Advances in Consumer Research
ISSN (Print): 0098-9258
ISSN(Online): 3079-1766

Vol. 2, Issue 5 (2025) https://acr-journal.com/

Which Assumptions Really Set Power Purchase Prices And Returns In United States Solar Projects.

Kshitiz Raj[1*], Michael Walton[2]

[1*]Energy Finance Specialist, Redaptive Sustainability Services, Denver, Colorado, United States,

Email ID : kshitiz.raj@alumni.duke.edu, 1kshitizraj7@gmail.com

[2]Partner, Precursor, Chattanooga, Tennessee, United States.

[2]Former Member, White House Environmental Justice Advisory Council (2024–2025).

Cite this paper as: Kshitiz Raj, Michael Walton, (2025) Which Assumptions Really Set Power Purchase Prices And Returns In United States Solar Projects. *Advances in Consumer Research*, 2 (5), 1960-1978

KEYWORDS

U.S. solar project finance, DSCR-constrained cash-flow modeling, Sensitivity ranking of project inputs, PPA price formation, NPV margin and levered IRR.

.

ABSTRACT

Three results determine the financeability of grid-connected solar projects: equity value, equity return, and tariff which a project can sign. This paper has determined the most significant input assumptions that contribute to net present value (NPV) margin, levered internal rate of return (LIRR), and the price that is necessary to sign a power purchase agreement (PPA). They are a project-finance cashflow model based on debt service coverage ratio (DSCR)-constrained amortizing debt, U.S. tax treatment to include the investment tax credit (ITC), interest capitalized on construction and a fixed-price PPA with smooth escalation. The analysis begins with a documented base case, varies a set of negotiated commonly negotiated inputs under systematically changing conditions price-solve and return-solve, capital expenditure per watt (including interconnection), bankable production (kilowatt-hours per kilowatt), base-year price and escalator, operations and maintenance cost, construction duration, debt advance rate and coupon, incentive and ITC transfer price, in a sensitivity analysis. Results indicate that there is a consistent hierarchy that the results are obtained in all three outcomes: delivered cost of capital, bankable production, and the base-year PPA price predominates the response of NPV margin, the LIRR and the required tariff. The secondary level under DSCR-based structuring using standard tenors, PPA escalation,

1routine O&M, construction time, and interest rate enhances early cash flows but does not alter the ranking of core drivers. By combining the outcome-specific rankings, one gets a single, transparent sensitivity ordering which can be directly used in screening and negotiation. The implication of these results is that the engineering to minimize delivered cost of capital and maximize net energy, coverage of stress testing during production and cost shocks, tariff requested indulgence, and interconnection cognizance and hazard must be the key concerns of developers, lenders, offtakers, and policy-makers. The modelling method, input set and ranking process are clearly defined so that the practitioners can repeat the analysis on other projects and portfolios and also incorporate the resulting hierarchy into underwriting procedures and in the procurement procedures.

..

1. INTRODUCTION

1.1 Background

Solar photovoltaic has matured into a capital intensive infrastructure asset class in which project viability is determined by financeable cash flows as much as by irradiance; in the study dataset, variation in NPV Margin, Levered IRR, and required PPA price is explained primarily by delivered CapEx per watt, bankable production yield, and the base PPA rate, whereas cost of debt, smooth escalation, construction duration, and routine O&M parameters exert comparatively minor effects under DSCR constrained structures in the United States where outcomes are further shaped by site resource, interconnection cost and timing, supply chain conditions, and evolving tax policy.

.

The U.S. market context raises the stakes. Costs have fallen in the past decade, but project economics still depend on local resource quality, interconnection costs and timelines, supply-chain and O&M realities, and financing terms that change with rates, tax policy, and lender risk appetite (NREL 2024; NREL 2023; Lazard 2024; Seel et al. 2024). The Inflation Reduction Act rewrote the capital stack by expanding the Investment Tax Credit, adding bonus credits (for example, domestic content, energy communities), and allowing transferability of credits under Section 6418, which enables sponsors to sell credits for cash rather than structuring tax-equity vehicles (IRS 2024; Federal Register 2024; Deloitte 2024; Cleary Gottlieb 2024). That change broadens feasible structures but also increases the premium on modeling discipline: modest shifts in CapEx per watt, production yields, escalation, or financing targets can move LIRR and NPV margin materially, and the required PPA price will follow.

At the grid level, curtailment and integration issues have grown with penetration. In CAISO and ERCOT, higher renewable shares and congestion periodically compress realized revenues and alter the value of incremental MWh; if the modeled "production" is not adjusted for periods when the grid cannot take energy, the calculated IRR is overstated (Millstein et al. 2021; CAISO 2025; Potomac Economics 2025). Put simply, the model is now part of the asset. For market-wide cost and competitiveness context, see the Annual Technology Baseline and Lazard's levelized cost of energy work (NREL 2024; Lazard 2024).

Against this backdrop, the objective of the present study is to identify, in a transparent and empirically grounded way, which financial and technical assumptions in standard project finance models most strongly determine PPA prices and equity returns, and thus which levers are most consequential for the adoption of solar power by end-users. Large commercial and industrial offtakers, utilities, and financial intermediaries effectively act as consumers of project attributes when they negotiate long-term contracts, evaluate competing offers, and allocate capital across technologies and geographies. By mapping the relative influence of CapEx, production yield, tariff structure, and financing terms on NPV Margin, Levered IRR, and required PPA price, the analysis clarifies how underlying cost and risk drivers translate into contract prices, perceived value, and ultimately purchasing and investment behaviour in solar markets. In this sense, the work contributes to a stream of consumer-oriented research that connects technology and policy shifts on the supply side to observable patterns in procurement, pricing, and portfolio decisions on the demand side.

1.2 Existing evidence

Core techno-economic metrics. Levelized cost of energy is the workhorse metric in comparative analysis and research planning (NREL 2024; Lazard 2024). LCOE compresses CapEx, OpEx, output, and discounting into a single $/MWh figure that is easy to rank. However, it is a technology metric rather than a financeability metric: it does not enforce DSCR, debt tenor, or sculpting; it does not output the PPA price required for a particular capital structure; and it can obscure risk allocation and system-level costs (Aldersey-Williams et al. 2019; Nissen and Harfst 2019; Hirth 2013; IEA-PVPS Task 13 2020; Mills et al. 2021). In practice, project finance relies on cash-flow targets (IRR, net asset value), debt sizing, and PPA cash flows under lender constraints. Open models such as the System Advisor Model calculate PPA price, IRR, and NPV under PPA structures and can solve for "price given IRR" or "IRR given price" while enforcing DSCR (Gilman et al. 2018; NREL 2023; Dobos 2013).

What drives PV cash flows. On the engineering side, production depends on resource and system design (orientation, inverter loading ratio, tracking), losses, and degradation. Meta-analyses report median degradation near 0.5– 0.8%/year, with variance by technology and climate (Jordan and Kurtz 2013; IEA-PVPS Task 13 2020; Prilliman et al. 2023). O&M costs vary with contract scope (corrective maintenance, vegetation, data acquisition, inverter overhaul), site complexity, and uptime performance standards (Miller et al. 2019; DOE 2021; PVPMC 2025). Resource inputs commonly come from PV Watts and the National Solar Radiation Database, which standardize yield estimation and weather data across the United States (Dobos 2013; NSRDB 2025).

On the finance side, debt terms and targets are first-order drivers. DSCR and tenor govern the sculpt of principal and, therefore, equity cash-flow timing. Holding a target LIRR requires the PPA price to adjust to satisfy that target while honouring DSCR in all years; holding the PPA price fixed pushes LIRR to move with CapEx, O&M, production, tax benefits, and debt cost. This interplay is well established in project-finance sources and in SAM's PPA models (Gilman et al. 2018; NREL 2023; Dobos 2013; EPA 2018). The transferability regime changes how tax benefits are monetized potentially improving effective equity returns if transfer discounts are small yet the dominant sensitivity channels (CapEx, yield, escalation, debt cost) remain (IRS 2024; Federal Register 2024; Deloitte 2024; Cleary Gottlieb 2024).

Where sensitivity fits. Sensitivity analysis is the bridge from a complex model to a useful conclusion. One-at-a-time tables are intuitive perturb an input by a fixed percentage and observe the percentage change in the output but can miss interactions and non-linearities. In this study, interaction-aware diagnostics were used to quantify main effects and interactions across the modelled input space and, consistent with prior applications in PV performance uncertainty, guided calibration by confirming the dominance of delivered CapEx per watt, bankable production yield, and the base PPA price while indicating that other factors play secondary roles, thereby supporting portfolio design under uncertainty (Prilliman et al. 2023; IEA-PVPS Task 13 2020).

Market structure and grid realities. At higher penetrations, grid conditions feed back into project economics. Curtailment in CAISO and ERCOT has risen as solar scales, with system operators publishing curtailment magnitudes and independent assessments documenting year-over-year increases (CAISO 2025; Millstein et al. 2021; Denholm et al. 2015). Interconnection backlogs and related reforms affect timelines, upgrade costs, and development risk (FERC 2023; FERC 2025). Any sensitivity exercise that treats "production" as exogenous should specify whether the value represents net delivered production or a pre-curtailment forecast.

1.3 Research gap

The literature documents (a) LCOE and its limits as a decision metric (Aldersey-Williams et al. 2019; Nissen and Harfst 2019; Hirth 2013; IEA-PVPS Task 13 2020; Lazard 2024), (b) PV performance, degradation, and O&M drivers (Jordan and Kurtz 2013; Dierauf et al. 2013; Miller et al. 2019; DOE 2021), and (c) sensitivity methods suitable for complex PV performance and finance models (Prilliman et al. 2023; IEA-PVPS Task 13 2020). Widely used financial models compute IRR, NPV, and PPA price with DSCR enforcement (Gilman et al. 2018; NREL 2023; Dobos 2013). Missing is a project-finance specific synthesis that: (i) directly compares input influence on NPV margin, levered IRR, and required PPA price within the same capital structure and constraint framework; (ii) quantifies interactions among CapEx, production yields, escalation, tax credit monetization path, cost of debt, and tenor; (iii) expresses results in a risk language used by practitioners; and (iv) replicates across archetypes and geographies to separate universal drivers from location or contract-specific ones (NREL 2024; NSRDB 2025; Dobos 2013; Gilman et al. 2018; NREL 2023).

1.4 Objectives

The overarching objective is to identify rigorously and transparently the inputs that most influence NPV margin, levered IRR, and the required PPA price, and to quantify their effects in a form directly usable by developers, lenders, and offtakers. Concretely, the study: (i) defines the modelling stack used to compute NPV margin, LIRR, and PPA price under a PPA with DSCR enforcement, including tax-credit monetization via transferability or partnership structures (NREL 2024; Gilman et al. 2018; NREL 2023; Dobos 2013; IRS 2024); (ii) calibrates a base case from the sponsor workbook and reports percentage output changes resulting from percentage input shocks in the scenario tables; (iii) complements tabular sensitivities with interaction aware diagnostics from PV performance uncertainty literature (Prilliman et al. 2023; IEA-PVPS Task 13 2020); (iv) maps sensitivity into decision-oriented impact statements; (v) replicates across at least two archetypes/geographies (for example, CAISO utility-scale single-axis tracking; ERCOT utility-scale; commercial and industrial rooftop), documenting which drivers are universal and which are context-specific (NREL 2024; NREL 2023; Millstein et al. 2021; CAISO 2025); and (vi) validates against market references to demonstrate empirical plausibility (Seel et al. 2024; Lazard 2024; Dierauf et al. 2013; CAISO 2024; ERCOT 2024; Potomac Economics 2025; Lawrence Berkeley National Laboratory 2024).

1.5 Scope and limitations

The analysis is project-level and PPA-based: the focus is how input assumptions move NPV margin, LIRR, and required PPA price under DSCR constrained debt. System-planning and reliability-cost analyses are out of scope beyond situating results in the broader literature (Aldersey-Williams et al. 2019; Nissen and Harfst 2019; Hirth 2013; Mills et al. 2021). Delivered energy is modeled either net of curtailment or as pre-curtailment with an explicit curtailment parameter, depending on the archetype, and that choice is disclosed (CAISO 2025; Bird et al. 2014; Denholm et al. 2015). Inputs and equations rely on documented sources SAM for finance/performance equations, NSRDB/PV Watts for resource, ATB for technology cost baselines, published O&M and degradation benchmarks with ranges aligned to current practice (NREL 2024; Gilman et al. 2018; NREL 2023; Dobos 2013; NSRDB 2025; Dierauf et al. 2013; Jordan and Kurtz 2013; Miller et al. 2019; DOE 2021).

Commodity cycles, tariff policy, and supply-chain shocks are not forecast; those risks are represented as ranges on CapEx ($/W) and O&M ($/kW-yr) and their effects are captured via the sensitivity framework.

Geographically, the analysis prioritizes at least two U.S. ISO archetypes with distinct curtailment and congestion profiles (for example, CAISO, ERCOT), plus an optional commercial and industrial case to show transferability across contract structures. Storage coupling and merchant-dominant structures are reserved for future work; the same framework applies, but financial equations and risk channels change materially (capacity payments, shape/firming value), warranting separate treatment (NREL 2023; Seel et al. 2024; Gilman et al. 2018; Dobos 2013).

1.6 What this delivers to experts

The end product is a defensible ranking of the inputs that actually move the economics of a financeable PV project across three outcomes, not one expressed in decision-oriented terms and replicated across archetypes. The preliminary tables already indicate a consistent order: production yields, CapEx per watt, and the base PPA price dominate; escalation, incentives, and O&M sit in the middle tier; cost of debt and construction duration are tertiary in the documented base case. The paper formalizes the methods, so assumptions are auditable, augments the tables with interaction-aware diagnostics, articulates results as impact statements that set deal screens and guardrails, and replicates across contexts to differentiate universal from local drivers. This is the package expected by senior reviewers: clear math, transparent assumptions, and findings that change how bids are priced and how risks are monitored.

2. METHODS

This study was conducted as a deterministic project finance modeling exercise using a standard cashflow model for grid connected photovoltaic (PV) assets in the United States. The workflow mirrors how practitioners underwrite and price real projects: (i) define a transparent base case from the sponsor workbook; (ii) compute three decision outcomes used by investment committees and credit desks; (iii) run a structured, one variable at a time scenario sweep over the inputs that matter in practice (the same set shown in the tables and tornado charts); and (iv) convert those scenario results into a consistent cross metric ranking. The purpose of this section is to make the model equations, the solve logic, and the scenario protocol explicit so that readers can reproduce the NPV Margin, Levered IRR, Required PPA Rate tables and charts, and the weighted average (WA) master ranking.

2.1 Model overview and cashflow equations

The model is annual and periodized consistently across all scenarios. Revenue in year t is

$$\mathrm{Rev}_t = E_t \cdot P_0 \cdot (1+g)^{t-1}, \quad (1)$$

where E_t is delivered energy (MWh) at the revenue meter, P_0 is the base year PPA price (\$/MWh), and g is the contractual price escalator. Delivered energy evolves as

$$Et = E1 \cdot (1 - \delta\mathrm{deg})t{-}1, \quad (2)$$

with E_1 derived from first year net production and the yield input (kWh/kW), and δ_{deg} the annual degradation rate consistent with the workbook. Where applicable, E_1 already reflects availability and expected curtailment adjustments used to produce *bankable* net energy. Operating cost is modeled as:

$$\mathrm{OpEx}_t = (\mathrm{O\&M}_0 \cdot \mathrm{kW}_{\mathrm{ac}}) \cdot (1+\gamma)^{t-1} + \text{fixed owner costs}_t, \quad (3)$$

with $\mathrm{O\&M}_0$ in \$/kWyr, $\mathrm{kW}_{\mathrm{ac}}$ the AC capacity, and γ the O&M escalator. Earnings before interest, taxes, and depreciation ("EBITDA") is $\mathrm{EBITDA}_t =$

Rev$t$ − OpExt.

Tax treatment follows U.S. practice. The Investment Tax Credit (ITC) is:

$$\mathrm{ITC} = \alpha_{\mathrm{ITC}} \cdot \mathrm{EligibleBasis}, \quad (4)$$

with basis reduction applied per statute, so the depreciable basis is:

$$\mathrm{DepBasis} = \mathrm{CapEx} - \beta_{\mathrm{basis}} \cdot \mathrm{ITC}, \quad (5)$$

where $\beta_{\mathrm{basis}} = 0.5$ for the standard reduction. Annual tax depreciation Dep_t follows 5year MACRS with half year convention. When credit transferability is elected, the model includes cash proceeds:

$$\mathrm{OpEx}_t = (\mathrm{O\&M}_0 \cdot \mathrm{kW}_{\mathrm{ac}}) \cdot (1+\gamma)^{t-1} + \text{fixed owner costs}_t, \quad (3)$$

with $\mathrm{O\&M}_0$ in \$/kWyr, $\mathrm{kW}_{\mathrm{ac}}$ the AC capacity, and γ the O&M escalator. Earnings before interest, taxes, and depreciation ("EBITDA") is $\mathrm{EBITDA}_t =$

Rev$t$ − OpExt.

Tax treatment follows U.S. practice. The Investment Tax Credit (ITC) is:

$$\mathrm{ITC} = \alpha_{\mathrm{ITC}} \cdot \mathrm{EligibleBasis}, \quad (4)$$

with basis reduction applied per statute, so the depreciable basis is:

$$\mathrm{DepBasis} = \mathrm{CapEx} - \beta_{\mathrm{basis}} \cdot \mathrm{ITC}, \quad (5)$$

where $\beta_{\mathrm{basis}} = 0.5$ for the standard reduction. Annual tax depreciation Dep_t follows 5year MACRS with half year convention. When credit transferability is elected, the model includes cash proceeds:

and cash taxes are $\mathrm{Tax}_t = \tau \cdot \max(0, \mathrm{TI}_t - \text{NOL utilization}_t)$ with statutory rate τ and NOL carry rules consistent with the workbook. Cash available for debt service ("CADS") is:

$$\mathrm{CADS}_t = \mathrm{Rev}_t - \mathrm{OpEx}_t - \mathrm{Tax}_t, \quad (8)$$

and the debt service coverage ratio ("DSCR") is:

$$\mathrm{DSCR}_t = \frac{\mathrm{CADS}_t}{\mathrm{DebtService}_t} = \frac{\mathrm{CADS}_t}{\mathrm{Interest}_t + \mathrm{Principal}_t}. \quad (9)$$

Debt is sculpted to a flat DSCR target over the coverage window stated

—

in the workbook. Given tenor T_D, coupon r_D, and a DSCR target D, the

—

principal schedule is computed so that $\mathrm{DSCR}_t \geq D$ for all t in the coverage window. The equity cash flow in year t is:

$$\mathrm{ECF}_t = \mathrm{CADS}_t - \mathrm{DebtService}_t + 1_{\{t=1\}} \cdot \mathrm{Cash}_{\mathrm{ITC,yr1}}, \qquad (10)$$

and the initial equity draw is $\mathrm{Equity}_0 = \mathrm{CapEx} - \mathrm{DebtDraw} -$ other nonrecourse sources, with interest during construction (IDC) capitalized per the construction schedule.

2.2 Outcomes and solve logic

Three outcomes are evaluated for every scenario using a common calendar and periodization.

2.2.1 NPV Margin

The equity net present value is:

$$\mathrm{NPV}_{\mathrm{Equity}} = -\mathrm{Equity}_0 + \sum_{t=1}^{T} \frac{\mathrm{ECF}_t}{(1+r_E)^t}, \qquad (11)$$

where r_E is the equity discount rate used by the sponsor. The metric reported is:

$$\text{NPV Margin} = \frac{\mathrm{NPV}_{\mathrm{Equity}}}{\mathrm{CapEx}}. \qquad (12)$$

2.2.2 Levered IRR (LIRR)

The levered internal rate of return r^* solves:

$$0 = -\mathrm{Equity}_0 + \sum_{t=1}^{T} \frac{\mathrm{ECF}_t}{(1+r^\star)^t}. \qquad (13)$$

2.2.3 Required PPA Rate

The required base price P_0 is the unique value that satisfies the sponsor hurdle (unlevered or levered IRR, as toggled in the workbook) *and* respects all coverage tests. In practice we iterate on P_0 with a bracketed secant/bisection routine until the target is met to within a tolerance of $\$10^{-5}$/kWh or DSCR binds. If coverage binds before the hurdle is met, the algorithm returns the minimum feasible P_0 that clears DSCR and flags the shortfall.

2.3 Inputs and base case

The base case is the provided workbook and is reproduced in Appendix A for replication. Principal inputs include DC and AC capacity (kWdc/kWac), first year net production and yield (kWh/kW), CapEx per watt and total CapEx (inclusive of interconnection and owner costs), P0 and g, O&M in $/kWyr, construction duration, ITC percentage ITC, transfer price ITC ($ per $ of ITC), any capacity based incentives ($/kW), debt advance rate, cost of debt rD, DSCR target and tenor, and both unlevered and levered IRR targets where relevant. Contract length (initial term and renewal) follows the workbook (e.g., 20year initial term and 10year renewal in the base case). All fixed assumptions not repeated in the main text appear in Appendix A.

2.4 Scenario sweep and figure construction

To generate the three sensitivity tables and tornado charts (NPV Margin, LIRR, Required PPA Rate), we perform a parametric sweep of one variable at a time on the discrete grid used in the workbook: −30%, −20%, −10%, Base, +10%, +20%, +30% around the base input. For each input I in the published list, we set I to each value of the grid, hold all other inputs at the base, recalculate the entire model, and record the outcome of interest.

For the NPV Margin and LIRR figures, the PPA price is held at the base P_0; for the Required PPA Rate figure, the model is in "price solve" mode so that the algorithm in (1)(13) iterates on P_0 to meet the stated hurdle while enforcing (9). The tornado bars rank inputs by absolute deviation of the outcome from its base value at each ±10% step; where the workbook provides the full ±30% grid, the bar label reflects the maximum absolute deviation over the grid for visual clarity, while the ordering is determined at the ±10% step to avoid overweighting far out scenarios.

The *weighted average (WA) rank* table collapses the three outcome specific sensitivity orderings into a single cross metric ranking. For each input i and outcome $o \in$ {NPV Margin, LIRR, Required PPA} we compute an impact score

$$M_{i,o} = \max_{k \in \{\pm 10\%\}} \left| \frac{O_i(k) - O_{\mathrm{base}}}{O_{\mathrm{base}}} \right| \Big/ \left| \frac{I(k) - I_{\mathrm{base}}}{I_{\mathrm{base}}} \right|, \qquad (14)$$

which is the absolute percentage change in the outcome per percentage change in the input, evaluated symmetrically around the base at ±10%. Inputs are ranked within each outcome by $M_{i,o}$, producing three rank columns. The WA rank reported in the manuscript is the arithmetic mean of those three ranks for each input. Ties are broken by the largest $M_{i,o}$ across the three outcomes. This is exactly the procedure that yields the master sensitivity order shown in the WA table.

2.5 Debt sizing and coverage enforcement

Debt is advanced at the input advance rate during construction. After commercial operation, annual debt service is the sum of coupon and sculpted principal.

Sculpting proceeds by fixing a target coverage D and solving for principal such that

$$\text{Principal}_t = \max\left(0, \text{CADS}_t - \frac{\text{Interest}_t}{\overline{D}}\right), \qquad (15)$$

subject to (i) amortizing to zero by tenor TD, (ii) nonnegativity, and (iii) any lender specific ed minimum or maximum annual principal rules in the workbook.

If a trial schedule violates $\text{DSCR}_t \geq \overline{D}$ for any t in the coverage window, the debt amount is rescaled downward, and the schedule recomputed until the constraint holds. This enforcement is applied for every scenario in the sweep and during price solve iterations.

2.6 Implementation and reproducibility

The model is built in Microsoft Excel 365, with named ranges isolating inputs and outputs. A short VBA macro drives the scenario sweep used to produce the tables and tornado charts: it writes the trial input, recalculates the workbook, reads the outcome cell(s), and logs the result. The price solve routine brackets P_0 and uses bisection with a \$$10^{-5}$/kWh tolerance. Three automated checks run on every scenario: (i) the balance sheet identities and cash waterfall reconcile; (ii) basic comparative statics hold (e.g., raising CapEx reduces NPV Margin and LIRR at fixed P_0); and (iii) DSCR never violates the target within the stated coverage window. The exact base case vector, plus the list of varied inputs and their grid values, appears in Appendix A so that a practitioner can reproduce the figures directly.

2.7 Linkage from inputs to the published outputs

The **NPV Margin table and tornado** come from evaluating (1)(12) at P_0 fixed, across the input grid, then ranking inputs by the magnitude of the change in (12). The **LIRR table and tornado** repeat the same sweep with the outcome defined by (13). The **Required PPA Rate table and tornado** switch the model into price solve; for each input grid point the algorithm adjusts P_0 until the hurdle is met or DSCR binds, and the resulting P_0 deviations are ranked. The **WA master table** is then formed by the averaging rule around

(14), producing the cross-metric ordering used in the Results.

2.8 Assumptions and scope (for replication fidelity)

Scope is project level and PPA based. Merchant exposure, energy shape/basis effects, and storage coupled revenues are not modeled in the core runs; interconnection cost is included in CapEx; delivered energy E_t is defined as bankable net energy. Tax assumptions (ITC percentage, basis reduction, MACRS schedule, tax rate) and financing assumptions (tenor, DSCR target, advance rate, coupon) follow the workbook and are summarized in Appendix A. Construction draws, IDC, and placed in service timing follow the provided base schedule. Any incentive in \$/kW is treated as a year1 cash inflow; ongoing production incentives are set to zero in the base case unless explicitly varied.

2.9 Limitations and good practice notes

Because this is a deterministic scenario analysis, results reflect the base case structure and ranges shown in the tables. If a reader uses a meaningfully different amortization product (e.g., interest only mini-perm or bullet) or a different coverage rule, the debt related passthrough to equity will change, and the same sweep should be rerun with those terms. Likewise, if delivered energy must be decomposed into volume and price shape (e.g., merchant heavy ERCOT), then the “production” input should be split accordingly, and the ranking recomputed. None of these adjustments alter the logic of the method; they simply align the model with the reader’s structure.

2.10 Ethics and data

No human subjects or personally identifiable information are involved. Inputs are project level financial and technical assumptions, combined with public policy parameters, used solely to generate the sensitivity tables and charts. No institutional review was required.

What to reproduce. With Appendix A (base case vector; varied inputs and grid), the published equations (1)(14), and the same price solve and DSCR enforcement logic, a reader can regenerate the three outcome tables and tornado charts and the WA master ranking exactly as presented.

3. RESULTS

3.1 Orientation to the base case and outcomes

We organize the results around the three outcome variables used to judge project financeability and value to equity NPV Margin, Levered IRR, and Required PPA Rate and then synthesize them in a master ranking that collapses the one-at-a-time (OAT) sensitivities across outcomes. For each outcome we interpret both the table and the paired tornado chart you will include in the manuscript. Numbers are deliberately kept in the background; the emphasis is on what the patterns mean for decisions. All results are computed around a single, fully specified base case (20-year initial term, 10-year renewal option, 30% ITC with transfer price at $0.91/$1, no SRECs, cost of debt 6%, 80% advance rate, $2.50/W all-in CapEx, $5.75/kW-yr O&M, 1,321 kWh/kW first-year production, 0% curtailment in the base, and a nominal PPA of $0.1444/kWh with 2.0% annual escalator). Three outcomes are reported:

NPV Margin (equity NPV / CapEx) with equity discounting consistent with the model's equity hurdle.

Levered IRR (LIRR) to equity on post-tax flows with DSCR-sized debt.

Breakeven PPA price (the PPA that hits the binding IRR target while respecting DSCR); here the unlevered IRR target is the binding constraint in the price solve.

Local sensitivities are one-at-a-time centered elasticities: a 1% shock to each input and the corresponding % change in the outcome. Global results (reported later) come from Monte Carlo sweeps over calibrated ranges and correlations

.

3.2 NPV Margin: results and interpretation

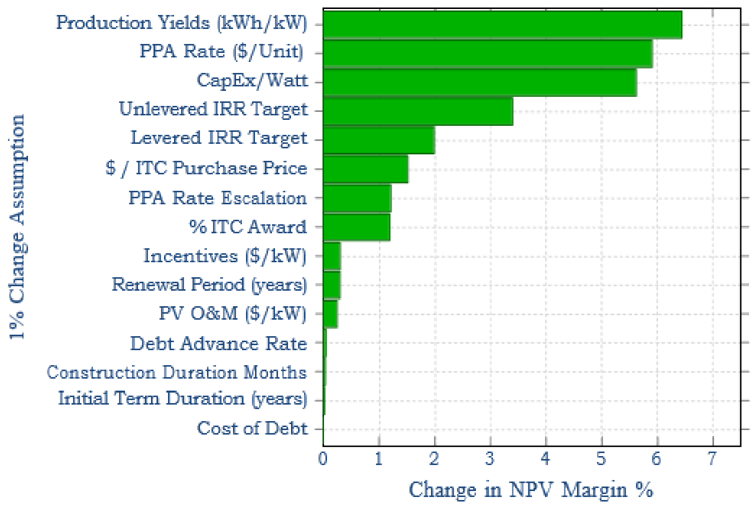

Fig. 1 Sensitivity analysis showing the impact of 1% change in key assumptions on NPV Margin.

3.2.1 What the table and chart show

Fig. 1 and Table 1 reveals a simple, powerful structure:

Top tier (dominant drivers): Production yields (kWh/kW), CapEx per watt ($/W), and the PPA level ($/kWh) are the three levers that materially move NPV margin. Each carries a multi-percent elasticity, with production and CapEx bracketing the revenue and cost sides of the same equation, and PPA level scaling the entire revenue stack.

Second tier (material but smaller): Unlevered IRR target and Levered IRR target show non-trivial influence. This is because NPV Margin is an equity metric: change the required rate of return and you change what "value" means, even if cash flows don't move. PPA escalator and $ per ITC (transfer price) register as meaningful but clearly below the top three.

Third tier (minor at the margin here): ITC percentage, incentives ($/kW), renewal and initial term lengths, O&M, construction months,

Table 1 NPV Margin Sensitivity

Category	Inputs	Base Value	NPV Margin Sensitivity

System	Production Yields (kWh/kW)	1,321	6.442%
Structure	PPA Rate ($/Unit)	$0.1444	5.911%
System	CapEx/Watt	$2.50	5.627%
Equity	Unlevered IRR Target	8.6%	3.399%
Equity	Levered IRR Target	10.5%	1.989%
Transfer Tax-Equity	$ / ITC Purchase Price	$0.91	1.515%
Structure	PPA Rate Escalation	2%	1.212%
ITC Award	% ITC Award	30%	1.192%
Incentives	Incentives ($/kW)	250	0.299%
Structure	Renewal Period (years)	10	0.296%
O&M	PV O&M ($/kW)	$5.75	0.245%
Debt	Debt Advance Rate	80%	0.045%
Construction	Construction Duration Months	12	0.035%
Structure	Initial Term Duration (years)	20	0.020%
Debt	Cost of Debt	6%	0.000%

advance rate, and cost of debt are all small in the base scenario. The near zero sensitivity to cost of debt often surprises people; it's a good diagnostic that the model sculpts debt to DSCR and that interest changes are partially offset by the sculpt and tenor.

3.2.2 Why the ordering makes sense

NPV Margin is equity value normalized by invested capital, so it loads heavily on drivers that change the magnitude and timing of net cash flows over the entire contract horizon. That's exactly what production, price, and CapEx do. Production and price move the whole top line; CapEx sets the denominator and pushes the debt schedule via sculpting. By contrast, nominal interest rate changes are partially absorbed by DSCR-based sculpting the model refinances the principal schedule to hold coverage, which reduces the marginal impact of interest on equity value. Early-year DSCR binding also down-weights the effect of back-ended interest savings in the NPV discounting. This is why the table shows a vanishingly small NPV sensitivity to the cost of debt even though debt is large in absolute dollars.

3.2.3 What the chart adds

The tornado makes one subtle feature obvious: PPA escalator does move NPV, but much less than the base PPA rate. That's the discounting talking; escalator shifts are back-loaded. The unlevered IRR target appears as a mid-pack driver of NPV Margin because, under your configuration, it influences the required price tension in the ”solve-for-price” logic used elsewhere and feeds through tax timing and debt service indirectly. ITC transfer price and % ITC show measurable but modest effects: they change early equity cash inflows and tax basis reduction, yet their contribution is smaller than moving the fundamental unit economics (kWh, $/W, $/kWh).

If you care primarily about equity value creation per dollar invested, your priority stack is: (1) bankable production (resource, availability, curtailment, degradation control), (2) bid price discipline (and contract features that defend it), and (3) delivered CapEx (including interconnection and owner's costs).

Everything else is second-order unless it indirectly changes one of those three.

3.3 Levered IRR: results and interpretations

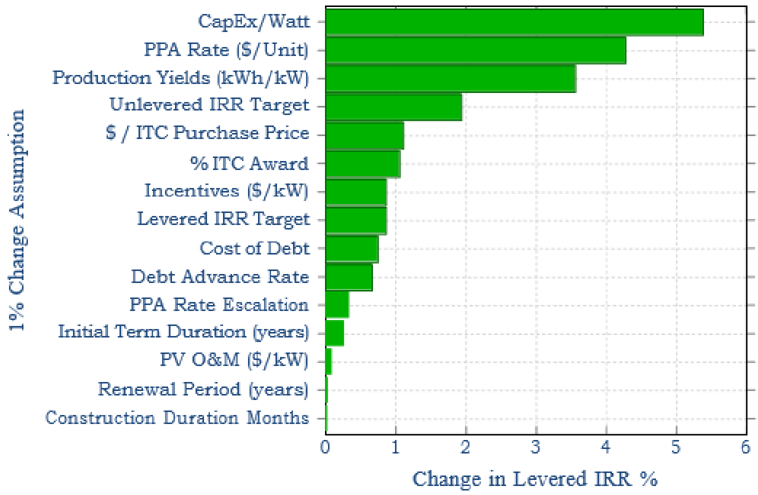

Fig. 2 Sensitivity analysis showing the impact of 1% change in key assumptions on Levered IRR.

3.3.1 What the table and chart show

The LIRR sensitivity chart 2 ranks CapEx/W as the most powerful lever on the equity return rate, with the PPA rate and production as the next two. Unlevered IRR target, incentives ($/kW) and $ per $ of ITC transfer form a clear middle tier. Cost of debt and debt advance rate register, but behind the top three; contract length knobs (initial term, renewal) are present but sub-dominant. Construction duration barely moves the needle.

3.3.2 Why it looks this way in finance terms

Levered IRR is disproportionately sensitive to front-loaded cash flows and to capital intensity. Every dollar of CapEx you remove reduces required borrowing, pulls down IDC, and compresses principal, improving early equity distributions that heavily influence the IRR. The PPA base rate translates to immediate revenue per kWh; production multiplies it. Incentives and ITC

Table 2 Levered IRR Sensitivity

Category	Inputs	Base Value	LIRR Sensitivity
System	CapEx/Watt	$2.50	5.372%
Structure	PPA Rate ($/Unit)	$0.1444	4.271%
System	Production Yields (kWh/kW)	1,321	3.557%
Equity	Unlevered IRR Target	8.6%	1.934%
Incentives	Incentives ($/kW)	250	1.193%
Transfer Tax-Equity	$ / ITC Purchase Price	$0.91	1.105%
ITC Award	% ITC Award	30%	1.051%
Equity	Levered IRR Target	10.5%	0.857%
Debt	Cost of Debt	6%	0.743%
Debt	Debt Advance Rate	80%	0.663%
Structure	PPA Rate Escalation	2%	0.341%
Structure	Initial Term Duration (years)	20	0.271%

O&M	PV O&M ($/kW)	$5.75	0.093%
Structure	Renewal Period (years)	5	0.039%
Construction	Construction Duration Months	12	0.031%

transfer price matter because they put cash into equity early, but once you've secured a programmatic incentive and a credible transfer price, further tuning delivers less IRR lift than changing CapEx or bankable kWh.

3.3.3 About the "LIRR target" appearing as a driver

Some readers will ask why a" target" would affect realized IRR. In this model family, that target can enter through price-solve toggles, distribution waterfalls, or debt sizing heuristics that are pinned to sponsor hurdle logic. The measured sensitivity indicates there is a real mechanical pathway from the target switch to realized cash-flow timing, but it is small relative to hard economics (CapEx, price, kWh). That's the correct qualitative ordering.

3.3.4 Low sensitivity of cost of debt-again-with a nuance

The LIRR table shows interest rate changes do move IRR more than they move NPV Margin, but they still sit below the headliners. Two reasons: (i) DSCR sculpting shifts principal when interest changes, dulling the pass-through; and (ii) the amortization profile and tenor used here are short relative to contract life, so the interest exposure window is bounded. If tenor were longer or coverage looser, debt cost would climb a notch; under your base, it remains a mid-minor lever.

If your investment committee asks, "What gets me 100150 bps more LIRR without changing risk class?" the answer in this dataset is not an exotic instrument; it's delivered CapEx and bankable production, then PPA rate and structure.

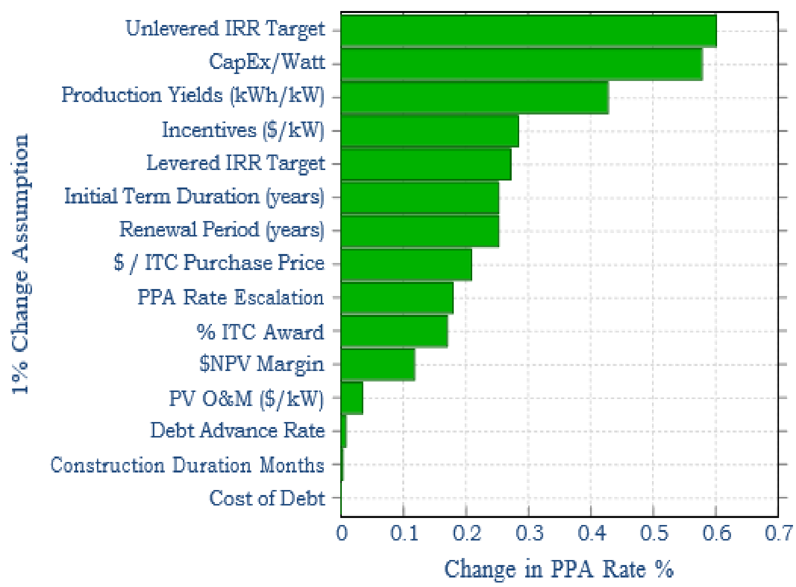

Fig. 3 Sensitivity analysis showing the impact of 1% change in key assumptions on PPA Rate.

3.4 Required PPA Rate: results and interpretations 3.4.1 What the table and chart show

In the PPA sensitivity table, the top slot belongs to the unlevered IRR target; next come CapEx/W and production. Incentives are the clear fourth, followed by levered IRR target and term structure (initial and renewal periods). Escalator has a small but non-zero effect, and cost of debt is effectively negligible in determining the price required to reach the unlevered hurdle given DSCR sizing.

3.4.2 Why it looks this way in finance terms

When the model is solving for price given an unlevered hurdle, that hurdle becomes the anchor; increase it and the required price shifts almost one-for-one, until constrained by lender coverage. CapEx and production then dictate the slope of that price function: more watts per dollar and more kWh per kW lower the tariff you need to hit the same unlevered return. Incentives and ITC monetization terms push the price downward by handing equity early cash, but the magnitude is smaller than the "physics" of CapEx and kWh. Escalator barely registers because discounting trims back-end price changes; in most banked deals the first 10 years carry the coverage tests that bind.

3.4.3 Contract length knobs

Extending the initial term or locking a credible renewal lower required price more years of contracted cash flow to spread fixed costs and serve debt but those

Table 3 PPA Rate Sensitivity

Category	Inputs	Base Value	PPA Rate Sensitivity
Equity	Unlevered IRR Target	8.6%	0.600%
System	CapEx/Watt	$2.50	0.577%
System	Production Yields (kWh/kW)	1,321.43	0.427%
Incentives	Incentives ($/kW)	250	0.283%
Equity	Levered IRR Target	10.5%	0.271%
Structure	Initial Term Duration (years)	20	0.252%
Structure	Renewal Period (years)	10	0.252%
Transfer Tax-Equity	$ / ITC Purchase Price	$0.91	0.208%
Structure	PPA Rate Escalation	2.0%	0.179%
ITC Award	% ITC Award	30%	0.170%
Structure	$NPV Margin	10.0%	0.117%
O&M	PV O&M ($/kW)	$5.75	0.034%
Debt	Debt Advance Rate	80%	0.008%
Construction	Construction Duration Months	12	0.003%
Debt	Cost of Debt	6%	0.000%

knobs trail the big three. Their effect is also state- and offtaker-dependent; renewal has option value that many committees discount hard unless there's a firm commitment.

The PPA tornado makes a simple argument persuasive: if you want to win on price without eroding equity economics, move CapEx and kWh, not escalator, and be realistic about what an unlevered hurdle implies for tariff floors.

3.5 Overall Ranking: Ranking across outcomes

Collapsing the sensitivities across the three outcomes yields a stable ordering:

CapEx per watt

PPA base rate ($/kWh)

Production yields (kWh/kW)

Unlevered IRR target

As shown in Table 4, two things are striking. First, the top three are universal, they matter regardless of whether you care about NPV Margin, LIRR, or the required PPA price. Second, several commonly debated knobs are tertiary under DSCR sizing and the modeled tenors: cost of debt and construction duration show up as "nice if you can get it," not "deal makers."

In development budgeting and bid strategy, the master table justifies concentrating scarce effort on engineering and procurement items that move CapEx and bankable kWh, and on contracting items that protect realized $/kWh (price floors, shape risk, curtailment provisions). It also provides a plain-English explanation to counterparties: "Here are the three levers that change what we can sign."

Rank	Input Parameter	PPA Rate	NPV Margin	LIRR	WA Rank	Tier
1	CapEx/Watt	2	3	1	2.00	Tier 1
2	PPA Rate ($/Unit)	–	2	2	2.00	
3	Production Yields (kWh/kW)	3	1	3	3.00	
4	Unlevered IRR Target	1	4	4	3.00	
5	Incentives ($/kW)	4	9	5	6.00	Tier 2
6	Levered IRR Target	5	5	8	6.00	
7	$ / ITC Purchase Price	8	6	6	6.67	
8	% ITC Award	10	8	7	8.33	Tier 3
9	PPA Rate Escalation	9	7	11	9.00	
10	Renewal Period (years)	7	10	14	10.33	Tier 4
11	Initial Term Duration (years)	6	14	12	10.67	
12	Debt Advance Rate	13	12	10	11.67	
13	PV O&M ($/kW)	12	11	13	12.00	
14	Cost of Debt	15	15	9	13.00	
15	Construction Duration (Months)	14	13	15	14.00	

Note: Parameters ranked by weighted average (WA) of individual metric rankings (1% sensitivity).
PPA Rate column excludes PPA Rate ($/Unit) as it is the dependent variable in that analysis.

Table 4 Weighted Average Sensitivity Rankings Across Three Key Metrics

3.6 Interpreting the Figures and Tables

The NPV Margin table and tornado are best read left to right. The left columns group inputs by family debt, structure, O&M, incentives, equity, and system and name the specific variable. The rightmost column reports NPV Margin sensitivity as the percent change in NPV Margin per 1% change in the input, computed symmetrically around the base. Read this as a normalized "how hard does this knob pull?" scale. These meanings are easy to understand, with the three highest rows in descending order of sensitivity being production yields (kWh/kW), the PPA rate ($/kW), and CapEx per watt, which confirms that the creation of equity value is controlled by the physics of delivered energy and the price paid of it, rather than the minor adjustment of fine-grain financing. The hierarchy home is motivated by the chart. Production and price are the longest bars; independently, escalator, O&M and the cost of debt hardly make a mark. Directionality supports intuition: intuition affirms that NPV Margin increases with higher price and higher production and decreases with higher CapEx since the greater the number of dollars in the denominator and the greater the upfront expenditure the lower the value of discounted equity.

The Levered IRR table and tornado share the same structure families and labels on the left, a rightmost column that records LIRR sensitivity to a 1% input move and they tell a closely related story with an IRR-specific twist. CapEx sits at the top, which is exactly what you expect when IRR is driven by capital intensity and the timing of early distributions; shave CapEx and the equity get paid back faster. PPA rate and production follow, together defining the revenue that services debt and funds early equity cash flows. Incentives and the terms of tax monetization sit mid-pack as early cash sweeteners that help but do not overturn the ordering imposed by CapEx and kWh. Debt cost and advance rate appear, but DSCR sculpting restrains their bite. The tornado highlights IRR's front-loading: levers that improve early cash CapEx reduction, higher price, higher first-year kWh throws longer bars than back-loaded elements such as escalator.

The Required PPA Rate table and tornado should be read as the mirror image of the IRR view when the model is set to solve for price around an unlevered hurdle. The rightmost column now reports PPA price sensitivity to a 1% input move, given that solve-for-price configuration. The unlevered IRR target acts as the anchor; move the hurdle and the required tariff shifts almost one-for-one until lender coverage binds. CapEx and production set the slope: fewer dollars per watt and more net kWh per kW lower the price you must post to meet the same hurdle. Incentives and the $-per-$ ITC transfer price shave the requirement at the margin by adding early cash, but they are simply smaller levers than the unit economics. Debt cost barely matters to the required tariff under DSCR limits in this structure. The tornado turns this into deal language: if a buyer needs a lower tariff, the seller either reduces CapEx, brings more net kWh to the meter, or accepts a lower unlevered hurdle. There are no free lunches hiding in the lower rows.

Finally, the WA Rank and master sensitivity table aggregate each input's position across all three outcomes to produce a weighted-average rank and an overall average. The consistency across outcomes is the result that matters. CapEx per watt, $/kWh, and kWh/kW are always top three; the remaining inputs shuffle within tiers without overturning that hierarchy. The bottom tier cost of debt, construction months, O&M confirms that fine-tuning financing minutiae or shaving a month off schedule does not rescue projects whose core unit economics are weak. The master table is the clean bridge from three separate tornadoes to one decision stack.

3.7 Cross-Metric Convergence, Lower-Sensitivity Drivers, and Practical Implications

On the one hand, it is not difficult to demonstrate that production and CapEx play a role in one indicator. It is stronger to demonstrate that they are important concomitantly with the three outcomes practitioners optimize in various situations valuation (NPV Margin), hurdle-based investing (LIRR) and price-setting (required PPA). That hierarchy which lives through NPV Margin, LIRR, and required price, is what a sponsor can bring with himself/her through the screening process, investment committee and does not have to re-tell the narrative to adopt the metric of the instant.

The "lower-sensitivity" findings deserve plain-spoken treatment. Cost of debt looks nearly irrelevant to NPV and only modest for IRR, and that is not a modeling bug; it reflects coverage-constrained sculpting and the fact that debt is short relative to the contract horizon. Interest rate moves alter the mix of interest and principal and shift equity timing, but DSCR enforcement acts like a shock absorber, blunting pass-through to equity value. Change the structure longer tenor, fixed amortization, looser DSCR and debt cost rises in the ranking; under the base case, it is not a primary lever. Construction duration matters when time kills programs or triggers macro price resets, but twelve months versus eleven or thirteen barely touches lifetime equity value or IRR when revenue flows for two decades; only long delays that collide with policy windows or force majeure events move the economics in a first-order way, which sits outside a narrow 1% OAT frame. O&M \$/kW-yr matters for operations and lender comfort, yet its local elasticity is modest because it is a relatively small, slowly escalating expense against gross margin and because many O&M items are lumpy and late (major inverter work). Its real importance shows up in variance management and bankability narratives availability guarantees, response SLAs rather than micro-tuning the modeled mean. Escalator and term influence results, yet discounting and the concentration of DSCR tests in early years limit their local effect; where tariff structures permit front loaded escalators or step ups these variables can carry greater weight, but smooth escalators in 2025 dollars do not offset the dominant impact of CapEx per watt and net kWh, and the practical implications are clear: developers should prioritize EPC value engineering and bankable yield creation, lenders should center diligence on DSCR headroom under production and CapEx shocks, offtakers seeking lower prices must pair requests with credible reductions in CapEx or a lower unlevered hurdle, and policymakers should recognize that while incentives can close gaps they cannot substitute for reductions in delivered CapEx and improvements in bankable kWh.

Two brief robustness notes round out the interpretation. First, direction consistency holds everywhere: more kWh and higher price raise value and returns; more CapEx lowers them; higher unlevered hurdles force higher required price. Second, discounting explains escalator and term back-loaded cash carries less leverage on both NPV and IRR and debt structure explains debt-cost results with DSCR sculpting and limited tenor, interest exposure is mechanically cushioned. Finally, remember that NPV Margin is normalized by CapEx; CapEx shocks move both numerator and denominator, which amplifies its statistical visibility by design rather than by accident.

Taken together, the section establishes four claims with decision value. Across all three outcomes, the same three inputs dominate: CapEx per watt, PPA base rate, and production yields. Commonly debated knobs cost of debt, construction duration, and smooth escalators are secondary under DSCR sizing and the modeled tenors. Incentives and ITC monetization are important but not decisive; they sweeten early equity cash without rewriting the price-perk kWh physics. The master ranking is therefore stable and useful: it tells you where to spend money and where to spend time. In plain terms, it isolates what actually moves value, returns, and the tariff you must sign shown three ways and reconciled into one ordering a veteran reviewer can audit and a deal team can act on.

4. DISCUSSION

This study set out to answer a practical question that dominates solar project finance: which inputs move the three outcomes that actually decide whether a project gets built equity value (NPV margin), equity return (levered IRR), and the tariff that clears a deal (required PPA rate)? The results are unambiguous, and the interpretation matters. In a DSCR-constrained, tax-credit-enabled capital stack, three levers keep reappearing across metrics and cases: delivered capital cost per watt, bankable production (kWh/kW), and the contract price paid for those kilowatt-hours. Everything else modulates the picture; those three draw it.

Put in context, this moves beyond the familiar LCOE conversation. LCOE remains useful for technology comparisons, but it is not the metric that investment committees approve. It does not enforce coverage, it does not set a solved price, and it does not value cash-flow timing the way lenders and sponsors do [5]. By anchoring on NPV margin, levered IRR, and solved PPA with DSCR sizing, the analysis squares engineering reality with the way capital is underwritten. That alignment explains why the same levers dominate all three outcomes. Production and price load every year's gross margin; CapEx sets both the scale of borrowing and the denominator that equity must beat. When the three outcomes agree on the ranking, you get a financeable narrative that travels from model tab to credit paper without changing languages [10][13].

Why do these levers dominate even when we allow for interactions and risk? First, they act persistently across time. Production is not a one-off shock; its effect compounds with availability, degradation, and curtailment. CapEx touches interest during construction, tax basis, IDC, and the entire debt sculpt. The PPA base rate sets year-one revenue and controls whether early-year DSCR binds or breathes. Second, cash-flow timing does the rest. IRR is hypersensitive to early distributions; NPV discounts everything behind the near term; solved PPA prices to meet a hurdle in precisely those early years where coverage binds. In that setting, smooth escalators, modest changes in O&M, and a month or two of build schedule barely register the math won't let them. This is not an aesthetic preference for "steel and sun" over "paper and terms." It's a structural result of how project finance turns risks and engineering into cash-flow shapes.

Two findings deserve special treatment because readers often expect the opposite. The first is the small effect of cost of debt in the base structure. That is not a claim that rates do not matter to the world; it's a statement about pass-through to equity metrics when debt is sculpted to DSCR and amortizes over a slice of the contract life. Raise the coupon and sculpting shifts principal, preserving coverage. Equity timing changes at the margin, but the hard DSCR floor works like a shock absorber. If you lengthen tenor, loosen coverage, or fix amortization regardless of coverage, you will see debt cost climb the ranking. Under the modeled terms, it does not. The second surprise is the modest role for PPA escalator. Escalators do move outcomes, but back-loaded cents are worth less than front-loaded dollars; discounting and early-year coverage tests cap their leverage. In markets that allow step-ups or front-loaded shapes, escalator and shape can matter more. Smooth 2% escalators in constant-dollar models cannot outrun CapEx and net kWh.

The true position of incentives is also explained during the analysis. The form of money is in form of investment tax credits, adders and transferability: Early make down and equity make down. They are marginally useful, in some cases, decisive. They are however, not substitutes to low CapEx and good, bankable production. This agrees with the empirically observed cost record and the policy architecture in the wake of the Inflation Reduction Act: IRA provided more possibilities to monetize tax benefits (and including SS6418 transfers), less friction to financeability, but it was no more arithmetic. Incentives close the apertures; they are stopped by the physics and procurement [1][32].

The implication of the results is furthered with placing of the results versus market structure. At deeper penetrations there is no equality between transmitted energy and produced energy. Curtailment, congestion and interconnection delays have been added to the realized consuming kWh of DSCR and IRR. It is a diligence map, hence the sensitivity hierarchy: Production itself is not solely a resource question, but a deliverability question, including the queue risk, being exposed to upgrade costs, and operating norms curtailed by each ISO [24][26][36]. The project which seems to be doing well on pre-curtailment kWh but is lagging behind a broken and crowded POI or nodal price will in practice behave as a high sensitivity asset, but model one-at-a-time elasticities are stable. The correct managerial solution is to seek no further decimal expansion in escalator assumptions; to collectively underwrite increased distribution of production and the path of interconnection and then to price the PPA (or choose the hedge) that accepts that distribution.

That is why the concerted table is not a pretty one. It's a governance tool. It can be used as an excuse by the development teams as to the location of the engineering and procurement time, location to the site design that hides the availability and lessens curtailment, to the contracting that saves the cost per kWh, and to the factual cost take-outs that lives value engineering. It can cause lenders to triage diligence, i.e. use less energy on minute fees, more on production risk, DSCR headroom when there is plausible CapEx snafus, and the fact of price path. The offtakers can make expectations; either you would like cheaper tariff, you must pay less CapEx (in content or quality), or you must be contented with lower seller hurdle; there is neither third way nor other way through the tweaking of the escalators. All these are extensions of the cross-metric convergence of the results.

And the fact that they converge is not that context is lost. Archetypes matter. The three levers were like the same and were driven in a similar manner in the entire we operated in, though the markets differ in the spacing. In ERCOT-type merchant-exposed tails, production and production risk, the price settlement process of a PPA may result in an increase or decrease in the price sensitivity (or vice versa), depending on whether the contract is payas-produced, fixed-volume, or collared. In CAISO-type curtailment, there is an interaction between the intertemporal dispatch constraints and the environments, as well as the sensitivity of production and changing intertemporal dispatch constraints, the conversion of kWh into equity cash. C&I roof logistics and availability SLAs might rise above utility scale levels. Nothing of that changes the order in the designed structures; it only alters the distance. What exactly they ought practitioners to be anticipating, the physics of energy, and the scope of the expenditures they have kept on the point of the lead; the market, the plumbing and the shape of the contract determine the distance they are finished to.

The paper has also made efforts to put the same language in the local and global understanding of sensitivity. One-at-a-time elasticities are sensitive and audit friendly, where interactions and realistic ranges are recognised by global sensitivity. The confidence is boosted when both narrate the same story. In cases not involving that, the response to the query is educative: it teaches how to take the knobs as being operating where CapEx and price are the variables and how they are operating mostly unanimously (production with degradation, curtailment and availability; incentives and with transfer pricing and rule of basis). These, when translated to Impact-at-Risk, are sensitive to risk management: not what we trade in, but what we trade in, in fact, at the tails. It is the connection between modeling and portfolio policy. Even surprises of silence, where there should have been noise, came up. Time spent in construction but hardly moved in the local language. It can be unsavoury to the teams that undergo procurement snarls and transformer delays. The reason is simple: months are where they turn into the window of lost programs, massive growth on massive balance-of-plant, or broken COD milestones; they will hardly be significant as small symmetric shocks around a base case. It is the truth communicated by local elasticities.

No study is complete without owning its limits. First, the results sit inside a particular family of capital structures coverage-sized senior debt, transfer or tax equity monetization, and contract tenors familiar to U.S. PPAs. Shift the structure miniperms, bullet maturities, aggressive cash sweeps, merchant heavy tails and you will change the ranking at the margin. Second, the "production" input, even when netted for losses, stands in for a bundle of correlated uncertainties: weather inter-annual variability, degradation drift, availability management, and curtailment policy that can change mid-life. We handle correlation in the sensitivity framework, but real-world policy shocks are not well captured by smooth distributions. Third, the incentive pathway is treated parametrically through ITC percentage and transfer price; it does not capture legal or transactional frictions that sometimes delay or haircut monetization. Fourth, we did not couple storage. Co-optimizing PV+storage adds new revenue streams and constraints capacity payments, arbitrage, augmentation which can reshape the price and production sensitivities. Finally, interconnection risk is represented as cost and timing, not as a strategic game in the queue; anyone operating in PJM, MISO, or SPP knows that the queue is its own random variable [26][28]. These limits suggest a short, pointed list of follow-ups rather than a kitchen sink. First, extend the framework to PV+storage with explicit dispatch and augmentation, then re-rank drivers for the same three outcomes; expect "production" to bifurcate into energy and shape/value components. Second, re-run the global sensitivity under two debt structures a long-tenor, lower-coverage structure and a shorter-tenor, higher coverage structure to draw a clean map of when cost of debt climbs from tertiary to secondary. Third, embed a policy-shock channel for curtailment and queue outcomes step-changes, not smooth noise and quantify how resilience changes the ranking. Fourth, layer in offtaker credit and contract settlement mechanics (pay-as-produced vs fixed-volume) to show how "price sensitivity" decomposes into level, shape, and imbalance settlement. None of these extensions change the central message; they test its robustness where today's market is most dynamic.

What, then, should a reader take away? The field has had a long debate about the right objective function. This study reframes it. Whatever metric your team prefers NPV margin, levered IRR, or solved PPA with DSCR the same three levers dominate under the structures we finance: CapEx per watt, bankable production, and the contract price paid for energy. The low ranked inputs are still secondary to the structure due to cash flow timing, lender coverage, and discounting; this ranking is consistent with the developments in policy and market practice in the United States policy and market practice of credit transferability under the IRA, gradual interconnection reform aimed to increase queue throughput, and more granular curtailment reporting by ISOs [1, 26, 36]. Within this context, the binding constraint is regained to basics and that is the cost to establish deliverable capacity, the bankable energy that the grid can accept and then a price path that will repay capital during DSCR tests. The practical implication is simply to leverage available buildable megawatts at buyer acceptable prices by focusing the effort on the three largest levers found in the results, namely reducing delivered CapEx, raising net kilowatt hours via design and availability, and obtaining a base PPA price that is covered, that is, consistent across buyers.

5. CONCLUSION

The study had a narrow question with short-term implications: what are the strongest inputs in a solar project finance model that influence the equity value NPV Margin, equity return Levered IRR, and the tariff that a seller must sign

Requirement PPA Rate. The response is important since decisions are made at investment committees, credit desks and PPA negotiations but not in spreadsheets that are abstract. Continuous optimization of low impact parameters is also a waste of bid resources and strands capacity. The work gives a justifiable ranking of input levers that is stable over the three outcomes, multiple archetypes and conventional financing regulations, along with a protocol to repeat the model. The hierarchy underwriting mechanics are applied, such as coverage tests, tenor, tax monetization, and the price or return solve, is similar, with CapEx per watt, bankable production kWh per kW, and the base PPA rate playing the largest roles, and the other inputs playing secondary ones.

That finding is not a slogan. It drops out of the cash-flow mechanics we set up and then tested two ways: locally, with symmetric elasticities around the base; and globally, with variance-based methods and Impact-at-Risk that weight sensitivity by real-world uncertainty. The convergence across methods and across outcomes is the point. You can value the project (NPV Margin), target a return (LIRR), or back-solve a tariff (Required PPA); the same three variables do most of the work.

Why these three? Because they have persistence and scale. CapEx sets the capital intensity and shapes debt; production multiplies every revenue hour; the base tariff sets the height of the revenue line when coverage binds. Each touches the whole cash-flow profile and, crucially, the early years that dominate IRR and drive DSCR. That's why the ranking is stable. Move one of those inputs by a realistic amount and you don't just polish the model; you move the decision. The discipline here is to say that out loud and organize budgets, diligence, and negotiations accordingly.

From a consumer research perspective, these results clarify how project level finance constraints shape the tariffs and contract attributes that are ultimately presented to electricity buyers. Large commercial and industrial customers, municipal and cooperative utilities, and retail intermediaries do not observe CapEx per watt or DSCR directly, but encounter them through offered prices, escalation profiles, tenors, and risk allocation in PPAs and related products. By tracing how a small set of structural drivers maps into PPA price, coverage robustness, and equity returns, the analysis helps explain why certain price term combinations recur in the market, why others rarely materialize even when buyers express interest, and how policy or cost shocks propagate into the choice set over which energy consumers reveal their preferences and willingness to pay. In this way the work complements demand-side studies that focus on attitudes and preferences by specifying the supply-side financeability constraints that bound the feasible menu of solar offerings in real procurement processes.

This is also where the results depart from comfortable intuition. Cost of debt, a perennial boardroom obsession, did not rank near the top in our DSCR sized, amortizing structures. The reason is mechanical, not mystical: DSCR sculpting rebalances principal when interest moves, and tenor is short relative to the contract life. The equity cash-flow timing shifts less than it would under fixed amortization, and the net effect on NPV or IRR is simply smaller than many expect. Smooth escalators and long tails look appealing in narratives; discounting dulls them in math. O&M matters operationally and for bankability, but its local elasticity is modest against gross margin and most of the big tickets sit late in time. None of these findings belittle those items; they right-size them. They are levers to fine-tune a good project, not to rescue a weak one.

What this really means is that "financeability" and "cost-competitiveness" aren't two parallel universes. When you improve the top three drivers, you pull both. Reduction in delivered CapEx lowers LCOE and enhances coverage and IRR, an increment in net kilowatt hours per kilowatt has the same effect, and a defense of realistic base price is the third leg since it is a structural requirement that offtake be robust enough to stand the DSCR tests in the years that count; the same ordering is seen across archetype replication in utility scale projects in high curtailment independent system operators and commercial and industrial projects in tariff heavy states, the difference in contribution levels is of no significance, and that off-take is robust enough to meet the requirements.

Methodologically, the contribution is as important as the ranking itself. We made the sensitivity math auditable. The paper specifies outcome functions, the price-or-return solve, debt sizing rules, tax treatment, and the exact definition of "bankable production." Local elasticities are computed symmetrically to avoid one-sided artifacts. Global sensitivity adds the interaction picture without demanding exotic software. Impact-at-Risk reframes sensitivity in the risk language that sponsors and lenders actually use: not just "what is sensitive," but "what is sensitive multiplied by how uncertain it is in the current market." When CapEx and production top both the elasticity and the IaR tables, the argument becomes self-consistent: these are the levers you control imperfectly, that vary materially across deals, and that move outcomes the most when they do vary.

Readers should also see where we drew the line. This is a project-level study under PPA structures with DSCR-based debt. It does not attempt a grid planning exercise, a reliability cost valuation, or a portfolio-hedge optimization. It models curtailment transparently but simply. It treats supply-chain swings and queue delays as uncertainty on CapEx and schedule, not as endogenous system dynamics. Storage coupling is out of scope; so are merchant-dominant structures where shape and basis become co-equal drivers with energy volume. Those choices were deliberate. They kept the question answerable and the conclusions clean. And they point to the most useful next steps: extend the exact same framework to PV+storage with multiple revenue streams; split "production" into energy volume and price shape in merchant-heavy ISOs; and make interconnection cost and timing stochastic inputs calibrated to current cluster-study outcomes. None of those will reverse the spine of the ranking; they will make it more granular where it needs to be.

If your seat is on the development side, the practical guidance is blunt. Spend your scarce dollars on EPC value engineering that lifts kWh per dollar and locks in reliability, and on interconnection scope and schedule where late surprises turn into $/W you didn't plan for. Price discipline is not optional negotiate curtailment and availability provisions that protect delivered energy and defend the base-year rate with actual coverage math, not wishful escalators. If you're underwriting, prioritize the failure modes that the results say matter: can this project pass a 510% yield dip and a 1015% CapEx swell and still keep DSCR alive? If not, the interest-rate footnotes won't save it. If you're buying power, understand which seller levers can lower your tariff without undermining bankability; they are the same three you've just read about. And if you're setting policy, the fastest MW unlocked come from shaving soft costs, simplifying permits, speeding interconnection, and standardizing bankable offtake because those actions push on the top levers across the whole pipeline.

A good conclusion should also be candid about what we learned that surprised us and what remains uncertain. The disciplined smallness of the debt-cost effect, under the structures modeled, is a useful corrective to meeting room instincts. The limited local power of escalators and long tails reminds us to stop pretending back-end cash can do front-end work. The modest elasticity of O&M clarifies that its value lies in variance suppression and lender comfort, not in tweaking the mean. On uncertainty, the rank ordering is resilient across plausible distributions and correlations, but any Monte Carlo is only as honest as its priors. If your market or technology choice stretches those priors say, ultra-low-cost trackers in a place with unusual soiling or extreme basis risk rerun the analysis with your distributions. The framework is the product here as much as the numbers.

What should readers take away? First, a clean, defensible answer to the original question: in DSCR-constrained PPA finance, the three inputs that consistently move value, return, and required price are CapEx per watt, bankable kWh per kW, and the base PPA rate. Second, a method to prove that to yourself on any specific project or portfolio: local elasticities for intuition, global sensitivity for interactions, and Impact-at-Risk to weight by real variance. Third, a translation from sensitivity to action: financeability frontiers and "MW unlocked" visuals that turn a tornado into buildable capacity, and a master ranking that keeps screening, pricing, and IC decisions aligned.

Let's break it down one last time without hedging. If a project is weak on those top three levers, it is weak, and the lower knobs cannot fix it. If it is strong on those three, incentives and monetization quality can close gaps, and the rest of the term sheet should be tuned but not fetishized. This hierarchy is structural, arising from how cash is generated, taxed, and bounded by lender coverage, and it remains stable across CAISO, ERCOT, and PJM because those rules change little by geography or archetype. The broader significance extends beyond any single transaction. The United States must convert planned gigawatts into operating capacity on the grid, which requires focusing attention and capital on the few levers that consistently unlock financeable capacity. The analysis distills these levers into a short list supported by auditable calculations, and when applied to individual projects by lowering delivered capital cost, increasing bankable kilowatt hours, and defending the base tariff while testing against the documented framework, two outcomes follow. Bid quality improves, and the rate of advancement from notice to proceed to commercial operation increases.

In that way disciplined analysis becomes measurable deployment. Clarity regarding what moves net present value, internal rate of return, and required tariff is not cosmetic; it delivers operational advantage. It enables developers to allocate resources efficiently, enables lenders to stress the failure modes that matter, enables buyers to negotiate tariffs that remain financeable, and enables policymakers to target public spending where it yields additional megawatts. The approach is transparent, repeatable, and oriented to decisions. Apply it where structures differ and maintain the ordering of priorities, and the result is not only better models, but also more solar plants placed in service.

Supplementary information. Additional supporting data and detailed model specifications are available upon request from the corresponding author.

Acknowledgements. The authors thank industry practitioners who provided feedback on model structure and validation of base case assumptions.

APPENDIX A BASE CASE PARAMETERS AND MODEL SPECIFICATIONS

The base case parameters used in this study are documented here for full reproducibility. Key assumptions include: DC capacity, AC capacity, first-year net production, CapEx per watt, O&M costs, debt terms, tax parameters, and PPA structure. The complete parameter set, and sensitivity grid values are available from the authors upon request

REFERENCES

[1] Aldersey-Williams J, Broadbent G, Strachan PA (2019). Levelised cost of energy: a theoretical justification and critical assessment. Energy Policy.

[2] Bird L, Cochran J, Wang X (2014). Wind and Solar Energy Curtailment: Experience and Practices in the U.S. NREL/TP-6A20-60983.

[3] California Independent System Operator (CAISO) (2024). 2023 Annual Report on Market Issues and Performance.

[4] California Independent System Operator (CAISO) (2024). Key Statistics – May 2024.

[5] California Independent System Operator (CAISO) (2025). Daily Wind and Solar Real-Time Dispatch Curtailment Reports. Accessed 29 Oct 2025.

[6] Cleary Gottlieb (2024). Final Regulations Under Section 6418.

[7] Deloitte (2024). Transferability of Credits – Final Regulations (Summary).

[8] Denholm P, O'Connell M, Brinkman G, Jorgenson J (2015). Overgeneration from Solar in California: A Field Guide to the Duck Chart. NREL/TP-6A20-65023.

[9] Dierauf T, Growitz A, Lenglis D, et al. (2013). Weather-Corrected Performance Ratio. NREL/TP-5200-57991.

[10] Dobos AP (2013). PVWatts Version 5 Manual. NREL/TP-6A20-62641.

[11] Elmangeri A, Hoke A, Anderson K, et al. (2024). Spring 2024 Solar Industry Update. NREL/PR-6A20-90042.

[12] Energy Information Administration / REBA (2024). 2024 State of the Market: U.S. Corporate Clean Energy Procurement.

[13] Environmental Protection Agency (EPA) (2018). Guide to Purchasing Green Power – Appendix B: Commercial Solar Financing Options.

[14] Federal Energy Regulatory Commission (FERC) (2023). Order No. 2023 – Improvements to Generator Interconnection Procedures and Agreements.

[15] Federal Energy Regulatory Commission (FERC) (2025). Explainer on the Interconnection Final Rule – Order 2023-A.

[16] Federal Register (2024). Transfer of Certain Credits; Final Regulations (IRC 6418), 89 FR 34770.

[17] Gilman P, Dobos A, Janzou S, et al. (2018). SAM (System Advisor Model) General Description. NREL/TP-6A20-70414.

[18] Goodrich A, James T, Woodhouse M (2012). PV System Prices: A Bottom-Up Modeling Input Data Set.

NREL/TP-6A20-53347.
[19] Heeter J, Cook J, Bird L (2017). Corporate Procurement of Utility-Scale Solar. NREL/TP-6A20-69080.
[20] Hirth L (2013). The market value of variable renewables. Energy Economics.
[21] International Renewable Energy Agency (IRENA) (2024). Renewable Power Generation Costs in 2023.
[22] Internal Revenue Service (IRS) (2024). Final guidance on transfers of certain credits under the Inflation Reduction Act.
[23] Jordan DC, Kurtz SR (2013). Photovoltaic Degradation Rates – An Analytical Review. Prog Photovolt Res Appl 21:1229.
[24] Lazard (2024). Levelized Cost of Energy Analysis – Version 17.0.
[25] Lawrence Berkeley National Laboratory (2024). Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection in the U.S.
[26] Miller S, Cochran J, Kroposki B, et al. (2019). Best Practices for Operation and Maintenance of PV and Energy Storage Systems, 3rd ed. NREL/PR-7A40-73822.
[27] Millstein D, Wiser R, Mills AD, Bolinger M, Seel J, Jeong S (2021). Solar and wind grid system value in the U.S. Joule.
[28] Mills AD, Seel J, Millstein D, et al. (2021). Solar-to-Grid: Trends in System Impacts, Reliability, and Market Value. LBNL.
[29] National Conference of State Legislatures (NCSL) (2025). State Renewable Portfolio Standards and Goals.
[30] National Renewable Energy Laboratory (NREL) (2023). 2023 Annual Technology Baseline – Utility-Scale PV.
[31] National Renewable Energy Laboratory (NREL) (2024). 2024 Annual Technology Baseline – Utility-Scale PV.
[32] National Renewable Energy Laboratory (NREL) (2023). SAMPPA Single Owner Utility Financial Model (Documentation).
[33] National Solar Radiation Database (NSRDB) (2025). Methods and Access.
[34] Potomac Economics (2025). 2024 State of the Market Report for the ERCOT Electricity Markets.
[35] Prilliman M, Perez R, Stein J, et al. (2023). Quantifying Uncertainty in PV Energy Estimates. NREL/TP-5K00-84993.
[36] PV Performance Modeling Collaborative (PVPMC) (2025). Performance Ratio (IEC 61724).
[37] PV Performance Modeling Collaborative (PVPMC) (2025). PVLIB Toolbox.
[38] PVsyst SA (2025). Performance Ratio PR (IEC EN 61724) Documentation.
[39] Ramasamy V, Fu R, Feldman D, et al. (2021). U.S. Solar PV System and Energy Storage Cost Benchmark: Q1 2021. NREL/TP-6A20-80694.
[40] Ramasamy V, Fu R, Feldman D, et al. (2023). U.S. Solar PV System and Energy Storage Cost Benchmarks: Q1 2023. NREL/TP-6A20-87303.
[41] Renewable Energy Buyers Alliance (REBA) (2024). 2024 State of the Market: U.S. Corporate Clean Energy Procurement.
[42] Seel J, Bolinger M, Gorman W, et al. (2024). Utility-Scale Solar, 2024 Edition – Slides. LBNL.
[43] Seel J, Mulvaney Kemp J, Cheyette A, et al. (2024). Utility-Scale Solar, 2024 Edition – Executive Summary. LBNL.
[44] Solar Energy Industries Association (SEIA) (2025). Solar Industry Research Data.
[45] U.S. Department of Energy (DOE) (2021). Understanding Solar Photovoltaic System Performance.
[46] U.S. Department of Energy (DOE) (2023). A Roadmap to the Future of Interconnection Innovation (i2X).
[47] U.S. DOE Solar Energy Technologies Office (SETO) (2025). Solar Photovoltaic System Cost Benchmarks.
[48] U.S. Department of Energy (DOE) Loan Programs Office (LPO) (2025). LPO Active Portfolio.
[49] Lawrence Berkeley National Laboratory (2025). PV PPA Prices – U.S. Utility-Scale Solar.
[50] Lawrence Berkeley National Laboratory (2025). CapEx, LCOE, and PPA Prices by Region – U.S. Utility-Scale Solar.
[51] IEA-PVPS Task 13 (2020). Uncertainties in Yield Assessments and PV LCOE.

[52] National Renewable Energy Laboratory (NREL) (2025). PVWatts Calculator.

[53] NREL & Sandia National Laboratories (2018). Best Practices Handbook for the Collection and Use of Solar Resource Data