark for climate policy [2–4].

Yet the path to net zero is far from straightforward. While the destination is broadly agreed upon, the roadmap remains contested [3]. Some place their confidence in technological breakthroughs, from full electrification to the large-scale deployment of renewable alternatives, while others argue that structural changes in consumption behaviours are equally indispensable. In practice, the energy transition will emerge from a calibrated combination of all available levers, technological and behavioural alike, activated to varying degrees depending on context and feasibility [5, 6].

To navigate this complexity, energy system modeling has emerged as a central methodological tool for both policymakers and investors. By integrating future trends, technology costs, resource availability, and behavioural assumptions into coherent quantitative frameworks, these models enable the systematic exploration of transition scenarios and the assessment of their long-term consequences [6]. A diverse ecosystem of models now exists, spanning different scales and levels of

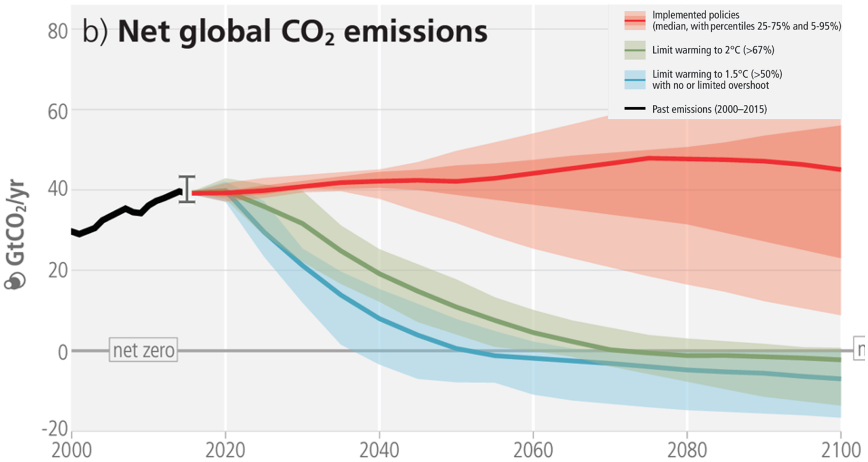

Figure 1.1: Net global CO_2 emissions under different IPCC illustrative mitigation pathways (IMPs). Each coloured band represents a distinct combination of mitigation strategies consistent with limiting global warming to 1.5 °C (blue), 2 °C (orange), or higher baselines (red). The shaded ranges show scenario variability within each pathway family. Source: [1].

complexity, from sectoral analyses to whole-country energy system optimisations, each offering a different lens through which the transition pathway can be examined and refined [5, 6]. Yet beyond their technical role, these models serve a deeper purpose: they are decision-support tools, and the decisions they inform are rarely made by a single actor with a single objective. Policymakers must navigate competing interests, political constraints, and societal trade-offs, a reality that calls for a portfolio of credible, diverse options from which informed choices can be made [7–9].

However, the predictive power of these models is fundamentally constrained by the assumptions on which they rest. Many existing energy planning studies treat future conditions as fixed and known, ignoring uncertainty altogether, making their solutions highly sensitive to deviations from the assumed trajectory and ultimately prone to failure [6, 7].

The real world offers a stark illustration of this vulnerability. In 2015, three energy companies (German RWE, Uniper, and French Engie) connected brand-new coal-fired power plants to the Dutch grid, representing a combined capacity of 3.4 GW and an investment of several billion euros, built on precise assumptions about what the European energy future would look like. A decade later, RWE and Uniper are claiming billions in compensation from Dutch taxpayers. Yet these losses were not the result of simple bad luck. As illustrated in Figure 1.2, they stem from the simultaneous materialisation of several unexpected events that had never been seriously integrated into the investment decision: the collapse of wholesale electricity prices driven by the rapid expansion of renewables, a sharp rise in European carbon prices, the accelerating competitiveness of solar and wind power, and ultimately a regulatory decision banning coal-fired generation by 2030 [9]. Each of these factors was uncertain in 2015. Together, they rendered these plants economically unviable before their construction costs had even been recovered.

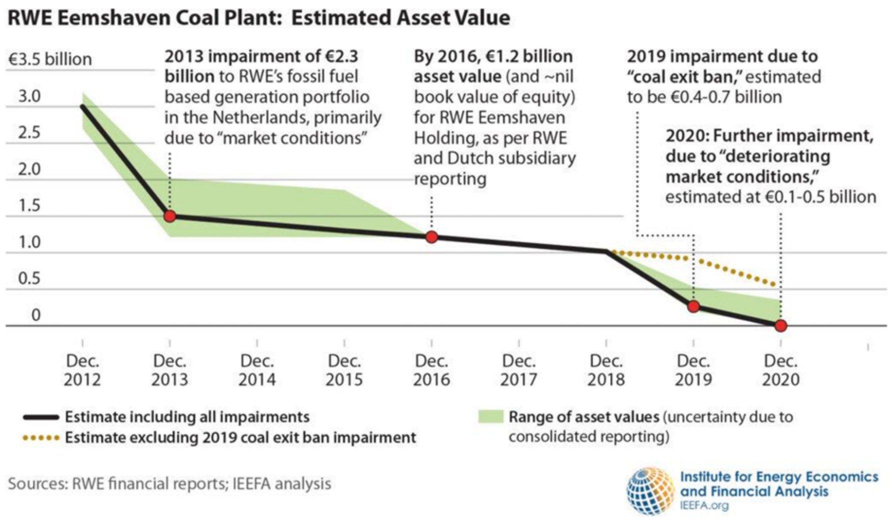

Figure 1.2: Estimated asset value of the RWE Eemshaven coal plant over time. Source: [9].

This case is emblematic of a broader challenge in energy planning. In the rare cases where uncertainty is considered at all, most models rely on probability distributions or sensitivity ranges over input parameters, technology costs, resource prices, demand levels. While useful, these approaches share a common blind spot: they focus on the ordinary and the predictable, and tend to underestimate the full spectrum of possible futures [10–12]. They are ill-equipped to handle a fundamentally different category of uncertainty; unexpected events, defined as *"events that are largely mispredicted or unpredicted, significantly deviate from expectations, and persist long enough to influence strategic decisions"* [7]. Such events may arise from technological breakthroughs, supply-chain disruptions, geopolitical conflicts, or regulatory shifts, and as the Eemshaven case illustrates, their simultaneous materialisation can be devastating for rigid, long-horizon investment decisions.

Despite their importance, unexpected events remain poorly represented in energy planning tools. Their low probability and the scarcity of data required to characterise them make them difficult to incorporate into standard uncertainty frameworks. This gap has motivated recent methodological developments, including the framework proposed by Coppitters et al. [7] to identify early-stage energy transition decisions that remain effective under unexpected events. The framework was applied to Belgium using EnergyScope Pathway (ESP), an open-source whole-energy system optimisation model covering power generation, heating, mobility, and non-energy demand simultaneously. ESP can operate in a myopic mode, where decisions are made sequentially at five-year intervals without foreknowledge of future disruptions. This limited-foresight formulation makes ESP particularly well suited for assessing the robustness of early-stage decisions: unexpected events emerge as genuine surprises rather than being anticipated from the outset, much like in the real world [5, 7].

Applied to Belgium, a particularly challenging case given its limited renewable potential and high population density, ESP provides the computational backbone on which this thesis builds. The accelerating urgency of the climate crisis, the central role of energy system models as decision-support tools, and the demonstrated capacity of unexpected events to derail even the most carefully planned transitions all paint a clear picture: robust, uncertainty-aware energy planning is not a theoretical nicety, it is a practical necessity. Yet the tools to achieve it remain underdeveloped, and the gap between what models currently offer and what decision-makers genuinely need is still wide open.

1.2 Related Work

The body of work most directly relevant to this thesis gravitates around two intertwined challenges: making energy system models robust to uncertainty, and exploring the near-optimal solution space to identify diverse, resilient transition pathways.

The treatment of uncertainty in energy planning has received growing attention over the past decade. Moret et al. [12] developed a systematic method to characterise input uncertainties in strategic energy models, showing that economic parameters, interest rates and resource prices, dominate the sensitivity of model outputs. Building on this, Moret et al. [11] proposed a robust optimisation framework for large-scale strategic energy models, capable of handling multiple uncertain parameters simultaneously in both the objective function and the constraints. Guevara et al. [10] pushed this further by modeling the temporal correlation of uncertain parameters through autoregressive and Markov chain processes, demonstrating improved out-of-sample performance on the Swiss energy system. More recently, Tan et al. [13] reviewed adjustable robust optimisation techniques for power systems, covering decomposition algorithms and multiparametric programming approaches applicable to long-term energy planning. These contributions establish a solid methodological foundation for uncertainty-aware energy planning, yet they remain largely focused on parametric uncertainty, leaving unexpected, poorly quantifiable disruptions underrepresented.

A natural platform on which to address this limitation is EnergyScope Pathway (ESP), an open-source whole-energy system optimisation model introduced by Limpens et al. [5] to plan regional energy transitions over a 30-year horizon. By jointly optimising investment decisions and hourly operations across power, heat, mobility, and non-energy demand, ESP captures the cross-sectoral interdependencies that narrower models miss. The model can operate under either perfect or myopic foresight, the latter making decisions sequentially in five-year intervals without anticipation of future conditions, a formulation far more representative of real-world planning. Rixhon explored this uncertainty-aware direction by applying reinforcement learning to ESP, training an agent to discover optimal CO_2 policy trajectories through sequential interaction with the model under parametric uncertainty [14]. Yet this line of work still operates within the realm of parametric uncertainty, leaving open the more disruptive question of how the transition holds up under genuinely unexpected events.

Coppitters et al. [7] addressed this question directly by introducing a framework to assess the impact of unexpected events on the energy transition. By injecting randomly sampled shocks onto the transition trajectory, representing disruptions such as failed technologies, geopolitical tensions, or social resistance, the framework identifies dealbreakers under perfect foresight, then reassesses those scenarios under myopic foresight to reveal how early-stage decisions hold up when disruptions arrive as genuine surprises. The study finds that a lack of electrofuel imports by 2050 is the principal dealbreaker, while accelerating renewable deployment emerges as the most robust early-stage policy. However, the diversification considered in that framework remains limited to selected early-stage decisions at the 2030 horizon, rather than to automatically generated multi-horizon near-optimal pathways. This leaves open the question of how structurally diverse transition trajectories, generated systematically across several decision horizons, perform under unexpected events.

This is precisely where the Modeling to Generate Alternatives (MGA) approach becomes relevant. Rather than seeking a single optimal solution, MGA systematically generates structurally diverse alternatives that remain within a specified cost tolerance of the optimum. Pedersen et al. [15] applied this framework to the European electricity system, showing that the near-optimal space is remarkably flat, a wide range of renewable deployment configurations can achieve near-identical costs, implying significant flexibility in system design. Esser et al. [8] demonstrated MGA's practical value in a real institutional setting, applying it to the energy system of a university campus, a scale admittedly smaller than a whole country, yet one where the stakes of the decision are concrete and the stakeholders directly involved. Their results show that at 10% above cost-optimum, no single heating technology is strictly required, illustrating that the near-optimal space can be genuinely rich even in constrained, real-world contexts. The natural next step is therefore to combine this diversity of near-optimal pathways with a robustness assessment under unexpected events.

Kchaou et al. [16] initiated this direction, first by combining MGA with interpretable Machine Learning (ML) to explore e-molecule import pathways, revealing a broad near-optimal design space with significant flexibility across energy carriers, and then by developing a more complete framework that integrates ESP, MGA, and a Reinforcement Learning algorithm based on Multi-Armed Bandit (MAB) to identify early-stage energy decisions that remain robust under unexpected disruptions. Applied to Belgium, the results show that the most resilient strategies share common features: a diversified energy portfolio, early wind deployment, faster adoption of battery electric vehicles, and reliance on e-fuel imports. This work constitutes the direct methodological foundation of the present thesis [17].

Taken together, these contributions trace a coherent intellectual trajectory: from quantifying parametric uncertainty, to building a whole-energy system platform capable of myopic decision-making, to identifying vulnerabilities under unexpected events, to exploring the near-optimal solution space and selecting the most robust pathways within it. This thesis positions itself at the frontier of that trajectory, extending the framework of Kchaou et al. [16] in a direction detailed in the following section.

1.3 Research Question & Case Study

The preceding sections have traced a clear arc: the urgency of the energy transition calls for long-horizon planning tools, yet the models available today remain ill-equipped to handle the full spectrum of uncertainty that characterises the real world, and in particular, the unexpected events that have proven capable of derailing even the most carefully designed investment strategies. What is missing is a framework that moves beyond the single cost-optimal trajectory, and instead identifies a portfolio of structurally diverse pathways that remain robust when disruptions materialise.

This thesis proposes a methodological extension along two axes. The first, inherited from Kchaou et al. [17] and retained here for continuity, is the shift from a single cost-optimal trajectory to a portfolio of near-optimal alternatives evaluated under unexpected events. The second, and the original contribution of this work, is the introduction of a multi-horizon decision tree that extends this framework from a single early-stage commitment to a sequential two-horizon structure: rather than fixing decisions at one point in time, the framework explicitly models how the near-optimal solution space evolves between 2035 and 2040, and constructs a portfolio of 12 structurally diverse trajectories that span both decision horizons. This extension captures the sequential nature of long-horizon investment planning, the fact that a 2035 commitment narrows but does not fully determine the 2040 choices available, and enables the identification of which decision horizon is more determinative for robustness under stress. The central research question is therefore:

How can we identify a multi-horizon portfolio of structurally diverse and robust energy transition pathways in the face of unexpected events?

This question unfolds along three axes: What is the structure of the near-optimal space of multi-period energy transition pathways? How can this space be clustered and represented as a tractable set of interpretable candidate pathways? And among these, which are most robust when unexpected events materialise?

The framework is applied to Belgium as a case study, building on data collected in previous works within the ESP research programme. Belgium presents a particularly demanding context for several reasons. It is a densely populated country with limited renewable energy potential, heavily constrained by land availability and geographical conditions. At the same time, Belgium's energy networks, electricity, gas, and oil , are well developed and uniformly distributed across the territory, which justifies the single-cell spatial resolution adopted by ESP: weather conditions are sufficiently homogeneous across the country to make this simplification reasonable without significant loss of accuracy [5, 7].

In practical terms, the Belgian system modelled here encompasses 6 end-use demand categories, 23 energy resources, and 112 technologies, spanning the full spectrum of the energy system from power generation and heating to mobility and non-energy uses. This breadth ensures that cross-sectoral trade-offs are properly captured, and that the pathways identified are grounded in a realistic representation of the system's constraints and possibilities [5].

Beyond its intrinsic relevance, Belgium can be viewed as an informative case for a broader class of industrialised economies: countries that are highly dependent on fossil fuel imports, face limited domestic renewable potential, and must nonetheless achieve carbon neutrality by 2050. Insights derived from the Belgian case are therefore likely to carry transferable lessons. While the case study is Belgian, the proposed framework can be applied to any regional energy system facing analogous planning challenges [5].

Chapter 2

Methodology

2.1 EnergyScope Pathway

EnergyScope Pathway (ESP) is an open-source, whole-energy system optimisation model designed to plan multi-decade energy transition pathways for regional energy systems. Applied here to Belgium, the model covers all major energy end-use sectors — electricity, low- and high-temperature heat, passenger and freight mobility, and non-energy demand — and jointly optimises investment decisions and operational dispatch across more than one hundred technologies spanning energy supply, storage, conversion, and cross-border imports and exports. By modeling the full system simultaneously, as illustrated in Figure 2.1, ESP captures the cross-sectoral trade-offs and interdependencies that narrower models may miss [5].

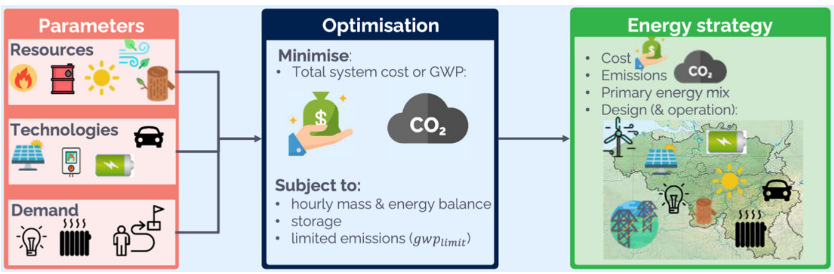

Figure 2.1: Structure of the EnergyScope Pathway (ESP) whole-energy system model applied to Belgium. Source: [5].

At its core, ESP is formulated as a linear programme (LP). In its single-period snapshot form, it minimises the total annualised system cost:

$$C_{\mathrm{tot}} = \sum_{j \in \mathcal{T}} \left[\tau(j)\, c_{\mathrm{inv}}(j)\, F(j) + c_{\mathrm{maint}}(j)\, F(j) \right] + \sum_{i \in \mathcal{R}} C_{\mathrm{op}}(i) \tag{2.1}$$

where $F(j)$ is the installed capacity of technology $j \in \mathcal{T}$, $c_{\text{inv}}(j)$ and $c_{\text{maint}}(j)$ are the specific investment and maintenance costs, and $C_{\text{op}}(i)$ is the operating cost of resource $i \in \mathcal{R}$. The annualisation factor $\tau(j)$ converts overnight investment costs into equivalent annual payments based on each technology's lifetime and a real discount rate of 1.5%. The LP structure is a deliberate design choice: convexity guarantees global optimality and keeps computation tractable at the scale of a whole national energy system.

Extending the snapshot to a full transition pathway requires capturing how investment decisions evolve over time. ESP divides the 2020–2050 horizon into a sequence of five-year phases, and the pathway objective becomes:

$$\min\ C_{\text{transition}} = C_{\text{tot,capex}} + C_{\text{tot,opex}} + C_{\text{GWP}} + \lambda \sum_{y} \delta_{\text{GWP}}[y] \tag{2.2}$$

where λ is a large penalty per tonne of excess CO_2-equivalent , set sufficiently high to render constraint violation prohibitively costly and effectively enforce the emission target as a hard constraint, and $\delta_{\text{GWP}}[y]$ is a slack variable on the emissions constraint at year y. As illustrated in Figure 2.2, the results of one phase, installed capacities, new investments, and decommissioned assets, are passed forward as initial conditions to the next, ensuring that the stock of infrastructure is consistently tracked throughout the transition. This objective is subject to four main sets of constraints: supply-demand balance equations ensuring that all energy needs are met at every time step; capacity and availability constraints governing the utilisation of each technology; an asset accounting identity linking installed capacity to new investments, inherited stock, and decommissioned units; and inertia constraints limiting how rapidly the system can evolve between consecutive phases, reflecting real-world investment lead times and infrastructure lock-in. It is precisely this last constraint that makes early-stage decisions consequential: a technology mix committed at H1 (2035) narrows the feasible space of all subsequent phases and cannot be easily undone [5, 17].

A critical modeling choice concerns how much of the future the optimiser is allowed to see. Under perfect foresight, the entire transition is optimised in a single pass with complete knowledge of future costs and constraints, an idealised but unrealistic assumption. This work adopts instead the myopic foresight formulation, in which the transition is optimised sequentially over a rolling ten-year horizon with a five-year overlap, as illustrated in Figure 2.3. At each step, only the conditions relevant to that window are visible to the optimiser; it has no knowledge of what lies beyond its current horizon. This sequential structure has a direct implication for how uncertainty is treated: a shock arising in 2040 is not anticipated when investment decisions are made in 2030, which better reflects real-world planning conditions [5, 9].

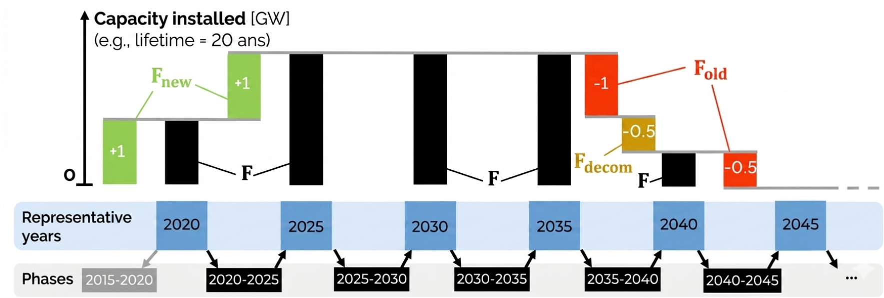

Figure 2.2: Asset accounting across consecutive five-year transition phases in ESP. Each phase begins with the infrastructure stock inherited from the previous phase (existing capacity), to which new investments are added and decommissioned units subtracted, yielding the installed base at the end of the phase. Source: [5].

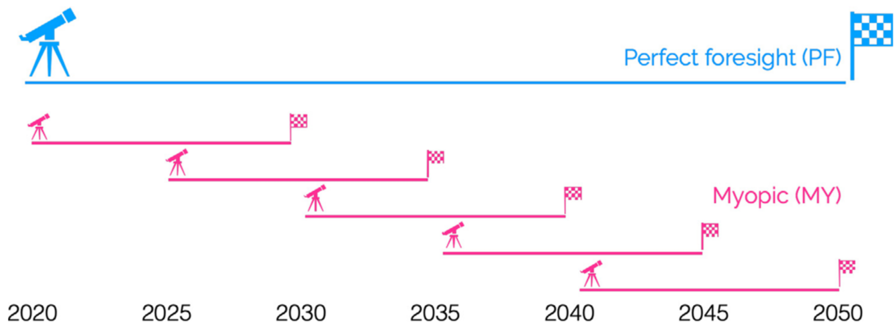

Figure 2.3: Comparison between perfect foresight and myopic foresight in ESP. Source: [5].

Because the myopic formulation does not optimise the entire 2020–2050 horizon in a single pass, the cumulative emissions budget is implemented through year-specific greenhouse gas limits rather than through a single intertemporal constraint:

$$e(t) = e_0\,(1 + (r + m)\,t)\,e^{-mt} \tag{2.3}$$

where $e_0 = 124{,}000$ $ktCO_2$-eq/yr is the 2020 emission level, r is the initial slope parameter (set to zero) and the decay parameter m is calibrated such that the cumulative emissions between 2020 and 2050 match a total budget of 1.5 $GtCO_2$-eq, ensuring carbon neutrality by 2050 while preserving flexibility in how the transition unfolds across intermediate years [5, 17].

A monthly temporal resolution is adopted, replacing the full hourly time series with twelve representative time-points per year. This substantially reduces computational cost, which is necessary given the repeated ESP optimisations required in this thesis. The consequence is that renewable intermittency, sub-monthly demand variations, and short-term balancing dynamics are not explicitly represented. As a result, the value of fast-cycling storage and operational flexibility may be underestimated. This limitation is acceptable here because the analysis focuses on long-term investment structures rather than operational dispatch. The results should therefore be interpreted as insights into structural transition robustness, not as detailed assessments of short-term system flexibility.

Following the calibration reported by Limpens et al. [5], the 2020 Belgian system includes approximately 5.94 GW of nuclear capacity, 5.90 GW of photovoltaic capacity, 2.73 GW of onshore wind, and 2.31 GW of offshore wind. From this baseline, the myopic formulation produces a reference cost-optimal trajectory, the least-cost pathway to carbon neutrality under the assumed future conditions, which serves as the starting point for the MGA framework developed in the following section. A comprehensive description of the model formulation, parameter calibration, and validation is provided in Limpens et al. [5].

2.2 Near-Optimal Solution Space

2.2.1 Modeling to Generate Alternatives

The reference trajectory produced by ESP in Section 2.1 represents the least-cost pathway to carbon neutrality under a fixed set of assumptions. While useful as a baseline, this single optimum carries a fundamental limitation: it reflects one particular vision of the future, optimised for one particular set of parameters. In practice, the true optimum cannot be known with certainty, model inputs are uncertain, structural assumptions are contestable, and the preferences of decision-makers are heterogeneous. Relying on a single solution is therefore not only epistemically naive, but also risky: a pathway tailored to one assumed future may perform poorly under a different one, and as established in Section 1.1, unexpected events are precisely the kind of disruption that can render such a pathway fragile or infeasible.

To address this limitation, this work draws on the Modeling to Generate Alternatives (MGA) framework [15, 17]. Rather than identifying one optimal solution, MGA explores the near-optimal feasible space by identifying structurally diverse solutions that remain within a specified cost increment of the optimum. This near-optimal feasible space $\mathcal{W}$ is formally defined as:

$$\mathcal{W} = \{\mathbf{x} \mid \mathbf{x} \in \mathcal{X}, \quad f(\mathbf{x}) \leq f(\mathbf{x}^*)(1+\varepsilon)\} \tag{2.4}$$

where $\mathcal{X}$ denotes the full feasible set defined by the model's constraints, $\mathbf{x}^*$ is the cost-optimal solution with objective value $f(\mathbf{x}^*)$, and ε is the user-specified slack parameter. A slack of $\varepsilon = 10\%$ is adopted in this work, consistent with values used in the energy transition literature [8, 16, 17]: this tolerance is large enough to capture structurally diverse alternatives, yet small enough to retain only economically credible pathways that decision-makers could realistically consider. A consolidated summary of all numerical parameters used throughout the whole pipeline of this work is provided in Table A.1 (Appendix A.1).

In practice, MGA is applied at the early-stage horizon of 2035. The first step is to run ESP under myopic foresight and extract the accumulated transition cost up to that year. An additional constraint is then imposed such that any alternative solution cannot exceed 110% of this reference cost at this horizon. The original objective function is replaced by a directional exploration over a set of carefully chosen decision variables, each run of ESP pushing the energy mix as far as possible in one direction while remaining within the near-optimal budget.

The choice of MGA variables is a critical design decision, as they define the dimensions along which the near-optimal space is explored. Six aggregated variables are retained, selected to reflect the strategic orientations that governments and policymakers can most directly influence: (1) solar PV capacity [GW], (2) wind capacity [GW] (onshore and offshore combined), (3) nuclear capacity [GW] (conventional and SMR), (4) battery electric vehicles [Mpkm/h], (5) e-fuel imports [GWh] (ammonia, methanol, and hydrogen), and (6) fossil resource imports [GWh] (coal, diesel, and natural gas). Together, these six dimensions characterise the fundamental strategic orientations of a transition pathway: the degree of electrification, the role of dispatchable low-carbon generation, the pace of mobility decarbonisation, and the balance between domestic production and imports [7, 17].

The full mapping between the six aggregated dimensions and their underlying ESP technologies or resources is reported in Table A.2 (Appendix A.2).

Exploring $\mathcal{W}$ by solving individual directional MGA problems produces a finite set of boundary points, but does not systematically characterise the full extent of the space. To obtain a structured approximation of the projected near-optimal space, the Modeling All Alternatives (MAA) algorithm is employed [15]. Since ESP is formulated as a linear programme, the near-optimal feasible set, and its projection onto the selected MGA variables, are convex polyhedra. MAA exploits this structure by iteratively constructing and expanding the convex hull of all near-optimal solutions found so far, solving at each iteration:

$$\min_{\mathbf{x}\in\mathcal{W}} \mathbf{n}\cdot\mathbf{x} \tag{2.5}$$

where $\mathbf{n}$ is a unit direction vector pointing towards unexplored regions of $\mathcal{W}$. The algorithm terminates when the relative growth in hull volume between successive iterations falls below a convergence threshold, indicating that further iterations add only marginal volume under the chosen convergence criterion as described by the figure 2.4. Once convergence is reached, the convex hull of $\mathcal{W}$ is densely sampled to produce a comprehensive set of 10 000 near-optimal energy mixes, providing the raw material from which representative pathways will be extracted in the following section. The full details of the MAA convergence criterion, the sampling procedure, and the associated implementation parameters are provided in Appendix A.2 [18].

The result is a structured approximation of all near-optimal energy mixes available to decision-makers at the H1 (2035) decision horizon. However, 10 000 candidate solutions cannot be directly communicated to policymakers or evaluated individually under thousands of unexpected-event scenarios. The next step is therefore to distil this high-dimensional cloud of solutions into a tractable set of interpretable, representative pathways, which is the purpose of the clustering procedure described in the following section.

2.2.2 Sampling and Clustering

The MAA procedure produces a convex hull approximation of $\mathcal{W}$ whose vertices and boundary points characterise the full extent of the near-optimal space. However, these boundary solutions alone do not represent the interior of $\mathcal{W}$: the vast majority of near-optimal energy mixes lie strictly inside the polyhedron, between the extremes. To obtain a dense and systematic representation of the projected near-optimal space, the convex hull is densely sampled using the Delaunay-based barycentric procedure described in Appendix A.2, yielding 10 000 interior points. This number reflects a compromise between geometric coverage of the six-dimensional projected space and the computational cost of the subsequent clustering procedure [19, 20].

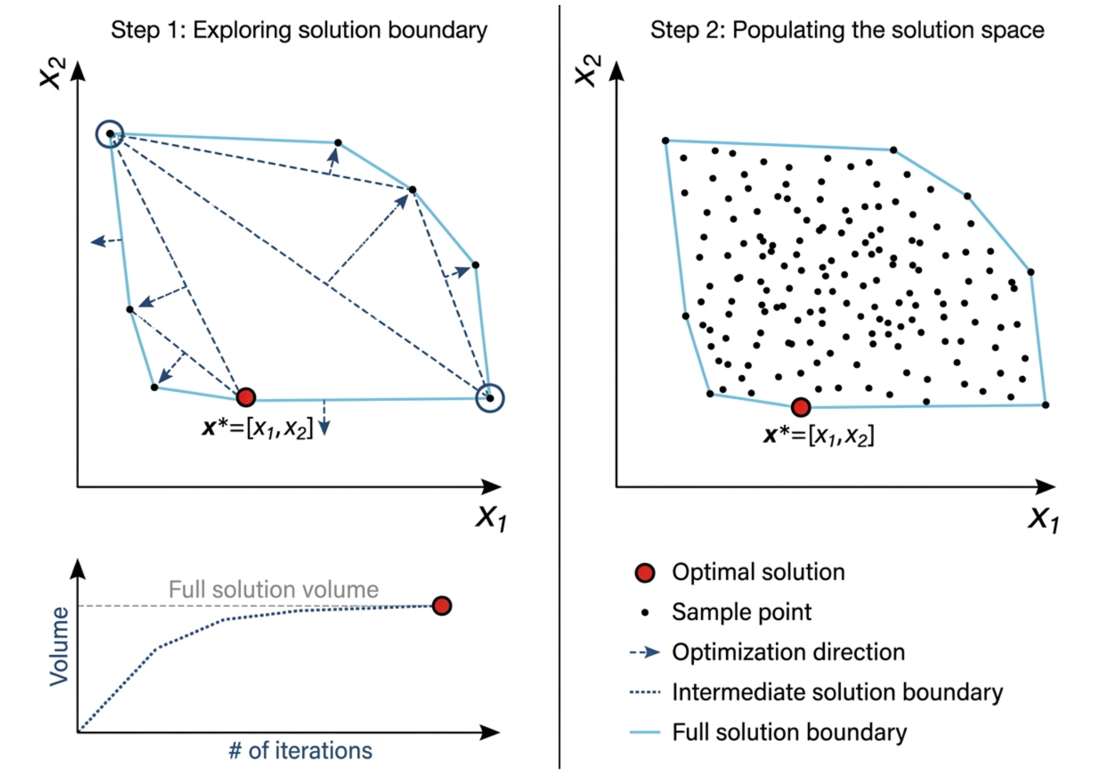

Figure 2.4: Step-by-step illustration of the Modeling All Alternatives (MAA) procedure applied to the near-optimal feasible space $\mathcal{W}$. At each iteration, a new direction $\mathbf{n}$ pointing towards an unexplored region of $\mathcal{W}$ is identified and an ESP optimisation is solved in that direction to find a new boundary point. The convex hull of all solutions found so far (grey polygon) is updated and expanded. The algorithm terminates when the relative growth in hull volume between successive iterations falls below a convergence threshold, indicating that $\mathcal{W}$ has been adequately characterised. Source: [15].

The result is a dataset of 10 000 points, each representing a valid near-optimal energy mix at the current horizon. Imposing all of them as candidate early-stage decisions would be computationally intractable: each candidate triggers a full ESP optimisation run for all subsequent horizons and all unexpected-event scenarios. A representative subset must therefore be selected. To this end, K-means clustering is applied, grouping the 10 000 points into K compact and well-separated clusters by minimising the total within-cluster inertia [21]:

$$J(\mathcal{C}, \boldsymbol{\mu}) = \sum_{k=1}^{K} \sum_{\mathbf{x} \in \mathcal{C}_k} \|\mathbf{x} - \boldsymbol{\mu}_k\|_2^2 \tag{2.6}$$

where $\boldsymbol{\mu}_k = |\mathcal{C}_k|^{-1} \sum_{\mathbf{x} \in \mathcal{C}_k} \mathbf{x}$ is the centroid of cluster k. Before clustering, the data are standardised using a zero-mean, unit-variance scaling, ensuring that all six MGA variables contribute equally to the distance metric regardless of their physical units and orders of magnitude. The algorithm is initialised ten times with different random seeds, retaining the solution with the lowest inertia.

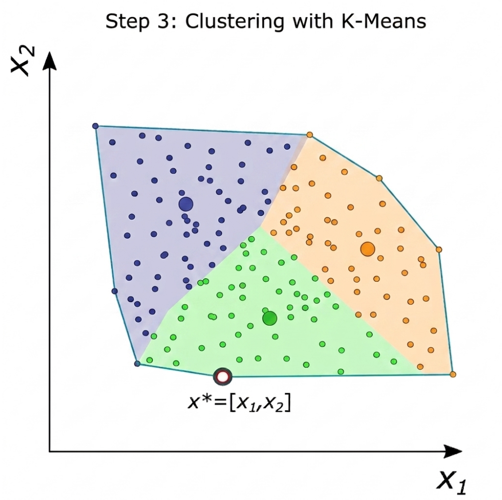

Figure 2.5: Schematic illustration of K-means clustering applied to the near-optimal feasible space $\mathcal{W}$ (grey region). The algorithm partitions the 10,000 sampled near-optimal energy mixes into K compact and well-separated clusters by minimising the total within-cluster inertia. Cluster boundaries are shown as coloured regions. Inspired by [15].

The number of clusters K is supported by three complementary considerations. The elbow method [22] and the silhouette score [23] provide data-driven diagnostics confirming that the selected values yield clusters that are both internally coherent and mutually well-separated; the corresponding figures are presented in Section 3.1, and the computation method is detailed in Appendix B.4. Computational tractability imposes an upper bound on K, since each additional cluster propagates through all subsequent horizons and unexpected-event evaluations: $K_{\text{H2}} = 3$ yields exactly $4 \times 3 = 12$ leaf trajectories, a value consistent with a meaningful UCB1 evaluation within 2 000 pulls. Finally, a uniform branching factor ensures that each H1 branch contributes equally to the final portfolio, preserving the comparability of leaf pathways across branches.

The centroid of each cluster constitutes a representative early-stage decision , a six-dimensional vector that summarises the strategic orientation of all near-optimal pathways in its neighbourhood. How these centroids are implemented as hard constraints in ESP, and how they are combined across horizons to form a complete portfolio of transition trajectories, is the subject of the following subsection.

2.2.3 Decision Tree

The clustering procedure yields K representative centroids at each horizon, each capturing a structurally distinct early-stage energy mix. To translate these centroids into actionable transition pathways, each centroid $\boldsymbol{\mu}_k^{(h)}$ is re-injected into the ESP model as a set of six hard constraints on the MGA variables at the corresponding horizon y_h:

$$\mu_{k,v}^{(h)} \cdot (1-\delta) \;\leq\; \text{MGA\_VAR}_v(y_h) \;\leq\; \mu_{k,v}^{(h)} \cdot (1+\delta) \qquad \forall\, v \in \{1, \ldots, 6\} \tag{2.7}$$

where $\delta = 1\%$ provides a narrow feasibility band around the centroid value to avoid numerical infeasibilities in the LP solver while faithfully reproducing the intended energy mix. Beyond the constrained horizon, the ESP myopic optimiser retains full freedom to complete the transition at minimal cost subject to the prevailing CO_2 constraints and the infrastructure stock inherited from the previous phase.

The full pipeline iterates sequentially across horizons. At H1 (2035), the MGA–MAA –sampling–clustering procedure is applied to the unconstrained near-optimal space, yielding $K_{\text{H1}} = 4$ representative centroids (full centroid values are reported in Appendix B.1). Each centroid is imposed via Equation (2.7) and the constrained ESP run is passed as fixed initial conditions to the next horizon. At H2 (2040), the entire procedure is repeated independently for each of the four H1 branches: the near-optimal space is re-explored conditional on the H1 constraints, sampled, and clustered into $K_{\text{H2}} = 3$ new centroids. This branching structure is illustrated in Figure 2.6.

Formally, each leaf trajectory $\mathcal{T}_{i,j}$ is characterised by the pair of centroids $(\boldsymbol{\mu}_i^{(\text{H1})}, \boldsymbol{\mu}_j^{(\text{H2})})$, where $i \in \{0, 1, 2, 3\}$ indexes the H1 branch (H1-i) and $j \in \{0, 1, 2\}$

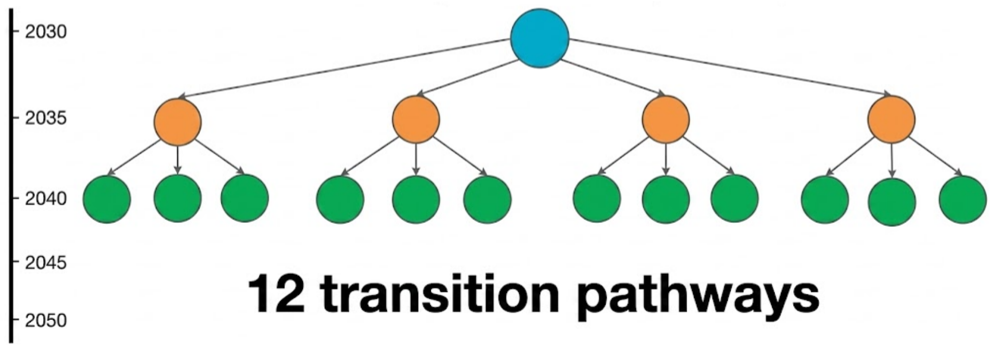

Figure 2.6: Transition decision tree yielding 12 candidate pathways labelled B0–B11. Four H1 centroids at the 2035 decision horizon each branch into three H2 centroids at the 2040 decision horizon ($4 \times 3 = 12$ leaves). Each pathway is characterised by a fixed sequence of six MGA decisions at both 2035 and 2040.

indexes the H2 branch (H2-j). For compactness, each trajectory is henceforth labelled B_n with $n = 3i + j$; the full correspondence is given in Table B.2. As illustrated in Figure 2.7, the original framework of Kchaou et al. [17] fixes decisions at a single early-stage horizon before evaluating robustness under unexpected events via a Multi-Armed Bandit algorithm. The present thesis extends this framework in two directions: first, by introducing a second decision horizon H2 (2040) that branches from each H1 centroid, capturing how the conditional near-optimal space evolves as the transition progresses; and second, by explicitly modeling the sequential nature of long-horizon investment planning through a structured decision tree rather than a flat set of alternatives.

The $4\times3 = 12$ resulting trajectories form a portfolio of near-optimal, structurally diverse transition pathways spanning the transition to 2050. Each trajectory is therefore defined by a consistent sequence of fixed strategic decisions at H1 and H2, while the remaining system evolution is optimised by ESP. This portfolio constitutes the input to the robustness evaluation described in the following section, where each trajectory is stress-tested under a large number of unexpected-event scenarios using the Multi-Armed Bandit algorithm [17].

2.3 Robustness Evaluation

2.3.1 Unexpected Events

Having constructed the decision tree of 12 candidate pathways, the robustness of each trajectory is evaluated against a broad range of unexpected events. Following the definition introduced in Chapter 1, an unexpected event is a development that

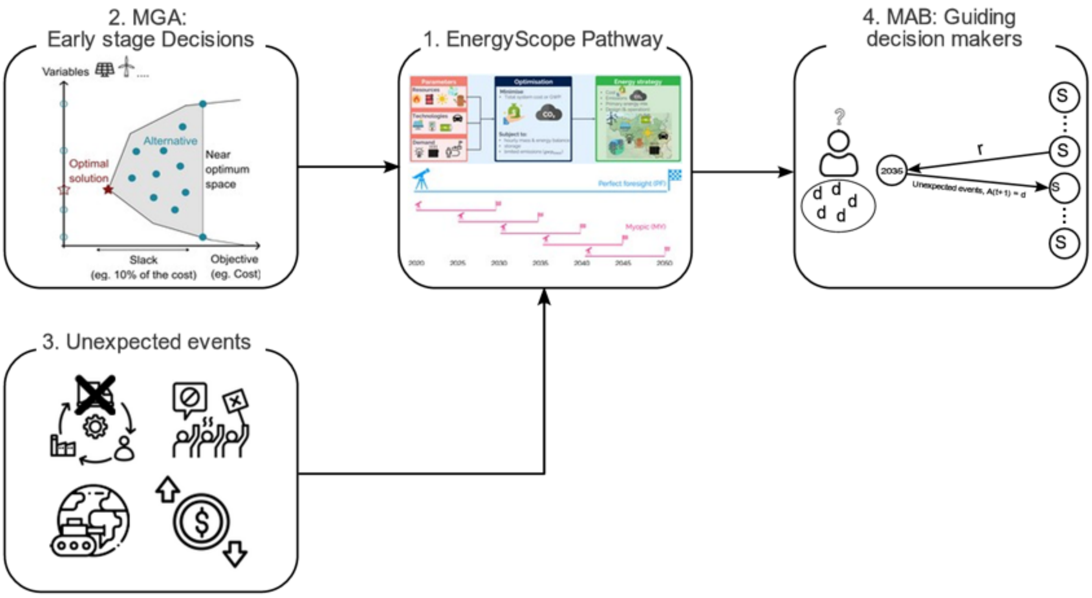

Figure 2.7: Overview of the MGA–MAB framework of Kchaou et al. [17], which serves as the methodological foundation of this thesis. The framework proceeds in three stages: (1) the MGA–MAA procedure explores the near-optimal feasible space at a single early-stage horizon and extracts representative energy mixes via K-means clustering; (2) each representative mix is imposed as a fixed strategic decision and the resulting trajectory is evaluated under unexpected-event scenarios via the ESP model; (3) a Multi-Armed Bandit (UCB1) algorithm allocates the evaluation budget adaptively to identify the most robust trajectories.

is largely mispredicted, significantly deviates from expectations, and persists long enough to influence strategic energy decisions.

The space of unexpected events is inspired by the taxonomy proposed by Coppitters et al. [7] and adopted in Kchaou et al. [17]. It comprises 47 parameters organised in six categories: availability of imported resources (electrofuels, biofuels, and electricity), societal resistance to change (PV, wind, heating renovations, and mobility), exchange-rate deterioration affecting Asian supply chains, upward shocks on final energy demand, availability of unicorn technologies (CCS, DAC, geothermal, ammonia-fired CCGT, and ammonia cracking), and the timing of a potential nuclear phase-out. The full parameter taxonomy, discretisation rules, and impact mechanisms are described in [7] and summarised in Appendix A.3.

With 47 parameters each taking discrete values across multiple levels (as detailed in Appendix A.3 and Coppitters et al. [7]), the total combinatorial space is astronomically large, on the order of tens of orders of magnitude beyond any tractable enumeration. A Sobol low-discrepancy sequence is therefore used to generate a pre-computed library of $N_s = 39{,}201$ quasi-random scenarios in $[0, 1]^{47}$, ensuring more uniform coverage of the parameter space than purely random sampling.

The present implementation departs from Coppitters et al. [7] in one important aspect, motivated by the structure of the decision tree. In the original framework, unexpected events are injected as hard constraints , setting technology availabilities to zero or capping capacities , directly onto a cost-optimal trajectory. In the present framework, however, the decision tree fixes structural choices at H1 and H2 simultaneously. Injecting a hard availability constraint on a technology whose capacity is simultaneously fixed by a centroid creates a direct conflict between constraints, producing guaranteed infeasibility that reflects a modelling artefact rather than genuine fragility.

To resolve this, all unexpected events are reformulated as cost perturbations: a Sobol value $p \in [0, 1]$ is mapped to a multiplicative cost factor $(1 + p \cdot \beta)$ applied to the relevant c_{inv} or c_{op} parameters, with $\beta = 9$ (maximum multiplier $\times 10$), interpreted as a severe stress-test parameter rather than a probabilistic estimate of future cost increases. Shocks are applied uniformly over 2040–2050, covering one centroid-constrained phase followed by two fully free optimisation phases in 2045 and 2050.

This formulation has a natural economic interpretation that connects directly to the Eemshaven case introduced in Section 1.1. Just as RWE's coal plant became economically unviable not because it was physically destroyed but because market conditions made it impossibly costly to operate, the soft-constraint approach signals disruption through prices rather than prohibitions. The LP solver steers away from technologies whose cost has exploded, substituting cheaper alternatives wherever inertia constraints allow. Trajectories whose committed 2040 infrastructure aligns with technologies that prove resilient under stress incur lower total costs; those that over-relied on disrupted technologies pay through stranded assets and expensive substitutions. The substitution response unfolds across the free 2045 and 2050 phases, where the myopic optimiser re-optimises without any MGA-variable bounds.

A small number of shocks retain their original hard-constraint formulation: the nuclear phase-out sets the capacity factor c_p to zero while leaving installed capacity intact, and unicorn technology availability blocks new installation via $f_{\max} = 0$. One parameter, avail_nuclear_smr, is neutralised throughout to avoid a direct conflict with the SMR centroid constraint. The full taxonomy of all 47 parameters, their AMPL targets, mechanisms, and maximum impacts is provided in Table A.3 (Appendix A.3).

2.3.2 Multi-Armed Bandit

The 12 candidate trajectories must now be evaluated for robustness against the unexpected-event scenarios introduced in Section 2.3.1. Evaluating every trajectory under every scenario is computationally prohibitive: with 12 arms and N_s = 39,201 Sobol scenarios, each requiring a full multi-horizon ESP solve, exhaustive enumeration is out of reach. This work therefore adopts a Multi-Armed Bandit (MAB) approach, which allocates the evaluation budget adaptively rather than uniformly [24, 25]. The analogy is a gambler facing a row of slot machines of unknown distribution: at each step, one arm is pulled and a reward observed, and the challenge is to balance exploitation of high-performing arms with exploration of under-evaluated ones. Here each arm is one of the 12 candidate trajectories, and each pull evaluates it under one randomly drawn scenario.

Arm selection follows the Upper Confidence Bound (UCB1) policy [25]:

$$a_t = \underset{a \in \{1,\ldots,12\}}{\arg\max} \left[\hat{\mu}_t(a) + \sqrt{\frac{\alpha \ln t}{N_t(a)}} \right] \tag{2.8}$$

where $\hat{\mu}_t(a)$ is the empirical mean reward of arm a, $N_t(a)$ its pull count, and $\alpha = 2$ the standard UCB1 exploration parameter [17, 25]. The first term exploits high-performing trajectories; the second ensures no arm is permanently neglected.

At each step, the selected trajectory a_t is evaluated under a scenario $\mathbf{s}_t$ drawn uniformly from the Sobol library. The reward is:

$$r_t(a_t, \mathbf{s}_t) = \begin{cases} \dfrac{C^*}{C_{\text{scenario}}(a_t, \mathbf{s}_t)} & \text{if the CO}_2 \text{ target is met by 2050,} \\ 0 & \text{otherwise} \end{cases} \tag{2.9}$$

where C^* is the unperturbed minimum cost. Rewards below 1 reflect the trajectory's combined intrinsic cost premium and shock-induced overrun; a reward of 0 flags infeasibility. This scalar simultaneously penalises cost inefficiency and infeasibility. To separate the two effects, the shock-induced degradation metric $(C_{\text{scenario}}/C_{\text{baseline}} - 1)$ is also reported in the results alongside three complementary signals: pull allocation, success rate, and mean cost under stress.

UCB1 concentrates the computational budget on the most promising candidates, making the evaluation tractable within 2 000 pulls while preserving methodological continuity with Kchaou et al. [17]. The trade-off is reduced precision on underperforming arms, which is acceptable when the objective is identification of robust candidates rather than exhaustive characterisation of all 12 pathways. The evaluation opens with an initialisation phase pulling each arm once, followed by the adaptive UCB phase. Full pseudocode is provided in Algorithm 1 (Appendix A.4); computational budget details are in Appendix A.5.

Chapter 3

Results

3.1 Trajectory Archetypes

The methodology of Chapter 2 produces a structured portfolio of transition pathways through the sequential application of MGA, clustering, and decision tree construction. Before stress-testing these pathways under unexpected events, it is worth characterising the pathways themselves: what the identified archetypes look like, how structurally distinct they are, and whether they genuinely span the near-optimal solution space. The answers provide the empirical basis for the robustness analysis that follows.

The selection of $K_{\mathrm{H1}} = 4$ at horizon H1 is supported by two data-driven diagnostics (the computation method is detailed in Appendix B.4). As shown in Figure 3.1, the elbow curve exhibits a clear inflection at $K = 4$, and the silhouette score reaches a global maximum at the same value, indicating that the four clusters are both internally coherent and mutually well-separated. The agreement of two independent criteria provides strong support for this choice.

The picture at horizon H2 is more nuanced. As shown in Figure 3.2, neither the elbow curve nor the silhouette score singles out a specific value of K across the four H1 branches: the elbow curves decline continuously without a sharp inflection, and the silhouette scores remain modest at $K = 3$ before rising again at higher values. This is consistent with the MAA convergence diagnostics reported in Appendix B.2: three of the four H1 branches reach the iteration cap without converging, and their conditional hull volumes at H2 (0.50–0.56) exceed both the converged branch (0.39) and the H1 volume (0.286). The conditional near-optimal space at H2 is therefore geometrically more intricate than at H1 and does not decompose naturally into a small number of well-separated clusters.

In the absence of a clear data-driven signal, the choice of $K_{\mathrm{H2}} = 3$ rests on the three structural considerations introduced in Section 2.2.2: uniform branching across

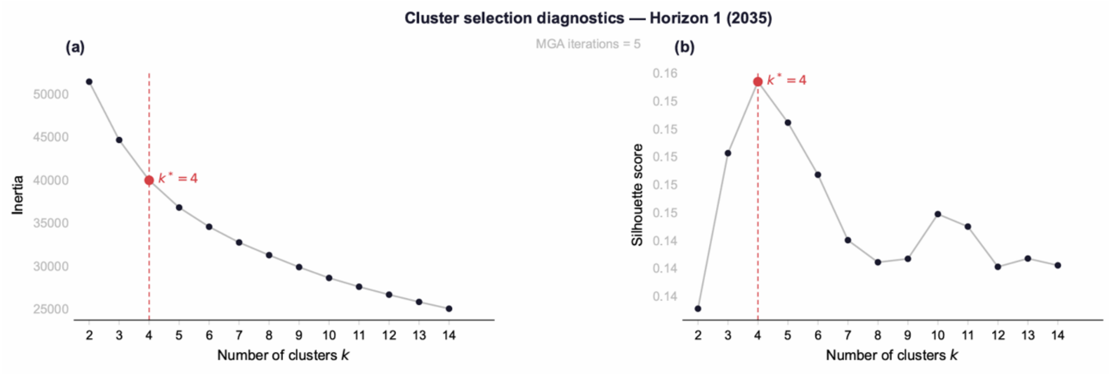

Figure 3.1: Cluster selection diagnostics at horizon H1 (2035): (a) elbow curve, showing total within-cluster inertia as a function of the number of clusters K (lower is better; the optimal K corresponds to the "elbow" where marginal improvement diminishes); (b) silhouette score, measuring the average ratio of inter-cluster separation to intra-cluster cohesion (higher is better; a score near 1 indicates well-separated, compact clusters).

the four H1 branches, computational tractability of the subsequent MAB evaluation, and consistency with the elbow diagnostics as the smallest value supported across all branches. With $K_{\mathrm{H2}} = 3$, the tree yields a portfolio of 12 leaf trajectories, a number compatible with a meaningful UCB1 evaluation within the 2 000-pull budget. This choice is grounded in tractability and branching uniformity, not in any claim that $\mathcal{W}$ physically contracts at H2. The MAA diagnostics in fact suggest the opposite for the conditional space, and the larger conditional hull volumes should not be read as an expansion of the global feasible set: they indicate only that the H1-conditioned slices explored at H2 have a more irregular geometry than the unconditional space at H1. Accordingly, $K_{\mathrm{H2}} = 3$ is best understood as a principled and transparent compromise, retained even where the silhouette score alone would point to a larger value.

With the cluster structure established, the 12 candidate trajectories can be read first through their economic envelope and then through their structural composition. Figure 3.3 shows the baseline transition cost of each pathway, the total cost incurred under unperturbed conditions, sorted from least to most expensive.

All 12 trajectories satisfy the $\varepsilon = 10\%$ near-optimality constraint, with baseline costs ranging from +3.0% above C^* (B2) to +7.3% (B9). This is a non-trivial result: the MGA procedure generates structurally diverse energy mixes by construction, and there was no guarantee a priori that all 12 centroid combinations would remain within the near-optimality bound. The spread in baseline costs also illustrates that structural diversity comes at a price, since trajectories committing to less

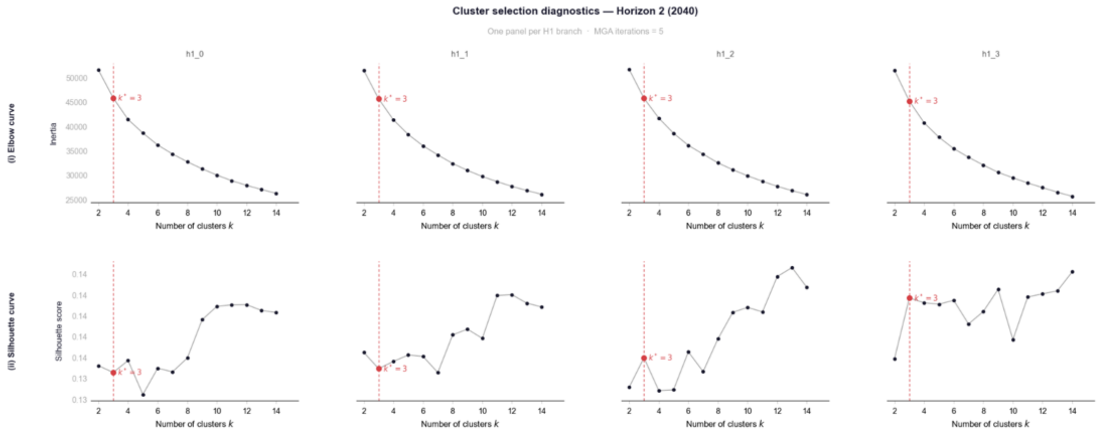

Figure 3.2: Cluster selection diagnostics at horizon H2 (2040), computed independently for each of the four H1 branches (one panel per branch): (i) elbow curves showing total within-cluster inertia vs. K; (ii) silhouette scores. Unlike at H1, neither criterion produces a clear signal: the elbow curves decline continuously without a sharp inflection, and silhouette scores remain modest at $K = 3$ before rising again at higher values.

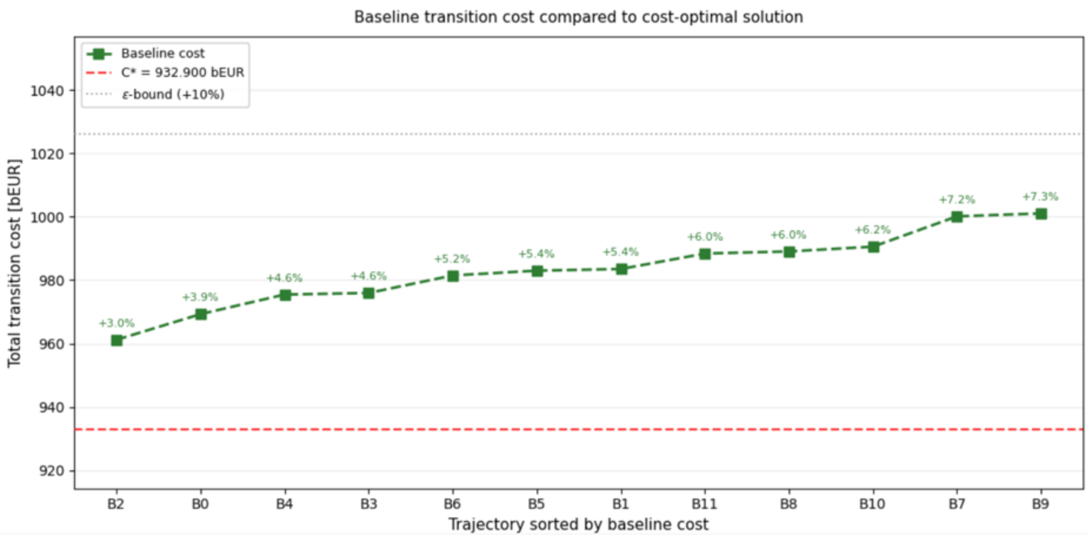

Figure 3.3: Baseline transition cost of the 12 near-optimal candidate pathways (B0–B11), sorted from least to most expensive.

cost-efficient technology mixes at H1 and H2 incur a higher baseline cost even before any unexpected event is considered. That premium is the cost of optionality: by accepting a modest economic overhead, decision-makers gain access to a broader set of structurally distinct pathways, each reflecting a different but politically implementable combination of technology choices. A single cost-optimal trajectory may be economically efficient yet politically fragile, whereas a portfolio of near-optimal alternatives offers the flexibility to adapt to the constraints and preferences of different institutional contexts.

Turning from cost to composition, Figure 3.4 tracks the evolution of each of the six MGA variables from H1 (2035) to H2 (2040) across all 12 pathways.

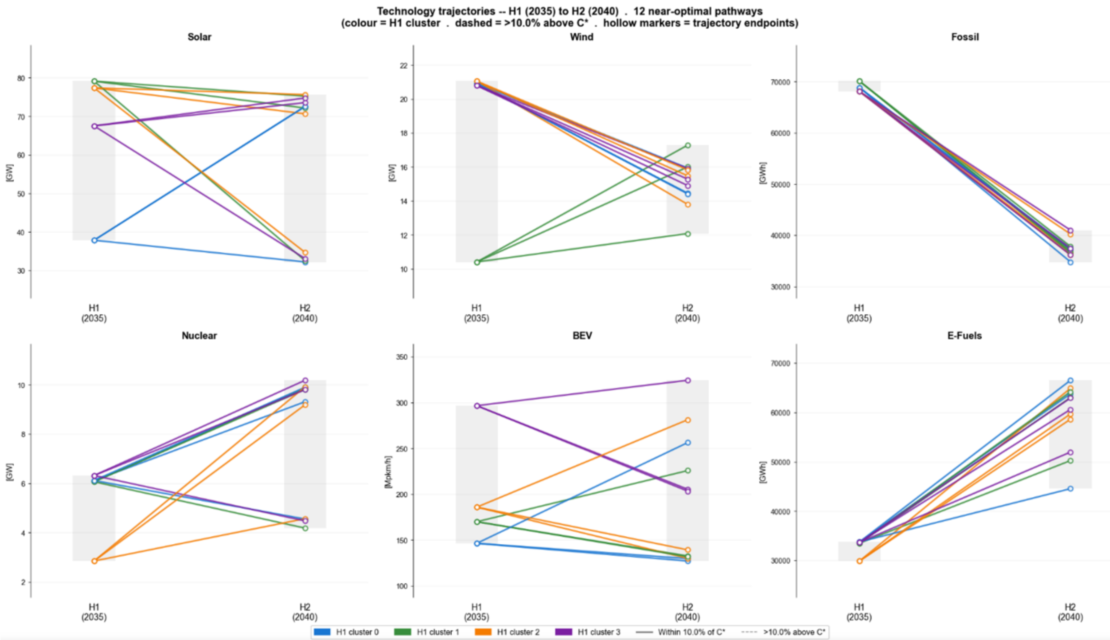

Figure 3.4: Evolution of the six MGA variables from horizon H1 (2035) to horizon H2 (2040) across the 12 near-optimal pathways (B0–B11). Each line connects the H1 centroid value (left) to the H2 centroid value (right) for one trajectory.

Two variables follow clear and uniform trends across all trajectories. E-fuel imports increase consistently from H1 to H2 regardless of the pathway, confirming that greater reliance on imported electrofuels is a near-universal feature of the near-optimal space at 2040. This is a direct consequence of the tightening CO_2 constraints, which push the system towards cleaner energy carriers as domestic decarbonisation options become more constrained. Fossil resources follow the opposite pattern, declining sharply and uniformly towards H2, consistent with the progressive elimination of carbon-intensive supply from the feasible solution space.

The remaining four variables show more nuanced dynamics. Wind capacity has one outlier at H1, trajectory H1-1 with distinctively low deployment (∼10 GW), while all other trajectories cluster at higher levels (∼20–21 GW). By H2 this polarisation resolves towards intermediate values across all pathways, suggesting that the near-optimal space at 2040 penalises extreme wind strategies in either direction. Nuclear follows the inverse pattern: starting from a relatively homogeneous cluster of low-to-medium values at H1, it bifurcates sharply at H2 into a high-nuclear group (around 9 GW) and a low-nuclear group (around 4 GW), reflecting the emergence of two fundamentally different long-term decarbonisation strategies. Solar PV polarises in a similar way: while H1 decisions span a range of solar intensities, the H2 configurations cluster towards either very high deployments (around 75 GW) or very low ones (around 35 GW), with few intermediate configurations. BEV adoption is the most evenly distributed at H2, with trajectories spread across low, medium, and high electrification levels, reflecting the genuine flexibility of the near-optimal space along this dimension.

Across horizons, then, the near-optimal space evolves along three distinct regimes: some dimensions converge (e-fuels, fossils), others polarise (nuclear, solar), and others remain open (BEV, wind).

The full structure of the decision tree is presented in Figure 3.5, which provides radar plots for each node, making the six-dimensional centroid of each archetype visible at a glance. Exact centroid values for all nodes are tabulated in Table B.1 (Appendix B.1).

A first grouping is defined by nuclear commitment at H2. Eight of the twelve trajectories converge towards high-nuclear configurations, visible as a pronounced extension along the Nuclear axis: B1, B2, B4, B5, B7, B8, B9, and B10. Within this group, solar deployment draws a sharp internal line. B1, B4, B7, and B9 combine high nuclear with high solar, producing a balanced and diversified profile, whereas B2, B5, B8, and B10 pair high nuclear with low solar and strong e-fuel reliance, yielding the more concentrated, import-dependent profile. The two sub-families share a similar radar shape and are therefore structurally close. This is a degree of redundancy that follows directly from applying the MAA procedure independently for each H1 branch, which can generate similar H2 configurations from different H1 starting points.

The wind dimension reinforces this visual similarity rather than adding a new distinction: as already seen in Figure 3.4, the H1 polarisation between high and low wind deployments collapses towards intermediate values at H2, compressing the differences along this axis at the second horizon and bringing otherwise distinct radar profiles closer together.

The 12 archetypes thus span the near-optimal solution space along three struc-

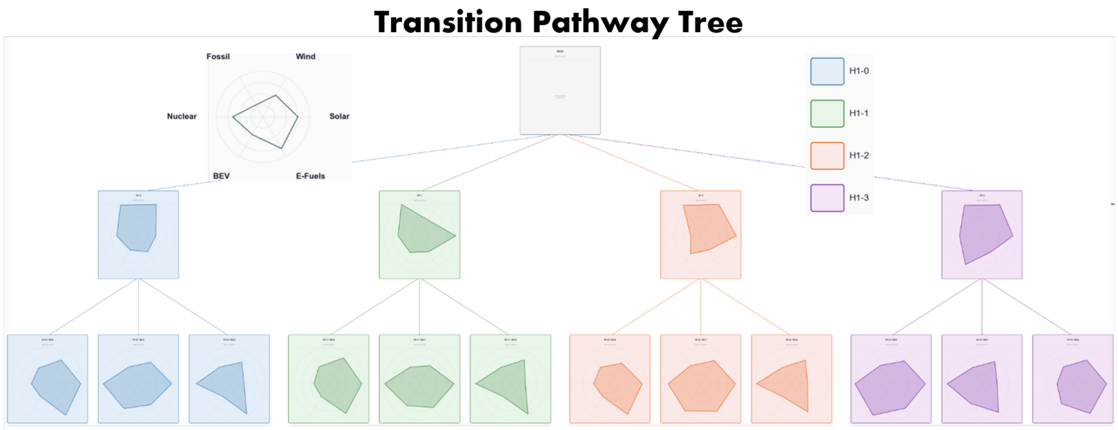

Figure 3.5: Transition pathway decision tree with radar plots at each node. Each radar shows the six-dimensional centroid of the corresponding archetype across the six MGA variables (axes, clockwise from top): solar PV [GW], wind [GW], nuclear [GW], battery electric vehicles [Mpkm/h], e-fuel imports [GWh], and fossil resources [GWh]. Exact centroid values are tabulated in Appendix B.1.

tural axes (solar intensity, nuclear commitment, and mobility electrification) while sharing a common trajectory of declining fossil dependency and growing e-fuel reliance. The MGA framework generates genuine structural diversity at H1, and the decision tree propagates it into 12 economically credible leaf trajectories. Whether this structural diversity translates into robustness diversity under unexpected events is the question addressed in the following section.

3.2 Robustness Evaluation Results

The Multi-Armed Bandit evaluation characterises the robustness of each candidate trajectory through four complementary signals reported across five empirical measures. The first signal is joint cost-and-feasibility performance under stress, captured by a composite of pull allocation and mean reward; the two are coupled by construction, since the UCB1 policy allocates pulls in proportion to observed mean rewards. The three remaining signals are mutually independent. The success rate captures feasibility probability independently of cost level. The average cost under unexpected events isolates cost performance conditional on feasibility, separating trajectories that fail often but cheaply from those that succeed often but expensively. Worst-case resilience is measured through both the standard deviation of rewards, which reflects performance stability across the full scenario distribution, and the Conditional Value at Risk (CVaR) [26], reported at the 10% level (CVaR_{10}) for the broader adverse tail and at the more extreme 5% level ($CVaR_5$) for the very worst scenarios; together they reveal whether a trajectory degrades gracefully or collapses entirely. The convergence of these four independent dimensions (feasibility, conditional cost, stability, and worst-case resilience) provides the cross-validation basis for the robustness conclusions of this work. Full statistical details, confidence intervals, and CVaR profiles are reported in Appendix B.5, and the complete arm ranking in Table B.5 (Appendix B.3).

3.2.1 Cumulative Pulls Allocation

Figure 3.6 presents the evolution of cumulative pulls per trajectory over the 2 000 MAB iterations. Under the UCB1 policy, pull allocation is not uniform: arms that consistently yield high rewards attract progressively more evaluations as the algorithm's confidence in their superiority grows, while underperforming arms are deprioritised. The colour scale, from dark green (most pulled) to dark red (least pulled), reflects the final pull ranking and makes the separation between the two groups immediately visible.

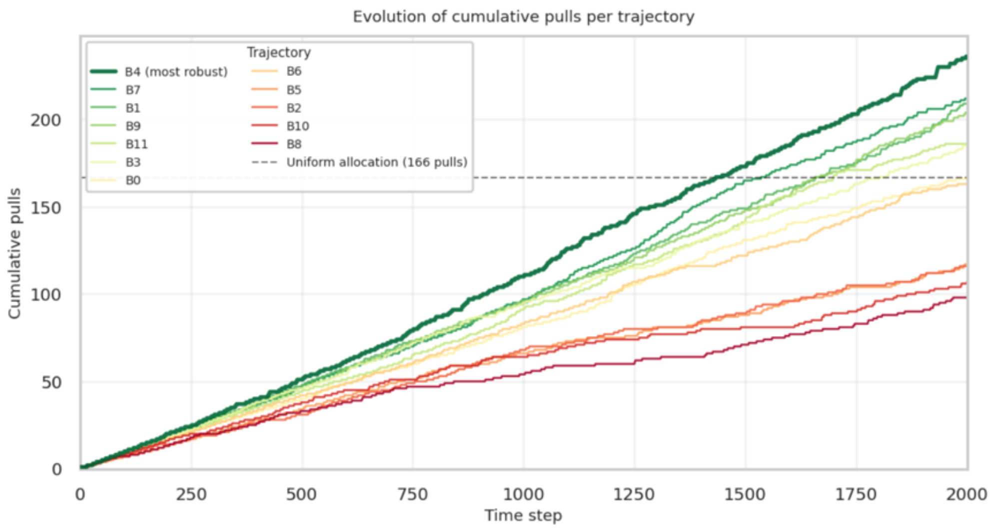

Figure 3.6: Evolution of cumulative pulls per trajectory over 2,000 MAB iterations under the UCB1 policy ($\alpha = 2$, $K = 12$ arms). Each line represents one of the 12 candidate trajectories (B0–B11). The colour scale runs from dark green (most evaluated: B4 with 236 pulls) to dark red (least evaluated: B8 with 98 pulls).

Two groups emerge clearly from the pull distribution. The green cluster, B4 (236 pulls), B7 (212), B1 (209), and B9 (204), all receive appreciably more than the uniform allocation baseline of 166 pulls, with their cumulative curves diverging progressively upward from the rest of the field. B4, B7, and B1 originate from three different H1 branches (H1-1, H1-2, and H1-0 respectively), yet independently converge towards a structurally similar configuration at H2: a balanced, high-solar and high-nuclear energy mix, shown in Figure 3.7. What unites them is a common structural signature, namely high solar deployment (~72–75 GW), high nuclear capacity (~9.2–9.9 GW), moderate wind (12–15 GW), and relatively contained e-fuel imports (44–58 TWh). No single dimension dominates; the portfolio is diversified. B9, from H1-3, joins this group by a different structural route: it relies on the highest nuclear deployment in the entire portfolio (10.18 GW) combined with the highest BEV adoption (324 Mpkm/h), two independent decarbonisation levers that reinforce each other's resilience.

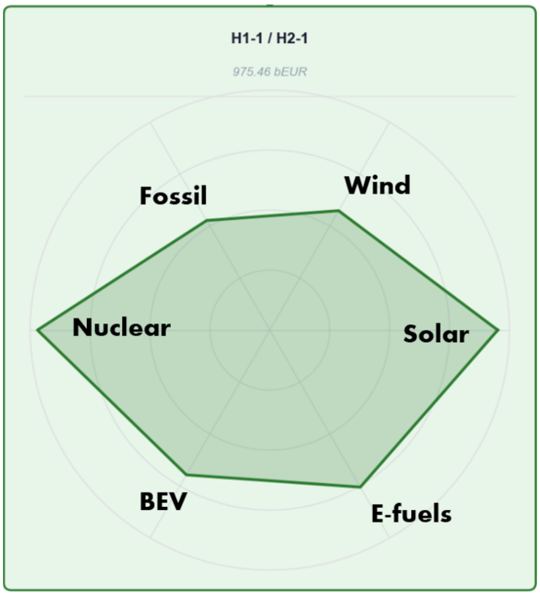

Figure 3.7: Radar profile of B4, the most robust trajectory.

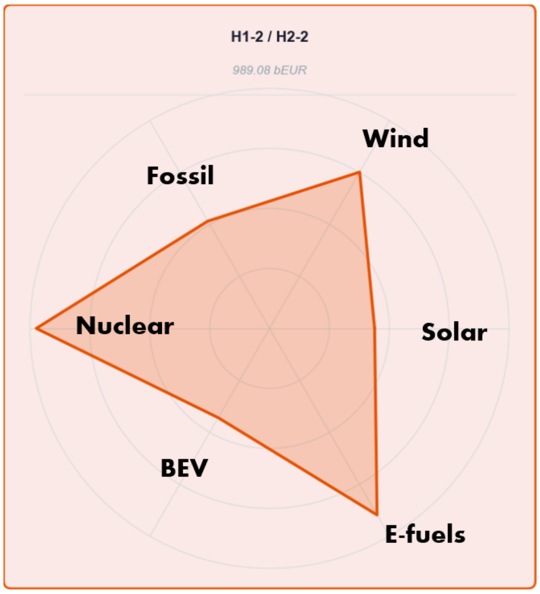

Figure 3.8: Radar profile of B8, the most fragile trajectory.

At the opposite end, the red cluster, B8 (98 pulls), B10 (107), B2 (116), and B5 (117), all fall well below the uniform baseline, their cumulative curves flattening progressively as the algorithm redirects evaluation effort towards the green group. B2 and B5, from H1-0 and H1-1 respectively, are structurally quasi-identical despite their different H1 branches: both commit to very low solar (~32 GW) combined with high nuclear and high e-fuel imports, producing the concentrated triangular radar profile of Figure 3.8. B8 (H1-2) and B10 (H1-3) share the same structural vulnerability, since low solar deployment leaves the system without a credible backup when nuclear is disrupted. UCB1 identified all four as underperforming early in the evaluation and progressively reduced their allocation, confirming a posteriori that the low-solar configuration, regardless of H1 origin, is consistently outperformed.

This visual separation is borne out statistically: every top-4 versus bottom-4 comparison of mean reward is significant at the 0.1% level, with non-overlapping confidence intervals between the two groups (full test statistics in Appendix B.5). For fragile arms the goal is to confirm underperformance rather than to characterise failure modes in detail, a task for which 98 pulls are entirely sufficient. A uniform allocation or a best-arm identification algorithm [24] would yield more balanced precision across all arms, which is a legitimate direction for future refinement; within the present computational budget, however, UCB1 is the most efficient instrument for the identification objective.

3.2.2 Feasibility Under Stress

Figure 3.9 presents the success rate of each trajectory, the fraction of evaluated scenarios under which the ESP model remains feasible, defined as the ability to satisfy the cumulative CO_2 budget constraint by 2050. Under the soft-constraint formulation adopted in this work, infeasibility arises exclusively when the combination of tightening carbon constraints and technology shocks renders the decarbonisation target physically unachievable: not because costs become prohibitive, but because the installed decarbonised capacity is insufficient to cover demand under the constraint.

Figure 3.9: Success rate (feasibility probability) of each trajectory under unexpected events, defined as the fraction of evaluated scenarios in which the ESP model meets the cumulative CO_2 budget constraint by 2050.

Five trajectories achieve a perfect success rate of 100%: B0, B1, B4, B7, and B9. The structural explanation is direct, as all five combine high solar deployment (∼70–75 GW) with high nuclear capacity (∼9–10 GW). Under a nuclear phase-out shock, which sets the nuclear capacity factor to zero while leaving installed capacity intact and thereby creating stranded assets, the 70–75 GW of solar provide enough alternative decarbonised production to keep the CO_2 budget within reach. At the energy-balance level modelled here, the redundancy between solar and nuclear is the operative protection mechanism: not a question of price, but of physical capacity available to decarbonise when one pillar fails.

The four trajectories with degraded success rates share the inverse structural profile. B2 (H1-0) and B5 (H1-1) reach 94% and 96.6% respectively: both commit to very low solar (∼35 GW) while relying on nuclear as their primary decarbonisation lever (∼9.3–9.8 GW). Under a nuclear shock, the residual 32 GW of solar is physically insufficient to compensate, and the CO_2 target becomes unachievable. B8 (H1-2, 91.8%) follows the same logic, compounded by the lowest nuclear deployment at 2035 in its H1 branch (2.85 GW), which reduces the cumulative decarbonised capital available before the shock materialises. B10 (H1-3, 91.6%) is a subtler case: despite its H1 origin, its H2 configuration reaches only 33 GW of solar, the same vulnerability as B2 and B5. The common thread is unambiguous. Trajectories that lean on a single decarbonisation pillar without a credible backup are structurally fragile, whatever their H1 origin.

3.2.3 Cost Under Unexpected Events

Figure 3.10 presents the arithmetic mean transition cost per trajectory over feasible scenarios under unexpected events, sorted from most robust to most fragile. The spread between the best (B4, +59% above C^*) and the worst (B5, +75%) amounts to 16 percentage points of cost overshoot, a genuine difference in how each trajectory absorbs disruption.

The top four trajectories (B4, B7, B1, B9) carry the same high-solar, high-nuclear signature identified in the pull analysis, and it is this diversified production base that limits cost overruns across a wide range of shock types. Among them, B4 (H1-1) is the most robust because it is the most balanced: e-fuel imports at their lowest in the group (50 TWh), BEV adoption at a moderate level (226 Mpkm/h), and no single dimension pushed to an extreme. B9 (H1-3) reaches comparable robustness through a different strategy, combining the highest nuclear deployment in the portfolio (10.18 GW) with the highest BEV adoption (324 Mpkm/h) to provide two structurally independent decarbonisation levers that compensate for one another when either is disrupted.

The bottom trajectories show the contrasting pattern. B5 (H1-1) and B8 (H1-2) are the worst performers (+75% and +74% respectively), sharing the same vulnerability: very low solar (∼33–35 GW) eliminates the renewable backup under a nuclear shock, while high e-fuel dependency (∼59–65 TWh) inflates costs under geopolitical supply disruptions. B2 (H1-0) shares the same fragility profile (+70%), confirming the structural redundancy between these three trajectories already noted in the archetype analysis. B10 (H1-3) is the most differentiated case within the

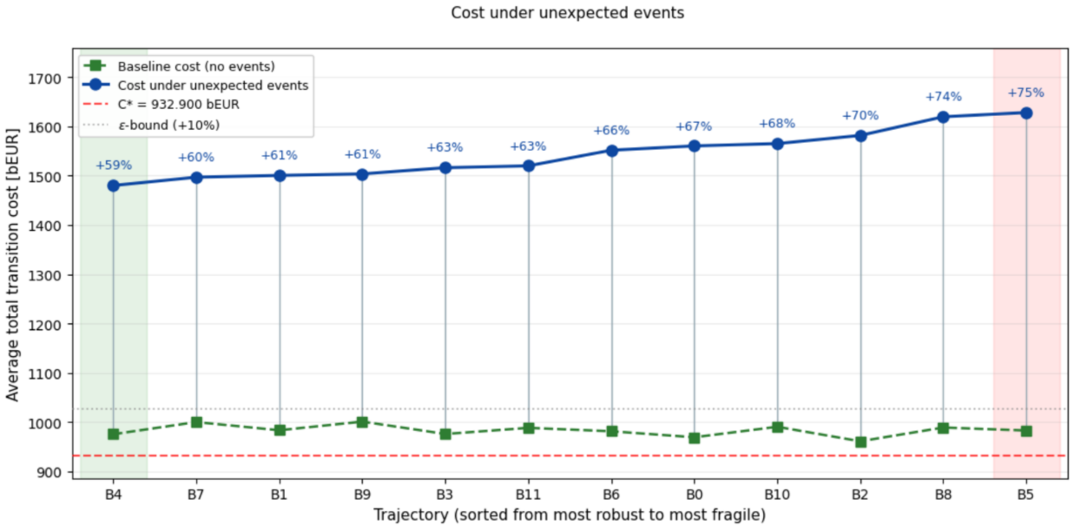

Figure 3.10: Arithmetic mean transition cost per trajectory over feasible scenarios under unexpected events (blue bars), compared to the baseline cost without events (green dashed line), sorted from most robust (left) to most fragile (right).

bottom group: its degradation relative to its own baseline (58%) is the lowest among the fragile trajectories, yet its mean reward ranks eleventh because its feasibility losses (SR = 91.6%) dominate the reward signal. This illustrates how cost degradation and reward diverge for trajectories where infeasibility, rather than cost overshoot, drives underperformance.

The worst-case dimension reinforces these conclusions. The CVaR_{10} on rewards, reported in Table B.6, places the top-4 trajectories between 0.494 and 0.512, while B2, B10, and B8 collapse to between 0.081 and 0.177, a more than fivefold gap. The asymmetry sharpens at the 5% tail (CVaR_5, Figure B.1): B2, B10, and B8 reach exactly zero, meaning their worst-case scenarios correspond to complete infeasibility rather than cost overruns, while the top-4 trajectories never reach zero regardless of the scenario drawn. The standard deviation of rewards tells the same story, with fragile arms exhibiting $\hat{\sigma}$ between 0.135 and 0.183, roughly twice the 0.067–0.080 range of the top-4: their performance is not only lower on average but substantially more erratic.

Two observations stand out before turning to policy implications. First, baseline cost and robustness are not correlated. B9, the most expensive trajectory without events (+7.3% above C^*), ranks fourth in robustness, while B2, the cheapest (+3.0%), ranks tenth by mean reward; optimising for baseline cost provides no guarantee of performance under disruption. Second, the robustness hierarchy is statistically established: confidence intervals on mean reward are non-overlapping between the top-4 group (lower bound ≥ 0.618) and the bottom-4 group (upper bound ≤ 0.590), and Welch t-tests confirm significance at the 0.1% level for all top-4 versus bottom-4 comparisons ($t \geq 3.94$, $p \leq 0.0001$) [27]. Within the mid-tier, by contrast, B3 versus B0 is not significant ($p = 0.245$), correctly reflecting their structural proximity. The agreement of four complementary signals, feasibility rate, conditional cost, performance stability, and worst-case CVaR, supplies the cross-validation basis for these conclusions (Appendix B.5).

3.3 Policy-Relevant Insights

The robustness evaluation of the preceding section was not designed to identify a single optimal transition pathway, but to give decision-makers a portfolio of credible, structurally diverse strategies whose relative merits can be weighed under deep uncertainty. This section distils the empirical findings into concrete guidance for the two decision horizons covered by the framework, examines the cost and composition of the resulting portfolio, and positions the results against the prior literature on uncertainty-aware energy planning.

3.3.1 Guidance Across the Two Decision Horizons

The first decision horizon concerns the energy mix to be committed by 2035. No single H1 branch dominates the robustness ranking: the four most robust trajectories are spread across all four H1 branches, with B1, B4, and B7 originating from H1-0, H1-1, and H1-2, and B9 from H1-3. This is a politically important finding, because it locates the decisive lever at the second horizon rather than the first: the robustness of the final portfolio is determined less by the 2035 decisions than by the 2040 ones. Decision-makers at H1 therefore retain genuine flexibility, and the framework identifies a range of near-optimal starting points, all compatible with robust long-term strategies. One H1 feature does matter, however. The H1-3 branch, distinguished by a very high BEV commitment (296 Mpkm/h) at 2035, produces markedly heterogeneous outcomes at H2: B9 ranks among the four most robust trajectories, whereas B10 falls into the fragile group, even though both inherit the same aggressive electrification at H1. An early and aggressive BEV deployment is thus no guarantee of robustness in itself; it narrows the subsequent solution space and leaves outcomes more sensitive to the 2040 decisions.

The 2040 horizon is where the robustness hierarchy is decided. B4, B7, and B1, from H1-1, H1-2, and H1-0 respectively, independently converge towards a high-solar, high-nuclear configuration at H2 and obtain near-identical robustness scores. That convergence is itself a strong signal: it suggests the configuration is not a local optimum specific to one trajectory, but a structural attractor of the near-optimal space under stress. Its advantage derives from physical redundancy, more precisely the complementarity of a variable renewable (solar) and a dispatchable low-carbon source (nuclear) that supplies two structurally independent decarbonised production levers at the energy-balance level. When one is disrupted, whether by a regulatory nuclear phase-out, a geopolitical supply shock, or a cost explosion in renewable technologies, the other remains available to maintain the decarbonisation trajectory, as confirmed by the 100% success rates of B4, B7, and B1. It should be acknowledged that the finding of H2 being more determinative than H1 is partly structural by design, since unexpected events are injected from 2040 onward (see Section 2.3.1). The result nonetheless remains actionable: given any 2035 commitment, the framework identifies which 2040 choices lead to robust outcomes.

A methodological caveat warrants explicit acknowledgement when interpreting the solar–nuclear redundancy finding. The ESP model operates at monthly temporal resolution, replacing the full hourly time series with twelve representative time-points per year and excluding fast-cycling storage technologies. Within this framework, the redundancy between solar and nuclear is established at the energy-balance level: 70–75 GW of solar capacity generates enough annual energy to compensate for a nuclear phase-out in terms of cumulative CO_2 budget compliance. Nuclear, however, provides dispatchable baseload generation, whereas solar is an inherently variable source subject to diurnal and seasonal intermittency that monthly resolution does not capture. Whether solar capacity of this magnitude could provide genuine firm operational backup for nuclear, once hourly dispatch constraints, winter production deficits, and the storage required to cover night-time and low-irradiance periods are accounted for, cannot be assessed within the current model. The robustness conclusions of this work therefore hold as energy-system-level findings under the ESP modelling assumptions; their operational translation would require validation at hourly resolution with explicit storage and grid-balancing constraints.

3.3.2 The Cost of Robustness and the Portfolio of Strategies

How much does robustness cost? The answer is reassuring. B4, the most robust trajectory, carries a baseline cost of only +4.6% above C^*, while B8, the most fragile, carries +6.0%, slightly more expensive without events yet dramatically less resilient under stress (+74% versus +59% above C^*). Within the near-optimal portfolio, choosing a robust trajectory over a fragile one carries no cost penalty: the most robust pathway is in fact cheaper than the most fragile, even though every near-optimal pathway sits a few percent above the strict cost-optimum. The benefit under disruption, by contrast, is substantial. The worst-case dimension reinforces this: under the 10% most adverse scenarios, the top-4 trajectories incur between 1.83 and 1.90 bEUR, against 1.92–2.13 bEUR for the bottom-4, a premium that materialises only when disruption strikes. These figures expose a cheap-but-fragile trap that planners should actively avoid, since the least expensive trajectory in the portfolio (B2, +3.0% above C^*) ranks tenth in robustness, whereas one of the most expensive (B9, +7.3%) ranks fourth with 100% feasibility. Optimising for baseline cost under assumed conditions is not merely insufficient; it can actively select for fragile strategies.

The central output of this framework is a portfolio of robust near-optimal strategies rather than a single recommended trajectory. The top-4 trajectories should not, however, be read as four radically different policy recipes: they are better understood as variants of a common robust family, all combining high solar deployment, high nuclear capacity, moderate wind, and relatively contained e-fuel imports. Their differences are second-order specifications within the same structural attractor rather than fundamentally distinct decarbonisation paradigms. B4 is the most balanced version, B7 places slightly more emphasis on solar and mobility electrification, B1 combines high nuclear capacity with moderate import dependence, and B9 relies more strongly on nuclear capacity and BEV adoption. The relevant policy insight is therefore not that decision-makers can freely choose among unrelated robust strategies, but that robustness is consistently associated with preserving multiple decarbonised production levers rather than relying on a single dominant pillar. Beyond the top-4, the portfolio also contains politically distinct alternatives whose vulnerabilities are explicitly identified. B3 (H1-1) and B6 (H1-2), which commit to roughly 4 GW of nuclear at H2 while maintaining high solar deployment (∼70–72 GW), rank fifth and seventh respectively, comfortably above the fragile group. Their vulnerability profile is structurally different: lacking nuclear as a firm decarbonised baseload, they compensate through higher e-fuel imports (∼63–65 TWh), which shifts their exposure from nuclear phase-out to geopolitical supply shocks. A decision-maker facing strong political opposition to nuclear can deliberately choose B3 or B6, not as a robustness optimum, but as a politically feasible compromise whose trade-offs are transparent. This is precisely the value of a portfolio over a single optimal solution: it confirms that all top-4 options are robust and backed by statistically significant evidence, while making the vulnerabilities of the remaining alternatives explicit.

The results condense into two operational principles for energy transition planning under deep uncertainty. First, build physical redundancy, not just cost efficiency: the defining weakness of the fragile trajectories is the absence of alternative decarbonised capacity when the primary technology fails. Early solar deployment acts as a feasibility insurance against nuclear phase-out scenarios at the energy-balance level, the two technologies being complementary in this role.

Second, avoid over-reliance on any single import-dependent technology: the fragile trajectories share a common dependence on imported e-fuels (∼60–65 TWh at H2), which amplifies cost overhead under geopolitical supply shocks and pushes worst-case costs above 2000 bEUR. Diversifying across domestic renewables, dispatchable low-carbon generation, and selective imports is the structural feature that separates robust from fragile strategies, regardless of the H1 decisions made in 2035. Applied to Belgium, these principles translate into a clear near-term priority: whatever the 2035 energy mix, the 2040 investment strategy should target a combined solar and nuclear deployment in the range of 70–75 GW and 9–10 GW respectively, while keeping e-fuel import dependency below roughly 55 TWh.

3.3.3 Relation to Prior Work

These findings broadly confirm and refine the conclusions of Kchaou et al. [17] and Coppitters et al. [7], while extending them to a multi-horizon setting. The portfolio diversification principle of Kchaou et al. [17] is borne out by B4, the most robust trajectory, which is precisely the most balanced, and by the general finding that robustness follows from preserving multiple decarbonised levers rather than any single one. The present multi-horizon analysis nonetheless adds three qualifications to their single-horizon results. Wind is not the dominant robustness signal at H2, where all trajectories converge towards intermediate values regardless of their H1 commitment, suggesting that wind matters more as an early-stage decision than as a 2040 lever. E-fuel imports exhibit a threshold effect absent from their work: moderate imports are beneficial, but trajectories exceeding ∼60 TWh turn fragile under geopolitical supply shocks. And aggressive BEV adoption, identified by Kchaou et al. [17] as a robust feature, proves insufficient on its own: B9 with high BEV (324 Mpkm/h) ranks fourth, whereas B10 with comparable BEV levels ranks eleventh by mean reward, so electrification translates into robustness only when paired with a complementary high-nuclear configuration.

The convergence with Coppitters et al. [7] is equally clear, although the comparison should be read with the caveat that the two studies do not share an identical definition and parameterisation of unexpected events. They identify the absence of electrofuel imports by 2050 as the principal dealbreaker and accelerated renewables as the most robust early-stage policy; both points align with the present results, the latter visible in the consistently high solar deployment (∼72–75 GW) of all top-4 trajectories. The nuclear-shock-without-solar-backup mechanism identified here is the multi-horizon equivalent of their dealbreaker concept, a physical infeasibility rather than a cost overrun, and the cost–robustness decoupling reported above directly parallels their expensive-but-robust accelerated-renewables policy.

Chapter 4

Conclusion & Future Work

This thesis addressed a central limitation of conventional energy transition planning: the tendency to rely on a single cost-optimal pathway despite the presence of deep uncertainty and unexpected events. The research question was therefore how a portfolio of structurally diverse and robust transition pathways can be identified when such disruptions are explicitly considered.

To answer this question, an integrated multi-horizon framework was developed and applied to the Belgian whole-energy system using EnergyScope Pathway. The framework combines Modelling to Generate Alternatives, k-means clustering and Multi-Armed Bandits. MGA is applied sequentially at two decision horizons, 2035 and 2040, to explore the near-optimal solution space and construct a decision tree of candidate trajectories. These trajectories are then evaluated under unexpected events using a UCB-based Multi-Armed Bandit algorithm, which directs the computational effort towards the most promising pathways.

The main methodological contribution of this work is the extension of the single-horizon MGA–MAB approach into a structured multi-horizon decision tree. This makes the sequential nature of long-term energy planning explicit: early decisions constrain future options, but do not fully determine them. The framework therefore provides a tractable way to compare not only individual energy mixes, but complete transition trajectories across successive decision stages.

Applied to Belgium, the results show that pathways with similar baseline costs can differ substantially in their robustness under unexpected events. The most robust trajectories are characterised by diversified low-carbon supply structures, in particular the combination of substantial solar and nuclear capacity, while more fragile trajectories tend to depend more strongly on externally constrained supply options such as e-fuel imports. The key message is therefore not that robustness is costless, but that the near-optimal space contains materially different robustness profiles. Energy planning should consequently assess portfolios of alternatives according to both cost and resilience criteria.

These findings must nevertheless be interpreted in light of important limitations. First, ESP is used here at monthly temporal resolution, which is sufficient to study long-term investment structures but not to validate operational feasibility at the hourly dispatch level. In particular, the solar–nuclear redundancy identified in this work should be tested with higher temporal resolution and explicit storage modelling. Second, the portfolio of 12 trajectories still contains some structural overlap, suggesting that the clustering procedure and the chosen near-optimality threshold may not fully capture the diversity of possible transition strategies.

Future work should therefore focus on four directions. First, the framework should be extended to hourly resolution with explicit storage and flexibility constraints. Second, a controlled comparison with the single-horizon framework of Kchaou et al. [17] would help isolate the added value of the multi-horizon structure. Third, a genuine recourse formulation, in which later decisions adapt to information revealed after earlier stages, would better represent sequential planning under uncertainty. Finally, the event space could be enriched with correlated and temporally structured shocks, and the framework could be applied to other national energy systems to test whether the Belgian conclusions generalise.

Overall, this thesis shows that robust energy transition planning should not be framed as the search for one optimal pathway, but as the comparison of several economically credible alternatives under disruption. In a context where the future is uncertain and infrastructure choices are difficult to reverse, preserving multiple independent low-carbon levers is not a secondary consideration; it is a necessary condition for credible long-term planning.

Appendix A

MAA–MAB Framework Details

This appendix documents the implementation details of the MAA–MAB pipeline: the numerical parameters (Section A.1), the MAA sampling procedure (Section A.2), the unexpected-event parameter space (Section A.3), the MAB evaluation algorithm (Section A.4), and the associated computational budget (Section A.5).

A.1 Framework Parameters

Table A.1 summarises the key numerical parameters of the MAA–MAB pipeline used in this thesis, and Table A.2 details the six MGA variable groups and their constituent technologies or resources.

Table A.1: Key parameters of the MAA–MAB framework.

Parameter	Symbol	Value	Role
MGA slack	ε	0.10	Near-optimal budget
MAA convergence threshold	$\varepsilon_{\text{conv}}$	0.05	Stopping criterion
MAA hard cap	—	5 iter.	Maximum iterations
Angle filter	$\theta_{\min}$	0.05 rad	Direction deduplication
Centroid tolerance	δ	0.01	$\pm 1\%$ constraint window
H1 clusters	K_{H1}	4	Branching at 2035
H2 clusters	K_{H2}	3	Branching at 2040
Leaf trajectories	K	12	MAB arms
Interior samples	N_{samp}	10 000	Delaunay sampling
UCB exploration param.	α	2	Confidence bound width [25]
MAB iterations	T	2 000	Total evaluation budget
Discount rate	i	1.5%	Annualisation (ESP) [5]

Table A.2: MGA variable groups and constituent technologies or resources.

#	Dimension	Unit	Components
1	Solar PV	GW	PV
2	Wind	GW	WIND_ONSHORE, WIND_OFFSHORE
3	Fossil imports	GWh	GAS, COAL, DIESEL
4	Nuclear	GW	NUCLEAR, NUCLEAR_SMR
5	BEV	Mpkm/h	CAR_BEV
6	E-fuel imports	GWh	H2_RE, METHANOL_RE, AMMONIA_RE

A.2 MAA Sampling Procedure

The MAA algorithm [15] iteratively expands the convex hull of the near-optimal space $\mathcal{W}$ by solving directional MGA subproblems. At each iteration, new search directions are derived as the outward face normals of the current convex hull. An angular filter of $\theta_{\min} = 0.05\,\text{rad}$ removes redundant directions that are too close to already-explored ones, and the total number of active directions is capped at 1 000 per iteration.

The algorithm terminates when the relative growth in convex hull volume between two consecutive iterations falls below 5%:

$$\varepsilon_{\text{conv}} = \frac{V^{(k)} - V^{(k-1)}}{V^{(k)}} < 0.05 \tag{A.1}$$

subject to a hard cap of five iterations. The convergence diagnostics for all horizons are reported in Appendix B.2.

Once convergence is reached, the convex hull of $\mathcal{W}$ is densely sampled using a Delaunay triangulation-based procedure [18, 20]. The hull is decomposed into simplices, and 10 000 interior points are drawn using randomly weighted barycentric coordinates:

$$\mathbf{x}_{\text{sample}} = \sum_{i=1}^{d+1} w_i\,\mathbf{v}_i, \qquad w_i \geq 0, \quad \sum_i w_i = 1 \tag{A.2}$$

where $\mathbf{v}_i$ are the vertices of a simplex. Weights are allocated proportionally to each simplex's volume relative to the total hull volume, ensuring uniform coverage of $\mathcal{W}$ without over-representing small simplices [19].

A.3 Unexpected Event Parameter Space

The full taxonomy of the 47 unexpected-event parameters is presented in Table A.3, following the categorisation of Coppitters et al. [7] and the adaptations described in Section 2.3.1.

A.4 MAB Evaluation Algorithm

Algorithm 1 provides the full pseudocode of the MAB robustness evaluation. The centroid constraints injected at each pull follow Equation (2.7). The unexpected-event parameterisation follows Section 2.3.1 and Appendix A.3. Full arm ranking results are provided in Table B.5.

Algorithm 1 Multi-horizon MAB robustness assessment (UCB1)

Require: transition tree with $K = 12$ leaf trajectories; Sobol library $\mathcal{L}$ of $N_s = 39{,}201$ scenarios; exploration parameter $\alpha = 2$; total iterations $T = 2{,}000$

Ensure: trajectories ranked by empirical mean reward $\hat{\mu}_k$

Phase 1 — Initialisation

for $k = 1$ **to** K **do**
  Draw ω_k uniformly from $\mathcal{L}$
  Inject centroid constraints for trajectory k (Eq. (2.7))
  Apply unexpected-event scenario ω_k to AMPL
  Solve ESP; record $C_k(\omega_k)$
  $r_k \leftarrow C^*/C_k(\omega_k)$ (reward = 0 if infeasible)
  $\hat{\mu}_k \leftarrow r_k; \quad N_k \leftarrow 1$
end for

Phase 2 — UCB exploration/exploitation

for $t = K + 1$ **to** T **do**
  $k_t \leftarrow \arg\max_k \left[\hat{\mu}_k + \sqrt{\frac{\alpha \ln t}{N_k}} \right]$
  Draw ω_t uniformly from $\mathcal{L}$
  Inject centroid constraints for trajectory k_t; apply ω_t
  Solve ESP; record $C_{k_t}(\omega_t)$
  $r_{k_t} \leftarrow C^*/C_{k_t}(\omega_t)$
  $N_{k_t} \leftarrow N_{k_t} + 1$
  $\hat{\mu}_{k_t} \leftarrow \hat{\mu}_{k_t} + \left(r_{k_t} - \hat{\mu}_{k_t} \right)/N_{k_t}$
end for
return $\{(k, \hat{\mu}_k, N_k)\}_{k=1}^{K}$ sorted by $\hat{\mu}_k$ descending

A.5 Computational Budget

Each MAB iteration requires solving a full multi-horizon ESP optimisation problem. The experiments were run on the author's personal computer, equipped with a

Table A.3: Summary of unexpected-event parameters, their AMPL targets, and maximum impact. All parameters are Sobol-sampled in $[0, 1]$ and discretised before injection. Shocks are applied exclusively over 2040–2050 (see Section 2.3.1).

Category	AMPL parameter	Mechanism	Max impact
Exchange rate (4 params)	c_{inv} (Asian tech)	Cumulative surcharge	+160% cumulative
PV supply constraints (4 params)	c_{inv}(PV)	Cost surcharge	×10 at 2040+
Wind supply constraints (8 params)	c_{inv}(WIND)	Cost surcharge	×10 at 2040+
E-fuel geopolitical tension (4 params)	c_{op} (e-fuels)	Operational cost multiplier	×10 at 2040+
Biofuel geopolitical tension (4 params)	c_{op} + avail (biofuels)	Dual cost + availability shock	×10 + avail → 0
Electricity import tension (4 params)	c_{op} + avail (elec.)	Dual cost + availability shock	×10 + avail → 0
Fossil price shock (4 params)	c_{op} (GAS, COAL, LFO)	Operational cost multiplier	×10 at 2040+
Energy demand growth (4 params)	end_demand	Cumulative multiplier	+60% cumulative
Renovation rate (4 params)	limit_renovation	Rate reduction	0.33 → 0
Unicorn technologies (6 params)	$f_{\max}$, c_{inv}	Availability gate ($f_{\max} \to 0$)	Full blockage
Nuclear phase-out (1 param)	c_p (NUCLEAR)	Capacity factor → 0	Full phase-out
Neutralised: SMR availability	—	Fixed at 0 (conflict with centroid)	—

Notes: All cost perturbations use $\beta = 9$ (factor ×10 at full shock intensity) applied uniformly over 2040–2050. The 2035 horizon is excluded from shock injection: the H1 centroid pins the six MGA dimensions at ±1%, making cost shocks there undiscriminating surcharges on fixed capacities rather than behavioural signals. Fossil price shocks use repurposed limit_mob_change parameters (see Section 2.3.1). Hard constraints are retained only for nuclear phase-out ($c_p \to 0$) and unicorn technology blockage ($f_{\max} \to 0$), which do not conflict with centroid constraints. Source: adapted from [7, 17].

13th Gen Intel(R) Core(TM) i5-1335U processor (10 cores, 12 logical processors) and 16 GB of RAM, using Gurobi through AMPL. One MAB iteration required approximately 22 seconds on average. The 2,000-iteration evaluation budget therefore corresponds to approximately 44,000 seconds, or 12.2 hours of total computation time. This computational cost motivated the use of an adaptive MAB procedure rather than an exhaustive evaluation of all trajectory–scenario combinations.

Appendix B

Supplementary Tables and Diagnostics

This appendix gathers the supplementary tables, diagnostics, and statistical analyses supporting the results of the main text: the cluster centroids and arm mapping (Section B.1), the MAA convergence diagnostics (Section B.2), the full MAB arm ranking (Section B.3), the cluster-selection criteria (Section B.4), and the statistical-significance and worst-case analysis (Section B.5).

B.1 Cluster Centroids of the Transition Pathway Tree

Table B.1 reports the full centroid values for all nodes of the transition pathway decision tree (MGA iterations = 5); costs are expressed in bEUR, and H2 nodes are conditioned on their parent H1 branch. Table B.2 gives the correspondence between the compact arm labels B0–B11 used throughout the analysis and their underlying (H1, H2) branch pairs.

B.2 MAA Convergence Diagnostics

Tables B.3 and B.4 report the MAA convergence diagnostics at H1 and H2. ε denotes the relative hull volume growth between successive iterations. Convergence requires $\varepsilon < 0.05$, subject to a hard cap of 5 iterations.

Table B.1: Cluster centroids for all nodes of the transition pathway tree (MGA iterations = 5). Costs in bEUR; capacity in GW except BEV (Mpkm/h) and resource variables (GWh).

Horizon	**Node**	**Cost [bEUR]**	**Solar [GW]**	**Wind [GW]**	**Fossil [GWh]**	**Nuclear [GW]**	**BEV [Mpkm/h]**	**E-Fuels [GWh]**
ROOT	ROOT	933	*unconstrained reference*					
H1 (2035)	H1-0	948	37.86	20.91	68 833	6.11	146.38	33 752
	H1-1	956	79.12	10.41	70 094	6.07	169.94	33 493
	H1-2	962	77.38	21.04	68 087	2.85	185.85	29 878
	H1-3	966	67.52	20.79	68 082	6.32	296.53	33 664
H2 | H1-0	H1-0 / H2-0	969	72.59	15.93	34 798	4.55	129.72	66 521
	H1-0 / H2-1	984	72.54	14.41	37 393	9.90	256.35	44 546
	H1-0 / H2-2	961	32.22	14.45	37 109	9.31	127.09	63 765
H2 | H1-1	H1-1 / H2-0	976	72.15	17.29	37 807	4.18	132.62	63 022
	H1-1 / H2-1	975	75.25	12.08	36 940	9.87	225.77	50 243
	H1-1 / H2-2	983	32.56	16.00	36 583	9.81	131.79	64 195
H2 | H1-2	H1-2 / H2-0	981	70.65	13.79	36 610	4.57	130.00	64 900
	H1-2 / H2-1	1000	75.65	15.47	40 240	9.19	281.28	58 632
	H1-2 / H2-2	989	34.69	15.82	36 172	9.92	139.24	59 700
H2 | H1-3	H1-3 / H2-0	1001	74.76	15.28	41 001	10.18	324.27	51 965
	H1-3 / H2-1	991	33.10	14.90	36 172	9.80	203.28	60 613
	H1-3 / H2-2	988	73.56	15.91	37 428	4.48	205.30	62 932

Table B.2: Correspondence between arm labels B0–B11 and their (H1, H2) branch pair.

Arm	**H1 branch**	**H2 branch**
B0	H1-0	H2-0
B1	H1-0	H2-1
B2	H1-0	H2-2
B3	H1-1	H2-0
B4	H1-1	H2-1
B5	H1-1	H2-2
B6	H1-2	H2-0
B7	H1-2	H2-1
B8	H1-2	H2-2
B9	H1-3	H2-0
B10	H1-3	H2-1
B11	H1-3	H2-2

Table B.3: MAA convergence diagnostics at H1 (2035).

Iteration	ε	**Vol**	**Status**
1	1.0000	0.001088	–
2	0.9680	0.034000	–
3	0.8634	0.248851	–
4	0.0997	0.276404	–
5	0.0339	0.286112	✓
2 999 MGA solutions → 10 000 interior points → 4 centroids			

✓ = convergence criterion met ($\varepsilon < 0.05$).

B.3 MAB Arm Ranking

Table B.5 reports the complete MAB–UCB arm ranking after 2 000 iterations ($\alpha = 2$, $K = 12$), sorted by mean reward.

B.4 Cluster Selection Diagnostics

The optimal number of clusters K^* is selected by evaluating two complementary criteria over $K \in \{2, \ldots, 14\}$ on the standardised 10 000-point MGA sample cloud. Before clustering, each MGA variable j is standardised:

$$\tilde{x}_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j} \tag{B.1}$$

and centroids are back-transformed via $\hat{\boldsymbol{\mu}}_c = \boldsymbol{\sigma} \odot \tilde{\boldsymbol{\mu}}_c + \boldsymbol{\mu}$.

The **elbow criterion** [22] identifies the point of diminishing marginal return in the within-cluster inertia:

$$W(K) = \sum_{c=1}^{K} \sum_{i \in \mathcal{C}_c} \|\tilde{\mathbf{x}}_i - \tilde{\boldsymbol{\mu}}_c\|_2^2 \tag{B.2}$$

The **silhouette score** [23] measures cluster cohesion $a(i)$ and separation $b(i)$ jointly:

$$s(i) = \frac{b(i) - a(i)}{\max(a(i),\, b(i))} \in [-1,\, 1] \tag{B.3}$$

with K^* maximising the average $\bar{s}(K)$. Both criteria use `sklearn.cluster.KMeans` with `n_init=10` and a fixed random seed; diagnostic figures are shown in Section 3.1.

Table B.4: MAA convergence diagnostics at H2 (2040), per H1 branch.

Branch	**Iter.**	ε	**Vol**	**Status**
H1-0	1	1.0000	0.001100	–
	2	0.9940	0.183635	–
	3	0.6083	0.468803	–
	4	0.0888	0.514492	–
	5	0.0846	0.562048	†
H1-1	1	1.0000	0.000983	–
	2	0.9943	0.171424	–
	3	0.5530	0.383478	–
	4	0.1607	0.456897	–
	5	0.1007	0.508038	†
H1-2	1	1.0000	0.000948	–
	2	0.9939	0.154897	–
	3	0.6224	0.410202	–
	4	0.1181	0.465148	–
	5	0.0737	0.502170	†
H1-3	1	1.0000	0.000828	–
	2	0.9923	0.107331	–
	3	0.6745	0.329719	–
	4	0.1429	0.384688	–
	5	0.0248	0.394468	✓

MGA solutions per branch: H1-0: 3 100, H1-1: 3 097, H1-2: 3 090, H1-3: 3 089
Each branch: 10 000 interior points → 3 centroids

✓ = convergence criterion met. † = hard cap reached without convergence ($\varepsilon \geq 0.05$). Only H1-3 converges strictly; H1-0, H1-1, and H1-2 terminate at the hard cap with final ε of 0.0846, 0.1007, and 0.0737 respectively, indicating that the conditional near-optimal space at H2 is geometrically more complex than at H1.

Table B.5: MAB–UCB arm ranking after 2 000 iterations ($\alpha = 2$, $K = 12$), sorted by mean reward. *Cost events* is the arithmetic mean of per-pull transition cost over feasible scenarios only.

Arm	Pulls	Mean reward	Success rate	Baseline	Cost events	Degradation [%]
B4	236	0.641	1.000	975	1480	51.8
B7	212	0.632	1.000	1000	1497	49.7
B1	209	0.630	1.000	984	1500	52.5
B9	204	0.628	1.000	1001	1503	50.2
B3	185	0.618	0.989	976	1516	55.3
B11	186	0.618	0.989	988	1520	53.8
B0	167	0.607	1.000	969	1560	61.0
B6	163	0.605	0.994	981	1551	58.2
B5	117	0.565	0.966	983	1628	65.6
B2	116	0.565	0.940	961	1581	64.6
B10	107	0.554	0.916	991	1565	57.9
B8	98	0.537	0.918	989	1619	63.7

Degradation: The top-4 trajectories are stable across both rankings; mid- and bottom-tier ordering shifts because degradation excludes the infeasibility penalty embedded in the reward, confirming that feasibility dominates the bottom-group separation

B.5 Statistical Significance and Worst-Case Analysis

Table B.6 extends the arm ranking with confidence intervals, standard deviations, and CVaR estimates. Table B.7 reports Welch t-tests for selected pairwise comparisons, and Figure B.1 presents the CVaR profiles visually [26].

On the choice of Welch's t-test. Welch's t-test is preferred over Student's because reward distributions are markedly heteroscedastic ($\hat{\sigma}_{B4} = 0.080$ vs $\hat{\sigma}_{B8} = 0.175$), causing Student's test to yield biased Type I error control under unequal variances. With $n \geq 98$ pulls per arm, the Central Limit Theorem ensures approximate normality of sample means. Samples are unpaired since the MAB allocates different Sobol scenarios to each arm; the test therefore captures both structural robustness differences and scenario variability. A methodological caveat applies: under adaptive UCB1 allocation, future pulls depend on previously observed rewards, so the i.i.d. assumption underlying standard confidence intervals and Welch's test is not strictly satisfied. As a result, reported p-values should be interpreted as indicative rather than fully confirmatory, since adaptive sampling may slightly exaggerate the separation between stronger and weaker arms. Given the magnitude of all top-4 vs bottom-4 statistics ($t \geq 3.94$, $p \leq 0.0001$), the minimum sample size of 98 pulls, and the corroborating CVaR, success-rate, and cost-degradation results, this does not affect the qualitative conclusions [27].

Table B.6: 95% confidence intervals, standard deviations, and CVaR estimates for all 12 trajectories after 2 000 MAB iterations, sorted by descending mean reward.

		Reward (all pulls, incl. infeasible)				*Cost | feasible*	
Arm	n	Mean	95% CI	Std	CVaR_{10}	Mean [bEUR]	CVaR_{10} [bEUR]
B4	236	0.641	±0.010	0.080	0.494	1480	1904
B7	212	0.632	±0.010	0.072	0.508	1497	1855
B1	209	0.630	±0.010	0.072	0.503	1500	1862
B9	204	0.628	±0.009	0.067	0.512	1503	1830
B3	185	0.618	±0.015	0.101	0.438	1516	1898
B11	186	0.618	±0.014	0.098	0.437	1520	1942
B0	167	0.607	±0.012	0.076	0.488	1560	1918
B6	163	0.605	±0.013	0.084	0.472	1551	1857
B5	117	0.565	±0.024	0.135	0.285	1628	2130
B2	116	0.565	±0.030	0.164	0.177	1581	2073
B10	107	0.554	±0.035	0.183	0.081	1565	1922
B8	98	0.537	±0.035	0.175	0.081	1619	2081

Mean reward: unconditional, includes zero-reward infeasible runs (Eq. 2.9). *$CVaR_{10}$ reward*: mean reward over the worst 10% of pulls; zero indicates complete infeasibility in the tail. *Cost mean and $CVaR_{10}$*: arithmetic mean and worst-10% mean cost, conditional on feasible runs only. *SR*: success rate. 95% CI $= \pm 1.96\,\hat{\sigma}/\sqrt{n}$. Top-4 lower bound ($\geq 0.618$) and bottom-4 upper bound (≤ 0.590) are non-overlapping.

Table B.7: Welch t-tests for selected pairwise comparisons of mean reward [27].

Comparison		n_1	n_2	t	p	
B4 vs B8	(best vs worst)	236	98	5.643	<0.0001	***
B9 vs B5	(top-4 vs bottom-4, boundary)	204	117	4.689	<0.0001	***
B9 vs B2	(top-4 vs bottom-4)	204	116	3.943	0.0001	***
B9 vs B10	(top-4 vs bottom-4)	204	107	4.029	0.0001	***
B3 vs B0	(mid-tier)	185	167	1.163	0.245	n.s.

*** $p < 0.001$; n.s. $p > 0.05$. The non-significance of B3 vs B0 reflects genuine structural proximity.

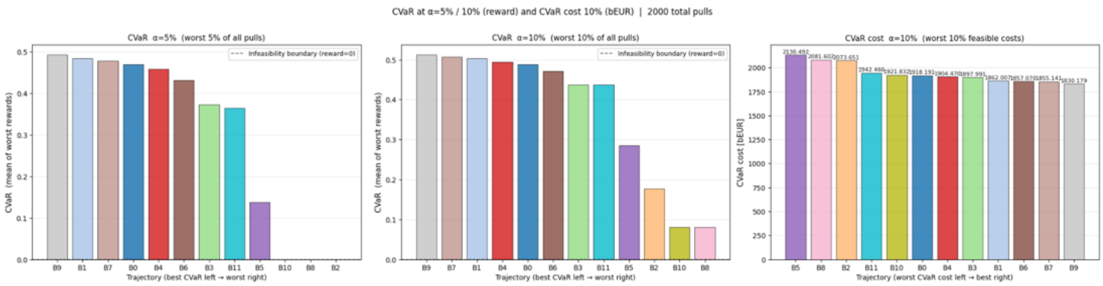

Figure B.1: Conditional Value at Risk (CVaR) profiles for all 12 candidate trajectories (B0–B11) after 2,000 MAB iterations, sorted by descending mean reward (most robust on the left). CVaR at level α measures the mean reward over the worst α% of pulls — a lower CVaR indicates worse performance in extreme scenarios. *Left*: CVaR_5 (mean reward over the worst 5% of pulls; reward$= C^*/C_{\text{scenario}}$ for feasible scenarios, 0 otherwise). *Centre*: CVaR_{10} (mean reward over the worst 10% of pulls). *Right*: CVaR on transition cost at $\alpha = 10\%$ (mean cost over the 10% most expensive feasible scenarios, in bEUR). B2, B8, and B10 reach $\mathrm{CVaR}_5 = 0$, indicating that their worst-case scenarios correspond to complete infeasibility rather than mere cost overruns. The top-4 trajectories (B4, B7, B1, B9) maintain strictly positive CVaR values across both tail levels, confirming robust worst-case performance.

Appendix C

Use of Generative AI Tools

Generative AI tools were used during the preparation of this thesis to support code development, report writing, and visualisation design.

Claude Code was used to assist with the generation, restructuring, and debugging of parts of the computational framework, starting from an initial code base provided by Mahdi Kchaou. The resulting implementation was reviewed, adapted, tested, and validated by the author before being used to produce the numerical results.

Claude was used to support the drafting, reformulation, and clarification of several sections of the report. The methodological choices, numerical results, interpretation, and final selection of the content remain the responsibility of the author.

Gemini was used to support the visual design of the figure illustrating the K-means clustering of near-optimal MGA samples (Figure 2.5) based on the methodological source by Pedersen et al. [15]. Generative AI tools were also used to suggest plot layouts and visual styles. The final figures and plots were selected and adapted by the author according to their relevance for communicating the methodology and results.

All cited sources were independently verified from their original publications. The author remains fully responsible for the accuracy, validity, and interpretation of the submitted work.

Bibliography

[1] IPCC. IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland. Technical report, Intergovernmental Panel on Climate Change (IPCC), July 2023. Edition: First.

[2] International Energy Agency. *Net Zero by 2050: A Roadmap for the Global Energy Sector*. OECD Publishing, 2021. Google-Books-ID: xBqszgEACAAJ.

[3] Nicholas Stern. *Why Are We Waiting?: The Logic, Urgency, and Promise of Tackling Climate Change*. The MIT Press, April 2015.

[4] United Nations. Paris agreement, 2015. United Nations Framework Convention on Climate Change.

[5] Gauthier Limpens, Xavier Rixhon, Francesco Contino, and Hervé Jeanmart. EnergyScope Pathway: An open-source model to optimise the energy transition pathways of a regional whole-energy system. *Applied Energy*, 358:122501, March 2024.

[6] Miguel Chang, Jakob Zink Thellufsen, Behnam Zakeri, Bryn Pickering, Stefan Pfenninger, Henrik Lund, and Poul Alberg Østergaard. Trends in tools and approaches for modelling the energy transition. *Applied Energy*, 290:116731, May 2021.

[7] Diederik Coppitters, Gabriel Wiest, Leonard Göke, Francesco Contino, André Bardow, and Stefano Moret. Identifying Dealbreakers and Robust Policies for the Energy Transition Amid Unexpected Events, February 2025. arXiv:2503.05735 [cs.CY].

[8] Katharina Esser, Jonas Finke, Valentin Bertsch, and Andreas Löschel. Participatory modelling to generate alternatives to support decision-makers with near-optimal decarbonisation options. *Applied Energy*, 395:126184, October 2025.

[9] Bart-Jaap Verbeek. Compensation for stranded assets?, April 2021.

[10] Esnil Guevara, Frédéric Babonneau, and Tito Homem-de Mello. Uncertainty dynamics in energy planning models: An autoregressive and Markov chain modeling approach. *Computers & Industrial Engineering*, 191:110084, May 2024.

[11] Stefano Moret, Frédéric Babonneau, Michel Bierlaire, and François Maréchal. Decision support for strategic energy planning: A robust optimization framework. *European Journal of Operational Research*, 280(2):539–554, January 2020.

[12] Stefano Moret, Víctor Codina Gironès, Michel Bierlaire, and François Maréchal. Characterization of input uncertainties in strategic energy planning models. *Applied Energy*, 202:597–617, September 2017.

[13] Tao Tan, Meng Yang, Rui Xie, Yuji Cao, and Yue Chen. Adjustable robust optimization with decision-dependent uncertainty for power system problems: A review. *Energy Conversion and Economics*, May 2025.

[14] Xavier RIXHON. *Exploration of uncertainty-aware energy transition pathways.* PhD Thesis, TU Delft, 2024.

[15] Tim T. Pedersen, Marta Victoria, Morten G. Rasmussen, and Gorm B. Andresen. Modeling all alternative solutions for highly renewable energy systems. *Energy*, 234:121294, November 2021.

[16] Mahdi Kchaou, Francesco Contino, and Diederik Coppitters. Revealing design archetypes and flexibility in e-molecule import pathways using Modeling to Generate Alternatives and interpretable machine learning, December 2025. arXiv:2512.23284 [eess.SY].

[17] Mahdi Kchaou, Diederik Coppitters, and Francesco Contino. Exploring robust early-stage decisions in energy transitions using near-optimal pathways and multi-armed bandits. In *Proceedings of the 36th European Symposium on Computer Aided Process Engineering (ESCAPE-36)*, Sheffield, UK, June 2026. Elsevier. Peer-reviewed conference proceeding, 21–24 June 2026.

[18] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)*, 22(4):469–483, December 1996.

[19] Robert L. Smith. Efficient Monte Carlo Procedures for Generating Points Uniformly Distributed Over Bounded Regions. *Operations Research*, 32(6):1296–1308, 1984.

[20] D. T. Lee and B. J. Schachter. Two algorithms for constructing a Delaunay triangulation. *International Journal of Computer & Information Sciences*, 9(3):219–242, June 1980.

[21] Abiodun M. Ikotun, Absalom E. Ezugwu, Laith Abualigah, Belal Abuhaija, and Jia Heming. K-means clustering algorithms: A comprehensive review, variants analysis, and advances in the era of big data. *Information Sciences*, 622:178–210, April 2023.

[22] Mengyao Cui. Introduction to the K-Means Clustering Algorithm Based on the Elbow Method. *Accounting, Auditing and Finance*, 1(1):5–8, October 2020.

[23] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, November 1987.

[24] Aleksandrs Slivkins. Introduction to Multi-Armed Bandits, April 2024. arXiv:1904.07272 [cs.LG].

[25] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2):235–256, May 2002.

[26] Alexander J. McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools.* Princeton University Press, Princeton, NJ, revised edition, 2015.

[27] Zhenqiu Lu and Ke-Hai Yuan. Welch's t test. pages 1620–1623. January 2010.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain
Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/epl**