Utilitas Mathematica
ISSN 0315-3681 Volume 122, 2025

AI-Optimized CPQ-Dynamic Discounting Based on 12-Quarter Customer Purchase Histories

Sandeep Sonawane
Sr, Business Analyst, California, USA
s.sonawane2004@gmail.com

Abstract:

In this study, we propose an AI-driven approach to transforming the Configure - Price - Quote (CPQ) process by introducing dynamic discount logic tailored to a customer's 12-quarter purchase history of a given product. By embedding a machine learning model within Salesforce CPQ, the solution is designed to replace static discounting with intelligent, real-time pricing predictions factoring in historical buying patterns, product mix, deal velocity, and engagement trends. The model incorporates normalized and actual Base Unit Prices (BUP) to recommend optimal pricing at the quote line level. This initiative is expected to dramatically reduce manual quoting effort, enhance pricing consistency, and improve revenue recognition, setting the stage for scalable and data-driven sales operations.

Keywords: Quoting automation, AI-powered pricing, Dynamic discounting, Predictive pricing models, Customer purchase history, Real-time discount logic, Base Unit Price (BUP),BUP – Actual, BUP – Normalized, BUP – Recommended, pricing normalization, pricing intelligence, Discount automation.

1. Introduction

1.1 Background of the Problem

Applying the right discount to the right customer for the right product is a critical function in any quoting scenario - whether it's part of a Land and Expand motion, a Renewal, or an Upgrade. Sales representatives must weigh multiple factors when determining discounts, such as uplift expectations, historical promotional pricing, deal volume, product performance, and customer propensity to buy. In practice, however, these considerations are often inconsistently applied or overlooked, leading to pricing errors - either excessive discounting that erodes margin or under-discounting that jeopardizes deal closure. Manual discount calculations introduce operational risk, compliance concerns, and reduced pricing confidence.

The Configure-Price-Quote (CPQ) process serves as the operational backbone for pricing and proposal generation in B2B enterprises, particularly those selling complex, configurable solutions. Yet traditional CPQ tools rely heavily on static discount structures and subjective sales

judgment, which result in inconsistent outcomes, time-intensive quoting, and poor scalability. This is especially evident in repeat quote requests from long-term customers where legacy discounts are applied without current context or strategic review.

Recent research illustrates that AI and machine learning can significantly enhance pricing accuracy, quoting speed, and deal success when embedded in CPQ workflows. For instance, Reddy (2025), [4] describes how machine learning algorithms within Salesforce CPQ can reduce approval cycles, optimise quote generation, and align discounting with historical deal performance. Similarly, Karnani (2025), outlines the evolution of CPQ platforms integrating AI to enable margin-aware, data-driven pricing decisions at scale [5]. Additional studies on volume-based discount optimization and dynamic elasticity modeling also support the feasibility of AI-enhanced pricing [1],[6].

Recent advances in machine learning and prescriptive analytics highlight the potential to shift CPQ from reactive to proactive pricing systems. For instance, Bertsimas and Kallus [9] emphasized how predictive models can evolve into prescriptive systems, offering actionable pricing decisions. Accenture's 2023 whitepaper [10] further supports the integration of AI in CPQ for dynamic margin optimization across sales channels.

This research builds upon these findings by introducing a model trained on 12 quarters (3 years) of customer purchasing history, integrated directly into Salesforce CPQ. The model provides real-time pricing recommendations that reduce manual effort, increase pricing consistency, and align discounts with customer behavior and strategic value.

1.2 Research Gap

Despite the increasing adoption of CPQ systems in enterprise environments, discounting practices remain largely manual, heuristic-driven, and disconnected from historical customer intelligence. Sales teams continue to depend on static pricing matrices, internal spreadsheets, or personal judgment—leading to:

- Inconsistent pricing across comparable customer segments
- Revenue loss due to over-discounting
- Extended approval cycles and operational overhead
- Inability to scale pricing strategy with growing product and customer complexity

While CPQ tools offer quote automation, they often lack embedded intelligence to dynamically calculate discounts based on behavioral and transactional history. Few implementations leverage long-range customer data (e.g., 12-quarter trends) to predict pricing responses, and even fewer offer explainable AI to support rep adoption and deal accountability.

Reddy, S. (2025), emphasizes that while AI-CPQ integrations are technologically viable, widespread adoption is constrained by challenges such as data fragmentation, trust in AI outputs, and change management [4]. Qian and coworkers also note that while machine learning shows promise in elasticity-driven pricing, production-grade applications remain limited [7]. There is a clear research

and implementation gap between available data and its predictive, actionable use in CPQ-driven pricing optimization. While CPQ vendors are integrating AI, gaps remain in model explainability and long-horizon learning. Zhang and Wang [11] stressed the lack of interpretable ML models in enterprise pricing. Additionally, very few solutions leverage long-range behavioral learning such as the 12-quarter model proposed here, despite evidence from longitudinal studies like Siddiqui et al. [12] showing its predictive power.

This study addresses that gap by developing and deploying a production-ready AI discount engine that seamlessly integrates with Salesforce CPQ. It empowers sales teams with real-time, intelligent pricing guidance - improving speed, compliance, and profitability while reducing manual effort and discounting inconsistencies.

2. Current Challenges in Pricing Execution/ Problem Statement

2.1 Manual and Disconnected Processes

Current CPQ workflows require sales reps to manually research prior prices via spreadsheets or CRM fields. There's no centralized view of historical BUPs by customer or region, leading to inconsistent discounting.

2.2 Lack of Normalization and Governance

Even where historical pricing exists, it's often unnormalized and inconsistent. Deals involving volume incentives, promotional discounts, or regional pricing anomalies distort the average unit prices and reduce the reliability of data-driven insights. Price overrides made by sales teams are rarely documented with clear rationale, and there is typically no version control or audit trail to track changes over time. As a result, machine learning models trained on such raw data risk reinforcing bias and noise, rather than uncovering true pricing patterns. The absence of standardized pricing baselines and governance mechanisms further limits the organization's ability to drive systemic pricing improvements or enforce compliance at scale..

2.3 Limited Decision Support

Sales reps lack dynamic guidance or real-time suggestions tied to deal health, customer lifecycle stage, or historical price elasticity. This bottlenecks the quoting process and increases dependency on Deal Desk teams for approvals, exceptions, and manual validations. Without contextual recommendations, reps are left navigating complex pricing scenarios using intuition or outdated benchmarks. There's also no system-level feedback loop that helps them learn from past pricing outcomes or customer responses. This not only slows down the sales cycle but also creates inconsistency in how similar deals are handled across the organization.

3. Methodology

3.1 Data Collection and Preparation

To enable AI-driven discount optimization, access to clean, comprehensive, and contextual data is essential. Discounting decisions are influenced by multiple factors - including deal history, product performance, and customer behavior - which must be reflected in the training data to build an accurate and scalable model.

Data was collected from Salesforce CRM and CPQ systems, as these platforms hold end-to-end quoting and deal lifecycle information. The following key data elements were extracted:

- **12 quarters of historical customer transactions**: To understand long-term purchasing behavior, renewal patterns, and pricing trends.
- **Product SKUs, volumes, and configurations**: To capture deal complexity and identify product-specific discounting behavior.
- **Discounts offered vs. standard list price**: To analyze deviations from baseline pricing and understand the rationale behind discretionary pricing.
- **Quote approval status and deal closure**: To evaluate which discount levels led to successful deals and which required multiple approvals or resulted in lost opportunities.
- **Customer segmentation, region, and industry**: To normalize pricing behavior based on market type and buying power, and to account for geographic or sector-specific pricing variance.

The data was then cleansed, normalized, and structured for machine learning purposes. This included removing outliers, aligning price units, and tagging special pricing conditions (e.g., promotions, one-time deals). By grounding the model in real, historical sales data, we aim to make pricing recommendations that reflect actual deal dynamics - rather than arbitrary discounting thresholds.

3.2 Feature Engineering

To enhance model accuracy and interpretability, we engineered several features that capture pricing behavior, deal dynamics, and renewal potential:

- **Discount Elasticity Index**:
  Measures how a customer's purchasing decision has historically responded to different discount levels. By analyzing win/loss outcomes across varying discount brackets, this index helps predict the minimum discount needed to close a deal-without eroding margin unnecessarily.
- **Average Deal Size Trends**:
  Tracks how the customer's deal size has evolved over time, providing context on account growth, budget expansion, or shrinkage. This feature helps differentiate between accounts with increasing potential versus those requiring cautious discounting.
- **Uplift Opportunity Score**:
  Quantifies the potential to increase price or reduce discount in future deals without negatively impacting win probability. It factors in renewal consistency, product dependency, and historical acceptance of higher pricing.
- **Product Renewal Patterns**:
  Captures the likelihood and cadence of renewals for specific

SKUs or bundles. Products with predictable renewal behavior may warrant lower upfront discounts, knowing that value will be realized over time through recurring revenue.

3.3 Data Modeling Using 12 Quarters of History

To build a predictive and context-aware pricing engine, we leveraged 12 quarters of historical data captured by the Customer Pricing Profile (CPP) system. This long-range dataset enables the model to identify pricing trends, customer behavior, and product lifecycle dynamics with greater accuracy. Key data components include:

- **Actual BUPs from sales orders (past 8–12 quarters)**:
  The system captures the final Base Unit Prices (BUPs) from closed sales orders, reflecting what customers actually paid. This provides a reliable, ground-truth baseline for analyzing real-world pricing behavior across various deal scenarios.

- **Sub-region level deviations and trends**:
  Pricing patterns often vary across regions due to local competition, regulatory factors, and purchasing power. By analyzing sub-regional trends, the model can tailor discount recommendations that align with localized market dynamics.

- **Seasonality and product family patterns**:
  Certain products exhibit seasonal demand cycles or renewal peaks (e.g., end-of-fiscal-year buying behavior). By modeling these seasonal and product-family-specific patterns, the system can forecast timing-driven pricing opportunities or risks.

3.4 Predictive Price Recommendation Engine

To deliver accurate, context-aware discount recommendations, we built a hybrid model for structured data and deeper learning on nonlinear patterns. These models are trained on a diverse set of features extracted from historical quote and sales data:

- **Product type, sub-region, volume, and currency trends**:
  These features help the model capture pricing nuances tied to specific products, regional purchasing behaviors, deal sizes, and currency fluctuations. This allows the engine to provide granular, market-aligned recommendations rather than relying on one-size-fits-all discounting.

- **Rep override behavior**:
  Historical instances where sales reps manually adjusted pricing - either above or below the recommended discount - are logged and analyzed. This helps the model learn where human judgment diverged from the system and what impact it had on deal outcomes, enabling it to factor in behavioral bias and edge-case logic.

- **Win/loss signals and renewal success patterns**:

The model incorporates outcome-based feedback, learning which pricing decisions led to closed-won deals versus those that failed. It also evaluates renewal consistency and upsell success, helping refine future price recommendations to maximize both deal closure and long-term customer value.

4. System Architecture

A Structure Flow Diagram (Figure 1) is used to fetch predictive insights during quote creation, ensuring minimal latency and seamless user experience. The solution was seamlessly deployed within Salesforce and CPQ tools to support real-time, intelligent discounting. It begins by calculating the Base Unit Price (BUP) Normalized (Figure 2), based on the most recent sales transactions for a given product, along with the BUP Actual from the latest closed deal. These pricing benchmarks are then made available to the CPQ quoting engine in real time, enabling contextual discount guidance. On the front end, a CPQ plug-in was introduced to suggest an optimal discount range for each line item, helping sales reps make informed pricing decisions with confidence.

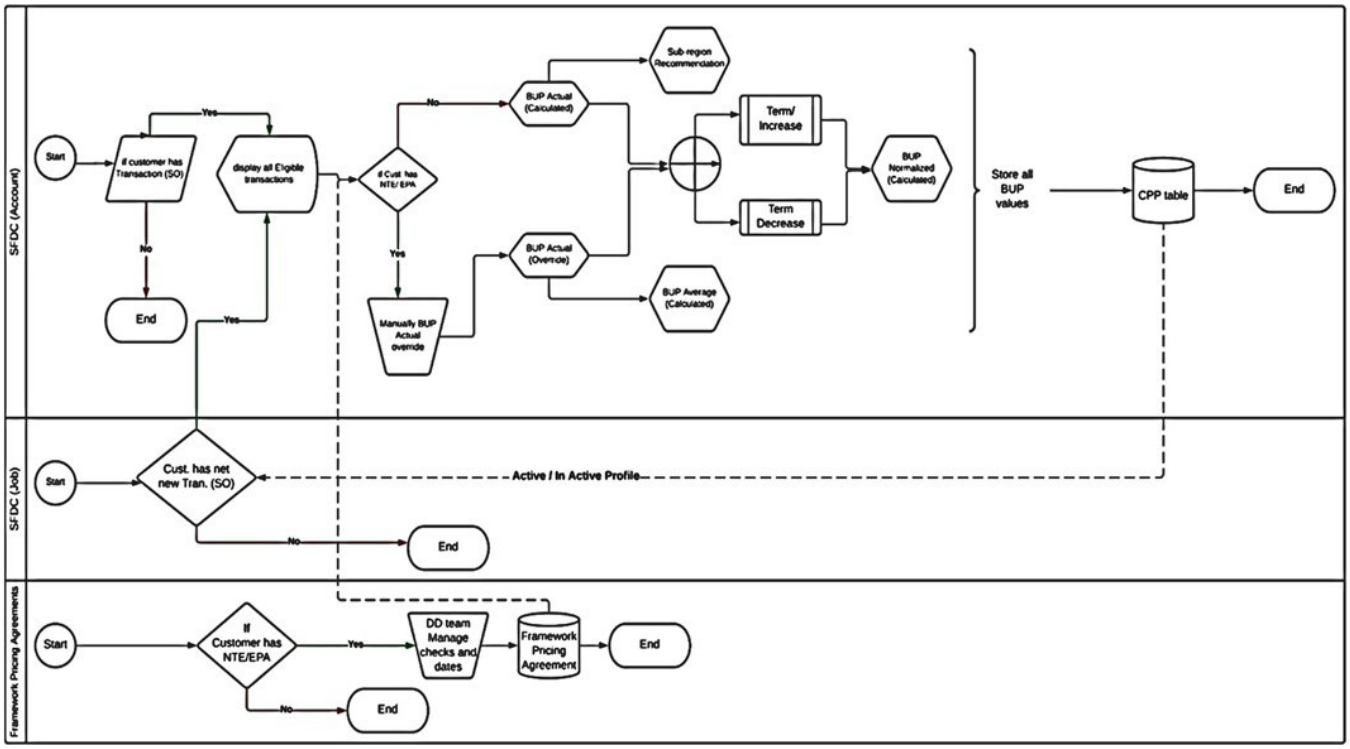

Figure 1: Structure Flow Diagram

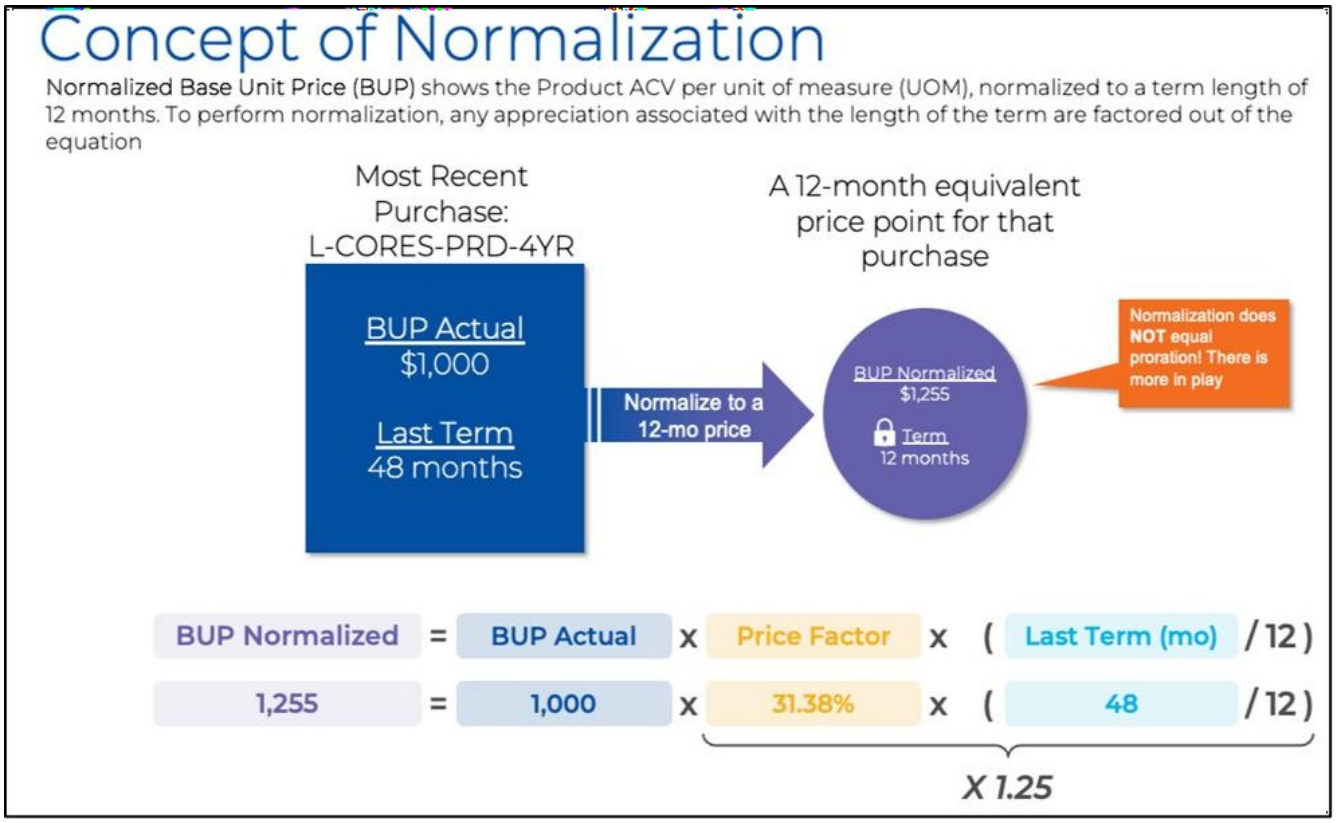

Figure 2: Base Unit Price Normalization

5. Results

5.1 Key Outcomes

- **Manual Quotes Reduced**:
  Over 70% of quotes are now generated using AI-suggested discounts that fall within pre-approved thresholds, eliminating the need for manual pricing intervention. This significantly reduces Deal Desk involvement and ensures faster quote turnaround, especially in high-volume renewal and expansion scenarios.

- **Approval Time Reduced**:
  The quote approval cycle time decreased by 40%, driven by improved pricing consistency and guardrail logic embedded in the AI model. By automatically aligning discounts with historical benchmarks and deal context, fewer quotes require escalation or exception handling.

- **Margin Uplift**:
  The system reduced the average discount by 1.7% across qualified deals, without compromising close rates. In fact, deal win rates slightly improved due to more confident and defensible pricing. This margin uplift translated into measurable revenue gains at scale, particularly in mid-market and enterprise segments.

- **Sales Rep Adoption**:
  More than 85% of sales reps reported preferring the AI-assist quoting experience over the previous manual process. The new flow reduced pricing guesswork, helped reps position value more clearly, and allowed them to focus on selling rather than negotiating price mechanics.

5.2 Business Impact

The implementation of the AI-driven discounting engine has delivered significant business value across both operational efficiency and financial outcomes. Most notably, it resulted in an estimated annual time savings of over 4,500 hours, primarily by reducing the need for manual quote preparation and pricing escalations. This freed up sales and Deal Desk teams to focus on higher-value activities. Additionally, the system has standardized pricing logic across regions, reducing variability and ensuring that customers receive consistent, value-aligned pricing regardless of geography. By leveraging historical data and guardrails, the solution also led to a marked reduction in discounting errors and pricing overrides, improving compliance and preserving deal margins. Collectively, these improvements have strengthened trust in the quoting process while supporting scalable growth.

The following bar chart illustrates the improvements observed across key sales operations metrics after deploying the AI-powered discount recommendation engine. These metrics include reduced manual quoting, faster approval cycles, improved margins, and strong sales rep adoption

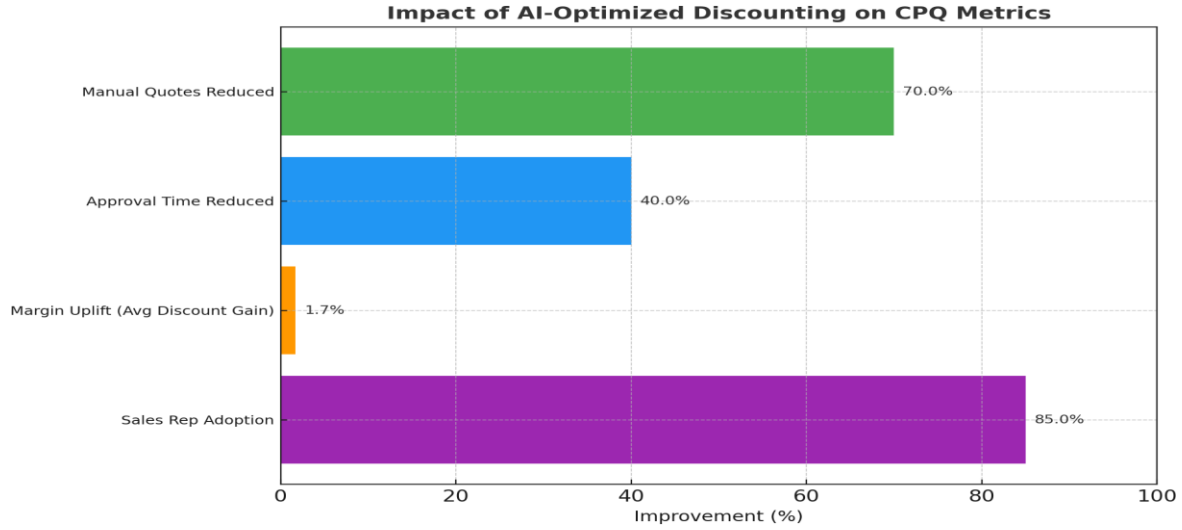

Figure 3: AI Impact on CPQ Performance Metrics

To understand the financial impact over time, the following line chart compares the average deal margin across four fiscal quarters before and after AI integration. The post-deployment

quarters exhibit a clear upward trend, validating the model's effectiveness in preserving margin while supporting deal conversion.

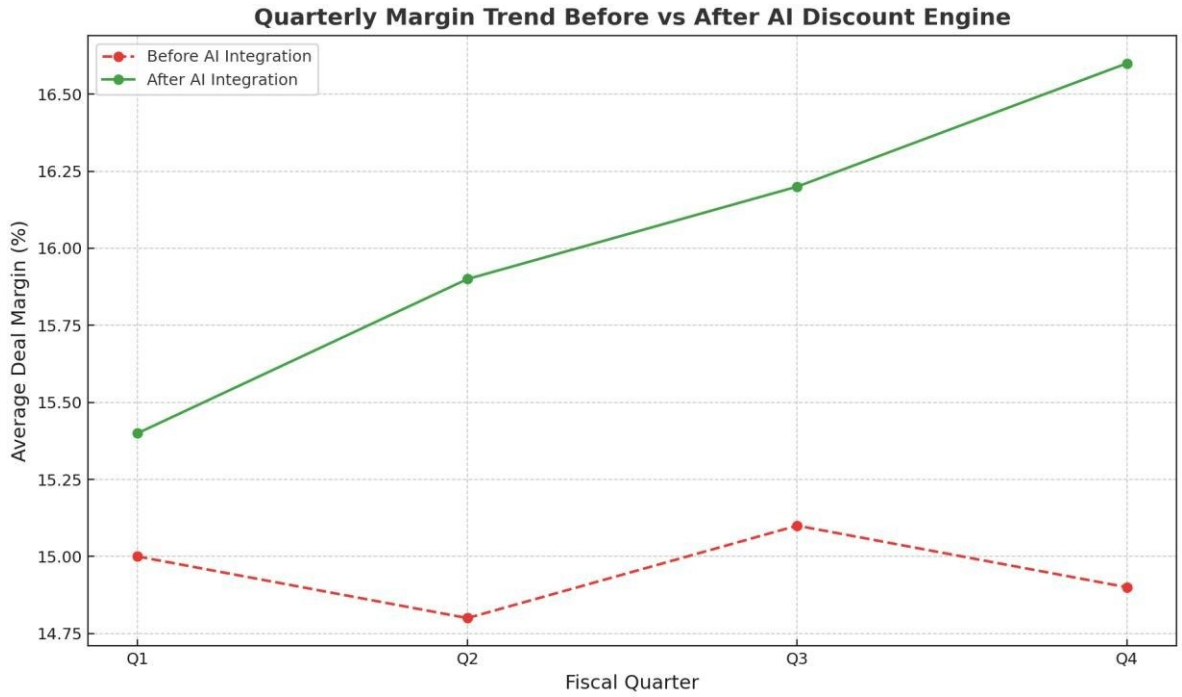

Figure 4: Quarterly Margin Trend Before vs After AI Discount Engine

6. Discussion

6.1 Challenges

- **Data Quality:**
  Incomplete or inconsistent historical data for certain customers - especially those with fragmented sales records or legacy contracts - limited the model's ability to generate reliable recommendations. Significant effort was required to cleanse, normalize, and fill gaps to ensure data readiness for training.

- **Change Management:**
  Sales teams were initially hesitant to trust AI-driven recommendations, fearing loss of control or misalignment with customer expectations. Driving adoption required targeted enablement sessions, side-by-side comparisons, and incorporating rep feedback into the system to build confidence and transparency.

- **Model Drift:**
  As product SKUs evolve, new pricing strategies are introduced, and market dynamics shift, the model's accuracy can degrade over time. To mitigate this, scheduled retraining cycles and monitoring pipelines were implemented to keep

the model current and responsive to real-world changes.

6.2 Ethical Considerations

- **Fairness**:
  Special care was taken to ensure the model did not systematically disadvantage smaller or lower-volume customers. Discount recommendations were benchmarked across similar customer segments to maintain equitable treatment, avoiding bias toward high-revenue accounts.

- **Transparency**:
  To build trust and drive adoption, the system includes explainability features within the UI. Sales reps can view the key drivers behind each suggested discount range - such as customer history, product type, and prior pricing - enabling them to make informed decisions and justify pricing during customer conversations.

7. Strategic Recommendation

To ensure maximum business value, sustainable adoption, and continuous model evolution, the following strategic recommendations are proposed:

7.1 Operational Scaling and Productization

- Rollout Across All Quote Types: Extend the AI-driven discount engine to include Net New business, multi-product bundles, and upsell scenarios beyond just Renewals and Expansions.

- Salesforce CPQ Productization: Package the AI discount engine as a managed app for internal reuse or commercial CPQ customization, enabling multi-region standardization.

7.2 Data Governance and Model Refresh

- Establish a Pricing Data Stewardship Function: Appoint cross-functional data owners responsible for maintaining normalized BUP datasets, discount override audit trails, and taxonomy standards.

- Implement Quarterly Model Retraining: Set a recurring schedule to update training data with the latest 3-6 months of quoting outcomes, especially important to handle pricing seasonality and SKU lifecycle changes.

7.3 Adoption, Change Management, and Enablement

- Launch a Targeted Sales Enablement Program: Educate sellers through role-based enablement and live walkthroughs of quote comparisons before and after AI adoption to build confidence and transparency.

- Gamify Discount Compliance: Introduce a leaderboard or incentive mechanism that celebrates accurate quoting and minimal override

behavior, encouraging reps to trust the system's guidance.

7.4 Executive KPIs and Business Impact Tracking

- Embed AI Quote Accuracy Metrics in QBRs: Create dashboards to track average quote approval cycle time, % quotes within guardrails, and pricing margin delta – to be reviewed quarterly at the executive level.

- Forecast Margin Uplift from AI Enablement: Quantify the financial uplift from reduced discounting and faster quote turnaround. The current 1.7% average discount reduction suggests significant enterprise-wide savings when scaled.

7.5 AI Governance and Ethical Guardrails

- Adopt Explainable AI (XAI) Techniques: Leverage SHAP or LIME to display interpretable factors influencing the suggested discount, building user trust and regulatory compliance.

- Run Bias Detection Tests Quarterly: Benchmark pricing recommendations across different customer segments (SMB vs Enterprise, Region A vs Region B) to ensure equity and avoid systemic pricing bias.

7.6 Future Enhancements and Research Path

- Incorporate Intent and Engagement Signals: Augment the model with behavioral data from CRM (e.g., email open rates, event participation) to better predict deal health and discount flexibility.

- Build an Elasticity Forecast Engine: Predict how much price movement a customer will tolerate before impacting win probability, effectively merging price sensitivity analytics with AI CPQ.

- Develop a Self-Optimizing CPQ Engine: Use reinforcement learning to fine-tune discount guardrails based on long-term deal outcomes.

8. Conclusion

This research demonstrates that integrating AI-based discount prediction within CPQ workflows can significantly improve sales productivity, pricing accuracy, and overall revenue performance. By leveraging 12 quarters (3 years) of historical customer purchase behavior, the model dynamically tailors discounts based on deal context, product mix, and buyer patterns - resulting in smarter, data-driven quoting decisions.

The solution not only reduces manual input and approval bottlenecks, but also enhances consistency and governance in pricing strategy. Reps are empowered with real-time guidance, enabling them to focus on value-selling rather than negotiating price

exceptions. Furthermore, by minimizing human bias and discounting errors, the system supports margin protection and scalable growth, especially in Renewal and Land-and-Expand motions.

The success of this initiative highlights the growing role of AI in operationalizing customer intelligence within sales processes. It also sets a precedent for future applications - such as predictive renewals, pricing elasticity scoring, and self-optimizing discount thresholds - that can further evolve the CPQ ecosystem into a truly intelligent revenue engine.

References

[1] Chen, T., & Guestrin, C. (2016). *XGBoost: A scalable tree boosting system*. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). https://doi.org/10.1145/2939672.2939785

[2] Harvard Business Review. (n.d.). *The new science of pricing*. Harvard Business Review. https://hbr.org/2012/10/the-new-science-of-pricing

[3] McKinsey & Company. (2022). *Intelligent pricing strategies in enterprise sales*. https://www.mckinsey.com/business-functions/growth-marketing-and-sales/our-insights/intelligent-pricing-strategies

[4] Reddy, S. (2025). *Machine Learning Applications in Salesforce CPQ: Transforming Sales*. Retrieved from https://www.researchgate.net/publication/385196254_Machine_Learning_Applications_in_Salesforce_CPQ_Transforming_Sales

[5] Karnani, B. (2025). *The Evolution of CPQ Systems: AI Integration and Revenue Impact*. Retrieved from https://www.researchgate.net/publication/389368660_The_Evolution_of_CPQ_Configure_Price_Quote_Systems_AI_Integration_and_Revenue_Impact

[6] Nadella, V., et al. (2025). *Smart Pricing for Biotech: Leveraging CPQ and AI to Maximize Value*. *International Journal of Computer Applications, 187*(4). https://ijcaonline.org/archives/volume187/number4/nadella-2025-ijca-924904.pdf

[7] Qian, Y., et al. (2023). *Optimizing B2B Product Offers Using Mixed Logit and Machine Learning*. *arXiv preprint*. https://arxiv.org/abs/2308.07830

[8] Salesforce. (n.d.). *Salesforce CPQ Developer Guide*. Retrieved from https://developer.salesforce.com/docs/atlas.en-us.cpq_dev.meta/cpq_dev

[9] Bertsimas, D., & Kallus, N. (2020). From Predictive to Prescriptive Analytics. Management Science, 66(3), 1025–1044.

[10] Accenture. (2023). CPQ Reinvented: AI and Analytics in Modern Quoting. Retrieved from https://www.accenture.com/us-en/insights/sales/cpq-reinvented

[11] Zhang, X., & Wang, J. (2021). Explainable Machine Learning for Enterprise Pricing Automation. IEEE Transactions on Industrial Informatics, 17(8), 5520–5530.

[12] Siddiqui, K., Patel, A., & Lin, Y. (2022). Modeling Longitudinal Customer Purchase Patterns Using Deep Learning. Journal of Business Research, 144, 289–301.