

Exploring Robust Early-Stage Decisions in Energy Transitions Using Near-Optimal Pathways and Multi-Armed Bandits

Mahdi Kchaou^{a*}, Diederik Coppitters^a, Francesco Contino^a

^a Institute of Mechanics, Materials and Civil Engineering (iMMC), Université catholique de Louvain (UCLouvain), Louvain-la-Neuve, Belgium

* Corresponding Author: mahdi.kchaou@uclouvain.be.

ABSTRACT

Although rare, unexpected events such as financial crises, geopolitical conflicts, and pandemics have reshaped reality in recent years. Despite their strong potential to affect the energy transition, such events are still largely overlooked in energy planning studies. Ignoring them can lead to poorly informed decisions that may jeopardize the transition. Identifying early-stage decisions that remain robust under unexpected events is therefore essential. To address this challenge, EnergyScope Pathway, a whole-energy system model with limited foresight, is applied to Belgium. To increase the likelihood of a successful transition, the Modeling to Generate Alternatives approach is used to diversify early-stage decisions in 2035. These alternatives are allowed to be up to 10% more expensive than the cost-optimal solution. However, the large number of alternative designs is difficult to navigate for decision makers. To address this, a decision-support framework based on the Multi-Armed Bandit framework is used to identify early-stage decisions that are most robust to future unexpected events. In this step, the remaining transition phases are optimized under unexpected events sampled within predefined impact ranges. The results show that, under normal conditions, there is a high degree of flexibility in the decision space for the 2030–2035 phase, with many technologies or resources that can be entirely omitted. However, robust early-stage decisions rely on a diverse energy generation portfolio, with a stronger emphasis on wind deployment, early mobility shifts toward battery electric vehicles, and the import of e-fuels. These insights can help decision makers steer the energy transition toward a robust path from the beginning. While Belgium is used as a case study, this framework is transferable to other contexts.

Keywords: Energy transition, Modeling to generate alternatives, Multi-armed bandit, Decision-making, Unexpected events

1. INTRODUCTION

The need for an energy transition has become increasingly urgent due to the impacts of climate change on the economy, food security, and human health [1]. To guide this transition, energy system modeling has emerged as a key tool for shaping policies and investment decisions. Despite the large number of studies on this topic, progress in the energy transition has been hindered by hesitation and inaction, largely driven by the many uncertainties surrounding energy projects [2]. Many existing energy planning studies ignore uncertainty altogether [3], making their solutions highly sensitive to

inaccuracies in future conditions and ultimately prone to failure. A striking example is the collapse of General Electric's thermal power business, which lost approximately \$193 billion between 2016 and 2018 due to incorrect assumptions favoring gas over renewables [4]. In the rare cases where uncertainty is considered, most models rely on local sensitivity analysis or probability-based approaches [5]. While useful, these methods fail to capture deeper and more complex forms of uncertainty, particularly those that are rare in nature but have significant consequences: unexpected events. These are largely mispredicted or unpredicted occurrences that deviate strongly from expectations and persist long enough to

influence strategic decisions [6]. Such events have already affected the energy sector in the past and are likely to do so again [7, 8]. They may arise from technological breakthroughs, supply-chain disruptions, global pandemics, geopolitical conflicts, or societal and political shifts. The COVID-19 pandemic, for instance, reshaped the global energy sector by altering fuel demand, accelerating the deployment of renewables, and highlighting vulnerabilities in global supply chains [9].

Despite their importance, unexpected events are poorly represented in energy planning tools [6]. Their low probability and the scarcity of data needed to characterize them make them difficult to incorporate into standard uncertainty analysis methods.

This paper addresses these challenges by proposing a novel framework to identify early-stage energy transition decisions that perform well under unexpected events.

The whole-energy, regional energy model EnergyScope Pathway [10] is applied to Belgium using a multistage formulation with limited foresight, ensuring that events occurring in later stages are not anticipated from the beginning but instead arise as surprises. The model optimizes the energy transition by minimizing costs, subject to constraints that ensure a carbon-neutral Belgium by 2050. The decision space is then expanded using the Modeling to Generate Alternatives (MGA) approach [11], which is applied to the early transition phase from 2030 to 2035 to generate a diverse set of near-optimal early-stage decisions. These decisions, such as not importing e-fuels or not relying on nuclear energy, increase the likelihood of a successful transition by providing more diverse starting conditions. Afterwards, these decisions are evaluated under unexpected events occurring in the remaining stages of the transition, from 2035 to 2050. Finally, a decision-making framework based on the Multi-Armed Bandit (MAB) algorithm [12] is used to identify the most robust early-stage decisions in 2035. Although the Belgian energy transition is used as a case study, the proposed framework is highly transferable and can be applied to other contexts, both within the energy sector and in other planning applications.

In this paper, we first present the EnergyScope Pathway model, the MGA approach, the generation of unexpected events, and the MAB approach (Section 2). The results section (Section 3) begins by illustrating the diversity of the near-optimal space, followed by the selection of representative early-stage decisions. The robust early-stage decisions are then identified using the MAB approach. Section 4 concludes the paper by summarizing the main takeaways.

2. METHODS

2.1 EnergyScope Pathway

The energy model used in this work is EnergyScope Pathway [10], a regional, whole-energy system model applied to the Belgian case. It covers most energy sectors, including electricity, heating, transport, and non-energy demand. It includes six demand categories and 112 technologies covering energy supply, storage and conversion, as well as import and export interactions with neighbouring countries. The model is formulated as a linear multistage optimization problem, with the objective of minimizing the total transition cost. In the pathway (multistage) formulation, the energy transition is divided into multiple phases, each with its own parameters and variables. The results of one phase are passed to the next as initial conditions, allowing the model to capture the evolution of investment decisions over time. The model can be used either with an hourly resolution, using a typical days approach [13], or with a monthly resolution, where each month is represented by a single time point. While the monthly formulation offers faster solution times, it is less detailed, as renewable intermittency and hourly demand variations are not explicitly captured. As a result, some technologies, such as storage technologies, are not included.

EnergyScope Pathway supports both perfect foresight and myopic foresight formulations. In the perfect foresight case, the entire transition is optimized in a single run, assuming full knowledge of future conditions, and decisions for all phases are optimized simultaneously. A global CO₂ emissions budget is imposed, allowing any decarbonization pathway as long as the final emissions target is met by the end of the transition. In the myopic foresight formulation, each phase of the transition is optimized on its own, sequentially, without knowledge of future conditions. Since each phase is optimized separately, a global CO₂ emissions constraint cannot be applied. Instead, a specific CO₂ emissions constraint is imposed for each one. These constraints follow an exponential decay formulation [14]:

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

where e_0 (124 MtCO₂/yr) is the initial emission level, r is the initial slope parameter (set to zero), and m is the decay parameter. The parameter m is calibrated such that the trapezoidal approximation of emissions, evaluated in five-year steps between 2020 and 2050, matches the cumulative emissions budget 1.5 GtCO₂ [15]. The resulting emissions levels at each five-year node are then used as constraints in the myopic optimization. In this work, the monthly and myopic options are chosen. Additional details on the model can be found in [6, 10].

2.2 Early stage decisions

EnergyScope Pathway optimizes the energy transition by minimizing costs, delivering a single cost-optimal solution. While informative, this approach is not enough,

as it lacks the diversity of solutions required for decision makers with their different preferences. Moreover, relying on a single solution is risky, as it may not perform well under structural uncertainties, particularly high-impact ones such as unexpected events. For instance, under normal conditions, a maximum deployment level may be assumed for a given technology. The cost-optimal pathway will then size its PV deployment to this maximum in the early stages of the transition. Under unexpected events, however, this maximum deployment could be reduced by half, leaving the optimal solution unable to cope with such situation. To address this limitation, we rely on the Modeling to Generate Alternatives (MGA) [11] approach to explore the near-optimal space of early-stage decisions for the transition phase 2030–2035. MGA allows a small increase (slack) in cost in exchange for greater solution diversity. In this work, a 10% slack (cost increase) is allowed during the 2030–2035 phase relative to the cost of that phase, and all solutions satisfying this condition are identified as near-optimal. As a result, by 2035, rather than a single optimal energy system, a set of diverse energy systems with near-equivalent economic performance is identified. Considering such alternative solutions can lead to better performance when unexpected events occur. These alternatives may include, for example, energy systems with lower reliance on PV, reduced fossil fuels imports to accelerate early decarbonization, or higher e-fuels imports during the initial transition stages. These choices, among many others, provide policymakers with a broader decision space, rather than locking decisions into a single pathway. The focus on the 2030–2035 period is motivated by the need to provide actionable guidance for decisions that need to be taken today or in the near future, helping to place the energy transition on the right track from the beginning.

Another key element in defining the MGA problem is the choice of the MGA variables. These variables are allowed to vary within the feasible space to explore the near-optimal solutions. Ideally, the MGA variables would include all investment decisions. However, due to computational limitations, only a limited number of variables can be considered. In this work, the selected MGA variables are: battery electric vehicles (BEV) [Mpkm/h], nuclear capacity [GW], PV capacity [GW], wind capacity (the sum of onshore and offshore wind) [GW], e-fuels imports (the sum of imported renewable methanol, ammonia, and hydrogen) [GWh], and fossil resource imports (the sum of imported coal, diesel, and gas) [GWh]. These variables correspond to the strategic decisions that governments and policymakers can directly influence. Using the MGA approach, the near-optimal space is defined as:

$$W = \{ x \in X \mid f(x) \leq f(x^*) (1 + \varepsilon) \} \quad (2)$$

where X is the feasible solution space, ε is the slack

parameter, $f(x)$ denotes the objective function, and x^* is the optimal solution. In this study, $f(x)$ corresponds to the cumulative cost of the transition up to 2035, and ε is set to 10%. As a result, all solutions that are at most 10% more expensive than the minimum-cost solution are accepted. The exploration of the near-optimal space follows the Modeling All Alternatives (MAA) approach developed by Pedersen et al. [16]. This method maps the entire near-optimal space by defining its convex hull, which is the smallest polyhedron containing all near-optimal solutions. Since the original problem formulation is linear, the near-optimal space is guaranteed to form a polyhedron. The method starts with an initial set of near-optimal solutions, from which a first convex hull is constructed using the Quickhull algorithm. In the next iteration, new near-optimal points are generated to increase the volume of the convex hull. Once the convex hull is fully defined, it is then sampled since all points within it are valid near-optimal solutions. After generating a large set of near-optimal solutions, using all of them as early-stage decisions becomes computationally expensive. Therefore, a representative subset is selected using the k-means clustering approach, which groups solutions with similar characteristics. The centroid of each cluster is chosen as a representative solution, and the number of clusters is determined using the elbow method [17].

2.3 Unexpected events

Following the work of Coppitters et al. [6], six categories of unexpected events are identified. The first category concerns the availability of imported resources, including electrofuels, biofuels, and electricity. The second category is resistance to change, which includes PV and wind deployment, low-temperature heating renovations, and private mobility changes. The third category relates to exchange-rate deterioration, and the fourth to increases in energy demand. The fifth category concerns unicorn technologies, including small modular reactors (SMR), carbon capture and storage (CCS), geothermal power, ammonia-fired combined cycle gas turbines (CCGT), and ammonia cracking. The final category is the nuclear phase-out. For the first four categories, each subcategory is characterized by four parameters, each corresponding to a specific year of impact: 2035, 2040, 2045, and 2050. Each parameter can take one of six values: 0, 20, 40, 60, 80, or 100, where 0 indicates no impact and 100 represents full impact in the corresponding year. For the fifth category, each subcategory is characterized by a single parameter indicating when the technology becomes available: 2040, 2045, 2050, or after that. Similarly, the nuclear phase-out is represented by a single parameter indicating when the phase-out occurs: 2035, 2040, 2045, 2050, or after that.

From this parameter space, 1000 unexpected-event scenarios are sampled using Sobol sampling to ensure

uniform coverage of the space. For further details on the selection, parametrization, and generation of unexpected events, the reader is referred to [6].

2.4 Multi-armed bandit

After generating a set of early-stage decisions for the year 2035 using MGA (Subsection 2.2), the objective is to identify the most robust ones of these decisions under the unexpected events defined in Subsection 2.3. Policymakers can take one of these decisions in the near future. However, evaluating all possible combinations of early-stage decisions and unexpected event scenarios is computationally expensive, even for simplified models such as the monthly model used in this work. To address this computational challenge, a reinforcement learning approach based on the Multi-Armed Bandit (MAB) algorithm is used [12].

The MAB problem can be illustrated by a gambler playing slot machines with multiple arms, each returning random rewards. At each time step, the gambler pulls an arm and gets a reward. Based on the history of rewards, the MAB algorithm suggests which arm to pull next in order to maximize the expected reward. In the context of this work, each early-stage decision for the year 2035 represents an arm, and the reward corresponds to the final outcome of the energy transition. Each time a specific early-stage decision is selected, the energy system in 2035 is fixed according to the decision, and the remainder of the transition is optimized under a randomly selected unexpected event scenario drawn from the 1000 generated scenarios. This is where the uncertainty in the reward arises. By balancing exploration and exploitation, the MAB approach identifies the most robust early-stage decisions without the need to evaluate all possible scenario–decision combinations.

In this MAB problem, the Upper Confidence Bound (UCB) policy [12] is used. The policy defines the logic behind the selection of the next arm at each time step. Let K denote the number of arms. At each time step t , an arm $a_t \in \{1, \dots, K\}$ is pulled, resulting in a random reward r_t . The UCB policy defines an objective that combines exploitation, through the use of empirical mean of the previous rewards, and exploration, through an explicit uncertainty term that accounts for how often each arm has been pulled. This objective is maximized to select the next arm a_{t+1} to pull.

The reward is typically defined within the interval $[0, 1]$. For that, the reward in this work is defined as follows:

$$r_t = \begin{cases} \frac{C_{\text{optimal}}}{C_{\text{scenario}}}, & \text{if the CO}_2 \text{ target is met by 2050} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where C_{optimal} is the minimum (optimal) cost, under normal conditions and C_{scenario} is the cost of the scenario defined by the combination of an unexpected event scenario and the selected early-stage decision.

A summary of the methodology is illustrated in Figure 1.

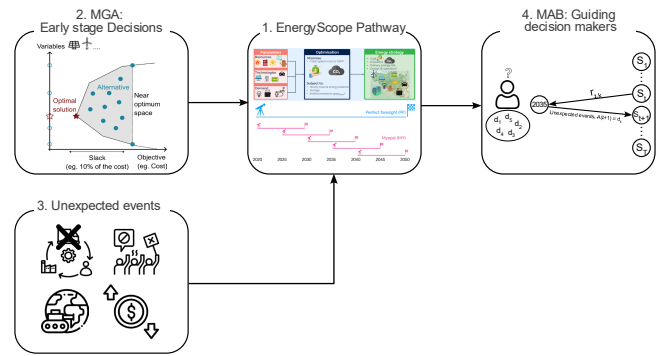


Figure 1: The methodology starts by setting up the energy transition model (1). Then, a set of alternative early-stage decisions for the year 2035 is generated (2), along with a set of unexpected-event scenarios (3). A MAB approach (4) is used to test the early-stage decisions against these unexpected events and identify the most robust ones.

3. RESULTS

3.1 Near-optimal space

The optimal early-stage decision, as well as the near-optimal decisions for the year 2035, are presented in Figure 2. The figure shows the probability mass functions (PMFs) of the near-optimal space together with the optimal values, indicated by red stars, for the MGA variables: BEV, nuclear capacity, PV capacity, wind capacity, e-fuels imports, and fossil resource imports.

The optimal solution is characterized by high nuclear and wind deployment, and fossil resource imports, medium level of integration for BEV and low PV deployment and e-fuels imports, as seen in the range of the near-optimal space. The near-optimal alternatives themselves span a wide range of values, indicating a high degree of flexibility and the existence of solutions that are structurally very different from the optimal one. For instance, most variables can reach zero or near-zero values, meaning that their corresponding technologies or resources can be entirely omitted (for example, some near-optimal systems have no wind deployment or almost no fossil fuels imports) while remaining at most 10% more expensive than the optimal system by 2035. However, the near-optimal solutions are not evenly distributed across this wide range. Looking at the PMFs of wind, PV, and nuclear capacity, distributions are skewed toward higher values, suggesting that higher deployment levels are more frequently favored among near-optimal solutions. This tendency does not always coincide with the optimal value, as seen for PV capacity. In contrast, the BEV PMF is normally distributed around the optimal value, indicating that the optimal BEV level is the most

Probability Mass Function of Near-Optimal Solutions

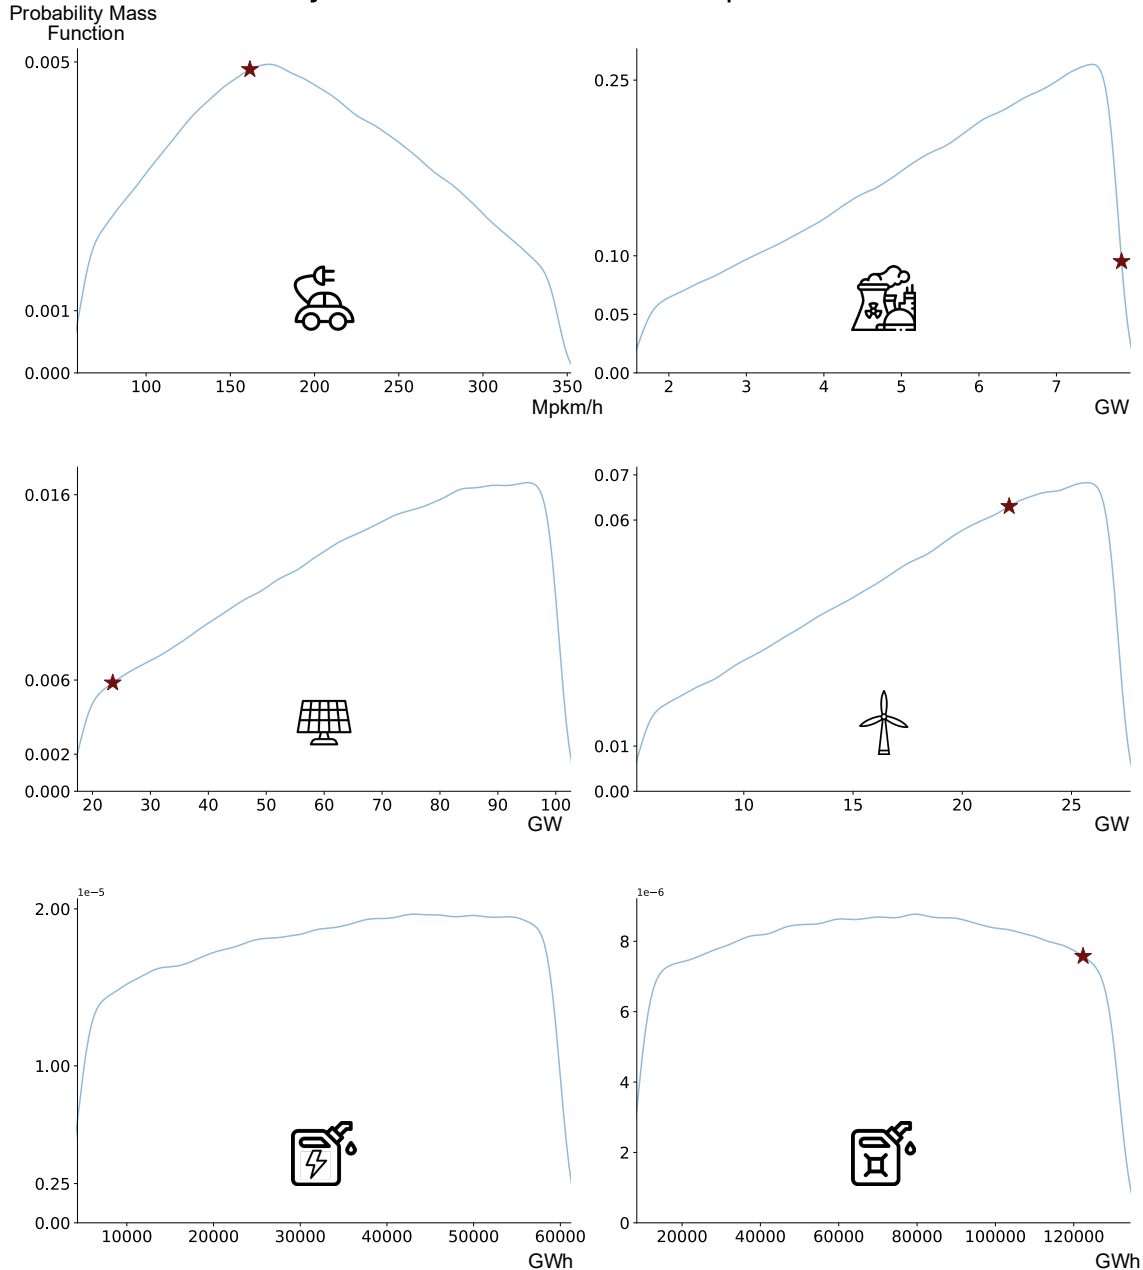


Figure 2. The probability mass function plots of the near-optimal space show a wide variety of feasible solutions. Despite this large space, solutions are not evenly distributed across the full range of values. PV, wind, and nuclear deployment exhibit strong skewness toward higher values. PV, wind, and nuclear deployment exhibit strong skewness toward higher values. BEV is normally distributed around its optimal value (indicated by the red star), whereas fuels imports are more uniformly distributed across the near-optimal space.

favored option. Finally, for fuels imports, both distributions are close to uniform, with a slight skew toward higher values for e-fuels imports, indicating that near-optimal solutions are evenly spread across these dimensions of the decision space. At this stage, it should be kept in mind that this flexibility concerns only early-stage decisions specific to the year 2035 of the transition.

3.2 Early stage decisions selection

The MAA method generates thousands of near-optimal early-stage decisions for the year 2035. However, retaining the full set of solutions to study their robustness would be computationally expensive, as discussed in Section 2.2. To address this, the near-optimal space is reduced using representative solutions, selected as the

centroids of clusters obtained through k-means clustering. The elbow method suggested dividing the space into 5 to 10 clusters: for better interpretability, the space was finally clustered into 5 clusters.

Figure 3 illustrates the resulting heatmap of early-stage decisions. All MGA variables are normalized to the [0, 1] interval, where a value of 0 represents the absence of the technology, while 1 corresponds to its maximum level. The optimal solution is shown in the first row, followed by the five cluster centroids that characterize the near-optimal solution space. Across all early-stage decisions, high deployment of wind, solar, and nuclear technologies is a consistent feature, with at least two of these technologies reaching high levels in every case. BEV integration is generally moderate, with two exceptions: a low level in decision 2 and a slightly higher level in decision 5. Fossil fuels imports are similar across decisions, with mostly mid-level imports and higher imports in the optimal solution. In contrast, e-fuels imports are generally moderate, except in the optimal solution, where they are comparatively low.

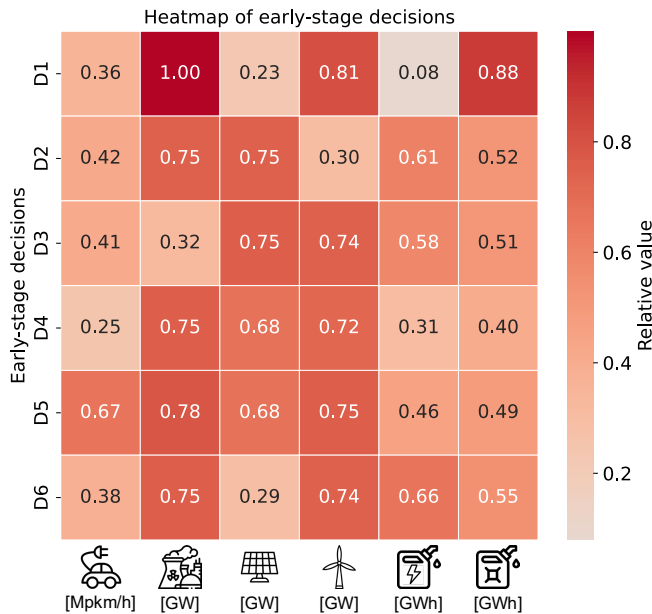


Figure 3. Heatmap of relative MGA variable values ([0, 1]) for the optimal solution and the five cluster centroids representing the near-optimal space. The figure shows consistent high deployment of wind, solar, and nuclear technologies across early-stage decisions, while BEV integration remains mostly moderate. Fossil fuels imports are present in all cases, with higher levels in the optimal solution, whereas e-fuels imports are generally moderate and lower in the optimal case.

The near-optimal early-stage decisions (decisions 2 to 6 in Figure 3) are then evaluated under normal conditions, without unexpected events. In 2035, the energy system is fixed according to each of these early-stage

decisions, and the remainder of the transition is optimized under a cost-minimization objective. All of the decisions lead to a successful transition, in the sense that the 2050 CO₂ target is met in all cases. However, the associated total transition costs (all of the transition, up until 2050) differ. As shown in Figure 4, all decisions result in higher total costs than the total optimal solution (red dashed line). Nevertheless, total costs remain within the 10% slack of the total optimal solution (grey dashed line): below 110% of the minimum total transition cost. In terms of specific values, decisions 2, 3, and 4 have similar total costs, around 960 b€. Decision 5 leads to the highest total cost (≈ 968 b€), while decision 6 leads to the lowest total cost (≈ 943 b€).




Figure 4. Total transition costs associated with the near-optimal early-stage decisions, compared to the optimal solution (red dashed line), under normal conditions. All near-optimal decisions result in higher total costs than the optimal case, while remaining below 110% of the minimum total transition cost (grey dashed line). Decision 2, 3 and 4 have similar costs, while decision 5 has the highest cost and decision 6 has the lowest cost.

3.3 Robust early stage decisions

The multi-armed bandit (MAB) method is used to identify the most robust early-stage decisions. The algorithm is run for 1000 time steps. At each step, a decision is selected (pulled) based on the UCB policy, and a random unexpected events scenario is applied starting from 2035. Figure 5 shows the cumulative number of pulls for each decision over time. Decision 6 clearly stands out, as it is selected most often, followed by decisions 3 and 5. In contrast, decisions 1 (the optimal decision for the 2030-2035 phase), 4, and 2 are selected least frequently. Similar conclusions can be drawn from the evolution of the average reward per decision, shown in Figure 6. Decision 6 reaches the highest average reward (≈ 0.49) (Eq. 3), followed by decisions 3 and 5 with average rewards of about 0.41. The weakest performing decisions are decisions 1, 4, and 2, with average rewards of approximately 0.35, 0.28, and 0.20, respectively.

One important observation is that all early-stage decisions share a common feature: they all have a mixed energy generation portfolio. However, what separates

well-performing decisions from poorly performing ones is the specific combination of technologies and resources that is chosen. When comparing decisions 2 and 3, both share similar characteristics, but decision 2 relies more on nuclear power, while decision 3 relies more on wind power. The better performance of decision 3 suggests that wind deployment provides greater robustness under uncertainty than relying mainly on nuclear power. Nuclear capacity can be entirely phased out during the transition, meaning that all previously installed nuclear capacity may eventually be dismantled. In contrast, even under resistance to further wind deployment, wind capacity installed earlier in the transition continues to be used.

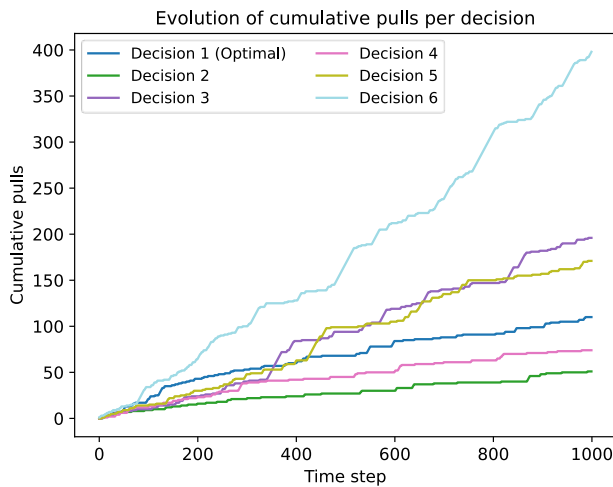


Figure 5. Cumulative number of pulls per decision over time. Decision 6 is selected most often, indicating the highest robustness, followed by decisions 5 and 3. Decisions 1, 4, and 2 are selected least frequently.

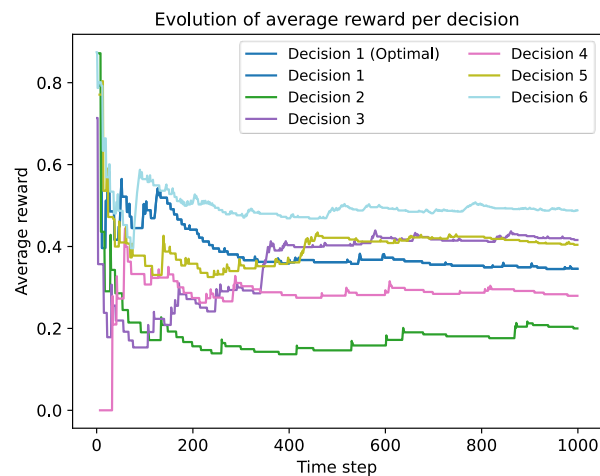


Figure 6. Evolution of the average reward per decision over time. Similar to the previous figure, decision 6 performs best, achieving the highest average reward over most time steps, followed by decisions 5 and 3. Decisions 1, 4, and 2 have the lowest final average rewards.

A similar pattern appears when comparing decisions 4 and 5. Although both have a diverse generation mix, decision 5 performs better because it includes higher BEV integration and higher e-fuels imports. This shows that having a mixed generation portfolio alone is not sufficient: early changes in mobility and use of e-fuels also play an important role.

Finally, decisions 1 and 6 also share some similarities, as both invest in the same technologies and resources, with the main difference being that decision 1 does not invest in e-fuels. Despite this similarity, their performance differs significantly. Decision 1 invests very heavily in a few technologies, with almost full nuclear and wind deployment and high fossil fuels imports. Decision 6 also invests strongly in nuclear, wind, and fossil fuels, but at more moderate levels, and additionally relies on e-fuels imports. This more balanced approach leads to better performance under uncertainty.

Overall, the results show that a diverse energy mix and strong wind deployment are essential, but they are not enough by themselves. Robust early-stage decisions also require early investment in e-fuels and mobility changes toward BEVs.

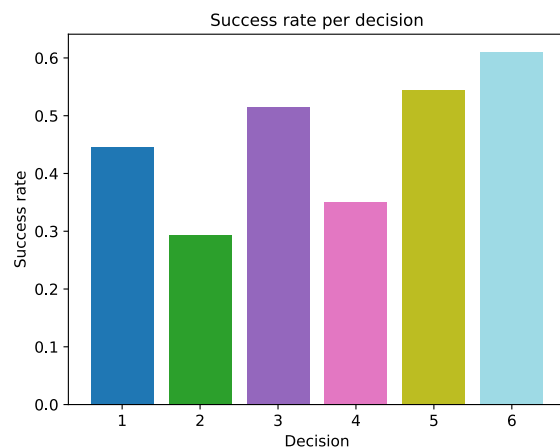


Figure 7. Success rate of each early-stage decision, defined as the ratio of successful trials meeting the 2050 CO₂ target to the total number of times each decision is selected. Decision 6 shows the highest success rate, followed by decisions 5, 3, 1, 4, and 2.

Digging deeper into the analysis of the results, the success rate is computed for each decision. The success rate is defined as the ratio of successful trials that meet the 2050 CO₂ target to the total number of times each decision is selected (Figure 7). In addition, the average total transition cost of the successful trials is calculated (Figure 8).

Consistent with previous results, decision 6 achieves the highest success rate at 61%, followed by decision 5 at 54%, decision 3 at 51%, decision 1 at 44%,

decision 4 at 35%, and decision 2 at 29%. When looking at the average total transition costs, decision 3 has the lowest cost at 1156 b€, followed by decision 6 at 1166 b€, decision 4 at 1172 b€, decision 1 at 1202 b€, decision 5 at 1255 b€, and decision 2 at 1372 b€. An important observation is that although decision 4 has one of the lowest average costs, it lacks robustness compared to decision 5, which has a higher average cost but performs better in terms of robustness. The same pattern appears when comparing decision 5 with decisions 1 and 3. Even though decisions 1 and 3 have lower average costs, they are less robust. This confirms that selecting the least costly option is not always the best strategy. Choosing alternatives that may be more expensive but ensure a robust transition leads to better outcomes.

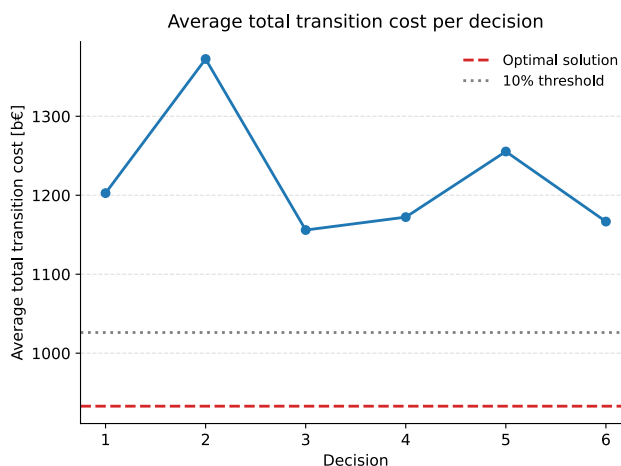


Figure 8. Average total transition cost of the successful trials for each early-stage decision. Although decision 3 exhibits the lowest average cost, it does not rank among the most robust options. Decision 6 combines relatively low costs with high robustness, while decision 5 shows higher average costs but stronger robustness compared to decisions 1, 3, and 4. This highlights that lower costs do not necessarily imply better robustness.

4 CONCLUSION

In this work, we explored the near-optimal space of early-stage decisions for the energy transition, focusing on the year 2035, and identified the most robust decisions under a broad range of unexpected events. To do so, we proposed a new decision-making framework that combines Modeling to Generate Alternatives (MGA) with a Multi-Armed Bandit (MAB) approach. This framework addresses the limitations of relying on a single cost-optimal pathway, which may perform poorly when unexpected events occur. In addition, the proposed approach reduces the computational effort required to identify robust early-stage transition strategies. Using the MGA approach on BEV adoption, nuclear, PV, and wind

deployment, as well as e-fuels imports and fossil fuels imports, the near-optimal space reveals a wide range of solutions that differ structurally while remaining close in cost. For the period 2030–2035, almost any technology or resource among the MGA variables can be entirely omitted without exceeding 110% of the minimum cost for that phase. However, near-optimal solutions are not evenly distributed across the solution space: solar, wind, and nuclear deployment are skewed toward high levels, BEV adoption is centered around the optimal value, and e-fuels and fossil fuels imports are more uniformly distributed. After selecting five representative early-stage decisions from the large near-optimal space using k-means clustering, the MAB analysis shows that robust early-stage decisions are characterized by diverse energy generation portfolios, with a strong emphasis on wind deployment. In addition, early mobility changes toward BEVs and the use of e-fuels imports play an important role in improving robustness. These insights can support decision makers facing highly uncertain and unpredictable futures by highlighting actions that can be taken today to place the energy transition on a more robust path from the start.

Future work will define early-stage decisions as sets of alternatives rather than single solutions (entire clusters instead of centroids). This would better reflect real-world decision-making, where policies often guide the system toward general directions rather than fixed technology levels.

ACKNOWLEDGEMENTS

This work was supported by the Wallonie-Bruxelles International - WBI (M.K) and the Fonds de la Recherche Scientifique - FNRS [CR 40016260] (D.C).

REFERENCES

1. 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.
2. Stern N. Why are we waiting?. The MIT Press (2015). <https://doi.org/10.7551/mitpress/10408.001.0001>
3. Yue X, Pye S, DeCarolis J, Li FGN, Rogan F, Gallachóir BÓ. A review of approaches to uncertainty assessment in energy system optimization models. *Energy Strategy Reviews* 21:204–217 (2018). <https://doi.org/10.1016/j.esr.2018.06.003>
4. General Electric misread the energy transition: A cautionary tale. URL

- <https://ieefa.org/resources/general-electricmisread-energy-transition-cautionary-tale> (accessed May 28, 2025).
5. Mavromatidis G, Orehounig K, Carmeliet J. A review of uncertainty characterisation approaches for the optimal design of distributed energy systems. *Renewable and Sustainable Energy Reviews* 88:258-277 (2018).
<https://doi.org/10.1016/j.rser.2018.02.021>
 6. Coppitters D, Wiest G, Göke L, et al. Identifying Dealbreakers and Robust Policies for the Energy Transition Amid Unexpected Events.
 7. McCollum DL, Gambhir A, Rogelj J, Wilson C. Energy modellers should explore extremes more systematically in scenarios. *Nat Energy* 5:104-107 (2020). <https://doi.org/10.1038/s41560-020-0555-3>
 8. Heinonen S, Karjalainen J, Ruotsalainen J, Steinmüller K. Surprise as the new normal – implications for energy security. *Eur J Futures Res* 5: (2017). <https://doi.org/10.1007/s40309-017-0117-5>
 9. World Energy Council, Netherlands. About Black Swans and Green Futures (2021).
<https://publications.tno.nl/publication/34638137/0FQCYm/WEC-2021-about.pdf>
 10. Limpens G, Rixhon X, Contino F, Jeanmart H. Energyscope pathway: an open-source model to optimise the energy transition pathways of a regional whole-energy system. *Applied Energy* 358:122501 (2024).
<https://doi.org/10.1016/j.apenergy.2023.122501>
 11. Brill ED Jr. The use of optimization models in public-sector planning. *Management Science* 25:413-422 (1979).
<https://doi.org/10.1287/mnsc.25.5.413>
 12. Lai TL, Robbins H. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6:4-22 (1985). [https://doi.org/10.1016/0196-8858\(85\)90002-8](https://doi.org/10.1016/0196-8858(85)90002-8)
 13. Thiran P, Jeanmart H, Contino F. Validation of a method to select a priori the number of typical days for energy system optimisation models. *Energies* 16:2772 (2023).
<https://doi.org/10.3390/en16062772>
 14. Victoria M, Zeyen E, Brown T. Speed of technological transformations required in europe to achieve different climate goals. *Joule* 6:1066-1086 (2022). <https://doi.org/10.1016/j.joule.2022.04.016>
 15. Rixhon X. Exploration of Uncertainty-Aware Energy Transition Pathways, Ph.D. thesis, Université Catholique de Louvain (2024).
 16. Pedersen TT, Victoria M, Rasmussen MG, Andresen GB. Modeling all alternative solutions for highly renewable energy systems. *Energy* 234:121294 (2021).
<https://doi.org/10.1016/j.energy.2021.121294>
 17. Ketchen DJ Jr, Shook CL. The application of cluster analysis in strategic management research: An analysis and critique. *Strat Mgmt J* 17:441-458 (1996). [https://doi.org/10.1002/\(sici\)1097-0266\(199606\)17:6<441::aid-smj819>3.0.co;2-g](https://doi.org/10.1002/(sici)1097-0266(199606)17:6<441::aid-smj819>3.0.co;2-g)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See <https://creativecommons.org/licenses/by-sa/4.0/>

