

How to Synthesize Data When You Don't Have It: A Deterministic Rule-Based Framework for Deriving Enterprise Metrics That No Source System Records

Santosh Nagnath Mehtre

Independent Researcher, Data Engineering, United States of America

ARTICLE INFO

Received: 02 Sept 2025

Revised: 18 Oct 2025

Accepted: 28 Oct 2025

ABSTRACT

Measures such as the cost to serve an individual customer, profit by distribution channel, emissions per unit of product, are regularly taken as enterprise decisions, which are not reported by any source system. Faced with this gap, the common response by practitioners is to give up on the analysis or resort to generative synthetic-data methods that attempt to address some other problem. The paper formalizes a third method deterministic, rule-based data synthesis whereby operative data is used combined with the explicit and versioned business rules to form measures that were never directly measured. The method is broken down into three modes of composition, direct fetch, driver-based allocation and manual percentage allocation. It has roots in activity-based costing, however, the framing focuses on cloud data warehouses and orchestrated transformation pipelines as opposed to purpose-built costing software. A worked example is an end-to-end tracing of profitability of retail-bank customer, a combination of direct revenue capture, branch, contact-centre and digital channel driver-based cost allocation, a regulatory capital charge based on risk-sensitive assets, and an overhead layer based on judges, resulting in a calculated per-customer profit amount that lives nowhere in the source system. A reference implementation to a commodity cloud stack is defined, in which business rules in business-owned configuration are independent of transformation logic in SQL and orchestration logic in directed acyclic graphs, causing one engine to derive any limited metric. The proposed four-condition applicability test separates the problems to which deterministic synthesis is appropriate, and those to which statistical, probabilistic or generative techniques are required, along with the governance controls making the judgement layer auditable, as opposed to arbitrary.

Keywords: Data Synthesis, Cost Allocation, Activity-Based Costing, Customer Profitability, Funds Transfer Pricing, Driver-Based Allocation, FP&A, Data Engineering

I. INTRODUCTION

Background

The synthesis of data that is not recorded by the system with the composition of data recorded by the system and business categories is rule-based data synthesis. The result is predictable: with the same input and rules, the same numbers will be obtained. No model, no probability distribution, no inference, just a proclaimed logic which should be checked and challenged by a reviewer. The issue is not new, many times these types of issues arise. The same gap was faced in the late 1980s by management accounting. The traditional cost systems were using a single volume based measure which was usually the direct labour hours to allocate overhead and generated product costs that were arithmetically defensible and economically meaningless [20]. The reaction of Cooper and Kaplan was an activity-based costing: determine the resources-consuming activities of an organization, locate a measurable driver of each, and assign the cost pool along that driver instead of a convenient average [1]. The next advancement time-driven activity-based costing, which gave the cost of capacity in units of time and consumption in units of time, had lessened the survey burden that had rendered first-generation applications unfriendly to sustain [2], [3]. Meanwhile the economics literature had worked out a formalized axiomatic treatment of cost and surplus sharing, and identified formal properties of cost sharing, including balance, symmetry, additivity, on which any plausible rule of cost sharing should be based [4]. The only difference is the location of the work. Activity-based costing was

utilized in special costing packages, which were monthly run by finance and were not actually visible to the engineering organization [21]. This allocation logic can now be executed in cloud data warehouses, is written in SQL, and the logic is under version control, much like application code, and is scheduled by the same scheduler as all other pipelines [22]. The fact that migration falls under the purview of the data engineer is why the technique should be rephrased in engineering lingo.

Research Aim

This study seeks to formalize deterministic rule-based data synthesis as a composable, implementable pattern of deriving enterprise metrics that are not in the underlying systems and to describe when it is the correct approach.

Objectives

- To study the evolution of the pattern of allocating costs to drivers and the formal characteristics of this allocation, and to determine the history of the suggested pattern.
- To assess customer profitability measurement as the canonical use, to incorporate distributional results as reported in the empirical literature.
- To differentiate between deterministic rule-based synthesis and generative synthetic data, the name of which they share but not their purpose.
- To suggest a reference implementation architecture and an applicability test that distinguishes between problems that can be solved with deterministic synthesis and problems that need statistical or generative algorithms.

Problem Statement

A good question is posed to an analyst.

- How much would it cost to serve this customer?
- How much carbon is used in each product?
- What is the actual profit by channel?

In most cases the answer is that there is no such data in place in any source system, in the warehouse, or as the result of an existing report. Neither of the existing reactions is satisfactory. The first is the abandonment: The question is set unanswered and the decision is made by intuition [23]. The second is inappropriate usage like the term synthetic data has been preemptively reidentified with generative models [9]. According to [10], authors who want to apply that approach find themselves producing statistically plausible fake records, handy to train their models and safe from privacy concerns. This and that difference is filled with an old, broadly used, and often neither described in engineering language, technique [24]. A bank can tell exactly how much earnings a customer brings it in the form of interest, fees, interchange or whatever, until it is tagged to a source system and is associated with the customer [25]. It is unaware of the cost of each customer so far as branch operations, contact-centre handling, digital channel capacity, compliance screening and regulatory capital are not individually assigned costs, but are defined as pools. The institution cannot determine which relationships are profitable, which are loss-making, and which are worth retaining without combining them into two categories. The identical gap in structure occurs in cost-to-serve analysis [7], channel profitability, and in product-level emissions accounting [14].

Novel Contribution

Management accounting literature forms a driver-based allocation as a costing field [1] to [3] and customer profitability as an application [5] and [8]. The economics literature forms an axiomatic approach to the concept of allocation [4]. The machine learning literature forms a generative approach to synthetic data [9] and [10]. This work incorporates the three operational modes into a unified composable model of three modes and shows how any bounded metric can be produced by changing the configuration of a single engine. This also provides an explicit test of applicability the lack of which is what causes practitioners to apply generative methods to deterministic problems.

II. LITERATURE REVIEW

Discuss driver-based cost allocation estimation

The contribution proved the systematic distortion of product cost resulting when overhead is allocated based on one volume-based measure of cost against another, which over-costs high volume simple and under-costs low volume complex products

and the solution lies in tracing costs using activity-specific drivers [1]. The time-driven allocation was refined to deal with the major implementation barrier, the cost of collecting activity surveys, by formulating capacity cost as time cost and consumption equations, as estimated time equations [2], [3].

The economics literature handles the same problem axiomatically and poses the question of what properties a rule of allocation should have, and which rules should uniquely have those properties. There must be a certain sum involved in the allocation; this is balance; there must be the same shares to the same parties; this is symmetry; there must be addition of a combined pool [4]. These properties actively concern the problem of the practitioner, since an allocation, that defies the balance, would defy cost and an allocation, which defies the symmetry, would be inexcusable in review. The utmost silence with respect to the case dealt with here is, however, the problem of allocation where there is no measurable driver in any sense and which must be replaced by judgement.

Analyze the profitability measurement of customers and distributional results

The canonical use of driver based synthesis and origin of the most consequential empirical result concerning it is customer profitability analysis. Using activity-based costing to customer relationships generates a cumulative curve of profitability of typical shape - steep up through the most profitable customers, flat in the middle and steep down through the unprofitable tail - often known as the whale curve [5]. This means that total profit often would well be more than a hundred per cent of realized profit, and some minority of relationships would be subsidizing a loss-making balance.

Empirical study of retail banking supports the trend, and extends its direction into a segmentation technique, examining the Nordic retail bank customer bases and defining profitability distribution as the key determinant in segmenting a customer base [6]. It is later developed that measurement can be extended to cost-to-serve under distribution environments, and that a graphical account of service and cost in mapping out profitability results in a significantly more complete account of profitability, as opposed to a traditional gross-margin account [7]. Contingency frameworks that are syntheses of the interdisciplinary literature would argue that the right complexity of a customer profitability measurement system is determined by the complexity of the environment and that it is not justified to uniformly elaborate measurement [8].

Differentiate deterministic synthesis and generative synthetic data

The expression synthetic data has been used since the introduction of generative adversarial networks to refer to records, namely artificially generated and statistically plausible, yet fictitious customers, transactions or patients, created to either train models or maintain privacy. The idea of survey work has made modalities such that the field is largely seen through the lenses of privacy, utility and fidelity, which are being asked to discover data that is sufficiently similar to real data in order to be useful without revealing information about any particular person [10].

The difference between it and the current topic is not terminological one but rather that between what the output is intended. Generative synthesis generates records that look like reality, and are judged based on being fidelities, deterministic synthesis generates measures based on reality, and judged based on traceability. Customer record created cannot be reconciled to a general ledger. An allocated cost could, and should, since it will be presented before an audit committee. To mix up the two, practitioners resort to probabilistic machinery with probabilistic arithmetic declared, and the separation should be a first-order design consideration, as in this study.

Determine governance issues with allocation based on judgement

In the absence of any defensible driver such as brand marketing cannot be attributed to any product, executive and corporate functions cannot be attributed to any customer, allocation must be based on stated judgement. The costing literature admits arbitrary allocation as a shortcoming instead of a planned element with its own controls, and the axiomatic literature does not provide much advice where a basis to measure does not exist at all [4].

The regulatory history is even more instructive, in that it establishes that artificial internal figures are manageable as opposed to being acceptable. Funds transfer pricing is an internal rate that a treasury functional purchases deposits to the funding businesses, and sells funds to lending businesses, is a completely synthesized price that has no external market analogue. The Basel Committee expects institutions to charge liquidity expenses, benefits and risks to each material business operation [11] and supervisors of the United States have provided specific interagency directions on the practice of funds transfer pricing of large institutions [12]. The attribution of regulatory capital is based on risk-weighted assets in a similar manner, but on a

quantity that is specified in the framework and calculated instead of observed [13]. Product level emissions accounting is needed outside financial services to define allocation procedure because at product granularity one cannot measure emissions directly [14]. In both instances where synthesized figures have a consequence and are open to external scrutiny, it proves that the layer of judgement is controllable and what it takes to govern the judgement.

Literature Gap

A significant body of literature has explored driver-based cost allocation, customer profitability measurement, synthetic data generation and governance mechanisms, but little has been done on how deterministic synthesis can be combined with judgement-based cost allocation in the presence of uncertain drivers. In the past, there is no system-based structure in executing traceability, transparency, and auditability of judgement-based allocations. Thus, the gap that will be covered by the study is the creation of a controlled estimation method with a level of analytical accuracy and governance need.

III. METHODOLOGY

Research Design

The research uses a design and specification as opposed to a controlled experiment. The synthesis pattern is broken down into component modes, the modes are assembled in a demonstrative derivation end to end, an implementation architecture is defined, and conditions of applicability have been derived. The example which has been worked in this section is illustrative, involving representative figures to show mechanism. It is not an empirical case study, neither asserting, nor suggesting, that its values are based on those of a particular institution.

Three Synthesis Modes




Fig. 1: The Three Modes of Synthesis

Direct fetch: The measure is already stored in a source system, marked at the necessary granularity and it is read without transformation. At the time, a customer pays interest on a personal loan, the transaction would be recorded in the loan servicing system under the account of the customer and the interest revenue by the customer is read directly. The only rule is that what can be read should not be assigned. This is mentioned since the distribution of information that is already on the necessary grain is one of the common and preventable errors.

Driver-based allocation: The aggregate is known yet not distributed across a dimension that is needed. A driver (a measurement associated with the aggregate, and cleaving clean along the dimension) is established and the aggregate is apportioned proportionately. A contact centre performing several products on a call basis, will be charged as a single amount, product driven by call volume:

Total contact centre cost		\$1,000
Calls by product (the driver)		
Credit Card	600 calls	
Personal Loan	300 calls	
Savings	100 calls	
Call mix	60% : 30% : 10%	
Allocated cost		
Credit Card	\$600	
Personal Loan	\$300	
Savings	\$100	

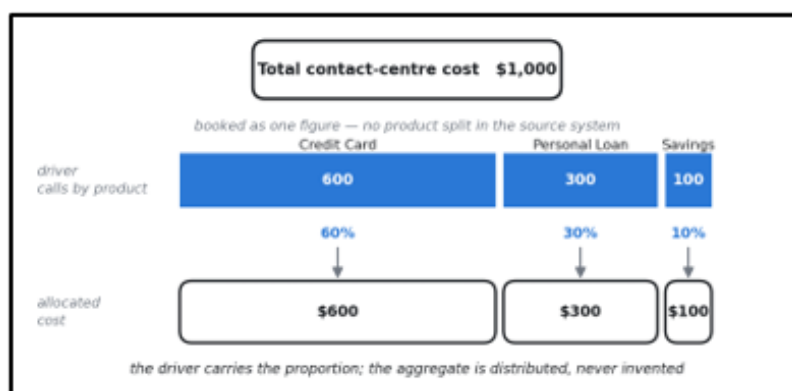


Fig. 2 The traces mechanism

Validity is controlled through selection of drivers. A driver is defensible when, the cost is, in fact, proportional to it. The call volume of the contact-center cost, the withdrawal volume of the ATM network cost, the score on transaction risk of compliance cost. An ill-chosen driver gives out arithmetically consistent but substantively false figures that the activity-based costing was designed to eliminate [1].

Manual percentage allocation: Some pools may not have a defensible driver. Brand marketing cost-TV, sponsorship, outdoor-cannot be identified with any product by any observable indicator. The split of the business is declared in such cases:

Brand marketing spend		\$10M
Declared		allocation
Credit Cards	35%	→ \$3.5M
Personal Loans	25%	→ \$2.5M
Savings	20%	→ \$2.0M
Wealth Management	20%	→ \$2.0M

These percentages are those that the business owns and no system knows, what products the companies are advertising, what types of people the brand advertisement is targeting, what the business strategic priorities are. Since they are not in code but configuration they are updated by the responsible business owner (without the involvement of the engineering mind) and the derivation indicates the modified split on the subsequent performance.

Mode is not an animated element but an element of design. Certain costs actually do not have a formulaic allocation and to claim that they do misrepresents. The governance needed by the mode includes a named approver, a specific revision cadence, documenting a rationale, and an audit trail of changes. The failure mode does not involve manual allocation, but uncontrolled

manual allocation that is reduced to a process of bargaining results as opposed to measuring. Most measures derived are all mode combinations. The method is the composition; the modes singly are not new.

Worked Example: Customer profitability of a Retail Bank.

The case is of a retail customer of this instance referred to as **Customer A** who had three products in a mid-sized bank.

Direct fetch — revenue

Savings balance	\$50,000	→ deposit margin	
	\$1,500		
Credit card balance	\$5,000	→ interest income	
	\$800		
		→ card interchange	\$200
Personal loan	\$30,000	→ interest income	
	\$3,000		

Total direct revenue			\$5,500

Each figure is tagged to the customer in a source system and no allocation is needed.

Driver-based allocation — cost

Branch visits	8	×	\$25 per visit	=	\$200
Contact center calls	3	×	\$15 per call	=	\$45
App logins	240	×	\$0.20 per session	=	\$48
Risk-weighted assets	\$28,000	×	10% capital cost	=	\$2,800

Driver-allocated cost					\$3,093

These costs are represented by the cost pools that grow as each driver. Branch operating cost is contributed by branch visits, contact-centre cost is contributed by calls, application session contributes to the digital channel cost. The cost of regulatory capital is driven by risk-weighted assets: capital is required against the loan under the Basel framework III regime [13] and the cost is charged where it occurs.

Manual percentage — corporate overhead.

Customer A warehouses in the category of mass retail.

Mass retail has been proclaimed to shoulder half of group overhead in business.

Customer A's share of segment overhead (revenue-weighted) is \$600.

There is no formula of attributes towards a specific customer in executive functions, legal, strategy and innovation. The business proclaims the basis, consents to it, and reforms it with the change of priorities.

Derived customer profit

Direct revenue	\$5,500
Driver-based cost	-\$3,093
Manual overhead	-\$600

Customer profit \$1,807

The resultant number is present in no source system. The measure is to be the determinant of relationship management, up-market and retention investment; gross revenue of \$5,500 of its own is not a very weighty matter of decision.

Reference Implementation Architecture




Fig. 3: The arrangement of Components

Business rules are stored as spreadsheets in a business-controlled document store such as definitions of drivers, proclaimed percentage partitions, mapping and part tables such that the responsible owner amends them, with no engineering intervention. In a cloud data warehouse [16], co-located with the source data and configuration tables are the logic allocating direct fetch, allocating driver-based, allocating declared split, and a final routine combining the resulting derived data. Object storage [17] stores executable artifacts and definition of orchestration, and the version control is in commit versioning with Git. Directed acyclic graph is orchestrated [18] in managed service [19], refreshing configuration, running the transformation routines, in dependency order, round trip sequential where there is dependence, and independent where there is independence, and realizing the derived dataset.

It is the separation of concerns that generalizes the architecture. Business rules are stored in configuration, logic of transformation in SQL, orchestration in definition of the DAG. Since these are independent, any limited measure is obtained out of the same engine: profitability of customers in the current period, cost per product in the next period, emissions in the next period, stock-keeping unit in the next period. The architecture remains unchanged, and the rules and output schema change.

Evaluation Strategy




Fig. 4: The test is mentioned as the test of decision

The trend is compared with four conditions of applicability, each of which is a condition and must be met:

- **Justifiable aggregate:** The starting cost, revenue or emissions pool is reliably captured somewhere even though it may not be at the required grain. It cannot allocate the total that is not reliable.
- **Honest drivers:** There is at least one measurable signal which is really correlated with the distribution required. Presentation as opposed to analysis is required where the driver must be invented.
- **Justifiable Production:** It will justify the figure before an auditor, regulator, board or customer and all units must justify the figure by rules, using inputs and configuration. Transparent models are not in a position to meet this.
- **Limited use case:** These are enumerable inputs, outputs and rules that can be customer profitability, cost per product, emissions per unit, channel profit and loss. Pattern discovery and open-ended forecasting are two different issues.

The following exclusions are associated with them. There is no useful aggregate, there is nothing which can be derived. In cases where the output of a non-linear relationship of inputs is essential, default propensity, likelihood of fraud, risk of churn, a statistical or machine learning model is necessary. Generative approaches are used where fictional records are required in order to learn models [9], [10]. In regions where uncertainty has to be modeled actuarial pricing, simulation and risk capital probabilistic techniques are needed. In situations where sub-second inference speed is a must, batch allocation architectures cannot be used.

IV. RESULT AND DISCUSSION

Derived Output of the Worked Example

Table 1: Stationary customer profit by synthesised mode

Mode	Component	Amount	Source of the value
Direct fetch	Deposit margin, interest, interchange	+\$5,500	Tagged in source systems
Driver-based	Branch, contact centre, digital channel	-\$293	Aggregate pools ÷ measured drivers
Driver-based	Regulatory capital on risk-weighted assets	-\$2,800	Basel III framework [13]
Manual percentage	Corporate overhead share	-\$600	Business-declared, approved
Derived	Customer profit	\$1,807	Exists in no source system




Fig. 5: Same of decomposition

The piece of writing shows the major statement. Direct fetch is the supplier of revenue side in totality. Most of the cost side, namely a total cost of 3,093, is provided by driver-based allocation. Of the balance of cost, 7% in the cost of labor, and the residue or balance of that, the manual percentage furnishes, and supplies, the rest of the cost, about 6% of it, not so much that it must be governed, nor so little that judgement comes to control the outcome. The latter two characters are not found anywhere on the source system unless they are the creation of some derived profit, which only exists in composition.

Distributional Consequence




Fig. 6: The shape

This is a constant experience under such derivation when carried through a full customer base once the whale curve of cumulative profitability has been followed such as the top decile or two contribute more than a hundred per cent of all the profit, which is a widely breaking middle, and then the breakers appear at the ends of the Swatters [5]. The pattern is supported by empirical analysis of the customer bases of retail banks and profitability distribution is established as a segmentation criterion [6]. The implication is that there is a significant difference in the gross revenue rank and profit rank, and providers with derived profitability have no choice but to deal with relationships, pricing and retention on the wrong ordering. They have not shown how to distribute, act upon it, or base products on it - and see dimension of margin compression without finding out what the cause of it is.

DISCUSSION

The framework gives rise to three observations.

First, the decisive design decision is to select a driver and not to be architecturally sophisticated. The three modes are arithmetically trivial; the theory work is to prove that an explicit driver is truly a product of the cost pool it represents. It is the failure mode that has been reported in two decades of costing literature [1], [3], and no quality of implementation can make up the difference.

Second, the most conceivable component that practitioners tend to hide and have the most need to regulate is the manual percentage mode. Making declared allocation a structured feature, with name ownership and defined upgrade frequency and auditability transforms accepted weakness into a managed one. Supervisory competence proof that synthesized internal figures can be governed that regulators can control [11], [12].

Third, the pattern is propelled beyond a single use case by the architectural isolation of rules, logic and orchestration. The marginal cost of deriving an extra metric is rule authorship and not system development where business rules are configured in a business-owned location [21].

BOUNDARY CONDITIONS AND LIMITATIONS

The pattern is only applicable where all of the four conditions are met; the exceptions therein are not incidental. The contingency literature further warns that the sophistication of measurement needs to be in tandem with environmental complexity and not enforced consistently maximized [8].

There are three limitations to this study.

1. The work is not empirical but illustrative: the figures chosen were to illustrate mechanism, and it is not said that they are based on a particular institution or that the profit derived has been checked against an external measurement.
2. There was no comparison of derived measures with a comparison of any other method and no quantification is provided of derivation error due to any misspecification of drivers which is the major risk in the framework but is not measured.
3. Lastly, the governance suggestions of the manual layer are prescriptive and not evaluated; they have not been experimented to understand their effectiveness in using allocation negotiation.

V. CONCLUSION AND FUTURE RESEARCH

Conclusion

A synthesis of data is not what makes an analysis come to an end. It will start a new analysis, where the answer is made up of what has been written, by rules that are stated, drivers that are justifiable, and judgement where judgement is the only good faith. This research has established that composition in three modes direct fetch, driver-based allocation and manual percentage allocation and has shown that a majority of consequential derived measures in large enterprises are made up of combinations of the three. It has demonstrated using an illustrative end to end derivation how a figure, representing customer profit, found in no source system, is constructed using figures which are and has identified a commodity architecture where the distinction between business rule and transformation logic and orchestration allows a single engine to compute any bound measure. The method needs neither generative models nor probabilistic inference nor a specific platform. It demands a clean data stack, few carefully chosen drivers, and control group ownership of the judgement layer. In cases where a hearer is informed that there is no data, that is no longer useful than to terminate the inquiry but to ask what there is out of which it can be produced.

FUTURE SCOPE

The former would be empirical validation like implementing the framework on a complete customer base in a cooperating institution and comparing the resulting profitability with measures independently built would validate accuracy, as opposed to mechanism. The second one is quantification of driver sensitivity, producing mechanisms to bracket the error in measured amounts which are explainable by driver misspecification, so that a derived amount can bear an acceptable confidence interval as opposed to a deceptive aura of precision. The third relates to the judgement layer that are the suggested governance controls in practice restrictive to the negotiation of allocation, and stated percentages can be partially substituted by guessed ones. The fourth is the extrapolation of the trend to new areas of cost allocation where allocation currently remains infantile, such as artificial intelligence infrastructure attribution, where inference and training costs need to be shared across products, teams and customers with no standards of driver conventions established.

VI. REFERENCES

- [1] Cooper, R. and Kaplan, R.S., 1988. Measure costs right: make the right decisions. Harvard Business Review, 66(5), pp. 96–103. Available at: <https://hbr.org/1988/09/measure-costs-right-make-the-right-decisions>

- [2] Kaplan, R.S. and Anderson, S.R., 2004. Time-driven activity-based costing. *Harvard Business Review*, 82(11), pp. 131–138. Available at: <https://hbr.org/2004/11/time-driven-activity-based-costing>
- [3] Kaplan, R.S. and Anderson, S.R., 2007. *Time-Driven Activity-Based Costing: A Simpler and More Powerful Path to Higher Profits*. Boston, MA: Harvard Business Review Press.
- [4] Moulin, H., 2002. Axiomatic cost and surplus sharing. In: *Handbook of Social Choice and Welfare*, Vol. 1. Amsterdam: Elsevier, pp. 289–357. DOI: 10.1016/S1574-0110(02)80010-8.
- [5] Kaplan, R.S. and Narayanan, V.G., 2001. Measuring and managing customer profitability. *Journal of Cost Management*, 15(5), pp. 5–15.
- [6] Storbacka, K., 1997. Segmentation based on customer profitability — retrospective analysis of retail bank customer bases. *Journal of Marketing Management*, 13(5), pp. 479–492. DOI: 10.1080/0267257X.1997.9964487.
- [7] Guerreiro, R., Rodrigues Bio, S. and Vazquez Villamor Merschmann, E., 2008. Cost-to-serve measurement and customer profitability analysis. *The International Journal of Logistics Management*, 19(3), pp. 389–407. DOI: 10.1108/09574090810919215.
- [8] Holm, M., Kumar, V. and Rohde, C., 2012. Measuring customer profitability in complex environments: an interdisciplinary contingency framework. *Journal of the Academy of Marketing Science*, 40(3), pp. 387–401. DOI: 10.1007/s11747-011-0263-4.
- [9] Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y., 2014. Generative adversarial networks. In: *Advances in Neural Information Processing Systems 27 (NIPS 2014)*. Preprint arXiv:1406.2661. Available at: <https://arxiv.org/abs/1406.2661>
- [10] Jordon, J., Szpruch, L., Houssiau, F., Bottarelli, M., Cherubin, G., Maple, C., Cohen, S.N. and Weller, A., 2022. Synthetic data — what, why and how? Preprint arXiv:2205.03257. Available at: <https://arxiv.org/abs/2205.03257>
- [11] Basel Committee on Banking Supervision, 2008. *Principles for Sound Liquidity Risk Management and Supervision*. Basel: Bank for International Settlements. Available at: <https://www.bis.org/publ/bcbs144.htm>
- [12] Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation and Office of the Comptroller of the Currency, 2016. *Interagency Guidance on Funds Transfer Pricing Related to Funding and Contingent Liquidity Risks*. Supervisory Letter SR 16-3. Available at: <https://www.federalreserve.gov/supervisionreg/srletters/sr1603.htm>
- [13] Basel Committee on Banking Supervision, 2011. *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Basel: Bank for International Settlements. Available at: <https://www.bis.org/publ/bcbs188.pdf>
- [14] World Resources Institute and World Business Council for Sustainable Development, 2011. *Greenhouse Gas Protocol Product Life Cycle Accounting and Reporting Standard*. Available at: <https://ghgprotocol.org/product-standard>
- [15] Microsoft Corporation, *Introduction to SharePoint*. Microsoft Learn. Available at: <https://learn.microsoft.com/en-us/sharepoint/introduction>
- [16] Google LLC, *BigQuery documentation*. Google Cloud. Available at: <https://cloud.google.com/bigquery/docs>
- [17] Google LLC, *Cloud Storage documentation*. Google Cloud. Available at: <https://cloud.google.com/storage/docs>
- [18] Apache Software Foundation, *Apache Airflow documentation*. Available at: <https://airflow.apache.org/docs/>
- [19] Google LLC, *Cloud Composer documentation*. Google Cloud. Available at: <https://cloud.google.com/composer/docs>
- [20] Bhagwat, V.B., 2025. Simplifying Payroll Balance Conversions in Payroll Systems Implementation through the Use of Generative AI.
- [21] Kumar Adabala, P., 2021. Optimizing ERP modernization: A smart data migration framework approach. *International Journal of Enhanced Research in Science, Technology & Management*, pp.61-72.
- [22] Narapareddy, V.S.R. and Yerramilli, S.K., 2022. Scaling the service now CMDB for distributed infrastructures. *International Journal of Engineering Technology Research & Management (IJETRM)*, 6(10), pp.101-113.
- [23] Pokala, H.K., 2022. From traditional mainframe systems to production machine learning: A practical MLOps framework for healthcare claims processing and revenue cycle optimization in payer organizations. *International Journal of Communication Networks and Information Security*, 14(2), pp.847-857.
- [24] Todupunuri, A., 2024. Develop machine learning models to predict customer lifetime value for banking customers, helping banks optimize services. *International Journal of Advanced Research and Interdisciplinary Scientific Endeavours*, 1(5), pp.275-282.
- [25] Prodduturi, S.M.K., 2023. Leveraging Big Data And Business Intelligence To Revolutionise Corporate Strategy. *International Journal for Research Trends and Innovation*, 8(7).