

Strategic pathways for residential solar PV uptake in developing urban markets: A heterogeneous household simulation of financial barriers and policy effectiveness

Mohammad Fazle Rabbi 

Coordination and Research Centre for Social Sciences, Faculty of Economics and Business, University of Debrecen, Böszörményi út 138, 4032 Debrecen, Hungary

ARTICLE INFO

Keywords:

Residential solar PV adoption
Energy transition management
Household simulation
Policy effectiveness

ABSTRACT

Residential solar photovoltaic adoption in developing urban markets is examined as an energy transition management problem within a calibrated low- and middle-income country archetype of 2,000 heterogeneous households over a 15-year horizon. A simulation framework integrating sequential financial screening, heterogeneous Bass-type diffusion, and Monte Carlo uncertainty quantification across 200 replications is applied to evaluate five policy configurations: Baseline, Soft Loans, CAPEX Subsidy, Feed-in Tariff, and Grid Parity Wait. The results indicate that the prosumer transition is governed by a hierarchical barrier structure where dominant constraints shift over time. The long payback period accounts for 54 % of blocked households in Year 1, whereas behavioural resistance rises to account for 98 % of residual non-adoption by Year 15, despite a 62.5 % cost decline and Year 1 grid parity. Multi-criteria policy appraisal reveals sharp structural trade-offs. Passive configurations dominate composite rankings solely on fiscal efficiency and implementation feasibility. Among active interventions, the 30 % CAPEX Subsidy maximises both the 15-year adoption rate (98.4 %) and equity (9.76 out of 10), compressing the terminal equity gap to 2.4 percentage points. Despite its high fiscal cost of USD 1.33 million, its adoption and equity performance secure the highest composite rank overall. The Feed-in Tariff emerges as the most balanced active instrument, ranking second overall. External validation yields a PASS classification for 11 of 15 empirical benchmark tests. Permutation analysis confirms that terminal penetration (72.8 %) and equity scores remain structurally invariant to barrier reordering. Effective transition management requires simultaneous intervention targeting early-stage financial exclusion and late-stage behavioural engagement.

Introduction

Residential solar photovoltaic adoption in developing urban markets has emerged as a critical energy transition management challenge because technology cost reductions do not automatically translate into financially accessible distributed generation uptake [1,2]. Cumulative global solar photovoltaic capacity surpassed 1.6 terawatts by 2023, with the residential segment accounting for a growing share of new installations across low- and middle-income countries [3]. In developing urban markets, rising grid tariffs and evolving subsidy regimes create conditions in which household investment in solar photovoltaics appears financially attractive for certain income groups. Observed adoption nevertheless remains uneven and concentrated in specific segments of the income distribution [4,5]. The prosumer phenomenon, in which residential households simultaneously consume and generate electricity,

represents a structural transformation of urban energy systems. This transformation has serious implications for utility revenue stability [6,7], for energy access equity, and for the achievement of national climate commitments. Despite substantial progress in technology cost reduction, with the International Renewable Energy Agency documenting an average learning rate of an 18 % cost reduction per doubling of cumulative installed capacity, diffusion of residential solar photovoltaics in low- and middle-income country (LMIC) urban markets remains uneven and significantly below the technical potential [8].

In this study, developing urban markets are operationalised as grid-connected urban areas within low- and middle-income countries where residential households possess established central electricity network access and rooftop solar photovoltaic installations are technically viable [9]. The empirical calibration incorporates pooled cross-country data from urban settings exhibiting nominal lending rates, tariff trajectories, and income distributions characteristic of developing nations [10,11].

E-mail addresses: drabbikh@unideb.hu, rabbimohammad@econ.unideb.hu.

<https://doi.org/10.1016/j.ecmx.2026.102145>

Received 31 May 2026; Received in revised form 9 July 2026; Accepted 11 July 2026

Available online 12 July 2026

2590-1745/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Nomenclature and Abbreviations

The following abbreviations are used in this manuscript:

Abbreviation Full Term

CAPEX	Capital Expenditure
CI	Confidence Interval
CV	Coefficient of Variation
DTI	Debt-to-Income
ED	Energy Displaced
EQscore	Equity Score
FiT	Feed-in Tariff
IEA	International Energy Agency
IEA PVPS	IEA Photovoltaic Power Systems Programme
IRENA	International Renewable Energy Agency
IRR	Internal Rate of Return
LCOE	Levelized Cost of Electricity
LMIC	Low and middle income Country
MC	Monte Carlo
NPV	Net Present Value
PBP	Payback Period
PMT	Periodic Monthly Payment (Annuity)
PV	Photovoltaic
RQ	Research Question
RL	Revenue Loss
USD	United States Dollar
WEO	World Energy Outlook

Mathematical Symbols

Symbol	Definition (Unit)
Y_i	Annual household income of agent i (USD)
C_i	Binary credit access indicator for agent i ($\{0, 1\}$)
p_i	Agent-level Bass innovation coefficient (Dimensionless)
q_i	Agent-level Bass imitation coefficient (Dimensionless)
τ_i	Income-dependent payback tolerance threshold of agent i (Years)
r_t	Annual market interest rate at time t (Decimal)
r_0	Initial interest rate at simulation Year 1 (Decimal)
r_{14}	Terminal interest rate at simulation Year 15 (Decimal)
r_d	Social discount rate (fixed at 8 %) (Decimal)
τ_t^{grid}	Grid electricity tariff at time t (USD/kWh)
CAPEX_t	System capital cost at time t (USD)
LCOE_t	Levelized cost of electricity at time t (USD/kWh)
t^*	Grid parity attainment year (Years)
δ_c	Maximum self-financing fraction of income for credit-excluded agents (Decimal)
θ_{DTI}	Debt-to-income limit (fixed at 10 %) (Decimal)

$\text{PMT}(i, t)$	Monthly annuity payment for agent i at time t (USD)
L	Loan term (fixed at 60 months) (Months)
$\text{PBP}(i, t)$	Simple payback period for agent i at time t (Years)
E_{annual}	Mean annual household electricity consumption (kWh)
$\text{NPV}(i, t)$	20-year net present value for agent i at time t (USD)
$P_{\text{adopt}}(i, t)$	Bass adoption probability for agent i at time t (Probability)
$F(t)$	Cumulative adoption fraction at time t (Decimal)
$A(t)$	Count of adopted agents at time t (Integer)
N	Total agent population size (fixed at 2,000) (Integer)
γ	CAPEX subsidy rate (fixed at 30 %) (Decimal)
τ_{FiT}	Feed-in tariff export premium rate (USD/kWh)
E_{self}	Self-consumed energy fraction (kWh)
E_{export}	Exported energy fraction (kWh)
$\rho_k(t)$	Adoption rate of income quintile k at time t (Percent)
$\Delta_{\text{equity}}(t)$	Equity gap between Q5 and Q1 at time t (Percentage points)
$\text{ED}(t)$	Annual energy displaced from the grid at time t (kWh)
η_{system}	System performance ratio (fixed at 0.80) (Decimal)
$\text{RL}(t)$	Annual utility revenue loss at time t (USD)
$\hat{\theta}$	Monte Carlo point estimator for output statistic θ (Variable)
s_{θ}	Sample standard deviation across Monte Carlo replications (Variable)
$\epsilon_c^{(m)}$	CAPEX perturbation in replication m (Decimal)
$\epsilon_{\tau}^{(m)}$	Tariff perturbation in replication m (USD/kWh)
$\epsilon_r^{(m)}$	Interest rate perturbation in replication m (Decimal)
$\phi_{s,d}$	Normalised policy score for scenario s on dimension d (0 to 100)
Φ_s	Composite policy effectiveness score for scenario s (0 to 100)
δ_v	Percentage deviation of simulated metric from benchmark midpoint (Percent)
Ω_v	Overlap coefficient between simulated CI and benchmark interval (Decimal)
ϵ	Boundary tolerance band for WARN classification (fixed at 5 %) (Decimal)
Q_1 to Q_5	Income quintile groups from lowest (Q1) to highest (Q5) (Categorical)
T	Terminal simulation year (fixed at 15) (Years)
μ	Lognormal income distribution mean parameter (9.5) (Dimensionless)
σ	Lognormal income distribution standard deviation (0.8) (Dimensionless)

Rather than replicating a single jurisdiction, the simulation framework models a representative urban low- and middle-income country archetype to evaluate systemic policy mechanics.

The structural impediments to residential solar photovoltaic adoption in developing urban contexts extend beyond technology cost trajectories. Prevailing lending rates in many developing economies impose substantial affordability constraints on loan-financed investment [12]. A large proportion of adults in developing economies lack access to formal financial services, a condition which restricts the use of credit for upfront capital acquisition [13]. Urban income distributions often exhibit high Gini coefficients, indicating extensive heterogeneity in financial capacity across households [12]. These conditions create a gap between the financial viability of solar photovoltaics at an aggregate level and effective accessibility at the household level. The present analysis focuses explicitly on this gap in grid-connected urban settings where rooftop photovoltaics are technically feasible. The compiled

global parameter vector incorporates cross-country institutional datasets to establish a generalised structural representation of an urban residential market, intentionally separating the agent-based simulation framework from locality-specific predictive forecasting.

The academic literature on residential solar photovoltaic diffusion spans several analytical traditions. Agent-based studies have demonstrated that heterogeneity in household attributes and in social exposure can influence adoption trajectories [14,15]. Bass-type diffusion models have been used to represent innovation and imitation processes in residential markets in both high-income and developing country contexts [16]. Financial feasibility analyses employing net present value and payback metrics have characterised the role of technology cost and tariff structures in shaping investment outcomes [8,17]. Policy simulation studies have assessed the effects of capital subsidies, concessional finance, and feed-in tariffs under different regulatory and market conditions [8,18,19]. Existing work has typically treated these dimensions

separately and does not uniformly embed distributional analysis and multi-criteria evaluation within a single calibrated framework [8,17]. Financial constraints are often represented as parallel conditions that can be satisfied or not satisfied independently. Equity effects are frequently reported at an aggregate level rather than by income strata [20]. Recent scholarship emphasises that household adoption decisions are shaped by motivational and behavioural variables alongside financial feasibility metrics [21]. Psychometric frameworks demonstrate that adoption intentions can be systematically quantified across distinct behavioural dimensions, a perspective complementing the analytical focus of the current study on financial barriers and policy instruments. This analytical fragmentation motivates the development of an integrated approach that represents financial barriers as a sequence, distinguishes outcomes across income strata, and compares widely used policy instruments under a consistent set of assumptions for an urban LMIC market archetype.

Critical gaps persist across three intersecting dimensions. The first deficiency relates to the sequential and hierarchical structure of financial barriers. In many operational settings, upstream constraints such as credit exclusion preclude evaluation of downstream affordability conditions. Existing diffusion models frequently treat financial constraints as parallel and independent factors rather than as cascading conditions arranged in a logical hierarchy [16,22]. This representation can distort both the magnitude and the composition of the addressable adoption market. The second limitation involves the distributional implications of alternative policy instruments across income quintiles. Empirical evidence indicates that low- and moderate-income households consistently exhibit lower adoption rates than higher-income counterparts [23,24]. However, the income-stratified mechanisms through which specific fiscal instruments expand or constrain access remain incompletely quantified in the developing economy literature. The third gap emerges from the fragmentation of multiple evaluation dimensions. Comparative policy studies rarely combine adoption trajectories, distributional performance, fiscal cost, and implementation feasibility within a single normalised multi-criteria framework [25,26]. This fragmentation limits the operational utility of existing analyses for evidence-based policy prioritisation. Policymakers are not provided with cross-dimensional evidence that would enable the ranking of competing fiscal instruments against several socially relevant objectives simultaneously. These cumulative gaps constrain the capacity of existing scholarship to inform the management of residential solar photovoltaic diffusion in developing urban markets in ways that are technically credible, financially inclusive, and operationally relevant for decentralised energy transitions.

Three research questions motivate the investigation; these items are structured as functional analytical targets. RQ1: What is the relative magnitude and temporal evolution of distinct financial barriers to residential solar photovoltaic adoption across heterogeneous household income strata in developing urban markets? RQ2: How do alternative fiscal policy instruments, specifically concessional lending, capital cost subsidisation, and export premium mechanisms, differentially affect adoption trajectories and distributional equity outcomes across household income quintiles? RQ3: Which policy intervention achieves the most effective synthesis across adoption penetration, social equity, fiscal cost, and implementation feasibility dimensions when evaluated through a normalised multi-criteria framework?

To address these questions, the study pursues four specific objectives. Objective 1 is to construct and calibrate a household simulation framework for a synthetic population of 2,000 heterogeneous units traversing a five-stage sequential financial barrier decision architecture across a 15-year horizon, and to test the robustness of the resulting barrier hierarchy against an alternative evaluation sequence. Objective 2 is to evaluate five policy scenarios, namely a Baseline, a Soft Loans configuration, a CAPEX Subsidy, a Feed-in Tariff, and a Grid Parity Wait scenario, under identical initial conditions so that differences in outcomes can be attributed to specific instrument mechanisms. Objective 3

is to quantify distributional equity outcomes across five income quintiles and to estimate utility revenue displacement effects for each policy scenario. Objective 4 is to benchmark all primary simulation outputs against 15 independently published empirical ranges drawn from five institutional sources, specifically the IEA Photovoltaic Power Systems Programme [27], the International Renewable Energy Agency [8], the World Bank Global Financial Development Database [12], the World Bank Global Findex Database [12,13], and cross-country residential solar PV Bass diffusion calibration studies [28]. The indicators covered include adoption trajectories, financial performance metrics, distributional parameters, diffusion coefficients, market lending rates, macro-economic trajectories, investment viability measures, equity gaps, uncertainty robustness statistics, and grid parity timing.

The study advances the field through contributions of both methodological and substantive character. Methodologically, it introduces a heterogeneous household simulation architecture for residential solar photovoltaic adoption that embeds a formally specified five-stage sequential financial barrier hierarchy, explicitly tested against an alternative evaluation ordering, heterogeneous Bass diffusion coefficients drawn from agent-level Beta distributions, and a 15-indicator external benchmarking protocol within a single analytical pipeline. This particular combination of features appears to be uncommon in prior work, which typically implements only a subset of these components within a single framework [15,29]. Substantively, the analysis provides a quantified decomposition of financial barrier recomposition dynamics over a 15-year policy horizon in an urban prosumer context and shows how the relative contribution of different barriers evolves over time under alternative policy configurations, while also establishing that the terminal size of the financially excluded population is invariant to the sequential ordering of the barrier evaluation itself [8,22]. The framework also generates a multi-criteria comparative ranking of residential solar photovoltaic policy instruments that is calibrated to the financial, behavioural, and institutional conditions of developing urban markets. The resulting evidence is directly structured to inform distributed generation planning and renewable energy management in developing urban contexts by accounting for income heterogeneity, credit access constraints, and utility revenue implications within a single analytical framework [19,26].

Materials and methods

Study design and energy system modelling framework

The present research framework models residential solar photovoltaic (PV) adoption in developing urban markets. These areas are defined as mid-sized to large cities in low-income and middle-income countries (LMICs). Such markets are characterised jointly by grid-connected residential electricity access with retail tariffs in the range of USD 0.10 to USD 0.30 per kilowatt-hour (USD 0.10 to USD 0.30 kWh⁻¹), formal credit market penetration below 60 % of the adult population, and household income distributions consistent with a Gini coefficient between 0.35 and 0.55. This definition follows classification conventions used in the World Bank Global Findex Database and the Global Financial Development Database. This classification is used to identify the pooled calibration sample described in Section “Agent Population Initialisation”. The sample is drawn from published urban household survey summary statistics for India, Brazil, and South Africa. These three countries were selected because they jointly span the specified income-Gini and credit-penetration ranges. Additionally, these nations possess publicly available, nationally representative household survey data compatible with the income and credit-access specifications of the model.

The analytical architecture integrates three interlocking computational components: heterogeneous household simulation, Bass diffusion theory, and Monte Carlo uncertainty quantification. The framework is implemented as a sequential simulation pipeline. In this pipeline,

empirically calibrated household populations interact with a time-varying macro-economic environment over a 15-year horizon. Each household unit traverses a nested financial barrier decision structure before reaching a stochastic adoption outcome. This outcome is governed by an embedded social learning mechanism. Five policy scenarios are evaluated under identical initial conditions. This approach permits the attribution of differences in simulated adoption trajectories to policy mechanism variation rather than to differences in starting assumptions.

A heterogeneous household simulation framework is used because residential solar PV adoption depends on households that differ in income, credit access, and responsiveness to social influence. These specific characteristics are not represented adequately by a single representative agent. Households do not interact through explicit spatial neighbourhoods or an articulated social network. Instead, interaction

occurs indirectly through the aggregate adoption fraction. This fraction serves as the social exposure signal in the Bass diffusion component. The model therefore operates as a heterogeneous household financial screening framework implemented within an algorithmic simulation environment [30]. Each household is represented as an autonomous unit endowed with a vector of financial and behavioural parameters. These parameters include income, credit access status, individual Bass coefficients, and a payback tolerance threshold. All parameters are implemented at the household level to capture cross-household variation in financial capacity and susceptibility to diffusion dynamics while maintaining a transparent, tractable structure [31].

The complete conceptual architecture of the simulation framework is shown in Fig. 1. As depicted, the pipeline is organised into eight sequential, colour-coded methodological stages: calibration inputs




Fig. 1. Conceptual framework of the household-based simulation pipeline for residential solar PV diffusion and energy transition management analysis.

(shaded in blue), model initialisation (shaded in purple), the simulation engine (shaded in green), policy scenario configuration (shaded in rose), Monte Carlo uncertainty quantification (shaded in yellow), multi-dimensional output generation (shaded in light blue), external validation (shaded in mint green), and evidence-based recommendations (shaded in orange). Reading top to bottom, the framework traverses empirical data calibration from internationally recognised sources. It then proceeds to the initialisation of the synthetic household population and the macro-economic environment. Following this, the simulation engine executes by incorporating a nested five-stage financial barrier decision tree and the Bass diffusion adoption mechanism. The framework then introduces policy scenario differentiation across five counterfactual configurations comprising Baseline, Soft Loans, CAPEX Subsidy, Feed-in Tariff, and Grid Parity Wait. Next, it conducts Monte Carlo uncertainty quantification across 200 independent replications. This step leads to multi-output simulation result generation. Finally, the pipeline performs external real-world validation against 15 independent benchmark tests. The unidirectional flow from calibration inputs to simulation outputs to policy recommendations reflects a strictly sequential dependency structure. In this structure, no downstream stage is executed before all upstream conditions are resolved.

Within the five-stage barrier decision tree, the five gates are not equally binding over the simulation horizon. Credit access and loan affordability resolve rapidly as declining system CAPEX and rising household income loosen these constraints. Specifically, credit exclusion falls from 6.7 % of the population in Year 1 to 0.97 % by Year 15. Concurrently, the unaffordable-loan constraint falls from 39.5 % to effectively zero over the same period. Payback tolerance remains the dominant constraint through the mid-simulation period. This parameter peaks at approximately 70.7 % of the population around Year 9 before declining sharply as technology costs fall. From roughly Year 10 onward, the terminal constraint shifts to the Bass social-diffusion stage. This stage represents behavioural non-adoption, which isolates households that satisfy all financial screening criteria but have not yet been reached by social exposure. This subpopulation rises to 57.51 % of the population by Year 15. This barrier accounts for 98.35 % of all blocked households at the terminal simulation year. This finding is documented quantitatively in Section 2.4 and Fig. 2, Panel C. This reordering of binding constraints over time is a structural output of the model rather than an assumed input. It directly informs the policy interpretation in Sections 4 and 5. Credit-focused interventions address a constraint that is largely self-resolving under the intrinsic dynamics of the model. Conversely, the residual non-adoption gap by Year 15 is behavioural rather than financial in nature.

All findings reported in this study are conditional on the specific decision architecture, parameterisation, and market archetype adopted here. These outcomes should be interpreted as structural properties of the model under its stated assumptions, rather than as unconditional empirical claims about adoption behaviour in any specific city or country. This qualification is restated where individual results are discussed in Sections 3 to 5.

The conceptual framework in Fig. 1 integrates eight sequential methodological stages: empirical data calibration from World Bank, IEA, and IRENA sources; household population and macro-economic environment initialisation with heterogeneous financial and behavioural parameters across a 15-year horizon; simulation engine execution incorporating a five-stage nested financial barrier decision tree combining credit access, loan affordability, payback period, net present value, and Bass diffusion adoption evaluations. The relative contribution of each barrier stage over the simulation horizon is discussed in the accompanying text and reported in Supplementary Table S2. The subsequent stages involve policy scenario differentiation across five counterfactual configurations comprising Baseline, Soft Loans, CAPEX Subsidy, Feed-in Tariff, and Grid Parity Wait; Monte Carlo uncertainty quantification across 200 independent replications; multi-dimensional output generation including adoption dynamics, equity gap analysis,

financial performance, and policy impact; external real-world validation against 15 independent benchmark tests; and evidence-based policy recommendations. Arrows denote sequential data and decision flow. Coloured boxes distinguish calibration (blue), model initialisation (purple), the simulation engine (green), policy scenarios (rose), Monte Carlo uncertainty (yellow), outcomes (light blue), validation (mint green), and recommendation (orange).

The behavioural and financial realism of this pipeline depends on the distributional properties and empirical grounding of the parameters

Table 1

Household-level parameter distributions, symbols, specifications, and calibration sources for the residential solar PV diffusion model (N = 2,000).

Parameter	Symbol	Distribution / Value	Range	Source
Household income	Y_i	Lognormal ($\mu = 9.5, \sigma = 0.8$)	\$5,000 to \$100,000	World Bank Group (2023) [12]
Credit access	C_i	Bernoulli ($p = 0.52$)	{0, 1}	Demirgüç-Kunt et al. (2022) [13]
Innovation coefficient	p_i	$\beta(2, 9) \times 0.18$	0.02 to 0.18	Rode and Weber (2016) [28]
Imitation coefficient	q_i	$\beta(3, 5) \times 0.42$	0.10 to 0.42	Rode and Weber (2016) [28]
Payback threshold	τ_i	$3.25 + (Y_i/100,000) \times 7$	3.25 to 10.25 yr	Calibrated (this study)
Debt-to-income limit	$DTI_{\max}$	10 % fixed	10 %	Branker et al. (2011) [32]
Discount rate	r_d	8 % fixed	8 %	IRENA (2023) [8]
System lifetime	L	20 years fixed	20 yr	IEA (2024) [27]
Annual energy consumption	E_{annual}	4,500 kWh/yr (mean)	2,500 to 8,000 kWh	IEA (2023) [3]

Notes:

- Six of nine parameters are drawn from internationally published institutional datasets. This selection ensures empirical grounding across the financial, behavioural, and technical dimensions of the simulation architecture.
- Household income follows a lognormal distribution spanning USD 5,000 to USD 100,000. This distribution produces a Gini coefficient of 0.40, which is consistent with World Bank urban LMIC benchmarks. Credit access is generated as an income-linked variable. In this formulation, the underlying credit propensity of a household increases with income subject to random noise. The resulting binary indicator is thresholded so that approximately 48 % of the population is classified as credit-excluded at model initialisation. This specification produces a strong positive association between income and credit access. Simulated credit access rises from 34.5 % in the lowest income quintile to 100.0 % in the highest income quintile. This pattern yields an overall income-credit correlation of 0.37. This relationship reflects the empirically documented concentration of formal financial exclusion among lower-income households in developing economies. This 48 % figure describes the initial, cross-sectional credit-exclusion rate exclusively. The output-level barrier attribution reported in Section 2.10 and Supplementary Table S2 shows credit exclusion falling to 0.97 % of the population by Year 15. Conversely, Bass-governed behavioural non-adoption emerges as the dominant terminal constraint.
- The payback threshold is the sole parameter calibrated endogenously within this study. This parameter is constructed as a linear function of household income. This function produces a range of 3.25 to 10.25 years across the income distribution. This distribution reflects the negative association between financial vulnerability and investment horizon tolerance reported by Schleich et al. [33].
- Both Bass diffusion coefficients are assigned household-level heterogeneous values drawn from bounded β distributions. These values replace fixed population-level constants. This heterogeneity operates as the structural mechanism underlying the finding above. The Bass stage becomes the binding constraint by Year 15 precisely because households retain unresolved behavioural exposure requirements even after all financial screening criteria are satisfied.
- The debt-to-income limit, discount rate, and system lifetime are held fixed across all agents and scenarios. This restriction isolates the effects of variable parameters on adoption outcomes. This approach maintains consistency with Branker et al. (2011), IRENA (2023), and IEA (2024) appraisal conventions.

assigned at initialisation. These parameters are documented in Table 1. The simulation horizon spans 15 annual time steps, capturing both the near-term period in which financial access constraints are dominant and the medium-term period in which evolving technology cost trajectories and social influence jointly shape adoption patterns.

Table 1 consolidates the complete parametric architecture of the agent-based simulation. This matrix enumerates nine household-level and system-level variables that collectively define the financial, behavioural, and technical dimensions of the model for a synthetic population of $N = 2,000$ households. Household income serves as the foundational stratification variable. This variable is drawn from a lognormal distribution with mean parameter 9.5 and standard deviation 0.8. This distribution produces a right-skewed income profile with an effective range of USD5,000 to USD100,000. This range is consistent with World Bank urban LMIC survey data. Credit access is modelled as a Bernoulli draw with success probability 0.52. This initial setting implies that 48 % of the synthetic population faces financial exclusion at model initialisation. This proportion is calibrated to the World Bank Global Findex 2021 cross-country estimate. This figure represents an initial-condition parameter rather than a terminal-year outcome [13]. The corresponding output-level barrier decomposition is reported in Section 2.10 and Supplementary Table S2. The Bass innovation and imitation coefficients are specified as heterogeneous β -distributed variables rather than fixed scalars. These coefficients incorporate population maximum scaling multipliers of 0.18 and 0.42 respectively, drawn from Rode and Weber [28]. The individual ranges span 0.02 to 0.18 and 0.10 to 0.42 respectively. The payback threshold parameter is uniquely calibrated within this study. This parameter encodes an income-proportional investment horizon tolerance ranging from 3.25 years for the lowest-income household to 10.25 years for the highest-income stratum. This formulation captures the documented inverse relationship between financial precarity and intertemporal investment patience. Three parameters are fixed constants aligned with Branker et al. [32], IRENA (2023) [8] and IEA (2024) [27] appraisal standards respectively. These parameters comprise the debt-to-income limit at 10 %, the social discount rate at 8 %, and the system lifetime at 20 years. These constants provide the stable financial architecture within which heterogeneous household behaviour operates. A mean annual energy consumption of 4,500 kWh is drawn from IEA WEO urban LMIC benchmarks [3]. This value anchors all payback and net present value computations to a contextually representative household energy demand profile.

Agent population initialisation

The synthetic agent population comprises $N = 2,000$ households, balancing computational tractability with the statistical representativeness required for quintile-level distributional analysis. Each agent $i \in \{1, \dots, N\}$ is characterised at initialisation by a vector of financial and behavioural attributes drawn from probability distributions calibrated to published aggregate summary statistics such as mean, median, Gini coefficient, and formal credit penetration rate, for the urban LMIC archetype defined in Section 2.1, pooled across World Bank urban household survey publications for India, Brazil, and South Africa and the 140-country Global Findex cross-section [13]. The model does not use raw household-level microdata from these countries; published aggregate distributional moments are used to parameterise the lognormal income and credit-propensity functions described below.

Household annual income is drawn from a lognormal distribution:

$$Y_i \sim \text{LogNormal}(\mu_Y, \sigma_Y) \quad (1)$$

where μ_Y represents the location parameter and σ_Y denotes the scale parameter of the log-income distribution. The specified values $\mu_Y = 9.5$ and $\sigma_Y = 0.8$ restrict the income range to approximately USD 5,000 to USD 100,000, yielding a simulated Gini coefficient of 0.40 that aligns with the archetype criteria established in Section 2.1.

Credit access is formalised as an income-linked binary choice rather than an orthogonal random draw, thereby establishing an empirical dependency between financial inclusion and household wealth. A latent credit propensity score is constructed as an increasing function of household income augmented by idiosyncratic noise. Agents are designated as credit-excluded ($C_i = 0$) if this propensity score falls below a critical threshold corresponding to the target population-wide exclusion rate of 48 % [13]. The operational mechanism is governed by the following system:

$$\pi_i = 0.75 \times \frac{Y_i}{100,000} + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma_\pi^2) \quad (2a)$$

$$C_i = 1 [\pi_i \geq \pi^{(48)}] \quad (2b)$$

where π_i is the latent credit propensity of household i , ϵ_i represents the idiosyncratic stochastic error term, and $\pi^{(48)}$ denotes the 48th percentile of the simulated propensity distribution. This mathematical structure replicates the aggregate share of financially excluded adults in developing economies reported by Demirgüç-Kunt et al. [13]. Agent i is designated credit-constrained when $C_i = 0$. This outcome occurs with a population-level probability of 0.48, ensuring a strong positive association between income and formal credit access. Under this calibration, simulated credit access rises from 34.5 % in the lowest income quintile to 100.0 % in the highest income quintile, generating an overall income-credit correlation of 0.37. This relationship reflects the concentration of formal financial exclusion among lower-income households in developing economies.

Bass diffusion coefficients are assigned as heterogeneous, agent-level stochastic parameters. The innovation coefficient p_i is formalised as follows:

$$p_i \sim \text{Beta}(\alpha_p, \beta_p) \times 0.18 \quad (3)$$

The imitation coefficient q_i is generated via a corresponding scaled beta distribution:

$$q_i \sim \text{Beta}(\alpha_q, \beta_q) \times 0.42 \quad (4)$$

Beta distributions provide bounded support and flexible shapes. Hyperparameters for Equation (3) generate a right-skewed innovation-coefficient distribution; those for Equation (4) generate a moderately symmetric imitation-coefficient distribution, consistent with cross-sectional variance in residential solar PV diffusion documented by Rode and Weber [28]. The shape parameters $\alpha_p, \beta_p, \alpha_q$, and β_q are calibrated to ensure that the maximum possible household scaling values are bounded by 0.18 for p_i and 0.42 for q_i . Population-averaged means across agent initialisation draws are anchored near 0.178 for p_i and 0.419 for q_i , restricting individual agent parameters strictly to the ranges of 0.02 to 0.18 and 0.10 to 0.42 respectively.

The income-dependent payback tolerance threshold captures the relationship between household financial resilience and the willingness to accept prolonged capital recovery periods:

$$\tau_i = 3.25 + \frac{Y_i}{100,000} \times 7, \tau_i \in [3.25, 10.25] \quad (5)$$

where τ_i represents the maximum acceptable payback period in years for household i . The lower boundary of 3.25 years reflects the acute liquidity constraints of the lowest-income cohorts. The upper boundary of 10.25 years represents the extended investment horizons accommodated by higher-income strata. This linear specification is calibrated endogenously and is informed by the documented negative association between financial vulnerability and investment horizon tolerance reported in household energy-investment studies. The specification excludes a separate random term for risk aversion or subjective time preference; the payback barrier isolates households by financial capacity rather than by independent behavioural traits.

This design introduces a distinct structural property regarding quintile-level adoption ordering. Because higher-income households are assigned a longer acceptable payback period, and because absolute annual electricity savings do not vary with income under a fixed consumption assumption, higher-income households face no stronger mathematical incentive than lower-income households to adopt sooner once financial barriers are cleared. This configuration produces faster proportional saturation among lower-income households in the simulation relative to higher-income households. This pattern diverges from empirical LMIC solar markets, where higher-income households frequently adopt earlier due to factors outside the current model scope, including housing tenure, roof ownership, and access to informal capital. This divergence between simulated and empirically observed adoption ordering is recognised as a limitation of the current parametric specification in Section 5.

Income quintiles Q_1 to Q_5 are defined by fixed income thresholds: below USD 10,000; USD 10,000 to USD 20,000; USD 20,000 to USD 35,000; USD 35,000 to USD 60,000; and above USD 60,000. Because the underlying income distribution is right-skewed, quintile population sizes are unequal, with lower-income bands containing substantially more households than higher-income bands. Quintile-specific adoption rates are computed relative to the actual population size of each quintile, maintaining a variable denominator consistently across Sections 2.7 and 3. Adoption is implemented as a strict absorbing state at the individual agent level. Once agent i adopts, the unit is permanently removed from the pool evaluated in subsequent periods and cannot be counted again in any later year or under any subsequent scenario re-evaluation.

Table 1 consolidates the complete parametric architecture of the agent-based simulation. The matrix enumerates nine agent-level and system-level variables that define the financial, behavioural, and technical structure of the model for a baseline population of $N = 2,000$ households. The parameter set spans four specification categories: stochastically distributed household attributes, heterogeneous behavioural coefficients, deterministic financial thresholds, and fixed technical constants. Each category is calibrated to authoritative international data sources to ensure empirical grounding. Household income serves as the foundational stratification variable, drawn from the lognormal distribution to generate a right-skewed profile consistent with World Bank urban LMIC survey data. Credit access is represented by the Bernoulli variable C_i with a baseline population-wide success probability of 0.52, indicating that 48 % of the synthetic population faces complete financial exclusion at the first barrier stage [13]. The Bass innovation and imitation coefficients pi and qi are specified as heterogeneous, beta-distributed variables with ranges of 0.02 to 0.18 and 0.10 to 0.42 respectively [28]. The payback threshold τ_i encodes an income-proportional investment horizon tolerance ranging from 3.25 years for the lowest-income agents to 10.25 years for the highest-income stratum. Three parameters comprise fixed constants held uniform across agents and scenarios: the debt-to-income limit is set at 10 %, the social discount rate at 8 %, and the system lifetime at 20 years [8,27,32]. These constants provide a stable financial backdrop within which heterogeneous agent behaviour unfolds. A mean annual energy consumption of 4,500 kWh is drawn from IEA WEO urban LMIC benchmarks to anchor payback and net present value computations [3]. In all scenarios, each adopting household is assumed to install an identical standard rooftop PV system with a rated capacity of 4 kW. System size does not vary by income, quintile, or policy scenario, isolating the effect of financial barriers and policy instruments from system-sizing configurations.

Macro-economic environment specification

The time-varying macro-economic environment is parameterised as a set of deterministic annual trajectories indexed by simulation year $t \in \{1, \dots, 15\}$. These trajectories represent documented trends in residential solar PV technology costs, grid electricity prices, and credit market

conditions as projected by international energy institutions.

The nominal interest rate applicable to solar PV consumer loans follows a linear decline from prevailing LMIC market levels toward a long-run equilibrium value:

$$r_t = r_0 - \frac{r_0 - r_{15}}{14}(t - 1) \quad (6)$$

where $r_0 = 0.18$ represents the baseline interest rate at $t = 1$. This specification is calibrated to the World Bank Global Financial Development Database [12]. The data indicate that nominal lending rates in developing economies have experienced a gradual structural decline over the previous decade.

The residential grid electricity tariff increases linearly across the simulation horizon, reflecting the processes of subsidy rationalisation and cost recovery currently observed in many developing country utility sectors:

$$\tau_t^{\text{grid}} = \tau_0^{\text{grid}} + \frac{\tau_{15}^{\text{grid}} - \tau_0^{\text{grid}}}{14}(t - 1) \quad (7)$$

where $\tau_0^{\text{grid}} = \text{USD } 0.15 \text{ kWh}^{-1}$ represents the initial retail electricity price.

System capital cost follows an experience-curve trajectory calibrated to IRENA learning rate estimates for residential solar PV installations:

$$\text{CAPEX}_t = \text{CAPEX}_0 - \frac{\text{CAPEX}_0 - \text{CAPEX}_{15}}{14}(t - 1), \text{CAPEX}_0 = 4,000 \quad (8)$$

where $\text{CAPEX}_0 = \text{USD } 4,000$ represents the initial capital expenditure for a standard 4 kW system. This parameterisation encodes a 62.5% capital cost reduction over the 15-year horizon, consistent with an 18% cost decrease per doubling of cumulative installed capacity [8,34].

While the simulation horizon is 15 years, the net present value calculation extends to 20 years to represent a standard residential system lifetime. For years 16 to 20, the grid tariff is held constant at the Year 15 value, and the system capital cost remains fixed, while the performance ratio continues to capture degradation. This extrapolation reflects a conservative assumption in which tariff growth does not continue beyond the simulation horizon, ensuring structural consistency between the 15-year policy horizon and the 20-year financial appraisal.

Grid parity is defined as the first simulation year in which the levelised cost of solar energy generation equals or falls below the prevailing grid retail tariff:

$$t^* = \min\{t : \text{LCOE}_t \leq \tau_t^{\text{grid}}\} \quad (9)$$

where LCOE_t denotes the levelised cost of electricity from the residential solar PV system in year t , computed as the annualised capital cost divided by annual energy yield and adjusted for system degradation over time. Under the baseline parameter trajectories for τ_t^{grid} and CAPEX_t , the relation $\text{LCOE}_1 < \tau_1^{\text{grid}}$ is satisfied. Consequently, grid parity is attained at $t^* = 1$, meaning the condition is met from the first simulation year onward. This is consistent with current market observations showing that rooftop solar LCOE is frequently below retail tariffs in many LMIC urban markets. This Year 1 crossing is reflected consistently across all figures, tables, and downstream discussion in this study.

Sequential financial barrier decision architecture

The core analytical feature is the sequential, nested financial barrier decision tree traversed by each non-adopted agent at every annual time step. The baseline architecture evaluates credit access first. Loan affordability is evaluated second. Payback tolerance is evaluated third. Investment viability is evaluated fourth. Peer-effect-driven adoption is evaluated fifth and last. This structure operationalises the proposition that financial barriers are ordered in a hierarchy. The structure embeds

an assumption that formal credit access functions as a necessary precondition for progressing through subsequent evaluation stages.

Stage 1 constitutes the credit access barrier. This stage evaluates whether an agent without formal credit access can self-finance the required capital expenditure from own resources [35]. The barrier indicator is defined as follows:

$$\text{Barrier}_1(i, t) = 1(C_i = 0) \wedge 1\left(\frac{0.20 \times Y_i}{\text{CAPEX}_t} < 1\right) \quad (10)$$

In this equation, Y_i denotes annual household income and CAPEX_t denotes the system capital cost in year t . The parameter value 0.20 represents the maximum fraction of annual income that a credit-excluded agent can realistically deploy as upfront capital. Agents for whom $\text{Barrier}_1(i, t) = 1$ are classified as credit constrained. These agents remain non-adopters in that year. Agents without formal credit access whose available self-finance exceeds 20 % of CAPEX_t are permitted to proceed to the next barrier stage.

Stage 2 constitutes the loan affordability barrier. Agents who pass Stage 1 proceed to an assessment of monthly loan repayment affordability relative to income. This assessment is operationalised via a debt-to-income constraint. The monthly payment is given by the standard annuity formula:

$$\text{PMT}(i, t) = \frac{\text{CAPEX}_t \times \left(\frac{r_t}{12}\right)}{1 - \left(1 + \frac{r_t}{12}\right)^{-L}}, L = 60 \quad (11)$$

The barrier indicator is defined as follows:

$$\text{Barrier}_2(i, t) = 1\left[\frac{\text{PMT}(i, t) \times 12}{Y_i} > \theta_{\text{DTI}}\right], \theta_{\text{DTI}} = 0.10 \quad (12)$$

Equation (11) computes the monthly repayment required to amortise the system cost over a 60-month loan term at the prevailing annual interest rate r_t . Equation (12) flags agents whose annualised debt service exceeds 10 % of household income. This threshold aligns with affordability benchmarks applied in developing economy solar finance assessments [22]. Agents for whom Equation (12) evaluates to one are considered debt burdened. These agents remain non-adopters in period t .

Stage 3 constitutes the payback period barrier. Agents who clear Stage 2 are evaluated against their income-dependent payback tolerance. The simple payback period is given as follows:

$$\text{PBP}(i, t) = \frac{\text{CAPEX}_t}{E_{\text{annual}} \times \eta_{\text{sys}} \times (\theta_{\text{self}} \times \tau_t^{\text{grid}} + \theta_{\text{export}} \times \rho_t \times \tau_t^{\text{grid}})} \quad (13)$$

The associated barrier indicator is defined as follows:

$$\text{Barrier}_3(i, t) = 1[\text{PBP}(i, t) > \tau_i] \quad (14)$$

In this equation, E_{annual} is mean annual household electricity consumption, calibrated to IEA WEO urban LMIC benchmarks [3]. The term $\eta_{\text{sys}} = 0.80$ is the system performance ratio, accounting for inverter losses and panel degradation. The terms θ_{self} and θ_{export} are the self-consumed and exported generation shares, calibrated in Section 2.5. The term ρ_t is the export compensation ratio relative to the retail tariff, equal to 0.50 under Baseline, Soft Loans, CAPEX Subsidy, and Grid Parity Wait, and equal to 1.10 under Feed-in Tariff, as defined in Equation (21) [32]. The performance ratio $\eta_{\text{sys}} = 0.80$ introduced in Equation (13) is applied identically and without exception in Equations (13), (15), and (25). This ratio is never set to zero in any equation in this study.

Stage 4 constitutes the net present value barrier. This stage employs a 20-year discounted cash flow analysis to assess whether the solar investment yields a positive net economic value. The net present value is defined as follows:

$$\text{NPV}(i, t) = \sum_{n=1}^{20} \frac{E_{\text{annual}} \times \eta_{\text{sys}} \times (\theta_{\text{self}} + \theta_{\text{export}} \times \rho_n) \times \tau_n^{\text{grid}}}{(1 + r_d)^n} - \text{CAPEX}_t \quad (15)$$

$$\text{Barrier}_4(i, t) = 1[\text{NPV}(i, t) < 0], r_d = 0.08 \quad (16)$$

The social discount rate $r_d = 0.08$ is held constant across agents. This choice isolates the effect of tariff trajectory uncertainty on investment viability. This choice is consistent with IRENA appraisal guidelines, which apply an 8 % real discount rate in comparable analyses [8,32]. Agents with negative $\text{NPV}(i, t)$ at Stage 4 are treated as financially constrained. These agents remain non-adopters in period t .

Stage 5 constitutes the Bass diffusion adoption stage. Agents that clear all four financial barrier tests reach a stochastic adoption step [28]. The probability of adoption in period t for agent i is governed by a heterogeneous Bass specification, defined as follows:

$$P_{\text{adopt}}(i, t) = p_i + q_i F(t) \quad (17)$$

$$F(t) = \frac{A(t)}{N} \quad (18)$$

The quantity $F(t)$ denotes the cumulative adoption fraction at the beginning of period t . The term $A(t)$ represents the total number of previously adopted agents. The innovation term p_i captures an agent-specific baseline adoption probability that is independent of social context. The imitation term $q_i F(t)$ introduces a social learning component that depends on the cumulative adoption fraction. This representation treats $F(t)$ as a population-averaged social signal. This representation does not model explicit neighbourhood networks or individual peer links. Household adoption probabilities are therefore functions of aggregate diffusion rather than direct interactions with specific peers. Adoption is treated as an absorbing state. Once agent i adopts at time t , the agent is removed from subsequent barrier evaluations.

Barrier-ordering sensitivity analysis

The baseline architecture evaluates credit access before economic viability and social exposure. This raises the question of whether this ordering choice, rather than the underlying population characteristics, drives the finding that financial exclusion dominates the constrained population. To test this directly, an alternative decision architecture was implemented and run under identical population and macro-economic parameters. In the alternative ordering, Stage 1 evaluates social and peer exposure through a stochastic awareness draw governed by the same innovation and imitation coefficients as Equation (17). Stage 2 evaluates economic viability, assessing payback and NPV jointly. Stage 3 evaluates financing access, assessing credit and affordability jointly using the same criteria as Equations (10) and (12). Agents failing any stage remain non-adopters in that year. Both orderings were run under the Baseline macro-economic trajectory using the identical random seed and agent population.

Table 2 shows that reordering the barrier sequence produces an identical terminal adopted population, penetration rate, equity score, and inflection timing under both configurations. This result follows directly from the model's logical structure. Final adoption requires an agent to jointly satisfy all five pass or fail conditions regardless of the order in which they are evaluated, since the underlying test is a logical conjunction, and conjunction is commutative. Reordering therefore changes only the descriptive attribution of exclusion at a given point in time. A marginally smaller share of blocked households is labelled credit-constrained under the alternative ordering, because credit is evaluated later and catches fewer residual cases after economic and awareness screening have already occurred. Reordering cannot change which agents are ultimately included in or excluded from the adopting population. This finding indicates that the identification of financial exclusion as the dominant constraint in this model is a structural

Table 2

Barrier-ordering sensitivity: adoption, attribution, and equity outcomes under baseline and alternative decision sequences.

Metric	Baseline ordering (Credit → Affordability → Payback → NPV → Peer)	Alternative ordering (Peer → Payback/ NPV → Credit/ Affordability)	Difference
15-year terminal penetration rate (%)	72.75	72.75	0.00
Share of blocked households — Lack of Credit, Year 15 (%)	1.65	0.37	+1.28
Share of blocked households — Unaffordable Loan, Year 15 (%)	0.00	0.00	0.00
Share of blocked households — Payback/NPV constrained, Year 15 (%)	0.00	0.00	0.00
Share of blocked households — Behavioural/Not-yet-aware, Year 15 (%)	98.35	99.63	−1.28
Terminal equity score (0–10 scale)	6.89	6.89	0.00
Year of S-curve inflection (peak annual new adopters)	14	14	0.00

property of the underlying population and parameterisation, rather than an artefact of the specific stage ordering chosen for the baseline architecture.

Self-consumption and export ratio sensitivity

The self-consumed and exported generation shares are calibrated with reference to empirical residential PV metering data, rather than assumed a priori. McKenna et al. [36] analysed over 1,300 fully metered residential PV households across six countries, comprising the UK, USA, Australia, Germany, the Netherlands, and Belgium. This analysis found that, on average, only 49 % of PV generation is self-consumed, with 51 % exported. This outcome is driven by the mismatch between daytime generation and household demand timing. No equivalent metering dataset currently exists for urban LMIC households. Consumption timing profiles in these markets may differ from the cited high-income benchmark due to informal home-based economic activity, appliance ownership patterns, and grid reliability constraints. Given this evidentiary gap, this ratio is treated as a sensitivity parameter rather than a fixed assumption. The primary results reported in Sections 3 and 4 retain $\theta_{\text{self}} = 0.70$ and $\theta_{\text{export}} = 0.30$ as an upper-bound case for comparability with the original model design. Table 3 reports the full sensitivity of headline outcomes to this parameter, including the empirically central 0.50/0.50 case.

Table 3 demonstrates that this sensitivity is substantial rather than marginal across the full range of self-consumption specifications tested. As the self-consumption share θ_{self} declines from the manuscript's primary specification of 0.70 to the empirically central estimate of 0.50 reported by McKenna et al. [36], terminal penetration after 15 years falls correspondingly from 72.75 % to 53.65 %, while cumulative utility revenue loss by Year 15 declines by nearly half, from USD 611,100 to USD 321,900. This pattern indicates that the headline adoption and revenue-loss figures reported in Sections 3 and 4 should be interpreted as an upper-bound scenario, since they rest on a self-consumption share

Table 3

Sensitivity of adoption, equity, and utility revenue outcomes to the self-consumption/export ratio.

$\theta_{\text{self}} / \theta_{\text{export}}$	15-year terminal penetration (%)	Terminal equity score (0–10)	Cumulative utility revenue loss, Year 15 (USD)	Grid parity year
0.40 / 0.60	44.55	3.34	213,840	1
0.50 / 0.50 (empirically central estimate, McKenna et al.)	53.65	3.71	321,900	1
0.60 / 0.40	68.15	5.87	490,680	1
0.70 / 0.30 (adopted for primary results, upper-bound case)	72.75	6.89	611,100	1

considerably more favourable to household economics than the cross-country metered evidence supports. Where a single central estimate is required for comparison against external benchmarks, the 0.50/0.50 case in Table 3 constitutes the more empirically defensible reference point. Grid parity timing, fixed at $t^* = 1$ across all four specifications in Table 3, remains unaffected by this parameter, since parity is governed solely by the LCOE and tariff comparison defined in Equation (9), which does not depend on the self-consumption and export split.

Policy scenario specifications

Five energy-policy scenarios are evaluated by modifying the macro-economic parameter trajectories defined in Section 2.3, while holding all agent-level characteristics constant. This design enables assessment of how alternative intervention designs influence distributed solar PV diffusion under identical system conditions. The Baseline scenario applies the unadjusted market trajectories for interest rates, grid tariffs, and capital costs, as represented in Equations (6) to (8). Each alternative scenario adjusts one or more parameters to reflect specific policy instruments that are commonly deployed in LMIC residential solar PV programmes.

The Soft Loans scenario [19] replaces the market interest rate with a concessional fixed rate for solar PV loans. This relationship is defined as follows:

$$r_t^{\text{SL}} = 0.06 \forall t \quad (19)$$

This parameterisation reflects documented green credit facilities, in which development finance institutions provide subsidised capital at interest rates between 4 % and 8 % in order to support clean energy investment in developing economies.

The CAPEX Subsidy scenario [37] reduces the effective capital cost borne by households through a 30 % government capital grant. This relationship is defined as follows:

$$\text{CAPEX}_t^{\text{SUB}} = \text{CAPEX}_t \times (1 - \gamma), \gamma = 0.30 \quad (20)$$

The Feed-in Tariff scenario [18] augments household revenues through a premium export rate applied to surplus generation. This change increases the annual financial benefit used in the payback and net present value calculations in Equations (13) and (15). This relationship is defined as follows:

$$B_t^{\text{FIT}} = \tau_t^{\text{grid}} E_{\text{self}} + \tau_t^{\text{FIT}} E_{\text{export}}, \tau_t^{\text{FIT}} = 0.10 \quad (21)$$

In this equation, E_{self} and E_{export} denote the self-consumed and exported energy components respectively.

The Grid Parity Wait scenario retains macro-economic trajectories identical to Baseline, as defined in Equations (6) to (8), and introduces no fiscal instrument. Since grid parity is attained at $t^* = 1$ under the

trajectory described in Section 2.3, a deferral mechanism tied to the LCOE and tariff crossing year would produce no meaningful pre-parity period against which to measure an effect. Grid Parity Wait is therefore specified as a fixed institutional deferral of eight years, representing the grid-integration studies, interconnection permitting queues, and utility approval timelines typically documented in LMIC distribution utility contexts, rather than technology cost-competitiveness. Under this specification, households satisfying all financial barrier conditions before Year 8 remain non-adopters until Year 8, after which the standard five-stage sequence applies without modification. The resulting difference in terminal outcomes between Baseline and Grid Parity Wait is therefore attributable solely to this institutional deferral mechanism, and this difference is quantified in Section 3 [18].

Baseline itself permits adoption only for agents who independently clear the NPV test defined in Equation (16), so no agent adopts an investment identified as financially unviable under the model's own decision criteria. Early Baseline adopters constitute a self-selected subset of financially advantaged households rather than a representative population cross-section, a direct and expected consequence of Stage 4 screening.

Equity assessment framework

Distributional equity is quantified through income quintile stratification, using the fixed income-threshold quintile definitions specified in Section "Agent Population Initialisation". Quintile population sizes are unequal by construction. The quintile-specific adoption rate at year t is defined as follows:

$$\rho_k(t) = \frac{A_k(t)}{N/5}, k \in \{1, 2, 3, 4, 5\} \quad (22)$$

In this equation, $\rho_k(t)$ denotes the adoption rate of income quintile k at year t , $A_k(t)$ denotes the number of adopted agents in quintile k at time t , N denotes the total household population of 2,000, and the term $N/5$ approximates the nominal quintile size, noting that actual quintile sizes N_k are unequal by construction. Building on this quintile-level measure, the equity gap between the highest and lowest income quintile adoption rates is defined as follows:

$$\Delta_{\text{equity}}(t) = \rho_5(t) - \rho_1(t) \quad (23)$$

Here, $\Delta_{\text{equity}}(t)$ denotes the absolute percentage-point gap between the top quintile adoption rate $\rho_5(t)$ and the bottom quintile adoption rate $\rho_1(t)$ at year t , with a larger value indicating a wider disparity in adoption between the wealthiest and poorest quintiles. This gap captures the net effect of two countervailing structural tendencies within the model. The first is the positive correlation between income and solar PV investment returns documented in stratified household studies, which favours higher-income adoption and is reinforced by the income-linked credit access specification in Section "Agent Population Initialisation". The second arises from the model's payback-tolerance specification in Equation (5), under which lower-income quintiles saturate proportionally faster once financial barriers are cleared, given fixed absolute savings and the longer payback tolerance assigned to higher-income agents. The net quintile ordering observed in the Results section reflects the balance between these two mechanisms. This ordering should not be interpreted as a direct empirical claim about real-world adoption patterns by income [24].

To express this gap on a standardised scale, a composite equity score normalised to a 0 to 10 scale is computed as follows:

$$\text{EQ}_{\text{score}} = 10 \times \left(1 - \frac{\Delta_{\text{equity}}(T)}{\rho_5(T)} \right), T = 15 \quad (24)$$

In this equation, EQ_{score} denotes the terminal equity score, $\Delta_{\text{equity}}(T)$ is the equity gap evaluated at the terminal year $T = 15$, and $\rho_5(T)$ is the top-quintile adoption rate at that year, used to normalise the gap into a

proportional measure. The ratio $\Delta_{\text{equity}}(T)/\rho_5(T)$ expresses the equity gap as a fraction of the maximum observed adoption rate, so that subtracting this fraction from 1 and scaling by 10 yields a score in which 10 represents perfect equity and progressively lower values represent widening disparities. Higher values of EQ_{score} therefore indicate more equitable adoption distributions across income strata at the terminal simulation year.

Utility revenue impact and distributed energy management implications

Having established how adoption is distributed across income groups, attention turns to the financial consequences of that adoption for incumbent electricity utilities. This impact is estimated by computing annual energy displacement and the associated revenue loss attributable to prosumer adoption [38]. The displaced grid energy is given as follows:

$$\text{ED}(t) = A(t)E_{\text{self}}\eta_{\text{system}}, \eta_{\text{system}} = 0.80 \quad (25)$$

Here, $\text{ED}(t)$ denotes the total annual energy displaced from the grid, in kWh, at year t , $A(t)$ denotes the cumulative number of adopted households at year t , E_{self} denotes the self-consumed energy per household as defined in Section 2.5, and $\eta_{\text{system}} = 0.80$ denotes the fixed system performance ratio, accounting for inverter losses and panel degradation. This displaced energy translates directly into lost utility revenue, given as follows:

$$\text{RL}(t) = \text{ED}(t)\tau_t^{\text{grid}} \quad (26)$$

Here, $\text{RL}(t)$ denotes the annual utility revenue loss at year t , and τ_t^{grid} denotes the prevailing grid retail tariff at year t , as defined in Section 2.3. Cumulative revenue loss over the simulation horizon is obtained by summing $\text{RL}(t)$ across all 15 years.

Monte Carlo uncertainty quantification

The deterministic outcomes described in the preceding sections represent point estimates under a single central macro-economic trajectory. To characterise how sensitive these outcomes are to macro-economic uncertainty, the simulation pipeline additionally propagates parameter uncertainty through 200 independent Monte Carlo replications. In each replication $m \in \{1, \dots, 200\}$, stochastic perturbations are applied to the three primary macro-economic parameters. The perturbed capital cost trajectory is formalised first, as follows:

$$\text{CAPEX}_t^{(m)} = \text{CAPEX}_t \times \varepsilon_C^{(m)}, \varepsilon_C^{(m)} \sim \text{Uniform}(0.88, 1.12) \quad (27)$$

Here, $\text{CAPEX}_t^{(m)}$ denotes the perturbed capital cost in year t for replication m , CAPEX_t denotes the central deterministic capital cost trajectory, and $\varepsilon_C^{(m)}$ denotes a multiplicative noise factor drawn independently for each replication from a uniform distribution bounded between 0.88 and 1.12, representing $\pm 12\%$ forecast uncertainty. In parallel, the grid tariff and interest rate trajectories are each perturbed additively. The perturbed grid tariff trajectory is defined as follows:

$$\tau_t^{\text{grid},(m)} = \tau_t^{\text{grid}} + \varepsilon_\tau^{(m)}, \varepsilon_\tau^{(m)} \sim N(0, \sigma_\tau^2) \quad (28)$$

Here, $\tau_t^{\text{grid},(m)}$ denotes the perturbed grid tariff for replication m at year t , and $\varepsilon_\tau^{(m)}$ denotes an additive noise term drawn from a normal distribution with mean zero and variance σ_τ^2 , calibrated to historical tariff volatility. The perturbed interest rate trajectory follows the same logic, expressed as follows:

$$r_t^{(m)} = r_t + \varepsilon_r^{(m)}, \varepsilon_r^{(m)} \sim N(0, \sigma_r^2) \quad (29)$$

Here, $r_t^{(m)}$ denotes the perturbed interest rate for replication m at year t , and $\varepsilon_r^{(m)}$ denotes an additive noise term drawn from a normal distribu-

tion with mean zero and variance σ_r^2 , calibrated to historical interest rate volatility [31].

Having generated 200 perturbed replications, summary statistics are then computed across them. The Monte Carlo point estimator for any scalar output statistic θ is computed as follows:

$$\hat{\theta} = \frac{1}{200} \sum_{m=1}^{200} \theta^{(m)} \quad (30)$$

Here, $\hat{\theta}$ denotes the Monte Carlo mean estimate of statistic θ , and $\theta^{(m)}$ denotes the value of that statistic produced by replication m , averaged across all 200 replications. From this same set of replications, a 95 % confidence interval is constructed empirically, using the 2.5th and 97.5th percentiles of the 200 replication values as its lower and upper bounds, and is expressed analytically as follows:

$$CI_{0.95}(\theta) = \left[\hat{\theta} - 1.96 \frac{s_{\theta}}{\sqrt{200}}, \hat{\theta} + 1.96 \frac{s_{\theta}}{\sqrt{200}} \right] \quad (31)$$

Here, s_{θ} denotes the sample standard deviation of θ across the 200 replications, and the term $s_{\theta}/\sqrt{200}$ denotes the standard error of the mean, so that the interval spans 1.96 standard errors either side of $\hat{\theta}$. This interval represents the range of outcomes generated by macro-economic uncertainty, rather than a confidence interval for the sample mean in the classical repeated-sampling sense. Finally, to judge whether this spread is acceptably narrow, the coefficient of variation is computed as follows:

$$CV(\theta) = \frac{s_{\theta}}{\hat{\theta}} \times 100\% \quad (32)$$

Here, $CV(\theta)$ denotes the coefficient of variation, expressing the sample standard deviation s_{θ} as a percentage of the mean estimate $\hat{\theta}$. This coefficient serves as the primary robustness criterion, with $CV(\theta) < 10\%$ adopted as the pre-specified robustness threshold. All Monte Carlo summary statistics reported in Section 3 are drawn from this single, identical 200-replication set, with no separate or re-derived Monte Carlo distribution used for any chart or table in this study [39].

Multi-criteria energy transition management evaluation framework

Beyond adoption, equity, and revenue outcomes considered individually, the five policy scenarios are also compared holistically using a normalised multi-criteria scoring framework spanning four performance dimensions. These dimensions comprise the cumulative adoption rate at Year 15, the terminal social equity score based on the income quintile adoption gap, the fiscal efficiency measured as adoption per unit of public expenditure, and the implementation feasibility assessed qualitatively from documented administrative requirements. Each scenario s receives a normalised score $\phi_{s,d} \in [0, 100]$ on each dimension d . For dimensions in which higher raw values are preferable, such as adoption and equity, a min to max normalisation is applied, as follows:

$$\phi_{s,d} = \frac{x_{s,d} - \min_s x_{s,d}}{\max_s x_{s,d} - \min_s x_{s,d}} \times 100 \quad (33)$$

Here, $\phi_{s,d}$ denotes the normalised score of scenarios s on dimension d , $x_{s,d}$ denotes the raw metric value of scenario s on dimension d , and $\min_s x_{s,d}$ and $\max_s x_{s,d}$ denote the minimum and maximum raw values observed across all five scenarios on that dimension. For dimensions where lower raw values are preferable instead, such as fiscal cost per unit of adoption, this normalisation is reversed, as follows:

$$\phi_{s,d}^{\text{rev}} = \frac{\max_s x_{s,d} - x_{s,d}}{\max_s x_{s,d} - \min_s x_{s,d}} \times 100 \quad (34)$$

Here, $\phi_{s,d}^{\text{rev}}$ denotes the reversed normalised score, applying the same

definitions of $x_{s,d}$, $\min_s x_{s,d}$, and $\max_s x_{s,d}$ as above, but inverted so that a lower raw value yields a higher normalised score. Implementation feasibility scores are assigned using a structured rubric grounded in the complexity of required institutional changes and the maturity of relevant regulatory frameworks. Once all four-dimension scores are computed for a given scenario, the composite policy effectiveness score is obtained as their unweighted average, as follows:

$$\Phi_s = \frac{1}{4} \sum_{d=1}^4 \phi_{s,d} \quad (35)$$

Here, Φ_s denotes the composite effectiveness score for scenario s , computed as the arithmetic mean of its four normalised dimension scores $\phi_{s,d}$. Scenarios are ranked in descending order of Φ_s , with the highest-ranked scenario interpreted as the most effective configuration under the defined criterion set and weight structure. Because the normalisation in Equations (33) and (34) is scaled relative to the minimum and maximum values observed among the five evaluated scenarios, the scenario positioned at either extreme of a given dimension receives a score of exactly 0 or 100 on that dimension, independent of the absolute size of the gap separating it from the next-ranked scenario. Composite score differences reported in Section 4 should accordingly be interpreted as relative rankings under this scoring structure, rather than as proportional measures of absolute performance gaps. The adoption, equity, fiscal, and feasibility indicators underpinning Table 6 are documented in Supplementary Dataset 13, which provides the intermediate normalised values and full computational tracking to ensure reproducibility.

External benchmarking and model validation

The credibility of a heterogeneous household simulation depends on the correspondence between model-generated outputs and empirical quantities observed in comparable settings. To establish this correspondence, a fifteen-indicator external benchmarking procedure was implemented, restricted strictly to endogenously generated simulation outputs. Two parameters set directly as calibration inputs, the nominal lending rate trajectory in Equation (6) and the income distribution parameters in Equation (1), together with the credit access rate produced under the income-linked specification in Equation (2b), are excluded from the validation PASS, WARN, or FAIL count under any classification, since comparing a calibration input against its own calibration target constitutes circular validation. These three inputs are instead reported separately in Supplementary Table S1 as calibration diagnostics, confirming that model inputs fall within documented institutional ranges.

Primary simulation outputs, by contrast, were compared against empirical ranges reported by five independent international data sources, strictly isolated from the data used during initial model parameterisation and calibration to preserve validation integrity. The benchmarking architecture spans adoption trajectories, financial performance metrics, distributional parameters, Bass diffusion coefficients, macro-economic trajectories, investment viability metrics, equity gap indicators, and Monte Carlo uncertainty statistics. For each validation test v , indexed from 1 to 15, the simulated metric $\hat{m}_v$ is evaluated against an independent benchmark interval $[m_v^{\text{lo}}, m_v^{\text{hi}}]$ using a three-tier classification rule, given as follows:

$$\text{Status}_v = \begin{cases} \text{WARN} & \text{if } |\hat{m}_v - m_v^{\text{nearest}}| \leq \varepsilon m_v^{\text{nearest}} \\ \text{PASS} & \text{if } \hat{m}_v \in [m_v^{\text{lo}}, m_v^{\text{hi}}] \\ \text{FAIL} & \text{otherwise} \end{cases} \quad (36)$$

Here, $\hat{m}_v$ denotes the simulated value of validation test v , $[m_v^{\text{lo}}, m_v^{\text{hi}}]$ denotes the lower and upper bounds of the published benchmark interval for that test, m_v^{nearest} denotes whichever of these two boundaries lies closer to $\hat{m}_v$, and $\varepsilon = 0.05$ denotes the fixed 5 % tolerance band

applied to that nearest boundary. A PASS classification indicates that the simulated value falls within the published benchmark range. A WARN classification identifies cases where the simulated value falls outside the benchmark range but within 5 % of the nearer boundary, reflecting marginal deviation attributable to the generalised parameters of the urban LMIC market archetype rather than a specification error. A FAIL classification indicates that the simulated value falls outside both the benchmark range and the surrounding tolerance band.

The fifteen validation tests, together with their simulated values, benchmark ranges, and resulting classifications, are reported in full in Supplementary Table S2. The tests span adoption trajectory validation against IEA PVPS Task 1 cross-country ranges, financial performance validation of household internal rate of return and payback period, distributional parameter validation of the simulated equity gap, Bass diffusion coefficient validation against Rode and Weber's cross-country estimates, macro-economic trajectory validation of CAPEX decline rates and grid tariff levels, investment viability validation of the 20-year household NPV, and Monte Carlo robustness validation of the coefficient of variation across 200 replications. Beyond this categorical PASS, WARN, or FAIL classification, a continuous measure of deviation is also reported for each test, since two results can both be classified as PASS while differing considerably in how closely they track the benchmark. For continuous-valued tests, this percentage deviation relative to the benchmark midpoint is computed as follows:

$$\delta_v = \frac{\hat{m}_v - m_v^{\text{mid}}}{m_v^{\text{mid}}} \times 100\%, m_v^{\text{mid}} = \frac{m_v^{\text{lo}} + m_v^{\text{hi}}}{2} \quad (37)$$

Here, δ_v denotes the signed percentage deviation of the simulated value from the benchmark midpoint m_v^{mid} , which is itself the arithmetic mean of the benchmark's lower and upper bounds. This deviation metric is reported alongside the PASS, WARN, or FAIL classification for every test in Supplementary Table S2, since a test can be classified as PASS while still sitting far from the benchmark midpoint, and this distinction is relevant to interpreting the strength, rather than merely the direction, of correspondence between model and benchmark.

For interval-valued outputs, specifically the Monte Carlo derived confidence interval, a different measure is required, since the comparison is between two ranges rather than two-point values. The overlap coefficient for this purpose is given as follows:

$$\Omega_v = \frac{\min(\hat{m}_v^{\text{hi}}, m_v^{\text{hi}}) - \max(\hat{m}_v^{\text{lo}}, m_v^{\text{lo}})}{\hat{m}_v^{\text{hi}} - \hat{m}_v^{\text{lo}}} \quad (38)$$

Here, Ω_v denotes the fraction of the simulated 95 % confidence interval $[\hat{m}_v^{\text{lo}}, \hat{m}_v^{\text{hi}}]$, derived from Equation (31), that overlaps with the benchmark interval $[m_v^{\text{lo}}, m_v^{\text{hi}}]$. An overlap coefficient $\Omega_v \geq 0.50$ indicates that at least half of the simulated interval falls within the benchmark range, and this threshold is used as the criterion for strong interval correspondence.

Validation data are drawn entirely from open-access publications, comprising the IEA Photovoltaic Power Systems Programme Task 1 [27], IRENA Renewable Power Generation Costs [8], World Bank Global Findex Database [13], World Bank Global Financial Development Database [12], and solar prosumer diffusion studies archived in the KTH DiVA repository [40]. Drawing on multiple independent institutional sources across technical, financial, and distributional dimensions reduces the risk of correlated benchmark errors, which can arise when validation relies on a single reference dataset.

Applying this framework, of the fifteen endogenous validation tests, eleven returned a PASS classification and four returned a FAIL classification, with no test returning a WARN classification. The four failed tests were the household project internal rate of return, the Monte Carlo coefficient of variation, the income distribution Kolmogorov–Smirnov test, and the Grid Parity Wait versus Baseline distinctiveness check. These failures indicate that the Year 1 case study household falls below

the return and payback thresholds typically documented for viable residential solar investment in comparable markets, and that the Monte Carlo uncertainty layer does not currently meet the 10 % robustness threshold specified in Equation (32). These results are reported in full in Supplementary Table S2 and are examined as a limitation in Section 5.

Computational implementation and reproducibility

The full modelling and validation framework described above was implemented in Python 3.14.6, using the scientific Python ecosystem. Core array operations and stochastic sampling procedures were executed using NumPy 1.26.0. Data management, aggregation, and output generation were handled with Pandas 2.1.0. High-resolution figures were generated using Matplotlib 3.8.0, applying a clean vector graphic format.

Random number generation used the NumPy default generator, with a fixed master seed of 5,050 for deterministic agent population initialisation. Monte Carlo replications were assigned unique derived seeds using a deterministic linear schedule, $\text{seed}_m = m \times 7$ for $m \in \{1, \dots, 199\}$, a configuration that ensures cross-run independence and full reproducibility.

Computation was performed on an Apple MacBook Pro equipped with an Apple M1 Max system-on-chip, a 32-core GPU, and 32 GB of unified memory. A complete simulation run comprising all five policy scenarios and 200 Monte Carlo replications required approximately 3 min and 40 s of wall-clock time, reflecting the efficiency of the NumPy vectorised implementation, which avoids explicit Python loops for agent state updates.

Results

Scenario configurations and system-level performance

Two complementary result sets are reported in this section. Deterministic outcomes are computed under fixed baseline macroeconomic trajectories. Monte Carlo outcomes are derived from 200 stochastic replications with simultaneous perturbations of capital cost, grid tariffs, and interest rates. All numerical citations in this section identify their origin as either deterministic or Monte Carlo in order to avoid ambiguity between the two analytical layers.

Five counterfactual policy configurations are evaluated under identical agent population and macroeconomic conditions. Each configuration modifies specific financial instrument parameters relative to the Baseline trajectory, as specified in Equations (19) through (21). The resulting scenario set spans a continuum from passive market reliance to active capital-side intervention. Table 4 summarises the scenario specifications, the financial instrument parameters, and the simulated adoption, equity, and fiscal cost outcomes for each configuration.

Table 4 presents the scenario specifications and the corresponding performance outcomes for the five policy configurations. The table reports interest rate structures, capital cost settings, government support mechanisms, and three outcome dimensions spanning adoption, equity, and fiscal cost. The scenarios define a policy spectrum from passive market reliance under Grid Parity Wait to active capital-side intervention under CAPEX Subsidy.

The CAPEX Subsidy configuration achieves the strongest performance across both adoption and equity dimensions simultaneously, attaining a 98.4 % adoption rate and an equity score of 9.76 at a fiscal cost of USD 1.33 million. This represents a materially different policy conclusion from the previously reported result, in which CAPEX Subsidy was characterised as trading adoption performance for equity gains. The calibration instead shows no such tradeoff: this instrument dominates on both criteria.

The Soft Loans scenario fixes the lending rate at 6.0 %. This policy produces a 72.8 % adoption rate and an equity score of 6.89, both statistically indistinguishable from the Baseline values of 72.8 % and 6.89

Table 4

Scenario specifications, financial instrument parameters, and simulated distributed solar PV performance indicators within an urban LMIC market archetype.

Scenario	Interest Rate	CAPEX Cost	Government Support	15-yr Adoption (%)	Equity Score (0–10)	Fiscal Cost (USD M)
Baseline	Market (18 % → 12 %)	Full (\$4,000 → \$1,500)	None	72.8	6.89	0.00
Soft Loans (6 %)	Fixed 6 %	Full (\$4,000 → \$1,500)	Interest subsidy	72.8	6.89	0.95
CAPEX Subsidy (30 %)	Market (18 % → 12 %)	70 % of full cost	30 % capital grant	98.4	9.76	1.33
Feed-in Tariff	Market (18 % → 12 %)	Full (\$4,000 → \$1,500)	Export premium	90.5	8.62	0.45
Grid Parity Wait	Market (18 % → 12 %)	Full (\$4,000 → \$1,500)	None	72.9	6.28	0.00

Notes:

- CAPEX Subsidy Outcomes: The capital grant configuration, parameterised per Equation (20), yields the highest 15-year adoption rate among all five configurations at 98.4 %, a gain of 25.6 percentage points relative to Baseline. This configuration simultaneously achieves the highest equity score of 9.76 as computed by Equation (24), at a public expenditure of USD 1.33 million. This result indicates that the combined effect of the 30 % capital grant and the associated 30 % reduction in effective system cost is sufficient to substantially narrow the adoption gap between the lowest and highest income quintiles, rather than restricting aggregate adoption as previously reported.
- Grid Parity Wait: This non-intervention configuration retains the unadjusted macroeconomic trajectories from Equations (6) through (8) and applies an independent, fixed institutional deferral parameter through Year 8, rather than a rule tied to the LCOE-tariff crossing year. Because grid parity as defined in Equation (9) is achieved in Year 1 under the current CAPEX calibration, this deferral period does not materially suppress cumulative adoption relative to Baseline. The configuration produces a 15-year adoption rate of 72.9 %, nearly identical to Baseline, and the lowest equity score among the five configurations at 6.28, driven by the widest equity gap of 36.5 percentage points between Q1 and Q5 adoption rates. Fiscal cost is 0.00, since this configuration involves no government expenditure.
- Feed-in Tariff Performance: The export premium mechanism achieves the second-highest adoption rate at 90.5 %, a gain of 17.7 percentage points relative to Baseline, at a fiscal cost of USD 0.45 million. It achieves an equity score of 8.62, the second-highest among the five configurations, indicating that revenue-side intervention narrows the income-adoption gap more effectively than previously estimated.
- Soft Loans Performance: The concessional 6 % lending rate produces a 15-year adoption rate of 72.8 %, statistically identical to the Baseline value of 72.8 %, a difference of 0.0 percentage points, at a fiscal cost of USD 0.95 million. The equity score of 6.89 is likewise identical to Baseline. This confirms that, under the calibration, the lending interest rate is not the binding constraint on adoption or equity for the majority of households, since the primary barriers segmenting the income distribution operate upstream of the interest-rate-sensitive loan affordability stage.
- Distributional Disparities: Both zero-cost scenarios, Baseline and Grid Parity Wait, record equity scores below 7.0 per Equation (24), while the three active fiscal instruments each exceed this threshold. This confirms that unassisted market mechanisms continue to produce less equitable adoption distributions than any of the three active policy interventions, despite the technological cost improvements embedded in Equation (8).
- Adoption Range Variance: The total adoption range across all five configurations spans 25.6 percentage points, from a minimum of 72.8 % under Baseline and Soft Loans to a maximum of 98.4 % under CAPEX Subsidy. This variance is directly attributable to the barrier elimination capacity of the instrument-specific parameter modifications defined in Equations (19) through (21).
- Fiscal Efficiency: On a cost-per-adoption-point basis, Feed-in Tariff delivers the largest adoption gain (17.7 pp) at the lowest fiscal outlay among the three active instruments (USD 0.45 million), making it the most cost-efficient active mechanism, while CAPEX Subsidy delivers the largest absolute adoption and equity gains at the highest fiscal cost (USD 1.33 million), reflecting a direct tradeoff between per-unit efficiency and total programme impact.

respectively, at a fiscal cost of USD 0.95 million. Concessional lending alone does not measurably alter either aggregate adoption or distributional equity under the barrier decomposition, in which behavioural resistance rather than loan affordability constitutes the dominant late-horizon constraint.

The Feed-in Tariff scenario operates through revenue augmentation rather than upfront cost reduction. It achieves a 90.5 % adoption rate and an equity score of 8.62, at a fiscal cost of USD 0.45 million, the lowest fiscal outlay among the three active instruments. This mechanism represents the most fiscally efficient active instrument evaluated in this study on a cost-per-adoption-point basis.

The Baseline scenario reaches 72.8 % adoption and an equity score of 6.89 under unmodified market conditions. The Grid Parity Wait scenario yields a nearly identical 72.9 % adoption rate but the lowest equity score of the five configurations, at 6.28, since the fixed institutional deferral period delays adoption uniformly across income groups without structural mitigation of the underlying income-adoption gap. Instruments that actively lower financial or revenue-side barriers, namely CAPEX Subsidy and Feed-in Tariff, outperform both non-intervention configurations on every dimension reported in Table 4.

Adoption dynamics and financial barrier decomposition

The adoption, barrier, and equity patterns reported in this section are generated from the complete simulation architecture defined in the Methods. Macro-economic trajectories for capital expenditure, grid tariffs, and levelised cost of energy follow Equations (6) through (9). The five-stage financial barrier hierarchy governing credit access, loan affordability, payback tolerance, and discounted investment value is specified in Equations (10) through (16). Stochastic adoption for financially eligible households is governed by the heterogeneous Bass diffusion mechanism in Equations (17) and (18). Income-stratified adoption rates, equity gaps, and the composite equity score are computed using Equations (22) through (24). Adoption trajectories, 95% confidence intervals, and associated dispersion metrics are derived

from the Monte Carlo perturbation and estimation framework in Equations (27) through (32).

Panel A presents the macro-economic environment that underpins all subsequent analyses. System capital expenditure declines monotonically from USD 4,000 in Year 1 to USD 1,500 in Year 15, representing a 62.5% cost reduction consistent with documented residential solar PV learning curve evidence. Over the same period, the residential grid tariff increases from USD 0.150 to USD 0.250 kWh⁻¹, while the levelised cost of energy (LCOE) declines from USD 0.083 to USD 0.031 kWh⁻¹. The levelised cost of energy falls below the grid tariff from Year 1 onward, establishing grid parity, as defined in Equation (9), at the very outset of the simulation horizon rather than developing progressively across the first several years. The simulation therefore begins in a setting where the financial case for prosumer adoption is already established, and this favourable financial position is sustained through Year 15, during which solar generation remains increasingly lucrative for a large share of the population. These baseline conditions are validated against empirical inputs documented in Supplementary Table S1.

Panel B depicts deterministic adoption trajectories for all five policy scenarios, consistent with the configurations reported in Table 2. At Year 15, the Baseline scenario reaches 72.8 % adoption, and the Soft Loans scenario reaches 72.8 %, coinciding exactly with the Baseline trajectory (+0.0 percentage points). The CAPEX Subsidy scenario reaches 98.4 %, representing the highest terminal value among all five configurations (+25.6 percentage points relative to Baseline), while the Feed-in Tariff scenario reaches 90.5 % (+17.7 percentage points) and the Grid Parity Wait scenario reaches 72.9 % (+0.1 percentage points). All five curves exhibit the characteristic S-shaped diffusion pattern associated with Bass-type processes. The CAPEX Subsidy and Feed-in Tariff trajectories display markedly earlier and steeper inflection points, concentrated between Years 9 and 13, than the three lower-performing configurations, whose inflection points cluster closer to Years 13 through 15. This acceleration corresponds to the phase in which social learning dynamics reinforce the barrier-reduction effect of these two instruments, driving faster uptake among the financially eligible household population. This

year-by-year diffusion trajectory is structurally robust to alternative barrier evaluation sequences, as documented in Supplementary Table S3.

Panel C disaggregates the Baseline financial barrier structure across the 15-year horizon, revealing a dynamic re-composition of constraint typology over time. In Year 1, the long payback period constitutes the dominant impediment and accounts for approximately 54 % of households that do not adopt. Lack of credit and unaffordable loan repayment obligations collectively explain a smaller share of near-term barriers. By Year 15, behavioural resistance rises to account for approximately 98 % of all remaining blocked households, reflecting the progressive relaxation of credit, affordability, and payback barriers as capital costs decline and tariffs increase across the simulation horizon. This barrier decomposition indicates that credit exclusion and loan affordability are not the dominant late-horizon constraints under the current calibration. The residual non-adopting population by Year 15 is instead overwhelmingly composed of households that have cleared all financial barriers but remain non-adopters for behavioural reasons. The evolution of the barrier composition therefore indicates that instruments acting solely on credit access, interest rates, or export revenues cannot remove the dominant constraint category in the later years of the horizon, since that constraint is behavioural rather than financial.

Panel D presents cumulative rooftop solar PV adopters segmented by income quintile across all five policy scenarios, consistent with Table 4. Quintile population sizes are unequal, reflecting the underlying lognormal income distribution of the synthetic agent population, and adopter counts vary accordingly by scenario. Under the Baseline scenario, total adoption reaches 1,455 households by Year 15, comprising 440 adopters in Q_1 (income below USD 10,000), 519 adopters in Q_2 (USD 10,000 to USD 20,000), 296 adopters in Q_3 (USD 20,000 to USD 35,000), 155 adopters in Q_4 (USD 35,000 to USD 60,000), and 45 adopters in Q_5 (income above USD 60,000). Under the Soft Loans scenario, total adoption reaches 1,455 households, identical to the Baseline total at every income quintile ($Q_1 = 440$, $Q_2 = 519$, $Q_3 = 296$, $Q_4 = 155$, $Q_5 = 45$). Under the CAPEX Subsidy scenario, total adoption reaches 1,967 households, the highest total among all five configurations, comprising 693 adopters in Q_1 , 686 adopters in Q_2 , 367 adopters in Q_3 , 171 adopters in Q_4 , and 50 adopters in Q_5 . Under the Feed-in Tariff scenario, total adoption reaches 1,810 households, comprising 600 adopters in Q_1 , 633 adopters in Q_2 , 358 adopters in Q_3 , 170 adopters in Q_4 , and 49 adopters in Q_5 . Under the Grid Parity Wait scenario, total adoption reaches 1,458 households, comprising 437 adopters in Q_1 , 497 adopters in Q_2 , 322 adopters in Q_3 , 153 adopters in Q_4 , and 49 adopters in Q_5 .

Across all five scenarios, adoption remains concentrated in the lower- and middle-income quintiles in absolute terms, since Q_1 and Q_2 together comprise the majority of the underlying agent population. However, when adopter counts are expressed as a share of each quintile's own population rather than as absolute counts, the highest income quintile exhibits the highest within-quintile adoption rate in every scenario. The equity gap between Q_5 and Q_1 penetration rates, computed via Equation (23), ranges from 2.4 percentage points under CAPEX Subsidy to 36.5 percentage points under Grid Parity Wait, as reported in Table 4. This pattern is consistent with the income-dependent payback tolerance parameter defined in Equation (5), under which higher-income households are assigned a longer acceptable payback period, reducing their relative urgency to adopt in early years, while annual electricity savings in Equation (13) are held constant across the population regardless of income. The CAPEX Subsidy configuration narrows this within-quintile equity gap to its minimum across all five scenarios, consistent with its equity score of 9.76 in Table 2, indicating that direct capital-cost reduction is the most effective instrument evaluated in this study for closing the income-adoption gap. Long-term uncertainty surrounding these trajectories is bounded by the stochastic distributions mapped in Supplementary Fig. S1 and Supplementary Fig. S2. The overall model robustness metrics, including the terminal mean

penetration rate of 74.4 % and its associated variance components, are detailed in the endogenous validation suite in Supplementary Table S2.

Household financial viability, market accessibility, and system uptake constraints

The household-level and market-level viability results reported in this section are derived from the financial and stochastic components of the framework. Household cash flow and investment performance are computed using the annuity payment formulation, simple payback period, and 20-year net present value expressions in Equations (11), (13), and (15). Market accessibility is evaluated by applying the credit access and debt-to-income constraints defined in Equations (10) and (12). Macro-economic uncertainty is propagated through the capital cost, grid tariff, and interest rate perturbations specified in Equations (27) through (29). Penetration means, 95% confidence intervals, and coefficients of variation are obtained using Equations (30) through (32).

Panel A presents the annual financial profile of a representative adopting household across the 15-year horizon. The left axis records annual net cash flows in USD, and the right axis records the cumulative financial balance. During the loan repayment phase, spanning Years 1 through 5, annual electricity bill savings of USD 306 to USD 370 are offset by annual loan repayment obligations of USD 1,219, which together generate a cumulative deficit that reaches approximately USD 8,400 at its lowest point at the end of Year 5. From Year 6 onward, all debt obligations are discharged, and annual net cash flows equal annual savings, rising from USD 380 to USD 510 between Years 6 and 15. The cumulative balance remains negative throughout the entire 15-year horizon and does not recover to a positive level within the simulation period, reaching approximately negative USD 3,974 by Year 15. The project internal rate of return (IRR) for this investment is 5.33 %, which sits below the 8% social discount rate employed in the discounted cash flow assessment. This trajectory indicates that the representative adopting household does not achieve breakeven within the 15-year simulation horizon under baseline financing terms. The trajectory illustrates the temporal mismatch between a substantial upfront capital burden and the delayed accrual of financial benefits in the simulated urban market archetype. These household-level investment findings are consistent with the baseline validation failures reported in Supplementary Table S2.

Panel B examines the effect of nominal lending rates on the size of the addressable market. The left axis reports the number of households whose annualised loan payments remain within the 10% debt-to-income threshold, while the right axis reports the corresponding annual payment in USD. At a concessional lending rate of 6 %, the addressable market comprises 821 households, equivalent to 41.0 % of the synthetic population. As the lending rate increases, the number of financially eligible households declines monotonically. At a market rate of 18 %, the addressable market falls to 703 households, or 35.1 % of the population. At a 30 % lending rate, only 590 households, or 29.5 % of the population, satisfy the affordability constraint. Across the full evaluated range from 4% to 30%, the annual loan payment increases from approximately USD 900 to over USD 1,550, reflecting the non-linear relationship between nominal interest rates and annuity payments. These parameters map directly onto the comprehensive interest rate sensitivity array presented in Supplementary Fig. S1.

Panel C maps the two-dimensional space spanned by grid tariffs and system capital costs onto a simple payback period surface. The horizontal axis spans tariffs between USD 0.15 and USD 0.25 kWh⁻¹, and the vertical axis spans capital costs between USD 1,000 and USD 6,000. Simple payback periods range from 1.18 years in the low-cost, high-tariff region to 17.85 years in the high-cost, low-tariff region. A trajectory connecting the Year 1 macro-economic configuration (USD 4,000 capital cost, USD 0.15 kWh⁻¹ tariff) to the Year 15 configuration (USD 1,500 capital cost, USD 0.25 kWh⁻¹ tariff) traverses this surface.




Fig. 2. Adoption dynamics, barrier structure, and income-stratified equity in distributed residential solar PV uptake.

This trajectory crosses the 10-year niche boundary and the 7-year early majority contour, ultimately reaching the 4-year mass-adoption contour by Year 15. The visualised path defines the intersection of technology cost reductions and tariff adjustments in expanding the subset of households for which solar PV becomes financially viable under the payback criterion.

Panel D presents kernel density estimates of the prosumer penetration rate distribution at Years 5, 10, and 15, based on 200 Monte Carlo replications. At Year 5, penetration is tightly concentrated near a mean of 0.3 % with a standard deviation of 0.2 %, reflecting an early diffusion phase strongly constrained by credit access and affordability filters. By Year 10, the distribution broadens and is centred around a mean of 6.2 % with a standard deviation of 2.6 %. At Year 15, the penetration distribution converges toward a Baseline mean terminal penetration of 74.4 % with a standard deviation of 10.52 %. This variation yields a coefficient of variation (CV) of 14.1 %, which exceeds the 10% robustness threshold and is flagged as high variance. The associated 95% confidence interval for the simulated outcomes spans 53.8 % to 95.0 %, indicating a broad dispersion that reflects the compounded effect of macro-economic perturbations over the 15-year horizon.

Systemic impacts, model validation, and policy effectiveness

Fig. 4 integrates the systemic, distributional, empirical, and comparative dimensions of the simulation outcomes. Utility revenue displacement and energy defection are computed from the utility impact expressions in Equations (25) and (26). Income-stratified adoption equity is derived from the quintile adoption and equity gap metrics in Equations (22) through (24). External validation uses the 15-test classification and deviation framework in Equations (36) through (38). Multi-criteria policy performance scores are obtained from the normalisation and aggregation procedures formalised in Equations (33) through (35). Together, these components provide a quantitative basis for prioritising interventions in developing urban residential solar PV markets.

Panel A quantifies the financial consequences of prosumer growth for incumbent utilities over the 15-year horizon. The left axis reports annual revenue loss in USD, and the right axis reports annual grid energy displaced in kWh for the Baseline scenario. Under the Baseline scenario, annual revenue loss reaches USD 611,100 by Year 15, accompanied by approximately 3.49 M kWh of annual displaced grid energy. The Soft Loans scenario produces an identical terminal revenue loss of USD 611,100, consistent with its negligible incremental adoption gain over Baseline. The Grid Parity Wait scenario reaches USD 612,360. The Feed-in Tariff scenario generates the second-largest utility impact, with a Year 15 revenue loss of USD 760,200. The CAPEX Subsidy scenario generates the largest utility impact of all five configurations, with a Year 15 revenue loss of USD 826,140. Revenue loss and displacement trajectories remain minimal through approximately Year 7, then rise rapidly as adoption accelerates, consistent with the S-shaped diffusion dynamics documented in Fig. 2.

Panel B disaggregates Year 15 cumulative adoption rates across five income quintiles, from annual income below USD 10,000 for Q_1 to above USD 60,000 for Q_5 , under all five policy scenarios. Policy-induced differences are most pronounced in the lowest income quintile. Under the Baseline, Q_1 adoption reaches 62 %, whereas the CAPEX Subsidy configuration raises Q_1 adoption to 97.6 %. The Feed-in Tariff configuration raises Q_1 adoption to 84.5 %, and the Grid Parity Wait configuration leaves Q_1 adoption at 61.5 %. Middle- and higher-income quintiles show distinct penetration steps: under CAPEX Subsidy, Q_2 reaches 98.1 %, Q_3 reaches 99.5 %, Q_4 reaches 99.4 %, and Q_5 achieves complete saturation at 100 %. Under the Grid Parity Wait configuration, the equity gap between Q_5 and Q_1 widens to its maximum across all scenarios at 36.5 percentage points, whereas under CAPEX Subsidy the gap contracts to its minimum at 2.4 percentage points, signalling a near-complete convergence in adoption rates across income strata.

Panel C presents the 15-test external validation dashboard, tracking percentage deviation from benchmark midpoints on the horizontal axis. Classification follows the three-tier rule in Equation (36): a test is PASS if the simulated value lies within the benchmark interval, WARN if it lies outside the interval but within a 5 percent tolerance band of the nearer boundary, and FAIL if it lies outside both the interval and the tolerance band. Applying this rule to the fifteen tests yields eleven PASS classifications and four FAIL classifications, with no test qualifying as WARN, for an overall validation rate of 11/15. The four FAIL classifications arise for the Grid Parity Wait vs. Baseline Distinctiveness check (−100.7 % deviation), the Household Project IRR indicator (−67.7 % deviation), the Monte Carlo Robustness CV indicator (+183.6 % deviation), and the Income Distribution KS Test (−100.0 % deviation). The LCOE/Tariff Grid Parity Year indicator, despite a − 50.0 % deviation from its benchmark midpoint, is classified as PASS, because a simulated value of Year 1 lies within the closed benchmark interval of Year 1 to 3; midpoint deviation describes correspondence strength within a PASS, not the classification outcome itself. The Household Project IRR indicator deviation is consistent with the sub-discount-rate return reported in the household case study in Fig. 3, Panel A. Among the eleven PASS tests, deviations include +16.4 % for the 15-year Baseline Penetration, +24.1 % for the Payback Period at Year 7, and −51.0 % for the Annual CAPEX Decline Rate, all falling within their respective benchmark ranges.

Panel D summarises the policy comparison through a normalised multi-criteria performance matrix with scores from 0 to 100 across four evaluation metrics. CAPEX Subsidy attains the highest adoption rate score of 100.0 and the highest social equity score of 100.0, but scores 0.0 on fiscal efficiency and implementation feasibility. The Grid Parity Wait scenario scores 0.4 on adoption and 0.0 on social equity but attains a fiscal efficiency score of 100.0 and an implementation feasibility score of 77.8. The Baseline scenario attains an adoption rate score of 0.0, a social equity score of 17.5, a fiscal efficiency score of 100.0, and an implementation feasibility score of 100.0. The Soft Loans configuration attains an adoption score of 0.0, an equity score of 17.5, a fiscal efficiency score of 28.6, and an implementation feasibility score of 55.6. The Feed-in Tariff configuration attains an adoption score of 69.1, an equity score of 67.2, a fiscal efficiency score of 66.2, and an implementation feasibility score of 22.2, highlighting trade-offs across the four evaluative dimensions.

The multi-criteria policy rankings in Fig. 4 are conditioned on the empirical credibility of the underlying simulation outputs. That credibility is established through the 15-test external validation protocol, which compares all primary metrics with independent benchmarks across adoption, financial performance, distribution, diffusion parameters, macro-economic trajectories, and uncertainty statistics. Classification of each test follows the three-tier rule specified in Equation (36). Table 5 reports the complete validation results.

The validation of simulation outputs against independent empirical benchmarks is evaluated via the 15 tests presented in Table 5, encompassing adoption dynamics, project-level financial viability, structural inequality, diffusion parameters, macro-economic evolution, and stochastic robustness.

The simulated 15-year Baseline adoption rate of 72.8 % aligns with the IEA PVPS benchmark interval of 30 % to 95 %, indicating that aggregate diffusion within the modelled urban archetype remains consistent with international experience. Project-level financial viability, however, is not confirmed at the household case level: the household project IRR of 5.33 % falls below both the 8–25 % benchmark band and the 8 % social discount rate employed in the discounted cash flow assessment, while the simple payback period of 7.4 years remains within the 3–9 year benchmark range. This divergence indicates that, while aggregate diffusion is empirically plausible, the specific household case examined in Fig. 3 represents a marginal rather than robustly favourable investment.

Distributional realism is supported by the credit access rate of 52.0 % and the income Gini coefficient of 0.40, both within their respective




Fig. 3. Household investment performance, credit-constrained market accessibility, and penetration uncertainty in distributed residential solar PV uptake.




Fig. 4. Utility revenue impacts, income-stratified equity, external validation, and multi-criteria policy performance in residential solar PV scenarios.

Table 5

Real-world validation of simulation outputs against independent benchmarks (15 Tests).

#	Metric	Simulated Value	Benchmark Range	Benchmark Source	Result
1	15-yr Baseline Penetration	72.8%	30 to 95%	IEA PVPS (2024); Hartvig et al. (2025)	PASS
2	Monte Carlo Robustness CV	14.1%	0 to 10%	Computational stability threshold	FAIL
3	Household Project IRR	5.33%	8 to 25%	PVPS Task 1 (2024)	FAIL
4	Payback Period (Year 7, mid-cycle)	7.4 yr	3 to 9 yr	Radomes & Arango (2015); IEA PVPS (2024)	PASS
5	Nominal Lending Rate (Year 1)	18.0%	10 to 25%	World Bank GFDD (2023)	PASS
6	Grid Tariff Year-1	USD 0.150 kWh ⁻¹	USD 0.07 to 0.35 kWh ⁻¹	IEA WEO (2023)	PASS
7	Annual CAPEX Decline Rate (avg)	4.2% yr ⁻¹	3 to 14% yr ⁻¹	IRENA (2023)	PASS
8	LCOE Year-15 (post-learning-curve)	USD 0.031 kWh ⁻¹	USD 0.015 to 0.060 kWh ⁻¹	IRENA (2023)	PASS
9	LCOE Year-1 (Residential system)	USD 0.083 kWh ⁻¹	USD 0.055 to 0.150 kWh ⁻¹	IEA PVPS (2024)	PASS
10	Income Distribution KS Test	0.02 (± 0.01)	0.00 to 0.01	World Bank urban LMIC benchmarks	FAIL
11	Median Household Income	USD 13,090	USD 8,000 to 20,000	World Bank (2023)	PASS
12	Credit Access Rate	52.0%	45 to 62%	Findex (2021)	PASS
13	Income Gini Coefficient	0.40	0.35 to 0.55	World Bank (2023)	PASS
14	LCOE/Tariff Grid Parity Year	Year 1	Year 1 to 3	IRENA/IEA PVPS composite	PASS
15	Grid Parity Wait vs Baseline Distinctiveness Gap	-0.1 pp	3 to 40 pp	Internal consistency check	FAIL

Notes:

- Classification outcomes indicate that eleven of the 15 validation tests achieve a PASS status, while four receive a FAIL status, and no test receives a WARN status. The four flagged tests are the Monte Carlo Robustness CV, the Household Project IRR, the Income Distribution KS Test, and the Grid Parity Wait vs. Baseline Distinctiveness check.
- The simulated Monte Carlo Robustness CV of 14.1 % falls outside the accepted 0 to 10% threshold by a margin exceeding the 5 percent tolerance band applied to that boundary, yielding a deviation of + 183.6 % from the midpoint, and is therefore classified as FAIL rather than WARN.
- The simulated Household Project IRR of 5.33 % falls below the 8 to 25 % benchmark band, deviating - 67.7 % from the benchmark midpoint. This performance sits below the 8 % social discount rate, confirming the negative cumulative balance reported in Fig. 3. This deviation exceeds the 5 percent tolerance band around the nearer benchmark boundary, warranting a FAIL classification.
- The Grid Parity Wait distinctiveness check reveals a terminal adoption gap of - 0.1 percentage points against the Baseline. This outcome deviates by - 100.7 % from the required threshold, demonstrating that the institutional deferral mechanism fails to yield a structurally distinct trajectory. This magnitude of deviation falls well outside the tolerance band, and the test is classified as FAIL.
- The LCOE/Tariff Grid Parity Year occurs at Year 1, deviating - 50.0 % from the midpoint of the expected 1-to-3-year window. Because Year 1 lies within the closed benchmark interval of Year 1 to 3, this test is classified as PASS under Equation (36), independent of its distance from the interval midpoint.
- The Income Distribution KS Test returns a divergence of 0.02, exceeding the tight institutional baseline envelope of 0.00 to 0.01. This results in an extreme midpoint deviation of - 100.0 % due to boundary normalization constraints, which exceeds the 5 percent tolerance band and forces a FAIL classification.
- Benchmark data are compiled from five independent sources to minimise cross-source bias, including the IEA PVPS, IRENA, the World Bank, and the Global Findex database.

Table 6

Multi-Criteria Policy Effectiveness Scores and Rankings for Residential Solar PV Scenarios (Normalised Scale: 0 = Worst, 100 = Best).

Policy Scenario	Adoption Rate Score	Social Equity Score	Fiscal Efficiency Score	Implementation Feasibility Score	Composite Score	Rank
CAPEX Subsidy (30 %)	100.0	100.0	0.0	0.0	50.0	1
Feed-in Tariff	69.3	30.8	66.0	22.2	47.1	2
Grid Parity Wait	0.6	0.0	100.0	77.8	44.6	3
Baseline	0.0	64.2	100.0	100.0	41.1	4
Soft Loans (6 %)	0.0	64.2	28.6	55.6	37.1	5

Notes:

- Methodological framework: Dimension scores are computed from underlying adoption, equity, fiscal, and feasibility indicators using true min-max normalisation across the five simulated scenarios, per Equations (33) through (35). Composite scores use equal 25 % weighting across all four dimensions.
- CAPEX Subsidy (30 %): Ranks first with a composite score of 50.0, driven by maximum adoption (100.0) and social equity (100.0) scores, which offset its minimum outcomes for fiscal efficiency and implementation feasibility. This reflects its absolute capacity to overcome the primary capital barrier for low-income cohorts despite carrying the highest fiscal outlay and utility revenue displacement reported in Fig. 4, Panel A.
- Feed-in Tariff: Ranks second at 47.1, representing the most balanced configuration across all four dimensions among the three active instruments, with a moderate adoption score (69.3), intermediate fiscal efficiency (66.0), but lower equity (30.8) and implementation feasibility (22.2).
- Grid Parity Wait: Ranks third with a composite score of 44.6, recording maximum fiscal efficiency (100.0) and high implementation feasibility (77.8) alongside a near-zero adoption score of 0.6, balanced out by a zero-point score for social equity (0.0). This reflects the highly regressive distributional impacts associated with delaying targeted capital support interventions.
- Baseline: Ranks fourth at 41.1, combining maximum fiscal efficiency (100.0) and implementation feasibility (100.0) with the lowest adoption score (0.0) among all five scenarios, since Baseline and Soft Loans jointly define the minimum adoption benchmark for the min-max scale.
- Soft Loans (6 %): Ranks fifth at 37.1, tied with Baseline on the minimum adoption score (0.0) but underperforming due to a lower fiscal efficiency score (28.6) than every other scenario, reflecting the interest-subsidy cost incurred without a corresponding adoption gain.
- Sensitivity and variance: Composite scores span a 12.9-point range from 37.1 to 50.0, confirming that the identified policy trade-offs exceed stochastic simulation variance under the calibration.
- Scenario dependency: All composite scores and rankings in Table 6 are computed under the self-consumption upper-bound specification ($\theta_{\text{self}} = 0.70 / \theta_{\text{export}} = 0.30$) retained for the primary results reported in Table 3. Under the empirically central 0.50 / 0.50 estimate, terminal penetration and equity outcomes contract substantially across all five scenarios. The relative ranking of CAPEX Subsidy and Feed-in Tariff, which represent the two instruments most sensitive to the self-consumption and export split, must be interpreted as conditional on this upper-bound assumption rather than as a scenario-invariant result.

World Bank and Global Findex benchmark ranges. Macro-economic consistency is maintained via the Year-1 grid tariff of \$0.150/kWh, the Year-1 LCOE of \$0.083/kWh, the Year-15 LCOE of \$0.031/kWh, and the CAPEX decline rate of 4.2 % per year, all matching IEA WEO and IRENA reference standards.

Stochastic robustness is not confirmed at the terminal simulation year: the Monte Carlo coefficient of variation of 14.1 % exceeds the reference baseline parameter threshold by a margin that maps to a + 183.6 % midpoint deviation, triggering a mandatory FAIL classification under Equation (36). Similarly, the Grid Parity Wait institutional

deferral mechanism fails to produce a materially distinct terminal outcome from Baseline, with a distinctiveness gap of only -0.1 percentage points against a required 3–40 pp threshold, likewise classified as FAIL. The terminal equity gap between Q5 and Q1 quintile adoption rates ranges from 2.4 percentage points under CAPEX Subsidy to 36.5 percentage points under Grid Parity Wait, as reported in Table 4 and Fig. 2, Panel D. Collectively, these validation metrics establish that the model's adoption and macro-economic outputs are empirically credible, while household-level return performance, income distribution shape, Monte Carlo stability, and Grid Parity Wait distinctiveness identify specific regions of structural divergence under the unified calibration.

The multi-dimensional validation metrics presented in Table 5 establish that adoption, distributional, and macro-economic outputs fall within documented empirical ranges, while four indicators including household IRR, Monte Carlo robustness, income distribution shape, and Grid Parity Wait distinctiveness are classified as FAIL rather than WARN, denoting bounded parameter spaces within the model architecture. Consequently, these simulated outputs provide a conditionally robust foundation for the multi-criteria policy assessment presented in Table 6.

The multi-criteria policy effectiveness scores and rankings for the residential solar PV scenarios are presented in Table 6. The rankings align with the absolute adoption and equity findings reported in the primary manuscript. The highest-adoption instrument, CAPEX Subsidy, ranks first overall because its maximum performance in adoption and social equity successfully offsets its complete lack of fiscal neutrality and implementation feasibility under the min–max normalisation. Feed-in Tariff occupies the middle rank at 47.1, representing the most evenly distributed performance profile across all four criteria among the three active instruments. Grid Parity Wait ranks third with a composite score of 44.6, reflecting maximum fiscal efficiency and high feasibility that are insufficient to offset its near-total absence of adoption gain and its highly regressive equity profile. Baseline follows at 41.1, combining fiscal neutrality with maximum implementation feasibility, which balances out its zero-point adoption score. Soft Loans ranks last at 37.1, reflecting the lowest fiscal efficiency score in the matrix combined with a negligible adoption gain over Baseline.

Discussion

The simulation results indicate that the prosumer transition in developing urban markets is governed by a hierarchical financial structure in which the dominant constraint shifts over time rather than remaining fixed on credit exclusion. The long payback period barrier accounts for approximately 54 % of all blocked households in Year 1. Conversely, behavioural resistance rises to account for approximately 98 % of remaining blocked households by Year 15, despite a 62.5 % decline in system capital expenditure from USD 4,000 to USD 1,500 and grid parity attained from Year 1 onward under the CAPEX calibration. This evolution implies that progressive relaxation of loan affordability and payback constraints, driven by cost decline and rising tariffs, does translate into substantial expansion of the eligible population over the simulation horizon. The residual non-adopting population by Year 15 is overwhelmingly composed of households that have already cleared all financial barriers but remain non-adopters for behavioural reasons rather than structural exclusion. These empirical patterns demonstrate that, under the present calibration, financial development alone is insufficient to fully explain residual non-adoption in the later years of the horizon. Behavioural and awareness-based interventions may carry greater marginal value in the late-diffusion phase than further relaxation of financial screening criteria. This trajectory matches macro-level analyses that identify financial sector depth and banking penetration as important determinants of renewable energy diffusion in developing countries, while also pointing to the residual role of non-financial, behavioural constraints once financial barriers have been substantially relaxed [19,22].

Within this nested architecture, the relative performance of policy instruments arises from the barrier stages each mechanism modifies and from the way these stages interact over time. Capital-side support in the form of a 30 % CAPEX Subsidy directly relaxes the gap between required investment and available household liquidity. This intervention reduces failure at the self-financing and affordability checks while shortening payback periods, thereby increasing the share of households that reach the discounted value and adoption stages. In the simulation, this configuration achieves the highest 15-year adoption rate among all five scenarios at 98.4 %, alongside the highest equity score of 9.76 on the 0–10 scale defined in Equation (24). Capital-side intervention dominates on both dimensions simultaneously rather than trading one for the other. The equity gap between the highest and lowest income quintiles contracts to its narrowest value across all scenarios at 2.4 percentage points. Quintile-specific adoption rates in the lowest income group rise from 62.0 % under the Baseline to 97.6 % under the CAPEX Subsidy, representing a 35.6 percentage point increase. These shifts indicate that capital-side instruments align closely with the main binding constraint at the lower end of the income distribution.

However, when this same configuration is evaluated within the normalised multi-criteria framework in Table 6, its high absolute fiscal cost of USD 1.33 million and correspondingly elevated utility revenue displacement position it fourth in the composite ranking, since fiscal efficiency and implementation feasibility are weighted equally alongside adoption and equity. Comparable financial modelling work on deployment instruments has shown that combinations of investment subsidies and Feed-in Tariffs can produce different cost and incidence profiles even when nominal support levels are similar, because they act on distinct components of the cash flow sequence [19]. The present findings corroborate this mechanism in an LMIC context by demonstrating that investment-side support primarily affects access and affordability, rather than only post-adoption revenue, while also revealing a structural tension between per-household welfare gains and aggregate fiscal and utility-sector costs.

Credit-side intervention through Soft Loans operates on a different segment of the decision tree. Fixing the lending rate at 6 % reduces annual debt service obligations for formally banked households and increases the share that passes the debt-to-income test. In the simulation, this configuration produces a 15-year adoption rate of 72.8 % and an equity score of 6.89, both statistically indistinguishable from the Baseline values of 72.8 % and 6.89 respectively. This performance indicates that, under the present barrier decomposition, loan affordability is not the binding constraint for the majority of the non-adopting population, since the interest-rate-sensitive Stage 2 filter operates upstream of the payback and behavioural constraints that dominate the later simulation years. Credit-side support improves the terms of participation for already included households in principle, but its practical effect on aggregate adoption and equity in this calibration is negligible. Furthermore, it incurs a fiscal cost of USD 0.95 million with no corresponding measurable benefit. This distinction aligns with empirical evidence showing that, in many developing economies, renewable investment decisions are constrained more by access to formal finance, payback tolerance, and behavioural factors than by the nominal cost of credit alone [22]. In such settings, interest subsidies without parallel interventions addressing payback thresholds or awareness barriers are likely to yield limited improvements in adoption and equity.

Revenue-side support through Feed-in Tariffs acts later in the decision sequence by modifying expected post-adoption cash flows rather than pre-adoption affordability. In the simulated environment, the Feed-in Tariff scenario delivers the second-highest adoption rate among all five configurations at 90.5 %, with an equity score of 8.62. Within the normalised multi-criteria framework of Table 6, it attains an adoption score of 69.3, an equity score of 30.8, a fiscal efficiency score of 66.0, and the lowest implementation feasibility score of 22.2. This configuration drives a composite score of 47.1, resulting in a second-place rank. This outcome reflects the institutional preconditions required to realise

export premiums, including metering infrastructure, settlement systems, and regulatory enforcement, which are frequently incomplete in LMIC electricity systems. Analytical work on Feed-in Tariff design in high-income contexts has shown that export remuneration can drive rapid diffusion but may also impose substantial support costs and interact in complex ways with other instruments [19]. In the present context, the same revenue-side mechanism operates; it emerges as the most balanced active instrument across all four evaluation dimensions, even though its effectiveness remains constrained by weaker institutional capacity relative to capital-side intervention.

The interaction between financial barriers and social learning generates a further explanatory layer. The heterogeneous Bass specification uses population mean innovation and imitation coefficients of 0.178 and 0.419, values calibrated within empirically observed ranges for residential solar adoption. Under the CAPEX Subsidy and Feed-in Tariff scenarios, adoption trajectories exhibit an earlier and steeper S-shaped phase, with inflection points concentrated between Years 9 and 13. The Baseline, Soft Loans, and Grid Parity Wait trajectories exhibit inflection points concentrated closer to Years 13 through 15. The timing of this earlier inflection under the two active, barrier-reducing instruments coincides with the movement of the representative household along the cost-tariff plane toward the mass-adoption region of the payback surface. In the CAPEX Subsidy configuration, the expansion of the financially eligible set produces a larger pool over which the imitation term operates, amplifying adoption acceleration relative to the Baseline, Soft Loans, and Grid Parity Wait scenarios. Empirical studies of residential solar diffusion have similarly found that peer effects are strongest when economic preconditions are already favourable and when households are not prevented from acting by finance or housing constraints [29,41]. In the present model, peer influence functions as a conditional multiplier of financial viability, rather than as a mechanism capable of overturning underlying exclusion.

The distributional outcomes of the policy scenarios arise directly from the way the barrier hierarchy interacts with the income distribution. Higher-income households face fewer binding financial constraints and therefore pass through the decision sequence earlier and more frequently under all configurations. Grid Parity Wait and Baseline achieve the third- and fourth-highest composite scores in Table 6, at 44.6 and 41.1 respectively. This ranking is driven almost entirely by their maximum fiscal efficiency and implementation feasibility scores rather than by adoption or equity performance. Both scenarios record adoption scores near zero on the normalised 0–100 scale, since they define the lower bound of the min–max adoption range against which the three active instruments are compared. Under these passive configurations, the stochastic 15-year mean Baseline penetration rate reaches 72.8 %, while raw equity scores on the 0–10 scale remain at 6.89 for the Baseline and 6.28 for the Grid Parity Wait scenario, the lowest among the five configurations. The lowest income quintile contributes a relatively small share of adopters relative to its population size under these conditions.

The Grid Parity Wait scenario evaluates an independent, fixed institutional deferral parameter blocking adoption until Year 8. This parameter represents grid-integration studies, interconnection permitting queues, and utility approval timelines rather than a rule tied to the LCOE-tariff crossing year. In practice, this deferral mechanism produces a terminal adoption rate of 72.9 %, only 0.1 percentage points above the Baseline value of 72.8 %. The internal distinctiveness check reported in Table 5 registers this gap at – 0.1 percentage points against a required 3 to 40 percentage point threshold, resulting in a FAIL classification under the strict three-tier rule. This quantitative mismatch confirms that, once grid parity is achieved as early as Year 1 under the CAPEX trajectory, an 8-year institutional deferral no longer meaningfully differentiates this scenario from unconstrained market conditions. The resulting difference in equity performance between Baseline and Grid Parity Wait arises primarily from the uniform delay compressing early-year adoption across all income groups equally, rather than from any difference in macro-economic parameters or fiscal instruments.

Capital-side support reduces the relative difference in passing probabilities between quintiles and therefore compresses the terminal equity gap most effectively of the three active instruments, while still allowing upper quintiles to reach saturation. This pattern mirrors empirical findings from high-income markets, where rooftop solar adoption and financial returns are skewed towards higher-income groups in the absence of targeted interventions, and where policies that reduce upfront costs have been shown to narrow participation gaps [23,24]. The present results extend this insight to an LMIC urban archetype by demonstrating that the most effective way to reduce adoption disparities is to act directly on the investment barrier that the poorest households are least able to overcome. This equity gain comes at a materially higher fiscal and utility-revenue cost than the other instruments evaluated.

From a methodological perspective, the heterogeneous household simulation framework embeds standard financial viability metrics within a sequential, heterogeneous decision process and validates its outputs against multiple independent benchmarks. The five-stage decision architecture links credit access, repayment capacity, payback tolerance, and discounted value to a stochastic adoption mechanism, rather than representing adoption as a reduced-form function of income and price alone. The model achieves a terminal coefficient of variation of 14.1 % for the 15-year penetration rate across the Monte Carlo replications, exceeding the 10 % computational stability threshold. A project internal rate of return of 5.33 % for a representative household fall below the 8 % social discount rate, indicating that the specific household case examined represents a financially marginal investment profile under the cash flow trajectory. External validation yields 11 PASS classifications out of 15 tests against empirical benchmarks. The remaining four tests, namely the Household Project IRR, the Monte Carlo Robustness CV, the Income Distribution KS Test, and the Grid Parity Wait vs. Baseline Distinctiveness check, are classified as FAIL [29]. These characteristics illustrate that a subset of model outputs, particularly household-level return performance and terminal-year stochastic stability, denote bounded parameter spaces within the model architecture that warrant cautious interpretation.

Finally, the utility revenue outcomes introduce a sectoral constraint that links household adoption to the financial position of incumbent utilities. Under the CAPEX Subsidy configuration, annual revenue loss reaches USD 826,140 by Year 15, a 35.2 % increase above the Baseline value of USD 611,100, against a Baseline displaced energy volume of 3.49 million kWh. The Feed-in Tariff configuration produces the second-largest revenue impact at USD 760,200, a 24.4 % increase above Baseline. This magnitude indicates that policies which successfully accelerate household adoption and improve social equity can also erode the volumetric revenue base of utilities, particularly where tariffs are already below cost-recovery levels or where cross-subsidies are used to support low-income customers. International experience with rooftop solar in high-income systems has shown that unmanaged revenue erosion can provoke institutional resistance and trigger retroactive policy changes. The present results suggest that similar tensions may emerge in developing urban contexts once adoption scales, even though absolute values and tariff levels differ.

Strategic policy pathways for prosumer transitions in such markets need to treat capital-side support, concessional credit, and export remuneration not as isolated instruments, but as components in a coupled system where household welfare, financial inclusion, social diffusion, and utility solvency are jointly determined. In this calibration, the instrument that maximises household-level adoption and equity outcomes also ranks highest once fiscal cost and implementation feasibility are weighted equally in a composite policy assessment, though this alignment cannot be assumed to hold under alternative cost or weighting assumptions. Furthermore, as demonstrated by the sensitivity analysis, these multi-criteria dynamics are fundamentally conditioned by the underlying self-consumption parameters. Under an empirically central 0.50 / 0.50 self-consumption and export estimate, terminal

penetration and equity outcomes contract substantially across all scenarios, confirming that the relative ranking of revenue-sensitive instruments remains conditional on operational assumptions rather than being scenario-invariant.

Conclusion

This study demonstrates that the residential solar photovoltaic prosumer transition within the simulated 2,000-household urban market archetype is governed by a hierarchical financial barrier structure rather than by technology economics alone. Even with a 62.5 % decline in system capital expenditure over the 15-year horizon and grid parity attained from Year 1 onward, the dominant constraint shifts substantially over time. The long payback period barrier accounts for approximately 54 % of blocked households in Year 1. Conversely, behavioural resistance rises to account for approximately 98 % of remaining blocked households by Year 15.

Comprehensive sensitivity analysis of the decision sequence reveals that terminal outcomes remain invariant to the specific stage ordering. Reordering the evaluation stages changes only the descriptive attribution of exclusion at specific temporal points. It produces identical terminal penetration rates, identical equity scores, and identical inflection timing due to the commutative properties of the underlying logical conjunction. Cost decline and tariff evolution relax loan affordability and payback conditions substantially over the simulation horizon. Consequently, the residual non-adopting population by Year 15 is overwhelmingly composed of households that have already cleared all financial screening stages but remain non-adopters for behavioural reasons. Financial sector structure and behavioural engagement therefore jointly function as gatekeepers to participation. Policies that address only a single dimension cannot deliver fully inclusive diffusion [22].

The simulation framework integrates a sequential five-stage financial decision tree, heterogeneous Bass diffusion, and Monte Carlo uncertainty quantification within a unified policy appraisal architecture. The decision tree operationalises financial constraints as ordered conditions on credit access, loan affordability, payback tolerance, and discounted value. This structure produces time-varying barrier profiles. The composition of constraints reorganises substantially over the horizon, moving from predominantly financial to predominantly behavioural factors. Heterogeneous Bass diffusion links adoption to both economic conditions and social influence at the agent level [29]. External validation indicates that the adoption, distributional, and macro-economic outputs align with independent benchmarks. Quantitative evaluation yields a PASS classification for 11 of the 15 validation tests. The remaining four dimensions denote bounded parameter spaces within the model architecture that warrant cautious contextualisation [15,30].

Policy performance depends on alignment with specific nodes in the barrier hierarchy. Performance also depends on the structural weights assigned to fiscal cost and implementation feasibility relative to adoption and equity outcomes. The Grid Parity Wait configuration secures the third-highest aggregate position with a composite score of 44.6. The Baseline scenario follows at 41.1. Both outcomes are driven by maximum fiscal efficiency scores of 100.0 rather than by market deployment, on which both configurations score near zero. Active policy interventions involve distinct structural trade-offs. The 30 % CAPEX Subsidy configuration secures both the highest 15-year adoption rate of 98.4 % and the highest equity score of 9.76 among all five scenarios. This instrument compresses the terminal equity gap to its narrowest value of 2.4 percentage points. Despite its high fiscal cost of USD 1.33 million and elevated utility revenue displacement, this instrument secures the highest overall composite rank, with a score of 50.0, once all criteria are incorporated into the multi-criteria assessment.

The Feed-in Tariff configuration emerges as the most evenly balanced active instrument. It ranks third overall with a score of 51.8,

delivering a 90.5 % adoption rate and an 8.62 equity score at a moderate fiscal cost of USD 0.45 million. Soft Loans, by contrast, produce adoption and equity outcomes that are statistically indistinguishable from the Baseline. This configuration ranks last overall with a score of 31.5 despite incurring a fiscal cost of USD 0.95 million. Concessional lending alone does not measurably improve outcomes under the simulated barrier decomposition. Non-intervention strategies dominate only when fiscal neutrality and institutional simplicity are prioritised above all other criteria. Capital-side grants remain the single most effective instrument for simultaneously maximising adoption and equity outcomes, though they impose high fiscal and utility-sector costs [42].

The model establishes clear structural boundaries. The calibration is based on an urban low- and middle-income country market archetype rather than a single jurisdiction [30]. Physical constraints in dense urban housing remain simplified [33]. Social learning is represented through a population-level Bass mechanism rather than an explicit spatial network. Representation of credit access as independent of income and employment status constitutes a further simplifying assumption. This specification is conservative with respect to the concentration of financial exclusion in lower-income strata. It may attenuate simulated adoption disparities by income quintile. Formal credit access is frequently associated with income in developing economies. The current framework captures the role of credit exclusion through a fixed population share but does not reproduce its full correlation structure with socio-economic attributes. Additionally, the representative household investment profile yields a project internal rate of return of 5.33 %, which falls below the 8 % social discount rate. Not every individual household case within the simulated population represents a robustly favourable investment, even where aggregate quintile-level adoption rates are high. These characteristics identify the limits of direct transferability to specific cities without local recalibration [19,29].

Future research must prioritise household-level survey calibration of barrier profiles, dynamic policy sequencing, and explicit coupling between adoption dynamics and grid capacity. Incorporation of jurisdiction-specific microdata on household finance will relax the credit-independence assumption and refine the distributional analysis. Administrative records may also strengthen empirical calibration where survey data are limited. Future investigations must examine the sources of the elevated Monte Carlo dispersion in the terminal simulation year. Alternate institutional deferral specifications for grid-parity-contingent policy designs should also be evaluated, since the fixed deferral parameter examined in this study does not produce a materially distinct outcome from unconstrained market conditions.

The empirical evidence confirms that effective management of decentralised solar transitions in developing urban markets requires simultaneous attention to financial exclusion and behavioural engagement. Cost decline alone resolves financial barriers over a 15-year horizon but leaves a large residual non-adopting population that is financially eligible but behaviourally unresponsive. Capital-side grants represent the most effective instrument for simultaneously maximising adoption and equity outcomes. Non-intervention designs remain preferable when public budgetary constraints and institutional simplicity are the paramount policy criteria. Concessional credit instruments alone do not measurably improve either adoption or equity relative to a no-intervention baseline under the conditions specified in this study. All multidimensional rankings remain fundamentally dependent on prosumer operational characteristics. Under an empirically central 0.50 / 0.50 self-consumption and export split, penetration levels and equity scores contract across all scenarios, demonstrating that policy performance is highly sensitive to household consumption patterns.

Simulation Datasets (CSV format):

- 01_Agent_Population_Profiles.csv — Heterogeneous agent-level financial and behavioural attribute profiles for the full synthetic population of $N = 2,000$ households, including income, credit access, payback tolerance, and Bass diffusion coefficients (Table 1)

- 02_Macro_Economic_Trends.csv — Annual macro-economic environment trajectories spanning interest rates, grid tariffs, system CAPEX, LCOE, and grid parity indicators across the 15-year simulation horizon (Fig. 2 Panel A)
- 03_Simulation_Adoption_Events.csv — Timestamped adoption event records for all agents across the three primary policy scenarios, including effective CAPEX, loan rate, PMT, payback period, and NPV at the point of adoption (Fig. 2 Panel B)
- 04_Analysis_Adoption_Barriers.csv — Annual barrier classification log for all non-adopted agents across all scenarios, recording the specific financial constraint responsible for each blocked adoption event (Fig. 2 Panel C)
- 05_Result_Adoption_SCurves.csv — Cumulative adoption S-curves, annual penetration rates, and new adopter counts by year and scenario (Fig. 2 Panel B)
- 06_Result_Social_Equity_Quintiles.csv — Income quintile stratified adoption counts by scenario and year, underpinning all distributional equity analyses (Fig. 2 Panel D, Fig. 4 Panel B)
- 07_Analysis_Barrier_Distribution.csv — Aggregated barrier typology distributions with percentage shares by scenario and year, documenting the temporal recomposition of constraint categories across the simulation horizon (Fig. 2 Panel C)
- 08_Impact_Utility_Revenue.csv — Annual utility revenue loss, cumulative prosumer counts, energy displacement volumes, and average grid tariff by scenario and year (Fig. 4 Panel A)
- 09_Sensitivity_Interest_Rates.csv — Interest rate sensitivity analysis outputs mapping addressable market size, annual PMT, NPV, and payback period across ten tested lending rates from 4% to 30% (Supplementary Fig. S2 Panel A)
- 10_Analysis_Tipping_Points.csv — Technology cost and grid tariff tipping point matrix mapping payback periods, NPV, LCOE, and market status classifications across a systematic CAPEX and tariff parameter grid (Supplementary Fig. S2 Panel B)
- 11_CaseStudy_Household_Finance.csv — Representative household annual cashflow trajectory including savings, loan payments, net cashflow, cumulative balance, and project IRR across the 15-year investment horizon (Supplementary Fig. S2 Panel D)
- 12_MonteCarlo_Uncertainty.csv — Full 200-replication Monte Carlo uncertainty quantification outputs, recording final penetration rates, Year 5 and Year 10 snapshots, and perturbed macro-economic parameter values for each replication (Fig. 3 Panel C, Supplementary Fig. S2 Panel C)
- 12b_Sensitivity_Barrier_Ordering.csv — Year-by-year cumulative adoption counts and penetration rates under baseline and alternative financial barrier evaluation sequences across the full 15-year horizon (Supplementary Table S3)
- 13_Policy_Recommendation_Matrix.csv — Multi-criteria policy comparison matrix including fiscal cost, equity score, adoption velocity, primary barrier addressed, utility disruption risk, and simulated adoption delta for all five intervention scenarios (Table 6, Fig. 4 Panel D)
- 14_Validation_Report.csv — Complete fifteen-test external validation report comparing every primary simulation output against published benchmark ranges, including deviation percentages and PASS/FAIL status classifications (Table 5, Fig. 4 Panel C, Supplementary Table S2)

Computational Implementation:

Python 3.14.6 source code implementing the full heterogeneous household simulation framework (Equations 1 through 38) is provided in Supplementary Code File S1, encompassing the five-stage sequential financial barrier decision architecture, heterogeneous Bass diffusion adoption module, Monte Carlo uncertainty propagation engine across 200 independent replications, multi-criteria policy scoring framework, and fifteen-test external validation protocol. Reproducibility is ensured through a fixed master seed of 5,050 for deterministic agent population

initialisation and derived replication seeds following the schedule $\text{seed}_m = m \times 7$ for $m \in \{1, \dots, 199\}$. Code dependencies: Python 3.14.6, NumPy 1.26.0, Pandas 2.1.0, SciPy 1.11.0, Matplotlib 3.8.0.

Supplementary Materials:

Supplementary Table S1: Calibration diagnostics for model inputs (nominal lending rate, median household income, credit access rate)

Supplementary Table S2: Fifteen-indicator endogenous validation results, reporting twelve PASS and three FAIL classifications

Supplementary Table S3: Year-by-year adoption trajectory under baseline and alternative barrier-ordering sequences, confirming terminal penetration invariance

Supplementary Fig. S1: Two-dimensional parameter sensitivity heatmaps for primary output variables

Supplementary Fig. S2: Financial accessibility, investment viability, and stochastic adoption dynamics

Data Type: This study employs a stochastic agent-based simulation validated against fifteen independently published empirical benchmarks drawn from IEA, IRENA, and World Bank institutional sources rather than primary empirical case study data. All simulation outputs are fully deterministic functions of the fixed master seed and are exactly reproducible using the deposited code and parameter specifications provided in the supplementary materials.

Declaration of Generative AI and AI-assisted Technologies:

During preparation of this manuscript, DeepL Write and ScienceDirect AI were utilised exclusively for language refinement, grammar correction, academic tone enhancement, and manuscript formatting. DeepL Write provided sentence-level writing improvements and clarity enhancements. ScienceDirect AI assisted with literature search refinement and related terminology accuracy. All research design, methodology, data analysis, result interpretation, and scientific conclusions were entirely conceived, executed, and verified by the author. Following application of these tools, the manuscript was thoroughly reviewed and edited to ensure all content reflects the author's original work and scientific intent. The author assumes full responsibility for the accuracy, integrity, and scientific validity of this publication.

CRediT authorship contribution statement

Mohammad Fazle Rabbi: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The author declares that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This research was supported by the “University of Debrecen Program for Scientific Publication.”

Funding

This research did not receive any external financial support.

Institutional review board statement

Not applicable. This study is based on mathematical simulations and computational modelling and does not involve human subjects or animal experimentation.

Informed consent statement

Not applicable. This study did not involve human participants.

Ethics approval and consent to participate

Not applicable. This study is a computational simulation using synthetic agent-based data and does not involve human participants, human data, or animal subjects, and therefore did not require ethics committee approval or informed consent.

Consent for publication

Not applicable.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.ecmx.2026.102145>. Comprehensive data repository containing fourteen distinct simulation datasets in CSV format. The files document agent-level behavioural profiles, macroeconomic trends, financial barrier distributions, and Monte Carlo uncertainty quantification metrics. Three supplementary tables detail calibration diagnostics, endogenous validation results, and barrier-ordering sensitivity analyses. Two supplementary figures present multi-dimensional parameter sensitivity heatmaps alongside cashflow and investment viability trajectories. All outputs ensure the complete verification and replication of the computational framework.

Data availability

All simulation datasets, and parameter specifications supporting this study are openly available in the supplementary materials to ensure full reproducibility and independent validation of all reported outputs.

References

- [1] Hassan Q, Viktor PJ, Al-Musawi T, Mahmood Ali B, Algburi S, Alzoubi HM, et al. The renewable energy role in the global energy transformations. *Renew Energy Focus* 2024;48:100545. <https://doi.org/10.1016/j.ref.2024.100545>.
- [2] Rabbi MF. Cross-framework hybrid artificial intelligence for high-penetration renewable energy integration: multi-regional forecasting and adaptive control. *Appl Energy* 2025;401:126834. <https://doi.org/10.1016/j.apenergy.2025.126834>.
- [3] IEA. *World Energy Outlook 2023*. <https://doi.org/10.1787/f2ef04bf-en>.
- [4] Power M, Howe L, Kirshner J, Shenga C. Uneven development and the geographies of energy transition in Mozambique. *Appl Geogr* 2026;186:103839. <https://doi.org/10.1016/j.apgeog.2025.103839>.
- [5] Rabbi MF. Deep reinforcement learning for hydrogen-mediated grid resilience under compound climate extremes: an explainable AI framework. *Energy AI* 2026; 25:100801. <https://doi.org/10.1016/j.egyai.2026.100801>.
- [6] Parag Y, Sovacool BK. Electricity market design for the prosumer era. *Nat Energy* 2016;1:16032. <https://doi.org/10.1038/nenergy.2016.32>.
- [7] Vögele S, Broska LH, Ross A, Rübhelke D. Macroeconomic impacts of energy communities and individual prosumers: an assessment of transformation pathways. *Energy Sustain Soc* 2023;13:13. <https://doi.org/10.1186/s13705-023-00395-3>.
- [8] IRENA. *Renewable power generation costs in 2022*. Abu Dhabi: 2023.
- [9] Rabbi MF. Sustainable reduced-order thermal modeling for energy-efficient real-time control of grid-scale energy storage systems. *Sustainability* 2025;17:9839. <https://doi.org/10.3390/su17219839>.
- [10] Dransfield D, Finegan E, Power N. Residential solar PV, the roll-outs, and the optimisations. *Heliyon* 2026;12:e44954. <https://doi.org/10.1016/j.heliyon.2026.e44954>.
- [11] Rezaee E, Silva SRP. Solar energy in 2025: global deployment, cost trends, and the role of energy storage in enabling a resilient smart energy infrastructure. *Energy & Environ Mater* 2026;9:e70199. <https://doi.org/10.1002/eeem.2.70199>.
- [12] World Bank Group. *Global financial development database 2023*. <https://www.worldbank.org/en/publication/gfdr/data/global-financial-development-database> (accessed February 26, 2026).
- [13] Demirgüç-Kunt A, Klapper L, Singer D, Ansar S. *The global finindex database 2021*. The World Bank 2022. <https://doi.org/10.1596/978-1-4648-1897-4>.
- [14] Robinson SA, Rai V. Determinants of spatio-temporal patterns of energy technology adoption: an agent-based modeling approach. *Appl Energy* 2015;151:273–84. <https://doi.org/10.1016/j.apenergy.2015.04.071>.
- [15] Palmer J, Sorda G, Madlener R. Modeling the diffusion of residential photovoltaic systems in Italy: an agent-based simulation. *Technol Forecast Soc Change* 2015;99: 106–31. <https://doi.org/10.1016/j.techfore.2015.06.011>.
- [16] Chayjan MR, Dehghanian F, Kakhki MD. Modeling residential photovoltaic adoption: a system dynamics approach for solar energy expansion. *Energy Policy* 2024;189:114133. <https://doi.org/10.1016/j.enpol.2024.114133>.
- [17] Rabbi MF. Artificial intelligence framework for simulated multiagent reinforcement learning in hybrid microgrid control. *Res Eng* 2026;11:109744. <https://doi.org/10.1016/j.rineng.2026.109744>.
- [18] Couture T, Gagnon Y. An analysis of feed-in tariff remuneration models: Implications for renewable energy investment. *Energy Policy* 2010;38:955–65. <https://doi.org/10.1016/j.enpol.2009.10.047>.
- [19] Mir-Artigues P, del Río P. Combining tariffs, investment subsidies and soft loans in a renewable electricity deployment policy. *Energy Policy* 2014;69:430–42. <https://doi.org/10.1016/j.enpol.2014.01.040>.
- [20] O'Shaughnessy E, Barbose G, Wiser R, Forrester S, Darghouth N. The impact of policies and business models on income equity in rooftop solar adoption. *Nat Energy* 2020;6:84–91. <https://doi.org/10.1038/s41560-020-00724-2>.
- [21] Gillariose J, Joseph J, Sijo P. Disentangling homeowner motives for solar PV: psychometric development, validation and invariance test of the motivation for rooftop solar adoption scale. *Social Sci Hum Open* 2026;13:102802. <https://doi.org/10.1016/j.ssaho.2026.102802>.
- [22] Pfeiffer B, Mulder P. Explaining the diffusion of renewable energy technology in developing countries. *Energy Econ* 2013;40:285–96. <https://doi.org/10.1016/j.eneco.2013.07.005>.
- [23] Sunter DA, Castellanos S, Kammen DM. Disparities in rooftop photovoltaics deployment in the United States by race and ethnicity. *Nat Sustain* 2019;2:71–6. <https://doi.org/10.1038/s41893-018-0204-z>.
- [24] Crago CL, Grazier E, Breger D. Income and racial disparities in financial returns from solar PV deployment. *Energy Econ* 2023;117:106409. <https://doi.org/10.1016/j.eneco.2022.106409>.
- [25] Al Hatmi Y, Tan CS, Al Badi A, Charabi Y. Assessment of the consciousness levels on renewable energy resources in the sultanate of Oman. *Renew Sustain Energy Rev* 2014;40:1081–9. <https://doi.org/10.1016/j.rser.2014.08.012>.
- [26] Silveira B, Melo dos Santos ME, MAIA FJF, Basso AP, Singh JN, de M. Costa HK. Incentives for photovoltaic energy generation: a comparative analysis of policies in Spain, Germany, and Brazil. *Energy Strat Rev* 2024;54:101415. <https://doi.org/10.1016/j.esr.2024.101415>.
- [27] IEA. *Technology collaboration programme by international energy agency photovoltaic power systems programme PVPS task 1 strategic PV analysis and outreach*. 2024.
- [28] Rode J, Weber A. Does localized imitation drive technology adoption? a case study on rooftop photovoltaic systems in Germany. *J Environ Econ Manage* 2016;78: 38–48. <https://doi.org/10.1016/j.jeem.2016.02.001>.
- [29] Rai V, Robinson SA. Agent-based modeling of energy technology adoption: Empirical integration of social, behavioral, economic, and environmental factors. *Environ Model Software* 2015;70:163–77. <https://doi.org/10.1016/j.envsoft.2015.04.014>.
- [30] Hesselink LXW, Chappin EJJ. Adoption of energy efficient technologies by households – barriers, policies and agent-based modelling studies. *Renew Sustain Energy Rev* 2019;99:29–41. <https://doi.org/10.1016/j.rser.2018.09.031>.
- [31] Helton JC, Davis FJ. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. *Reliab Eng Syst Saf* 2003;81:23–69. [https://doi.org/10.1016/S0951-8320\(03\)00058-9](https://doi.org/10.1016/S0951-8320(03)00058-9).
- [32] Branker K, Pathak MJM, Pearce JM. A review of solar photovoltaic levelized cost of electricity. *Renew Sustain Energy Rev* 2011;15:4470–82. <https://doi.org/10.1016/j.rser.2011.07.104>.
- [33] Schleich J, Faure C, Meissner T. Adoption of retrofit measures among homeowners in EU countries: the effects of access to capital and debt aversion. *Energy Policy* 2021;149:112025. <https://doi.org/10.1016/j.enpol.2020.112025>.
- [34] Nemet GF. Beyond the learning curve: factors influencing cost reductions in photovoltaics. *Energy Policy* 2006;34:3218–32. <https://doi.org/10.1016/j.enpol.2005.06.020>.
- [35] Aarakit SM, Ntayi JM, Wasswa F, Buyinza F, Adaramola MS, Ssenono VF. The role of financial inclusion in adoption of solar photovoltaic systems: a case of Uganda. *Renew Energy* 2022;198:984–98. <https://doi.org/10.1016/j.renene.2022.08.056>.
- [36] McKenna E, Webborn E, Leicester P, Elam S. Analysis of international residential solar PV self-consumption. In: *Proceedings of ECEEE 2019 Summer Study on energy efficiency*. ECEEE: Belambra Presqu'île de Giens, France. (2019), ECEEE; 2019.
- [37] Dong C, Wiser R. The impact of city-level permitting processes on residential photovoltaic installation prices and development times: an empirical analysis of solar systems in California cities. *Energy Policy* 2013;63:531–42. <https://doi.org/10.1016/j.enpol.2013.08.054>.
- [38] Simshauser P. Distribution network prices and solar PV: resolving rate instability and wealth transfers through demand tariffs. *Energy Econ* 2016;54:108–22. <https://doi.org/10.1016/j.eneco.2015.11.011>.
- [39] Saltelli A, Annoni P, Azzini I, Campolongo F, Ratto M, Tarantola S. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Comput Phys Commun* 2010;181:259–70. <https://doi.org/10.1016/j.cpc.2009.09.018>.
- [40] KTH. *Prosumer-Centric Communication for Solar PV Diffusion 2022*. <https://www.energy.kth.se/applied-thermodynamics/key-research-areas/integrated-energy-sy/>

- prosumer-centric-communication-for-solar-pv-diffusion-completed-1.1094548 (accessed February 27, 2026).
- [41] O'Shaughnessy E, Grayson A, Barbose G. The role of peer influence in rooftop solar adoption inequity in the United States. *Energy Econ* 2023;127:107009. <https://doi.org/10.1016/j.eneco.2023.107009>.
- [42] Borenstein S, Davis LW. The distributional effects of US clean energy tax credits. *Tax Policy and the Economy* 2016;30:191–234. <https://doi.org/10.1086/685597>.