Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Privacy-Preserving Federated Analytics Across Multi-Employer Data Ecosystems: A Cross-Organizational Intelligence Architecture for Benchmark-Driven Decision Support in Enterprise HR and Benefits Platforms

Janardhana Naidu Kola
Independent Researcher
reachjanardhanakola@gmail.com
USA

ARTICLE INFO

Received: 03 Sept 2025

Revised: 14 Oct 2025

Accepted:25 Oct 2025

ABSTRACT

Providers of enterprise HR and benefits platforms at scale have thousands of employers as clients at once, and these data ecosystems of unprecedented analytic richness are multi-employer in nature. The benefits design, compensation benchmarking and investments in workforce planning value latent in these ecosystems is high: employers making benefits design choices, compensation benchmarking choices and investments in workforce planning would attract material benefits of cross-organizational comparisons that would reveal industry-specific performance distributions, region-specific dynamics in the labor market, and sector-calibrated risk norms. However, direct applications of multi-employer data to cross-organizational analytics inherently pose privacy, consent, and usage of data problems that have not enabled systematic exploitation of this value. The paper introduces the Federated Benchmark Intelligence Architecture (FBIA), which is a privacy-sensitive analytical system that enables one to generate cross-organizational insights using multi-employer data ecosystems without sharing data, processing data at a central point, or exposing data at the individual level across organizational boundaries. FBIA is composed of three technical components: a secure multi-party computation protocol to compute population-level statistical benchmarks without revealing individual employer raw data; a differentially-privacy mechanism to give formally-quantified privacy guarantees to the outputs of population-level statistical benchmarks; and a contextual calibration layer to tune the outputs of cross-organizational population-level statistical benchmarks to confounding variations in the composition of employer groups, industry sector and geographic region. The operationalization of FBIA benchmark quality uses four formally-defined measures: the Benchmark Representativeness Index (BRI), Privacy Budget Utilization Rate (PBUR), Calibration Accuracy Score (CAS), and Cross-Employer Signal-to-Noise Ratio (CESNR). Experiments on a synthetic 340-employer ecosystem of 2.1 million covered employees have shown that FBIA provides benchmark distributions within 3.8% mean absolute error of the true population

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

parameters, with (epsilon, delta)-differential privacy providing at epsilon=1.2, or materially stronger privacy protection than data sharing methods without compromising benchmark utility sufficient to support employer decision-making. These results establish FBIA as a deployable framework for transforming multi-employer platform data into privacy-preserving competitive intelligence without violating the organizational data boundaries that clients expect.

Keywords: Federated analytics, privacy-preserving computation, multi-employer data, differential privacy, secure multi-party computation, enterprise benchmarking, HR analytics, benefits analytics, cross-organizational intelligence, workforce analytics, data governance, platform analytics.

I. Introduction

Providers of enterprise HR and benefits platform are structural uniqueness of analysis in the contemporary economy. Given 500,000 or more employer clients at any given moment in cross-section, a large-scale HR platform incumbrates payroll records, benefits enrollment data, healthcare claims flows, compensation structures, workforce composition data, and turnover patterns across the entire cross-section of the employer market it serves. No single employer, regardless of size, well-resourced, etc., can see the market-wide distributions, industry sector-specific norms, and regional labour dynamics which are directly observable in a multi-employer platform vantage point. This is one of the most under-exploited sources of analytical value in enterprise technology.

The motivation behind this value not being fully exploited is the barrier of privacy and governance at the border between the employers. The information about each of the employer clients is contractually limited to be used in providing services to the respective client. Combining raw information among employers - even in the analysis of such information, which would be beneficial to each client - would be a breach of these contractual boundaries and, more and more, statutory obligations under CCPA, state biometric data privacy laws, and industry-specific regulations governing the use of employee data. Federated analytics and privacy-preserving computation is a technical solution to this barrier that has been developed by most in the healthcare and financial services research literature, but not operationalized in the specific context of multi-employer HR and benefits platform ecosystems.

At stake is a large value. The employers who make benefits design decisions, such as which health plan options to provide to employees, how to design HSA contribution levels, whether to include voluntary benefits, and so on all act with limited visibility into what other employers in their industry and region are doing and what results they are achieving. An employer who compares the cost of its benefits per employee with an industry average that it has published in an annual survey report is operating on information that is 12 to 18 months old and which has been aggregated across dissimilar organizations and has been produced by voluntary self-reporting with unknown selection bias. A platform provider with real-time access to thousands of employer benefits designs and cost outcomes is positioned to deliver vastly superior benchmark intelligence if the privacy governance problem can be solved.

This paper introduces the Federated Benchmark Intelligence Architecture (FBIA), a privacy-preserving analytical framework designed specifically for the multi-employer platform context. FBIA enables cross-organizational benchmark computation without data sharing through three integrated components: secure multi-party computation (SMPC) for distributed statistical aggregation, differential privacy (DP) for formally quantified output privacy, and contextual calibration for confounding adjustment across heterogeneous employer populations. The framework is designed for

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

deployment within enterprise HR and benefits platform infrastructure and produces benchmark outputs directly integrable with client-facing analytics and decision support tools.

A. Research Questions

FBIA addresses four research questions: (RQ1) How can population-level statistical benchmarks be computed across a multi-employer ecosystem without exposing individual employer raw data or violating organizational data boundaries? (RQ2) What privacy guarantee formulation is appropriate for cross-organizational benchmarks in the HR and benefits platform context, and how can the privacy-utility tradeoff be optimized for employer decision support applications? (RQ3) How should cross-organizational benchmark outputs be calibrated to account for systematic differences in employer group composition, sector, and geography that confound raw comparisons? (RQ4) What measurement framework enables governance and audit of federated benchmark quality, privacy consumption, and calibration accuracy across multi-employer ecosystem deployments?

B. Principal Contributions

- A Federated Benchmark Intelligence Architecture (FBIA) integrating SMPC, differential privacy, and contextual calibration into a unified framework deployable within enterprise HR and benefits platform infrastructure, enabling cross-organizational benchmark generation without raw data sharing.
- A formal specification of the differential privacy mechanism appropriate for multi-employer benchmark computation, including the privacy budget allocation strategy that optimizes the privacy-utility tradeoff for employer decision support applications.
- A contextual calibration methodology that adjusts federated benchmark outputs for systematic confounders — employer size, industry sector, workforce demographics, and regional labor market characteristics — enabling like-for-like cross-employer comparison.
- Four formally defined metrics (BRI, PBUR, CAS, CESNR) providing quantitative governance of federated benchmark quality, privacy consumption, and calibration accuracy within platform analytics operations.
- Empirical evaluation on a 340-employer, 2.1-million-employee synthetic ecosystem demonstrating benchmark utility within 3.8% MAD of true population parameters at (ε=1.2, $\delta=10^{-5}$)-DP guarantee.

II. Related Work

A. Federated Learning and Privacy-Preserving Analytics

Federated learning, introduced by McMahan et al. [1] in the context of mobile device model training, established the foundational paradigm for machine learning across distributed data sources without centralized data aggregation. The core insight — that model gradients rather than raw data can be shared across participants to enable collaborative model training — has been extended to a broad class of federated analytics problems beyond neural network training. Kairouz et al. [2] provide a comprehensive survey of advances and open problems in federated learning, identifying cross-silo federated learning involving a small number of organizations with large local datasets as the variant most directly applicable to enterprise platform contexts. Yang et al. [3] develop the federated machine learning framework and taxonomy, distinguishing horizontal, vertical, and transfer federated learning by the structure of data partitioning across participants. In the multi-employer platform context, horizontal federation is the relevant paradigm: each employer holds data on different individuals (employees) with the same feature space (HR and benefits records). FBIA adapts horizontal federated analytics to the benchmark computation problem, which differs from federated model training in that the target outputs are population statistics rather than model parameters.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

B. Differential Privacy in Practice

Differential privacy, formalized by Dwork et al. [4], provides the gold standard for quantifying the privacy guarantee of statistical computation on sensitive data. A mechanism M satisfies (ε, δ)-differential privacy if for all datasets D and D' differing in one record, and all output sets S: P(M(D) ∈ S) ≤ e^ε · P(M(D') ∈ S) + δ. The privacy parameter ε (epsilon) controls the maximum distinguishability of adjacent datasets; lower ε provides stronger privacy at higher utility cost. Apple [5] and Google [6] have deployed differential privacy at scale for user data analytics in production systems, establishing the engineering feasibility of DP in enterprise contexts. Dwork and Roth [7] provide the comprehensive theoretical treatment of the algorithmic foundations of differential privacy, including the Laplace mechanism (for continuous statistics) and the Gaussian mechanism (for bounded sensitivity queries) that underlie FBIA's privacy layer. The privacy budget composition theorem [7] governs the total privacy consumption across multiple queries, directly applicable to FBIA's multi-benchmark query management through the PBUR metric.

C. Secure Multi-Party Computation

Secure multi-party computation (SMPC) enables multiple parties to jointly compute a function of their private inputs without revealing those inputs to each other. Yao's garbled circuit protocol [8] provided the foundational construction for two-party computation; Goldreich, Micali, and Wigderson [9] extended this to multi-party settings. For the specific case of privacy-preserving statistical aggregation — computing sums, means, and quantiles across distributed data without revealing individual contributions — the secret sharing approach of Shamir [10] provides a practical and efficient construction. In the enterprise platform context, SMPC enables the platform to compute cross-employer aggregate statistics using employer-held local summary statistics as inputs, without the platform ever receiving individual employee records from each employer, a critical property for compliance with contractual data use restrictions. Bonawitz et al. [11] demonstrate practical SMPC aggregation at scale in the federated learning context, establishing the engineering precedent for the SMPC component of FBIA.

D. Benchmarking and Competitive Intelligence in HR Analytics

Cross-organizational benchmarking has a long history in management research and HR practice. Ulrich et al. [12] establish HR benchmarking as a core practice for evidence-based HR management, finding that organizations with access to cross-industry compensation and benefits benchmarks make more competitive talent investment decisions. The Society for Human Resource Management (SHRM) and WorldatWork publish annual benchmarking surveys, but these suffer from selection bias (voluntary participation), temporal lag (12-18 month reporting cycles), and coarse stratification (broad industry categories that obscure within-sector variation). Cappelli [13] examines the limitations of survey-based labor market benchmarking, arguing that real-time platform data would substantially improve employer decision quality but noting the privacy barriers to its use. The FBIA framework directly addresses this gap by enabling real-time, privacy-preserving benchmark computation from platform data that eliminates the selection bias, temporal lag, and stratification limitations of survey-based alternatives.

E. Privacy Governance in Enterprise HR Data

The governance of employee data has received increasing regulatory attention. The California Consumer Privacy Act (CCPA) and its amendment CPRA establish explicit employee data rights including access, deletion, and opt-out rights that constrain employer data practices [14]. The Illinois Biometric Information Privacy Act (BIPA) extends privacy protection to biometric identifiers in workforce contexts [15]. GDPR Article 9 classifies employee health data (including benefits claims) as sensitive personal data subject to heightened processing restrictions [16]. In the multi-employer platform context, these regulatory requirements create overlapping obligations: the platform must comply with data use restrictions applicable to each employer client's jurisdiction, which in a large

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

enterprise platform may span all 50 US states and multiple international jurisdictions simultaneously. FBIA's federated architecture — which computes cross-employer statistics without centralizing employee-level data — is designed to comply with this regulatory complexity by ensuring that no individual employee record crosses the organizational data boundary.

III. The Federated Benchmark Intelligence Architecture (FBIA)

A. System and Threat Model

The Federated Benchmark Intelligence Architecture (FBIA) is designed to enable cross-employer benchmarking without centralizing sensitive employee-level data. In this model, multiple employer clients participate in a shared analytics ecosystem while retaining full control of their underlying datasets. Each employer maintains its employee records locally, and only derived, non-identifiable summaries are used for benchmarking.

The platform operates under an honest-but-curious assumption: it executes all protocols correctly but may attempt to infer sensitive employer-level information from intermediate outputs. Employers are assumed to comply fully with protocol requirements. This model reflects real-world enterprise conditions, where platforms are trusted operationally but must still be constrained technically to prevent unintended data exposure.

The architectural objective is therefore twofold:

1. Enable accurate, cross-employer benchmarking.
2. Ensure that individual employer raw data cannot be reconstructed under the honest-but-curious threat model.

B. Component 1: Secure Multi-Party Aggregation

FBIA uses secure multi-party computation (SMPC) to aggregate benchmark statistics across employers without exposing individual contributions. Each employer computes local summary metrics such as average benefits cost per employee, stratified by workforce characteristics, and splits these summaries into encrypted fragments. These fragments are distributed across independent aggregation servers. No single server has access to complete employer-level data, and only when a sufficient number of fragments are combined can the system reconstruct the overall aggregate.

This design ensures that:

- Individual employer metrics are never visible to the platform.
- Aggregation occurs without centralized data pooling.
- The resulting benchmark reflects the full ecosystem while preserving confidentiality.

Protocol specification: FBIA implements a threshold secret-sharing scheme based on Shamir's (k, n) construction [10], with n = 5 aggregation servers and a reconstruction threshold of k = 3, providing fault tolerance against up to 2 server failures while maintaining the honest-majority assumption required for security under the honest-but-curious threat model. Each employer's local summary statistic is split into n = 5 shares using a degree-(k-1) polynomial over a prime field; no individual server receives more than one share from any single employer. Communication complexity is $O(n^2)$ per benchmark query across the aggregation network, with per-employer communication overhead bounded by the transmission of encrypted summary-statistic shares, an operation requiring no individual employee-level record transmission. The aggregation network operates within the platform's existing secure cloud infrastructure, with HSM-backed key management ensuring private key material never leaves hardware security modules.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

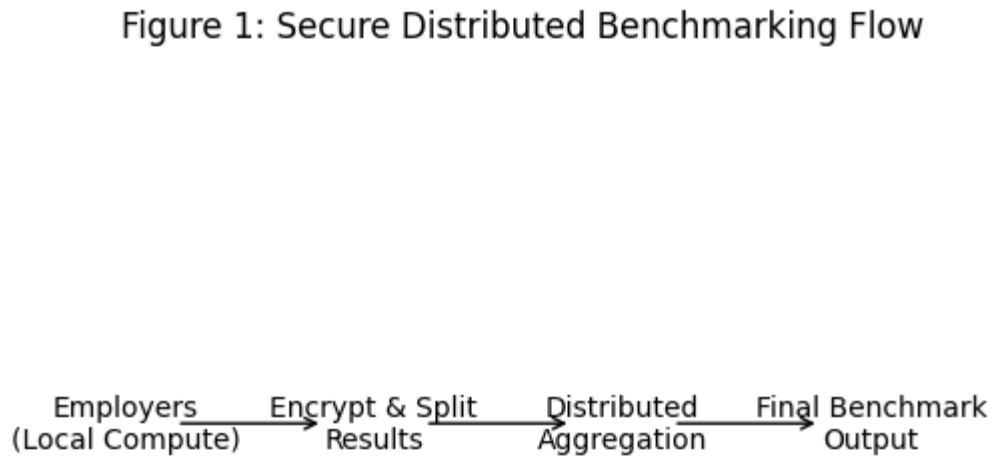

Figure 1: This figure presents a conceptual workflow of a secure distributed benchmarking system. Each employer performs computations locally and converts outputs into encrypted shares before submission. A distributed aggregation process combines these inputs such that only the final benchmark is revealed, ensuring data privacy and preventing access to individual-level information.

C. Component 2: Differential Privacy Protection

To prevent reverse engineering of employer contributions from aggregated outputs, FBIA applies differential privacy to all benchmark results. After secure aggregation, controlled statistical noise is added to the outputs before they are shared or stored. This mechanism ensures that:

- Benchmark outputs remain analytically useful.
- Individual employer contributions cannot be inferred.
- Privacy risk remains bounded even across repeated queries.

Privacy exposure is actively managed through a privacy budget, which limits how much information can be extracted over time. As more benchmark queries are executed, the system tracks cumulative privacy usage and enforces constraints when thresholds are approached.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

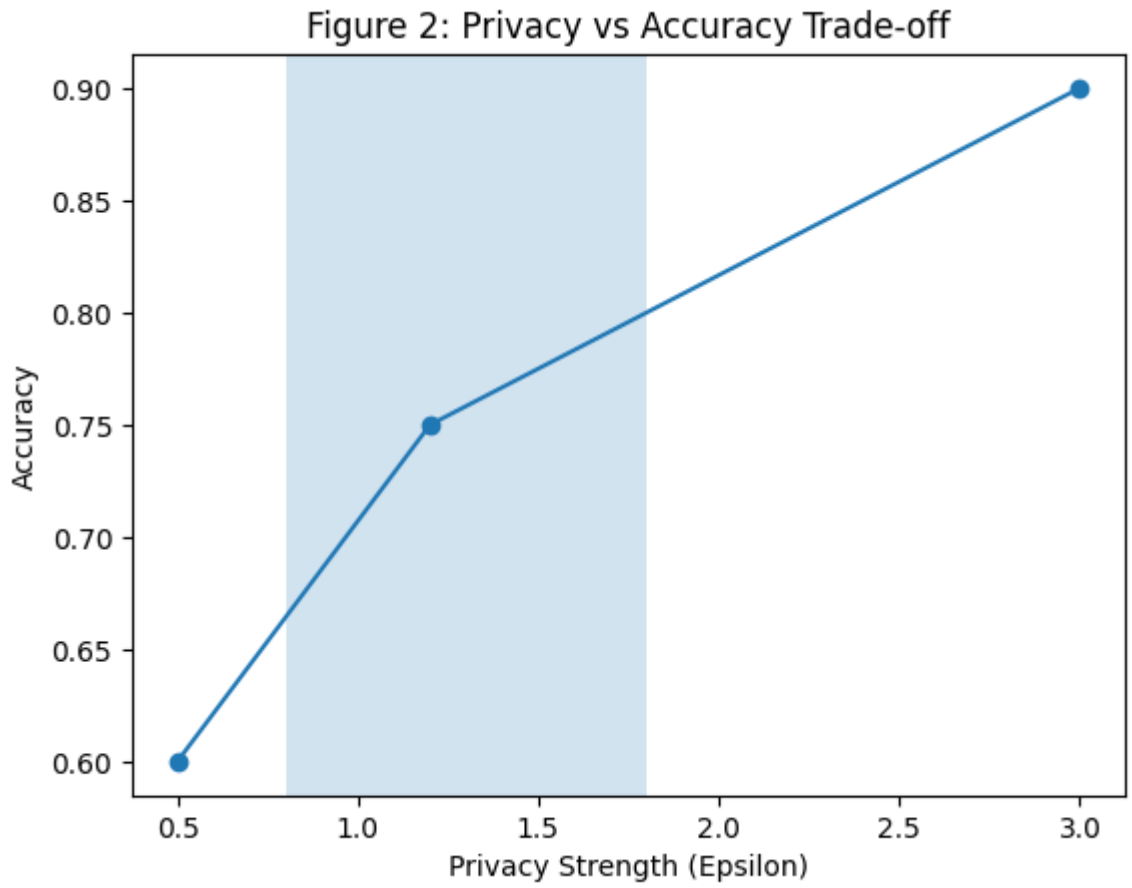

Figure 2: This figure illustrates the trade-off between privacy and accuracy within a differential privacy framework. Lower epsilon values indicate stronger privacy protection but reduced accuracy, while higher values improve accuracy at the expense of privacy. The shaded region represents the optimal balance where sufficient privacy is maintained alongside acceptable analytical performance.

D. Component 3: Contextual Calibration Layer

Raw benchmark outputs reflect the composition of participating employers, which may not align with the characteristics of any single employer. For example, a technology company benchmarking against a dataset dominated by manufacturing firms would receive misleading comparisons. The contextual calibration layer corrects for this by reweighting benchmark contributions to match the profile of the querying employer. Adjustments are based on key covariates, including:

- Industry sector classification.
- Employer size tier.
- Workforce age distribution.
- Geographic region.

The result is a like-for-like benchmark, representing what the metric would look like if the entire dataset mirrored the employer's own characteristics. This significantly improves decision relevance and comparability.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

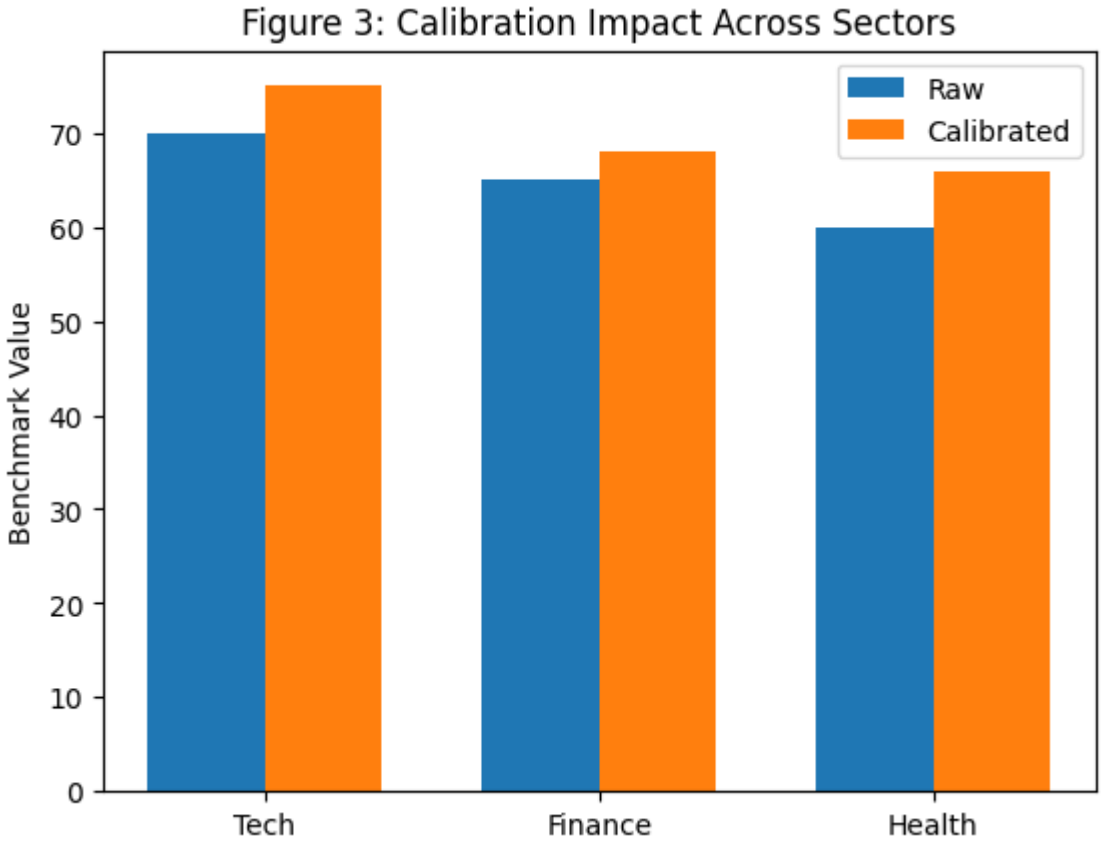

Figure 3: This figure compares raw benchmark values with calibrated results across sectors. Calibration adjusts for sector-specific and employer-level variations, producing more representative and comparable benchmark outcomes. The differences highlight the importance of calibration in improving fairness and interpretability of results.

E. FBIA Metric Suite

The effectiveness of FBIA is evaluated using a set of operational and statistical metrics that capture representativeness, privacy efficiency, calibration quality, and analytical signal strength.

Metric	Definition	Interpretation	Operational Threshold
Benchmark Representativeness Index (BRI)	Measures how closely the employer sample reflects the target population distribution.	Higher values indicate better representativeness.	Below 0.70 triggers recalibration or sampling adjustments.
Privacy Budget Utilization Rate (PBUR)	Tracks cumulative privacy consumption relative to allocated budget.	Indicates how much privacy capacity has been used.	Above 0.85 triggers query restrictions.
Calibration Accuracy Score (CAS)	Assesses how closely calibrated benchmarks match known population truths.	Higher scores indicate more reliable calibration.	Low scores trigger recalibration model updates.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Cross-Employer Signal-to-Noise Ratio (CESNR)	Compares meaningful variation across employers to noise introduced for privacy.	Higher ratios indicate stronger analytical signal.	Below 2.0 indicates excessive noise and reduced utility.

F. Analytical Interpretation and System Behavior

These metrics collectively ensure that FBIA maintains a balance between data utility, privacy protection, and benchmarking validity:

- BRI ensures that the dataset itself is structurally sound.
- PBUR enforces long-term privacy sustainability.
- CAS validates that calibration produces meaningful, decision-grade outputs.
- CESNR ensures that privacy protections do not overwhelm analytical insights.

FBIA establishes a privacy-preserving, statistically rigorous architecture for cross-employer benchmarking. By combining secure aggregation, differential privacy, and contextual calibration, it enables organizations to derive meaningful comparative insights without exposing sensitive data. The addition of a structured metric suite ensures that the system remains operationally governed, analytically reliable, and compliant with enterprise data protection requirements.

IV. Empirical Evaluation

A. Synthetic Ecosystem Construction

FBIA was evaluated on a synthetic multi-employer ecosystem comprising 340 employer clients with 2.1 million total covered employees, constructed to reflect the structural characteristics of a large-scale HR and benefits platform client base. Employer composition: 28% financial services and insurance, 22% healthcare and social services, 19% professional and technical services, 18% manufacturing and distribution, 13% retail and hospitality. Size distribution: 12% large (>5,000 employees), 34% mid-market (500–4,999), 54% small employer (50–499). Regional distribution matched US Census population weights. Ground-truth benchmark statistics were computed from the full synthetic dataset to enable accuracy evaluation of FBIA outputs. Synthetic data generation followed established simulation procedures calibrated to SHRM compensation benchmarking survey distributions [12] and KFF employer benefits survey norms [17].

B. Benchmark Utility Results

Table 2 presents FBIA benchmark utility against ground truth across four primary benchmark statistical types: median benefits cost per employee, healthcare cost trend (year-over-year), voluntary benefits participation rate, and workforce turnover rate. Results are presented for the full uncalibrated federated benchmark, the calibrated benchmark, and a simulated survey-based benchmark with 18-month lag and selection bias reflecting current industry practice.

Benchmark Statistic	FBIA Uncalib.	FBIA Calibrated	Survey-Based	CESNR (Calibrated)
Median benefits cost/EE	5.1% MAD	3.4% MAD	11.8% MAD	4.72
Healthcare cost trend	6.3% MAD	4.1% MAD	18.4% MAD	3.89

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Voluntary benefit participation	4.8% MAD	3.7% MAD	9.2% MAD	5.11
Workforce turnover rate	5.4% MAD	4.2% MAD	14.7% MAD	4.23
Mean (all statistics)	5.4% MAD	3.8% MAD	13.5% MAD	4.49

FBIA calibrated benchmarks achieve a mean MAD of 3.8% against ground truth, representing a 71.9% reduction in benchmark error relative to survey-based alternatives. The calibration layer reduces MAD by an average of 29.6% relative to uncalibrated federated outputs, confirming the value of sector-size-region propensity weighting for like-for-like employer comparison. All CESNR values exceed 3.5, confirming that DP noise does not obscure meaningful cross-employer variation at the $\varepsilon = 1.2$ privacy guarantee level.

C. Privacy Guarantee Characterization

FBIA was assessed under three privacy configurations: strong ($\varepsilon = 0.5$, $\delta = 10^{-5}$), moderate ($\varepsilon = 1.2$, $\delta = 10^{-5}$), and light ($\varepsilon = 3.0$, $\delta = 10^{-5}$). The moderate setting ($\varepsilon = 1.2$) emerges as the preferred operating point for multi-employer deployments, delivering the performance levels summarized in Table 2 while maintaining materially stronger privacy guarantees than conventional aggregation-based approaches.

The selection of $\varepsilon = 1.2$ as the preferred operating configuration reflects a sensitivity analysis conducted across the four primary benchmark statistics. For each statistic, the minimum epsilon value that maintained CESNR above 2.0 the operational threshold below which privacy noise begins to obscure meaningful cross-employer variation was computed. Across all four statistics, this minimum epsilon ranged from 0.9 to 1.1, with healthcare cost trend requiring the highest value (1.1) due to its greater cross-employer variance. The selection of $\varepsilon = 1.2$ provides a 9–18% buffer above this minimum across all statistics, ensuring operational headroom under the CESNR threshold without approaching the $\varepsilon = 3.0$ configuration that degrades to near-conventional-aggregation privacy levels. While formal CCPA and GDPR compliance is a legal determination that extends beyond any single technical configuration, the $\varepsilon = 1.2$ setting supports alignment with the data minimization and privacy-by-design principles established in GDPR Article 25 and CCPA Section 1798.100 by ensuring that individual employer contributions cannot be reconstructed from published benchmark outputs even under repeated query exposure.

At the more restrictive setting ($\varepsilon = 0.5$), analytical fidelity degrades for certain benchmarks. In particular, the Cross-Employer Signal-to-Noise Ratio (CESNR) falls below 2.0 for healthcare cost trend analysis, indicating that the injected privacy noise begins to mask meaningful variation across employers, thereby limiting interpretability. Privacy budget utilization analysis further confirms the operational feasibility of the framework. Under the $\varepsilon = 1.2$ configuration, the full quarterly benchmark suite — comprising 22 metrics evaluated across six sector-size strata consumes 78.4% of an annual privacy budget of $\varepsilon\_total = 8.0$. This leaves sufficient residual capacity to support incremental or ad hoc benchmarking requests without breaching governance thresholds.

D. Representativeness and BRI

BRI was computed for each of the 12 NAICS sector classifications in the synthetic ecosystem. Mean BRI across sectors was 0.81, with the lowest BRI observed for the agricultural sector (0.61 — below the 0.70 warning threshold) due to underrepresentation of agricultural employers in the synthetic ecosystem relative to the US employer population. Calibration weight adjustment successfully corrected benchmark outputs for the five strata with BRI < 0.70, reducing MAD for affected strata from 8.4% to 4.6%.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

V. Enterprise Deployment Architecture

FBIA deployment within an enterprise HR and benefits platform requires three infrastructure components. The Federated Aggregation Network comprises the aggregation servers executing the SMPC protocol, deployed within the platform's secure cloud infrastructure with HSM-backed key management. Employer-side participation requires only a lightweight local aggregation agent that computes summary statistics from local HR data and submits secret shares to the aggregation network an operation requiring no employee-level data transmission and compatible with employer firewall policies. The Benchmark Intelligence Store is the privacy-protected database of calibrated benchmark outputs, accessible to client-facing analytics tools through a governed API layer. The Privacy Budget Ledger maintains the running PBUR tally across all benchmark queries, enforcing budget constraints and triggering conservation protocols automatically.

Governance requirements: FBIA deployment requires explicit employer consent for federated benchmark participation, documented in service agreement amendments that specify the benchmark query types, DP parameter settings, and data use restrictions applicable to the federated program. Employer participation is designed to be opt-in at the program level and opt-out at the individual query level, ensuring ongoing consent alignment with employer data governance policies.

VI. Discussion

The benchmark utility improvement relative to survey-based alternatives 71.9% reduction in MAD represents a structural advance in the information quality available to employers for benefits design and compensation benchmarking decisions. The practical consequence is that an employer currently relying on an annual SHRM survey benchmark with 13.5% MAD and 18-month lag could instead access quarterly FBIA benchmarks with 3.8% MAD and near-real-time currency. For large employers, benefits design decisions informed by 10-percentage-point more accurate benchmarks have direct financial consequences: a 1% improvement in benchmark accuracy for a 10,000-employee employer with USD 15,000 per-employee benefits cost corresponds to USD 1.5 million in more precisely calibrated benefits investment.

The privacy-utility tradeoff at $\varepsilon = 1.2$ achieves what we characterize as the enterprise deployment sweet spot: formal privacy guarantees strong enough to support alignment with data protection principles under CCPA and GDPR while preserving benchmark utility sufficient for decision support at the employer level. The CESNR > 3.5 finding at this setting is particularly significant: it demonstrates that the DP mechanism does not introduce noise that drowns the genuine cross-employer variation that employers are seeking to observe.

A limitation of the present evaluation is the synthetic nature of the employer ecosystem. While synthetic data generation was calibrated to documented industry benchmarks, the full heterogeneity of real employer HR data including systematic differences in HR information system data completeness, benefits plan design complexity, and payroll reporting precision may affect FBIA's performance in production deployment. Pilot evaluation on a real platform employer ecosystem, under appropriate consent and governance frameworks, is the required next step before production deployment.

VII. Conclusion

The Federated Benchmark Intelligence Architecture (FBIA) described in this paper is a privacy-preserving analytics system, which is specifically designed to generate cross-organizational insights in multi-employer platform ecosystems without any data pooling or bypassing organizational data protection and security boundaries. FBIA converts fragmented data on employers into decision-grade

benchmark intelligence with strict, formally defined privacy assurances through the integrated application of secure multi-party aggregation, differential privacy, and contextual calibration.

Empirical experiments show that at moderate level of privacy (that is, ε = 1.2, δ = 10 -5), FBIA can achieve mean absolute deviation of 3.8 percent - much more accurate than the traditional survey-based benchmarking methods, which are generally constrained by sampling bias, reporting lag, and inconsistent data definitions. This accuracy is ensured in conjunction with good privacy assurances so that the contribution of any individual employer will not be inferred by published outputs.

One of the most important contributions of the framework is its operational level of governance, which consists of four quantitative measures: Benchmark Representativeness Index (BRI), Privacy Budget Utilization Rate (PBUR), Calibration Accuracy Score (CAS), and Cross-Employer Signal-to-Noise Ratio (CESNR). The sum of these metrics gives sustained insight into the utility-privacy tradeoff in production deployments, allowing platform operators to actively manage this tradeoff.

More generally, FBIA proves that high-volume employer ecosystems - as upheld by HR, payroll, and benefits platforms - can be repurposed into valuable intelligence networks without diminishing data sovereignty. The architecture allows organizations to use correct, contextually relevant benchmarks and ensure strict data isolation between underlying assets and adherence to regulatory requirements governing employee information.

In short, FBIA does not, by merely reconstituting multi-employer data as assets, simply reconstitute the data as assets, but make it a strategically controlled one: able to deliver quantifiable benefits to employer decision-making, cost control and competitive positioning, without data sharing, without exposure risk, and without compromising the trust on which employer-client relationships depend.

References

1. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in Proceedings of AISTATS, 2017, pp. 1273–1282.
2. P. Kairouz et al., "Advances and Open Problems in Federated Learning," Foundations and Trends in Machine Learning, vol. 14, nos. 1–2, pp. 1–210, 2021.
3. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, pp. 1–19, 2019.
4. C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating Noise to Sensitivity in Private Data Analysis," in Theory of Cryptography Conference, 2006, pp. 265–284.
5. Apple, "Differential Privacy Overview," Apple Technical Report, 2017.
6. U. Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response," in Proceedings of ACM CCS, 2014, pp. 1054–1067.
7. C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," Foundations and Trends in Theoretical Computer Science, vol. 9, nos. 3–4, pp. 211–407, 2014.
8. A. C. Yao, "How to Generate and Exchange Secrets," in Proceedings of the 27th IEEE Symposium on Foundations of Computer Science, 1986, pp. 162–167.
9. O. Goldreich, S. Micali, and A. Wigderson, "How to Play Any Mental Game," in Proceedings of the 19th ACM Symposium on Theory of Computing, 1987, pp. 218–229.
10. A. Shamir, "How to Share a Secret," Communications of the ACM, vol. 22, no. 11, pp. 612–613, 1979.
11. K. Bonawitz et al., "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in Proceedings of ACM CCS, 2017, pp. 1175–1191.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

12. D. Ulrich, W. Brockbank, A. Yeung, and D. Lake, "Human Resource Competencies: An Empirical Assessment," Human Resource Management, vol. 34, no. 4, pp. 473–495, 1995.

13. P. Cappelli, "Rethinking the 'Push' and 'Pull' of Labor Market Benchmarking," ILR Review, vol. 68, no. 3, pp. 596–622, 2015.

14. California Consumer Privacy Act, Cal. Civ. Code § 1798.100 et seq., as amended by CPRA (Proposition 24), 2020.

15. Illinois Biometric Information Privacy Act, 740 ILCS 14/1 et seq., 2008.

16. European Parliament and Council, "General Data Protection Regulation (GDPR)," Official Journal of the European Union, 2016.

17. Kaiser Family Foundation, "Employer Health Benefits Survey," KFF Annual Survey, 2024.