



Traceable Normative Computation as a Research Architecture

Proactive regulatory discovery with JudgeAI

Research study Revised 20 September 2026

Abstract

JudgeAI makes a normative decision a computable, traceable research object. This study examines how autonomous orchestration uses that object to identify empirical questions, investigate consequences, and construct regulatory alternatives. The analysis develops decision regions, sufficient empirical information, and alternative expansion, then tests these mechanisms in the United Arab Emirates. The original corpus contains 9,012 scenario configurations, 44,316 evaluated worlds, and 92,700 pairwise comparisons. Further experiments replay 700 cases, resolve comparison boundaries through 1,149 additional native-engine evaluations, and examine every subset of four research variables. Removing inter-actor propagation changes the selected set in 176 of 700 cases; replacing the comparator with a sum changes it in 392. Three inputs determine selection throughout the tested grid. Expanding a two-option cooling menu to six evidence-informed alternatives produces strictly preferred outcomes in 128 of 3,456 conditional worlds. A separate broad-search experiment adds 12 series containing 285 scenario records across 16 threat directions. It examines the earlier task of selecting what to investigate, follows cardiovascular and climate-related branches into local mechanisms, and recomputes all directions after adding human contexts. Six of 21 local actions are conditionally preferred, although all absolute Threat Index profiles remain indistinguishable at the specified tolerance. Together, the studies connect research-agenda formation, empirical refinement, alternative discovery, and repeated normative comparison. The regulatory applications concern recovery services, medical-electricity continuity, and executable financing. These results identify mechanisms through which systematic inquiry can improve normative decisions; the broad search establishes a reproducible research cycle rather than an empirically identified ranking of fundamental threats.

Keywords: computational legal theory; normative ontology; traceability; autonomous orchestration; regulatory discovery; empirical decision dependence; research-agenda formation; counterfactual analysis.

1. Research question and contribution

A system that produces a plausible legal proposal has not yet identified what would make that proposal better. In normative work, revision requires a relationship between an intervention, its consequences for affected people, and the grounds on which competing interventions are compared. More fluent argument, additional citations, or greater agreement among generated answers do not themselves specify that relationship. JudgeAI addresses it by making normative choice fully computable and traceable within an explicit representation of reality.

The research question is: *how can a computable normative decision become the basis for sustained autonomous investigation that discovers better regulatory alternatives?* The object of improvement is the decision and its empirical foundations. The agent investigates which circumstances matter, which causal

relationships determine the ordering of alternatives, and which changes in implementation could produce a preferable world. The normative engine carries out the comparison under its computational legal theory. The interaction is reciprocal. Computation gives empirical research a target: a consequence, dependency, or alternative capable of changing the result. Research gives computation a more developed object: better specified inputs, additional causal distinctions, and interventions that were absent from the initial menu. A technology, contractual arrangement, or organisation of service delivery can therefore enter the analysis before it has a familiar legislative name.

This study makes five contributions. First, it specifies the information that a normative trace must expose for this process to operate. Second, it formalises empirical sufficiency and decision robustness without reducing normative choice to an aggregate monetary score. Third, it measures the consequences of removing dependencies, participants’ burdens, and alternative designs from otherwise identical comparisons. Fourth, it follows the computation through to the content of regulation: the object purchased, the conditions of eligibility, the evidence of performance, and the financing sequence. Fifth, a separate broad-search experiment examines how an initial threat map generates empirical research directions and how local findings return to the common model. This additional experiment extends the architecture to the formation of the task itself.

The UAE provides the application setting. Existing UAE duties and service arrangements define the counterfactual. Empirical studies identify mechanisms and technical possibilities. The discovery process does not depend on finding a law elsewhere that can be copied. Its starting point is a consequential change in the represented world.

2. Why traceability changes the capacity for improvement

2.1 The normative decision as an executable research object

Let $\mathcal{O}$ denote the normative ontology, E the empirical record, A the available interventions, and x the consequence parameters supported or investigated by that record. The computational representation supplies a world for each intervention,

$$W_a = \Phi(\mathcal{O}, E, x, a), \quad a \in A, \quad (1)$$

and the normative engine returns a preferred set

$$\mathcal{C}(E, A) = \arg \min_{a \in A}^{\preceq_N} V(W_a, W_0), \quad (2)$$

where W_0 is the shared reference world, V denotes the participant-level consequence representation, and $\preceq_N$ is the engine’s normative comparison. A set is retained when several alternatives are equivalent under that comparison. The argmin notation presupposes a well-defined preferred set in the comparison under study; it does not assert that every tolerance-based implementation induces a strict total order. The broad-search experiment therefore retains and audits pairwise relations separately (Section 5).

The empirical representation and normative evaluation form one decision procedure. A change in evidence changes the consequences on which the normative ordering operates. The empirical record includes whose circumstances change, the event and time horizon being represented, the conditional relationships involved, and the implementation of each alternative. These elements determine the content of the normative choice, including which protections take priority in the concrete case.

For autonomous research, the trace must connect a selected alternative to the participant-level consequences responsible for its position; those consequences to their direct or propagated origins; and those origins to empirical claims, assumptions, and intervention specifications. This produces an addressable research object. The agent can ask whether a decisive relationship has the sign assumed, whether an apparent

benefit reaches the intended person, whether financing prevents implementation, or whether an additional alternative removes the adverse pathway.

A textual explanation can communicate these relationships. An executable trace additionally permits them to be changed and recomputed. That property is what supports systematic self-improvement: the system can inspect and experimentally vary the grounds of its own answer.

2.2 Decision-relevant information and a necessity result

Let Θ be the set of empirical states currently compatible with an inquiry and let $C(\theta)$ be the preferred set at state θ . A representation $z(\theta)$ is *decision-sufficient* on Θ if

$$z(\theta_1) = z(\theta_2) \implies C(\theta_1) = C(\theta_2) \quad \text{for all } \theta_1, \theta_2 \in \Theta. \quad (3)$$

This criterion asks whether the information retained by the system preserves every distinction that changes the decision.

Proposition 1 (Indistinguishable observations). *If two empirical states have the same observed representation but different preferred sets, no decision procedure using only that representation can recover the correct preferred set in both states.*

Proof. A deterministic procedure receives identical input in the two states and therefore returns identical output. That output cannot equal two different preferred sets. Randomisation does not supply the missing distinction and cannot ensure correct recovery in both states. $\square$

The practical consequence is a research obligation: identify and observe the missing distinction. In the medical-continuity example below, separate distributions of outage duration and reserve capacity are identical in two constructed populations, while the probability of losing the essential service differs. The required research is a joint measurement on the same device and event. More precise marginal averages cannot recover that distinction.

Traceability is therefore operational infrastructure for empirical learning. It tells the agent what its current representation fails to distinguish and connects that distinction to the final normative choice. This is the foundation for targeted investigation, including investigation of evidence that could overturn the current answer.

2.3 Decision regions and robust conclusions

Define a decision region $R_S = \{\theta : C(\theta) = S\}$ for a preferred set S . If all states still compatible with the evidence lie in R_S , the preferred set is identified on that evidence. Further refinement inside the same region can improve consequence estimates while preserving the choice. If the compatible states cross a boundary, the trace indicates which empirical distinction can resolve it.

There is a second useful question when the exact preferred set remains variable. An intervention is robustly preferred on Θ when it belongs to

$$A_{\text{robust}}(\Theta) = \bigcap_{\theta \in \Theta} C(\theta). \quad (4)$$

A non-empty intersection identifies an option that remains preferred throughout the stated uncertainty set. An empty intersection directs research to the conflicting states. Both operations support inquiry using finite scenarios and partial information.

For a fixed representation and alternative set, when valid new evidence narrows Θ to a non-empty subset Θ' , an already identified choice remains identified, and $A_{\text{robust}}(\Theta) \subseteq A_{\text{robust}}(\Theta')$. These are consequences of set inclusion. Discovery of a previously omitted mechanism is a different operation: it revises the

representation and requires recomputation. JudgeAI’s trace separates refinement of an existing empirical question from expansion of the questions being asked.

2.4 Forming the research agenda

The same architecture can investigate which normative problem to pursue before a particular intervention is selected. Let G_t be a map of threat directions and let X_t calibrate a common consequence space for the jurisdiction. Perturbations associated with a direction $g \in G_t$ reveal which actor-level consequences and empirical assumptions can alter the comparison. The resulting research state can be written as

$$\mathcal{S}_t = (\mathcal{O}_t, E_t, X_t, A_t, U_t), \quad q_t = \Pi(G_t, \mathcal{S}_t, \tau_t), \quad \mathcal{S}_{t+1} = \text{Update}(\mathcal{S}_t, \text{Investigate}(q_t)), \quad (5)$$

where U_t records investigated uncertainties, τ_t is the computational trace, and Π denotes the orchestrator’s selection of the next inquiry. Contextual updates enrich the representation of reality while retaining the normative foundation. The revised state must be evaluated again across the shared alternatives.

This operation separates research-agenda formation from regulatory adoption. A direction can merit investigation because its effects are sensitive to a missing observation, because it exposes a systemic dependency, or because evidence makes a consequential mechanism testable. Its selection for inquiry does not establish its final normative priority. In the broad-search experiment below, Π is implemented by recorded agent research operations; an optimal information-acquisition policy is not assumed.

3. Relationship to contemporary AI research

The relevant comparison concerns where the feedback for improvement comes from and what is revised. Self-Refine uses an LLM to generate, critique, and revise its own output [1]. JudgeAI’s research cycle adds an independently executable object: the represented consequences and their normative comparison. An objection can be turned into a changed world, a missing empirical coordinate, or an additional intervention and then evaluated under the same rule.

Decision-focused learning connects prediction to downstream optimisation, training predictive models for the decisions that their outputs support [2]. Targeted active learning goes further in the direction of research selection: Filstroff et al. select information to reduce uncertainty about the optimal decision [3]. Causal active learning investigates interventions through an explicit causal model [4]. These works establish important methodological connections. JudgeAI applies the connection to a structured normative object with participant-specific consequences, an explicit normative ordering, and a changing set of legal and implementation alternatives.

Anthropic’s July 2026 study examines values expressed across models and languages, using axes extracted from conversational responses [5]. Its stated approach develops context-sensitive judgment, consistent with the broader account in Claude’s constitution [6]. JudgeAI places the grounds of a concrete normative comparison in a computable representation. Empirical changes can be traced through that representation to the changed choice, and the agent can investigate those dependencies directly. The distinction is the location and form of the evaluative operation: expressed judgment in generated responses versus an executable comparison of represented consequences.

The distinctive contribution studied here is the integration of these operations around a computable normative decision. Ontology specifies what a consequential claim refers to; computation determines how it enters a comparison; orchestration investigates it and constructs alternatives. The comparison results reported below isolate these mechanisms on shared inputs.

Table 1: Research approaches and the object through which improvement occurs.

Approach	Object revised or studied	Feedback used for improvement
Iterative self-feedback	Generated answer	Critique and revised generation
Decision-focused learning	Predictive parameters	Performance of downstream decisions
Decision-aware active learning	Information acquisition	Uncertainty about the optimal decision
Contextual judgment and value analysis	Behaviour and expressed values	Training, contextual responses, behavioural measurement
JudgeAI normative inquiry	Empirical world representation, alternatives, and implementation	Traced consequences and repeatable normative comparisons

4. Experimental design and computational basis

4.1 Ontology and native engine

The UAE inquiry examines relationships among participants, their circumstances, and the consequences of interventions. The ontology connects employment, accommodation, essential services, and financing within a common analytical structure. It supplies research directions that cross conventional sectoral boundaries and retains the empirical conditions needed to establish each consequential relationship.

The simulations use JudgeAI’s normative engine to propagate consequences and compare alternative worlds. Each comparison uses a shared baseline and the same represented participants and outcomes. Inputs retain their measured or scenario status, and the normative criterion remains fixed while factual assumptions and intervention designs vary.

The engine evaluates participant-level consequence profiles under JudgeAI’s computational legal theory. It preserves the significance of adverse consequences and the distribution of improvements across participants. This enables the controlled experiments below: the same alternatives can be compared with and without particular dependencies, with financing consequences included or omitted, and against an aggregate calculation. The differences identify which features of the normative representation materially affect the result.

4.2 Intervention corpus and controlled experiments

Table 2: Original UAE scenario corpus. Counts are computational configurations and evaluations.

Series	Configurations	Worlds	Pairs
Ontology-linked frontier	2,100	6,300	6,300
Medical service continuity	3,456	17,280	34,560
Recovery and implementation	3,456	20,736	51,840
Total	9,012	44,316	92,700

The mechanism analysis uses the 700 worker-recovery configurations from the frontier series. It varies four decision-relevant inputs: the direct effect, the inter-actor response, the payer consequence, and a common implementation multiplier. These inputs correspond to distinct empirical questions about the intervention and its consequences.

The new analysis contains four computational tests. First, all 700 cases are replayed through the native engine, with exact equality of the stored profile vectors and preferred sets. Second, three component ablations are evaluated on their identical inputs. Third, all 16 subsets of the four research variables are checked for decision sufficiency. Fourth, a dense sensitivity study adds 1,149 native comparisons, each

evaluating three worlds and their three pairs. The alternative-expansion test then reuses the 3,456 recorded recovery states and compares restricted menus under the native comparator.

The constructed scenarios specify conditional states of the world. Empirical studies enter with their measured endpoints and populations; they determine relevant mechanisms, plausible interventions, and the next measurements. Scenario frequencies below describe the enumerated design. This study uses controlled computational comparisons to identify how the architecture changes what can be investigated and selected.

4.3 Separate design for broad threat exploration

The additional experiment starts with 16 threat directions in a shared UAE consequence representation: cardiovascular/metabolic burden, cancer, air pollution, infection and antimicrobial resistance, mental/neurological health, heat, water, energy, food/ecosystems, flooding/coastal exposure, injury, violence/conflict, material deprivation, healthcare, institutional protection, and the information environment. Global climate and systemic-risk assessments help define this scope [15, 16, 17]. They do not directly supply individual consequence probabilities.

The experiment contains 12 series and 285 world records, including 24 repeated reference and unchanged-control records (Table 3). These records are kept separate from the intervention corpus and are not independent empirical observations. The numerical basis and direct-response assumptions are versioned within this experiment. It uses an isolated research copy with a documented correction to the scaling of direct updates; the normative comparator and TI formulae are retained. Appendix B records the verification scope.

Table 3: Additional broad-search experiment. Counts are series and stored scenario records.

Experiment family	Series	World records
Increasing/decreasing burdens, 1/10/30-year horizons	3	102
Equal initial dose, 10/30-year horizons	2	36
Compound resource stress	1	18
Baseline-calibration and response sensitivity	4	88
Macro recomputation after contextual refinement	1	18
Local alternatives and controls	1	23
Total	12	285

Threat directions transform shared consequence coordinates rather than create independent copies of harm. Overlapping assumed reductions are combined using the largest declared reduction unless an interaction is separately specified. Horizons describe distinct prospective states; annual risks are not repeatedly added. Equal-dose perturbations control the inserted magnitude, not the real-world effectiveness of different interventions. Numerical intensities and responses retain their scenario status.

5. From a threat space to an empirical research agenda

5.1 Broad comparison identifies several research branches

The native comparison first examines complete-field profiles of positive and signed changes by actor, then changes in TI relative to the common anchor, and finally absolute TI profiles. The numerical tolerance is $\varepsilon = 10^{-8}$. Absolute profiles are also compared separately as a diagnostic; this does not replace the native sequence.

In proportional-load experiments at 1-, 10-, and 30-year horizons, the native comparison selects the mental/neurological direction, whereas the absolute-TI diagnostic selects the cardiovascular/metabolic direction. Equal initial doses at 10 and 30 years instead select infection and antimicrobial resistance under the native comparison. Thus, the broad stage reveals dependence on the perturbation design rather than a single established hierarchy of threats.

The cardiovascular branch is retained for empirical investigation because its absolute-profile result and available observations support a specific next question. The climate/resource branch is retained as a systemic dependency linking environmental exposure, essential services, and individual protection. It is not presented as the numerical winner of the main comparison. Mental/neurological and infectious threats remain in the research queue. This branching records the agent's research strategy; it does not establish that the chosen branches are the optimal use of research resources.

5.2 Evidence changes the local question

The cardiovascular inquiry proceeds from severe disease outcomes to the hypertension-care cascade. The WHO UAE profile reports 2024 estimates of 35% hypertension prevalence among adults aged 30–79; among adults with hypertension, 37% are diagnosed, 31% treated, and 18% controlled [18]. These denominators describe population states, not longitudinal transitions in one cohort. Existing programmes, including Nabdak [19], make the next question more specific: which additional organisation of confirmation, treatment, follow-up, and reassessment could improve completion of care relative to current practice?

A UAE randomised telepharmacy study involving 239 patients in 16 pharmacies informs investigation of that mechanism [20]. It tests a service package rather than a new payment rule; the 12-month between-group systolic-pressure comparison has $p = 0.06$. The inquiry therefore produces a conditional care-completion alternative and requests linked pathway data. It does not infer a mortality reduction or a regulatory effect from the trial's surrogate outcomes.

The climate inquiry moves from heat to effective cooling, water access, and safe restoration following a local failure. Existing Abu Dhabi heat and accommodation requirements define the comparator [21, 22]. Network reliability, premises equipment, and protection actually reaching a person are distinct. The DEWA reliability statistic used in the medical-continuity analysis also illustrates this distinction: average network performance does not identify local cooling failure or exposure duration. The resulting candidate combines temporary protection with verified restoration and explicit repair-worker consequences. A residual protection gap must be established in the relevant setting.

5.3 Context refinement returns to the common calculation

The empirical descent adds six conditional human contexts: a patient with uncontrolled hypertension; a person affected by cooling and water failure; household providers with and without a material reserve; another patient whose care is displaced; and a repair worker. They are reference situations, not observed population cohorts. Primary consequence states increase from 20 to 40. All 16 macro directions are then recomputed on the expanded basis before evaluating the local alternatives.

Financing enters through the consequences borne by a person. A payment cannot be converted directly into a TI effect without an account of reserves, obligations, consumption, care, or other affected conditions. Reliability is held fixed between local alternatives, while effects and burden incidence vary. Zero-benefit, adverse-effect, care-displacement, and repair-exposure controls test the assumed mechanisms.

Of 21 substantive local actions, six are preferred to the current state, two are tied, and 13 are worse. The local series contains 23 worlds after adding the current state and its unchanged copy. A zero-benefit option with a small burden ties with the reference; a larger burden makes it worse. Both adverse-effect controls

also lose. Benefits can therefore justify some modelled burdens, while their incidence can reverse the comparison.

Table 4: Selected results of the separate broad-search and context-refinement experiment.

Condition	Native comparison	Absolute TI diagnostic
Proportional burden changes, three horizons	Mental/neurological direction	Cardiovascular/metabolic direction
Equal-dose perturbations, two horizons	Infection/antimicrobial resistance	A profile ties with every other; some other pairs differ
All directions after context refinement	Material deprivation	All 18 profiles pairwise tied
Final local alternatives	6 preferred, 2 tied, 13 worse than the reference	All 23 profiles pairwise tied

The selected cooling/water scenario assumes reductions of 25% in adverse thermal intensity, 17.5% in water-related intensity, and 12.5% in impaired recovery, with an additional material-threat intensity of 0.005 for the provider with reserves. Its first decisive complete-field difference is -1.53995×10^{-7} , approximately -15.40ϵ . These are conditional inputs and a numerical comparison, not estimated intervention effects or monetary costs. All absolute TI profiles remain pairwise indistinguishable, so the preference does not establish a distinguishable reduction in absolute minimax TI.

The revised macro selection changes to material deprivation. This demonstrates sensitivity to the expanded representation, not an observed change in the jurisdiction’s fundamental priority. Absolute levels from differently constituted bases cannot measure societal improvement. In several non-leading macro alternatives, pairwise checks also reveal non-transitive ties and strict preference cycles. Leading native candidates pass the no-pairwise-loss check; the remaining linear display is not interpreted as a strict ranking. No such cycles occur in the final local set.

5.4 What the additional experiment establishes

The experiment completes a bounded cycle from general threats through empirical questions and local alternatives back to the common model. Its numerical propagation uses direct state perturbations and ontology geometry; there are no active typed empirical causal transitions in this experimental layer. The sources establish observations, mechanisms, and existing duties, while the stated response magnitudes remain hypotheses.

The result supplies a next measurement programme: linked hypertension pathways, event-level cooling and water failures, actual temporary protection, restoration conditions, and the joint incidence of costs and displaced services. These measurements can revise the shared calibration and candidate effects, followed by another full comparison. Optimal research prioritisation, population-wide threat calibration, and real-world efficacy remain unestablished by this finite experiment. Its contribution is the traceable formation and revision of research tasks within the architecture tested by the intervention studies that follow.

6. How empirical dependencies change regulation

6.1 Empirical conditions change the preferred intervention

The frontier experiment compares the reference arrangement, a targeted intervention with additional enterprise protection, and a broad intervention. The direct effect and payer consequence are held fixed while the worker-to-enterprise relationship changes. Three recorded conditions produce three different preferred regulatory configurations.

Table 5: Changes in the represented relationship change the preferred regulation.

Condition	Consequential relationship	Native selection
Adverse	Both interventions introduce an adverse enterprise consequence	Reference arrangement
Intermediate	Additional enterprise protection changes the consequences of targeted delivery	Targeted intervention with protection
Favourable	Broad delivery improves worker and enterprise outcomes	Broad intervention

The calculation identifies an empirical research task: establish the actual relationship between worker recovery and the enterprise’s operating conditions. Inquiry can then distinguish the conditions supporting targeted protection, broader delivery, or retention of the reference arrangement. Each change in choice is traceable to the consequences that distinguish the alternatives.

A separate payer sensitivity test changes the research direction. In the tested three-option menu, an added payer threat makes the reference preferred. The next substantive question becomes whether a different financing or implementation arrangement can deliver the intervention while avoiding the identified adverse consequence. Such an arrangement enters as an additional world for comparison.

6.2 Distribution of consequences affects the result

A recorded case gives the targeted intervention a larger aggregate reduction in the represented threats while the native criterion selects broad delivery. The interventions distribute their benefits differently between the worker and the enterprise. Their preferred ordering therefore depends on the participant-level normative comparison, rather than the total alone.

In another recorded case, the targeted intervention benefits the worker while introducing an adverse enterprise consequence. The aggregate calculation selects that intervention; the native comparison retains the reference. Together, the examples show why the consequences for each participant remain available throughout the inquiry. Research can investigate the particular dependency responsible for the divergence and construct an intervention that addresses it.

6.3 Component ablations

The direct-only ablation evaluates consequences before their propagation between participants. The payer ablation sets the payer consequence to zero while preserving the remaining coordinates and comparator. The scalar ablation sums the same post-propagation coordinates and minimises that sum. Each ablation retains the same alternatives and scenario inputs.

Table 6: Changed preferred sets under controlled ablations, compared with full native evaluation.

Ablation	Changed / 700	Mechanism isolated
Remove inter-actor propagation	176 (25.1%)	Secondary consequences alter the choice
Omit payer consequence	384 (54.9%)	Distribution of financing consequences matters
Replace normative vector with a sum	392 (56.0%)	Participant-level ordering differs from aggregation

In 388 of the 392 scalar disagreements, the preferred sets are disjoint. The remaining four are changes in equivalence structure. These are deterministic comparisons of representations on the same grid. They




Figure 1: Controlled component comparisons on the same 700 scenario inputs. Bars count changes in the preferred set when a specified part of the evaluation is removed or replaced.

locate the work performed by the ontology-linked propagation and normative comparator: the system’s conclusion depends on specific relationships and participant-level effects that simpler calculations discard.

7. From sensitivity to a programme of empirical inquiry

7.1 What the system needs to know

For a chosen subset Q of the four parameters, group all 700 states that agree on Q . A state is resolved by Q when every member of its group has the same preferred set. This directly implements Equation (3). The test identifies which combinations of inputs determine the preferred set throughout the scenario grid.

Table 7: Information sufficient to identify the preferred set in the scenario grid.

Coordinates known	Resolved states / 700	Largest remaining number of choice sets
Direct effect a	0	3
Inter-actor response λ	0	4
Payer consequence p	280	4
Common multiplier f	0	4
a and λ	80	3
a and p	364	3
λ and p	380	2
a , λ , and p	700	1

The result is more informative than a global sensitivity ranking. Payer information alone resolves the 280 positive-payer cases in this menu. Other states require linked information about the response and direct effect. The smallest state-specific certificates use one coordinate in 280 states, two in 212, and three in 208. None requires all four.

The common multiplier provides a negative control: it changes consequence magnitudes while leaving the preferred sets unchanged throughout the tested grid. This result distinguishes a change in the estimated size of effects from a change in the choice itself. The system can direct further inquiry toward variables that distinguish regulatory alternatives, while retaining information about scale for implementation planning. The observed invariance concerns the common multiplier used in this experiment; technology-specific delivery conditions remain separate research variables.

7.2 From a numeric coordinate to an empirical request

The trace translates each decision-relevant input into a concrete empirical research request. In this case, it identifies a worker, the actual recovery period, the employer relationship, and the operational or financial endpoint. The required observations join the worker's recovery conditions to absence, continuity, output, and enterprise cash constraints over a common horizon. Alternative explanations include shift patterns, seasonal heat, workforce changes, and intervention costs.




Figure 2: An ontology-bound research question. The relationship between worker recovery and enterprise activity identifies the subject of empirical inquiry. Financing consequences remain represented separately. The diagram presents the case relationships at the level of the research question.

A payer coordinate generates a different request: the identity of the payer, contractual cash dates, mandatory obligations, available liquidity, displaced services, and the distribution of costs passed to others. The ontology connects expenditure to its consequences for the payer's liquidity, obligations, and affected services.

The research cycle can also discriminate between conflicting studies. If the disagreement leaves all compatible states inside the same decision region, it can be recorded without suspending the choice. If it crosses a boundary, the system has an explicit reason to investigate design, population, endpoint, or causal identification more deeply. The significance of the conflict is established through its effect on the normative result.

8. Discovery through expansion of regulatory alternatives

8.1 Evidence changes the object of intervention

The recovery investigation connects the actual sleep environment, access to appropriate care, and the continuity of economic activity. Existing Abu Dhabi fatigue provisions already address night-worker accommodation and sleep conditions [7]. The additional regulatory object developed through the inquiry is a verified recovery service: the actual delivered conditions, technical maintenance, and appropriate access to assessment, linked to a financeable delivery arrangement.

The empirical literature separates mechanisms that a general recommendation to improve sleep would combine. A randomised study of 452 adults found that the night-sleep interventions added approximately 27 minutes, without a detectable productivity gain from the pooled night treatment [8]. A 15-person crossover

study found that a conductive mattress topper increased sleep by 21.4 minutes under the tested 32-degree condition, while the neutral-temperature condition did not show the same duration benefit [9]. A station-randomised programme involving 1,189 firefighters combined education, screening, and referral; disability and injury-related absence averaged 1.4 versus 2.6 24-hour pay periods per person per year [10].

These results identify different intervention objects and endpoints. Thermal relief, sleep duration, access to assessment, and occupational absence require their own causal links. The local Abu Dhabi study of 36,515 full-time workers supplies relevant joint domains, including sleep, housing satisfaction, work conditions, and health [11]. Together, the evidence develops a set of candidate mechanisms and a specification for targeted local measurements.

The six compared configurations are the existing regime, additional general air conditioning, an efficient retrofit, a device targeted to overheated sleepers, diagnosis and recovery-care access, and a retrofit combined with that access. Each world specifies whether a recovery deficiency exists, whether the technology works, whether the sleep environment is preserved, whether an assessed need and referral capacity exist, and whether the alternative can actually be financed. Employer, payer, and additional clinical consequences remain represented separately.

8.2 Testing the value of a richer alternative set

Two restricted menus were evaluated against the full six-option set. The first retained only the existing regime and additional general air conditioning. The second retained every individual intervention while removing the combined retrofit-and-access option. Consequences were held fixed; the native comparator determined whether an omitted configuration was strictly preferable to the best retained alternative.

Table 8: Strict improvements discovered by expanding the alternative set.

Starting menu	Improved states / 3,456	Concrete witness
Reference and general cooling	128	Savings-backed retrofit is executable where general cooling is not
Reference and all individual interventions	5	Joint recovery and access deficits favour the combined package

The combined-deficiency witness is particularly instructive. With effective implementation, available referral, upfront finance, and equal additional side effects, the four individual interventions are equivalent under the comparison: each addresses one of two specified deficits. The combined package addresses both and becomes the unique preferred option. When funding is restricted to verified energy savings, the efficient retrofit is the executable preferred intervention. The research therefore changes both the substantive package and the conditions under which it should be chosen.

Within the ordered comparison setting used for these menu tests, adding alternatives can preserve the existing best value or improve it, provided the original alternatives and their consequences are retained. The observed strict improvements establish this effect on the tested menus; the statement is not an assertion of a total order for every tolerance-based comparison discussed in Section 5. This property supports systematic exploration beyond the initial policy menu. An empirically motivated intervention can enter, be compared, and displace that menu's best option.

8.3 The resulting legal mechanism

A recovery-service contract can specify the actual recovery interval, the deficiency to be remedied, the performance condition to be maintained, and the measurement needed for acceptance. Bidders can propose building improvements, equipment, operational arrangements, or a combination. Clinical access is

included when an assessed unmet need and delivery capacity make it consequential. Energy performance is measured at a preserved recovery service, preventing apparent savings obtained by degrading the contracted outcome.

The legal object thereby follows the causal object. Procurement pays for a maintained improvement in conditions rather than treating the purchase of a named device as sufficient performance. This supports innovation because suppliers can propose different means of achieving the verified result. Comparison remains open to those different means, their adverse consequences, and their financing requirements.

9. Medical continuity and the need for joint empirical data

9.1 Deriving the relevant endpoint

For a person dependent on an electrical medical device, the relevant object is the continued operation of that device under the required load. The network’s average reliability is a contextual fact. DEWA reports 0.82 customer-minutes lost in 2025 [12]. The scenario model therefore investigates event-level supply, including building-internal faults, usable reserve, and time to a functioning compatible substitute.

Let D be outage duration, B tested reserve autonomy, and T time until substitute supply is operational. Under uninterrupted switching and working compatible equipment, a loss of the essential electrical service occurs when

$$H(D, B, T) = \mathbf{1}\{\min(D, T) > B\}. \quad (6)$$

The gap duration is $[\min(D, T) - B]_+$. This physical relation identifies a joint empirical object. It asks about the same event, load, reserve state, and functioning substitute, rather than three unrelated averages.

9.2 A constructive demonstration of missing information

Consider two equally weighted event types with duration values of 30 and 120 minutes, reserve values of 60 and 180 minutes, and substitute time of 120 minutes. Both constructed populations have identical duration and reserve marginals and the same means: 75 and 120 minutes. Their joint pairing differs.

Table 9: Identical marginal data, different continuity consequences. Each event has weight one half.

Joint population	Event 1: (D, B)	Event 2: (D, B)	$\Pr(H = 1)$
Aligned reserve	(30, 60)	(120, 180)	0
Crossed reserve	(30, 180)	(120, 60)	0.5

This is a concrete instance of Proposition 1. A procedure observing only those marginals cannot distinguish a population with no service loss from one in which half the constructed events cause a gap. JudgeAI’s consequence model states exactly which join is required. The empirical request includes fault origin, duration, actual operating load, available reserve at the event, switching performance, and time to functioning replacement supply.

9.3 From a fault mechanism to regulatory design

The original continuity experiment compares existing arrangements, network prevention, faster assisted supply, end-use reserve, and a combined option. In the recorded building-fault example, the outage lasts 120 minutes, current reserve is 15 minutes, and substitute supply becomes operational after 120 minutes. A tested upgrade supplies 240 minutes of autonomy. With the additional payer and device-risk coordinates set to zero in this conditional world, end-use reserve and the combined option are preferred. Network prevention alone leaves the internal fault unresolved.

Technical field evidence also changes the delivery object. A study of solar battery provision for electricity-dependent medical devices reported operational problems involving power sufficiency and components; support was material to functioning use [13]. Thus installed capacity is only one part of the service specification. Compatibility, actual-load testing, switching, maintenance, replacement, and available assistance enter the implementation model.

The resulting regulatory proposal procures *maintained medical-electricity continuity at the point of use*. Its verified output is functioning service through the relevant interruption, with the responsibilities needed to preserve that output. The calculation identifies what must be achieved, while technical inquiry discovers feasible ways of achieving it.

10. Financing is an intervention on the causal pathway

An intervention with a beneficial technical effect can remain unexecuted because the implementer cannot finance its initial cost. Financing must therefore be represented before the comparison of realised worlds. Let K be capital expenditure, L initially available private liquidity, and G public support. Self-financing and reimbursement after completion require

$$L \geq K. \quad (7)$$

An advance available before the expenditure changes the private condition to

$$L \geq K - G, \quad (8)$$

while requiring the corresponding public liquidity. The financing analysis derives these conditions from dated cash flows, retaining the separate parties' requirements.




Figure 3: Exact cash-timing boundary normalised by actual capital cost K . The shaded region identifies projects made executable by a cleared advance, with the same eventual support amount. The figure describes liquidity feasibility; participant consequences remain inputs to normative comparison.

The result identifies a legal difference between nominally equal subsidies. A reimbursement may leave the initial barrier intact; an advance or contractually financed delivery can remove it. The research then follows the public obligation, supplier exposure, performance risk, and effect on other services. A financing arrangement changes the feasible delivery world and the distribution of consequences simultaneously.

A savings-financed retrofit introduces a second boundary. With counterfactual annual energy expenditure B_e , verified saving fraction s , annual contract payment P , and separately charged maintenance, verifica-

tion, financing, and replacement costs M, V, F, R , positive customer cash savings require

$$P < sB_e - M - V - F - R. \quad (9)$$

The expenditure comparison must preserve the contracted service, occupancy, and tariff basis. DEWA’s reported Jafza retrofit provides an engineering context with 32.6% energy savings [14]; the inequality uses the actual project’s verified saving fraction and costs. It consequently supplies a bid-evaluation question and a pricing frontier rather than a universal technology preference.

This analysis changes the proposed legal instruments: staged procurement, performance acceptance, advance availability, savings-backed execution, and allocation of maintenance and shortfall responsibility. These are causal features of an implementable intervention. They belong in the compared world before the normative choice is finalised.

11. From computational findings to a regulatory package

The UAE package translates the inquiry into a proposed enabling framework with 13 articles and model procurement provisions. Its structure follows three computed objects: recovery conditions, essential-service continuity, and executable finance. The following table shows the relationship between a result and the corresponding legal design decision.

Table 10: Trace from computation to regulatory construction.

Computational finding	Information or intervention required	Regulatory construction
Secondary enterprise effect changes the preferred scope	Joint recovery and enterprise outcomes	Targeting and delivery conditions tied to the affected relation
Thermal and access deficits can favour a combined option	Separate assessment of both deficits and joint delivery	A package contract permitting retrofit plus care access
Building faults survive network-only intervention	Device-level reserve, fault origin, compatible supply	Maintained point-of-use continuity service
Support after completion can leave an initial financing gap	Cash timing and each party’s liquidity	Executable advance, milestone, or supplier-financed delivery
A common scale factor leaves ordering unchanged	Research into distinctions that cross decision boundaries	Evidence requirements focused on choice-relevant variables

The framework permits experimental procurement directed at a defined empirical question and subsequent delivery conditional on verified performance. Its proposed obligations state the service to be maintained, the actors responsible, the evidence required, and the financing arrangement. The available technology can evolve while the verified objective and participant-level consequence accounting remain explicit.

This is how new regulatory possibilities arise. A familiar sectoral starting point can be transformed into a different legal object: from room equipment to recovery service, from network reliability to maintained device continuity, or from nominal support to a cash schedule that makes delivery possible. The novelty demonstrated is the derivation and comparison of those designs from the analysed conditions and mechanisms. Each design has a computational and empirical reason for entering the package.

12. What sustained self-improvement means in this architecture

The system can improve along three connected dimensions. *Empirical refinement* specifies a relevant effect, subgroup, causal relationship, or uncertainty more accurately. *Representational refinement* captures a consequential distinction such as fault origin, an actual sleep interval, or the identity of the person bearing a financing burden. *Alternative expansion* discovers a technical, contractual, or institutional intervention and evaluates it under the same normative apparatus.

These operations accumulate because the trace preserves their relationship to the decision. A new source does not merely lengthen an explanation. It can change the parameter range, distinguish populations previously pooled, invalidate an assumed effect, or motivate an additional intervention. The resulting recalculation states whether the choice changes and why. If the choice remains stable, the system can establish the range of conditions supporting that stability.

Research-agenda formation extends the cycle to an earlier stage. The broad-search experiment begins with several threat directions, identifies questions for empirical descent, and returns new contexts to the whole calculation. This differs from simply refining a selected intervention: the subject of further inquiry can itself change. The observed cycle supports that operational capability, while comparative evidence about the quality of research prioritisation remains a further requirement.

The empirical programme is endogenous to the normative problem. A missing variable becomes a research target because it distinguishes worlds with different normative outcomes. A proposed invention becomes relevant because its feasible effect could move an option into a preferred region. A financing reform becomes relevant because the represented technology is otherwise unavailable. This relationship gives the autonomous orchestrator a continuing direction of work.

Sustained inquiry improves the decision through concrete research operations: resolving consequential contradictions, measuring joint dependencies, testing decision boundaries, and evaluating additional interventions. The improvement is embodied in the represented world and the available regulation, with the grounds of each change preserved. The engine remains deterministic for the same complete inputs, while the agent's research changes the information and alternatives submitted to it.

The most consequential architectural property is that normative reasoning can participate in a cumulative scientific process before the eventual social outcomes of a law are observed. The system can already test dependence on evidence, robustness across specified worlds, completeness of an intervention's implementation, and improvements obtainable by expanding the alternative set. Subsequent real-world evidence can enter the same structure and revise the decision again.

13. Why these results support the potential to outperform human decision-making

The experiments support a concrete proposition: JudgeAI can potentially reach a better-grounded normative decision than a human decision process by systematically discovering and evaluating consequential relationships, empirical distinctions, and intervention designs that the latter leaves unresolved. The comparison concerns the quality of the resulting decision under an explicit normative theory. Its basis is how adequately the decision represents the relevant reality and which feasible alternatives it considers.

13.1 The connection between the tests and decision quality

The component experiments show that omitted consequences can change regulation. Removing interactor propagation changes the preferred set in 176 of 700 scenarios; omitting payer consequences changes it in 384. A decision process that leaves these relationships unexamined can therefore reach a different conclusion even when it begins with the same intervention and direct-effect evidence. Traceability makes the consequential omission identifiable, and recomputation makes its effect on the decision observable. These counts measure the effects of analytical omissions in the experiment, rather than human error rates.

The information-sufficiency test supplies the next link. Knowledge of the direct effect alone resolves none of the 700 states, while three linked inputs resolve all of them. More precise knowledge of an isolated benefit can therefore leave the normative question unanswered. JudgeAI can identify the additional relationship that must be investigated and test whether acquiring it settles the choice. The constructed continuity example makes the same point physically: identical marginal data can conceal different risks of losing an

essential service. Recovering the relevant joint relationship improves the empirical basis on which any normative criterion must operate.

Alternative expansion provides direct evidence of improved choice within the comparison. In 128 of 3,456 scenario states, the enlarged menu contains an option strictly preferred to the best option in the original two-option menu. In five states, a combined intervention improves on every available individual intervention. Here the improvement is a better-ranked alternative, found while holding each represented option's consequences and the normative criterion fixed. A human decision process confined to the narrower menu could not select that alternative. JudgeAI's research architecture provides a way to discover it, specify its implementation, and bring it into the comparison.

The aggregate-score experiment establishes a different point: preserving participant-level consequences changes 392 selections. This is evidence that the normative representation performs substantive work. The assessment of improvement follows the stated computational legal theory; disagreement with an aggregate score alone is not a measure of superiority. Empirical completeness and normative evaluation remain connected throughout the argument.

13.2 A cumulative advantage in normative inquiry

The potential advantage over human decision-making lies in making these operations systematic and cumulative. Once a dependency has been represented, subsequent alternatives can be evaluated against it. Once a missing joint observation has been identified, research can target that observation. Once a promising alternative has been constructed, it can be compared with the existing menu across the represented conditions. Earlier findings remain connected to the decision as the inquiry develops.

A human analyst can also discover these relationships and proposals. JudgeAI turns their investigation into a repeatable computational process that autonomous orchestration can sustain across many conditional worlds. The 92,700 pairwise comparisons in the original corpus establish the scale at which this process was exercised; the controlled tests establish why particular distinctions and additional alternatives matter. Computational volume becomes useful when it preserves consequential distinctions, locates research needs, or reveals a preferable implementable intervention.

The broad-search extension adds evidence about the organisation of inquiry, rather than a new human-performance benchmark. Its conditional local preferences and sensitivity findings do not establish the most important threat or the best research allocation. The potential advantage discussed here remains grounded in the demonstrated component effects and alternative-expansion results, with the additional cycle showing how such investigation can be initiated and revised.

The resulting scientific claim is mechanism-based: **JudgeAI can potentially outperform a human normative decision process when its traceable inquiry achieves a more adequate representation of consequential reality and discovers better alternatives under the same normative criterion.** The tests demonstrate the component mechanisms of that advantage. They show how an omitted relationship changes a decision, how a specific empirical distinction resolves it, and how an expanded design space produces a strictly preferred option. Autonomous orchestration connects those mechanisms into sustained improvement of the normative decision itself.

14. Conclusion

JudgeAI's computable and traceable normative decision provides an operational foundation for autonomous regulatory discovery. The experiments show three distinct effects. Inter-actor and payer consequences materially change the choice; targeted empirical information can identify it without resolving every numerical uncertainty; and an expanded intervention set can supply a strictly preferred solution that a narrower menu cannot express.

The separate broad-search experiment extends the architecture to discovering which problem to investigate. Sixteen threat directions generate several empirical branches; local inquiry adds human contexts; and the entire represented set is recalculated. The resulting preferences remain conditional, and the absolute TI diagnostic does not distinguish the final local worlds. This finding identifies the scope of the achieved result: an executed, inspectable cycle of research-agenda formation and revision.

The applied result is a set of legal designs derived from the mechanisms that make them effective and executable: verified recovery services, maintained medical-electricity continuity, and financing aligned with delivery. Ontology locates the consequential relations, the normative engine compares their realised worlds, and orchestration investigates and develops the possibilities exposed by those comparisons.

These findings explain the potential to improve on human normative decision-making: JudgeAI can turn consequential omissions and unexplored alternatives into explicit research tasks, then incorporate their results into a consistent comparison. The advantage grows through the substance of the inquiry: a more adequate account of reality, a richer set of feasible interventions, and a decision recalculated on that basis.

Full traceability makes the decision available for continued research. Empirical inquiry can consequently improve the substance of regulation itself: what is protected, which intervention is chosen, who delivers it, how it is financed, and under which conditions it should be adopted.

A. Experimental verification

The intervention replay, boundary and component experiments used JudgeAI’s normative engine with a fixed computational configuration. Identical complete inputs returned the same preferred sets, and equivalent best alternatives were retained. The 700-case replay reproduced the earlier recorded results. The dense sensitivity study added 1,149 comparisons, each evaluating three worlds. Replay and new sensitivity evaluations together comprised 5,547 world evaluations and 5,547 within-run pairs.

The component comparisons retained the same alternatives and factual inputs while changing the specified evaluation component. The information-sufficiency test examined all subsets of four research variables and identified the smallest combinations that resolved the choice in each tested state. The alternative-expansion analysis retained the same represented consequences while comparing restricted menus with the full menu. These controls connect each reported difference to the analytical operation under examination.

B. Verification of the broad-search extension

The additional corpus retains the 12 final series, their shared references, input versions, full profiles, pairwise comparisons, empirical records, and the reasons for each contextual revision. Its 285 records include repeated controls and are not pooled with the original intervention corpus. The broad and local stages have separate versioned bases; comparisons within each final series use the same actors, routes, and initial states.

Before those comparisons, a scale-consistency check identified a mismatch between direct updates and non-empty baseline cells. An isolated research copy corrected update scaling and its audit; the normative comparator and TI formulae were retained. After correction, all 21 mappings of 20 initial outcomes responded monotonically. Separate checks covered 420 reachable actor–endpoint coordinates, with no clipping across 1,680 non-empty calibrated cells. The maximum scale-identity error was 6.94×10^{-18} . Reverse representations of three already bidirectional geometric links distinguished water, energy, and sanitation states. These numerical and representational checks do not identify empirical causal effects.

The unchanged control reproduces the reference across all compared profiles. Macro and local comparisons after contextual refinement share the same basis. Pairwise audits preserve the macro cycles and local ties reported in Section 5. Consequently, this extension is reproducible as a conditional computational ex-

periment; neither its precision nor its source bibliography establishes the scenario responses as observed intervention effects.

C. Worked research requests

Table 11: Research requests derived from the traced comparisons.

Decision dependency	Data to obtain jointly	How the result enters the next comparison
Worker recovery and enterprise consequence	Actual sleep period and conditions; shifts; absence and output; operating costs and cash obligations over a common horizon	Refine the signed consequence; compare additional protection, different scheduling, targeted delivery, or broader coverage
Recovery package composition	Thermal deficiency; assessed unmet care need; available referral; technology performance; package-specific delivery and adverse effects	Determine whether individual or combined interventions remove the represented deficits
Essential electrical continuity	Fault origin; duration; device load; reserve state and ageing; switching; operational substitute time	Construct event-level loss and gap distributions for each delivery alternative
Financing incidence	Payer and implementer liquidity; dated obligations; contractual support; displaced expenditure; pass-through	Recompute executable worlds and separate consequences for every payer and beneficiary
Savings-backed delivery	Baseline service; occupancy and tariffs; metered use; contract payment; maintenance, finance and replacement costs	Test financing availability while preserving the contracted recovery outcome

References

- [1] Madaan, A., et al. (2023). Self-Refine: Iterative Refinement with Self-Feedback. *NeurIPS*. arXiv:2303.17651.
- [2] Wilder, B., Dilkina, B., and Tambe, M. (2019). Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization. *AAAI*, 33, 1658–1665. doi:10.1609/aaai.v33i01.33011658.
- [3] Filstroff, L., Sundin, I., Mikkola, P., Tiulpin, A., Kylmäoja, J., and Kaski, S. (2024). Targeted Active Learning for Bayesian Decision-Making. *Transactions on Machine Learning Research*. Article and publication record.
- [4] Zhang, J., Cammarata, L., Squires, C., Sapsis, T. P., and Uhler, C. (2023). Active Learning for Optimal Intervention Design in Causal Models. *Nature Machine Intelligence*, 5, 1066–1075. Author manuscript.
- [5] Anthropic (2026). Claude’s values across models and languages. 13 July. Research article.
- [6] Anthropic (2026). Claude’s Constitution. Official document; accessed 20 September 2026.
- [7] Abu Dhabi Public Health Centre (2024). ADOSH-SF Code of Practice 25.0: Driver Fatigue Prevention, version 4.0. Official code, sections 3.2–3.5.
- [8] Bessone, P., et al. (2021). The Economic Consequences of Increasing Sleep Among the Urban Poor. *Quarterly Journal of Economics*. doi:10.1093/qje/qjab013.
- [9] Pasquier, F., et al. (2026). Protective Effect of a High Heat Conductivity Mattress Topper on Sleep During Heat Night. *Journal of Sleep Research*; first published online 24 July 2025. doi:10.1111/jsr.70137.
- [10] Sullivan, J. P., et al. (2017). Randomized, Prospective Study of the Impact of a Sleep Health Program on Firefighter Injury and Disability. *Sleep*. doi:10.1093/sleep/zsw001.
- [11] From good sleep to health and to quality of life: a path analysis of determinants of sleep quality of working adults in Abu Dhabi (2023). *Sleep Science and Practice*. doi:10.1186/s41606-023-00083-3.
- [12] Dubai Electricity and Water Authority. World-Class Results, 2025 performance series. Official performance record; accessed 20 September 2026.
- [13] Evaluation of Solar-Powered Battery Systems for Individuals Using Electricity-Dependent Medical Devices in Puerto Rico Following Hurricane Maria (2021, online). *Disaster Medicine and Public Health Preparedness*. doi:10.1017/dmp.2021.114.
- [14] Dubai Electricity and Water Authority (2021). DP World: enhancing cooperation in energy efficiency improvement projects. 25 May. Official project report.
- [15] IPCC (2023). *Climate Change 2023: Synthesis Report*. Contribution of Working Groups I, II and III to the Sixth Assessment Report. Official report.
- [16] UNDRR (2025). *Global Assessment Report on Disaster Risk Reduction 2025*. Official report.
- [17] United Arab Emirates (2024). *Third Nationally Determined Contribution (NDC 3.0)*, Section 6.2.1, pp. 43–45. UNFCCC submission.
- [18] World Health Organization (2025). *Hypertension country profiles 2025*. United Arab Emirates, printed p. 275 (PDF p. 183); estimates for 2024. Country profiles.
- [19] Federal Competitiveness and Statistics Centre. *Advancing the SDGs: UAE Voluntary National Review 2024*, p. 21, Nabdak programme. Official report.
- [20] Ibrahim, O. M., Al Meslamani, A. Z., Ibrahim, R., Kaloush, R., and Al Mazrouei, N. (2022). The impact of telepharmacy on hypertension management in the United Arab Emirates. *Pharmacy Practice*, 20(4), 1–11. Publisher’s article.
- [21] Abu Dhabi Public Health Centre (2024). ADOSH-SF Code of Practice 11.0: Safety in the Heat, version 4.0, Sections 3.1–3.2. Official code.
- [22] Abu Dhabi Public Health Centre (2024). ADOSH-SF Code of Practice 18.0: Employer Supplied Accommodation, version 4.0. Official code.