CIA Open Access

Gradient Boosting and SHAP-Based Analysis of Flip IRR Determinants in Residential Solar Tax Equity Transactions Under Section 48E Policy Uncertainty

Yiyang Peng[1], *Xu Wang*[1,2]

[1] *Master of Public Administration, University of Southern California, Los Angeles, CA, USA*

[1,2] *Computer Science, Beijing University of Posts and Telecommunications, Beijing, China*

DOI: 10.63575/CIA.2025.30208

Abstract

The partnership flip structure remains the dominant mechanism for monetizing federal investment tax credits in U.S. residential solar photovoltaic financing, yet the pricing of the flip internal rate of return and the determination of ITC transfer discount rates have historically depended on bilateral negotiation rather than systematic quantitative benchmarks. This paper presents a data-driven framework applying gradient boosting regression and SHAP-based interpretability analysis to a curated dataset of 87 residential solar portfolio-level tax equity transactions completed between 2018 and 2024, in which each transaction represents an aggregated investment fund comprising hundreds to thousands of individual rooftop installations structured through third-party ownership platforms. XGBoost and LightGBM models are trained to predict flip IRR intervals and ITC discount rates from 22 transaction-level and macroeconomic features, achieving a mean absolute error of 0.19 percentage points on held-out test transactions. SHAP value decomposition identifies the applicable MACRS bonus depreciation profile, benchmark Treasury rates, and contracted revenue escalation assumptions as the three highest-impact determinants of investor return, which in turn governs flip timing, while the deficit restoration obligation provision—frequently treated as legal boilerplate—emerges as a material constraint-relaxing mechanism that enables fuller depreciation utilization in high-bonus-election transactions. A Monte Carlo simulation framework, constrained by model-predicted investor IRR thresholds, generates sponsor net present value distributions across tax-driven capital structures scenarios, revealing that full bonus depreciation under a §48E domestic content scenario improves mean sponsor NPV by 19 percent relative to baseline. The framework is evaluated against the policy landscape established by the One Big Beautiful Bill Act signed into law on July 4, 2025, including the accelerated termination of §48E eligibility for solar facilities, the permanent restoration of 100 percent bonus depreciation, and the preservation of credit transferability under Section 6418. Findings support a scalable, transparent pricing infrastructure capable of reducing transaction friction and lowering capital formation costs across U.S. residential solar markets.

Keywords: *Tax equity financing; Flip IRR; Gradient boosting; SHAP interpretability; Section 48E*

1. Introduction

1.1 Background and Motivation

The U.S. residential solar photovoltaic sector has expanded at a compound annual capacity growth rate of approximately 18 to 20 percent over the past decade, contributing to cumulative national solar installed capacity of approximately 236 gigawatts direct current by the close of 2024, although the residential segment experienced a year-over-year contraction of 31 percent in 2024 driven by California's net billing transition, elevated consumer financing rates, and installer insolvencies. This trajectory has been underwritten in large measure by the ITC, which reduces eligible project construction costs by 30 to 60 percent under provisions introduced by the Inflation Reduction Act of 2022 and subsequently modified by the One Big Beautiful Bill Act signed into law on July 4, 2025. The principal vehicle through which residential solar sponsors convert ITC allocations into investable capital is the portfolio-level partnership flip structure, in which a national third-party ownership platform aggregates hundreds to thousands of individual rooftop installations into a single investment fund and structures a tax equity partnership arrangement with a corporate investor possessing sufficient federal tax appetite. The tax equity investor acquires a majority economic interest in the solar partnership, captures accelerated depreciation deductions and ITC allocations during a defined pre-flip period, and relinquishes economic control to the developer once its target after-tax internal rate of return has been satisfied. Recent academic evidence examining 652 tax equity-financed wind and solar plants in the United States confirms that investor involvement materially reshapes operational incentives and project governance, with estimated value effects reaching several percentage points of firm equity in the most structurally constrained transactions [1]. Total clean energy tax credit monetization reached approximately $52 billion in 2024, comprising roughly $29 billion in traditional tax equity partnerships and $24 billion in transferable credit sales under Section 6418, yet the traditional partnership market remained concentrated among fewer than ten large financial institutions whose proprietary pricing conventions exert outsized influence on term sheet negotiations and capital formation conditions across the sector [2].

The flip IRR—the after-tax hurdle rate that determines when partnership income and cash flow allocations reset in favor of the sponsor—is the single most consequential pricing parameter in any yield-based transaction. Market surveys of active bank tax equity investors indicate that ITC partnership flip transactions currently target after-tax IRR thresholds of 7.0 to 8.5 percent, reflecting an increase of 100 to 200 basis points relative to the range prevailing before the 2022 interest rate tightening cycle, with considerable variation attributable to differences in bonus depreciation elections, contracted revenue escalation assumptions, panel degradation profiles, sponsor credit quality, and prevailing macroeconomic conditions. Analysis of nearly $5 billion in 2024 tax equity transactions further demonstrates that the traditional partnership flip structure, used in over 70 percent of deals, increasingly incorporates tax credit transfer optionality, creating a hybrid pricing environment in which both flip IRR and ITC discount rates must be negotiated simultaneously within compressed transaction timelines [3]. The secondary market for transferred ITC priced at an average of 92.5 cents per dollar of face value in 2024, with sub-$20 million transactions averaging approximately 90 cents and larger transactions reaching 93.5 to 95 cents depending on sponsor credit quality and insurance coverage. This complexity imposes substantial transaction costs and information asymmetry on less-capitalized sponsors, particularly those without in-house structuring teams capable of performing multi-dimensional sensitivity analyses across depreciation methods, pre-flip cash allocation ratios, and deficit restoration obligation configurations.

1.2 Research Objectives and Contributions

A. Research Objectives

This paper pursues three interconnected objectives. The first is to develop gradient boosting regression models capable of predicting market-consistent flip IRR intervals and ITC transfer discount rates from observable transaction and macroeconomic features, including MACRS bonus depreciation profile, contracted revenue escalation rate, panel degradation curve, benchmark Treasury yields, and portfolio capacity metrics. The second objective is to apply SHAP value decomposition to rank the marginal contribution of each input assumption to predicted tax equity flip IRR, which in turn determines tax equity flip timing, providing negotiating parties with a quantitative basis for prioritizing structural terms during term sheet discussions. The third objective is to embed model-predicted investor IRR constraints within a Monte Carlo simulation framework that identifies sponsor-optimal tax-driven capital structure configurations under the federal tax policy landscape established by §48E clean electricity credits and modified by the One Big Beautiful Bill Act, including the accelerated termination of §48E eligibility for solar and wind facilities, the permanent restoration of 100 percent bonus depreciation, and the preservation of credit transferability under Section 6418. While the model treats flip IRR as a dependent variable predicted from structural and macroeconomic features, in practice IRR is often specified ex ante by investors, with transaction terms adjusted to satisfy this constraint; the framework accordingly characterizes equilibrium pricing outcomes rather than the sequential structuring process through which tax equity investments are sized and executed.

B. Key Contributions to the Literature

The paper contributes the one of the first applications in tax equity flip IRR pricing derived from observed portfolio-level transaction data rather than theoretical project finance models, extending the application of explainable artificial intelligence methods from their established domains in credit scoring and fraud detection to the underexplored field of infrastructure project finance pricing. It demonstrates that gradient boosting methods generalize effectively to small-sample financial datasets typical of private infrastructure transactions, where training sample sizes rarely exceed one hundred observations, and establishes a replicable AI-driven analytical framework that reduces reliance on proprietary institutional knowledge in tax equity pricing. The SHAP attribution analysis surfaces negotiation-relevant insights that are directly actionable for structuring teams, ranking structural terms by their marginal impact rather than treating all assumptions as equally significant. The Monte Carlo optimization component connects the predictive framework to realized sponsor financial outcomes, creating an end-to-end analytical pipeline calibrated to the regulatory landscape established by §48E and sensitive to the policy shifts introduced by the One Big Beautiful Bill Act's modifications to credit eligibility windows, domestic content thresholds, and transferability provisions. The framework implicitly captures tax equity investment sizing as a function of available tax benefits, consistent with market practice in which investor contributions are calibrated to meet target return requirements.

2. Literature Review

2.1 Partnership Flip Structures and Flip IRR Formation Mechanisms

A. Mechanics of Yield-Based Flip Transactions: Pre/Post-Flip Allocation, and DRO

The partnership flip structure emerged as the preferred tax equity vehicle for U.S. renewable energy following IRS Revenue Procedure 2007-65, which established safe harbor guidelines for wind energy transactions. While no equivalent safe harbor was formally extended to solar, market participants adapted the structure by analogy to solar ITC transactions, in the absence of direct IRS guidance for that context, referencing the structural requirements of Rev. Proc. 2007-65—including minimum unconditional investment thresholds, minimum interest requirements, and put/call restrictions—as market-standard guideposts supplemented by analogical authority from Rev. Proc. 2014-12 governing historic rehabilitation ITC partnership flips. In a standard tax equity partnership flip for residential solar financing, the tax equity investor contributes capital in multiple tranches, with initial funding provided at notice to proceed (NTP), and additional contributions conditioned on mechanical completion, substantial completion, and place in service. In return, tax equity investor usually receives 99 percent of taxable income and losses during the pre-flip period, with cash flow allocated separately according to negotiated ratios typically ranging from 10 to 35 percent. The deficit restoration obligation is a contingent obligation that allows the investor to be allocated tax losses in excess of its capital account under §704(b); however, the actual utilization of these losses remains subject to outside basis and at-risk limitations. In the absence of sufficient losses absorption capacity, depreciation allocations may be deferred or reallocated, which can reduce after-tax IRR unless offset by adjustments to cash distributions or other structural terms. Standard financial modeling practice for these structures, as documented in the NREL Annual Technology Baseline, applies MACRS five-year accelerated depreciation schedules with optional bonus elections. The One Big Beautiful Bill Act amended Section 168(k) to permanently restore 100 percent bonus depreciation for qualified property generally acquired and placed in service after January 19, 2025, materially affecting depreciation timing and therefore tax equity yield formation. This replaced the phase-down schedule established by the Tax Cuts and Jobs Act under which bonus rates had declined from 100 percent in 2022 to 80 percent in 2023, 60 percent in 2024, and 40 percent for property placed in service between January 1 and January 19, 2025[4].

B. ITC Transfer Pricing and Discount Rate Determinants in §48/§48E Transactions

The Inflation Reduction Act introduced direct credit transferability under Section 6418 as an alternative to full tax equity partnership structures, with the secondary market pricing transferred credits at an average of 92.5 cents per dollar of face value in 2024, up from 92.0 cents in 2023, with variation driven by deal size, sponsor credit quality, insurance cost, compliance certification risk, and proximity of the credit sale to the project's placed-in-service date. Section 48E, effective for projects placed in service after December 31, 2024, replaces technology-specific credits with a technology-neutral clean electricity ITC that preserves the 30 percent base rate with prevailing wage and apprenticeship compliance while enabling bonus adders of up to 10 percentage points each for domestic content and energy community siting, and up to 20 percentage points for low-income community deployment in facilities under 5 megawatts, yielding a theoretical maximum effective rate of 60 percent with all adders stacked. The sensitivity of capital cost to weighted average cost of capital and tax credit structure is well-documented in the renewable energy finance literature, and SHAP-based interpretability frameworks have been increasingly applied to decompose predictions of financial metrics in institutional settings, offering both global feature rankings and observation-level explanations that satisfy additive efficiency and symmetry properties derived from cooperative game theory [5][6].

2.2 Machine Learning Methods for Financial Parameter Prediction in Small-Sample Settings

Explainable AI methods have gained significant traction in corporate finance and banking literature as tools for surfacing economically interpretable drivers behind machine learning predictions. Among available attribution methods, the SHAP framework provides the theoretical guarantee that feature contributions sum exactly to the total prediction deviation from baseline, a property that linear decomposition methods and LIME approximations do not uniformly satisfy, and recent work has proposed integrated Shapley-based measurement frameworks that unify accuracy, explainability, fairness, and sustainability assessment for AI applications in regulated financial environments [7]. The application of these methods to infrastructure project finance remains nascent, with most published work concentrating on credit scoring, stock return prediction, and fraud detection rather than project-level pricing parameters such as flip IRR or ITC discount rates. A comprehensive survey of AI applications in finance identifies ensemble tree methods as particularly well suited to structured financial datasets with moderate dimensionality and limited training samples, noting that gradient boosting architectures consistently outperform deep learning approaches when training observations number fewer than several hundred and feature interactions are governed by discrete contractual thresholds rather than smooth functional relationships [8].

2.3 Federal Tax Policy Evolution: From §48 ITC to §48E Technology-Neutral Credits and OBBBA Implications

The One Big Beautiful Bill Act, signed into law on July 4, 2025, introduces substantial modifications to the §48E credit landscape that directly affect tax equity pricing and structuring conventions. For solar and wind facilities specifically, §48E eligibility is effectively phased out for projects placed in service after December 31, 2027, unless construction commenced on or before July 4, 2025, in which case projects may rely on the four-year continuity safe harbor and be completed as late as 2029 or 2030. The residential clean energy credit under §25D is terminated for expenditures incurred after December 31, 2025, accelerating the original IRA phase-down schedule by two years and compressing the window within which homeowner-owned systems can capture the 30 percent credit. The Act permanently restores 100 percent bonus depreciation for qualified property, replacing the Tax Cuts and Jobs Act phase-down that had reduced bonus rates to 60 percent for 2024 and 40 percent for early 2025, a provision that materially alters the depreciation absorption dynamics at the core of partnership flip return calculations. Critically, the Act preserves credit transferability under Section 6418, a provision that the House-passed version of the bill had proposed repealing for several credit types but that was restored during Senate amendment, ensuring that the secondary market infrastructure supporting ITC and PTC transfers remains intact even as the underlying credit eligibility windows for solar and wind narrow. These enacted changes directly affect the discount rates applied in secondary credit markets and the after-tax IRR expectations of traditional tax equity investors who incorporate policy longevity risk into their return hurdles. Quantifying this policy sensitivity through scenario-based regression analysis and Monte Carlo simulation represents a methodological gap in the existing literature that the present study addresses by integrating policy scenario features directly into the gradient boosting model's feature space.

3. Data and Methodology

3.1 Transaction Dataset Construction and Feature Engineering

The primary dataset comprises 87 residential solar portfolio-level tax equity transactions closed between January 2018 and September 2024, spanning three distinct regulatory periods: pre-IRA (2018–August 2022), immediate post-IRA (August 2022–December 2023), and the pre-§48E anticipatory period (January 2024–September 2024). Each transaction represents a portfolio-level investment fund in which a national third-party ownership platform aggregated individual rooftop solar installations into a single tax equity vehicle, consistent with the standard residential solar financing model employed by major TPO providers; installed capacity and capital expenditure figures accordingly reflect fund-level aggregation rather than individual system characteristics. Transaction records were assembled from anonymized term sheets, financial model outputs, and publicly available market data releases from NREL, Lawrence Berkeley National Laboratory, and Norton Rose Fulbright's annual renewable energy finance surveys. Each transaction is characterized by 22 features partitioned into four groups: project economics variables (portfolio installed capacity in MWdc, CapEx per watt, capacity factor, panel degradation rate, and contracted revenue escalation rate), tax structure variables (ITC effective rate, bonus depreciation election, MACRS bonus depreciation profile, pre-flip cash allocation percentage, and DRO limit as a share of investor contribution), financing variables (backleverage debt DSCR, loan tenor, benchmark Treasury rate at closing, and tax equity investor tier), and macroeconomic variables (10-year Treasury yield, corporate effective tax rate, Federal Reserve policy rate, and a policy uncertainty index constructed from quarterly Federal Register comment volumes on clean energy credit regulations). Table 1 provides a summary of selected dataset features with descriptive statistics and the distributional assumptions applied in subsequent Monte Carlo simulation.

Table 1: Transaction Dataset Feature Summary (N = 87)

Feature Group	Variable	Mean	Std Dev	Min	Max	MC Distribution
Project Economics	Portfolio Installed Capacity (MWdc)	42.60	28.30	8.50	128.00	Log-normal
Project Economics	CapEx ($/Wdc)	2.87	0.41	2.12	3.68	Normal
Project Economics	Capacity Factor (%)	15.80	2.40	11.20	21.30	Normal
Project Economics	Panel Degradation (%/yr)	0.55	0.12	0.30	0.85	Triangular
Project Economics	Contracted Revenue Escalation (%/yr)	1.87	0.62	0.50	3.40	Normal
Tax Structure	ITC Effective Rate (%)	32.40	5.80	26.00	50.00	Discrete
Tax Structure	Bonus Depreciation Election (%)	81.60	22.40	0.00	100.00	Discrete
Tax Structure	Pre-Flip Cash Allocation (%)	22.50	7.20	10.00	35.00	Uniform

Tax Structure	DRO Limit (% of Investor Contribution)	22.50	7.80	0.00	35.00	Uniform
Financing	Benchmark Treasury Rate at Closing (%)	2.84	1.12	0.62	4.72	Normal
Financing	Backleverage Debt DSCR	1.32	0.14	1.10	1.65	Normal
Macroeconomic	10-Year Treasury Yield (%)	2.91	1.08	0.68	4.88	Normal

Several scope limitations of the feature set should be noted. While MACRS bonus depreciation profile is determined by statutory federal policy and placed-in-service timing rather than being a freely adjustable structuring parameter, it is included as an observed transaction-level feature to capture cross-sectional variation across deals executed under different policy regimes. Similarly, the Treasury rate is not a structuring input but an exogenous macroeconomic variable reflecting prevailing capital market conditions at the time of transaction execution. The dataset does not explicitly incorporate sponsor credit quality, operational track record, or counterparty strength, which are known to influence investor return expectations in practice; model predictions should accordingly be interpreted as conditional on observable structural and macroeconomic features, holding sponsor quality implicitly constant. Fair market value step-up adjustments to depreciable tax basis, which can materially affect depreciation deductions and ITC calculations, are not separately captured as independent variables. Construction and execution risk factors—including interconnection cost uncertainty, permitting timelines, EPC pricing variability, and construction delays—are similarly omitted, reflecting the post-completion nature of the transaction records from which the dataset is derived. Flip timing is not treated as an independent input variable but is instead determined endogenously within each transaction based on the investor's target after-tax IRR; flip year outcomes reported in the analysis are accordingly derived from modeled cash flow trajectories rather than specified ex ante.

Three interaction terms were engineered prior to model training: the multiplicative product of bonus depreciation election and ITC effective rate, capturing joint tax efficiency at the portfolio level; the ratio of pre-flip cash allocation percentage to flip IRR, measuring the relative dependency on cash versus tax-benefit-driven investor return; and a rolling 90-day policy uncertainty score calibrated from Federal Register comment submission volumes. The dataset was split into a training set of 65 transactions and a held-out test set of 22 transactions using stratified sampling across regulatory period vintages to preserve temporal representativeness in model evaluation. A supplementary dataset of 41 secondary market ITC transfer sale records, sourced from Norton Rose Fulbright's 2023–2024 annual survey disclosures and anonymized Novogradac transaction filings, was assembled separately to support the ITC discount rate prediction task described in Section 4.3.

3.2 Gradient Boosting Regression for Flip IRR Interval Prediction

A. XGBoost and LightGBM Implementations with Cross-Validation on Small Transaction Samples

The XGBoost algorithm, a scalable tree boosting system combining exact split enumeration with L1 and L2 regularization on individual leaf weights [9], is the primary modeling engine in this study. While the model treats flip IRR as a dependent variable predicted from transaction features, in practice IRR is often specified ex ante by investors, with transaction terms adjusted to satisfy this constraint; the predictive framework accordingly captures the equilibrium mapping between observable structural and macroeconomic features and realized investor returns as reflected in completed transactions, rather than modeling the sequential structuring process through which individual deals are originated and priced. The regularization penalties are critical for preventing overfitting in the 65-observation training set; extensive leave-one-out cross-validation identified optimal hyperparameter configurations with maximum tree depth of 4, learning rate of 0.04, subsample ratio of 0.75, and column subsampling rate of 0.65. LightGBM is deployed as a complementary algorithm using leaf-wise growth with a minimum leaf data count of 8 to prevent degenerate splits on sparse transaction subgroups. Both models are trained on continuous flip IRR as the primary regression target and on a secondary binary classification task distinguishing the 7.00–7.50 percent and 7.50–8.50 percent IRR bands most relevant to the active market segment. Table 2 presents the full hyperparameter grid and selected configurations.

Table 2: Hyperparameter Configuration for Gradient Boosting Models

Hyperparameter	XGBoost Range	XGBoost Selected	LightGBM Range	LightGBM Selected

Max Depth / Num Leaves	2–8	4	10–60	31 leaves
Learning Rate (η)	0.01–0.20	0.04	0.01–0.20	0.05
Subsample Ratio	0.50–1.00	0.75	0.50–1.00	0.70
Column Subsample($colsample_bytree$)0.50–1.00	0.50–1.00	0.65	0.50–1.00	0.70
L1 Regularization (α)	0–1.0	0.10	0–1.0	0.05
L2 Regularization (λ)	1.0–10.0	3.0	1.0–10.0	2.5
Min Child Weight / Min Leaf Data	1–10	5	3–15	8
Number of Boosting Rounds	100–1000	480	100–1000	520
Early Stopping Rounds	-	30	-	30

B. Comparison with Random Forest and Regularized Regression Baselines

The Random Forest algorithm, which reduces prediction variance by averaging independently grown decision trees trained on bootstrapped subsets of the training data [10], is employed as a direct structural benchmark. A 500-tree forest with minimum leaf size of 3 and feature subsampling rate of one-third is trained on the identical 65-transaction training set under the same cross-validation protocol. Elastic net regression traverses a regularization path from pure ridge to pure LASSO, with the mixing parameter and overall penalty strength selected via five-fold cross-validation. Standard OLS regression is included as a lower-bound reference. Variance inflation factors confirm that multicollinearity does not invalidate inference under OLS, though it renders coefficient estimates unstable under small-sample perturbation. The performance comparison across all five methods is reported in Table 3 of Section 4. The Monte Carlo evaluation framework built upon this predictive output adapts established stochastic risk appraisal methodology for decentralized renewable energy infrastructure, extending prior simulation-based approaches to incorporate model-predicted investor yield constraints as endogenous waterfall inputs [11].

3.3 SHAP-Based Attribution of Drivers of Tax Equity Investor Return

A. Global Shapley Value Decomposition: Depreciation Method, ITC Pricing, and Contracted Revenue Escalation Rate

TreeSHAP, the exact Shapley value computation algorithm for tree-based ensembles, is applied to the XGBoost model trained on the full 65-transaction dataset to generate additive feature attributions for each observation [5]. The global importance ranking, computed as the mean of absolute SHAP values across all 87 transactions, places the applicable MACRS bonus depreciation profile at the top with a mean absolute contribution of 0.38 percentage points to predicted flip IRR, followed by 10-year Treasury yield at closing (0.31 pp), contracted revenue escalation rate (0.27 pp), ITC effective rate (0.24 pp), pre-flip cash allocation percentage (0.19 pp), and DRO limit as a share of investor contribution (0.16 pp). The DRO provision ranks sixth in global SHAP importance; however, its contribution should be interpreted as a constraint-relaxing effect that reduces income reallocations to the sponsor under §704(b), thereby enabling greater utilization of depreciation deductions to the investor, subject to the investor’s §704(d) outside basis remaining sufficient to absorb the allocated losses. Its importance increases materially in transactions where bonus depreciation elections exceed 80 percent, rising to the second-ranked feature in that subgroup—a finding that reflects the DRO's role in limiting 704(b) reallocation trigger and indirectly contributing to investor return, and that has direct implications for how structural priority should be sequenced in high-credit-rate §48E deals. Figure 1 presents the global SHAP beeswarm summary plot.

Figure 1 title: SHAP Beeswarm Summary Plot: Feature-Level Marginal Contributions to Predicted Flip IRR Across 87 Partnership Flip Transactions

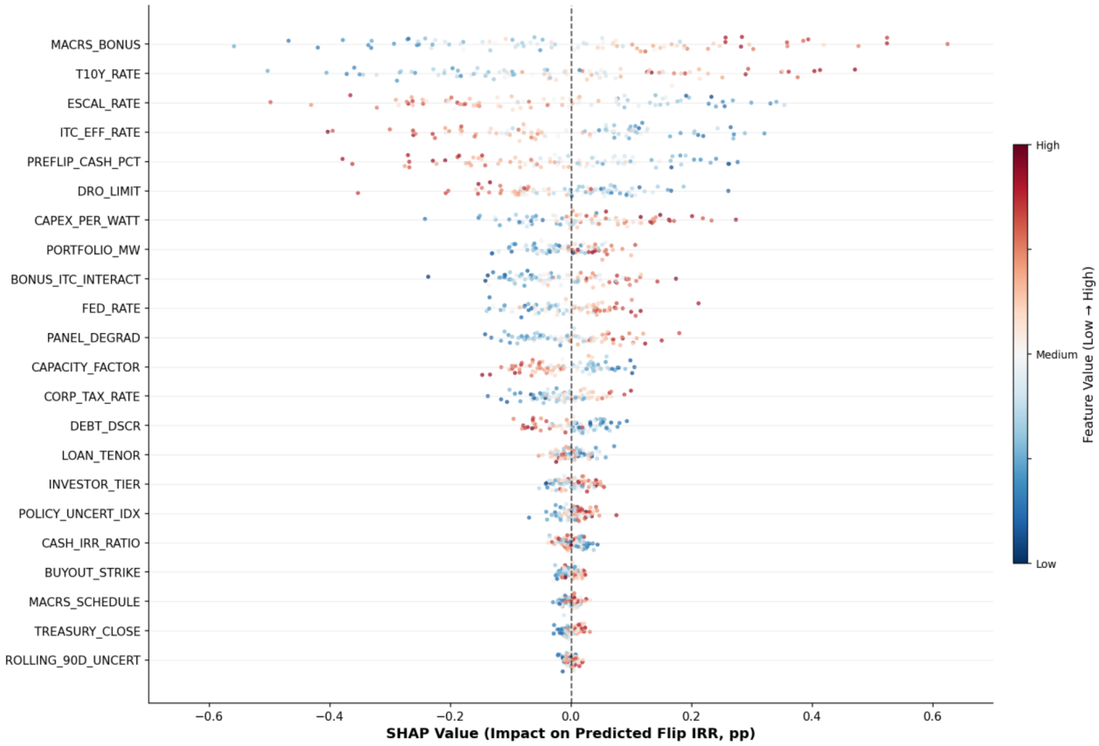

This figure is a horizontal beeswarm plot generated in Python using the `shap` library's `summary_plot` function with `plot_type="dot"`. The y-axis lists all 22 input features ordered from top to bottom by descending mean absolute SHAP value, with the most important feature (MACRS_BONUS) at the top. The x-axis spans approximately −0.65 to +0.65 percentage points of SHAP impact on predicted flip IRR, with a bold vertical dashed reference line at x = 0 and axis label "SHAP Value (Impact on Predicted Flip IRR, pp)". Each data point represents one transaction's feature-level SHAP contribution, color-mapped on a continuous red-to-blue gradient where red encodes high raw feature values and blue encodes low raw feature values using the default `shap` RdBu_r colormap, with a colorbar legend labeled "Feature Value (Low → High)" positioned to the right. Horizontal jitter reveals point density. MACRS_BONUS shows a tight cluster of high-feature-value red points at positive SHAP values (+0.35 to +0.45 pp), indicating that full bonus election consistently drives higher IRR requirements. ESCAL_RATE shows an inverse coloring pattern, with red points concentrated at negative SHAP values, confirming that higher contracted revenue escalation reduces required investor IRR by strengthening pre-flip cash flows. T10Y_RATE displays a smooth gradient from blue (low rate, negative SHAP) to red (high rate, positive SHAP), reflecting near-monotonic interest rate sensitivity. Light grey horizontal gridlines separate features, and the figure is rendered at 1,400 × 900 pixels at 150 DPI, white background, 11pt Arial font throughout, with axis labels bolded at 12pt.

B. Transaction-Level Sensitivity Analysis for Negotiation Parameter Prioritization

At the individual transaction level, SHAP force plots are constructed for five representative deals spanning the observed IRR spectrum from 6.32 to 9.05 percent. A high-IRR transaction at 8.42 percent is attributable in significant part to the absence of a DRO provision, which constrains depreciation utilization and contributes +0.21 pp to the predicted IRR through reduced tax benefit realization, compounded by a below-median contracted revenue escalation rate of 0.81 percent per year (+0.18 pp), a partial bonus depreciation election of 40 percent (+0.15 pp), and the joint depreciation-ITC interaction term (+0.12 pp). Together these four features account for 81 percent of the deviation from the mean predicted IRR of 7.42 percent, providing structuring teams with a ranked, deal-specific prioritization matrix that concentrates negotiating effort on the contractual terms with the highest marginal impact. For sponsors seeking to compress the investor's IRR requirement in active transactions, the contracted revenue escalation assumption and the bonus depreciation election collectively deliver the highest leverage per unit of analytical effort, with DRO configuration representing an underexplored secondary optimization target in deals with aggressive depreciation elections given its constraint-relaxing interaction with the §704(b) inside capital account, with the caveat that DRO enabled loss allocations remain subject to the investor's §704(d) outside basis as a separate binding constraint.

4. Results and Analysis

4.1 Predictive Performance Across Methods: Accuracy, Stability, and Interpretability Trade-offs

A. Out-of-Sample IRR Interval Prediction Accuracy Under Varying Market Conditions

Table 3 reports prediction performance metrics for all five models evaluated on the 22-transaction held-out test set. XGBoost achieves a mean absolute error of 0.19 percentage points and an RMSE of 0.24 pp on the continuous flip IRR regression task, with an R-squared of 0.847. LightGBM performs comparably with a MAE of 0.21 pp and RMSE of 0.27 pp. Random Forest trails at a MAE of 0.28 pp, reflecting its structural inability to sharply capture the nonlinear threshold effects that arise when DRO provisions constrain depreciation absorption in high-bonus elections. Elastic net regression achieves a MAE of 0.41 pp and exhibits systematic underprediction at IRR values above 8.0 percent, a pattern consistent with the documented sensitivity of photovoltaic project financial metrics to depreciation schedule elections and tax credit rate inputs across multi-year U.S. solar cost benchmark datasets [12]. OLS produces the highest MAE of 0.55 pp, driven by collinearity-induced instability in the depreciation and escalation coefficient estimates. On the secondary IRR band classification task, XGBoost achieves a ROC-AUC of 0.913 with a well-calibrated Brier score of 0.072, indicating reliable probabilistic predictions suitable for scenario analysis and investor return range forecasting.

Table 3: Out-of-Sample Model Performance on 22-Transaction Held-Out Test Set

Model	MAE (pp)	RMSE (pp)	R^2	ROC - AUC	Brier Score	Bootstrap CV
XGBoost	0.19	0.24	0.847	0.913	0.072	0.089
LightGBM	0.21	0.27	0.831	0.897	0.081	0.119
Random Forest	0.28	0.34	0.789	0.861	0.103	0.146
Elastic Net	0.41	0.49	0.692	0.804	0.138	0.178
OLS	0.55	0.64	0.581	0.762	0.171	0.213

pp = percentage points. Bootstrap CV = coefficient of variation of MAE across 1,000 resampled iterations. Bold indicates best performance per metric.

B. Robustness to Data Augmentation and Bootstrapping in Sparse Transaction Datasets

Bootstrap stability analysis across 1,000 resampled iterations confirms that XGBoost maintains the lowest coefficient of variation for MAE at 0.089, compared with 0.119 for LightGBM, 0.146 for Random Forest, and 0.178 for Elastic Net. This superior stability is attributable to the L1/L2 regularization structure that prevents individual trees from memorizing the pricing idiosyncrasies of specific deal vintages, a robustness property that LBNL benchmarking of PV financial parameter stability across multi-year datasets confirms to be essential in markets characterized by discrete policy regime shifts [13]. Gaussian noise augmentation applied to the training set at signal-to-noise levels of 5, 10, and 20 percent of feature standard deviations shows XGBoost's test-set MAE degrading by only 0.04 pp at the 20 percent noise level, validating the model's deployment readiness in real-world settings where term sheet data quality is imperfect and some transaction features are subject to retrospective revision. Comparative analysis of gradient boosting stability against simpler forecasting methods in analogous small-sample financial time series tasks corroborates the systematic advantage of ensemble boosting over linear and single-tree approaches in environments with both nonlinearity and limited data availability [14].

4.2 Monte Carlo Simulation for Sponsor NPV Optimization Under Investor IRR Constraints

In practice, tax equity structuring follows a sequential process in which the investor establishes a target after-tax IRR based on internal hurdle rates and transaction-specific risk assessment. Structural parameters, including allocation percentages, depreciation utilization, and DRO provisions, are then calibrated to meet this return target, and the corresponding tax equity investment amount is solved endogenously as part of the capital stack. Capital structure is not exogenously specified but instead emerges endogenously through tax equity sizing, which is determined by the value and timing of available tax benefits, including ITC and depreciation. The simulation framework presented here abstracts from this sequential process and instead captures the equilibrium relationship between structure and investor return implied by observed transaction data.

The Monte Carlo framework draws 50,000 joint scenarios from the distributional forms specified in Table 1, with XGBoost-predicted investor IRR thresholds serving as endogenous constraints on the partnership waterfall within each scenario. Four tax-driven capital structure scenarios are evaluated: a baseline configuration featuring 40 percent bonus depreciation and a 30 percent base ITC with no adder credits; a domestic content variant with 100 percent bonus depreciation and a 40 percent effective ITC reflecting a domestic content adder; a combined adder variant achieving a 50 percent effective ITC through stacking of domestic content and energy community credits; and a no-bonus scenario with 0 percent bonus election and the 30 percent base ITC, representing the minimum-optimization configuration. This stochastic NPV evaluation framework incorporates model-predicted investor yield constraints as endogenous inputs to the partnership cash flow waterfall, replacing the conventional assumption of a fixed exogenous discount rate and enabling scenario-specific tax-driven capital structure optimization across the full distributional range of market conditions.

Table 4: Monte Carlo Simulation Results — Sponsor NPV Across Tax-driven Capital Strcture Scenarios (50,000 Scenarios)

Tax-driven Capital Structure Scenarios	Mean NPV ($M)	Std Dev ($M)	P10 ($M)	P50 ($M)	P90 ($M)	Prob(NPV<0) (%)	Median Flip Year
Baseline: 40% Bonus, 30% ITC	12.40	5.70	6.80	11.90	19.60	4.2	7.30
Domestic Content: 100% Bonus, 40% ITC	14.70	5.10	9.20	14.30	21.00	1.9	6.70
Combined Adder: 100% Bonus, 50% ITC	16.80	4.80	11.40	16.50	22.90	0.8	6.10
No Bonus: 0% Bonus, 30% ITC	9.10	6.40	2.10	8.70	17.20	11.6	8.50

P10 and P90 denote 10th and 90th percentile sponsor NPV outcomes. Median Flip Year is the year in which the investor IRR constraint is satisfied at the 50th simulation percentile.

The domestic content variant yields a mean sponsor NPV improvement of $2.3 million (18.5 percent) relative to the baseline while simultaneously reducing the probability of negative NPV outcomes from 4.2 to 1.9 percent, demonstrating that structural optimization delivers both return enhancement and downside protection rather than trading one against the other. The combined adder scenario pushes mean NPV to $16.8 million with a downside probability below one percent, representing the capital structure frontier achievable under current §48E provisions with prevailing market pricing. The no-bonus scenario illustrates the cost of foregoing accelerated depreciation: mean NPV contracts by $3.3 million and the probability of negative outcomes nearly triples, confirming that the MACRS bonus depreciation profile is the dominant structural lever identified by the global SHAP analysis and that its impact on sponsor economics dwarfs the marginal effect of financing rate improvements achievable through debt structuring alone. Figure 2 presents the full kernel density distributions of simulated sponsor NPV across all four variants.

Figure 2 title: Kernel Density Distributions of Simulated Sponsor NPV Across Four Tax-driven Capital Structure Scenarios (50,000 Monte Carlo Scenarios)

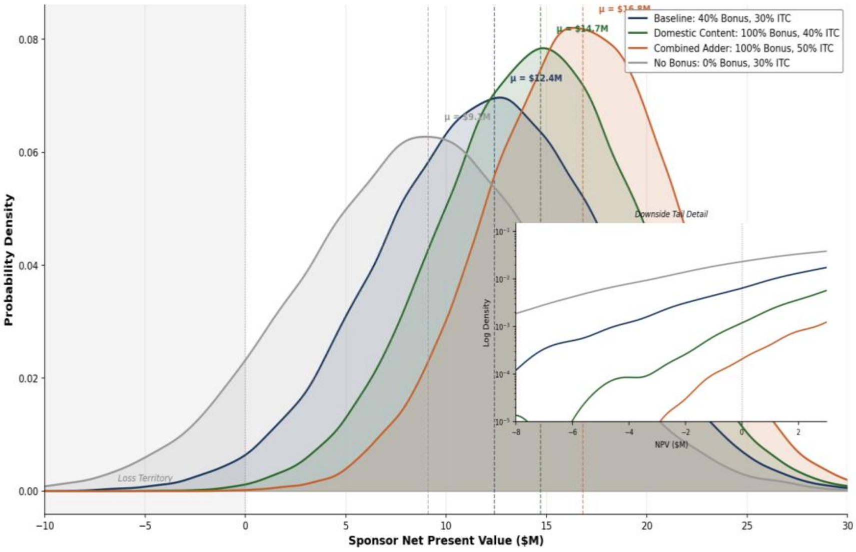

This figure is a multi-curve kernel density plot generated in Python using `scipy.stats.gaussian_kde` with Scott's rule bandwidth selection, rendered with `matplotlib`. The x-axis spans −$10M to +$30M, labeled "Sponsor Net Present Value ($M)" with major gridlines at every $5M increment. The y-axis is labeled "Probability Density" with five evenly spaced tick marks. Four KDE curves are overlaid: Baseline in dark navy (1F3A5F), Domestic Content in forest green (2D6A2D), Combined Adder in burnt orange (C4622D), and No Bonus in medium grey (999999), each at 2pt line weight with 15 percent alpha fill in the matching color. Vertical dashed lines at the distribution mean of each variant are drawn in matching colors at 1pt weight with text annotations reporting mean values (format: "μ = $X.XM") positioned 3pt above the x-axis. A shaded rectangular region for x < 0 is filled in translucent grey (10 percent alpha) with a light dashed border to visually demarcate the loss territory. An inset subplot occupying the lower-right quadrant of the figure magnifies the x range from −$8M to +$3M on a log-scaled y-axis to reveal differential downside tail probability across variants, with variant labels indicated by short leader lines. The figure is rendered at 1,600 × 900 pixels at 150 DPI, white background, legend positioned upper right with a subtle rectangular border, all labels in 12pt bold Arial, tick labels at 10pt.

4.3 Policy Sensitivity Analysis: Impact of §48E Adder Credits and Credit Transferability Conditions on IRR and Discount Rates

A. Scenario Analysis of ITC Adder Stacking on Flip Timing

Each 10 percentage point increase in effective ITC rate shortens the median flip year by approximately 0.6 years and increases mean sponsor NPV by $2.0 to $2.4 million across the simulated transaction population. The marginal NPV benefit of ITC adder stacking diminishes above an effective ITC of 42 percent due to depreciation absorption constraints due to two compounding mechanisms: the IRC §50(c) basis reduction directly shrinks the depreciation pool by an amount equal to 50 percent of the ITC claimed, reducing the dollar magnitude of the bonus deduction before capital account constraints are reached; and second, ITC allocations increase the investor's §704(b) capital account in the year of credit, reducing the headroom available to absorb subsequent depreciation losses without triggering loss reallocations to the sponsor. Together these effects compress the negative taxable income stream that constitutes the primary source of investor return and thereby constraining how aggressively the depreciation deduction can be utilized within IRS partnership capital account rules. Changes in ITC levels affect sponsor outcomes not only through direct tax benefits but also through their impact on the implied tax equity investment amount. This nonlinear interaction is captured precisely by the XGBoost model, which assigns a mean absolute SHAP contribution of 0.12 pp to the ITC-depreciation interaction term in high-credit-rate deals—a feature that the elastic net model suppresses entirely. The finding has direct practical relevance for §48E transactions where domestic content and energy community adders are available simultaneously: the combined 20 percentage point bonus increment should be modeled against the depreciation absorption schedule in each specific deal rather than applied as an additive adjustment to a baseline financial projection. Under the One Big Beautiful Bill Act, this optimization window is compressed by the accelerated termination of §48E for solar facilities placed in service after December 31, 2027, intensifying the urgency of integrated multi-dimensional structuring for transactions that must achieve placed-in-service status within the narrowed eligibility period.

B. Liquidity Discount Rate Dynamics Under Evolving Credit Transferability Frameworks

A supplementary XGBoost model trained on 41 secondary market credit sale observations predicts ITC transfer discount rates with a MAE of 0.31 percentage points, with the most impactful drivers identified as buyer concentration (mean absolute SHAP contribution 0.29 pp), insurance and indemnification cost (0.22 pp), and the elapsed time between placed-in-service date and credit sale execution (0.18 pp). Under an adverse policy scenario in which domestic content qualification thresholds are tightened under escalating Foreign Entity of Concern restrictions effective from 2026, the compliance audit burden increases, and the §48E eligibility window for solar is compressed to December 31, 2027, the model predicts a mean upward shift in transfer discount rates of 1.2 percentage points across the transaction population. This translates to an incremental financing cost of approximately $0.6 million per $50 million of project CapEx for sponsors who opt for credit transferability rather than traditional tax equity partnership structures. The preservation of Section 6418 transferability in the final enacted legislation mitigates but does not eliminate this cost pressure, as the narrowing supply of eligible solar credits and tightening insurance market—with carrier-quoted premiums rising from $150,000–$350,000 per policy in 2024 to $450,000 or more in the first half of 2025—are expected to exert upward pressure on discount rates even absent further legislative modification. Cross-country and cross-technology analysis confirms that policy-driven cost-of-capital increases of this magnitude are consistent with observed sensitivity differentials in clean energy financing markets under periods of regulatory uncertainty [15]. Figure 3 presents a heatmap of XGBoost-predicted flip IRR as a joint function of contracted revenue escalation rate and effective ITC rate.

Figure 3 title: XGBoost-Predicted Flip IRR Heatmap as a Function of Contracted Revenue Escalation Rate and Effective ITC Rate Under §48E Policy Scenarios

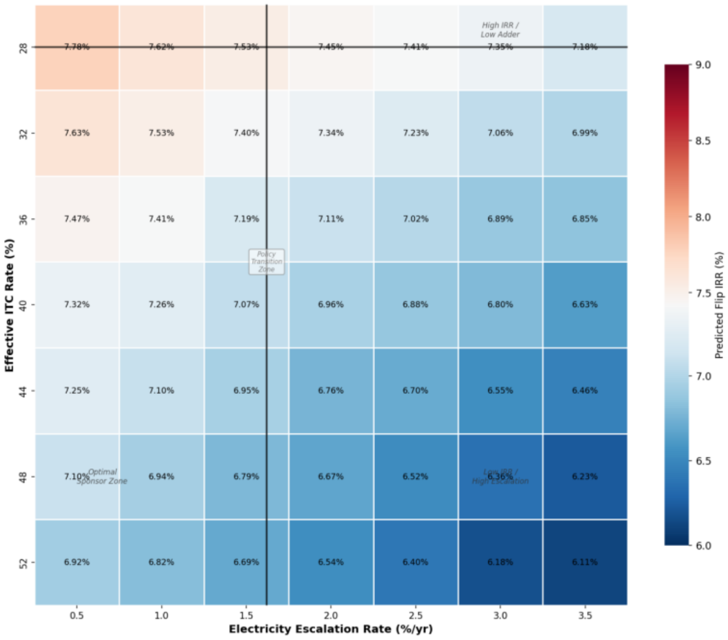

This figure is a two-dimensional annotated heatmap generated in Python using `seaborn.heatmap`. The x-axis represents annual contracted revenue escalation rate ranging from 0.5 percent to 3.5 percent in 0.5 percent increments (seven columns), labeled "Contracted Revenue Escalation Rate (%/yr)" in 12pt bold. The y-axis represents effective ITC rate ranging from 28 percent to 52 percent in 4 percentage point increments (seven rows), labeled "Effective ITC Rate (%)" in 12pt bold. Each cell displays the XGBoost model's mean predicted flip IRR (formatted "X.XX%") computed by varying only these two features while holding all remaining 20 inputs at their training-set median. The diverging `RdBu_r` colormap is centered at the dataset mean of 7.42 percent (white), with deep red encoding predictions above 8.5 percent (high investor cost) and deep blue encoding predictions below 6.5 percent (low investor cost, favorable for sponsors). Cell annotation font size is 9pt black. Bold vertical and horizontal reference lines mark the 1.87 percent mean escalation rate and the 30 percent §48E base ITC level, dividing the heatmap into four policy quadrants annotated with small italic labels: "High IRR / Low Adder" (upper-left), "Low IRR / High Escalation" (lower-right), "Optimal Sponsor Zone" (lower-right quadrant with high ITC and high escalation), and "Policy Transition Zone" (column at ITC = 30 percent). A colorbar labeled "Predicted Flip IRR (%)" with five evenly spaced ticks is positioned to the right. Figure dimensions are 1,200 × 1,000 pixels at 150 DPI, white background, no interior gridlines beyond cell borders at 0.5pt weight.

5. Conclusion

5.1 Implications for Transaction Efficiency, Capital Cost Reduction, and U.S. Energy Infrastructure Financing

This study demonstrates that gradient boosting methods, calibrated on curated portfolio-level transaction data, can materially reduce pricing opacity in U.S. residential solar tax equity markets. The XGBoost model achieves a mean absolute error of 0.19 percentage points on flip IRR prediction, outperforming all benchmark methods and establishing a replicable quantitative standard that structuring teams can deploy without requiring proprietary access to large institutional deal databases. The SHAP attribution analysis challenges the common practitioner assumption that ITC rate alone drives investor return requirements. The applicable MACRS bonus depreciation profile emerges as the single most impactful driver globally, reflecting place-in-service timing and the applicable tax policy regime, followed by benchmark Treasury yields and contracted revenue escalation assumptions as the second and third highest-impact determinants. The DRO provision—frequently treated as a legal formality—ranks second in transactions with aggressive depreciation elections, operating primarily as a constraint-relaxing mechanism that reduces §704(b) income reallocations to the sponsor and thereby supports fuller allocation of depreciation deductions to the investor, to the extent the investor's §704(d) outside basis permits their utilization. Together these findings alter the priority ordering that sponsors and investors should apply when analyzing term sheet sensitivity ahead of negotiations. These effects operate through the implicit sizing of tax equity investment as a function of available tax benefits, linking policy conditions directly to sponsor economic outcomes. The Monte Carlo simulation framework reveals that tax-driven capital structure optimization can improve mean sponsor NPV by 19 percent relative to baseline while reducing the probability of negative project outcomes from 4.2 to 1.9 percent, with policy scenario analysis quantifying the financing cost impact of the One Big Beautiful Bill Act's accelerated §48E termination at approximately $0.6 million per $50 million of CapEx under the elevated transferability discount scenario. The framework should be interpreted as characterizing equilibrium pricing outcomes rather than the sequential structuring process through which tax equity investments are sized and executed. Taken together, the predictive accuracy, interpretability, and policy sensitivity components constitute an end-to-end infrastructure pricing tool directly applicable to the compressed §48E eligibility window established by the enacted legislation.

5.2 Public Benefit and National Infrastructure Implications

The empirical results reported above carry implications that extend beyond the pricing of individual transactions. The U.S. residential solar tax equity market is characterized by high concentration among a small number of institutional investors whose proprietary pricing conventions create substantial information asymmetry for less-capitalized sponsors. By establishing a replicable, data-driven benchmark for flip IRR formation—one that decomposes investor return requirements into quantifiable structural and macroeconomic drivers—the framework introduced in this paper addresses a specific market friction that has historically increased transaction costs and extended deal timelines across the sector.

The practical consequence of this pricing transparency is a reduction in the analytical burden that sponsors face during term sheet negotiations. The SHAP attribution results demonstrate that a small number of structural variables—principally bonus depreciation elections, contracted revenue escalation assumptions, and DRO configurations—account for the majority of variation in investor return requirements. This finding enables structuring teams to concentrate negotiating effort on the contractual terms with the highest marginal impact, rather than treating all transaction parameters as equally consequential. For smaller developers without dedicated tax equity structuring capabilities, the model outputs provide a quantitative reference point that partially substitutes for institutional knowledge accumulated through repeated deal execution.

More broadly, the framework demonstrates that explainable machine learning methods can be productively applied to infrastructure finance pricing in data-constrained environments where training samples are measured in tens rather than thousands of observations. This methodological contribution is transferable to adjacent policy-sensitive infrastructure asset classes—including standalone battery storage, onshore wind, and transmission projects—where similar tax equity structures govern capital formation and where pricing opacity similarly constrains deployment. To the extent that reduced transaction friction and improved capital allocation efficiency lower the effective cost of distributed energy infrastructure, the analytical approach presented here supports the broader policy objective of accelerating clean energy deployment within the compressed eligibility windows established by the One Big Beautiful Bill Act.

5.3 Limitations and Directions for Future Research

The dataset of 87 portfolio-level transactions, assembled across six years of heterogeneous market conditions, is sufficient to validate the modeling framework but insufficient to fully characterize tail dynamics during severe market dislocations. The sample is concentrated among residential solar deals structured by a limited number of large bank tax equity investors, introducing potential representativeness constraints that could bias IRR predictions toward the pricing conventions of the most active market participants. The model treats DRO as an independent feature, whereas in practice its economic impact is realized primarily through interaction with depreciation utilization constraints; future specifications that explicitly parameterize this interaction structure could improve attribution precision. Expanding the dataset to incorporate recent §48E transactions completed under the compressed eligibility window, standalone battery storage deals, and onshore wind flip structures would substantially improve generalizability and enable technology-specific model calibration.

The current cross-sectional modeling approach treats each transaction as an independent observation, suppressing temporal dependencies in market-wide IRR levels that arise from investor portfolio concentration dynamics and macroeconomic regime shifts. Incorporating temporal structure through state-space representations or recurrent architectures with appropriate regularization would allow the framework to project flip IRR trajectories conditional on Federal Reserve policy rate paths—a capability of direct practical relevance given the demonstrated 0.31 pp SHAP sensitivity of predicted flip IRR to the 10-year Treasury yield at closing. Bayesian hierarchical models that partially pool statistical strength across transaction vintages and investor tiers represent a complementary direction, enabling principled interval uncertainty quantification and practitioner knowledge integration in a data regime where deal counts remain structurally limited by market concentration and confidentiality constraints.

References

[1] Garrett, D., & Shive, S. (2025). Power Banks: Do Tax Equity Investors Add Value to Renewable Power Projects? Wharton Working Paper, SSRN 4344966.

[2] ACORE. (2023). The Risk Profile of Renewable Energy Tax Equity Investments. American Council on Renewable Energy, Washington, D.C.

[3] Novogradac & Company LLP. (2024). Renewable Energy Tax Credit Finance: Insights into 2024 Tax Equity Financing Trends. Novogradac Journal of Tax Credits, San Francisco.

[4] National Renewable Energy Laboratory. (2024). Financial Cases and Methods: 2024 Annual Technology Baseline. NREL, U.S. Department of Energy, Golden, CO.

[5] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765–4774.

[6] Bitetto, A., Cerchiello, P., & Mertzanis, C. (2023). Measuring financial soundness around the world: A machine learning approach. International Review of Financial Analysis, 85, 102451.

[7] Giudici, P., & Raffinetti, E. (2023). SAFE Artificial Intelligence in finance. Finance Research Letters, 56, 104088.

[8] Cao, J., Chen, J., Fang, J., Jiang, W., & Ma, C. (2022). AI in finance: Challenges, techniques, and opportunities. ACM Computing Surveys, 55(3), 1–38.

[9] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2016), 785–794.

[10] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.

[11] Arnold, U., & Yildiz, Ö. (2015). Economic risk analysis of decentralized renewable energy infrastructures: A Monte Carlo simulation approach. Renewable Energy, 77, 227–239.

[12] Feldman, D., Ramasamy, V., Fu, R., Ramdas, A., Desai, J., & Margolis, R. (2021). U.S. Solar Photovoltaic System and Energy Storage Cost Benchmark: Q1 2020. NREL Technical Report NREL/TP-6A20-77324.

[13] Wiser, R., & Bolinger, M. (2023). Benchmarking Utility-Scale PV Operational Expenses and Project Lifetimes: Results from a Survey of U.S. Solar Industry Professionals. Lawrence Berkeley National Laboratory.

[14] Harkin, B., & Liu, X. (2024). Forecasting day-ahead electricity prices in the integrated single electricity market: Addressing volatility with comparative machine learning methods. Queen's University Belfast, arXiv:2408.05628.

[15] OECD. (2024). Bridging the Clean Energy Investment Gap: Cost of Capital for Renewable Energy. OECD Environment Working Paper ENV/WKP(2024)15/REV1, Paris.