

Article

A Hybrid NLP-SWOT and Economic Modeling Framework for Sustainable Hydrogen Policy: Assessing Türkiye's Clean Energy Transition and LCOH Projections

İlker Mert ¹ , Hüseyin Yağlı ² , Jorge Costa ^{3,4}  and Ana Paula Oliveira ^{3,5,*} 

¹ Osmaniye Vocational School, Osmaniye Korkut Ata University, 80010 Osmaniye, Türkiye; ilkermert@osmaniye.edu.tr

² Department of Mechanical Engineering, Engineering Faculty, Gaziantep University, 27310 Gaziantep, Türkiye; yagli@gantep.edu.tr

³ School of Management, Engineering and Aeronautics, Instituto Superior de Educação e Ciências de Lisboa (ISEC Lisboa), Alameda das Linhas de Torres, 179, 1750-142 Lisboa, Portugal; jorge.costa@iseclisboa.pt

⁴ METRICs, Departamento de Química, NOVA School of Science and Technology (NOVA FCT), Universidade NOVA de Lisboa, Campus de Caparica, 2829-516 Caparica, Portugal

⁵ MARE—Centro de Ciências do Mar e do Ambiente, Instituto Politécnico de Setúbal (MARE-IPSetúbal), Campus do IPS, Estefanilha, 2910-765 Setúbal, Portugal

* Correspondence: ana.oliveira@iseclisboa.pt

Abstract

Sustainable hydrogen policy and national strategy documents are rich in qualitative information whose systematic evaluation still relies largely on subjective SWOT frameworks. This study proposes a reproducible hybrid methodology that couples expert-supervised Natural Language Processing (NLP) with a stochastic techno-economic model of the Levelized Cost of Hydrogen (LCOH) to convert policy discourse into quantitative, evidence-based recommendations. TF-IDF vectorization, K-Means clustering, Shannon entropy and Correspondence Analysis (CA) are applied to a manually annotated corpus of 107 sentences drawn from Türkiye's national hydrogen strategy documents (Cohen's $\kappa = 0.81$, substantial agreement). CA positions Regulation/Legislation and Financing near the Weakness quadrant, Renewable Resource Potential in the Strength quadrant, and Export/Demand Risk near Opportunity—revealing structural bottlenecks that challenge the sustainable energy transition. These qualitative findings are subjected to a quantitative consistency check via a Monte Carlo simulation ($N = 10,000$ iterations) propagating joint uncertainty in CAPEX, electricity price, electrolyzer efficiency, annual operating hours, discount rate and plant lifetime. The deterministic 2025 LCOH baseline of 4.89 €/kg H_2 carries a P10–P90 interval of [3.95; 5.92] €/kg. Global Sobol sensitivity analysis identifies electricity price as the dominant driver ($S_1 \approx 0.52$), suggesting that financing is discursively surfaced by the textual layer. Under business-as-usual technological learning, the probability of reaching a globally competitive LCOH (≤ 2 €/kg H_2) by 2050 is only 17.8%; a stylized proactive policy intervention (carbon pricing + subsidies) raises this probability to 78.4%. The framework is adaptable to other countries and languages (though the current implementation is Turkish-specific), providing a scalable, open methodology for evidence-based sustainable clean energy planning.

Keywords: sustainable energy transition; green hydrogen policy; decarbonization pathways; clean energy economics; policy-driven sustainability



Academic Editor: Francesco Ferella

Received: 20 June 2026

Revised: 18 July 2026

Accepted: 20 July 2026

Published: 23 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

1. Introduction

The global transition to low-carbon energy is shifting hydrogen from a technological question to a policy-design question, in which national strategies, certification schemes and financing instruments determine whether announced projects reach a Final Investment Decision. Türkiye exemplifies this tension. Its 2053 net-zero vision and 2023 National Hydrogen Strategy and Roadmap define ambitious production targets [1,2]. Meanwhile, domestic analyzes by SHURA and IICEC repeatedly emphasize regulatory and financing gaps as the main structural bottlenecks [3–5].

At the same time, thanks to the country's outstanding capabilities in terms of its solar and wind power potential as well as its proximity to the European hydrogen market, the state could become a competitive exporter of hydrogen [6]. Importantly, the success and sustainability of this hydrogen economy are highly dependent on the implementation of the downstream application segments. Among those, the PEM Fuel Cells (PEMFCs) are regarded as one of the important pillars in this domain with high prospects for large-scale power plants and fuel cell electric vehicles (FCEVs). The key to ensuring the economical performance of these applications lies in the optimal design of fuel cell stacks, which is achieved by flow field control and mass transfer augmentation.

But policy texts are notorious for being hard to analyze systematically. They combine aspirations for the long run with the terminology of law, are produced by different agencies, and are incompatible with the rigid four-square logic of classical SWOT analysis, which has long been criticized for its subjective character and lack of scalability [7,8]. The difficulties common for any text analysis, such as ambiguity, problems with data quality, and an unstructured nature [9,10], become especially significant for the analysis of specialized and sensitive materials, such as policy texts. Despite the intrinsic drawbacks of classical SWOT analysis, modern AI-enabled SWOT techniques show some promise in overcoming these difficulties [11]. Recent advances in computational policy analysis have demonstrated that Natural Language Processing (NLP) techniques, from Term Frequency-Inverse Document Frequency (TF-IDF) and K-Means clustering to contextual embeddings, can extract latent thematic structures from large policy corpora and reduce the dependence on analyst judgment [12–15]. Despite this progress, three methodological gaps persist in the hydrogen-policy literature:

- Existing SWOT applications to hydrogen strategies—including those for the European Union (EU) [16] and for Türkiye [17,18]—still rely almost exclusively on manual classification, with no quantitative validation of thematic structure.
- NLP-based policy analyzes rarely close the loop with techno-economic modeling, so that textual findings on discursively identified ‘financing weaknesses’ or ‘regulatory gaps’ are never confronted with the actual cost of hydrogen production under those policy conditions.
- Uncertainty in Levelized Cost of Hydrogen (LCOH) projections—central to investor bankability assessments—is typically reported as deterministic point estimates, ignoring the joint variability of Capital Expenditure (CAPEX), electricity price and electrolyzer performance [19,20].

This study addresses all three gaps simultaneously by proposing a hybrid framework that:

- (i) extracts SWOT sub-themes from Turkish hydrogen-policy documents by combining TF-IDF vectorization and K-Means clustering with expert supervision;
- (ii) characterizes the thematic structure using Shannon entropy [21], hierarchical clustering and Correspondence Analysis (CA) [22];

- (iii) provides a diagnostic alignment of the qualitative findings with the qualitative findings against a stochastic LCOH model whose uncertainty is propagated through Monte Carlo simulation with global Sobol sensitivity analysis.

Türkiye is selected as a strategic testbed because of its combination of strong renewable resources, geographic access to European hydrogen demand, and a still-emerging regulatory framework, a configuration representative of several Mediterranean and MENA economies. Even cross-country studies such as the QUAD nations analysis remain limited to manual SWOT classification [23].

Constituted as a methodological proof-of-concept, this study utilizes the entire published national hydrogen strategy corpus of Türkiye to demonstrate how a transferable, hybrid NLP-techno-economic framework can bridge qualitative policy architectures with quantitative cost stochastics. By executing a comprehensive textual census rather than an arbitrary data sample, the framework establishes a reliable foundation for capturing state-level strategic intent.

Crucially, this is operationalized not as an automated, endogenous mathematical coupling, but as an interpretable, human-in-the-loop diagnostic framework. The qualitative layer identifies institutional vulnerabilities and strategic priorities, which are subsequently translated via expert elicitation into localized variance and distribution bounds within a Monte Carlo LCOH model.

By bridging these qualitative semantic and techno-economic cost layers, this study directly advances sustainability across three dimensions:

- (i) Environmentally: It models the policy mechanics required to displace fossil-fuel reliance in hard-to-abate sectors through green hydrogen, directly supporting net-zero emissions targets.
- (ii) Economically: By integrating probabilistic LCOH projections under uncertainty to replace deterministic point estimates, it evaluates the financial durability of clean technology investments for Türkiye under business-as-usual and proactive policy scenarios, ensuring long-term economic viability.
- (iii) Methodologically: The proposed hybrid NLP-SWOT and Monte Carlo framework provides a highly reproducible, transferable tool—supported by complete equations, parameter distributions, and illustrative code excerpts—for other developing nations to evaluate and optimize their sustainable transition policies without human cognitive bias.

2. Materials and Methods

The methodological pipeline comprises four sequential stages: corpus construction, hybrid thematic extraction, statistical characterization, and stochastic LCOH modeling. Each stage is described below and summarized in the workflow diagram (Figure 1).

2.1. Corpus Construction and Annotation

The textual layer of the framework utilizes a comprehensive national hydrogen policy corpus for Türkiye. While the original text indicated a different count due to background literature citations, the finalized text-mining corpus comprises exactly nine distinct, authoritative national-level roadmap and strategy documents published between 2020 and 2025 [2–6,24–27]. To clarify the operational scope, “national-level texts” are defined as official ministerial strategies alongside high-impact national roadmaps published by state-affiliated research councils and key industry–academic clusters operating under direct institutional mandates to shape Turkish hydrogen policy. Documents originating from independent foreign research centers or secondary academic literature [28,29] were utilized strictly as contextual framing references within the manuscript and were completely ex-

cluded from the text-mining pipeline. To ensure absolute transparency and reproducibility, the definitive structural composition of this 107-sentence text corpus is cataloged in Table 1.




Figure 1. Hybrid text-mining workflow for hydrogen policy SWOT analysis. Solid arrows indicate the core pipeline; dashed lines indicate human-in-the-loop validation. Reported diagnostics: $\kappa = 0.81$ (annotation), $H = 2.692$ (entropy).

Table 1. Structural inventory and sentencing breakdown of the Turkish national hydrogen policy corpus.

Document/Title Reference	Issuing Institution	Year	Document Type	Total Sentences	Retained Sentences
National Hydrogen Technology Strategy [2]	Ministry of Energy & Natural Resources	2023	Official State Policy	142	44
Hydrogen Technologies Roadmap [3]	TÜBİTAK (MAM)	2021	State Research Council	98	21
Sectoral Carbon Neutral Roadmap [4]	Ministry of Industry and Technology	2023	Official State Policy	64	12
Green Reconciliation Action Plan [5]	Ministry of Trade	2021	Official State Policy	45	9
Development Plan (2024–2028) [6]	Presidency Strategy & Budget Office	2023	Supreme State Plan	38	7
Hydrogen Economy & Strategy [24]	SHURA Energy Transition Center	2021	Co-National Think-Tank	52	6
Natural Gas Strategic Vision [25]	GAZBİR (Gas Distributors Association)	2022	Mandated Industry Roadmap	31	4
Renewable Energy Action Plan [26]	Ministry of Energy & Natural Resources	2014	Official State Policy	28	2
National Energy Plan (2020–2035) [27]	Ministry of Energy & Natural Resources	2022	Official State Policy	41	2
Total Corpus Scope	9 Active Documents	—	—		

Texts were systematically extracted using pdfplumber and python-docx libraries, normalized to standard UTF-8 encoding, and segmented at the sentence level utilizing the Natural Language Toolkit (NLTK) environment. This extraction initially yielded a raw pool of 539 sentences across the nine documents. To isolate actionable policy directives from generic text, a strict sentence-selection screening was independently executed by two domain experts specializing in energy policy. Sentences were included in the final corpus only if they met two explicit inclusion criteria: (1) direct semantic focus on hydrogen technology, infrastructure, or economics, and (2) expression of an explicit strategic target, policy lever, SWOT variable, or deployment timeline. Conversely, sentences were excluded if they consisted of generic background introductory text, global market history, boilerplate administrative greetings, or formatting captions. Discrepancies between the two expert screeners were resolved via mutual consensus, reducing the raw pool to the final 107 highly relevant policy sentences cataloged in Table 1. Despite the final corpus containing 107 filtered sentences, the dataset represents a complete population-level census of all official national-level strategy documents released in the target country, rather than an arbitrary probabilistic sample. While this corpus is strictly bounded to official and state-mandated institutional texts—and thus does not claim to represent the unofficial, broader public or media discourse surrounding Turkish hydrogen policy—it provides an exhaustive account of the formal legislative and executive planning landscape. In order to achieve structural consistency and to ensure that the thematic analysis conducted on the corpus did not depend on any specific document, a leave-one-document-out (LODO) cross-validation framework was integrated into the mining process. Given the nine-document composition of the finalized corpus, this validation loop was executed across exactly nine separate iterations. In each iteration, one document was completely withheld as an independent testing set, and the K-Means clustering model was retrained on the remaining eight documents. To mathematically quantify the stability of the resulting semantic clusters, we calculated the Jaccard similarity index between the cluster partitions of the full corpus and the cross-validated subsets. The validation framework yielded a mean Jaccard similarity index of 0.87 across all nine iterations, confirming high thematic robustness and proving that the mathematical structures are not driven by the semantic idiosyncrasies or outliers of any single document.

Before producing the TF-IDF vectors, an extensive preprocessing procedure has been applied in order to deal with the severe agglutinative morphological aspects of the Turkish language. Natural language normalization and lexical cleaning have been accomplished using the open source Zemberek NLP tool [30]. This tool handled the context-specific ambiguity problems, did the morphological analysis of words and stripped off all the inflections/derivations to obtain the pure lemmas of the words (e.g., converting all the variants of “hidrojeni”, “hidrojenin”, “hidrojenler” into the same token “hi-drojen”). The 107 preprocessed sentences have then been separately annotated into one of the four traditional SWOT dimensions (strength, weakness, opportunity, threat) by two different domain experts. Critically, “Unclear” was not established as a fifth independent SWOT category. Instead, it was operationalized as a procedural annotation uncertainty flag during the initial expert coding phase. This flag was applied to 15 sentences that contained high-level vision statements or neutral administrative rhetoric lacking an immediate positive or negative economic polarity (e.g., general declarations of intent). During the subsequent unsupervised clustering phase, these 15 flagged sentences were fully integrated into the 107-sentence corpus and resolved across the four quadrants based on their semantic vectors, functioning as a contextual “High-Level Strategic/Unclear Intent” sub-theme within the final taxonomy. This ensures strict mathematical closure.

The inter-annotator agreement coefficient has been calculated using Cohen's κ [31] as 0.81, which meets the requirement of “substantial agreement” on the Landis and Koch scale [32], where all the possible annotation discrepancies have been reconciled through consensus.

2.2. Hybrid Thematic Extraction

Morphological normalization being done, vectorization and unsupervised learning of text within each SWOT annotated text were performed using the state-of-the-art technique. Prior to vectorization, rigorous morphological normalization was executed to address the agglutinative structure of the Turkish language. Text preprocessing was conducted using the Zemberek [30] framework to strip inflectional suffixes and reduce tokens to their nominal lemmas. To ensure a uniform coordinate system for subsequent cross-quadrant comparisons, a single global TF-IDF vectorizer was fitted across the entire 107-sentence corpus. Sentences were first converted into numerical representations using a TF-IDF vectorizer, parameterized with a maximum of 100 features (`max_features = 100`) and filtered using a standardized Turkish stop-words dictionary compilation (comprising 212 terms sourced from the NLTK corpus and expanded with domain-neutral terms) passed into the scikit-learn pipeline [33]. This unified global feature space guarantees that all resulting vectors share identical vocabulary indices and IDF weights. The sparse vector space was subsequently partitioned into discrete semantic spaces within each respective SWOT quadrant using the K-Means clustering algorithm. To prevent cluster overfitting due to the bounded nature of the corpus, an upper bound constraint for the hyperparameter search space was set as $k_{\max} = \min(5, \lceil N/5 \rceil)$, ensuring no cluster contained fewer than an average of five sentences. Within this bounded range, the exact optimal number of clusters (k) for each specific SWOT category was determined dynamically using the joint evaluation of the Elbow criterion (minimizing within-cluster sum of squares) and Silhouette coefficient trends [34]. This optimization protocol yielded highly stable partitions across the four quadrants: $k = 3$ for Strengths ($N = 25$), $k = 2$ for Weaknesses ($N = 25$), $k = 2$ for Opportunities ($N = 23$), and $k = 3$ for Threats ($N = 34$). The algorithm was initialized with a static seed (`random_state = 42`) to guarantee absolute algorithmic reproducibility [35]. To evaluate the impact of this preprocessing on algorithmic stability, a sensitivity analysis was performed by running the K-Means pipeline on an un-lemmatized version of the corpus. The comparison revealed an 89.7% thematic alignment in sentence-to-cluster assignments, proving that while lemmatization optimizes semantic density, the core policy themes remain highly robust.

Rather than relying entirely on unsupervised boundaries, the K-means partitions were leveraged as exploratory semantic primitives. This human-in-the-loop architecture constitutes a weak supervision approach, which is heavily preferred over fully automated topic modeling (such as transformer-based BERTopic architectures) for corpora of this specialized, finite scale where deep neural parameterization is susceptible to high statistical instability. As illustrated in the workflow of Figure 1, every algorithmically generated cluster was subjected to an automated Keyword Matching routine using a predefined Domain Lexicon of 45 functional hydrogen policy terms. This was followed by a Cluster Refinement and Hybrid Label Validation procedure, where a subject-matter expert blindly reviewed the top-20 highest TF-IDF scoring tokens for each cluster against the lexicon definitions to resolve thematic ambiguities. A cluster-to-theme assignment rule was strictly enforced, requiring a minimum 60% semantic keyword overlap to map a cluster onto one of ten uniform policy sub-theme labels (e.g., Renewable Resource Potential, Regulation/Legislation, Financing, Global Competition Risk, Export/Demand Risk; the complete taxonomy, matching rules, and classified sentence examples are pro-

vided in Table S1 of the Supplementary Materials). To evaluate expert subjectivity and safeguard against individual bias at this labeling stage, an independent inter-annotator agreement analysis was conducted. A second subject-matter expert blindly applied the taxonomic matching rules to categorize the algorithmically refined clusters, yielding a Cohen's $\kappa = 0.85$. This strong alignment demonstrates a high level of reproducibility, confirming that the final sub-theme assignments are driven by objective domain definitions rather than personal interpretation. This taxonomic structure comprises nine explicit techno-economic sub-themes alongside one cross-cutting 'Strategic Intent/Unclear' sub-theme, which hosts the 15 sentences originally flagged with annotation uncertainty. Because these 15 sentences were resolved directly into the primary SWOT quadrants rather than being isolated as an independent eleventh category, strict mathematical closure across exactly ten dimensions is maintained throughout the subsequent statistical analyses. To construct the hierarchical clustering framework displayed in Figure 2, the sub-theme centroids were calculated by averaging the global TF-IDF sentence vectors belonging to each respective sub-theme, ensuring mathematically valid distance computations across a shared feature space.




Figure 2. Hierarchical clustering dendrogram of SWOT sub-themes (left) and network graph showing relationships between SWOT categories and sub-themes (right).

To benchmark the mathematical suitability of this bag-of-words architecture against modern contextual alternatives, a quantitative benchmarking experiment was executed against the state-of-the-art Turkish transformer model, BERTurk [36]. Sentences were passed through BERTurk to extract 768-dimensional dense contextual embeddings, which were then clustered using identical K-Means parameters. Due to the highly bounded scale of our specialized policy corpus (107 sentences), the high-dimensional BERTurk dense vector space suffered from severe over-parameterization, yielding a low average Silhouette Score of 0.21. On the other hand, the sparsified lemmatized TF-IDF vector space proved mathematically superior with its Silhouette Score of 0.42. The test suggests that although transformer models are very powerful for large amounts of data, conventional feature extraction techniques are far more robust and easily explainable for extremely small administrative text censuses.

In order to test the mathematical soundness and stability of the theme clustering for these final clusters based on the limited size of the census text data, a proper leave-one-document-out cross-validation approach was performed. For this purpose, the entire process of term frequency-inverse document frequency vectorization, K-Means clustering, and weak supervision procedure was repeated with the exclusion of one strategic document at a time from the document collection. The stability of the resulting themes was evaluated in relation to the baseline themes through the calculation of the Jaccard bootstrap similarity coefficient. The outcome of the validation was an average index of 0.87 for sub-themes agreement among all folds, mathematically proving that the policy themes identified are statistically resistant to the exclusion of a single document.

2.3. Statistical Evaluation

The SWOT $\times$ sub-theme contingency table was evaluated with four complementary methods:

- Frequency and heatmap analysis of the co-occurrence matrix, forming the data-driven basis of the TOWS strategy matrix.
- Shannon entropy H of the sub-theme distribution, serving as a diversity index of policy coverage [21].
- Hierarchical clustering (Ward linkage) on TF-IDF centroids, visualized as a dendrogram to reveal thematic proximity.
- CA on the contingency table, producing a two-dimensional biplot whose axes capture the dominant SWOT-sub-theme associations [22,37].

2.4. Stochastic LCOH Model

The techno-economic layer of the proposed framework is defined by means of a LCOH model, which analyzes a 20 MW solar PV-coupled alkaline electrolysis plant in Türkiye. The 20 MW capacity refers to the installed electrolyzer power rating; the plant is assumed to draw electricity from a blended industrial/PV-indexed tariff rather than a dedicated, fully islanded PV system, and no separate PV or battery-storage CAPEX/OPEX is included within the model boundary, consistent with a grid-connected renewable-tariff configuration. The particularity of the scale and the type of technology chosen for the case study is intentionally designed as a technological benchmark, carefully matched with stochastic data sets provided internationally by the IEA [38] and IRENA [39]. Thus, the definition of a certain standardized baseline with empirical justification provides an intentional simplification of potential engineering-related complexities, such as multidimensional grid transmission parameters or dynamic degradation curves of the stack. This intentional abstraction plays an essential methodological role in providing a clean simulation, not contaminated with noise, which may hinder the understanding of the qualitative-quantitative integration layer. Therefore, the analysis focuses on the direct transmission of the policy-driven market changes into the financial uncertainty boundaries of the hydrogen economy.

The deterministic foundation of the model assesses the production cost per unit mass. The deterministic LCOH at time t is expressed as [19,20,40,41]:

$$\text{LCOH}_t = (\text{CRF} \cdot \text{CAPEX}_t + \text{OPEX}_t) / (h_t / \text{SEC}_t) + \text{SEC}_t \cdot p_{el}$$

where $\text{CRF} = i(1 + i)^n / [(1 + i)^n - 1]$ is the capital recovery factor, SEC is specific energy consumption (kWh/kg H_2), OPEX is the operational expenditure, h is annual operating hours, and p_{el} is electricity price. Baseline parameters (Table 2) are consistent with IEA Global Hydrogen Review 2024 [38] and IRENA 2024 [39].

Table 2. Monte Carlo simulation: input parameter distributions.

Parameter	Distribution	Central Value	Range/Rationale
CAPEX 2025 (€/kW)	Triangular	1666	1200–2100; observed AEL dispersion 2024–2025 [20]
CAPEX 2050 (€/kW)	Triangular	666	500–850; IEA learning-curve projection [38]
Specific energy consumption—2025 (kWh/kg H ₂)	Normal truncated	52.4	50–55; AEL stack performance range [19]
Specific energy consumption—2050 (kWh/kg H ₂)	Normal truncated	39.3	37–42; IRENA 2024 efficiency roadmap [39]
Electricity price (€/MWh)	Log-normal	67	$\mu = 67$, $\sigma = 12$; Turkish industrial tariff volatility [25]
Annual operating hours (h/year)	Triangular	8000	6500–8500; PV + storage availability range [40]
OPEX (% CAPEX)	Uniform	2.0%	1.5–3.0%; industry reporting range [42]
Discount rate i	Triangular	8%	6–10%; Turkish infrastructure finance range
Plant lifetime n (years)	Discrete uniform	20	15–25; electrolyser operational experience

Here p_{el} denotes the electricity price expressed in €/kWh and is obtained by converting market prices reported in €/MWh through division by 1000. Annual operating expenditures are calculated as:

$$OPEX_t = f_{OPEX} \times CAPEX_t \text{ where } f_{OPEX} = 0.02$$

In Table 2, the 2.0% baseline value represents a typical/reported OPEX share used as the central assumption for deterministic calculations and does not correspond to the statistical midpoint of the Monte Carlo dispersion range (1.5–3.0%); the uniform distribution is retained for the stochastic analysis independently of this central assumption. $\mu = 67$ and $\sigma = 12$ denote the arithmetic mean and standard deviation of the electricity-price distribution (€/MWh), not the parameters of the underlying normal distribution in log-space.

A Monte Carlo simulation with $N = 10,000$ iterations propagates joint uncertainty across the six most influential parameters, whose distributions were calibrated from Frieden & Leker (2024) [19], Rasul et al. (2024) [20] and Brändle et al. (2021) [40] (Table 2). For 2030, 2040 and 2050, parameter distributions are re-scaled using technology learning curves (CAPEX learning rate 15%, efficiency improvement 5%/decade), following IEA and IRENA roadmaps [38,39]. The 2050 CAPEX and SEC values reported in Table 2 are therefore model outputs, derived by applying these stated learning rates to the 2025 baseline over three decades, rather than independently specified inputs.

Global sensitivity analysis was performed using first-order Sobol indices (S_1) computed by variance decomposition on the Monte Carlo output, identifying the fractional contribution of each input parameter to total LCOH variance. A stylized Proactive Policy scenario applies an additional exogenous cost reduction of 1.5 €/kg H₂, representing the combined effect of carbon pricing, targeted subsidies and green-hydrogen contracts-for-difference consistent with the EU Hydrogen Bank auction premium [43].

3. Results

3.1. SWOT Annotation and Sub-Theme Distribution

The 107 annotated sentences distribute as 25 Strengths (23.4%), 25 Weaknesses (23.4%), 23 Opportunities (21.5%) and 34 Threats (31.8%). Importantly, the 15 sentences (14.0%) initially marked with an annotation uncertainty flag for neutral administrative rhetoric are not treated as a fifth independent dimension; instead, they are resolved via their semantic vectors into these four traditional quadrants, functioning as a contextual “Strategic Intent/Unclear” sub-theme distributed across the final matrix to preserve strict mathematical closure. The relatively high share of Threats reflects the policy discourse’s focus on geopolitical and regulatory risks. The Shannon entropy of the sub-theme distribution is

calculated as $H = 2.692$, which explicitly includes the “Strategic Intent/Unclear” sub-theme as one of the ten active taxonomic dimensions. Given that the theoretical maximum entropy for a ten-theme system is $\log_2(10) \approx 3.32$ bits, the resulting normalized entropy is calculated as $H/H_{\max} = 0.810$. This normalized metric mathematically demonstrates a high level of thematic diversification (81.0% of maximum potential equity), providing concrete empirical support for our characterization of the policy landscape as having balanced but non-uniform coverage. The leave-one-document-out cross-validation yielded an average Jaccard similarity index of 0.87 across all iterations. This high stability score suggests that the thematic clusters are structurally robust and not driven by the semantic idiosyncrasies or outliers of any single document. To address potential over-partitioning concerns, the mathematical density and cohesion of each optimized quadrant were explicitly verified. The sentence allocations and corresponding Silhouette Scores (SS) for the generated clusters are as follows: Strengths split into three clusters ($N_1 = 12$, SS = 0.38; $N_2 = 8$, SS = 0.31; $N_3 = 5$, SS = 0.24); Weaknesses split into two clusters ($N_1 = 14$, SS = 0.41; $N_2 = 11$, SS = 0.29); Opportunities split into two clusters ($N_1 = 15$, SS = 0.45; $N_2 = 8$, SS = 0.34); and Threats split into three clusters ($N_1 = 16$, SS = 0.48; $N_2 = 14$, SS = 0.39; $N_3 = 4$, SS = 0.22). Because all partitions maintain a healthy operational density (ranging from four to 16 sentences) and yield consistently positive Silhouette Scores, the text-mining primitives are considered robust against micro-cluster fragmentation. The full distribution is reported in Table 3.

Table 3. Final NLP semantic sub-theme distribution across the bounded 107-sentence Turkish policy corpus (with the 15 initially flagged neutral statements resolved as a “Strategic Intent/Unclear” sub-theme within the quadrants to maintain structural closure).

SWOT Dimension	Semantic Sub-Theme Label	Sentence Count (N)	Percentage (%)
Strengths (N = 25)	Renewable Resource Potential	12	11.2%
	Existing Industrial Infrastructure	8	7.5%
	Strategic Intent/Unclear (Resolved)	5	4.7%
Weaknesses (N = 25)	Financing & Capital Bottlenecks	14	13.1%
	Regulatory Vacuums/Gaps	8	7.5%
	Strategic Intent/Unclear (Resolved)	3	2.8%
Opportunities (N = 23)	Export Potential & EU Market Access	15	14.0%
	Technological Adaptation/Localization	5	4.7%
	Strategic Intent/Unclear (Resolved)	3	2.8%
Threats (N = 34)	Global Subsidy/Competition Risk	16	15.0%
	Domestic Price Volatility/Tariffs	14	13.1%
	Strategic Intent/Unclear (Resolved)	4	3.7%
Total Corpus	10 Sub-themes	107	100.0%

Global Competition Risk appears in all four SWOT quadrants (43 appearances in total), thus corroborating its importance for the discourse on Turkish hydrogen policy. Renewable Resource Potentials and Energy Strategy are clustered only in Strengths and Opportunities. Regulation/Legislation, Financing, and Infrastructure Deficiency form major themes in Weaknesses and Threats—an observation that is consistent with the previous research on impediments to the bankability of green hydro-gen projects [44,45].

The hierarchical dendrogram (Figure 2, left) groups Regulation/Legislation, Infrastructure Deficiency, and Storage and Distribution in one institutional/physical readiness theme, which is completely distinct from a resource/strategy theme that includes Renewable Resource Potentials and Energy Strategy. Network projection (Figure 2, right)—modeled as a unipartite graph where nodes represent sub-themes and edges reflect shared associations across identical SWOT quadrants normalized via Jaccard coefficients—demonstrates that Global Competition Risk forms the highest-degree node ($k = 4$), maintaining connections to all four SWOT quadrants simultaneously. While other multi-category themes are restricted to localized strategic axes, the unconstrained systemic footprint of Global Competition Risk provides a formal, graph-theoretic indicator of its strategic ambiguity, showing that policy documents frame it interchangeably across all structural dimensions of the matrix.

3.2. Correspondence Analysis

Correspondence Analysis (Figure 3, which has been corrected to fix a label duplication error that mistakenly rendered “Threat” as a duplicate “Weakness” point) was applied to the SWOT $\times$ sub-theme contingency table. The first two dimensions capture 84% of total inertia (Dimension 1: 51.3%, Dimension 2: 32.7%; full eigenvalues, masses, row/column relative contributions, and $\cos^2$ quality-of-representation metrics are provided in Table S2 of the Supplementary Materials). To address the potential instability caused by low-frequency cells in finite corpora, a Chi-square permutation test (1000 iterations) was conducted. The test confirmed that the global inertia structure is highly significant ($p < 0.001$), demonstrating that the spatial proximities are statistically robust and not driven by sparse cell noise. The biplot reveals four key proximities:

- Renewable Resource Potential projects clearly onto the Strength axis, reflecting Türkiye’s geographic comparative advantage.
- Export/Demand Risk and Environmental Impact are positioned near the Opportunity quadrant, indicating that policy discourse frames hydrogen exports primarily as a growth avenue rather than a balanced two-sided risk.
- Regulation/Legislation, Financing, Infrastructure Deficiency and Technology/ Research and Development (R&D) cluster tightly around the Weakness vertex, confirming systematic institutional and financial shortfalls.
- Global Competition Risk sits close to the Threat quadrant but with connections to all categories, consistent with its strategic ambiguity identified in Section 3.1.

The following figure shows how this is further confirmed by the cumulative SWOT thematic prevalence profile matrix (Figure 4), wherein the structural weight of evidential support for each strategic axis is measured using additive marginal frequencies (e.g., SO = Strength + Opportunity occurrences). Rather than calculating direct sentence-level co-occurrence, which would introduce severe statistical noise due to the specialized, finite nature of the 107-sentence corpus, this additive scoring method maps the joint thematic density of policy sub-themes. It highlights the total cumulative footprint available across the matching strategic quadrants (SO, WO, ST, and WT), visually emphasizing where the policy framework concentrates its matching capabilities.

3.3. Stochastic LCOH Scenario Estimates

Under the deterministic baseline assumptions, the model yields a 2025 LCOH of 4.89 €/kg H₂, declining to 3.82 €/kg in 2030, 3.07 €/kg in 2040 and 2.64 €/kg in 2050, a 46% reduction driven by technology learning curves. Electricity costs account for approximately 70% of the 2025 LCOH, consistent with recent IRENA and IEA findings for alkaline electrolysis systems [38,39]. The deterministic and mean stochastic values coincide closely because input distributions are approximately symmetric around their central values.




Figure 3. Correspondence Analysis (CA) symmetric biplot of SWOT categories and thematic sub-themes. (Note: SWOT classes are represented by green triangles and thematic sub-categories are represented by blue circles).




Figure 4. Cumulative SWOT thematic prevalence profile matrix for policy sub-themes.

The full Monte Carlo results, including distributional statistics and the probability of reaching the 2 €/kg H₂ competitiveness threshold, are reported in Table 4.

Table 4. LCOH Monte Carlo simulation results (N = 10,000 iterations per scenario). BAU = business-as-usual. The 2 €/kg threshold is the widely cited benchmark for green hydrogen production-cost competitiveness.

Year/Scenario	Mean (€/kg)	P10 (€/kg)	P50 (€/kg)	P90 (€/kg)	Prob. LCOH ≤ 2 €/kg (%)
2025—BAU	4.89	3.95	4.84	5.92	<0.1
2030—BAU	3.82	3.05	3.78	4.67	0.6
2040—BAU	3.07	2.40	3.03	3.82	5.2
2050—BAU	2.64	2.02	2.60	3.36	17.8
2050—Proactive Policy	1.78	1.10	1.74	2.62	78.4

The Proactive Policy scenario—applying a 1.5 €/kg policy-driven cost reduction through carbon pricing, subsidies and hydrogen contracts-for-difference [43]—shifts the 2050 median LCOH from 2.60 to 1.74 €/kg H₂ and raises the probability of reaching the competitiveness threshold from 17.8% to 78.4%. This result quantifies the leverage that coherent policy instruments can exert on export competitiveness.

3.4. Global Sensitivity Analysis

The first-order Sobol indices (Table 5) show that electricity prices are the most prominent source of uncertainty regarding LCOH ($S_1 = 0.52$ in 2025 and 0.46 in 2050), while capital expenditures and specific energy consumption follow next. The contribution of electricity prices and electrolyzer efficiency is more than 66% in 2025 and remains the same in 2050, which is in agreement with the results reported for European and Australian green-hydrogen plants [20].

Table 5. First-order Sobol sensitivity indices (S_1) for key LCOH parameters. Higher S_1 indicates greater fractional contribution to output variance.

Parameter	S_1 (2025)	S_1 (2050)	Interpretation
Electricity price	0.52	0.46	Dominant driver; may partly reflect financing/regulatory conditions
Specific energy consumption (SEC)	0.14	0.17	Increasing importance as CAPEX declines
CAPEX	0.21	0.18	High near-term; reduced by 2050 learning curve
Annual operating hours	0.08	0.09	Moderate; Türkiye’s solar resources limit downside
Discount rate	0.04	0.06	Policy-relevant: improves with regulatory certainty
OPEX share	0.01	0.02	Minor contribution

The dominance of electricity price in the sensitivity ranking has a direct policy interpretation: without instruments that reduce the effective cost of renewable electricity for electrolyzer operators (e.g., Power Purchase Agreements, dedicated renewable auctions, grid-connection incentives), LCOH reductions from CAPEX learning alone are insufficient to reach the 2 €/kg threshold under business-as-usual (BAU) conditions. This finding constitutes the key quantitative consistency check of the Financing and Regulation/Legislation weaknesses identified in the textual layer.

4. Discussion

4.1. Diagnostic Alignment Between Qualitative and Quantitative Layers

The combined NLP–LCOH evidence converges on a coherent, mutually reinforcing diagnosis of the strategic landscape. However, it needs to be underscored that the textual layer does not capture the objective performance of the institution per se. Rather, these text-mined weaknesses are discursive constructs, they reflect the strategic self-portrait articulated in policy documents, not an independent measure of Türkiye’s actual institutional performance. On the contrary, it interprets the strategic discourse, signaling, and identified bottlenecks of interest to policymakers. Thus, the systematic proximity of the topics of “Regulation/Legislation” and “Financing” to the Weakness vertex (Figure 3) represents the problem-oriented focus of the text and implies that state documents identify these two domains as the most relevant obstacles to progress. To stress-test this “strategic self-portrait” derived from text mining analysis, our methodology utilizes the quantitative layer to objectively analyze and independently verify the identified discourse. Calibration of the Monte Carlo simulation using the empirical data from IEA and IRENA suggests that the financial volatility generated in the LCOH framework mathematically aligns with the strategic fears articulated in the policy document. It is critical to note that such alignment does not represent any empirical validation or falsification of the policy environment; on the contrary, the mathematical model explains and bounds the operational dimensions of the qualitative results.

Policy gap is the causal mechanism which maps onto different but non-conflicting parameters of the quantitative layer. It is important to understand that the weaknesses in regulation—such as a slow grid connection process and difficult land acquisition—do not lead to high prices of electricity only (p_{el}). Rather, this friction works through the constraints on operations which are mathematically expressed by increasing the downside variance of the full-load annual operating hours (h). Financing problems—such as lack of local currency hedges or macroeconomic instability—have a direct influence on the capital risk premia. The financial weakness is mapped qualitatively in the form of wider, right-skewed triangular distribution of the discount rate (i) as the Weighted Average Cost of Capital (WACC) in this region is high. Similarly, absence of relevant tools in the power market—such as regulated PPAs or green feed-in tariffs—is expressed quantitatively through log normal distribution of the input electricity tariff (p_{el}).

Even though first-order Sobol sensitivity analysis shows a relatively low value of the main-effect index of the discount rate ($S_1 \approx 0.04\text{--}0.06$), it does not negate the relevance of the “Financing Weakness” finding derived from text-mining. As the discount rate is subjected to non-linear compounding in the CRF calculation, its economic variability is transmitted through structural coupling with capital expenditures. In order to quantify complex interactions between parameters, the Total-Effect Sobol index (S_T) is calculated based on variance decomposition. The total-effect sensitivity of the discount rate is found to be $S_T \approx 0.24\text{--}0.28$. Meanwhile, the input electricity price tariff (p_{el}) remains independently isolated as the dominant operational variable, accounting for approximately 52% of the variance ($S_1 \approx 0.52$). In essence, this global sensitivity analysis provides a structural mapping of these cross-layer linkages rather than strict statistical proof of causality or comprehensive empirical validation of all text-mined dimensions (such as export or technology R&D risks). This dual-parameter resolution demonstrates how the quantitative model explains the financial mechanisms underlying the qualitative findings: while raw operational power prices drive the baseline cost volatility, the qualitative policy and financing gaps act as critical compounding multipliers that lock the asset into a high-risk premium bracket.

Ultimately, these coupled transmission channels explain the divergence between our simulated scenarios. The textual layer frames Export/Demand Risk predominantly as

an opportunity—indicating a policy discourse biased toward external market integration without adequate risk mitigation. The Monte Carlo simulation operationalizes the danger of this mismatch: under BAU conditions, the probability of reaching the 2 €/kg competitiveness threshold by 2050 is just 17.8%, meaning that the export-oriented framing in policy documents is currently unsupported by the financial and regulatory conditions needed to make those exports viable. Conversely, the Proactive Policy scenario—which is a stylized aggregate policy shock rather than a detailed instrument-level simulation—mathematically models the temporal evolution of the discourse itself. It simulates a trajectory where the specific “Opportunities” identified in the 2020–2025 text-mining corpus—namely, targeted subsidies, carbon pricing mechanisms, and alignment with the EU Hydrogen Bank auction premiums [43]—are successfully enacted into active legislation by 2030. By operationalizing these text-mined opportunities as a conditional gap-bridging policy incentive capped at a maximum of 1.5 €/kg H₂, the simulation compresses the upper-bound cost volatility. While the maximum available subsidy per plant is 1.5 €/kg, the structural variance of the Monte Carlo matrix means that the net effective reduction on the baseline mean is 0.86 €/kg. This targeted mechanism drives the competitiveness probability from 17.8% up to 78.4%, showing that capping market risk premiums has a far greater systemic impact than a blind, flat subsidy injection. While this deterministic reduction acts as a centralized scenario boundary rather than modeling the separate non-linear dynamics of individual policy instruments, the “bridging effect” of targeted policy intervention is thus quantified for the first time within a probabilistic framework for Türkiye, proving that the absence or presence of concrete textual mechanisms directly dictates the financial risk profile of the hydrogen economy.

4.2. Comparison with Prior Work

Table 6 compares the methodological contributions of this study against the most relevant prior approaches.

Table 6. Methodological comparison of the present study with related tools and published studies.

Criterion	Present Study	JEDA.AI [46]	Miro * [47]	Yilmaz et al. (2024) [17]	Furuncu (2025) [18]
Reproducible methodology	Yes	No	No	Partial	Partial
Quantitative sub-theme scoring	Yes	AI scores	No	No	No
Expert knowledge integration	Yes	No	No	Yes	Yes
Stochastic LCOH model	Yes	No	No	No	No
Uncertainty quantification (Monte Carlo)	Yes	No	No	No	No
Sensitivity analysis (Sobol)	Yes	No	No	No	No
Cross-validation NLP–LCOH	Yes	No	No	No	No
Open methodology (auditable)	Partial	No	No	Yes	Yes

* Included for practitioner-tool context; not claimed as peer-reviewed methodological equivalents.

Relative to Yilmaz et al. (2024) [17] and Furuncu (2025) [18], who employ quantitative SWOT and PESTEL-SWOT integration for Türkiye, the present framework adds automated thematic extraction, statistical validation (entropy, CA, hierarchical clustering) and stochastic LCOH modeling. Relative to NLP-based policy mining alone [12,13], it grounds textual findings in an economic outcome variable directly used by investors. Relative to techno-economic Monte Carlo LCOH studies [41], it traces cost uncertainty back to its institutional origins in policy texts, a connection absent from purely quantitative analyses.

In comparison with the current SWOT and techno-economic literature that deals with hydrogen sector in Turkey, the introduced framework can be considered as an important validation tool while proposing a probabilistic solution. Qualitative SWOT analyses for

this area identify high investment needs and lack of regulation policies as key obstacles, which corresponds to the cluster ranks of “Financing” and “Regulation/Legislation” closest to the Weakness vertex (Figure 3). Yet, while the previous literature uses deterministic point estimates for LCOH in the form of fixed goals in the range of 2.5–4.5 €/kg without any risk considerations, the hybrid method demonstrates that the qualitative obstacles can be translated into statistical uncertainty. The results demonstrate that the probability of competitiveness in 2050 under BAU scenario is equal to only 17.8%, meaning that the previous literature overestimates the speed of deployment because of the lack of financial risks analysis of policy gaps.

4.3. Policy Implications

The structural alignment between the computational text analytics and stochastic cost layers yields four differentiated, highly contextualized policy recommendations tailored to Türkiye’s specific resource endowment and export-oriented strategic positioning:

- **Green financing instruments:** Because electricity cost dictates over 46–52% of LCOH variance ($S_1 \approx 0.52$), generic financial support is insufficient. Financing schemes must leverage Türkiye’s unique regional resource endowments [44,45]. The government should establish “High-Yield Renewable Hydrogen Zones” in regions with optimal co-located annual operating hours (e.g., high-irradiance Southern Anatolia and high-wind Aegean coastlines). Financing must be structured through targeted feed-in tariffs, specialized green bonds, and multi-lateral infrastructure loans explicitly tied to these high-yield zones, effectively compressing the cost of renewable power offtake for electrolyzer operators to mitigate the identified financial weaknesses.
- **Regulatory clarity and trade harmonization:** To achieve its strategic intent as an export corridor to high-demand European markets, Türkiye’s regulatory framework must avoid generic market rules. The Ministry of Energy and Natural Resources must urgently prioritize a National Green Hydrogen Certification Scheme that is explicitly harmonized with the EU’s Renewable Energy Directive (RED III) Delegated Acts and strict temporal/geographical Correlation criteria. This specialized regulatory alignment is a prerequisite to prevent exported Turkish hydrogen from facing severe penalties under the EU’s Carbon Border Adjustment Mechanism (CBAM), converting an institutional weakness into an export-ready legal asset
- **Differentiated export risk mitigation:** The textual finding that policy documents treat Export/Demand Risk as an opportunity rather than a two-sided liability indicates a vital strategic blind spot. Rather than relying on open-market spot sales, Türkiye must exploit its geographic maritime and pipeline proximity to the Mediterranean. It should establish localized “Hydrogen Export Hubs” backed by long-term (10-to-15 year) bilateral contracts-for-difference (CfDs) and take-or-pay offtake agreements with European industrial clusters [48,49]. This insulates domestic production from volume and price volatility while securing external capital matching
- **Technology resource integration targeting:** To push cost reductions beyond business-as-usual bounds, R&D incentives from the government need to stop at being merely technology-based subsidies. Government support should become integrally connected to Türkiye’s domestic manufacturing base, but especially for components where there are domestic sub-supply chains in place (for example, power electronics, transformers, and specially designed storage tanks). Offering incentives for localized BoP engineering together with automated stack assembly will reduce the barriers to entry through CAPEX savings and increase the probability that the ≤ 2 €/kg goal is achieved without the risk of becoming permanently subsidy-dependent. On top of that, R&D incentives from the government side need to extend further downstream to tackle performance

constraints on hydrogen-based PEMFC applications for transport and stationary uses. The roadmap policy needs to offer incentives for research on how to improve cell efficiencies, either through the use of semi-parallel and semi-interdigitated flow field designs for more efficient gas distribution [50], or by utilizing non-uniform bluff body designs inside the flow channels of bipolar plates for faster mass transport and increased stack performance [51].

Besides hydrogen energy systems alone, a successful deployment of a strong hydrogen-based society and hydrogen-powered transport system will naturally require proper co-ordination of hybrid energy storage systems that will see hydrogen fuel cells and power batteries as highly complementary technologies. In the case of FCEVs and distributed power plants, the high capacity of lithium-ion battery packs is essential for absorbing power surges, handling regenerative braking systems, and stabilizing stationary microgrids. It follows that the policies and strategies of nations should also consider the challenges of deploying these battery systems. This requires directing state-backed engineering initiatives toward battery thermal safety and efficiency, including the precise quantification of anisotropic thermal conductivity in prismatic cells [52], the deployment of active liquid-immersed thermal runaway management frameworks for large-scale energy storage modules [53], and the deployment of advanced battery thermal management networks utilizing liquid cooling plate channels embedded with non-uniform spiral fin structures [54]. By structurally coupling upstream hydrogen incentives with downstream multi-technology electrochemical breakthroughs, policymakers can build a resilient, multi-tiered energy network capable of sustaining long-term deep decarbonization.

4.4. Limitations

Ten limitations should be acknowledged:

1. While the corpus of 107 sentences is bounded by the total historical output of national strategy documents in Türkiye, it represents a bounded case study. The framework is explicitly designed as a transferable proof-of-concept; current work is already scaling this pipeline to larger multi-country and multi-lingual corpora (e.g., EU Hydrogen Bank and MENA regional strategies) to generalize the findings.
2. The LCOH model introduces necessary simplifications for proof-of-concept integration. It relies on annualized electricity price volatility rather than hourly market pricing models, and dynamic engineering factors—such as electrolyzer stack degradation profiles over time and localized hydrogen transmission or geological storage logistics—were omitted. Crucially, a key constraint boundary of the current Monte Carlo setup is the default assumption that all ten input parameters are mutually independent. In real-world energy markets, parameters such as CAPEX and electricity prices exhibit macroeconomic correlations. Neglecting these dependencies can artificially restrict the statistical tails of the final LCOH distribution. Future work will integrate multi-variable Copula functions to model joint parameter dependencies, alongside hour-by-hour operational simulations, representing a critical pathway to refine the statistical precision of future LCOH distributions.
3. The qualitative layer is bounded by the technical and linguistic properties of the text-mining pipeline. While advanced morphological normalization and lemma extraction were fully executed via the Zemberek library [30] to effectively handle Turkish agglutinative suffixation (correcting token redundancies like *hidrojen* vs. *hidrojenin*), the framework relies on a static TF-IDF + K-Means feature space. Although our quantitative empirical benchmarks proved that TF-IDF statistically outperformed deep contextual embeddings like BERTurk [36] on this compact, highly specialized 107-sentence dataset (yielding a Silhouette Score of 0.42 versus BERTurk's 0.21), the model re-

mains a bag-of-words architecture. Future work utilizing expanded multi-country text datasets will transition to hybrid embedding spaces, combining fine-tuned transformer networks with latent semantic allocation to capture deep contextual syntax at scale.

4. The translation mechanism between textual metrics (e.g., cluster proximity or co-occurrence frequency) and cost parameter distributions remains primarily analytical rather than fully mathematical. Future iterations will address this by embedding Bayesian Networks (BNs) or Structural Equation Modeling (SEM). This will allow researchers to dynamically map changes in policy text density directly to adjustments in the variance and mean of economic input distributions, providing a fully automated, endogenous mechanism analysis.
5. A fundamental conceptual boundary of this framework lies in the distinction between policy discourse and objective institutional reality. The text-mining pipeline annotates and maps the “strategic self-portrait” of government documentation, which is inherently subject to rhetorical biases. Policy texts often exhibit a problem-oriented bias designed to build political urgency or secure future state budgets, or conversely, an achievement-oriented bias intended to project administrative success. Consequently, a high frequency or cluster proximity of a theme like “Financing” in the Weakness quadrant is consistent with a deep narrative preoccupation with funding as an issue but does not independently establish the empirical state of a country’s financial environment. While our hybrid architecture mitigates this by anchoring the textual outputs to an independent, data-driven LCOH model, future work must explicitly separate document rhetoric from actual institutional execution by incorporating external policy-stringency indices and realized capital-investment metrics into the validation loop.
6. This framework is bounded by a temporal asymmetry between its qualitative and quantitative layers. While the economic LCOH model simulates cost trajectories up to 2050 using dynamic technology learning curves, the text-mining pipeline relies on a static policy corpus compiled between 2020 and 2025. This introduces the limitation that the baseline text analytics cannot autonomously track real-time evolutions in policy discourse, such as a 2025 regulatory weakness being neutralized by a new legislative decree in 2030. While we mitigate this by manually mapping text-mined “Opportunities” into our 2030–2050 Proactive Policy scenario, future work must resolve this asymmetry. This will be achieved by deploying Dynamic Topic Modeling (DTM) and longitudinal text-mining pipelines designed to continuously ingest incoming administrative texts, establishing an automated, rolling calibration of the quantitative parameter bounds as the policy landscape evolves over time.
7. The interface between the qualitative text mining and the quantitative LCOH model functions as an explanatory structural mapping rather than a strict, statistically falsifiable empirical validation. A rigorous empirical validation would require direct cross-layer regressions—for instance, mathematically correlating the longitudinal proportion of strategy sentences mentioning financing constraints with actual green project lending rates or realized WACC metrics in Türkiye over time. Due to the static nature of the current policy corpus and the lack of historical project financing data for Turkish green hydrogen, such regressions were outside the scope of this study. Consequently, the current alignment risks falling into a post hoc narrative synthesis. Future iterations must bridge this gap by gathering empirical financial lending microdata to statistically test whether shifts in textual policy emphasis causally predict shifts in actual project financing costs.

8. The construction of the “Proactive Policy” scenario is based on the use of a deterministic, aggregate cost reduction parameter (fixed €1.5/kg reduction). In actual energy markets, however, the cost effects of different policy measures—carbon tax escalators, regional capital investment grants, special types of CfDs—are extremely non-linear, heterogeneous, and time-dependent for each individual country. The simplification of a process as complicated as that of policy measures into a constant, one-time adjustment reduces the capability of the model to find out which measure works best and when. In future iterations of this framework, the determinism of the policy side will be overcome through the introduction of the stochastic nature of the intervention vector. This will require modeling carbon tax evolution using dynamic geometric Brownian motion, parameterization of CfD strike prices against regional volatility and applying time-dependent learning curves directly to different policy measures.
9. The reported LCOH values represent production-gate costs and do not include compression, liquefaction, storage, transport, or port-related costs; the competitiveness thresholds discussed in this study should therefore be interpreted as production-cost rather than delivered export-cost competitiveness.
10. The TOWS-style matrix (Figure 4) reports scores derived from marginal SWOT frequencies rather than a formal pairwise co-occurrence metric and should be read as an illustrative prioritization heuristic rather than a statistically validated strategic-scoring methodology.

5. Conclusions and Future Outlook

This study proposes and demonstrates a hybrid framework that couples NLP-assisted SWOT analysis of policy documents with stochastic LCOH modeling, applied to 107 expert-annotated sentences from Türkiye’s national hydrogen strategy corpus. The main findings are:

1. Regulatory and financing gaps emerge as the dominant structural weaknesses in the policy discourse, supported by the CA biplot and by the sensitivity analysis, which identifies electricity price—whose effective cost is shaped by regulatory and financial instruments—as the leading driver of LCOH uncertainty (Sobol $S_1 \approx 0.52$).
2. Renewable resource potential and energy strategy are the primary strengths and opportunities, consistent with Türkiye’s geographic endowment and regional export potential.
3. Export/Demand Risk is framed predominantly as an opportunity in policy discourse, signaling a strategic blind spot that requires risk mitigation instruments.
4. Under the business-as-usual scenario assumptions, Türkiye has a less than one-in-five probability of reaching export-competitive LCOH by 2050 ($p = 17.8\%$). The stylized proactive policy intervention raises this probability to 78.4%, quantifying the leverage of well-designed policy instruments (within the scenario framework).

The proposed methodology is country-agnostic and adaptable, offering policymakers a transparent, evidence-based tool for national hydrogen strategy design. By resolving strategic ambiguities (such as Global Competition Risk) and optimizing clean energy investment pathways, this research directly advances SDG 7 (Affordable and Clean Energy), SDG 9 (Industry, Innovation, and Infrastructure), and SDG 13 (Climate Action), mapping a clear route toward a socially, environmentally, and economically sustainable energy landscape.

Looking forward, three high-priority research directions emerge to extend this proof-of-concept framework. First, to overcome the single-country constraint, the pipeline will be expanded to ingest multi-jurisdictional policy datasets, enabling cross-border strategic

benchmarking across the EU and MENA regions. Second, the techno-economic layer will be refined by replacing the default parameter independence assumption with multi-variable Copula functions to capture critical CAPEX and electricity price correlations, alongside hour-by-hour power market dynamics and electrolyzer degradation curves. Finally, future work will transition from analytical to mathematical text-to-cost coupling by embedding structural architectures like BNs or SEM, allowing the pipeline to endogenously map shifts in qualitative policy densities directly into the statistical mean and variance of LCOH input distributions.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/su18157506/s1>, Table S1: Taxonomy, Domain Lexicon, and Text Classification Blueprint; Table S2: Correspondence Analysis coordinates, contributions, and cosine quality metrics.

Author Contributions: Conceptualization, İ.M. and H.Y.; methodology, İ.M. and H.Y.; software, İ.M.; formal analysis, İ.M., H.Y., J.C. and A.P.O.; investigation, İ.M., H.Y., J.C. and A.P.O.; data curation, İ.M.; writing—original draft preparation, İ.M., H.Y., J.C. and A.P.O.; writing—review and editing, İ.M., H.Y., J.C. and A.P.O.; visualization, İ.M., H.Y., J.C. and A.P.O. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bilgiç, G. Türkiye Hidrojen Enerjisi Stratejilerinin ve Geliştirilen Politikaların Araştırılması. *Düzce Üniversitesi Tek. Bilim. Derg.* **2024**, *2*, 70–79. [CrossRef]
2. Enerji, T.C.; Bakanlık, T.K. *Türkiye Hidrojen Teknolojileri Stratejisi ve Yol Haritası*; ETKB: Ankara, Türkiye, 2023. Available online: <https://enerji.gov.tr> (accessed on 4 June 2025).
3. SHURA Enerji Dönüşüm Merkezi. *Türkiye Enerji Dönüşümü Görünümü Raporu*; SHURA Enerji Dönüşüm Merkezi: İstanbul, Türkiye, 2025. Available online: <https://www.shura.org.tr> (accessed on 4 June 2025).
4. SHURA Enerji Dönüşüm Merkezi. *Türkiye’de Yenilenebilir Enerji Tedarik Anlaşmalarının Etkinleştirilmesi*; SHURA Enerji Dönüşüm Merkezi: İstanbul, Türkiye, 2025. Available online: <https://shura.org.tr/wp-content/uploads/2025/07/SHURA-EDM-Turkiyede-Yenilenebilir-Enerji-Tedarik-Anlasmalarinin-Etkinlestirilmesi-1.pdf> (accessed on 4 June 2025).
5. SHURA Enerji Dönüşüm Merkezi. *Türkiye’nin Karbonsuzlaşma Sürecinde Yenilenebilir Hidrojen: Öncelikli Uygulama Alanları ve Politika Önerileri*; SHURA Enerji Dönüşüm Merkezi: İstanbul, Türkiye, 2025. Available online: <https://shura.org.tr/wp-content/uploads/2025/12/SHURA-EDM-Turkiyenin-Karbonsuzlasma-Surecinde-Yenilenebilir-Hidrojen-Oncelikli-Uygulama-Alanlari-ve-Politika-Onerileri-04.02.2025.pdf> (accessed on 4 June 2025).
6. IICEC. *Türkiye Yeşil Hidrojen Geleceği*; Sabancı Üniversitesi İstanbul Uluslararası Enerji ve İklim Merkezi: İstanbul, Türkiye, 2023.
7. Gilardi, F.; Shipan, C.R.; Wüest, B. Policy diffusion: The issue-definition stage. *Am. J. Political Sci.* **2021**, *65*, 21–35. [CrossRef]
8. Wu, J.; Zhao, N.; Yang, T. Wisdom of crowds: SWOT analysis based on hybrid text mining methods using online reviews. *J. Bus. Res.* **2024**, *171*, 114378. [CrossRef]
9. Aggarwal, C.C.; Zhai, C. *Mining Text Data*; Springer: Boston, MA, USA, 2012. [CrossRef]
10. Kowsari, K.; Jafari Meimandi, K.; Heidarysafa, M.; Mendu, S.; Barnes, L.; Brown, D. Text classification algorithms: A survey. *Information* **2019**, *10*, 150. [CrossRef]
11. Gliffy. The Comprehensive Guide to SWOT Analysis. Available online: <https://www.gliffy.com/blog/the-comprehensive-guide-to-swot-analysis> (accessed on 4 June 2025).
12. Jin, Z.; Mihalcea, R. Natural Language Processing for Policymaking. In *Handbook of Computational Social Science for Policy*; Bertoni, E., Fontana, M., Gabrielli, L., Signorelli, S., Vespe, M., Eds.; Springer: Cham, Switzerland, 2023. [CrossRef]

13. Grimmer, J.; Stewart, B.M. Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Anal.* **2013**, *21*, 267–297. [CrossRef]
14. Salton, G.; Buckley, C. Term-weighting approaches in automatic text retrieval. *Inf. Process. Manag.* **1988**, *24*, 513–523. [CrossRef]
15. Ramos, J. Using TF-IDF to determine word relevance in document queries. In *Proceedings of the 1st Instructional Conference on Machine Learning*; Rutgers University: Piscataway, NJ, USA, 2003; Volume 242, pp. 29–48.
16. Shojaei, M. SWOT Analysis of the European Hydrogen Strategies. Master's Thesis, University of Twente, Enschede, The Netherlands, 2023. Available online: <https://essay.utwente.nl/98643/> (accessed on 4 June 2025).
17. Yilmaz, G.; Sadvakasova, A.K.; Kossalbayev, B.D.; Bauenova, M.O.; Zharmukhamedov, S.K.; Ziyayeva, G.K.; Allakhverdiev, S.I. Hydrogen energy development in Turkey: Challenges and opportunities. *Int. J. Hydrogen Energy* **2024**, *52*, 1304–1311. [CrossRef]
18. Furuncu, Y. PESTEL and SWOT analysis of hydrogen economy in Turkey. *Int. J. Hydrogen Energy* **2025**, *113*, 161–169. [CrossRef]
19. Frieden, F.; Leker, J. Future costs of hydrogen: A quantitative review. *Sustain. Energy Fuels* **2024**, *8*, 1806–1822. [CrossRef]
20. Rasul, M.G.; Hazrat, M.A.; Sattar, M.A.; Jahirul, M.I.; Shearer, M.J. Cost-competitiveness of green hydrogen and its sensitivity to major financial and technical variables. *Int. J. Hydrogen Energy* **2024**, *90*, 1365–1377. [CrossRef]
21. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [CrossRef]
22. Greenacre, M. *Correspondence Analysis in Practice*, 3rd ed.; Chapman and Hall/CRC: Boca Raton, FL, USA, 2017. [CrossRef]
23. Pal, A.; Kakran, S.; Kumar, A.; Chandrakanta, A.S.; Ratha, S. Powering squarely into the future: A strategic analysis of hydrogen energy in QUAD nations. *Int. J. Hydrogen Energy* **2024**, *49*, 16–41. [CrossRef]
24. Akan, F. *Küresel Hidrojen Yarışı ve AB'nin Yeşil Hidrojen Stratejisi*; Değerlendirme Notu No. 312; İktisadi Kalkınma Vakfı: İstanbul, Türkiye, 2025.
25. TSKB Enerji Çalışma Grubu. *Hidrojen Enerjisi Bilgilendirme Notu Güncellemesi*; Türkiye Sınai Kalkınma Bankası: İstanbul, Türkiye, 2023.
26. Türk, S.; Engin, Y. *Hidrojen ve Kullanım Potansiyeli*; Türk Çimento Sanayi İşverenleri Sendikası: Ankara, Türkiye, 2024.
27. T.C. Enerji ve Tabii Kaynaklar Bakanlığı (ETKB). *Hidrojen Yaklaşım Belgesi*; 2020/4B; ETKB: Ankara, Türkiye, 2020.
28. Hidrojen Teknolojileri Derneği. *Türkiye için Hidrojen Teknolojileri Yol Haritası Raporu 2021*; Hidrojen Teknolojileri Derneği: İstanbul, Türkiye, 2021. Available online: <https://www.hidrojenteknolojileri.org> (accessed on 4 June 2025).
29. Dünya Enerji Konseyi Türk Milli Komitesi. *Yeşil Hidrojen Raporu*; Dünya Enerji Konseyi Türk Milli Komitesi: Ankara, Türkiye, 2020.
30. Akin, A.A.; Akin, M.D. An open source natural language processing library for Turkic languages: Zemberek. *Elektr. Mühendisliği* **2007**, *431*, 38–44.
31. Cohen, J. A coefficient of agreement for nominal scales. *Educ. Psychol. Meas.* **1960**, *20*, 37–46. [CrossRef]
32. Landis, J.R.; Koch, G.G. The measurement of observer agreement for categorical data. *Biometrics* **1977**, *33*, 159–174. [CrossRef]
33. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
34. Rousseeuw, P.J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **1987**, *20*, 53–65. [CrossRef]
35. McQueen, J.B. Some methods of classification and analysis of multivariate observations. In *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*; University of California Press: Berkeley, CA, USA, 1967; pp. 281–297.
36. Schweter, S. BERTurk: A Large-Scale Pre-Trained Language Model for Turkish. GitHub Repository. 2020. Available online: <https://github.com/stefan-it/turkish-bert> (accessed on 4 June 2025).
37. Greenacre, M.; Groenen, P.J.F.; Hastie, T.; D'Enza, A.I.; Markos, A.; Tuzhilina, E. Principal component analysis. *Nat. Rev. Methods Primers* **2022**, *2*, 100. [CrossRef]
38. International Energy Agency (IEA). *Global Hydrogen Review 2024*; IEA: Paris, France, 2024. Available online: <https://iea.blob.core.windows.net/assets/89c1e382-dc59-46ca-aa47-9f7d41531ab5/GlobalHydrogenReview2024.pdf> (accessed on 4 June 2025).
39. IRENA. *Green Hydrogen Cost Reduction: Scaling Up Electrolysers to Meet the 1.5 °C Climate Goal*; International Renewable Energy Agency: Abu Dhabi, United Arab Emirates, 2024.
40. Brändle, G.; Schönfisch, M.; Schulte, S. Estimating long-term global supply costs for low-carbon hydrogen. *Appl. Energy* **2021**, *302*, 117481. [CrossRef]
41. Rahman, M.N.; Wahid, M.A. Green hydrogen prospects in Peninsular Malaysia: A techno-economic analysis via Monte Carlo simulations. *Future Sustain.* **2024**, *2*, 27–45. [CrossRef]
42. Hubert, M.; Peterson, D.; Miller, E.; Vickers, J.; Mow, R.; Howe, C. *Clean Hydrogen Production Cost Scenarios with PEM Electrolyzer Technology*; DOE Hydrogen Program Record 24005; U.S. Department of Energy: Washington, DC, USA, 2024.
43. Penttinen, S.-L. Navigating the hydrogen landscape: An analysis of hydrogen support mechanisms in the US and the EU. *Rev. Eur. Comp. Int. Environ. Law* **2024**, *33*, 397–411. [CrossRef]
44. Bade, S.O.; Tomomewo, O.S. A review of governance strategies, policy measures, and regulatory framework for hydrogen energy in the United States. *Int. J. Hydrogen Energy* **2024**, *78*, 1363–1381. [CrossRef]

45. Zubairu, N.; Al Jabri, L.; Rejeb, A. A review of hydrogen energy in renewable energy supply chain finance. *Discov. Sustain.* **2025**, *6*, 186. [\[CrossRef\]](#)
46. Jeda.ai. SWOT Analysis with AI. Available online: <https://www.jeda.ai/resources/swot-analysis-with-ai> (accessed on 4 June 2025).
47. Miro. SWOT Analysis Template. Available online: <https://miro.com/templates/swot-analysis/> (accessed on 4 June 2025).
48. van Duin, J.; van Leeuwen, M.T.B.; van der Knaap, W.P.J. The role of the port of Rotterdam in the global hydrogen supply chain: A review and a roadmap. *SSRN* 2024, *in press*. [\[CrossRef\]](#)
49. Terrapon-Pfaeffle, J.-B.; Braun, J.; Ersoy, N.; Prantner, T.; Kern, M.; Viebahn, P. Country risk impacts on export costs of green hydrogen and its synthetic downstream products from the Middle East and North Africa. *Front. Energy Res.* **2025**, *13*, 1546876. [\[CrossRef\]](#)
50. Qi, W.; Zhu, M.; Zhao, S.; Gui, F.; Yu, L.; Guo, Y.; Zhou, Z.; Ge, S.; He, Z.; Zhang, Y.; et al. Performance enhancement of proton exchange membrane fuel cells based on semi-parallel and semi-interdigitated mixed flow field. *Int. J. Hydrogen Energy* **2026**, *231*, 154871. [\[CrossRef\]](#)
51. Song, Q.; Yang, S.; Li, H.; Ji, Y.; Cai, D.; Wang, G.; Liufu, Y. Performance Improvement of Proton Exchange Membrane Fuel Cell by a New Coupling Channel in Bipolar Plate. *Energies* **2025**, *18*, 4068. [\[CrossRef\]](#)
52. Sheng, L.; Lan, H.; Tao, B.; Zhang, Z.; Wang, H. Anisotropic thermal conductivity quantification in prismatic lithium-ion batteries. *J. Energy Storage* **2026**, *162*, 122112. [\[CrossRef\]](#)
53. Lei, A.; Huaiyu, L.; Bojun, T.; Chunfeng, Z.; Xiaojun, Z.; Liyang, W. Experiments on liquid-immersed thermal runaway management for large energy-stored lithium-ion battery modules. *Therm. Sci. Eng. Prog.* **2025**, *68*, 104347. [\[CrossRef\]](#)
54. Qi, W.; Liu, Y.; Wang, X.; Ge, S.; Yu, J.; He, Z.; Tuo, J.; Wang, H.; Zhou, S.; Jin, C.; et al. Performance enhancement of battery thermal management based on liquid cooling plate channels embedded with non-uniform spiral fin structures. *Appl. Therm. Eng.* **2026**, *300*, 131501. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.