applied sciences

MDPI

Article

Determining Social License to Operate: Design of Experiments, Uncertainty Characterization and Historical Consistency Assessment of a Hybrid Simulation Model for a LatAm Green Hydrogen Project

Giovanni Maria Ferraris [1,2,*], Marco Gotelli [2] and Francesco Longo [3]

[1] SIM4Future, Via Trento 43, 16145 Genova, Italy
[2] Department of Mechanical, Energy, Management and Transportation Engineering (DIME), University of Genoa, Via Opera Pia 15, 16145 Genova, Italy; marco.gotelli@unige.it
[3] Department of Mechanical, Energy and Management Engineering (DIMEG), University of Calabria, Via Pietro Bucci, Cubo 45C, 87036 Rende, Italy; f.longo@unical.it
* Correspondence: giovannimaria.ferraris@edu.unige.it

Abstract

Uruguay's largest green-hydrogen initiative—an integrated e-fuels complex proposed for the Department of Paysandú—is approaching its final investment decision amid a public debate centered on water use, land occupation and community acceptance. Building on a peer-reviewed formulation in which social acceptance acts as an operational moderator of the local economy, this article reports the experimental, uncertainty and corroboration program of a hybrid System Dynamics and agent-based territorial impact model: a replicated factorial design, two global sensitivity campaigns, a 1000-replication Monte Carlo campaign, policy isolation studies, a feasibility assessment under editable thresholds, and a 2010–2022 historical backcast. In the simulations, capacity raises income only jointly with external demand (interaction: 21.2% of variance), while employment responds primarily to the public-investment rate (55.1%). Two behavioral parameters—the rates at which community trust builds and erodes—rank highest for every output (PRCC 0.74–0.96); as these are judgment parameters, the quantitative conclusions are conditional on their declared ranges, and their empirical calibration through a longitudinal design is identified as the priority. Acceptance declines monotonically across the sampled capacity range and the outcome distribution is bimodal, so policy is read as shifting the probability of the favorable branch. The backcast reproduces observed income conservatively and the official human-development index to half a percent. A model-conditional feasible configuration emerges—one train, phased construction, labor peace, early cost decline—delivered as a decision tool with editable thresholds.

Keywords: system dynamics; agent-based modeling; modeling and simulation; social acceptance; socio–economic impact assessment; global sensitivity analysis; energy transition; e-fuels; Uruguay

Check for updates

Academic Editor: José A. Orosa García

Received: 15 September 2026
Revised: 10 October 2026
Accepted: 11 October 2026
Published: 22 October 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

Uruguay's national hydrogen roadmap targets, by 2040, an annual production of about one million tons of green hydrogen and derivatives, revenues near USD 1900 million per year, and some 30,000 qualified jobs [1]. The largest single initiative in that program is an integrated e-fuels complex proposed for the Department of Paysandú (121,838 inhabitants):

dedicated renewable generation, an electrolysis block of roughly 1 GW, and synthesis units for e-methanol and e-gasoline, with a projected investment of about USD 4000 million. The project's rationale is decarbonization: green hydrogen and its synthetic fuels are among the few credible vectors for abating emissions in long-haul transport and other hard-to-electrify uses, and Uruguay—whose power system already runs almost entirely on renewables sources—is positioning this second energy transition as an export strategy [1]. That global climate purpose, however, does not exempt the project from local scrutiny; if anything, it sharpens it. The analyzed project is pre-Final Investment Decision (FID)—the environmental license was granted in July 2025 and a Memorandum of Understanding with the Government was signed in December 2025—which defines both the opportunity and the constraint of the research: it is precisely before the final investment decision that simulation can inform the choices still open, and precisely then that several data one would wish for do not yet exist.

The public debate concentrates less on technical feasibility than on territorial sustainability in its environmental and social dimensions alike: the consumption of water and the occupation of land after the severe 2022–2023 drought, physical resources whose cumulative use is explicitly tracked by the model against a reference envelope derived from the project's licensing filing; the memory of the binational pulp-mill conflict on this same river (the Uruguay River), adjudicated at the International Court of Justice [2]; and the institutional question—sharpened by the development-economics literature [3–5]—of whether a capital-intensive export project will operate as an inclusive growth engine or as an extractive enclave. The national capacity-building study warns explicitly that enclave-type configurations generate growing and legitimate rejection [6]. The social-acceptance literature, meanwhile, has established that acceptance is multidimensional, path-dependent and locally constructed [7,8], but its hydrogen-specific strand remains dominated by static survey instruments, with little dynamic modeling of how acceptance evolves and interacts with project economics [9,10]—the gap this research program addresses.

A peer-reviewed antecedent by the authors' group formulated, for this same case, social acceptance as a bounded state variable that builds with perceived information, erodes with perceived risk, and moderates both effective investment and institutional delay [11]. The present article reports the study that formulation called for: the complete experimental, uncertainty and validation program executed on the full territorial impact model built around it, a hybrid of System Dynamics (SD) and agent-based modeling (ABM) components in the tradition of simulation-based decision support for complex and non-ordinary systems [12–16]. The article's main contributions are as follows:

- a design of experiment (DoE) program: 1000-replication Monte Carlo under formal convergence discipline, replicated 2^6 factorial with analysis of variance (ANOVA), two-regime global sensitivity by quasi-Latin-Hypercube with partial rank correlation, and paired policy isolation studies, executed on a verified implementation and reproducible from archived exports;
- the characterization of the model's central emergent structure: a rise-and-collapse of the simulated social license in plant scale—an inverted-U pattern implied by the structure of the acceptance channels rather than by calibration, and exhibited across the sampled capacity range—together with the structurally bimodal outcome distribution it induces and the branch-probability reporting it demands;
- a two-part external corroboration feasible for a pre-FID case: a 2010–2022 historical backcast of the growth and human-development machinery against institutional series, and a documentary traceability audit of the official parameter set;
- a feasibility instrument in which nine capacity-policy configurations are scored against five editable value thresholds, yielding a model-conditional feasible configuration.

Following this introductory section, Section 2 summarizes the model, Section 3 offers the experimental and validation design, and Section 4 presents the results. Section 5 is added for historical corroboration, Section 6 discusses the mechanisms, limitations, and the decision-maker's interpretation, and Section 7 presents the conclusions.

Although executed on a hydrogen case, the framework—an endogenous social-license moderator, a designed experimental program, branch-probability uncertainty reporting and an editable-threshold feasibility tool—is transferable as a whole to any large, contested infrastructure, from mining and pulp to wind farms and data centers, wherever a project's economics depend on a license that the project itself erodes or earns.

This study builds directly on reference [11], which introduced, for this same case, the formulation of social acceptance as a bounded state that moderates effective investment and institutional delay. The antecedent contributed to that formulation; the present manuscript contributes the full territorial model built around it—five interacting subsystems, two agent components, sourced parameterization—and the complete experimental, uncertainty and corroboration program, none of which appears in the antecedent. The gap addressed is specific: hydrogen-acceptance research remains dominated by static survey instruments, with little dynamic modeling of how acceptance evolves and interacts with project economics. Five research questions organize the work: (1) by how much would the plant lift departmental income above its counterfactual trend, and under which market conditions; (2) how many jobs would it generate, and which levers govern that number; (3) how would human development respond; (4) what cumulative environmental pressure would operation exert; and (5) how would social acceptance evolve, and how would its evolution provide feedback on the first four answers.

2. The Model

2.1. Architecture

The model advances annually over a thirty-year horizon at the scale of the Department of Paysandú, initialized on institutional data (census population [17]; national income and development components as declared proxies [18–20]). Five interacting SD subsystems compose the core:

(1) plant production and economics: production constrained by installed capacity, a switchable construction ramp and external demand, with sourced price and cost paths and resulting profit;
(2) fiscal recycling: taxes on plant profit on the traditional agro–economy and on effective private investment, converted into public services through a first-order institutional delay;
(3) education, competencies and human development: services improve adult schooling, schooling maps into workforce competencies that drive plant operating efficiency, and the complete UNDP index machinery is computed at the departmental scale;
(4) environment: cumulative water withdrawal and land occupation measured against a reference envelope constructed from the one-train licensing filing (the 'one-train reference envelope'; Appendix A.5);
(5) social acceptance and perception: the moderator subsystem, with a three-stock compartmental perception panel partitioning the population into negative, indifferent and positive stances.

The division of labor between the two formalisms follows the phenomena: accumulation with feedback—trust, income, education, cumulative resource use—is the natural domain of System Dynamics, while discrete, state-dependent, history-carrying events are naturally expressed as agents with statecharts. Two agent components complete the architecture accordingly. The labor-conflict agent is a statechart with two states (idle; on-strike):

episode onset is stochastic, the duration distribution is declared, and, during the strike, the agent forces plant production to zero in the SD core and injects a shock into the perception panel. The country-risk agent is a two-state regime process (calm; stressed) with exponentially distributed sojourn times, in the tradition of regime-switching models [21]; its state selects the discount conditions entering the effective-investment computation. Coupling is bidirectional: the agents act on production, perception, and discounting and read the simulation clock and configuration in return. Two autoregressive disturbances (demand-price conditions; general economic conditions) carry persistent stochastic weather, with variance-matched bounded innovations whose properties are summarized in Section 4.4. The general architecture scheme is represented in Figure 1.

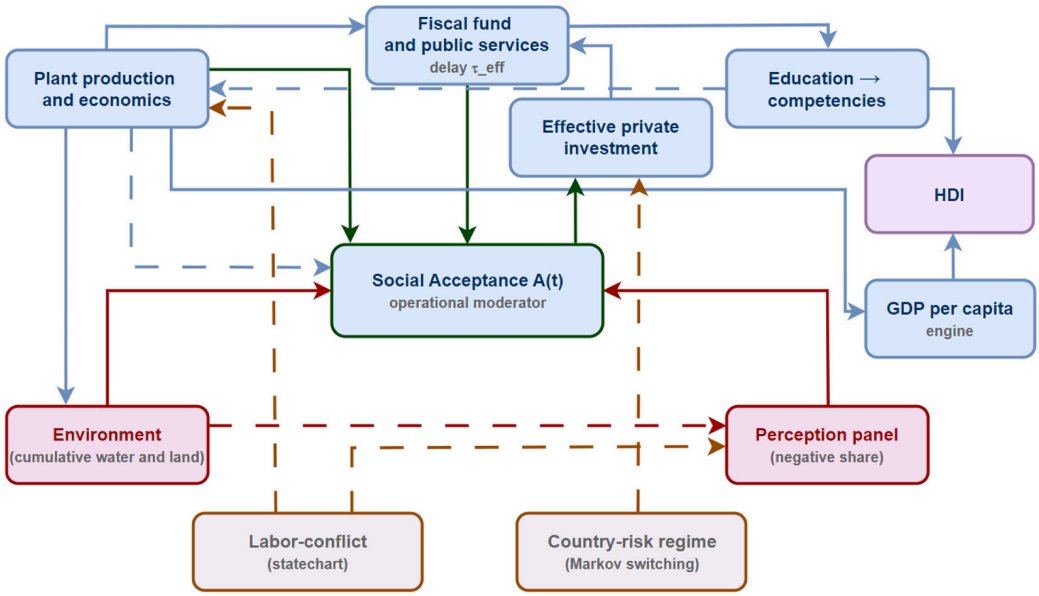


Figure 1. Model architecture with numbered couplings. Blue boxes are SD subsystems (production and profit, fiscal fund and services under the institutional delay, education and competencies, GDP engine, and private investment), with HDI output, risk-side accumulators, and acceptance moderator with its Information and Risk channels. Green arrows are moderation channels of equations (2) and (3), including the investment gate closing the long social-return loop; dashed brown arrows are agent couplings (strike: production halt and perception shock; country risk: discount conditions on effective private investment).

We note that these processes could equivalently be encoded as stochastic event structures within SD: the model is best described as a System Dynamics core with agent-based event components, in the multimethod practice of the simulation literature. The model boundary follows the decision problem: mechanisms feeding the five research questions are included and anchored to institutional data or to the documented public debate; mechanisms that would not change the answers within the horizon or that cannot yet be honestly parameterized (demography and migration, exchange-rate dynamics, construction-phase employment, the investor's own return) are excluded with declared entry paths.

2.2. Acceptance as an Operational Moderator

Following reference [11], acceptance A(t) is a stock bounded in [0, 1]:

$$\frac{dA}{dt} = \gamma \cdot I(t) - \delta \cdot R(t) \quad (1)$$

where the Information channel I(t) aggregates, with saturating terms, the benefits the community can verify—public services delivered, jobs created, an accident-free production

record—and the Risk channel R(t) aggregates the perceived threats: cumulative environmental pressure against the one-train reference envelope, the negative-stance population share from the perception panel, and a diplomatic-tension term activated by production visibility, anchored to the adjudicated binational precedent [2]. The rates γ and δ are the model's central behavioral parameters. Acceptance acts on the economy through two bounded multipliers:

$$I_{eff}(t) = I_{plan}(t) \cdot \min(1, \quad \max(0, A(t))) \tag{2}$$

$$\tau_{eff}(t) = \tau_0 \cdot [1 + k \cdot \max(0, \quad 0.5 - A(t))] \tag{3}$$

where equation (2) gates effective private investment, equation (3) stretches the institutional delay when acceptance falls below one half, a threshold nonlinearity through which low acceptance becomes self-reinforcing. Through equations (2) and (3), the social license is an operating condition of the simulated economy: every economic outcome is contingent on the trajectory of A(t).

2.3. The Growth Engine and the Counterfactual

$$\frac{dP}{dt} = \max\left(0, \quad \left[g \cdot P + k_a \cdot \frac{\Pi(t)}{N}\right] \cdot \left(1 - \frac{P}{P_{\max}}\right)\right) \tag{4}$$

In equation (4) GDP per capita P(t) follows compound growth with logistic damping, into which plant and traditional-sector profits are injected; g is the exogenous departmental trend—a deliberately conservative placeholder pending the official departmental activity series, whose conservatism Section 5 demonstrates—and the separation between the compounding engine and the profit injection is what makes the no-plant counterfactual meaningful: removing the plant removes the injection, not the engine.

2.4. Parameter Epistemology and Verification

Every parameter is catalogued in one of four declared categories: *official* (institutional measurements, audited one by one against their sources as census, income series, UNDP components, the filing's capacities and water intensities, statutory tax rates, the published sector wage awards [17,20,22,23]); *literature* (values or ranges from peer-reviewed or agency sources [24–26]); *calibrated* (tuned to reproduce an observed behavior and documented with the match); and *judgment* (declared assumptions where no measurement exists—including γ and δ—each anchored to an interpretable quantity and submitted to the global sensitivity analysis of Section 4.3, so that the model itself reports which judgments its conclusions depend on).

Construction followed an incremental verification doctrine: every structural addition validated by targeted runs against hand-computed expectations, extreme-condition tests at subsystem edges, switchable extensions proven to reproduce prior results bit-identically in their off position under fixed seeds, and platform primitives treated as part of the model.

The doctrine's yield is a nine-defect register, reported in full in the underlying research documentation; its emblematic entry is a platform random-number primitive found to return its second argument deterministically—identified by matching the frozen state to an analytic fixed point to seven decimal places on two independent artifacts (declarative model source and generated code)—and repaired by uniform, variance-matched innovations whose exact first-two-moment invariance underwrites every means-, variance-, and rank-based method used below.

2.5. Computational Research Support and Its Perimeter

The study's statistics were run on a dedicated computational layer, documented so that the support is verifiable rather than asserted: an independent Python replica-

tion of the documented equations—written from the documents, not from the platform implementation—cross-validates the model as a clean-room second reading; a Monte Carlo engine drove the 1000-replication campaigns under formal convergence discipline [27,28]; a factorial engine built the 2^6 design, proved its orthogonality live, estimated the Yates effects and ran the ANOVA against pure error; a sensitivity engine generated the quasi-Latin-Hypercube design and the partial-rank-correlation machinery; and a backcast engine executed the 2010–2022 historical comparison. The perimeter is declared with equal clarity: machine learning does not currently operate within the model or its analysis, the simulation is mechanistic SD with agent-based components, the statistics are classical, and no result of this article depends on a trained model. The architecture is, however, deliberately prepared for the coupling: the discussion charts where learning methods would extend it, each behind its activation trigger.

3. Design of Experiment and Validation Program

Table 1 states the program as a designed whole: five experimental instruments, each answering a question the others cannot, plus the validation line. All experiments run under a standard protocol (thirty years; construction ramp off unless stated; strike channel on; fixed seeds where single trajectories are reported), and every number reported below is reproducible from the archived experiment exports through the statistical toolkit and ~~a~~ an interactive, live-formula decision-support environment of the research program, in which every input assumption is exposed for re-examination and every downstream result recalculates on change.

Table 1. The experimental and validation program.

Instrument	Research question addressed	Question answered	Size
Reference scenarios	(1)–(5)—baseline anchors for all five	What do the futures look like at the three capacity anchors?	Three trajectories + policy variants
Monte Carlo with convergence discipline [27,28]	(1)–(5)—uncertainty of all five	How dispersed are outcomes, and are the estimators stable?	1000 replications, all stochastics active
Full factorial 2^6 with replication (Yates/ANOVA) [29,30]	(1), (2), (5)—drivers and interactions	Which levers and interactions drive each key performance indicator (KPI), with significance?	64 cells × 5 replicates = 320 runs
Quasi-LHS + PRCC, two campaigns [31,32]	(5)—and its feedback on (1)–(2)	Which of fourteen judgment parameters matter—and does the ranking depend on scale?	500 runs at base + 500 at full capacity
Policy isolation studies	(2), (5)—the levers	What do phasing, conflict and cost timing change, other things being equal?	Paired runs per lever
Historical backcast (corroboration) [33]	(1), (3)—machinery corroboration	Does the growth and development machinery reproduce observed history without the plant?	2010–2022, GDP and HDI

Six factors were selected to span the three families of forces the model distinguishes: the investor's choices, the government's levers, and the environment neither control.

Factor A is installed capacity, contrasted at the two anchors that frame the real decision: the announced first-train scale of 84.63 million kg of hydrogen per year against the full nominal capacity of roughly 300 million, a figure that was not taken from any

announcement but derived by synthesis stoichiometry from the officially filed product mix, so that the design's most consequential factor rests on the filing itself rather than on communication material.

Factors B and C are the government's two fiscal levers, the public-investment rate and the profit-tax rate, while factor D is external demand growth, contrasted between a stagnant and a buoyant world market for e-fuels (2% versus 8% per year); factor E is the base institutional delay with which allocated resources become delivered services, and factor F switches the labor-conflict channel off and on so that the cost of strikes is estimated within the design rather than assumed outside it. Crossing the six factors at two levels each yields sixty-four scenario cells, and each cell was run five times under different stochastic seeds: this replication is not redundancy but the source of the pure-error estimate against which every effect is F-tested, giving the significance statements of Section 4.2 a noise floor measured from the experiment itself.

One further property of the design is deliberately kept verifiable rather than asserted: the orthogonality of the design matrix—the guarantee that every effect is estimated free of contamination from every other—is maintained as a live proof inside the same interactive environment, with each pairwise column product computed in place and visibly summing to zero.

Where the factorial design interrogates the decision levers, the sensitivity campaigns interrogate the model's own assumptions. Fourteen parameters of the 'judgment' category—the quantities for which no measurement yet exists, from the acceptance rates to the channel weights and the induced-employment intensity—were varied simultaneously across their declared plausible ranges rather than one at a time, so that influence is measured in the presence of every other uncertainty, as global sensitivity practice requires [31,32].

The sampling scheme is a lattice quasi-Latin-Hypercube design of five hundred runs, whose space-filling quality was itself verified rather than presumed: the maximum pairwise correlation among the sampled parameter columns is 0.046, indicating low collinearity of the sampled inputs—a sampling property, distinct from the exact orthogonality of the factorial design matrix of Section 3.

Influence is then measured by partial rank correlation coefficients: both parameter and output are rank-transformed, each is regressed on the ranks of all remaining parameters, and the coefficient correlates the two residual series, isolating each parameter's own monotone contribution, cleaned of the others', and robust to the nonlinearities a threshold-gated model necessarily contains.

The entire campaign was executed twice, at first train and at full nominal capacity, so that Section 4.3 can report not only which parameters matter but whether the ranking itself survives the change of scale.

The backcast removes the plant, initializes the machinery on observed 2010 conditions, runs to 2022 with parameters fixed in advance, and measures misfit by the root mean square percentage error (RMSPE):

$$RMSPE = 100 \cdot \sqrt{\frac{1}{n} \cdot \sum_{i=1}^{n} \left(\frac{\left(y_i^{sim} - y_i^{obs} \right)}{y_i^{obs}} \right)^2} \quad (5)$$

reported against the observed compound growth rate, the pattern-oriented standard for models of this class [13,33].

4. Results

4.1. Reference Futures

The program's anchors are three reference futures: the no-plant counterfactual, the plant at the announced first-train scale, and the plant at its full nominal capacity. Table 2 reports their year-30 states under the standard protocol; the values are single trajectories under fixed stochastic seeds—the deterministic skeleton of the analysis—with the full uncertainty treatment deferred to Section 4.4.

Table 2. Year-30 state of the three reference scenarios under the standard protocol (single trajectories, fixed seeds).

KPI (year 30)	No plant	One-train scale	Full capacity
GDP per capita [USD]	35,739	48,543 (+35.8%)	61,415 (+71.8%)
Total employment [jobs]	≈120 (induced baseline)	3262	5153
Human Development Index	0.861	0.917	0.931
Social acceptance A(30)	–	0.396	≈0 (collapse)
Effective private investment [USD, cumulative]	0	171.3 M	negligible (0.42)
Environmental pressure index [× one-train reference envelope]	0	2.03	6.70
Negative-stance population share	5%	43%	94%
Net plant wage bill [UYU/yr]	0	+609.8 M	+1502.6 M

Reading the table correctly requires knowing how its baseline and its extremes were constructed. First, the counterfactual column: the roughly 120 jobs it carries are not residual noise but the honest floor of the comparison—induced employment that arises because the traditional agro–economy keeps paying taxes and those taxes keep funding services even with no plant—and the income trend beneath that column is deliberately conservative, under-tracking thirteen years of observed national history in every single year (Section 5).

Every gain in the table is therefore measured against a baseline built to err in the modest direction: whatever bias the anchors contain understates rather than inflates the plant's merit. The with-plant income figures, in turn, are disciplined by the market: production follows the lesser of ramped capacity and external demand, so capacity beyond what the world absorbs adds nothing until demand grows into it (the structural seam from which the capacity–demand interaction of Section 4.2 emerges as the second driver of income variance).

The environmental pressure index deserves its own reading before the social one: the index is a measure of cumulative resource pressure in the strict sense—water withdrawn and land occupied, expressed against the one-train reference envelope—so its values of 2.03 and 6.70 state that the plant's physical resource use runs at roughly twice and at nearly seven times that reference, independently of anyone's perception of it. The social consequences of that pressure, routed through the Risk channel, come on top of the environmental fact, not instead of it.

The table's most arresting entry, finally, is the USD 0.42 of cumulative effective private investment at full capacity, and it must be read for what it is: not a forecast of a negligible amount, but the arithmetic of a multiplier driven to zero. Gross investment appetite exists at that scale—profits are enormous—but the acceptance gate of equation (2) multiplies it by a license that has collapsed, and the marginal private dollar is entirely deterred. That such a collapse is a realistic mode of failure, and not a modeling curiosity, is a matter of record

rather than conjecture: on this same river, a licensed and technically sound industrial project produced a decade of conflict and an international lawsuit [2], and the mining literature documents licensed projects halted outright by community rejection [8].

4.2. Factorial Evidence: Who Drives What

The chart's visual structure (Figure 2) carries the finding before any number is read: the capacity bar towers at over half the total variation, but immediately beneath it, the interaction bar A_capacity × D_demandGrowth stands essentially as tall as the demand main effect itself, the graphical signature of a system in which capacity does not pay unconditionally, but jointly with the market that absorbs it.

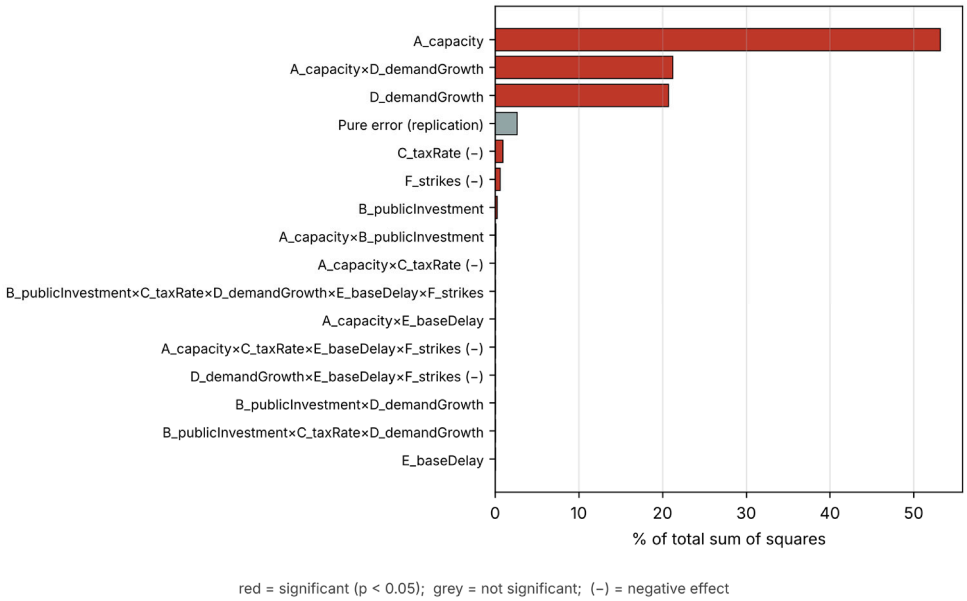


Figure 2. Effects Pareto on GDP per capita (2^6 factorial, 320 runs; red = significant at 0.05; (−) marks negative effects).

Two further features certify the design's quality: the gray pure-error bar (the replication noise floor) sits far below every leading effect, so the red significance marks rest on ample signal-to-noise, and the (−) markers identify the value-destroying terms—the tax rate mildly and the strike channel materially—whose mechanisms are unpacked in Sections 4.2 and 4.5.

Figure 3 is the mirror image that makes the central trade-off visible at a glance: the longest bar is again production capacity but now flagged (−)—the scale reduces effective private investment through the acceptance gate of equation (2)—while the longest positive bar is the public-investment rate, the lever that rebuilds the license through visible services. The two opposing bars, both significant far beyond the 0.05 criterion, are the graphical statement of the tension that this article prices: more capacity, less effective private money, unless the policy compensates.

Read column-wise, Table 3 exposes the asymmetry that organizes the rest of the article: the four hierarchies are not variations on one story but two different regimes of causation.

The purely economic KPI obeys the market—capacity and demand, alone and interacting, account for some 95% of income variation, leaving policy levers in the residual—whereas the socially mediated KPIs obey policy and its tension with scale: employment answers first to the public-investment rate, and the two license-linked outputs answer to capacity negatively and to public investment positively, the same two levers with reversed roles.

The sign reversal of factor A down the table—the largest positive driver of income and the largest negative driver of acceptance and effective investment—is the factorial fingerprint of the inverted-U that Section 4.6 develops: within one designed experiment, the scale simultaneously builds the economy and erodes the license that the economy's private channel requires. A practical corollary follows for the decision-maker: no single configuration maximizes all four columns, so the feasibility question of Section 4.7 is necessarily one of thresholds and trade-offs, not of optima.

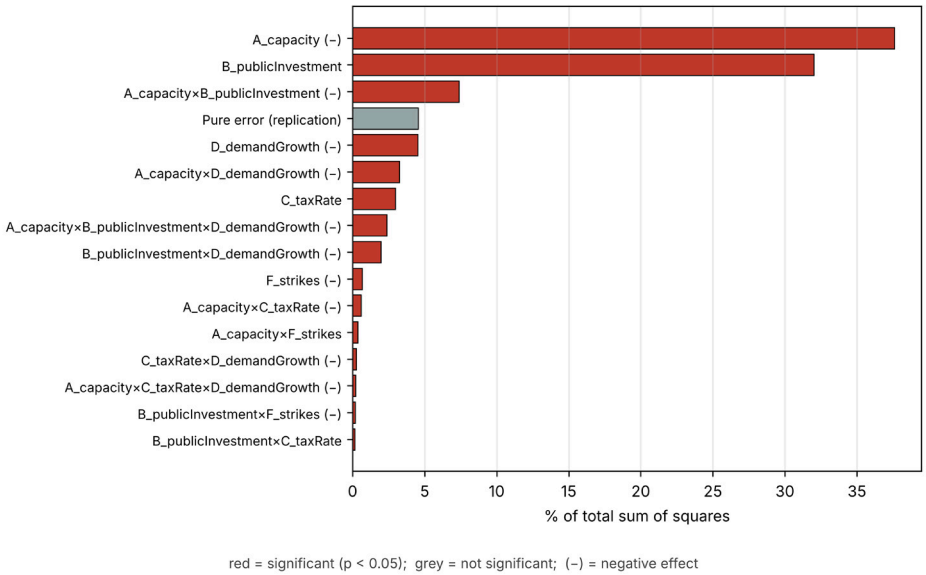


Figure 3. Effects Pareto on effective private investment: the central trade-off—capacity negative, public investment positive—with statistical significance.

Table 3. Factorial variance hierarchies by KPI (share of total sum of squares; all leading effects significant at $p < 0.05$ against pure error).

KPI	Variance hierarchy (Total SS share)	Reading
GDP per capita	A 53.2% > A × D 21.2% > D 20.7%	capacity pays only jointly with demand
Employment	B 55.1% > A 11.8% > C 10.7%	jobs are a policy outcome: the fiscal lever explains five times more variance than capacity
Social acceptance	A 47.8% (−) > B 26.2% (+)	scale erodes the license; visible services rebuild it
Effective private investment	A 37.6% (−) > B ≈ 32% (+)	the central trade-off: more capacity, less effective private money—unless policy compensates

4.3. Global Sensitivity: Two-Parameter Rule

Figure 4 presents the results as a tornado diagram: the standard display of sensitivity analysis, in which one horizontal bar per parameter (here, its partial rank correlation with year-30 acceptance), is stacked in decreasing order of magnitude, so that the ranking of influence is read directly from the vertical order and the funnel-like taper that gives the chart its name. On this acceptance-specific ranking, the leading channel weights necessarily score high alongside the two rates—they multiply the very flows that form the stock—so the dominance claim of this section is the cross-output pattern documented below, not the presence of a gap on any single chart.

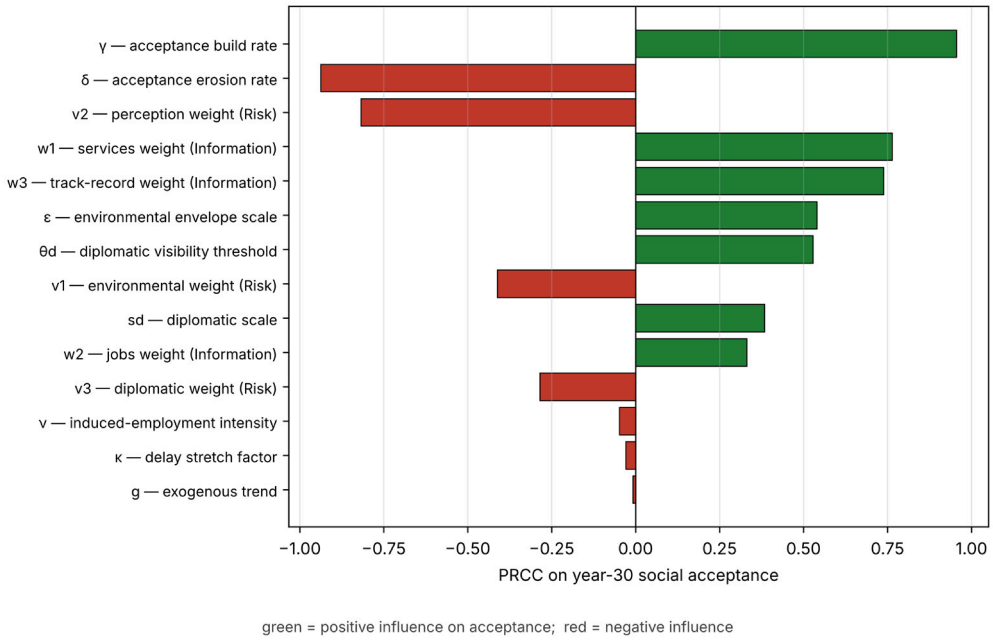


Figure 4. PRCC tornado for social acceptance (base capacity, 500-run quasi-LHS): the acceptance build and erosion rates dominate (green positive, red negative).

The sign pattern beneath is a coherence check the model passes in full view: every information-channel weight (w_1 services, w_2 jobs, w_3 track record) pulls green, every Risk-channel weight (v_1 environment, v_2 perception, v_3 diplomacy) pulls red, and the two scale parameters read mechanistically: a larger reference environmental envelope (ε) relieves relative pressure and supports the license, a higher diplomatic visibility threshold (θ_d) delays that channel's activation with the same effect. The near-zero bars at the bottom carry their own lesson: the induced-employment intensity ν, the delay stretch k, and the exogenous trend g are negligible for this KPI while dominant elsewhere (ν on employment, g on income), a reminder that sensitivity rankings are output-specific, and that the pair at the top is exceptional precisely because its dominance holds across every KPI and both capacity regimes.

The ranking identifies the calibration priority. The pair (γ, δ) carries partial rank correlations of 0.74–0.96 on every KPI at both capacity regimes: the speed at which trust forms and erodes gates the entire simulated economy because acceptance multiplies the private channel and everything downstream inherits it. The induced-employment intensity ranks second (Partial Rank Correlation Coefficient (PRCC) 0.85 on employment), the exogenous trend owns GDP—as its denominator role requires—, the diplomatic-channel parameters test negligible at both regimes for a structural reason the model itself supplies (below the visibility threshold at base capacity, saturated at full capacity, so marginal variation moves nothing in either regime), and everything else is second-order across its plausible range. The two-regime comparison indicates robustness of priorities: the ordering barely moves between one train and full capacity, so the field recommendation that follows—a single community survey calibrating γ and δ, in the social-license survey protocols of the reference [8]—does not depend on which future materializes.

One qualification applies to the full-capacity campaign: 61.8% of its 500 replications end at $A(30) \leq 0.01$, with 174 runs numerically identical at the lower boundary, so the full-capacity PRCC on acceptance itself is read qualitatively, and the two-regime comparison above rests on the KPIs unaffected by this degeneracy: income, employment and investment (per-KPI table in Appendix A.6).

Their empirical calibration requires a longitudinal design rather than a single cross-section: two survey waves separated in time identify the net drift γ-δ given the model's computed channel values between waves; separation of the two rates uses heterogeneity of exposure, with benefit-delivery periods informing γ and risk-salient episodes informing δ; and the public-hearing records of the licensing process provide a retrospective, dated series of stance expressions extending the panel backwards at near-zero cost.

4.4. Uncertainty: Stable Economics, Two Social Futures

With every stochastic channel active, the economic core is remarkably stable (Figure 5): GDP per capita shows a coefficient of variation near 6.9%, employment and fiscal accumulation show similar single-digit dispersions, and the convergence panel indicates that the estimators, not only the estimates, behave [27,28]. Two outputs behave completely differently: social acceptance and effective private investment hold coefficients of variation of 74–79% at any replication count, while their means converge normally. This is the empirical signature of a two-branch mixture.

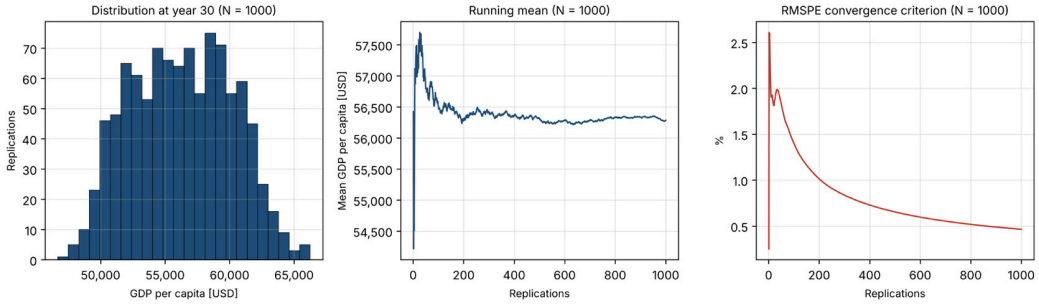


Figure 5. Monte Carlo panel for GDP per capita (1000 replications, all stochastic channels active): outcome distribution, running mean, RMSPE convergence criterion.

For a bimodal outcome Y with branch probability π and branch moments (μ_1, σ_1^2), (μ_2, σ_2^2):

$$E[Y] = \pi \cdot \mu_1 + (1 - \pi) \cdot \mu_2 \quad (6)$$

$$Var[Y] = \pi \cdot \sigma_1^2 + (1 - \pi) \cdot \sigma_2^2 + \pi(1 - \pi) \cdot (\mu_1 - \mu_2)^2 \quad (7)$$

and when the branch separation is large the variance is dominated by the separation term, which no replication budget reduces: the dispersion is an intrinsic property of the system, not an estimation error.

The mechanism is structural: the acceptance integrator has no interior restoring force of its own, the threshold in equation (3) makes low acceptance self-reinforcing, and the multiplicative gate of equation (2) couples the branches to the economy (the alternative-stable-state configuration of reference [34]). Three reporting rules follow and are adopted throughout: report modes, quantiles, and the branch share π, never the mean alone (which sits in the valley between the modes, describing almost no actual future); read policy as the art of moving π; and treat the stubborn social-variable dispersion as the model's statement that this system has 'two futures'.

4.5. Policy Isolation: Phasing, Conflict, Cost Timing

Three paired studies isolate one lever each, other things being equal.

Phasing: activating the construction ramp (production zero during the filed 2026–2029 build, then capacity in 25% steps every two years) raises year-30 acceptance at base scale by 59.2% and employment by 14% at a GDP cost of 5.1%—a quantified, favorable trade of income for license and jobs—and cannot redeem the wrong scale; in the simulations, no

phasing schedule renders the four-train configuration socially feasible under the declared assumptions, with the cumulative impact still reaching 6.6 reference envelopes; pacing is policy, not absolution [35].

Conflict: two strike episodes covering roughly 7% of operating time cost 26.5% of employment, 36.5% of effective private investment, 21% of fiscal accumulation and 14% of year-30 acceptance; destruction far beyond lost output because each halt interrupts the compounding of trust; in the model's accounting, the value destroyed by strike episodes stands at multiples of the wage stake, an argument available to every party.

Cost timing: replacing the stepped cost decline with a smooth learning-curve descent of identical endpoints raises acceptance by 13.9% and effective private investment by 22.9% (GDP +3.3%, employment +11.8%); earlier margins fund visible benefits exactly when trust compounds fastest, so the timing of cost decline is a social-license variable, not only a competitiveness one, and the stepped default is thereby certified as conservative.

4.6. The License in Plant Scale: Rise and Collapse

The central scale result assembles the program's lines of evidence. Within a scenario, time shows a race between delivered benefits and cumulative erosion: at base scale acceptance dips deep (to 0.054 by year 16: worse before better, the signature of benefits that must traverse institutional delays) before the information channels catch up and the license recovers, battered but alive; at full capacity erosion runs at four times the speed and the trajectory never turns.

Across scale, the capacity-resolved Monte Carlo evidence of Figure 6 measures the decline on a continuum: over the sampled range (84.63–170.5 M kg/yr) the median A(30) falls monotonically from 0.32 near one-train scale to below 0.01, and in the highest capacity bin 100% of the replications lie below 0.05 (the collapse is complete already at two-train scale). The four-train scenario endpoint and the structural argument below anchor the range beyond the sample; a dedicated sweep over (170, 300) M kg/yr is left as future work, so the peak's precise location under specific policy settings is not identified by the present evidence.

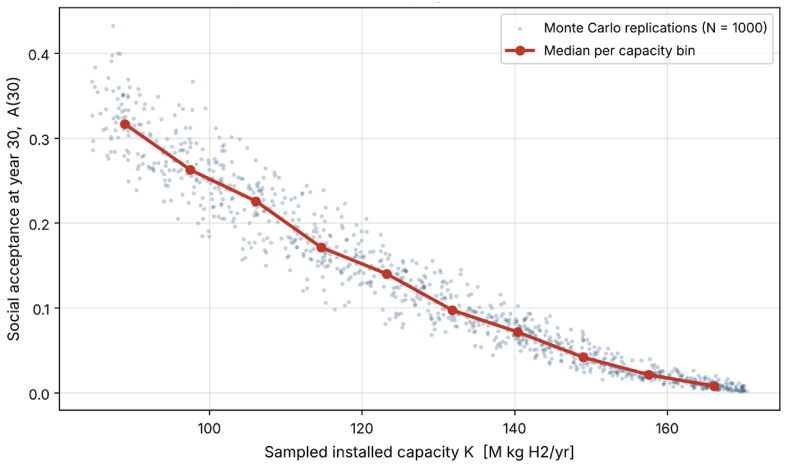


Figure 6. Year-30 acceptance across the sampled capacity range: 1000 Monte Carlo replications with installed capacity ~ U (84.63, 170.5) M kg/yr (points) and median per capacity bin (line).

The figure's geometry (Figure 7) carries three messages of its own. The vertical gap between the two endpoint series at one-train scale—from 0.396 under the standard protocol to 0.735 under the phased, conflict-free configuration (industrial peace)—is the policy dividend made visible as distance: nearly doubling the license's year-30 level using only the levers the project controls. The convergence of both curves to the same zero at full

capacity is the complementary message, as well as the sharper one: policy raises acceptance substantially at one-train scale but does not alter the large-scale decline; no configuration of phasing and peace survives the four-train scale because the collapse is driven by the cumulative-burden structure that policy does not touch. And the left anchor at K = 0 is plotted as reference only: with no plant there is no license to hold, so the inverted-U proper is the rise-and-fall over the operational range.

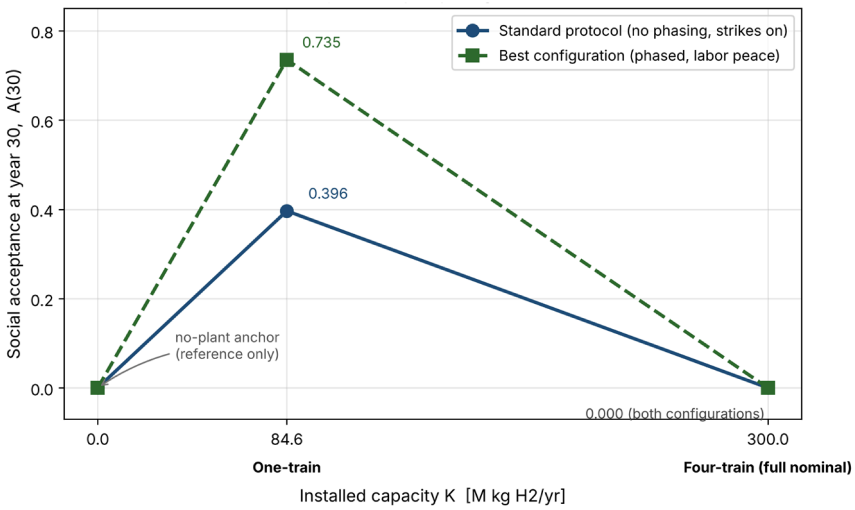


Figure 7. Year-30 acceptance endpoints of the archived scenario configurations at the three capacity anchors, under the standard protocol and the best configuration (phased construction, labor peace). Endpoints are fixed-seed trajectories; the no-plant anchor is a reference point, not an observation of the license.

Importantly, the shape is structurally implied by the channel structure, not calibrated into it: every term of the Information channel is bounded above by its weight, so verified benefits saturate while the environmental term of the Risk channel accumulates in proportion to capacity-scaled operating history against an envelope anchored to one-train scale; hence, for a sufficiently large capacity, erosion dominates permanently, acceptance must decline whatever the parameter values, and calibration decides only where the peak sits, never whether it exists. The result is the quantitative, structural form of the limits-to-growth archetype [36] and of the extractive-enclave warning of the national literature [3,6].

The plotted endpoints of Figure 7 are fixed-seed trajectories; under full stochastic replication the descending branch is precisely the region where the branch probability π of Section 4.4 transitions from one to zero: the inverted-U and the bimodality are the deterministic and probabilistic faces of the same structure.

4.7. Feasibility Under Editable Thresholds

The program closes with a decision instrument rather than a verdict. Nine configurations—the three capacity anchors crossed with the phasing and industrial-peace settings—are scored against five feasibility thresholds: a minimum income uplift over the counterfactual ($\geq$25%), a surviving social license (A(30) $\geq$ 0.5), a cumulative environmental pressure within the accepted ceiling (pressure index $\leq$ 3, in units of the one-train reference envelope), a minimum employment gain ($\geq$1500 jobs) and a minimum effective private investment ($\geq$100 M USD). The defaults are defensible round anchors, not scientific constants: each threshold encodes a value judgment about what 'feasible' must mean for this territory, and for precisely this reason, all five are exposed as editable inputs of the interactive decision-support environment, every verdict recalculating live on any change—so that a decision-maker who weighs the criteria differently generates the corresponding

alternative scenario assessment in seconds. The analysis fixes the trade-offs exactly; the standard remains the reader's.

The environmental criterion requires one clarification of terms. The reference envelope is not a legal cap on cumulative use: the license regulates intake rates and occupied area, which the plant respects at all scales (the intake is of the order of 0.01% of river flow); the index aggregates cumulative water withdrawn and land occupied over the horizon against a reference constructed from the one-train filing. Values above 1 therefore signal cumulative pressure beyond that reference, not illegality, and the threshold of 3 is a value judgment about tolerable cumulative pressure, exposed as editable precisely because it is not a legal or physical line (construction of the index in Appendix A.5).

Table 4 reports the verdicts under the defaults.

Table 4. Feasibility verdicts of the nine capacity-policy configurations under the default thresholds (Section 4.7); acceptance endpoints from the archived scenario exports.

Configuration	Verdict (default thresholds)	Binding failures
Base, no ramp (±strikes)	NOT FEASIBLE	acceptance (0.396 < 0.5)—all economic thresholds pass
Base + ramp, strikes on	FEASIBLE	– (acceptance 0.631 > 0.5)
Base + ramp, no strikes	FEASIBLE (best: A 0.735; ≈5060 jobs; 391 M USD investment)	–
Full capacity, any configuration	NOT FEASIBLE	acceptance ≈ 0; environment 6.6–6.7 ≫ 3—social binds before technical

One structural limitation qualifies the full-capacity rows: the model contains no hard-stop channel (no injunction, license revocation, blockade or forced shutdown, the mechanisms that halted the real projects cited above). Full-capacity economic outputs are therefore upper bounds under a continued-operations assumption, and the observation that economic thresholds do not bind is conditional on it; a conflict-escalation channel with production interruption is a designed extension.

Three structural readings organize the verdicts beyond the row level. First, the economic thresholds never bind: every with-plant configuration clears the income, employment and investment criteria comfortably at base scale, and the full-capacity cases fail no economic test that the social collapse has not already caused: the binding constraints of this project are social and environmental, never technical or economic, which is the feasibility-language restatement of the article's central claim. Second, at base scale the license criterion does all the discriminating work, and the policy levers move exactly that margin: without phasing the configuration misses the threshold by a tenth (A(30) = 0.396), while phasing it clears it with a comfortable margin even under the strike regime (0.631); protecting industrial peace as well lifts it to 0.735: three endpoints that quantify, inside a single instrument, what the isolation studies of Section 4.5 and the scale analysis of Section 4.6 established mechanically. Third, the verdicts are robust in the direction that matters because sensitivity to the thresholds themselves is asymmetric: a decision-maker who relaxed the acceptance criterion to 0.35 would rescue the unphased base case—legitimately, if that is their standard, and the instrument makes such a choice explicit rather than buried in an analyst's assumption—but no defensible movement of any threshold rescues full capacity, which misses the acceptance criterion by essentially its whole range and exceeds the environmental one by more than a factor of two, simultaneously. The feasible region is therefore robust to reasonable variation of the default thresholds: relaxing the acceptance criterion can admit the unphased base case, whereas no defensible adjustment of any

threshold renders the full-capacity configurations feasible, as the sensitivity of the verdicts is asymmetric in the precautionary direction.

Under the defaults and the declared assumptions, the model-conditional feasible configuration is specific: one-train scale, phased construction, protected labor peace, accelerated early cost decline, and services made visible early. The historical justification for taking the pattern seriously was given in Section 4.1: on the same river and in the regional mining record, licensed projects have been halted by community rejection [2,8]. The model does not predict such events; it quantifies the conditions that generate them.

5. Historical Corroboration and Traceability

Regarding income (Figure 8), the structure is validated, with a placeholder proven conservative. At the fitted national rate, the engine tracks the observed series with a RMSPE of 4.7%: compound growth with logistic damping, correct curvature, no drift, and the COVID-19 year visible as a residual exactly where a model without epidemics should show one (observed income drops by 7.4% in 2020 while the simulation passes smoothly). This run is a calibration demonstration, not validation: with one parameter fitted to the same series it reproduces, it establishes that the growth structure is adequate to the observed history (structure adequacy, nothing more). At the departmental placeholder the RMSPE is 7.7% with the decisive property: the error is one-sided, the simulation under-tracking the data in every single year; so the counterfactual errs consistently in the prudent direction, and every uplift of Section 4 is measured against a deliberately modest baseline. This execution is a demonstration of conservatism, as the parameter was fixed in advance, so the one-sided under-tracking establishes the direction of the counterfactual's bias rather than its accuracy.

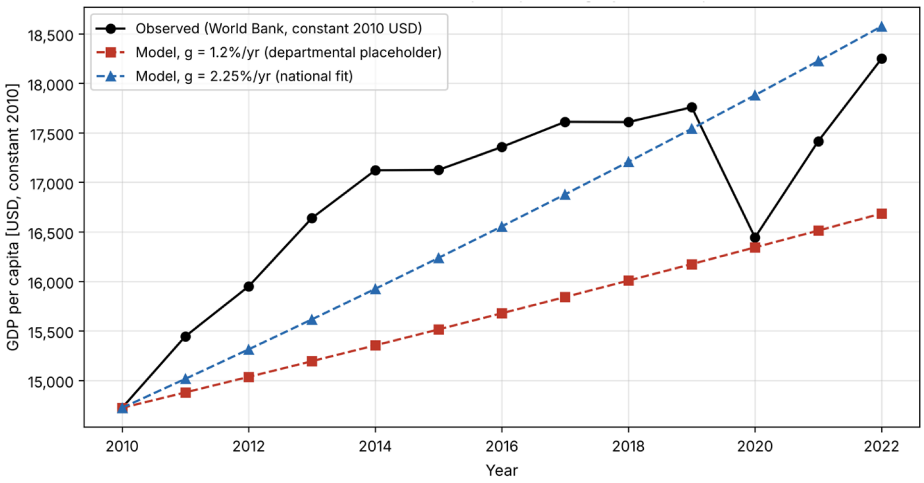


Figure 8. Backcast 2010–2022, GDP per capita: observed national series (World Bank, constant 2010 USD [20]) versus the model's growth engine at the departmental placeholder (1.2%/yr) and at the fitted national rate (2.25%/yr).

On human development (Figure 9), thirteen years are observed at half a percent. Fed the 2010 components and left to run, the UNDP machinery reproduces the observed trajectory with a RMSPE of 0.50% (RMSE 0.0040), the largest residual at the 2021 pandemic dip that the model rightly does not produce. This exercise is a genuine retrodiction of the machinery: the 2010 component values and the UNDP formula set were fixed ex ante, and nothing was tuned to the historical trajectory.

The traceability audit completes the external corroboration feasible pre-FID: the official parameter set is audited item-by-item against institutional sources—census, licensing filings (with plant capacity cross-derived by synthesis stoichiometry from the filed product mix),

statutory rates, the wage awards of the sector councils [22], the real-wage record [37] that brackets the model's emergent 0.70–0.90%/yr—and third-party consistency holds against the capacity-building study, the international agencies' cost and demand bands [6,24–26], and the workforce-readiness assessment [38].

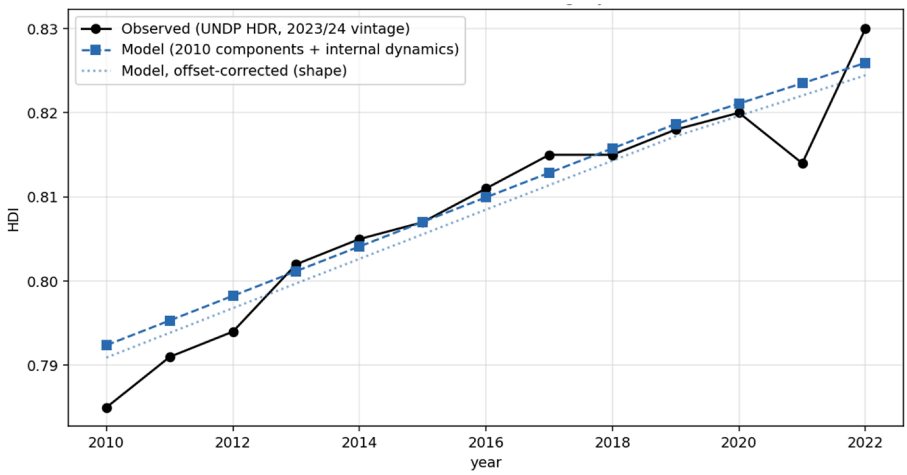


Figure 9. Backcast 2010–2022, HDI: observed UNDP trajectory (2023/2024 vintage [18]) versus the simulated UNDP machinery run from the 2010 components.

What remains genuinely unvalidatable today is with-plant regime behavior: the strongest test the field allows—completed-case validation on the pulp-mill analog across the same river, operating since 2007—awaits ex-post socio–economic studies that remain unpublished; the research declares the gap and repeats the public recommendation that they be commissioned.

One boundary of this exercise must be stated plainly: all three comparisons run at national level, because the departmental series do not yet exist. National corroboration therefore supports the functional machinery—the growth and human-development structures behave as the observed institutions do—while the transfer to the Department (national proxies as departmental initial conditions) remains a declared assumption, mechanical to replace when the official departmental series become available.

6. Discussion

'Structural' results follow from functional form and hold for any admissible parameter values: the decline of the license at sufficiently large scale (saturating benefit terms against a cumulative burden, Section 4.6), the bimodal outcome structure under threshold gating (Section 4.4), and the concave scaling of employment and of demand-capped income (Sections 4.1 and 4.2). Quantitative results—endpoint values, the 0.42 figure, the boundaries of the feasible configuration, and the isolation-study percentages—are conditional on the declared judgment ranges, above all on γ and δ, and are labeled as such. All policy-relevant statements in this article are conditional results of a declared model, offered as structured decision support, not as established properties of the Paysandú project.

The evidence supports the program's central methodological claim: once social acceptance is modeled as an operating condition rather than a survey snapshot, the economics of the project become contingent on the social trajectory, and the contingency is not a nuance but the first-order structure of the outcome space, bimodal in the social variables and inverted-U in scale. Four readings follow.

Institutionally, the full-capacity future is the extractive enclave rendered in numbers—output-large, license void, ninety-four percent negative stance, thirty years of horizon attracting virtually no effective private investment (USD 0.42, cumulatively)—while the

base-scale phased futures show the institution still open; the levers of the feasible configuration are readable as an operational definition of institution-making, converting the plant's presence into broadly shared, verifiable benefit [3–5].

Environmentally, the results separate the physical from the perceived without conflating them. At base scale, cumulative water and land use remain in the neighborhood of the one-train reference envelope (index 2.03); at full capacity, they exceed it by nearly sevenfold, a statement about resource sustainability that stands on its own, prior to any social response. In sustainability terms, the divide is between weak and strong: every with-plant future satisfies the weak criterion (income and human development rise at all scales), but only the futures inside that configuration satisfy the strong one—economic gain, a surviving license and cumulative environmental pressure within the accepted ceiling, jointly—so that full capacity purchases weak sustainability at the price of strong. The decarbonization framing completes the picture rather than dissolving it: the plant's global climate contribution is real, but a project whose purpose is planetary sustainability cannot fund that purpose through local environmental overdraft, and in this model, the social license is precisely the mechanism through which the territory audits that exchange.

Methodologically, the hybrid approach delivered what no static description can: no survey measures a curve, no fixed-coefficient model produces an inverted-U from monotone channels, and the emergent-then-proven character of the central result—its tell-tale signature (parameters that raise acceptance also raising environmental impact) appeared in the sensitivity analysis before the causal loop was drawn—is evidence about the structure, not about the modelers' intentions.

For the decision-maker, the program compresses to the following: the department gains substantially in every feasible future; the model-conditional feasible configuration is narrow and specific; the correct object of policy under this uncertainty is the branch probability π; the cheapest high-value act available today is measurement (a two-wave survey calibrating γ and δ, per the design of Section 4.3; the public-hearing records of the licensing process are a zero-cost first tranche); and the thresholds that define feasibility remain, properly, editable.

Limitations bound the claims and define the agenda: the counterfactual trend is a conservative placeholder pending the official departmental series (its substitution is engineered to be mechanical); population is fixed at census, so the distributional question—who captures the benefit—awaits a demographic module, with the capacity-building literature indicating that several required profiles do not yet exist locally [6,38]; the acceptance rates are declared judgment pending field measurement: no community survey or stakeholder-perception dataset exists yet for the Department on this project, so none was used, and its collection through the longitudinal design of Section 4.3 is the first planned implementation step of the research program; the price path proxies the value realized across the e-fuels chain, audited stoichiometrically for order-of-magnitude plausibility; bounded innovations understate rare extreme single years by a known, conservative margin; construction-phase employment, exchange-rate dynamics and the investor's own return sit outside the perimeter, each with a designed, triggered entry path. These are measurement debts more than modeling problems, and they are prioritized: behavioral calibration first, institutional series second, the completed-case record third. Three designed extensions chart the simulation-artificial intelligence (AI) coupling ahead, each with its trigger: a surrogate metamodel of the simulator—a Gaussian-process or gradient-boosted emulator [39,40]—would make the branch probability π computable across the full policy space in a short time, turning the feasibility tool into an interactive map; Bayesian calibration in the framework of Kennedy and O'Hagan [41] will convert the judgment parameters into measured posteriors when the community survey delivers its data; and, with a fast surrogate available, policy search

methods [42] could seek the configurations that maximize π under constraints. In this way, the surrogate waits for engineering, calibration waits for the survey, and research waits for the surrogate.

7. Conclusions

This manuscript set out to answer a question that static instruments cannot: what happens to a territory over thirty years when the largest industrial project in its history arrives carrying not only capital and capacity but a social license that must be continuously earned? To answer it, the article reported the complete experimental, uncertainty and validation program of a hybrid SD and ABM of Uruguay's largest green-hydrogen project, in which social acceptance operates as an economically binding moderator—a state of the system that gates private investment and stretches institutional delivery, rather than a survey figure recorded beside the model.

One rule governs the reading of everything that follows: the findings summarized here are simulation results, conditional on the declared model structure and on the judgment ranges of the behavioral parameters. The empirical evidence in this article is confined to the historical consistency of the growth and human-development machinery and to the documentary traceability of the official inputs; the with-plant social dynamics have no empirical counterpart yet and await the field measurement that the research agenda awaits.

The program's findings compose a single arc rather than a list. The economics of the plant are real but conditional at every step: productive capacity is remunerated only jointly with the demand that absorbs it, and employment answers to the public-investment lever far more than to plant scale, so that the two headline promises of the project, income and jobs, depend on the world market and on the government, respectively, before they depend on the investor. The condition beneath both, however, is social: the simulated license rises and then collapses with plant scale—a pattern structurally implied by the channel asymmetry and exhibited across the sampled capacity range (Section 4.6)—surviving at first-train scale and collapsing at full capacity, where the acceptance gate closes the private channel almost entirely (Table 2); and because the license's dynamics are threshold-gated, the outcome space is structurally bimodal, which changes what rational policy means: the correct object of intervention is not the average future, which describes no future anyone will live, but the probability of the license-holds branch. These statements rest on the corroboration feasible for a pre-FID case—its growth and human-development machinery reproduce thirteen years of institutional history, conservatively on income and to half a percent on the official index—and through a sensitivity verdict that doubles as a research agenda: two behavioral parameters, the rates at which community trust builds and erodes, dominate every output at both scales, a dominance that Section 4.3 converts into the field recommendation heading the research agenda.

What emerges for the decision is accordingly not a verdict but a model-conditional feasible configuration: modest in scale, gradual in construction, protected from labor conflict and front-loaded in its cost decline (Section 4.7), a configuration that satisfies every feasibility criterion, while no configuration of the same levers rescues full capacity. The instrument delivering that configuration keeps its feasibility criteria open to the decision-maker's own weighing: the analysis fixes the trade-offs exactly, and what counts as feasible remains—as it should—a political choice.

The contribution is therefore two-fold. For the research community, it is a documented, sourced, verified, and reproducible experimental program executed on a live pre-FID case as a demonstration that the social license can be modeled as an operating condition, and that doing so changes the answers—on a template applicable well beyond hydrogen. For

the public decision, it is something rarer: the possibility of seeing the department's 'two futures'—and the levers that select between them—before choosing.

Author Contributions: Conceptualization, G.M.F., M.G. and F.L.; methodology, G.M.F., M.G. and F.L.; software, G.M.F., M.G. and F.L.; writing—original draft preparation, G.M.F., M.G. and F.L.; writing—review and editing, G.M.F., M.G. and F.L. 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 data supporting this article—model implementation, parameter tables and ranges, archived experiment exports (replicated factorial design, global sensitivity campaigns, Monte Carlo campaigns, 2010–2022 backcast), analysis scripts and the interactive results workbook—are available from the corresponding author on reasonable request. Public deposit is deferred because these materials are integral part of research program still in progress; designed model extensions are under active development, and further publications based on the same materials are in preparation. Requests for the purposes of scientific verification or peer review will be accommodated.

Acknowledgments: The authors wish to honor the memory of Prof. Agostino G. Bruzzone, a leading figure in Modeling and Simulation, who passed away in late 2025. His memory, spirit, and devotion to research will endure through his contributions, which remain foundational to this manuscript and to the broader research community.

Conflicts of Interest: Author Giovanni Maria Ferraris was employed by the company SIM4Future. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

M&S	Modeling and Simulation
SD	System Dynamics
ABM	Agent-Based Modeling
AI	Artificial Intelligence
KPI	Key Performance Indicator
ANOVA	Analysis of Variance
DoE	Design of Experiment
GDP	Gross Domestic Product
HDI	Human Development Index
PRCC	Partial Rank Correlation Coefficient
RMSPE	Root Mean Square Percentage Error
RMSE	Root Mean Square Error
UNDP	United Nations Development Program
USD	United States Dollar
UYU	Uruguayan Peso

Appendix A. Model Specification, Parameter Ranges and Supplementary Sensitivity Results

This appendix collects the specification details that the main text references: the functional forms of the acceptance channels, the weight choices and their rationale, the fourteen judgment parameters with their sampled ranges, the initial conditions, the construction of the environmental pressure index, the per-KPI sensitivity results at both capacity regimes, and the agent parameters with their epistemic categories. Its purpose is transparency:

together with the materials described in the Data Availability Statement, it allows the reader to reconstruct every equation and every sensitivity figure of the article.

Appendix A.1. Acceptance Channels

The Information and Risk channels of equation (1) are weighted sums of saturating terms: $I(t) = w_1 \min(1,S/S_{ref}) + w_2 \min(1,E_T/E_{ref}) + w_3 \min(1,Q/Q_{ref})$, where S is delivered public services, E is plant-linked employment, and Q is the accident-free production track record. $R(t) = v_1 \cdot \min(1, M_e(t)/\varepsilon) + v_2 \cdot P_n(t)/N + v_3 \cdot D(t)$, where M_e is the environmental pressure index of Appendix A.5, P_n is the negative-stance population share from the perception panel, N is the department's resident population (census), and D is the diplomatic-tension term, activated when production visibility exceeds the threshold θ_d.

Appendix A.2. Weights

Information-channel weight

w_1–w_3 = 0.34/0.33/0.33: the composite-indicator equal-weighting default absent elicited preferences, summing to one so that the Information ceiling is unity.

Risk-channel weight

v_1–v_3 = 0.5/0.5/0.3: an equal pair on the two always-active threat families, with the episodic, threshold-gated diplomatic term weighted below them.

The Risk weights deliberately sum to 1.3, above the Information ceiling of 1.0, encoding the documented asymmetry of risk perception; without this choice the license collapse of Section 4 would be structurally impossible.

Appendix A.3. Judgment Parameters and Sampled Ranges

Table A1. The fourteen judgment parameters varied in the global sensitivity campaigns: baseline value, sampled range (minimum–maximum of the archived quasi-Latin-Hypercube design, 500 runs per campaign) and anchor of the range choice. Baselines marked ‹§2› are the scenario defaults reported with the model equations.

Parameter	Baseline	Sampled range	Anchor/rationale
γ—acceptance build rate [1/yr]	‹§2›	0.0754–0.3	judgment; longitudinal calibration pending
δ—acceptance erosion rate [1/yr]	‹§2›	0.0752–0.3	judgment; longitudinal calibration pending
w_1—services weight (Information)	0.34	0.17–0.51	equal-weighting default (OECD/JRC)
w_2—jobs weight (Information)	0.33	0.165–0.495	equal-weighting default
w_3—track-record weight (Information)	0.33	0.166–0.495	equal-weighting default
v_1—environmental weight (Risk)	0.50	0.25–0.749	equal pair, always-active threats
v_2—perception weight (Risk)	0.50	0.25–0.75	equal pair, always-active threats
v_3—diplomatic weight (Risk)	0.30	0.15–0.449	episodic, threshold-gated term
k—institutional delay stretch	‹§2›	1–4	judgment; delay literature bracket
ε—environmental envelope scale	1.0	3.02–12	one-train filing = 1 by construction
ν—induced-employment intensity	‹§2›	3.52×10^{-6}–1.4×10^{-5}	judgment; regional multiplier band
g—exogenous income trend [1/yr]	0.012	0.00601–0.018	prudent historical anchor (Section 5)
θ_d—diplomatic visibility threshold	0.80	0.4–1.2	structural (visibility gating)
s_d—diplomatic scale	‹§2›	1–4	structural (tension scaling)

Appendix A.4. Initial Conditions

A(0) = 0.5, the midpoint of the unit interval, the maximal-ignorance default: no measurement of community stance exists yet, so the initial state encodes no directional judgment and sits exactly at the activation threshold of the moderating multipliers, so the license is earned or lost from the first year; population, income and human-development components are initialized on the institutional 2010–2011 baseline values (2011 census; 2010 income and HDI series) reported in Section 2.

Appendix A.5. Environmental Pressure Index

$M_e(t)$ aggregates cumulative water withdrawn plus land occupied over the horizon, normalized by a reference envelope constructed from the one-train licensing filing (the filed annual intensities integrated over the horizon, plus the filed footprint). It is a normative stringency yardstick, not a compliance boundary: the license itself regulates intake rates and occupied area, which the plant respects at all scales.

Appendix A.6. Per-KPI Sensitivity Results at Both Regimes

Table A2. Partial rank correlation coefficients of the fourteen judgment parameters on each KPI, at base capacity (upper value) and at full nominal capacity (lower value). In the full-capacity campaign 61.8% of the 500 replications end at A(30) $\leq$ 0.01 (174 are numerically identical at the lower boundary), so the full-capacity PRCC on acceptance is read qualitatively.

Parameter	GDP pc	Employment	Acceptance	Priv. investment
γ	+0.39 *+0.36*	+0.85 *+0.61*	+0.96 *+0.74*	+0.91 *+0.71*
δ	−0.73 *−0.66*	−0.76 *−0.69*	−0.94 *−0.89*	−0.88 *−0.87*
w_1	+0.12 *+0.10*	+0.50 *+0.38*	+0.76 *+0.35*	+0.65 *+0.35*
w_2	+0.11 *+0.25*	+0.26 *+0.26*	+0.33 *+0.39*	+0.31 *+0.36*
w_3	+0.22 *+0.18*	+0.42 *+0.13*	+0.74 *+0.17*	+0.57 *+0.12*
v_1	−0.13 *−0.19*	−0.26 *−0.26*	−0.41 *−0.44*	−0.29 *−0.43*
v_2	−0.43 *−0.34*	−0.54 *−0.31*	−0.82 *−0.62*	−0.70 *−0.58*
v_3	−0.07 *−0.09*	−0.12 *−0.15*	−0.28 *−0.27*	−0.20 *−0.26*
k	−0.05 *+0.00*	+0.01 *+0.06*	−0.03 *+0.03*	−0.03 *+0.03*
ε	+0.18 *+0.34*	+0.33 *+0.21*	+0.54 *+0.67*	+0.43 *+0.64*
ν	+0.02 *−0.01*	+0.85 *+0.76*	−0.05 *+0.00*	−0.05 *+0.02*
g	+1.00 *+1.00*	−0.04 *−0.03*	−0.01 *−0.05*	+0.02 *−0.01*
θ_d	+0.13 *+0.10*	+0.24 *+0.02*	+0.53 *+0.26*	+0.39 *+0.22*
s_d	+0.12 *+0.15*	+0.20 *+0.09*	+0.38 *+0.26*	+0.29 *+0.23*

Legend: base capacity in roman numerals; full capacity in italics; the four KPIs indicated are those discussed in the main text, while the archived matrices include all seven.

Appendix A.7. Agent Parameters

Strike agent: episode frequency and duration—*Judgment*, pending sectoral labor-conflict statistics; the feasibility analysis brackets the channel by running configurations with it enabled and disabled. Country-risk agent: sojourn rates and premium levels of the two regimes—*Literature/Official*, anchored to public country-risk data.

References

1. MIEM. *Hoja de Ruta del Hidrógeno Verde y Derivados en Uruguay*; Ministerio de Industria, Energía y Minería: Montevideo, Uruguay, 2023.
2. International Court of Justice. *Pulp Mills on the River Uruguay (Argentina v. Uruguay), Judgment of 20 April 2010*; ICJ: The Hague, The Netherlands, 2010.
3. Acemoglu, D.; Johnson, S.; Robinson, J.A. The Colonial Origins of Comparative Development: An Empirical Investigation. *Am. Econ. Rev.* **2001**, *91*, 1369–1401. [CrossRef]
4. van der Ploeg, F. Natural Resources: Curse or Blessing? *J. Econ. Lit.* **2011**, *49*, 366–420. [CrossRef]
5. Hirschman, A.O. *The Strategy of Economic Development*; Yale University Press: New Haven, CT, USA, 1958.
6. GIZ. *Producción del Hidrógeno Verde y sus Derivados en Uruguay—Recomendaciones para el Fomento de Capacidades*; Energy Partnership Uruguay-Alemania, Deutsche Gesellschaft für Internationale Zusammenarbeit: Montevideo, Uruguay, 2025.
7. Wüstenhagen, R.; Wolsink, M.; Bürer, M.J. Social acceptance of renewable energy innovation: An introduction to the concept. *Energy Policy* **2007**, *35*, 2683–2691. [CrossRef]
8. Moffat, K.; Zhang, A. The paths to social licence to operate: An integrative model explaining community acceptance of mining. *Resour. Policy* **2014**, *39*, 61–70. [CrossRef]
9. Scovell, M.D. Explaining hydrogen energy technology acceptance: A critical review. *Int. J. Hydrogen Energy* **2022**, *47*, 10441–10459. [CrossRef]
10. Emodi, N.V.; Lovell, H.; Levitt, C.; Franklin, E. A systematic literature review of societal acceptance and stakeholders' perception of hydrogen technologies. *Int. J. Hydrogen Energy* **2021**, *46*, 30669–30697. [CrossRef]

11. Ferraris, G.M.; Giovannetti, A.; González Chagas, S.; Gotelli, M.; Gutiérrez, S.; Kreimerman, R.; Mauttone, A.; Solina, V.; Tonelli, F. Social Analysis Modeling with System Dynamics Approach in a Uruguayan Case of Green Hydrogen Production. *Energies* **2026**, *19*, 1352. [CrossRef]
12. Forrester, J.W. *Industrial Dynamics*; MIT Press: Cambridge, MA, USA, 1961.
13. Sterman, J.D. *Business Dynamics: Systems Thinking and Modeling for a Complex World*; McGraw-Hill: Boston, MA, USA, 2000.
14. Borshchev, A.; Filippov, A. From System Dynamics and Discrete Event to Practical Agent Based Modeling. In Proceedings of the 22nd International Conference of the System Dynamics Society, Oxford, UK, 25–29 July 2004.
15. Bruzzone, A.G.; Ferraris, G.M.; Giovannetti, A.; Gotelli, M. Modeling & Simulation for an investigation process to support decision making in a non-ordinary context. In Proceedings of the 36th European Modeling and Simulation Symposium (EMSS 2024, I3M), Tenerife, Spain, 18–20 September 2024.
16. Bruzzone, A.G.; Sinelshchikov, K.; Di Matteo, R. Population Behavior, Social Networks, Transportations, Infrastructures & Urban Simulation for Decision Makers. In Proceedings of the 30th European Modeling and Simulation Symposium (EMSS 2018, I3M), Budapest, Hungary, 17–19 September 2018.
17. INE. *Censo de Población y Viviendas 2023*; Instituto Nacional de Estadística: Montevideo, Uruguay, 2023.
18. UNDP. *Human Development Report 2023/24—Breaking the Gridlock. Statistical Annex (HDI Trends)*; United Nations Development Programme: New York, NY, USA, 2024.
19. UNDP. *Human Development Report 2025—Technical Notes*; United Nations Development Programme: New York, NY, USA, 2025.
20. World Bank. National Accounts Data, Constant GDP per Capita for Uruguay [NYGDPPCAPKDURY], Retrieved via FRED Federal Reserve Bank of St. Louis. Available online: https://fred.stlouisfed.org/series/NYGDPPCAPKDURY (accessed on 3 August 2026).
21. Hamilton, J.D. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* **1989**, *57*, 357–384. [CrossRef]
22. BPS—Instituto de Seguridad Social. *Consejos de Salarios, Grupo 7, Expedientes 2024-13-2-0003608 y 2024-13-2-0002510 (Laudos)*; BPS—Instituto de Seguridad Social: Montevideo, Uruguay, January 2025.
23. Ministerio de Ambiente. *Observatorio Ambiental Nacional—Planta e-Combustibles HIF Paysandú*; Ministerio de Ambiente, EIA: Montevideo, Uruguay, 2025.
24. 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, 2020.
25. IEA. *Global Hydrogen Review 2025*; International Energy Agency: Paris, France, 2025.
26. RMI. *Hydrogen Reality Check: Distilling Green Hydrogen's Water Consumption*; Rocky Mountain Institute: Basalt, CO, USA, 2023.
27. Law, A.M. *Simulation Modeling and Analysis*, 5th ed.; McGraw-Hill: New York, NY, USA, 2015.
28. Kleijnen, J.P.C. *Design and Analysis of Simulation Experiments*, 2nd ed.; Springer: Heidelberg, Germany, 2015.
29. Yates, F. *The Design and Analysis of Factorial Experiments*; Imperial Bureau of Soil Science: Harpenden, UK, 1937.
30. Montgomery, D.C. *Design and Analysis of Experiments*, 10th ed.; John Wiley & Sons: Hoboken, NJ, USA, 2019.
31. McKay, M.D.; Beckman, R.J.; Conover, W.J. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics* **1979**, *21*, 239–245. [CrossRef]
32. Marino, S.; Hogue, I.B.; Ray, C.J.; Kirschner, D.E. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *J. Theor. Biol.* **2008**, *254*, 178–196. [CrossRef] [PubMed]
33. Barlas, Y. Formal aspects of model validity and validation in system dynamics. *Syst. Dyn. Rev.* **1996**, *12*, 183–210. [CrossRef]
34. Scheffer, M.; Carpenter, S.; Foley, J.A.; Folke, C.; Walker, B. Catastrophic shifts in ecosystems. *Nature* **2001**, *413*, 591–596. [CrossRef] [PubMed]
35. Flyvbjerg, B. What You Should Know About Megaprojects and Why: An Overview. *Proj. Manag. J.* **2014**, *45*, 6–19. [CrossRef]
36. Meadows, D.H.; Meadows, D.L.; Randers, J.; Behrens, W.W., III. *The Limits to Growth*; Universe Books: New York, NY, USA, 1972.
37. INE. *Anuario Estadístico Nacional 2024, Cuadros 3.4.15–3.4.17 (Índice Medio de Salarios Reales)*; Instituto Nacional de Estadística: Montevideo, Uruguay, 2024.
38. Institute of the Americas. *Latin America Energy Transition Workforce Readiness Assessment*; Institute of the Americas: La Jolla, CA, USA, 2022.
39. Sacks, J.; Welch, W.J.; Mitchell, T.J.; Wynn, H.P. Design and Analysis of Computer Experiments. *Stat. Sci.* **1989**, *4*, 409–423. [CrossRef]
40. Rasmussen, C.E.; Williams, C.K.I. *Gaussian Processes for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006.
41. Kennedy, M.C.; O'Hagan, A. Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B* **2001**, *63*, 425–464. [CrossRef]
42. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*, 2nd ed.; MIT Press: Cambridge, MA, USA, 2018.

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.