

Article

From Cooperative Attribution to Dispute-Aware Settlement in Integrated Charging–Swapping V2G Stations

Wendi Wang ¹, Shufan Wang ¹, Yong Wang ², Jie Tan ², Liuting Gu ¹ and Yumin Yao ^{3,*}

¹ Nanjing Suyi Industrial Co., Ltd., Nanjing 210008, China; wwdseu@163.com (W.W.); 13645151376@163.com (S.W.); gult@js.sgcc.com.cn (L.G.)

² Nanjing Huaqun Energy Group Co., Ltd., Nanjing 210008, China; njgdjwy@js.sgcc.com.cn (Y.W.); tanj@js.sgcc.com.cn (J.T.)

³ School of Computer Science and Engineering, Central South University, Changsha 410083, China

* Correspondence: yaoyumin@csu.edu.cn

Abstract

Fair allocation in an integrated charging–swapping station depends on both how an event defines value and how evidentiary uncertainty propagates into settlement. Existing vehicle-to-grid (V2G) studies commonly treat contribution attribution and trusted transaction recording as separate problems. We formulate an event-conditioned cooperative game in which service requirements, available capabilities, and a service prior shape the characteristic function before Shapley attribution. An evidence–confidence vector further links measured performance to attributable performance. A calculation–provenance certificate and an evidence–dependency matrix then connect attributed amounts to component-level settlement states under dispute. Under the stated budget and dependency conditions, the formulation establishes conservation and dispute locality. In reproducible synthetic cases, event conditioning produced distinct allocations across three service types and a smaller normalized core deficit than the equal-weight game in 86.4% of sampled events. A position-balanced estimator reduced Shapley RMSE by factors of 2.6–3.4 relative to random sampling and 1.6–1.9 relative to antithetic sampling, although position-stratified sampling remained more accurate. Missing dependency edges released unsupported components, whereas spurious edges froze otherwise releasable amounts. The framework therefore connects event-specific cooperative attribution with evidence-dependent settlement and identifies dependency specification as a first-class condition of dispute-aware V2G revenue allocation.

Keywords: vehicle-to-grid (V2G); Shapley value; revenue allocation; charging station; battery swapping; blockchain; smart contract; trustworthy computing; dispute-aware settlement



Academic Editor: Michael Short

Received: 10 September 2026

Revised: 3 October 2026

Accepted: 6 October 2026

Published: 15 October 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Vehicle-to-grid (V2G) technology turns electric vehicles (EVs) from passive loads into controllable resources for peak shaving, renewable-energy integration, balancing, and other flexibility services [1]. Recent field and data-driven studies show that grid-interactive charging can operate under network constraints and at urban scale [2,3]. The economic value of this flexibility nevertheless depends on more than delivered energy. Mobility constraints, battery costs, response speed, availability, tracking accuracy, and service reliability all shape the value perceived by EV owners, aggregators, and grid-side actors [4–7].

This attribution problem becomes sharper in an integrated charging–swapping station. A single station may coordinate photovoltaic generation, a stationary battery energy storage system (BESS), bidirectional V2G fleets, smart charging loads, and a battery swapping station (BSS). These resources provide different combinations of power, energy, timeliness, and tracking performance. Co-optimization can determine how they should be dispatched [8,9], but it does not determine how the benefit created by a completed event should be attributed among them.

Cooperative game theory offers a principled answer to the attribution problem. The Shapley value assigns each participant its expected marginal contribution over all coalition-entry orders [10]. It has been applied to V2G carbon benefits [11], energy communities [12], fairness-aware aggregation [13], integrated energy systems [14], and charging or swapping incentives [15,16]. These studies clarify how a given coalition value can be distributed. They do not, however, remove a prior modeling question: whether the characteristic function represents the technical objective of the event. A Shapley allocation may satisfy its axioms while rewarding the wrong capability if an energy-centric value function is used for a fast-response event, or a response-centric value function is used for sustained energy support.

A second question arises after attribution. Blockchain and smart-contract studies have improved the traceability of V2G trading, charging records, and scheduling decisions [17–19]. Permissioned ledgers can make recorded states difficult to alter [20,21], and recent platforms have demonstrated blockchain-anchored flexibility markets and auditable proof-of-charge receipts [22,23]. Yet an immutable record does not establish which data snapshot, metric definition, algorithm version, and execution state produced an attributed amount. It also does not determine which payment components remain supported when only part of the evidence is disputed. Attribution and transaction recording therefore leave a common problem unresolved: how the evidentiary validity of an allocation should propagate into the state of settlement.

The central argument of this paper is that fair post-event settlement is not a property of the allocation rule alone. It depends jointly on how event-specific value is defined and how evidentiary uncertainty is mapped to releasable or frozen payment components. We address this problem through a unified event-level chain from cooperative valuation to dispute-aware settlement.

The methodological contribution is defined explicitly as follows:

- **Event-conditioned valuation.** Service requirements, available capabilities, and a service-semantics prior determine the characteristic function before Shapley attribution; evidence-adjusted service vectors make measurement support part of attributable performance.
- **Evidence-dependent settlement.** A binary dependency relation maps disputed evidence to component-level freeze states. Under non-negative component amounts and an explicit budget constraint, the formulation proves conservation and dispute locality, while versioned releases provide idempotent and append-only settlement updates.
- **Auditable computation.** A position-balanced Shapley estimator and a structured provenance certificate connect the numerical result to its inputs, algorithm, execution state, and output.

The independent theoretical innovation is therefore the composition of these operators into one event-level settlement map: the event specification changes the cooperative game, evidence confidence changes the effective service before attribution, and the same evidence state is propagated through a component dependency algebra after attribution. This gives a formal bridge between valuation and dispute handling that is not supplied by a standalone Shapley rule, a blockchain record, or a payment ledger. The numerical study examines

how these constructions alter allocation, approximation error, and settlement behavior in reproducible synthetic cases.

Positioning relative to the related literature. Table 1 positions the proposed formulation against closely related V2G and integrated-energy allocation studies. The comparison follows four analytical dimensions: heterogeneous resources at a common allocation level, adaptation to the service event, binding to an evidence record, and component-level dispute propagation. Prior work has developed important parts of this chain, but these parts are usually studied as allocation, transaction integrity, or market operation problems. The present study examines their connection within a single event-level attribution-to-settlement formulation.

Table 1. Positioning of the proposed framework against related V2G and integrated-energy allocation work; the dominant allocation mechanism of each work is given in parentheses. “Hetero.” denotes multi-resource heterogeneity at the same allocation level; “Event-adapt.” denotes an allocation rule that changes with the service event type; “Evidence-bound” denotes binding of the result to a frozen data and algorithm record; “Componentised” denotes evidence-dependent local settlement freezing. A checkmark indicates that the feature is explicitly addressed; a dash indicates that it is not addressed; partial indicates partial coverage.

Work	Application (Allocation Rule)	Hetero.	Event-Adapt.	Evidence-Bound	Componentized
Yu 2024 [11]	V2G carbon benefits (Shapley)	✓	—	—	—
Li 2026 [15]	V2G scheduling (mod. Shapley)	—	—	—	—
Wang 2026 [14]	Integrated energy (semi-tensor)	✓	—	—	—
Liang 2025 [13]	Prosumer fairness (aggregation)	✓	partial	—	—
Volpato 2024 [12]	Energy community (multi-rule)	✓	—	—	—
Chen 2024 [19]	V2G trading (smart contract)	—	—	✓	—
Hossain 2024 [18]	V2V energy trading (PKI protocol)	—	—	✓	—
Fan 2024 [24]	P2P energy trading (BPET framework)	partial	—	✓	partial
Khan 2024 [25] (review)	EV P2P trading (review)	—	—	—	—
Shi 2025 [16]	Truck swapping stations (adaptive Shapley)	✓	partial	—	—
Dagoumas 2026 [22]	DER marketplace (blockchain-anchored)	—	—	✓	—
Sejfuli-Ramadani 2026 [23]	Charging receipts (proof-of-charge)	—	—	✓	—
This work	Charging-swapping V2G (event-cond. v_e + Shapley)	✓	✓	✓	✓

Study Scope and Operational Context

The formulation addresses post-event attribution after an external market or contractual process has determined the event benefit pool. Dispatch optimisation, bidding, and participant recruitment are treated as upstream processes. The workflow is informed by operational concerns in integrated charging-swapping services, including heterogeneous metering, data correction, reconciliation, and dispute handling. The numerical cases are reproducible and synthetic, while operational calibration is considered in Section 5.4.

2. Framework Formulation

The methodology is presented in five parts: the event and evidence representation (Section 2.1), the event-conditioned cooperative value (Section 2.2), the Shapley attribution and its structured approximation (Section 2.3), the calculation-provenance construction, and the evidence-dependent settlement formulation. Every construction is defined at the event level and versioned with the event specification, so that a valuation is always tied to an explicit, reproducible state of its inputs.

2.1. Event and Evidence Representation

Let the participant set of flexibility event e be

$$N_e = \{1, 2, \dots, n_e\}. \quad (1)$$

A participant can represent a stationary BESS, a V2G vehicle fleet, an operator-controlled smart-charging portfolio, a battery-swapping load, or another dispatchable station resource. Each participant is linked to one or more physical assets, but settlement is performed at the participant level after asset-level measurements have been aggregated and verified.

An event is represented by

$$e = (T_e, s_e, \mathbf{r}_e, \mathbf{a}_e, N_e, \mathcal{D}_e, \theta_e, \rho_e), \quad (2)$$

where T_e is the event time window, s_e denotes the service type, $\mathbf{r}_e$ is the requested service vector, $\mathbf{a}_e$ is the available-capability vector, N_e is the participant set, $\mathcal{D}_e$ denotes the frozen input data snapshot, θ_e denotes the valuation algorithm and parameter version, and ρ_e denotes the settlement-rule version. The event-level representation prevents measurements from different event windows or different rule versions from being silently mixed.

Evidence-Adjusted Service Representation

For participant i , the *physical-performance service vector* is

$$\mathbf{x}_{i,e}^{phys} = [x_{i,e}^{P,phys}, x_{i,e}^{E,phys}, x_{i,e}^{T,phys}, x_{i,e}^{Q,phys}]^\top, \quad (3)$$

where the four components are, respectively, the normalized delivered power, the normalized delivered energy, the normalized response timeliness, and the normalized trajectory-tracking quality, as observed from the station measurement stack. The *evidence-confidence vector* is

$$\mathbf{c}_{i,e} = [c_{i,e}^P, c_{i,e}^E, c_{i,e}^T, c_{i,e}^Q]^\top \in [0, 1]^4, \quad (4)$$

where each component aggregates the data completeness, the time-synchronisation quality, the meter-calibration status, the data-source provenance, and the cryptographic-signature verification of the corresponding physical measurement. The *effective service vector* is then

$$\mathbf{x}_{i,e}^{eff} = \mathbf{x}_{i,e}^{phys} \odot \mathbf{c}_{i,e}, \quad (5)$$

where $\odot$ denotes element-wise multiplication. The valuation is computed on $\mathbf{x}^{eff}$ rather than on $\mathbf{x}^{phys}$ alone, so that a contribution is credited only when both the physical response and the evidentiary support are present. This decomposition is what allows the framework to bind the value computation to the evidence state without conflating physical performance with measurement quality. The proposed framework does not require these four dimensions to be universal; additional dimensions can be introduced as long as their definitions, data sources, directions, and normalization versions are frozen with the event. The framework deliberately does not propose a universal estimator for $\mathbf{c}_{i,e}$: the evidence-confidence vector is a versioned input generated by the deployment system's evidence-quality policy (completeness, time-synchronization, calibration, provenance, and signature checks per dimension); Section 3.2 verifies that the allocation responds continuously and monotonically to its degradation.

All four physical dimensions are converted to dimensionless scores before Equation (5), with larger values meaning better service. In the synthetic cases, power and energy are clipped delivered-to-request ratios, $x^P = \text{clip}(p^{del}/p^{req}, 0, 1)$ and $x^E = \text{clip}(E^{del}/E^{req}, 0, 1)$. Timeliness is an inverse response-time score, $x^T = \text{clip}((t_{\max} - t^{resp})/(t_{\max} - t_{\min}), 0, 1)$, so that a smaller physical response time gives a larger score. Tracking quality is $x^Q = \text{clip}(1 - \text{nMAE}, 0, 1)$, where nMAE is the normalized trajectory error. The bounds, reference values $(t_{\min}, t_{\max})$, clipping rule, and measurement windows are versioned in

the metric-definition root $R_M^{(v)}$; hence, the same normalization is used for all participants in one event and cannot be changed after attribution.

2.2. Event-Conditioned Cooperative Value

The event specification consists of two vectors. The *event requirement* $\mathbf{r}_e \in \mathbb{R}_{\geq 0}^{|D|}$ and the *available capability* $\mathbf{a}_e \in \mathbb{R}_{> 0}^{|D|}$ are both provided as inputs to the valuation at event time and are versioned with the event snapshot. The dimension set $D = \{P, E, T, Q\}$ corresponds to power, energy, timeliness, and tracking quality, but the framework is agnostic to the choice of dimensions; additional dimensions such as ramping, state-of-charge, or network location can be added as long as their definitions, data sources, and normalization versions are frozen with the event.

The *technical constraint pressure* on dimension d is the logarithmic ratio

$$p_{e,d} = \log\left(1 + \frac{r_{e,d}}{a_{e,d} + \epsilon}\right), \quad (6)$$

where $\epsilon > 0$ prevents division by zero. The logarithmic form is preferred over the raw ratio $r_{e,d}/(a_{e,d} + \epsilon)$ because the raw ratio is unbounded as $a_{e,d} \rightarrow 0$, which makes the downstream softmax brittle when a capability is temporarily unavailable. The log form saturates gracefully and remains well-defined across the full range of $r_{e,d}$ and $a_{e,d}$.

Before mixing the four pressure components into a single weight vector, we apply a small calibration: the raw log-pressure already lives on a comparable scale across dimensions, and a z-score standardization was found to compress all dimensions toward zero and effectively suppress the service-semantics prior. The implementation therefore uses the raw log-pressure directly in the softmax, and the *service-semantics prior* $\boldsymbol{\pi} \in \mathbb{R}_{> 0}^{|D|}$ enters as a multiplicative factor:

$$w_{e,d} = \text{softmax}_d(\beta p_{e,d} + \log \pi_d) = \frac{\pi_d \exp(\beta p_{e,d})}{\sum_{k \in D} \pi_k \exp(\beta p_{e,k})}, \quad \sum_{d \in D} w_{e,d} = 1, \quad (7)$$

where $\beta > 0$ controls the sensitivity to the per-dimension scarcity term. The prior $\boldsymbol{\pi}$ encodes the operational fact that some dimensions are intrinsically more important for a given service type regardless of the instantaneous scarcity level: a fast-response event is *by definition* more about timeliness than about sustained energy, and a peak-shaving event is *by definition* more about power than about tracking quality. Without the prior, the weight vector would be entirely determined by the realized scarcity, which can produce counter-intuitive weights when one dimension is simultaneously abundant and operationally important. With the prior, the event weight vector is the product of *service semantics* and *real-time scarcity*:

$$\text{service semantics} + \text{real-time scarcity} \rightarrow \mathbf{w}_e. \quad (8)$$

In the representative cases below, peak shaving, fast response, and energy support are characterized by $\boldsymbol{\pi} = (1, 1, 1, 1)$, $(0.6, 0.4, 1.4, 1.0)$, and $(0.4, 1.4, 0.4, 0.8)$, respectively. The pressure vector $(\mathbf{r}_e, \mathbf{a}_e)$ is propagated through Equations (6) and (7) with $\beta = 3$ to obtain the event weight vector. Section 3.1 examines the response of this mapping across 500 synthetic events and through controlled counterfactuals. The prior $\boldsymbol{\pi}$ remains an operator- and policy-defined input that is versioned with the event specification.

Coalition Value

For coalition $S \subseteq N_e$, the coalition value is decomposed into a separable per-dimension component and a small set of physically meaningful interaction terms:

$$v_e(S) = \sum_{d \in D_e} w_{e,d} u_d \left(\frac{\sum_{i \in S} x_{i,e}^{eff,d}}{r_{e,d}} \right) + \sum_{(d,k) \in \mathcal{C}} \eta_{dk} u_{dk}(S), \quad (9)$$

where $\mathcal{C}$ is a small set of dimension pairs that the operator declares to be physically coupled, and $D_e = \{d \in D : r_{e,d} > 0\}$ is the *active dimension set*: dimensions on which the event places no requirement ($r_{e,d} = 0$) are excluded from the value construction, which keeps the normalization $\sum_{i \in S} x_{i,e}^{eff,d} / r_{e,d}$ well-defined. The single-dimension utility is a continuous, saturating concave function with a hard clip:

$$u_d(z) = 1 - \exp(-\alpha_d \min(z, z_{\max})), \quad z_{\max} \geq 1, \quad (10)$$

which replaces the hard cap $\min(\cdot, 1)$ used in the preliminary version of this framework. The utility is strictly increasing and concave on $[0, z_{\max}]$ and constant beyond $z_{\max}$ (it is C^0 but not C^1 at $z = z_{\max}$); it has three desirable properties: (i) it is strictly increasing up to $z_{\max}$, so early resources have a non-trivial marginal value; (ii) it is concave on $[0, z_{\max}]$, so diminishing returns are first-class; and (iii) it is bounded, so the total coalition value remains finite even if a coalition oversupplies a dimension. The parameter $\alpha_d > 0$ is the steepness of the utility on dimension d , with the convention that a smaller α_d corresponds to a more elastic demand on that dimension.

The interaction term u_{PE} captures the physical coupling between power and energy: a coalition that supplies enough power for the duration of the event but insufficient energy (or vice versa) cannot fully deliver the response. We define

$$u_{PE}(S) = \min \left(\frac{\sum_{i \in S} x_{i,e}^{eff,P}}{r_{e,P}}, \frac{\sum_{i \in S} x_{i,e}^{eff,E}}{r_{e,E}}, 1 \right), \quad (11)$$

weighted by a small coupling coefficient η_{PE} so that the interaction adds at most a modest correction to the separable sum. The framework is intentionally conservative in the number of interaction terms: most flexibility events in integrated charging–swapping stations admit a single power–energy interaction, and additional couplings can be added at the cost of additional operator-supplied coefficients.

2.3. Shapley Attribution and Approximation

The exact Shapley value of participant i is then

$$\phi_i(v_e) = \sum_{S \subseteq N_e \setminus \{i\}} \frac{|S|!(n_e - |S| - 1)!}{n_e!} [v_e(S \cup \{i\}) - v_e(S)]. \quad (12)$$

If the allocation pool assigned to contribution is B_e^{contrib} , the contribution component is

$$A_{i,e}^{\text{contrib}} = B_e^{\text{contrib}} \frac{\phi_i}{\sum_{j \in N_e} \phi_j}. \quad (13)$$

Equation (12) makes the participant-level marginal attribution explicit. The standard Shapley efficiency property guarantees that the value of the grand coalition is fully distributed with respect to the defined game v_e .

The sum over dimensions in Equation (9) is deliberately additive after normalisation. It represents separable marginal service benefits: one unit of normalised power, energy, timeliness, or tracking quality contributes according to its event weight, while the concave utility captures diminishing returns within that dimension. Timeliness and tracking quality are therefore not assumed to be physically identical; they are measured on distinct, common-

direction scales and are additive only at the benefit-accounting layer. Cross-effects that are material for a deployment are represented explicitly by u_{dk} (the present study uses the power–energy term u_{PE}). This assumption makes the contribution of each dimension auditable and can be relaxed by adding a versioned interaction term rather than by silently changing the sum.

If all effective service volumes vanish, the Shapley vector is zero and the normalisation in the contribution allocation is defined by the following convention:

$$q_{i,e} = \begin{cases} \phi_i / \sum_j \phi_j, & \sum_j \phi_j > \epsilon_\phi, \\ 0, & \sum_j \phi_j \leq \epsilon_\phi, \end{cases} \quad A_{i,e}^{\text{contrib}} = B_e^{\text{contrib}} q_{i,e}. \quad (14)$$

When the second branch applies, the pre-established contribution pool is not silently attributed: it is recorded as an unreleased reserve $B_e^{\text{reserve}} = B_e^{\text{contrib}} - \sum_i A_{i,e}^{\text{contrib}}$ and is handled by the contract’s fallback rule (for example, returned to the market or held pending verified service). The same convention covers an all-zero evidence–confidence vector and makes the attribution rule in Equation (14) well-defined without inventing service value.

Exact calculation requires evaluating exponentially many coalitions. The permutation form of the Shapley value is

$$\phi_i = \mathbb{E}_\pi [v_e(P_i^\pi \cup \{i\}) - v_e(P_i^\pi)], \quad (15)$$

where P_i^π is the set of participants preceding i in permutation π and π is uniform over the set Π_{N_e} of all permutations of N_e . We use a variance-reduction strategy designed for heterogeneous event participants and for saturating coalition functions of the form of Equation (9).

Definition 1 (Event-Oriented Position-Balanced Estimator). *Let $\pi \sim \text{Unif}(\Pi_{N_e})$ and let C_k denote the cyclic-shift operator that maps a permutation to its k -th rotation. A balanced batch of size n_e is the multiset $\{C_0(\pi), C_1(\pi), \dots, C_{n_e-1}(\pi)\}$. Given M such batches, the estimator is*

$$\hat{\phi}_i = \frac{1}{M n_e} \sum_{m=1}^M \sum_{k=0}^{n_e-1} \Delta_i(C_k(\pi_m)), \quad \Delta_i(\pi) = v_e(P_i^\pi \cup \{i\}) - v_e(P_i^\pi). \quad (16)$$

Proposition 1 (Unbiasedness of the Event-Oriented Position-Balanced Estimator). *For any $i \in N_e$, $\mathbb{E}[\hat{\phi}_i] = \phi_i$.*

Proof. For any fixed cyclic shift C_k , the image $C_k(\pi)$ is uniform over Π_{N_e} whenever π is. Therefore $\mathbb{E}[\Delta_i(C_k(\pi))] = \mathbb{E}[\Delta_i(\pi)] = \phi_i$ for every k . The estimator is a finite average of these equal-mean terms, so its expectation is also ϕ_i . $\square$

The estimator is therefore theoretically unbiased. The cyclic construction guarantees exact position balance within each complete batch: each participant occupies each coalition–entry position exactly once per batch. Whether this structural balance reduces variance depends on the covariance structure induced by the coalition-value landscape; no general covariance sign is assumed here, and the variance effect is therefore evaluated empirically in Section 3.3.

For operational use, we augment the estimator with a confidence-interval stopping rule based on the *independent batch means*. Define the batch mean for batch m as

$$b_{m,i} = \frac{1}{n_e} \sum_{k=0}^{n_e-1} \Delta_i(C_k(\pi_m)), \quad (17)$$

so that different batches correspond to independent random base permutations. After M batches, the standard error of $\hat{\phi}_i = \frac{1}{M} \sum_{m=1}^M b_{m,i}$ is computed from the M independent batch means.

$$SE_i(M) = \frac{s(b_{\cdot,i})}{\sqrt{M}}, \quad (18)$$

where $s(b_{\cdot,i})$ is the sample standard deviation of the M batch means for participant i . The worst-case relative half-width is

$$R(M) = \max_{i \in N_e} \frac{t_{0.975, M-1} SE_i(M)}{|\hat{\phi}_i| + \epsilon}, \quad (19)$$

where $t_{0.975, M-1}$ is the 0.975 quantile of the Student- t distribution with $M - 1$ degrees of freedom; because the standard error is itself estimated from the M batch means, the Student- t quantile is used instead of the normal approximation 1.96. The sampling loop terminates as soon as $R(M) \leq \tau$ for a target $\tau \in (0, 1)$ (default $\tau = 0.05$), and reports the smallest M that satisfies the criterion together with the resulting $\hat{\phi}$. This is an auditable empirical stopping criterion, not a claim of fixed-sample 95% coverage under optional stopping: repeated monitoring of an ordinary Student- t interval can alter its unconditional coverage. The criterion is useful for transparent computation because the batch means, threshold, and stopping batch are recorded in the certificate. Note that the standard error is computed from the M independent batch means rather than from the Mn_e correlated rotation samples, which would underestimate the true variance.

Section 3.3 empirically compares the proposed estimator against plain random, anti-thetic, and position-stratified baselines on RMSE and runtime at equal marginal-evaluation budget; the batch-mean confidence criterion is demonstrated separately through the adaptive stopping rule.

2.4. Calculation Provenance

A fair allocation is insufficient if the result cannot be reproduced from the exact data state used at the time of calculation. Therefore, the raw operational database is separated from the frozen event snapshot. Let

$$R_e^{(v)} = H(\text{canonicalize}(\mathcal{D}_e^{(v)})) \quad (20)$$

denote the snapshot root for version v , where $H(\cdot)$ is a cryptographic hash function. High-frequency data remain off-chain; only digests and indices need to be anchored in a permissioned ledger.

The framework distinguishes *proof-of-charge* platforms, which establish that a charging record is authentic and tamper-evident [23], from *valuation provenance*, which records how a particular attributed monetary amount was computed from the data, metric, algorithm, and parameter state. The proposed certificate is the latter. Formally, the calculation certificate for event e at version v is the ordered 9-tuple

$$\mathcal{C}_e^{(v)} = \{H(\mathbf{E}_e), R_D^{(v)}, R_N^{(v)}, R_M^{(v)}, H(\mathcal{A}), H(\Theta_e), s^{(v)}, R_T^{(v)}, R_O^{(v)}\}, \quad (21)$$

with the following interpretation:

- $H(\mathbf{E}_e)$ —the digest of the event state $(T_e, s_e, \mathbf{r}_e, \mathbf{a}_e)$: the time window, the service type, the requirement vector, and the available-capability vector.
- $R_D^{(v)}$ —the frozen data-snapshot root, covering both the raw measurements and the operator-declared evidence-confidence vector $\mathbf{c}_{i,e}$.
- $R_N^{(v)}$ —the participant-set root: the Merkle root of the canonicalized identifier, role, and service-vector list of N_e .
- $R_M^{(v)}$ —the metric-definition root, binding the dimension set D , the normalization constants, the utility parameters $(\alpha_d, z_{\max}, \eta_{dk})$, and the measurement-to-normalization rules.

- $H(\mathcal{A})$ —the digest of the *valuation-algorithm* artifact (exact enumeration or structured sampling), so that the same computation can be re-executed byte for byte; the settlement-mechanism artifact $H(\mathcal{A}_{\text{set}})$ is anchored in the same way but lives as a row of the dependency matrix (Section 2.5).
- $H(\Theta_e)$ —the parameter digest of all valuation parameters that are not part of the event state, namely π_e , β , α_d , z_{max} , and η_{dk} ; the event state and the parameter state are thereby committed separately and without overlap.
- $s^{(v)}$ —the structured execution descriptor, recording the computational mode (exact enumeration, or a named stochastic estimator together with its randomisation state and sample count). The descriptor must truthfully characterise the execution that produced $R_O^{(v)}$; Section 4.6 verifies that even a minimal change to the recorded execution state breaks certificate verification.
- $R_T^{(v)}$ —the calculation-trace root, namely the Merkle root of the per-step partial results during the valuation.
- $R_O^{(v)}$ —the output root, the Merkle root of the per-participant contribution vector and the per-component settlement vector.

The nine components distinguish the questions that arise in an attribution audit: which event was evaluated, which data and participants were included, which metric and algorithm versions were applied, how the calculation was executed, and which trace produced the output. The certificate therefore represents valuation provenance rather than transaction integrity alone. A prototype implementation of canonicalization, hashing, and Merkle combination accompanies the paper. Section 4.6 evaluates the construction through five single-field mutation checks. A permissioned ledger can provide an external immutable anchor for the resulting digests, while detailed telemetry and calculation traces remain off-chain. Figure 1 summarizes the resulting connection between source systems, frozen evidence, cooperative attribution, commitment, and dispute handling.




Figure 1. Proposed event-level valuation and evidence-bound settlement framework. Five planes connect the source and station systems (Plane 1), data quality and snapshot freezing (Plane 2), and event valuation with componentized settlement (Plane 3) to the permissioned evidence anchor (Plane 4) and the audit, dispute, and reconciliation layer (Plane 5). Solid arrows trace the forward data and commitment flow, the dashed line marks the permissioned-ledger boundary, and the red dashed path shows a dispute z propagating through the dependency structure B into a component-level freeze mask m . High-frequency telemetry remains off-chain; only evidence commitments and settlement state are anchored.

2.5. Evidence-Dependent Settlement

The gross event benefit B_e is decomposed into components:

$$A_{i,e} = A_{i,e}^{\text{base}} + A_{i,e}^{\text{contrib}} + A_{i,e}^{\text{cost}} + A_{i,e}^{\text{adj}}, \quad (22)$$

where the terms represent, respectively, a base/guaranteed component, contribution-based allocation, verified cost compensation, and a non-negative bonus component. The settlement components used in Equations (25) and (26) satisfy $A_{i,g} \geq 0$. A penalty is stored separately as a non-negative debit $A_{i,e}^{\text{pen}} \geq 0$ and the participant's net payable amount is $A_{i,e}^{\text{net}} = A_{i,e} - A_{i,e}^{\text{pen}}$; it is not represented by a negative settlement component. Thus, U_i and F_i are releasable and frozen monetary amounts, respectively, while SCR and FR are the corresponding aggregate ratios of actual gross payable funds; penalty recovery is reconciled in a separate ledger. Equation (22) is not merely financial bookkeeping: each component is attached to its own evidence dependencies.

Let m be the number of evidence items (frozen data records, metric roots, algorithm/parameter digests, and calculation-trace commitments) and g the number of settlement components for the event. The dependency structure is encoded in a binary dependency matrix

$$B \in \{0,1\}^{m \times g}, \quad B_{jg} = 1 \iff \text{component } g \text{ depends on evidence item } j. \quad (23)$$

B can be derived from a directed acyclic dependency graph, but the matrix form is more convenient for both analysis and implementation.

A dispute over evidence item j is encoded by a binary vector $\mathbf{z} \in \{0,1\}^m$ with $z_j = 1$ if and only if item j is disputed. The set of disputed-dependent components is then the column-wise OR of B along the disputed rows:

$$\mathbf{m} = \mathbb{I}(B^\top \mathbf{z} > 0) \in \{0,1\}^g, \quad m_g = 1 \iff \exists j: B_{jg} = 1 \wedge z_j = 1. \quad (24)$$

The participant-level releasable and frozen amounts are

$$U_i = \sum_g (1 - m_g) A_{i,g}, \quad F_i = \sum_g m_g A_{i,g}, \quad (25)$$

so that $U_i + F_i = A_i$ for every participant.

The columns of B are settlement-component types and the rows are event-wide evidence items. Consequently, a dispute in one row freezes that component for every participant whose $A_{i,g}$ is supported by the row. This common freeze is intentional: the contribution component is a Shapley allocation whose participants are mutually coupled through the same event snapshot and valuation algorithm. If evidence is genuinely participant-local, the row is split into participant-specific items (for example, “telemetry-V2G fleet”) and the same equations then freeze only the affected participant's component. The granularity is therefore explicit in the evidence-item definition rather than implicit in the matrix algebra.

Throughout, the versioned component decomposition is assumed to satisfy the *settlement budget constraint*

$$\sum_{i \in N_e} \sum_{g \in G_e} A_{i,g}^{(v)} = B_e^{(v)}, \quad (26)$$

i.e., the component amounts of version v exhaust the versioned event pool exactly; this constraint is an explicit input to the conservation property (Proposition 2).

Let $K_e^{(v)}$ denote the *allocation commitment* of version v : the Merkle root of the canonicalized allocation matrix $\{A_{i,g}^{(v)}\}$ over all participants and components. A unique release

identifier is derived from the event, the settlement version, the participant, the released component set, the released amount, and the allocation commitment:

$$\text{release_id} = H\left(H(\mathbf{E}_e) \| v \| i \| \mathcal{G}_{\text{released}} \| U_i \| K_e^{(v)}\right). \quad (27)$$

Binding U_i and $K_e^{(v)}$ into the identifier means that the same event, version, participant, and component set with an altered amount yields a different identifier; the identifier therefore commits to the event, the version, the participant, the component set, the amount, and the settlement state together. The deterministic hash guarantees identical identifiers for identical canonical inputs; distinct inputs are expected to produce distinct identifiers under the usual collision-resistance assumption for H . Equation (27) therefore supports idempotent retries, but the downstream payment system must still enforce uniqueness and reject a second execution of an already-consumed release identifier.

Finally, the settlement mechanism itself (the executable mapping that applies Equations (24) and (25)) is an evidence item like any other: its digest $H(\mathcal{A}_{\text{set}})$ is anchored alongside the valuation-algorithm artifact $H(\mathcal{A}_{\text{val}})$, and both appear as rows of B . A dispute over the *valuation* algorithm freezes only the contribution component, whereas a dispute over the *settlement engine* freezes every component; Section 3.4 exercises both.

2.6. Formal Settlement Properties

The matrix-driven settlement is supported by four properties stated as formal propositions. They are the contract that the framework offers to a station operator.

Proposition 2 (Conservation). *For every event e and every version v ,*

$$\sum_{i \in N_e} (U_i^{(v)} + F_i^{(v)}) = \sum_{i \in N_e} A_{i,e}^{(v)} = B_e^{(v)}. \quad (28)$$

Proof. Direct from the definitions: $U_i = \sum_g (1 - m_g) A_{i,g}$ and $F_i = \sum_g m_g A_{i,g}$, so $U_i + F_i = \sum_g A_{i,g} = A_i$ for every participant, and $\sum_i A_i = B_e^{(v)}$ by the settlement budget constraint of Equation (26). $\square$

Proposition 3 (Dispute Locality). *If a settlement component g does not depend on any disputed evidence item, i.e., $\sum_j B_{jg} z_j = 0$, then $m_g = 0$ and the component is not frozen. Equivalently, $m_g = 0$ whenever the row-projection $\text{supp}(B \mathbf{1}_g) \cap \text{supp}(\mathbf{z}) = \emptyset$, where $\mathbf{1}_g$ is the indicator of column g .*

Proof. Direct from Equation (24): $m_g = 1$ requires at least one j with $B_{jg} = 1$ and $z_j = 1$. $\square$

Proposition 4 (Idempotent Release). *The mapping*

$$(e, v, i, \mathcal{G}_{\text{released}}, U_i, K_e^{(v)}) \mapsto \text{release_id} \quad (29)$$

is a deterministic function of its six arguments. Two release instructions with identical arguments therefore produce identical identifiers, and any downstream settlement system can use the identifier as a deterministic unique key.

Proof. $H(\cdot)$ is a deterministic function on the canonicalised concatenation of its inputs, which is itself canonical; the inputs include the released amount U_i and the allocation commitment $K_e^{(v)}$, so two release instructions coincide if and only if the event, version, participant, component set, amount, and allocation state all coincide. $\square$

Proposition 5 (Append-Only Correction). *If late telemetry or an accepted dispute changes the frozen snapshot, a new version $v + 1$ is created and the historical version v is preserved verbatim. The monetary correction for participant i is*

$$\Delta A_i^{(v+1)} = A_i^{(v+1)} - A_i^{(v)}, \quad (30)$$

and the corresponding release instruction carries the version tag $v + 1$ so that the new identifier does not collide with the old one.

Proof. The new certificate $C_e^{(v+1)}$ is computed independently from the new snapshot, and the version tag enters `release_id` through the second argument v , so the two identifiers are distinct even if the released component set coincides. $\square$

Together, the four properties guarantee that the settlement is conservative (no funds created or destroyed by disputes), local (no spurious freeze), idempotent (safe against external retries), and append-only (no destructive recalculation). Section 3.4 reports the quantitative dispute-scenario experiments; a prototype implementation accompanies the paper in the supplementary reproducibility package.

3. Numerical Evaluation Design

The experiments are organized into four groups, each answering a distinct methodological question; all groups use the same physical-performance service vectors and the same coalition-value parameters, so that the results are mutually consistent. Because the cases are synthetic, the parameterized service vectors are not treated as observations from which an empirical effect size can be inferred. Instead, the design separates model response from parameter artifact through fixed-input counterfactuals, one-factor-at-a-time sweeps, confidence degradation experiments, and comparisons with energy-proportional and static-Shapley baselines. These controls identify which reported differences are generated by event conditioning or evidence adjustment and which are simply inherited from the deliberately ordered resource vectors. All experiments were executed on a single commodity x86-64 CPU (16 logical cores) with a fixed random seed; the complete reproduction script, the generated result files, and the pinned software environment are provided in the supplementary reproducibility package. Appendix A provides the synthetic inputs and reference dependency matrix.

3.1. Event-Conditioning and Counterfactual Design

The event-conditioning analysis examines how the weight map of Equations (6) and (7) responds to technical pressure and the service-semantics prior. Three complementary tests separate internal consistency from the allocation consequences of the mapping. Table 2 summarizes these three analyses.

Weight-response internal-consistency check. For each dimension $d \in \{P, E, T, Q\}$, the Pearson correlation between the pressure $p_{e,d} = \log(1 + r_{e,d} / (a_{e,d} + \epsilon))$ and the resulting weight $w_{e,d}$ is computed across 500 random events. Because $w_d = \text{softmax}_d(\beta p + \log \pi)$ is monotonically increasing in p_d when the remaining inputs are held fixed, this analysis checks the numerical implementation of the mapping. The measured values are $\text{corr}(p_P, w_P) = 0.78$, $\text{corr}(p_E, w_E) = 0.67$, $\text{corr}(p_T, w_T) = 0.73$, and $\text{corr}(p_Q, w_Q) = 0.80$. Figure S1 overlays the closed-form response on the synthetic-event scatter.

Service-prior sensitivity. The unconstrained sweep mixes prior and pressure variation. We therefore use a controlled counterfactual in which r is fixed at the peak-shaving representative value and a varies in $[0.3, 1.5]$. The between-cluster centroid distance is 0.299, the average within-cluster spread is 0.073, and their ratio is 4.09. Figure S2 projects

the samples and service-type centroids onto the first two principal components. This test isolates how the specified service priors differentiate the weight space; it does not estimate the priors from operational observations.

Controlled counterfactual. Two effects compete in the unconstrained sweep. A larger w_P increases the importance of power in $v(S)$ and favors the power-strongest resource, whereas a larger r_P reduces the normalized satisfaction x_P^{eff}/r_P . Figure S3 shows their combined effect. The controlled analysis fixes r and varies a , thereby isolating the pathway from capability pressure to allocation. In the strict one-factor-at-a-time (OFAT) case, only a_P varies over $[0.30, 1.50]$, while a_E , a_T , and a_Q remain fixed. Figure 2 therefore traces the computational pathway $a_P \rightarrow p_P \rightarrow w_P \rightarrow \phi_{\text{BESS}}$ without co-variation in the remaining capability dimensions. Figure S5 compares the combined, pressure-only, and prior-only mappings on the same sweep.

Table 2. Event-conditioning analyses. The weight-response internal-consistency check reports the correlation between $p_{e,d}$ and $w_{e,d}$ over 500 synthetic events. The service-prior sensitivity analysis reports centroid separation under fixed r and varying a . The controlled counterfactual varies only a_P , with the remaining capability dimensions fixed.

Analysis Type	Quantity	Value
Weight response	$\text{corr}(p_P, w_P)$	0.78
Weight response	$\text{corr}(p_E, w_E)$	0.67
Weight response	$\text{corr}(p_T, w_T)$	0.73
Weight response	$\text{corr}(p_Q, w_Q)$	0.80
Prior sensitivity	between-cluster centroid distance	0.299
Prior sensitivity	average within-cluster spread	0.073
Prior sensitivity	cluster-separation ratio	4.09
Controlled OFAT	BESS share at $a_P = 0.30$	$\approx 36.9\%$
Controlled OFAT	BESS share at $a_P = 1.50$	$\approx 35.0\%$




Figure 2. One-factor-at-a-time counterfactual in which only a_P varies over $[0.30, 1.50]$. The event requirement and the remaining capability dimensions are fixed. The BESS share and w_P decrease together, tracing the computational pathway from capability pressure to marginal attribution.

3.2. Allocation Sensitivity, Evidence Adjustment, and Coalition Properties

Two allocation baselines are compared with the proposed Event-Adaptive Shapley on the 4-resource representative service vector under random multiplicative noise on the physical-performance components: (i) Energy proportional and (ii) Static Shapley with equal weights. For each rule, we record (a) the mean allocation under 200 noise realizations,

(b) the per-resource rank stability (standard deviation of the rank over the 200 realizations), and (c) the maximum coalition excess (core deficit), defined as

$$\epsilon_{\text{core}} = \max_{S \subset N} \left[v(S) - \sum_{i \in S} \phi_i \right], \quad (31)$$

which measures the largest coalition-level shortfall of the Shapley allocation. The quantity distinguishes marginal-contribution attribution from coalition-stable allocation. The reported ratio normalizes ϵ_{core} by $v(N)$. In the representative peak-shaving event, the adaptive game gives $\epsilon_{\text{core}} = 0.446$, or 41.8% of $v(N)$, whereas the equal-weight game gives $\epsilon_{\text{core}} = 0.576$, or 53.7%. These quantities compare the Shapley outcomes of two different characteristic functions; they characterize how event conditioning changes the game rather than establishing core membership. A nucleolus or core-constrained rule would be required when coalition stability is itself the allocation objective.

All three rules preserve the dominant share ordering under most perturbations, but occasional BESS/V2G rank exchanges occur under both the static and the event-adaptive Shapley, and occasional exchanges between smart charging and swap load under the energy-proportional rule; event conditioning moderately increases this sensitivity (per-resource rank std ≈ 0.10 for the adaptive rule vs. ≈ 0.07 for static). The cost is a small rank fluctuation; the benefit is a meaningful response to the event context. The percentage of perturbations that change each resource's nominal rank makes the sparse rank changes directly visible.

Two additional analyses examine the cooperative and evidentiary dimensions of the formulation. First, the normalized core deficits of the event-conditioned and equal-weight games are compared across the same 500 synthetic events. Second, the V2G fleet's confidence vector is reduced across $c \in \{1.0, 0.9, 0.8, 0.7, 0.5\}$ under all-dimension, energy-only, and timeliness-only degradation. The latter analysis quantifies how evidence quality enters marginal attribution through the effective-service vector.

3.3. Approximation Benchmark

The 12-participant scalability case (two BESS, four V2G, three smart-charging, three swap-load, with role-specific service vectors perturbed by 10% multiplicative Gaussian noise and clipped to $[0.005, 0.4]$) is used as a tractable exact-reference benchmark. Four estimators are compared against the exact Shapley value:

1. **Plain random permutation Monte Carlo.**
2. **Antithetic permutations:** each random base permutation is averaged with its reverse.
3. **Position-stratified sampling:** for every participant and every entry position k , predecessor coalitions are sampled conditional on that position. Since B full permutations correspond to Bn_e marginals, $K = B/n_e$ samples are allocated to each of the n_e^2 participant–position strata, giving the same total marginal-evaluation budget. This implements the position/cardinality stratification principle used in the Shapley sampling literature [26].
4. **Event-Oriented Position-Balanced** (proposed): M full cyclic batches, each of size n_e , so that every participant occupies every entry position once per complete batch.

For each estimator and each total-permutation budget $B \in \{720, 1440, 2880, 5760, 11520\}$ (equivalently $M = B/n_e$ complete cyclic batches for the proposed estimator), we record the RMSE against the exact Shapley vector and the per-replicate runtime. Two-sided Welch's t -tests compare the proposed estimator against plain random MC at each budget. The confidence-interval stopping rule in Equation (19) is exercised separately on the 4-resource representative

case to demonstrate that the algorithm can self-terminate once the worst-case relative half-width falls below 5%.

3.4. Dependency and Provenance Evaluation

The dependency-matrix settlement is exercised under six synthetic dispute scenarios on a CNY 100,000 event, plus a no-dispute reference scenario. The scenarios and the component affected by each dispute are summarized in Table 3:

1. Local cost-evidence dispute.
2. Contribution-wide telemetry dispute.
3. Valuation-algorithm dispute (the valuation-artifact row $H(\mathcal{A}_{\text{val}})$ feeds only the contribution component).
4. Rule-version dispute (a single bonus/penalty component is disputed).
5. Valuation-snapshot dispute (the event inputs used by the valuation).
6. Settlement-mechanism dispute (the settlement-artifact row $H(\mathcal{A}_{\text{set}})$ feeds all four components; used as the global-dependency scenario).
7. No dispute (reference).

Table 3. Synthetic dispute scenarios and their event-wide evidence scope. A component is frozen for all participants supported by the disputed evidence row; participant-local evidence is represented by splitting that row by participant.

Scenario	Disputed Evidence Item	Components Frozen
Local cost-evidence dispute	verified cost record	cost
Contribution-wide telemetry	data snapshot	contribution
Valuation-algorithm dispute	valuation artifact $H(\mathcal{A}_{\text{val}})$	contribution
Rule-version dispute	bonus/penalty rule	adjustment
Valuation-snapshot dispute	event and valuation inputs	contribution
Settlement-mechanism dispute	settlement artifact $H(\mathcal{A}_{\text{set}})$	base, contribution, cost, adjustment
No dispute (reference)	none	none

For each scenario, we report four ratios:

- **Settlement Continuity Ratio (SCR)** = $\sum_i U_i / \sum_i A_i$, the share of the grand pool that remains releasable.
- **Freeze Ratio (FR)** = $\sum_i F_i / \sum_i A_i$, the share of the grand pool that is frozen; conservation implies $\text{SCR} + \text{FR} = 1$.
- **False Freeze Ratio (FFR)** = (amount frozen despite having no dependency on the disputed evidence) / $\sum_i A_i$, the share of the pool that is frozen unnecessarily.
- **Erroneous Release Ratio (ERR)** = (dispute-dependent amount incorrectly released) / (total dispute-dependent amount), the share of genuinely disputed-dependent money that leaks out.

Under a correctly specified dependency matrix, Propositions 3 and 2 imply $\text{FFR} = \text{ERR} = 0$. The correctly specified cases therefore check the implementation of the freeze algebra. The more informative cases perturb the dependency structure itself: a missing edge ($B_{jg} : 1 \rightarrow 0$) and a spurious edge ($B_{jg} : 0 \rightarrow 1$) are evaluated against the reference relation. The calculation-provenance certificate is examined separately through five single-field mutations of the certified state.

4. Results

The results follow the conceptual chain of the framework. We first examine how event conditions reshape the cooperative game and how evidence quality enters attribution. We then evaluate the computational support for Shapley estimation. Finally, we analyze how evidence dependencies determine settlement states under correctly specified and misspecified relations.

4.1. Event Conditions Reshape the Cooperative-Game Geometry

The analyses in Table 2 first confirmed the numerical response summarized by Equation (8). The log-pressure and service-prior mapping then produced three differentiated characteristic functions (Table 4). Peak shaving concentrated value on power ($w_P = 0.758$), fast response concentrated value on timeliness ($w_T = 0.824$), and energy support concentrated value on energy ($w_E = 0.858$). Because these weights enter $v_e(S)$ before attribution, the event definition changes the marginal-contribution landscape rather than post-processing a fixed allocation. Figure 3 shows the resulting participant shares.

Table 4. Event weights and resulting event-adaptive Shapley allocations for the representative 4-resource case (%).

Event	w_P	w_E	w_T	w_Q	BESS	V2G	Smart Charge	Swap Load
Peak shaving	75.84	10.50	6.83	6.83	36.36	28.86	20.63	14.15
Fast response	3.18	3.26	82.39	11.17	43.16	33.27	15.65	7.93
Energy support	3.39	85.78	2.21	8.62	37.62	28.76	20.03	13.60




Figure 3. Event-adaptive Shapley share for the three event types. The BESS dominates in the fast-response event because the saturating coalition value rewards fast dispatch, while in the peak-shaving and energy-support events the allocation is closer to the static-Shapley baseline.

Table 5 and Figure S7 in the Supplementary Materials compare the proposed event-adaptive Shapley with two baselines on the peak-shaving event: energy-proportional allocation and static Shapley with equal weights. The weighted-proportional rule is not a third baseline in this version; the comparison is intentionally limited to these two interpretable alternatives and the proposed rule.

Table 5. Allocation comparison for the representative peak-shaving event (%).

Resource	Energy Prop.	Static Shapley	EA Shapley
BESS	39.68	34.49	36.36
V2G fleet	28.57	29.73	28.86
Smart charging	19.05	21.61	20.63
Swap load	12.70	14.17	14.15

4.2. Evidence-Adjusted Attribution and Coalition Properties

Table 6 reports rank sensitivity and coalition excess for the three allocation rules. The dominant ordering was retained in most of the 200 perturbations. The BESS and V2G fleet exchanged ranks in two cases under event-conditioned Shapley and in 1 case under

equal-weight Shapley. Smart charging and swap load exchanged ranks once under energy-proportional allocation. The corresponding rank standard deviations were approximately 0.10, 0.07, and 0.07 for the affected resource pairs. Event conditioning therefore retained the dominant ordering while remaining responsive to changes in the service vectors.

The event-conditioned and equal-weight games also produced different coalition-excess profiles under Equation (31). In the representative event, the event-conditioned game gave $\epsilon_{\text{core}} = 0.446$, or 41.8% of $v(N)$, whereas the equal-weight game gave $\epsilon_{\text{core}} = 0.576$, or 53.7%. These values do not establish core membership because each allocation is evaluated in its own characteristic function. A normalized deficit of 0.418, for example, means that at least one coalition can identify a shortfall equal to 41.8% of the grand-coalition value relative to the sum of its Shapley shares; it is a diagnostic of possible coalition dissatisfaction, not money that is missing from the conserved pool. In practice, a large deficit warns that a subset of resources may prefer to renegotiate or withdraw, so operators should pair Shapley attribution with a stability-oriented rule (such as the nucleolus or a core-constrained allocation) when participation stability is critical. Figure 4 displays the rank-change percentages as a heatmap.

Table 6. Allocation sensitivity and coalition excess for the three rules. Rank standard deviations are reported in the order (BESS, V2G, Smart charging, Swap load) over 200 perturbations. The heatmap in Figure 4 reports the corresponding rank-change percentages; the absolute maximum coalition excess and its normalised value are reported separately.

Allocation Rule	Rank Std. (per Resource)	ϵ_{core}	$\epsilon_{\text{core}}/v(N)$
Energy proportional	(0, 0, 0.07, 0.07)	—	—
Equal-weight Shapley	(0.07, 0.07, 0, 0)	0.576	53.7%
Event-conditioned Shapley	(0.10, 0.10, 0, 0)	0.446	41.8%




Figure 4. Percentage of the 200 physical-performance perturbations that change each resource's nominal rank under the three allocation rules. The heatmap makes sparse BESS/V2G and smart-charging/swap-load exchanges visible even when the dominant ordering is preserved.

Across the same 500 synthetic events, the normalized core deficit of the event-conditioned game had a median of 0.344 (IQR 0.298–0.400), compared with 0.398 (IQR 0.357–0.441) for the equal-weight game. The event-conditioned deficit was smaller in 86.4% of cases. The remaining deficits were substantial, which distinguishes the proposed marginal-contribution interpretation from a coalition-stability criterion.

The evidence-confidence analysis links provenance quality to marginal attribution before any settlement dispute is raised (Table 7; Figure 5). Reducing the V2G fleet's confidence vector from 1.0 to 0.5 changed its share from 28.86% to 19.00% under all-dimension

degradation, to 27.96% under energy-only degradation, and to 28.54% under timeliness-only degradation. The response followed the event weights: energy- and timeliness-only changes had a smaller effect in the power-dominated peak-shaving event. Within the formulation, evidence quality therefore modifies the attributed physical service continuously rather than acting only as a post-allocation acceptance flag.

Table 7. Evidence-confidence sensitivity on the representative peak-shaving event: V2G-fleet allocation share (%) as its evidence-confidence vector is degraded from $c = 1.0$ to $c = 0.5$ under three degradation modes.

Degradation Mode	$c = 1.0$	$c = 0.9$	$c = 0.8$	$c = 0.7$	$c = 0.5$
All dimensions (provenance loss)	28.86	27.31	25.59	23.65	19.00
Energy only (completeness)	28.86	28.74	28.60	28.43	27.96
Timeliness only (time synchronisation)	28.86	28.83	28.79	28.74	28.54




Figure 5. Evidence-confidence sensitivity: V2G-fleet allocation share vs. its evidence-confidence level under the three degradation modes of Table 7. The response is continuous and monotone in every mode.

4.3. Sensitivity to the PE-Coupling Weight η_{PE}

The PE-interaction term in Equation (11) carries weight η_{PE} in the coalition value. To assess its influence, η_{PE} is varied over $\{0, 0.05, 0.10, 0.20, 0.40\}$ (Table 8). The grand-coalition value increases linearly with η_{PE} , and the BESS share increases because the BESS contributes strongly to both dimensions. The slower resources lose a small share as the coupling grows. The default $\eta_{PE} = 0.10$ lies in a range where the interaction remains a modest correction to the separable value. Figure S6 visualises the sweep.

Table 8. Sensitivity of the event-adaptive Shapley allocation to the PE-coupling weight η_{PE} on the representative peak-shaving event (%).

η_{PE}	$v(N)$	BESS	V2G Fleet	Smart Charging	Swap Load
0.00	0.966	35.72	28.97	20.94	14.37
0.05	1.016	36.05	28.91	20.78	14.26
0.10	1.066	36.36	28.86	20.63	14.15
0.20	1.166	36.89	28.77	20.37	13.97
0.40	1.366	37.72	28.63	19.97	13.68

4.4. Computational Support for Shapley Attribution

Table 9 and Figure 6 report the approximation-accuracy results at equal total-permutation budget B . Random and antithetic sampling draw B independent permutations directly; the position-stratified estimator allocates $K = B/n_e$ samples to each of the n_e^2 participant–position strata, so that all estimators consume the same Bn_e marginal evaluations; the position-balanced estimator draws B/n_e base permutations and uses the n_e cyclic rotations of each. A genuine position-stratified estimator materially changes the interpretation relative to a plain-random baseline: it achieves the lowest RMSE at every tested budget, consistent with the variance-reduction literature [26,27]. The proposed position-balanced estimator remains substantially better than plain random sampling by a factor of 2.6–3.4 and better than antithetic sampling by a factor of 1.6–1.9, but its RMSE is 1.8–2.2 times that of the fully position-stratified estimator. The position-balanced method should therefore be viewed as a simple, batch-structured estimator with an elementary unbiasedness proof and a natural batch-mean stopping rule, not as a universally optimal Shapley sampler. The runtime of the proposed estimator at the largest budget is ~ 0.17 s per replicate, still within an interactive budget. Runtime comparisons exclude the coalition-value-table construction that is shared by all four estimators and precomputed once per event; the position-stratified baseline additionally precomputes the full per-stratum marginal populations. The reported runtimes therefore measure estimator execution only and are not intended to establish the deployment-time superiority of any baseline.




Figure 6. RMSE convergence for the four estimators on the 12-participant case (log-log scale).

Two-sided Welch's t -tests gave $p < 10^{-8}$ for the difference in mean RMSE between the position-balanced and random estimators at every tested budget. In the 4-resource representative case, the stopping rule of Equation (19) met the selected threshold after $M = 30$ independent balanced batches, with a worst-case relative half-width of $0.048 < \tau = 0.05$.

Table 9. Approximation accuracy: RMSE (mean $\pm$ SD over 20 replicates) on the 12-participant scalability case, at equal marginal-evaluation budget. The position-balanced column uses $B/n_e = B/12$ base permutations (each contributing n_e cyclic rotations); the position-stratified column allocates $K = B/n_e$ samples to each of the n_e^2 participant–position strata. Bold marks the lowest RMSE per row, which is attained by the position-stratified estimator at every budget.

<i>B</i> (Permutations)	Random	Antithetic	Position-Stratified	Position-Balanced
720	0.00350 $\pm$ 0.00076	0.00207 $\pm$ 0.00060	0.00060 $\pm$ 0.00010	0.00108 $\pm$ 0.00021
1440	0.00258 $\pm$ 0.00076	0.00150 $\pm$ 0.00034	0.00042 $\pm$ 0.00008	0.00090 $\pm$ 0.00020
2880	0.00200 $\pm$ 0.00052	0.00110 $\pm$ 0.00024	0.00027 $\pm$ 0.00006	0.00059 $\pm$ 0.00015
5760	0.00114 $\pm$ 0.00023	0.00072 $\pm$ 0.00021	0.00021 $\pm$ 0.00006	0.00044 $\pm$ 0.00010
11,520	0.00089 $\pm$ 0.00018	0.00051 $\pm$ 0.00011	0.00016 $\pm$ 0.00004	0.00031 $\pm$ 0.00009

4.5. Evidence Dependencies Determine Settlement Locality

Table 10 and Figure 7 show how the dependency relation maps evidence disputes to settlement states under the properties of Section 2.6. Conservation held in every case, with $SCR + FR = 1$. Under the reference dependency matrix, $FFR = ERR = 0$ by construction. A cost-evidence dispute left 85% of the pool releasable, while a rule-version dispute left 95% releasable. Telemetry, valuation-snapshot, and valuation-algorithm disputes each froze the 70% contribution component. A dispute over the settlement mechanism propagated to all four components and froze the full pool. The comparison shows that settlement locality is determined by the structure of the evidence dependencies, not by the existence of an immutable record alone.

Table 10. Settlement Continuity Ratio (SCR), Freeze Ratio (FR), False Freeze Ratio (FFR), and Errorneous Release Ratio (ERR) for six dispute scenarios and the no-dispute reference. Short labels are used in the plot for readability; the full scenario names and evidence scopes are given in Table 3. The synthetic event pool is CNY 100,000. Under the reference dependency matrix, conservation gives $SCR + FR = 1$ and locality gives $FFR = ERR = 0$.

Dispute Scope	SCR	FR	FFR	ERR
Local cost-evidence dispute	0.85	0.15	0.00	0.00
Contribution-wide telemetry	0.30	0.70	0.00	0.00
Valuation-algorithm dispute	0.30	0.70	0.00	0.00
Rule-version dispute	0.95	0.05	0.00	0.00
Valuation-snapshot dispute	0.30	0.70	0.00	0.00
Settlement-mechanism dispute	0.00	1.00	0.00	0.00
No dispute (reference)	1.00	0.00	0.00	0.00

These cases verify that the implementation follows the specified dependency relation. The next analysis asks how the settlement state changes when that relation is itself incorrect.

4.6. Dependency Misspecification and Provenance Integrity Checks

Table 11 reports the single-entry misspecification experiment. Deleting the dependency edge from the valuation-snapshot row to the contribution column ($B_{0,1} : 1 \rightarrow 0$) leaves the freeze mask blind to a genuine dependency: under a valuation-snapshot dispute, the misspecified engine releases the CNY 70,000 contribution component that truly depends on the disputed snapshot, giving $ERR = 1.00$. Conversely, inserting a spurious edge from the cost-evidence row to the base component ($B_{2,0} : 0 \rightarrow 1$) freezes the CNY 10,000 base component under a cost dispute although it does not depend on the cost evidence, giving $FFR = 0.10$. The two fault modes are cleanly separated:

$$\text{missing edge} \Rightarrow ERR > 0, \quad \text{spurious edge} \Rightarrow FFR > 0, \quad (32)$$

with the reference relation producing $\text{FFR} = \text{ERR} = 0$. The two failures expose the boundary of dispute locality: a missing edge releases an unsupported component, whereas a spurious edge withholds a supported component. Correct dependency specification is therefore part of the settlement model rather than an implementation detail.




Figure 7. SCR, FR, and FFR for the dispute scenarios in Table 10. ERR is zero under the reference dependency matrix and is omitted from the bars. FFR is zero in every scenario; therefore, the green bars have zero height and coincide with the horizontal axis.

Table 11. Effect of dependency-matrix misspecification. FFR and ERR are evaluated against the reference dependency relation while the settlement mechanism uses the indicated matrix.

Case	SCR	FR	FFR	ERR
Correct B (valuation-snapshot dispute)	0.30	0.70	0.00	0.00
Correct B (cost-evidence dispute)	0.85	0.15	0.00	0.00
Missing edge: snapshot $\rightarrow$ contribution	1.00	0.00	0.00	1.00
Spurious edge: cost $\rightarrow$ base	0.75	0.25	0.10	0.00

The provenance check used the position-balanced run terminated by the stopping rule of Section 2.3. Its execution descriptor, batch-mean trace, and estimated attribution formed the certified state. Honest recomputation reproduced the certificate, whereas each of the five single-field mutations changed at least one digest. The tested fields were a meter value, β , the algorithm artifact, a participant identifier, and the recorded batch count. These checks show that the certificate binds the tested calculation state to a trusted reference. External authentication or ledger anchoring remains the mechanism that protects that reference from replacement.

5. Discussion

5.1. Fairness Is Conditional on the Definition of Value

The first implication of the study is that cooperative fairness is conditional on the characteristic function. Shapley axioms determine how a given coalition value is attributed, but they do not determine whether that value represents the technical objective of the event. A fixed energy-centric game can undervalue fast or accurately tracking resources, while a response-centric game can undervalue sustained energy delivery. By placing event conditions upstream of attribution (Section 2.2), the proposed formulation treats the definition of value as part of the fairness problem rather than as a fixed input to it. The saturating utility in Equation (10) represents diminishing returns, while the effective-service relation in Equation (5) makes evidentiary support part of attributable performance.

The numerical results illustrate this distinction at two levels. Event weights changed the participant shares across peak-shaving, fast-response, and energy-support cases. They

also changed the coalition-excess profile: the event-conditioned game had a smaller normalised deficit than the equal-weight game in 86.4% of the sampled events. This comparison does not imply that Shapley lies in the core, because each weight vector defines a different game. Instead, it shows that event modelling affects both individual marginal contributions and coalition-level dissatisfaction. Where coalition stability is the primary objective, the nucleolus or a core-constrained rule should be evaluated alongside Shapley [12].

5.2. Trust Is a Compositional Property of Settlement

The second implication is that settlement trust is compositional. Transaction immutability and calculation provenance answer different questions. An immutable record establishes persistence, whereas the certificate of Equation (21) identifies the event, evidence, metric, algorithm, execution state, trace, and output associated with an attributed amount. The dependency relation of Equation (23) adds a further layer by determining which payment components remain supported when part of that evidence is contested. Together, Sections 2.4 and 2.5 connect reproducibility with dispute propagation.

This distinction connects two research streams that are commonly treated separately. V2G allocation studies have extended Shapley attribution with carbon, supply-security, participation, and adaptive contribution factors [11,13–16]. Charging cost and dynamic pricing are also part of the economic context: on-demand dynamic charging prices can shift both the value of flexibility and the incentives to delay or discharge, as discussed by Hussain et al. [28]. In the present framework, price and verified cost enter the event benefit pool and the cost component, while the event-conditioned game allocates the contribution component; a dynamic-pricing model can therefore be attached upstream without changing the evidence-dependent settlement algebra. Blockchain-based studies have strengthened transaction integrity and proof of charging, while reviews emphasise scalability and implementation challenges [22,23,29]. The present formulation links these streams by making evidentiary support an explicit determinant of both attributed value and settlement state. A missing dependency edge releases an unsupported component, whereas a spurious edge withholds a supported component. Dispute locality is therefore a property of the dependency structure, not of blockchain anchoring alone.

To make the pricing implication quantitative, we added an upstream dynamic-price sensitivity experiment to the reproducibility script. The verified-cost component was multiplied by 0.6, 0.8, 1.0, 1.2, and 1.4 while the event definition, evidence-confidence vector, dependency matrix, and attribution rule were held fixed. The resulting CSV file and Figure 8 show the corresponding verified cost, residual contribution pool, and net settlement pool. This experiment is intentionally an economic sensitivity rather than a new pricing optimizer: it demonstrates how an on-demand price signal can change the amount available for settlement, while the downstream evidence and dispute algebra remains unchanged. The synthetic multipliers are stress-test values and should not be interpreted as estimates of a field tariff response.

5.3. From Allocation Outputs to Revisable Settlement States

Versioned certificates and append-only corrections turn an allocation from a static output into a revisable settlement state. Proposition 5 and Equation (30) express the correction as a new version and a monetary delta without overwriting the historical calculation. High-frequency telemetry can remain off-chain, while compact commitments preserve the link between data versions and settlement versions [21]. Proposition 4 provides a deterministic release key for exact retries, although enforcement of payment state remains a responsibility of the downstream financial system.




Figure 8. Upstream dynamic-price sensitivity for the reference CNY 100,000 pool. The verified-cost multiplier changes the cost and residual contribution amounts, while the event-conditioned attribution and evidence-dependent release rules are held fixed.

The position-balanced estimator supports this chain computationally. Its contribution is not a new definition of the permutation representation in Equation (15), but the batch structure of Definition 1. Proposition 1 establishes unbiasedness, and Equation (17) supplies independent batch means for uncertainty monitoring. Position-stratified sampling remained more accurate in the benchmark, which clarifies the trade-off: position balancing favours a simple event-level execution and audit structure rather than minimum RMSE under every budget. More advanced stratified or quasi-Monte-Carlo estimators remain compatible with the valuation and settlement formulation [26,27,30].

The resulting audit trail also has an economic interpretation. V2G participation depends on availability, battery degradation, penalties, remuneration, and perceived fairness [6,7]. Field trials further show that participant understanding and trust influence continued engagement [31]. Evidence-linked attribution does not replace bidding, dispatch, or participant communication, but it makes the physical and evidentiary basis of a reward available for reconciliation.

5.4. Boundary Conditions and Research Agenda

The valuation layer raises a calibration problem. The service prior π , temperature β , utility steepness α_d , and coupling coefficient η_{PE} are policy or operator inputs. Their visibility and versioning make the allocation auditable, but do not make the selected values empirically correct. The synthetic cases identify how these parameters affect the game; operational event histories are required to calibrate them. Further work should also extend the current power–energy interaction to ramping, state-of-charge, duration, and network-location effects.

The cooperative layer raises a strategic problem. Shapley attribution provides a marginal-contribution interpretation, not individual rationality or coalition stability. The present model also does not resolve misreporting, collusion, selective telemetry omission, or strategic dispute initiation. These behaviors motivate a broader comparison among Shapley, nucleolus, core-constrained allocation, and incentive-compatible settlement rules.

The evidence layer raises a governance problem. The framework treats the evidence-confidence vector and dependency matrix as versioned inputs, but it does not automatically

derive either of them. Their construction requires data-lineage rules, metric ownership, dispute procedures, and independent audit. The certificate tests reported here establish integrity relative to a trusted reference, not identity, key security, or production-ledger performance. Field validation should therefore combine operational telemetry, observed dispute records, authenticated anchoring, and participant-facing reconciliation. These steps would test whether the formal connection between event value and evidentiary support improves settlement practice at operating stations.

6. Conclusions

This study advances the argument that fair post-event allocation is not determined by the allocation rule alone. It depends jointly on how the event defines coalition value and how evidentiary uncertainty propagates into settlement. The proposed framework addresses both questions by placing service requirements, capabilities, and evidence confidence inside the cooperative game, and by linking the resulting attribution to component-level evidence dependencies.

The numerical study supports three conclusions within the analyzed synthetic cases. Event conditioning changed the marginal-contribution structure and produced a smaller normalized core deficit than the equal-weight game in 86.4% of sampled events. Evidence-confidence degradation changed attribution continuously according to the affected event dimensions. Dependency misspecification separated into two structurally different failures: missing edges released unsupported components, whereas spurious edges withheld supported amounts. The unbiased position-balanced estimator and calculation-provenance certificate provide computational and auditability support for this chain.

Together, these results extend V2G revenue allocation from static contribution sharing to a versioned process in which the validity and releasability of payment components remain explicitly connected to their supporting evidence. Operational calibration, authenticated dependency governance, and field studies are the next steps for determining how this connection affects deployed settlement and participant trust.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/en19204885/s1>, Figure S1: Weight-response analysis (main-text Section 3.1): per-dimension scatter of $p_{e,d}$ vs. $w_{e,d}$ for 500 random events, with the closed-form softmax curve overlaid at the average log-prior. Correlations: $P = 0.78$, $E = 0.67$, $T = 0.73$, $Q = 0.80$; Figure S2: Service-prior sensitivity (main-text Section 3.1): weight samples and the three service-type centroids projected onto the first two principal components of the weight simplex, computed on the controlled counterfactual. Between/within separation ratio = 4.09; Figure S3: Unconstrained random sweep (main-text Section 3.1): BESS share vs. w_p for 500 random events, coloured by service type. The scatter is not monotonic; the pattern reflects the competing effects of event weighting and requirement normalisation; Figure S4: Joint- a controlled counterfactual (main-text Section 3.1): BESS share vs. the available capability a (all four dimensions varied together), with r fixed at the peak-shaving representative value; Figure S5: Weight-map design comparison on the unconstrained sweep: combined vs. pressure-only vs. prior-only. The three designs show similar within-type spread; the combined design is used in the main results because the prior encodes service semantics that the pressure map alone cannot; Figure S6: Sensitivity of the event-adaptive Shapley allocation to the PE-coupling weight η_{PE} on the representative peak-shaving event (numerical values in Table 8 of the main text); Figure S7: Allocation comparison for the representative peak-shaving event: energy-proportional, static Shapley, and event-adaptive Shapley (numerical values in Table 5 of the main text).

Author Contributions: Conceptualization, W.W. and Y.Y.; methodology, W.W. and Y.Y.; software, W.W.; validation, W.W., S.W., Y.W., J.T., and Y.Y.; formal analysis, W.W. and Y.Y.; investigation, W.W., S.W., Y.W., J.T., and L.G.; resources, Y.W., J.T., and L.G.; data curation, W.W. and S.W.; writing—

original draft preparation, W.W. and Y.Y.; writing—review and editing, W.W., S.W., Y.W., J.T., L.G., and Y.Y.; visualization, W.W.; supervision, Y.Y.; project administration, W.W. and Y.Y.; funding acquisition, W.W. and Y.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Science and Technology Project for Provincially Administered Industrial Units of State Grid Jiangsu Electric Power Co., Ltd., grant number JCNJ2026002.

Data Availability Statement: All numerical results reported in this study are generated from reproducible synthetic cases. The problem formulation and workflow design are informed by the authors' operational experience in integrated charging–swapping services; however, no confidential operational measurements are used in the present study. The source code, generated result files, pinned environment, and supplementary input tables are provided in the public repository: <https://github.com/jerryao/v2g> (accessed on 3 October 2026).

Acknowledgments: The authors thank the operational and technical staff involved in integrated charging–swapping station management for their discussions on metering, event reconciliation, and settlement workflow requirements. No confidential operational data from these discussions are used in the numerical experiments reported in this paper.

Conflicts of Interest: Authors W.W., S.W., and L.G. were employed by Nanjing Suyi Industrial Co., Ltd., and authors Y.W. and J.T. were employed by Nanjing Huaqun Energy Group Co., Ltd., whose operational activities are related to the application context investigated in this study. The authors declare that the research was conducted in the absence of any other commercial or financial relationships that could be construed as a potential conflict of interest. As the owner of the charging–swapping station, the funder, State Grid Jiangsu Electric Power Co., Ltd., provided the operational context and proposed the specific research directions addressed in this study.

Appendix A. Synthetic Inputs and Reference Dependency Matrix

Appendix A and Table A1 give the complete synthetic inputs used for the representative four-resource cases. All service scores are dimensionless and already normalized according to Section 2.1; the evidence-confidence vector is $(1, 1, 1, 1)$ unless a sensitivity experiment states otherwise. The random-event experiments use seed 12345, while the estimator benchmark uses the seeds recorded in the reproduction script.

Table A1. Representative synthetic service vectors, event requirements, capabilities, priors, and valuation parameters.

Item	Values	Notes
Resource order	BESS; V2G fleet; Smart charging; Swap load	rows in X^{phys}
X^{phys} (columns P, E, T, Q)	$(0.50, 0.50, 0.50, 0.490); (0.36, 0.36, 0.342, 0.324); (0.24, 0.24, 0.120, 0.180); (0.16, 0.16, 0.048, 0.096)$	clipped dimensionless service scores
Peak shaving ($\mathbf{r}, \mathbf{a}$)	$(0.95, 0.50, 0.30, 0.30); (0.50, 1.00, 1.00, 1.00)$	power-dominated event
Fast response ($\mathbf{r}, \mathbf{a}$)	$(0.30, 0.50, 0.95, 0.60); (1.00, 1.00, 0.50, 0.90)$	timeliness-dominated event
Energy support ($\mathbf{r}, \mathbf{a}$)	$(0.50, 0.95, 0.30, 0.50); (1.00, 0.50, 1.00, 0.80)$	energy-dominated event
Service priors π	Peak: $(1, 1, 1, 1)$; Fast: $(0.6, 0.4, 1.4, 1.0)$; Energy: $(0.4, 1.4, 0.4, 0.8)$	multiplicative softmax prior
Evidence confidence $\mathbf{c}_i$	$(1, 1, 1, 1)$	clean reference; sensitivity uses $c \in \{1, 0.9, 0.8, 0.7, 0.5\}$
Valuation parameters	$\beta = 3; \alpha = (2.5, 2.0, 3.0, 3.0); z_{\max} = 1.5; \eta_{PE} = 0.10; \epsilon = 10^{-9}$	fixed unless ablated
Settlement pool shares	base 10%; contribution 70%; cost 15%; adjustment 5%	CNY 100,000 reference pool

The reference dependency matrix B has rows ordered as (valuation snapshot, data snapshot, cost evidence, rule version, metric definition, valuation algorithm, settlement mechanism) and columns ordered as (base, contribution, cost, adjustment):

$$B = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}. \quad (A1)$$

The six dispute cases and the no-dispute reference in Table 3 are generated by setting one of these rows to one in $\mathbf{z}$. Equation (A1) fixes the row and column order of the reference matrix. This appendix, the source code, and the generated CSV files together provide the synthetic-data overview requested for independent reproduction.

References

1. Kempton, W.; Tomić, J. Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy. *J. Power Sources* **2005**, *144*, 280–294. [\[CrossRef\]](#)
2. Malkova, A.; Striani, S.; Zepter, J.M.; Marinelli, M. Distributed Hierarchical Control with Cost Optimization and Priority-Based Dispatch for Workplace EV Charging: A Field Study. *Energies* **2025**, *18*, 5581. [\[CrossRef\]](#)
3. Li, K.; Li, X.; Xiong, Z.; Tao, S.; Zhao, G.; Jiang, Y.; Qi, H.; Zhang, Y. Unlocking vehicle-to-grid potential of load shifting in China's megacities considering comprehensive real-world behaviors. *Nat. Commun.* **2025**, *16*, 10087. [\[CrossRef\]](#) [\[PubMed\]](#)
4. Chen, D.; Jing, Z.; Tan, H. Optimal Bidding/Offering Strategy for EV Aggregators under a Novel Business Model. *Energies* **2019**, *12*, 1384. [\[CrossRef\]](#)
5. Zoltowska, I. Risk Preferences of EV Fleet Aggregators in Day-Ahead Market Bidding: Mean-CVaR Linear Programming Model. *Energies* **2025**, *18*, 93. [\[CrossRef\]](#)
6. De Caro, F.; Graber, G.; Calderaro, V.; Vaccaro, A.; Galdi, V. Exploring economic profit of EV aggregators in supplying flexibility services through V2G. *Sustain. Energy Grids Netw.* **2025**, *43*, 101851. [\[CrossRef\]](#)
7. Zhang, C.; Kitamura, H.; Goto, M. A New Framework of Vehicle-to-Grid Economic Evaluation: From Semi-Systematic Review of 132 Prior Studies. *Energies* **2025**, *18*, 3088. [\[CrossRef\]](#)
8. Zhan, S.; Zhou, Y.; Feng, D.; Wang, H.; Fang, C. Energy optimization dispatch based on two-stage and multi-objective comparative analysis for PV and BESS integrated fast charging stations with V2G. *IET Renew. Power Gener.* **2024**, *18*, 4426–4438. [\[CrossRef\]](#)
9. Alberizzi, J.C.; Callioni, F.; Pérez Estévez, M.A.; Renzi, M. Optimized integration of charging and battery swapping stations for peak load and cost reduction. *Energy* **2026**, *342*, 139683. [\[CrossRef\]](#)
10. Shapley, L.S. A Value for n -Person Games. In *Contributions to the Theory of Games II*; Kuhn, H.W., Tucker, A.W., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; pp. 307–317.
11. Yu, B.; Lei, X.; Shao, Z.; Jian, L. V2G Carbon Accounting and Revenue Allocation: Balancing EV Contributions in Distribution Systems. *Electronics* **2024**, *13*, 1063. [\[CrossRef\]](#)
12. Volpato, G.; Carraro, G.; Dal Cin, E.; Rech, S. On the Different Fair Allocations of Economic Benefits for Energy Communities. *Energies* **2024**, *17*, 4788. [\[CrossRef\]](#)
13. Liang, B.; Yang, J.; Wen, F.; Wang, L.; Dong, Z.Y. Fairness considered aggregation mechanism for consumers and prosumers in electricity distribution networks. *Electr. Power Syst. Res.* **2025**, *240*, 111285. [\[CrossRef\]](#)
14. Wang, X.; Zhao, H.; Dong, H. A Shapley value profit allocation method for integrated energy system with hydrogen-electric hybrid vehicle aggregator via matrix semi-tensor product. *Energy* **2026**, *346*, 140297. [\[CrossRef\]](#)
15. Li, X.; Mao, X.; Tang, N.; Zhu, Y.; Lin, J. Multi-Agent Deep Reinforcement Learning for V2G Scheduling and Incentive Mechanism to Enhance Distribution Network Loadability. *Energies* **2026**, *19*, 3260. [\[CrossRef\]](#)
16. Shi, P.; Ni, G.; Jin, R.; Wang, H.; Wang, J.; Chen, X. Economic Operation Optimization for Electric Heavy-Duty Truck Battery Swapping Stations Considering Time-of-Use Pricing. *Processes* **2025**, *13*, 2271. [\[CrossRef\]](#)
17. Xue, F.; Chang, K.; Li, W.; Wang, Q.; Zhao, H.; Zhang, H.; Ni, Y.; Xia, W. Blockchain Smart Contract-Enabled Secure Energy Trading for Electric Vehicles. *Energies* **2022**, *15*, 6733. [\[CrossRef\]](#)
18. Hossain, M.S.; Rodine, C.; Tsiropoulou, E.E. A Blockchain and PKI-Based Secure Vehicle-to-Vehicle Energy-Trading Protocol. *Energies* **2024**, *17*, 4245. [\[CrossRef\]](#)
19. Chen, Y.; Lei, X.; Niu, S.; Jian, L. Trustworthy V2G scheduling and energy trading: A blockchain-based framework. *eTransportation* **2024**, *22*, 100376. [\[CrossRef\]](#)
20. Andoni, M.; Robu, V.; Flynn, D.; Abram, S.; Geach, D.; Jenkins, D.; McCallum, P.; Peacock, A. Blockchain technology in the energy sector: A systematic review of challenges and opportunities. *Renew. Sustain. Energy Rev.* **2019**, *100*, 143–174. [\[CrossRef\]](#)

21. Androulaki, E.; Barger, A.; Bortnikov, V.; Cachin, C.; Christidis, K.; De Caro, A.; Enyeart, D.; Ferris, C.; Laventman, G.; Manevich, Y.; et al. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. In Proceedings of the Thirteenth EuroSys Conference, Porto, Portugal, 23–26 April 2018; pp. 1–15. [\[CrossRef\]](#)
22. Dagoumas, A.S.; Laitsos, V.; Makrygiorgou, D.I.; Makrygiorgou, D.; Rigas, F.; Koltsaklis, N.; Roukanas, S. DFLEX: A Blockchain-Anchored Marketplace for Direct Participation of Small Distributed Energy Resources in Transmission and Distribution System Operator Flexibility Markets—Evidence from a One-Year Island Pilot. *Energies* **2026**, *19*, 3627. [\[CrossRef\]](#)
23. Sejfuli-Ramadani, N.; Angelkoska, V.; Idrizi, F.; Rakovic, V.; Ramadani, E.; Risteski, A. A Reproducible Blockchain-Anchored Proof-of-Charge Platform for Auditable EV Charging Receipts. *Future Internet* **2026**, *18*, 423. [\[CrossRef\]](#)
24. Fan, C.; Khazaei, H.; Musilek, P. BPET: A Unified Blockchain-Based Framework for Peer-to-Peer Energy Trading. *Future Internet* **2024**, *16*, 162. [\[CrossRef\]](#)
25. Khan, S.; Amin, U.; Abu-Siada, A. P2P Energy Trading of EVs Using Blockchain Technology in Centralized and Decentralized Networks: A Review. *Energies* **2024**, *17*, 2135. [\[CrossRef\]](#)
26. Castro, J.; Gómez, D.; Molina Ferragut, E.; Tejada, J. Improving polynomial estimation of the Shapley value by stratified random sampling with optimum allocation. *Comput. Oper. Res.* **2017**, *82*, 180–188. [\[CrossRef\]](#)
27. Mitchell, R.; Cooper, J.; Frank, E.; Holmes, G. Sampling Permutations for Shapley Value Estimation. *J. Mach. Learn. Res.* **2022**, *23*, 1–46.
28. Hussain, A.; Lu, Q.-C.; Qureshi, K.N.; Javeed, K. On-demand dynamic charging pricing strategy for Electric Vehicles. *Egypt. Inform. J.* **2026**, *33*, 100921. [\[CrossRef\]](#)
29. Aoudia, M.; Alaraj, M.B.M.; Abu Waraga, O.; Mokhamed, T.; Abu Talib, M.; Bettayeb, M.; Nasir, Q.; Ghenai, C. Toward better blockchain-enabled energy trading between electric vehicles and smart grids in Internet of Things environments: A survey. *Front. Energy Res.* **2024**, *12*, 1393084. [\[CrossRef\]](#)
30. Castro, J.; Gómez, D.; Tejada, J. Polynomial calculation of the Shapley value based on sampling. *Comput. Oper. Res.* **2009**, *36*, 1726–1730. [\[CrossRef\]](#)
31. Lucas-Healey, K.; Jones, L.; Sturmberg, B.C.P.; Ransan-Cooper, H. Participation and sensemaking in electric vehicle field trials: A study of fleet vehicle-to-grid in Australia. *Energy Res. Soc. Sci.* **2024**, *107*, 103343. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.