

Proof of Scientific Contribution: Eligibility-Gated Scoring and Reserve-Aware Compensation for Human Biological Data

Christopher Crecelius

`chris@axondao.io`

AxonDAO

The AxonDAO Research Collective

`community@axondao.io`

AxonDAO

Research Article

Keywords: data valuation, biomedical data, contributor compensation, consent, weighted-product scoring, reserve-aware allocation, decentralized science, desc,

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11413192/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare potential competing interests as follows: The corresponding author, Christopher Crecelius, discloses his role as Founder and CEO of AXDT Inc. and a financial interest in AxonDAO and its associated utility token. The co-author, The AxonDAO Research Collective, is composed of AxonDAO members who likewise hold the token and therefore share this financial interest. All authors thus hold a financial interest that could be perceived to influence the work. The framework is analyzed as a research method, not presented as an investment offer.

Proof of Scientific Contribution

Eligibility-Gated Scoring and Reserve-Aware Compensation for Human Biological Data

Christopher Crecelius^{1,*} The AxonDAO Research Collective

¹AXDT Inc. / AxonDAO

*Correspondence: chris@axondao.io

16 October 2026

Abstract

Compensation for human biological data requires distinct decisions about permission, task-related usefulness, license price, and distribution of proceeds. We present Proof of Scientific Contribution (PoSC), a methodological framework that connects these decisions without equating a dimensionless score with a market price. A license-specific eligibility gate precedes a weighted-product score over quality, verification, rarity, longitudinal completeness, and cohort relevance. Available proceeds are allocated either proportionally or through a reserve-first rule that distributes surplus after meeting agreed contributor reserves. We derive conditions for bounded scoring, budget conservation, reserve feasibility, and allocation invariance. We show that ranking-preserving score transformations can alter payments, that consistent currency conversion leaves exact reserve feasibility unchanged, and that accepted duplicates can increase a submitter’s share despite budget conservation. A currency-quantized extension rounds reserves upward and apportions surplus while preserving the available budget. A dependency-free reference implementation passes 29 deterministic tests, including 11,304 enumerated allocation configurations and 125 scoring configurations. These checks concern the mathematical specification and are not evidence of biological utility, operational verification, market performance, or adversarial security. We outline separate validation requirements for those claims. The contribution is an auditable synthesis of scoring and conditional compensation rules, with explicit assumptions, failure boundaries, and executable checks. The framework does not determine market-clearing prices, establish incentive compatibility, or guarantee payment for unsold contributions.

Keywords: data valuation; biomedical data; contributor compensation; consent; weighted-product scoring; reserve-aware allocation.

1 Introduction

Research-oriented data markets must decide what may be used, what appears useful for a particular task, what a buyer pays, and how contributors are compensated. These decisions are related but not interchangeable. Proposals to treat data as labor motivate contributor compensation [1]; data-valuation methods address the attribution of predictive utility [2, 3]; and private-data pricing studies distinguish buyer payments from compensation to individuals [4]. A practical framework must state which of these functions it performs and which remain external requirements.

For human biological data, a record’s availability does not establish permission for every research purpose. Likewise, a favorable quality score does not determine a transaction price or establish

a contributor’s entitlement to a particular amount. An accounting estimate of collection costs cannot ensure that a willing buyer exists. These distinctions become consequential when an allocation rule is described as a scientific valuation or a guaranteed production-cost floor.

Proof of Scientific Contribution (PoSC) addresses this specification problem. The framework separates an eligibility gate, a factor-based score, the realized proceeds of a defined license, and the resulting contributor payment. It combines established mathematical primitives rather than proposing a new general theory of data value. Its intended application is contributor compensation for biological and behavioral data, but the allocation equations do not require a blockchain, a token, or any particular marketplace implementation.

We study three questions. **RQ1:** What properties follow from an eligibility-gated weighted product, and when can equivalent rankings produce different allocations? **RQ2:** Under what conditions can contributor reserves and budget conservation hold simultaneously, including fees and currency quantization? **RQ3:** Which changes to contribution granularity preserve allocation, and why does normalization alone fail to prevent gains from accepted duplicates? These questions are answered by explicit definitions, derivations, counterexamples, and deterministic software checks.

The contribution has three parts: a license-scoped specification that separates policy inputs from measured outcomes; a joint analysis of conditional guarantees and failure cases; and an executable reference for inspecting the arithmetic independently. We do not claim priority for weighted products, proportional allocation, reserve-first distributions, or remainder apportionment. The value of the composition is its explicit accounting contract and its separation of mathematical correctness from empirical and operational validity.

Compensation is treated as a design objective, not as proof of an economic welfare gain. Models of data markets show that privacy externalities can depress prices while encouraging excessive sharing [5]. Whether a payment mechanism increases desirable participation, improves research, or transfers costs to others therefore requires a separately specified population, counterfactual, and welfare criterion.

2 Related work

2.1 Attribution, prices, and individual compensation

Data Shapley assigns value to training observations through their marginal effects on a defined learning utility [2]. Jia et al. derive efficient exact computation for specific nearest-neighbor utilities and explicitly discuss distributing a buyer’s payment to patients whose records contribute to a model [3]. Individual compensation in a biomedical setting is consequently established prior framing. PoSC does not introduce that objective.

Li et al. develop a framework connecting query prices, privacy loss, and payments to data owners, including compensation and cost-recovery requirements [4]. Agarwal et al. analyze a marketplace for task-dependent data and address challenges arising from replication and combinatorial value [6]. These studies already distinguish market pricing from contributor payments. PoSC’s narrower focus is a factor-based allocation policy whose eligibility and reserve conditions are made explicit. It neither supplies an auction nor inherits the incentive properties of a different market mechanism.

A Shapley value depends on the coalition utility, learning procedure, and reference data. Beta Shapley generalizes the weighting of marginal contributions [7]; Data Banzhaf studies robustness to stochastic utility estimates [8]; and a subsequent analysis revisits the KNN utility used in

nearest-neighbor valuation [9]. OpenDataVal reports task-dependent differences among valuation methods rather than a uniformly best algorithm [10]. Accordingly, agreement with one attribution reference is insufficient to establish a universal measure of scientific contribution.

2.2 Price estimation, consent, and privacy

Hedonic analysis models prices through product characteristics under a specified economic setting [11]. In a data application, the licensed bundle, permitted purpose, access duration, exclusivity, and other terms would be relevant candidate attributes. An observed transaction price is not, by definition, the buyer’s maximum willingness to pay. No transaction-price model is fitted in this article.

Dynamic consent treats permission as a continuing, purpose-sensitive process [12]. Differential privacy instead constrains information disclosure by a randomized mechanism [13]. The latter does not certify measurement authenticity or valid authorization. PoSC therefore treats consent, provenance, acquisition quality, and disclosure controls as distinct interfaces. A score cannot replace the evidence required at any of these interfaces.

Table 1: Positioning relative to selected primary approaches. This comparison is focused on the functions used in the proposed framework, not an exhaustive novelty review.

Prior approach	Established function	PoSC boundary
Data Shapley and KNN-Shapley [2, 3]	Attribution of specified predictive utility.	A factor-based policy index is not a Shapley value or a utility estimate by definition.
Private-data pricing [4]	Separate buyer prices and individual payments, with privacy-related compensation.	Reserves are declared inputs; no calibrated privacy-loss payment function is provided.
Algorithmic data markets [6]	Pricing and allocation under task dependence and replication challenges.	Market clearing, truthfulness, and replication resistance are not solved by normalization.
Dynamic consent and differential privacy [12, 13]	Purpose-sensitive permissions and controlled disclosure, respectively.	Eligibility consumes evidence from these processes; the arithmetic does not implement them.

3 Methods: framework specification

3.1 Research design and unit of analysis

The study is a deductive analysis of a parameterized framework. We define admissible inputs, derive consequences of the allocation equations, construct counterexamples to stronger claims, and check a reference implementation against a finite, fully enumerated input grid. The constructed examples are mathematical fixtures, not sampled biological records or observed transactions. No human-participant dataset or predictive-performance benchmark is analyzed.

A candidate data unit must have a declared boundary: for example, an acquisition episode, a specified interval of a wearable stream, or one curated assay. The unit is not automatically a row, wallet, identity, or individual. A license ℓ defines permitted use, the included units, access conditions, the relevant time interval, and payment terms. Let $\mathcal{J}_\ell$ be its finite, nonempty set of eligible units. Membership and the payment inputs are fixed for each allocation calculation.

Multiple legitimate units belonging to one contributor may be aggregated after calculation; duplicate content must be addressed independently of identifier uniqueness.

Table 2: Notation. All monetary inputs in an allocation use one currency and accounting snapshot.

Symbol	Domain or units	Meaning
$E_{i\ell}$	$\{0, 1\}$	Eligibility of unit i for the specified license.
$x_{i\ell}$	$[0, 1]$	Declared value of quality dimension d .
$b_{m(i)}, w_d$	Dimensionless	Positive modality reference weight and nonnegative exponent.
$S_{i\ell}$	Dimensionless	Policy score; not an observed price or utility.
$U_{i\ell}, a_{i\ell}$	Dimensionless	Usage weight and effective allocation weight $S_{i\ell}U_{i\ell}$.
G_ℓ, P_ℓ	Money	Gross proceeds and net amount available to contributors.
$F_{i\ell}, \text{Pay}_{i\ell}$	Money	Agreed reserve and allocated payment.
B_ℓ, D_ℓ	Money; dimensionless	Summed reserves and summed allocation weights.
q, N	Money; integer	Currency quantum and available number of quanta.

3.2 Eligibility and factor definitions

The binary gate $E_{i\ell}$ represents the outcome of specified consent, provenance, and modality-appropriate acquisition or calibration checks for the license. Failing the gate prevents inclusion in $\mathcal{J}_\ell$; merely assigning zero reward would not prevent unauthorized use. Eligibility must be reconsidered when purpose, permissions, or relevant evidence changes.

Among eligible units, the score uses quality Q , verification V , rarity R , longitudinal completeness L , and cohort relevance K . These are declared proxies, not measurements of scientific contribution by definition. The gate enforces indispensable conditions; V represents a graded assessment above those conditions. A high graded score cannot repair a failed authorization check.

Table 3: Measurement interfaces for an empirical instantiation. The present analysis is conditional on these inputs and does not establish their biological validity.

Factor	Intended meaning	Information that a concrete application must specify
Q	Fitness of acquisition for the declared use	Quality-control variables, normalization, measurement error, thresholds, and missingness.
V	Reliability of verification evidence	Evidence source, independent reference assessment, false acceptance and rejection, and uncertainty.
R	Task-relevant underrepresentation	Reference population, frequency estimator, estimation time, and treatment of erroneous outliers.
L	Completeness of a specified longitudinal interval	Expected schedule, valid observations, censoring, and treatment of nonlongitudinal modalities.
K	Compatibility with a defined cohort or task	Cohort definition, relevance measure, fitting data, and separation from final evaluation outcomes.

Measurement protocols must distinguish an inapplicable dimension from an unmeasured or failed dimension. Correlated factors can count the same evidence repeatedly. Neither common numerical ranges nor a modality multiplier establish comparability between different kinds of biological data. Factor mappings, exponents, and modality weights must therefore be published as part of the policy and assessed for sensitivity.

3.3 Eligibility-gated weighted-product score

Let $\mathcal{A} = \{d : w_d > 0\}$ be a nonempty active set, with finite exponents and positive, finite $b_{m(i)}$. Define

$$S_{i\ell} = E_{i\ell} b_{m(i)} \prod_{d \in \mathcal{A}} x_{id\ell}^{w_d}. \quad (1)$$

Dimensions with zero exponents are omitted, avoiding an implicit 0^0 convention. An indispensable condition must be enforced by the gate or represented by an active factor. Zero active factors yield zero score. Missing active measurements require an explicit, evaluated handling rule; they are not automatically replaced by one.

For eligible units with strictly positive active factors,

$$\log S_{i\ell} = \log b_{m(i)} + \sum_{d \in \mathcal{A}} w_d \log x_{id\ell}. \quad (2)$$

The product implements a policy in which a factor approaching zero suppresses the index. It is a weighted geometric mean times $b_{m(i)}$ only when the active exponents sum to one. Equal unit exponents and equal normalized exponents are distinct specifications. They may preserve ordering within a fixed modality but change allocation shares.

The term *contribution score* names this index, not an established scientific measurement. When used only to rank units, the score can be interpreted ordinally. When its ratios determine payments, its numerical scale has a cardinal policy role. Section 4.3 makes the distinction explicit. A learned predictor of task utility is a separate model and is not a calibrated version of Equation (1) unless the mapping and its constraints are specified.

3.4 Allocation and reserve policy

Let $U_{i\ell} \in [0, 1]$ be a prespecified, auditable usage weight. Usage affects allocation rather than the underlying score. Define

$$a_{i\ell} = S_{i\ell} U_{i\ell}, \quad D_\ell = \sum_{i \in \mathcal{I}_\ell} a_{i\ell}, \quad B_\ell = \sum_{i \in \mathcal{I}_\ell} F_{i\ell}. \quad (3)$$

All reserves and effective weights are finite and nonnegative. Net proceeds P_ℓ are the amount actually available to the contributor allocation after declared external deductions. They are not the buyer's gross payment.

For $P_\ell \geq 0$ and $D_\ell > 0$, proportional allocation is

$$\text{Pay}_{i\ell}^{\text{prop}} = P_\ell \frac{a_{i\ell}}{D_\ell}. \quad (4)$$

This rule contains no reserve guarantee. The reserve-aware rule instead allocates

$$\text{Pay}_{i\ell}^{\text{res}} = F_{i\ell} + (P_\ell - B_\ell) \frac{a_{i\ell}}{D_\ell}, \quad P_\ell \geq B_\ell, \quad D_\ell > 0. \quad (5)$$

The reserve is attached to an included, licensed contribution. A unit with zero effective weight can therefore receive its reserve but no surplus. This convention must be agreed before licensing. If proceeds are below summed reserves, settlement is rejected under the reserve guarantee. If $D_\ell = 0$ and $P_\ell = B_\ell$, paying the reserves is a defined boundary extension. Positive surplus with $D_\ell = 0$ is rejected rather than silently redistributed under another policy.

A reserve may be negotiated, normative, or informed by documented costs. These interpretations must not be conflated. A candidate ledger is

$$C_{i\ell}^{\text{acct}} = c_{i\ell}^{\text{human}} + c_{i\ell}^{\text{capture}} + c_{i\ell}^{\text{verify}} + c_{i\ell}^{\text{settle}}. \quad (6)$$

Each item needs a payee, evidence, currency and date, recurrence rule, and share attributable to the license. Costs deducted externally must not also be reimbursed as the same contributor cost. Repeated licensing must distinguish a new-use payment from repayment of an already recovered acquisition expense. Assigning ordinary living expenditure to a data record would be an additional distributive assumption, not a demonstrated marginal production cost. The formal results do not require $F_{i\ell} = C_{i\ell}^{\text{acct}}$ and do not validate any particular reserve amount.

3.5 Price estimation and attribution interfaces

If actual transactions become available, a candidate price model is

$$\log G_\ell = \beta_0 + \sum_k \beta_k h_k(z_{\ell k}) + \epsilon_\ell, \quad G_\ell > 0, \quad (7)$$

where $z_{\ell k}$ describes the bundle and license terms. This is a proposed statistical specification, not an identity; transformations, selection effects, unsold offers, and held-out prediction require attention. The price coefficients β_k are distinct from the score exponents w_d [11].

For a specified coalition utility ν on n units, a task-specific Shapley reference is

$$\psi_i(\nu) = \sum_{T \subseteq \mathcal{J} \setminus \{i\}} \frac{|T|!(n - |T| - 1)!}{n!} [\nu(T \cup \{i\}) - \nu(T)]. \quad (8)$$

Such values can be negative and depend on the chosen utility [2, 3]. Equations (4) and (5) do not allocate using ψ_i and are not Shapley allocations. Predicting ψ_i from factors does not transfer Shapley axioms to the resulting payments. A future attribution-based payment rule would need its own treatment of negative values, normalization, and reserves.

4 Results: analytical properties and specification checks

The following results concern the stated mathematical model. They are conditional statements, not observations about participant behavior or biological utility. For readability, the license index is omitted where one fixed license is understood.

4.1 Score bounds and local sensitivity

Proposition 1 (Boundedness and conditional monotonicity). *Under Equation (1), $0 \leq S_i \leq b_{m(i)}$. For an eligible unit with positive active factors,*

$$\frac{\partial \log S_i}{\partial \log x_{id}} = w_d > 0. \quad (9)$$

Holding all other inputs fixed, an active factor tending to zero drives the score to zero.

Proof. Each active factor power lies in $[0, 1]$, so their product does also. Differentiating Equation (2) gives the elasticity. The limit follows from $x^{w_d} \rightarrow 0$ as $x \rightarrow 0$ for $w_d > 0$. $\square$

This result establishes formula behavior. It does not establish that increasing a factor improves task utility. In particular, rewarding rarity is a policy decision whose scientific value depends on the reference population and whether the rarity reflects useful information rather than error.

4.2 Budget conservation, reserves, and uniqueness

Proposition 2 (Conditional allocation guarantees). *For $D > 0$, Equation (4) gives nonnegative payments summing to P . If $P \geq B$, Equation (5) gives $\text{Pay}_i \geq F_i$ and $\sum_i \text{Pay}_i = P$. No fully budget-conserving allocation can meet every reserve when $P < B$.*

Proof. The weights a_i/D are nonnegative and sum to one. Summing Equation (5) gives $B + (P - B) = P$; its surplus term is nonnegative. If every payment were at least its reserve when $P < B$, summation would imply $P = \sum_i \text{Pay}_i \geq B$, a contradiction. $\square$

Proposition 3 (Uniqueness under the stipulated surplus policy). *Fix $P \geq B$ and $D > 0$. Require exact budget conservation, payment of each reserve, zero surplus for a zero-weight unit, and one common surplus per unit of positive weight. Equation (5) is the unique allocation satisfying these requirements.*

Proof. The requirements imply $\text{Pay}_i = F_i + \lambda a_i$ for every i , with $\lambda \geq 0$. Summation yields $P = B + \lambda D$, hence $\lambda = (P - B)/D$. $\square$

The uniqueness is conditional on the chosen proportional-surplus policy; it is not a characterization of a universally fair allocation. For a constructed example with weights (1, 3), reserves (\$10, \$20), and net proceeds \$50, the rule pays (\$15, \$35). Conversely, proportional allocation with weights (1, 99) and the same reserves and proceeds pays only \$0.50 to the first unit, despite an aggregate budget sufficient to cover both reserves.

4.3 Ranking invariance is not allocation invariance

Proposition 4 (Common-scale characterization). *Let a and $\tilde{a}$ be nonnegative weight vectors with positive totals D and $\tilde{D}$. They produce identical proportional shares if and only if $\tilde{a}_i = \lambda a_i$ for every i , for one $\lambda > 0$.*

Proof. Equality $\tilde{a}_i/\tilde{D} = a_i/D$ implies $\tilde{a}_i = (\tilde{D}/D)a_i$, including zero entries. Conversely, a common positive multiplier cancels on normalization. $\square$

For fixed reserves and a strictly positive surplus, the same condition characterizes identical reserve-aware payments. At zero surplus, all such payments equal reserves regardless of positive weights. A monotone transformation need not be a common rescaling: weights (1, 4) and (1, 2) have the same ordering but proportional shares (1/5, 4/5) and (1/3, 2/3). Consequently, rank agreement alone cannot justify exchanging scoring conventions in a compensation system.

For positive factors and usage, define $\pi_i = a_i/D$. With membership and usage fixed, the local sensitivity to a shared score exponent is

$$\frac{\partial \pi_i}{\partial w_d} = \pi_i \left(\log x_{id} - \sum_j \pi_j \log x_{jd} \right). \quad (10)$$

This follows by differentiating the normalized exponential weights. An exponent change redistributes shares relative to the current weighted average log factor. Modality-specific rescaling can similarly alter cross-modality compensation even when every within-modality ordering is unchanged. These are reasons to evaluate allocation sensitivity, not only ranking accuracy.

4.4 Fees and currency consistency

Suppose gross proceeds are G , the proportional external fee is $\phi \in [0, 1)$, and fixed external deductions are $c \geq 0$. If these are the only deductions,

$$P = (1 - \phi)G - c, \quad P \geq B \iff G \geq \frac{B + c}{1 - \phi}. \quad (11)$$

Thus correctly calculated net and gross feasibility checks are equivalent. For $c = 0$, the excess over the naive gross threshold B is $B\phi/(1 - \phi)$ in monetary units. The fraction $\phi/(1 - \phi)$ alone is a relative markup, not a monetary shortfall.

Let $p > 0$ be the monetary value of one alternative settlement unit at a common conversion time. In exact arithmetic,

$$P/p \geq B/p \iff P \geq B. \quad (12)$$

Dividing both proceeds and reserves by the same positive conversion rate leaves feasibility unchanged and scales all exact payments consistently. Fixed-token revenues against monetary reserves, different conversion times, liquidity costs, and denomination-specific rounding are distinct mechanisms; they are not covered by this invariance result.

4.5 Contribution granularity and accepted duplicates

Proposition 5 (Partition conservation before rounding). *Replace a unit with weight a and reserve F by subunits with $\sum_k a_k = a$ and $\sum_k F_k = F$. Keep other inputs and net proceeds fixed. Whenever Equation (5) is defined before and after replacement, its aggregate payment to these subunits equals the original unit's payment, and other payments are unchanged.*

Proof. Both D and B are unchanged. Summing the subunit payments gives $\sum_k F_k + (P - B)\sum_k a_k/D = F + (P - B)a/D$. $\square$

The condition concerns conserved weight and reserve, not arbitrary splitting of raw records. A factor estimator can assign more total weight to split records, and a per-record reserve can increase total obligations. Such changes violate the assumptions. Currency rounding can also break exact partition invariance; it is not implied by the discrete extension below.

For a counterexample to duplication resistance, take one weight $a > 0$ and aggregate other weight $H > 0$. If $m > 1$ copies are accepted under different identifiers and each receives weight a , the submitter's proportional share increases by

$$\frac{ma}{ma + H} - \frac{a}{a + H} = \frac{(m - 1)aH}{(ma + H)(a + H)} > 0. \quad (13)$$

The total budget remains conserved. This does not assert that a particular deployment accepts duplicates; it proves that simple normalization is not a defense. It also explains why identifier uniqueness in an arithmetic implementation must not be advertised as semantic deduplication or Sybil resistance.

4.6 Currency-quantized reserve-aware settlement

Let $q > 0$ be the currency quantum and let available proceeds be Nq , where N is a nonnegative integer. Define the payable reserve in quanta by

$$f_i = \lceil F_i/q \rceil, \quad M = N - \sum_i f_i. \quad (14)$$

Reject the allocation when $M < 0$. For $D > 0$, set $t_i = Ma_i/D$, $h_i = \lfloor t_i \rfloor$, and $r = M - \sum_i h_i$. Give one additional quantum to each of the r largest fractional remainders $t_i - h_i$, resolving ties through a declared canonical order. If $e_i \in \{0, 1\}$ records that assignment, the payment is

$$\text{Pay}_i^{(q)} = q(f_i + h_i + e_i). \quad (15)$$

For $D = 0$, pay qf_i only when $M = 0$; otherwise reject positive surplus as undefined.

Proposition 6 (Discrete budget and reserve guarantees). *For an accepted allocation, Equation (15) sums to Nq and satisfies $\text{Pay}_i^{(q)} \geq F_i$. When $D > 0$, its deviation from the continuous quota based on the rounded reserves is strictly less than one quantum per unit:*

$$\left| \text{Pay}_i^{(q)} - q(f_i + t_i) \right| < q. \quad (16)$$

Proof. Because $\sum_i t_i = M$, the integer $r = \sum_i (t_i - h_i)$ satisfies $0 \leq r < n$. Assigning exactly r remainder quanta makes the total $q(\sum_i f_i + M) = Nq$. Reserve protection follows from $qf_i \geq F_i$ and nonnegative surplus. Each rounded surplus quota is its floor or ceiling; an integral quota is not assigned an unnecessary extra quantum because positive remainders suffice for the r assignments. Its error is therefore less than one quantum. The zero-weight boundary follows directly. $\square$

The feasibility threshold must use the rounded reserves. For example, two reserves of \$0.005 sum to \$0.01 in continuous arithmetic, but each requires one cent when $q = \$0.01$; one cent of proceeds is then insufficient. The error bound in Equation (16) is relative to the *rounded-reserve* quota, not the original unrounded formula. This apportionment does not establish strategy-proofness or partition invariance; canonical identifiers and contributor-level rounding policies require separate governance.

4.7 Reference implementation and verification results

Supplementary Software S1 implements the scoring formula, exact proportional and reserve-aware allocation, gross-fee thresholds, and integer-quantum settlement. Monetary inputs and effective allocation weights use Python’s exact rational arithmetic. The scoring routine uses floating-point logarithms for positive factors, handles active zeros explicitly, and rejects positive products that underflow. Exact allocation is conditional on the supplied numerical weights and does not eliminate measurement or serialization error in those weights.

The software validates nonempty unit sets, unique identifiers, nonnegative monetary inputs and weights, valid fee ranges, reserve feasibility, and denominator boundaries. It does not verify consent, authenticate acquisitions, identify duplicate content, train a predictive model, or execute a financial transaction. Its constructed examples are independent of any previously reported simulation or calibration exercise.

The deterministic grid enumerates $n \in \{1, 2, 3\}$ units, each weight in $\{0, 1, 2, 3\}$, each reserve in $\{0, 1, 2\}$, and integer surplus in $\{0, 1, 2, 3, 4, 5\}$. With currency quantum one, the number of configurations is

$$6 \sum_{n=1}^3 4^n 3^n = 11,304. \quad (17)$$

Both exact and quantized allocation are checked for each configuration. Cases with zero total weight and positive surplus must reject. Accepted cases must conserve the budget, meet every reserve, and remain within the declared rounding error. A separate scoring grid varies three

factors over five values, giving $5^3 = 125$ configurations. Targeted tests cover invalid inputs, fees, common-scale invariance, partition conservation, accepted duplicates, fractional-reserve rounding, and deterministic ties.

Table 4: Executed deterministic specification checks. These counts describe constructed test configurations, not independent biological observations or a statistical validation sample.

Check	Observed result
Named automated tests	29 passed
Failures, errors, and skips	0, 0, and 0
Enumerated allocation configurations	11,304 checked
Accepted configurations satisfying checked invariants	11,109
Expected zero-weight, positive-surplus rejections	195
Scoring configurations	125 checked
Optimized-Python rerun	29 passed; same grid counts
Execution environment	Python 3.13.5; Linux x86_64

The supplied result reports contain source hashes and environment details. Reproduction requires only the standard library and the command `python verify_reference.py`; an optimized run uses `python -O verify_reference.py`. No random seeds or downloaded datasets are required. The finite checks complement the proofs but neither exhaust arbitrary-precision input space nor constitute an independent implementation audit.

For n units and d active factors, scoring and continuous allocation require $O(nd + n)$ arithmetic operations after inputs are available. Sorting a full ranking or the remainder order costs $O(n \log n)$. These are operation counts, not bit-complexity or runtime guarantees; rational-arithmetic cost depends on operand sizes. Acquisition, factor estimation, eligibility assessment, privacy protection, and semantic deduplication are excluded and may dominate an operational system’s cost.

5 Discussion

5.1 Interpretation and scientific scope

The analytical results answer the three research questions within the specified model. The product is bounded and monotone in active factors, but preserving a ranking is weaker than preserving payment shares. Reserve-aware allocation conserves the available budget and pays agreed floors only when the corresponding net or currency-quantized feasibility condition holds. Conserving both weight and reserve under partition preserves exact aggregate allocation, whereas accepted copies with additional weight can benefit their submitter.

These results support a transparent accounting specification, not a completed data market. Rarity, quality, and cohort relevance remain measurement constructs requiring evidence. In particular, payout concentration does not demonstrate concentration on scientific merit: a Gini coefficient or top-share statistic describes payments, not the utility of the underlying records. Likewise, a floor protects an agreed amount conditional on sale; it does not establish a fair price, efficient participation, or demand for unsold contributions.

The name *Proof of Scientific Contribution* identifies the framework. The proofs establish the stated algebraic properties, not authenticity of biological measurements or a blockchain consensus

property. The policy does not establish incentive compatibility, coalition stability, privacy compliance, or normative fairness. This boundary is material because conditional correctness of arithmetic cannot be transferred to the truth of its inputs.

5.2 Empirical validation agenda

An empirical study should first specify the contribution unit, factor measurements, target task, primary outcome, evaluation population, and comparison rules. Fixed factor exponents, learned predictors, and alternative utility-based values should be named as distinct methods. A minimal comparative design includes verification-only and quality-only rankings, equal-product and additive scores, random selection, and task-appropriate attribution baselines [7, 8, 10]. Ablations should separate indispensable eligibility controls from graded verification and isolate each additional factor. A nonlongitudinal dataset cannot establish the usefulness of L .

Development and evaluation must separate candidate data, any reference data used to construct valuation targets, model-selection data, and final test outcomes. Splits must respect participant, site, acquisition, or temporal dependencies where applicable. Cross-fitting a factor predictor prevents training directly on its own held-out target but does not automatically remove dependencies introduced by global coalition values. Studies should state whether they evaluate transductive ranking within one fixed pool or generalization to previously unseen contributors. Scaling, factor definitions, oracle construction, model tuning, and retention thresholds must not use final evaluation outcomes.

A simulated verification flag can test responses to assumed information about corruption; it cannot validate a real verifier. Such studies should publish corruption type, prevalence, and the flag’s sensitivity and specificity separately, include a verification-only baseline, and distinguish simulated from independently measured evidence. Attribution agreement and downstream utility should also be distinguished: matching a KNN-Shapley reference does not guarantee optimal subset selection for another model or endpoint [3, 9].

Reporting should include absolute downstream performance, paired differences, all randomization-level results, and uncertainty for clearly defined estimands. Variability across seeds on one dataset is not evidence of generalization to independent cohorts. A claim that fitted weights outperform fixed weights needs a prespecified practically meaningful gain; a claim of equivalence needs a justified equivalence margin and appropriate analysis. Failure to establish superiority alone does not establish equivalence. No such comparative performance claim is made here.

5.3 Operational, economic, and ethical requirements

A deployment needs a participant-understandable permission process, versioned license records, and procedures for changes, withdrawal, and recourse. Restricting future access and reversing past disclosure are different operations. Dynamic consent provides relevant prior framing but is not implemented by a binary variable [12]. Consent breadth is not used as evidence of scientific merit in the score; any separately negotiated rights must remain explicit in license terms.

An operational security assessment must specify attacker capabilities and test replay, accepted duplicate content, identity splitting, collusion, forged provenance, compromised acquisition devices, and synthetic records that obtain favorable measured factors. Assigning low weights to attackers by construction tests attenuation, not successful detection. Database or payment integration additionally requires consistency under permission changes, atomic settlement, replay protection, error recovery, and prevention of duplicate payments. None of these is provided by the reference arithmetic.

If differential privacy is used, the release mechanism, adjacency relation, parameters, composition, and outputs must be specified [13]. Scores and payment amounts themselves may disclose sensitive information. An attestation of recorded attributes does not independently establish that the original acquisition was truthful. A combined consent, verification, and disclosure-control design therefore requires end-to-end assessment rather than an inference from the presence of one technology.

Cost-informed reserves require documented local costs, defined payees, and repeated-license accounting. Cross-border applications require suitable local evidence and a declared conversion policy; a cost assumption from one country is not automatically transferable to another. Market validation would additionally require transaction and unsold-offer data, participant and buyer behavior, and a defined welfare analysis that considers third-party effects [5]. Potential incentives to overproduce rewarded records, broaden permissions, or strategically alter measurements remain risks to examine, not observed outcomes of this study. A data score must not be interpreted as a person’s worth or entitlement to care.

5.4 Limitations and threats to validity

Table 5: Boundaries of the present contribution.

Dimension	Present limitation	Implication
Construct validity	Factor reliability and mapping to task utility are not established.	A mathematically coherent score may rank scientifically unhelpful units highly.
External validity	No participant dataset or operational market is evaluated.	Transfer to genomics, wearables, longitudinal cohorts, or international settings remains untested.
Security and incentives	Inputs, license membership, and weight formation are treated as supplied.	Arithmetic guarantees do not establish truthfulness, consent validity, or resistance to accepted duplicates.
Computational verification	Finite deterministic fixtures and one reference implementation.	Passing checks is not exhaustive verification or an independent audit; it is not a performance benchmark.
Economic interpretation	Reserves and prices are external inputs rather than estimated welfare measures.	Conditional cost coverage does not prove fair pricing, sufficient demand, or welfare improvement.

The elementary components and deductions are not presented as a new general valuation theorem. The framework’s substantive application value depends on whether transparent factor measurements and allocation policies can improve outcomes relative to simpler alternatives. That remains an empirical question separate from the accounting consistency established here.

6 Conclusion

PoSC formalizes a separation among permission, factor-based scoring, license proceeds, and contributor compensation. Its reserve-aware rule has explicit feasibility conditions, a policy-conditional characterization, and a currency-quantized implementation that preserves budget and agreed floors for accepted inputs. The analysis also identifies important limits: rank-preserving score changes can alter compensation, normalization does not prevent accepted-duplicate gains, and currency conversion must be separated from genuine exposure and rounding effects.

The reference implementation makes these arithmetic claims inspectable and reproducible through deterministic fixtures. It does not supply the missing evidence for biological utility, operational verification, security, or market performance. PoSC should therefore be evaluated as an analytical and methodological framework whose measurement and economic interfaces require independent application-specific validation.

Declarations

Data availability. No empirical biological or transaction dataset is analyzed in this article. All constructed examples and finite input grids are generated by Supplementary Software S1.

Code availability. Supplementary Software S1 accompanies the manuscript and contains the complete reference implementation, test runner, normal and optimized execution reports, source hashes, and reproduction instructions. It requires no third-party Python dependencies. No public repository deposit is claimed.

Ethics and consent. The work reported here consists of mathematical analysis and constructed software fixtures. It does not include participant recruitment, clinical interventions, or analysis of identifiable or individual-level health records. This statement does not establish an ethics determination or permission for a future deployment.

Competing interests. The corresponding author, Christopher Crecelius, discloses his role as Founder and CEO of AXDT Inc. and a financial interest in AxonDAO and its associated utility token. The framework is analyzed as a research method, not presented as an investment offer.

AI assistance. ChatGPT (OpenAI) assisted with manuscript restructuring, literature discovery, mathematical exposition, and development of the reference code and its tests. The software checks are documented in Supplementary Software S1. Responsibility for the submitted manuscript, its scientific claims, and final approval remains with the human authors.

References

- [1] I. Arrieta-Ibarra, L. Goff, D. Jiménez-Hernández, J. Lanier, and E. G. Weyl. Should We Treat Data as Labor? Moving beyond “Free”. *AEA Papers and Proceedings*, 108:38–42, 2018. doi:10.1257/pandp.20181003. Open conference manuscript.
- [2] A. Ghorbani and J. Zou. Data Shapley: Equitable Valuation of Data for Machine Learning. *Proceedings of ICML*, PMLR 97:2242–2251, 2019. <https://proceedings.mlr.press/v97/ghorbani19c.html>.
- [3] R. Jia, D. Dao, B. Wang, F. A. Hubis, N. M. Gürel, B. Li, C. Zhang, C. Spanos, and D. Song. Efficient Task-Specific Data Valuation for Nearest Neighbor Algorithms. *Proceedings of the VLDB Endowment*, 12(11):1610–1623, 2019. doi:10.14778/3342263.3342637. Open full text.
- [4] C. Li, D. Y. Li, G. Miklau, and D. Suciu. A Theory of Pricing Private Data. *Proceedings of ICDT*, 2013. <https://people.cs.umass.edu/~miklau/assets/pubs/markets/li2013theory.pdf>.
- [5] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Too Much Data: Prices and Inefficiencies in Data Markets. *American Economic Journal: Microeconomics*, 14(4):218–256, 2022. doi:10.1257/mic.20200200. Open author manuscript.
- [6] A. Agarwal, M. Dahleh, and T. Sarkar. A Marketplace for Data: An Algorithmic Solution. *Proceedings of the ACM Conference on Economics and Computation*, 2019. doi:10.1145/3328526.3329589. Open manuscript: <https://arxiv.org/abs/1805.08125>.

- [7] Y. Kwon and J. Zou. Beta Shapley: a Unified and Noise-reduced Data Valuation Framework for Machine Learning. *Proceedings of AISTATS*, PMLR 151:8780–8802, 2022. <https://proceedings.mlr.press/v151/kwon22a.html>.
- [8] J. T. Wang and R. Jia. Data Banzhaf: A Robust Data Valuation Framework for Machine Learning. *Proceedings of AISTATS*, PMLR 206:6388–6421, 2023. <https://proceedings.mlr.press/v206/wang23e.html>.
- [9] J. T. Wang and R. Jia. A Note on “Efficient Task-Specific Data Valuation for Nearest Neighbor Algorithms”. *arXiv:2304.04258*, 2023. <https://arxiv.org/abs/2304.04258>.
- [10] K. F. Jiang, W. Liang, J. Zou, and Y. Kwon. OpenDataVal: a Unified Benchmark for Data Valuation. *NeurIPS Datasets and Benchmarks*, 2023. Open manuscript: <https://arxiv.org/abs/2306.10577>.
- [11] S. Rosen. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82(1):34–55, 1974. doi:10.1086/260169. Open university-hosted copy.
- [12] J. Kaye, E. A. Whitley, D. Lund, M. Morrison, H. Teare, and K. Melham. Dynamic consent: a patient interface for twenty-first century research networks. *European Journal of Human Genetics*, 23:141–146, 2015. doi:10.1038/ejhg.2014.71. <https://pmc.ncbi.nlm.nih.gov/articles/PMC4130658/>.
- [13] C. Dwork and A. Roth. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science, 9(3–4):211–407, 2014. doi:10.1561/04000000042. <https://www.cis.upenn.edu/~aaroht/Papers/privacybook.pdf>.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [AxonDAOPoSCSupplementarySoftwareS1.zip](#)