LLM AND MCP BASED AUTOMATED DEAL PRICING NEGOTIATION USING MULTI MODAL MARGIN FORECASTING AND PRICING SCENARIO SIMULATION

Chirag Soni [1], Swati Shah [2], Avi Reddy [3], Mahesh Toro [4],
Rishabh Rishi Sharma [5], Kiran R [6]

[1] Lead AI Product Manager, PayPal, Bangalore, India
[2] Head (Product) – AI PSP, PayPal, San Jose, USA
[3] Sr Director (Product), PayPal, San Jose, USA
[4] Sr Director (Engineering), PayPal, San Jose, USA
[5] Staff Software Engineer, PayPal, San Jose, USA
[6] Sr Software Engineer, PayPal, San Jose, USA

ABSTRACT

Enterprise deal negotiation continues to present persistent challenges in modern business environments. The process itself remains highly manual, relying majorly on individual expertise rather than on broad, data-driven analysis. This approach becomes increasingly untenable as the volume and complexity of pricing scenarios grow, and as organizations face heightened competitive and operational pressures. Negotiators often base decisions on limited historical pricing or isolated financial data, overlooking emerging factors such as market mood, regulatory changes, customer willingness to pay, and peer benchmarking. The proliferation of AI agents has opened new opportunities for automating complex business processes. This paper presents our work on enhancing end-to-end deal negotiation through the integration of multiple AI systems via a Model Context Protocol (MCP) server. Our approach combines traditional machine learning with large language models to provide multi-modal margin forecasting and pricing scenario simulation, which serve as critical inputs for negotiation decisions. We demonstrate how consolidating financial health assessment, market sentiment analysis, pricing intelligence, and margin forecasting through a unified MCP framework can significantly improve negotiation outcomes while reducing cycle times. The system addresses key challenges in sales operations where human negotiators often miss critical data points due to time constraints and information silos across departments.

KEYWORDS

Pricing Optimization, MCP Server, GenAI, LLMs, Artificial Intelligence, Multi Modal Forecasting

1. INTRODUCTION

Deal Negotiation is a time taking and inconsistent process dependent on expertise of sales teams. Especially when there are multiple inputs to the negotiation process, we as human beings cannot comprehend every input that should go into the negotiation and miss out on important parameters that can drive the deal negotiation in positive direction. While trying to negotiate a better pricing to increase margins, we would be looking at a limited set of data which is available during that time, however a lot of metrics that could prove critical during negotiation might be totally overlooked. For example while upselling our products to an existing business customer, our sales teams may do a deep research on their historical engagement and come up with an understanding around whether higher discounts could be provided. However they may be overlooking that customer's current market performance and external indicators like sentiments and financial health at the given point in time, which could totally drive the discussion in another direction. Another major input here will be to generate margin forecasting based on multiple scenarios like sensitivity to price changes, and optimal price point that works for both the parties, floor pricing which will indicate the lowest we can go in the negotiation, prevailing market conditions for the industry they come from, macroeconomic indicators, and several other factors.

It becomes humanly impossible to gather this data quickly at the time Merchant initiates the conversation because this time period becomes critical in the sales lifecycle. There are multiple teams that need to get involved to gather the data and provide a single consolidated view for the sales teams to continue their negotiation journey. The final output, or interpretation of the data, depends on individual and will not have a consistency which affects the negotiation based on who is carrying it out.

2. METHODOLOGY

2.1. Drivers of Deal Negotiation

We will first look at the most important metrics that are needed to carry out a negotiation. In the next section we will discuss how to consolidate everything and provide guidance to Sales teams based on information collated.

- Current Pricing Information : for on boarded Businesses using your products, how has their prior negotiation been on products already configured. This acts as a guiding framework for negotiation progress and takes into account the current pricing agreed upon for all products and services
- Current Financial Health : this indicates how a given Business is performing on Financial metrics like Sales, Operating Expenses, EBITDA, Revenue etc at that given point in time compared to historical trends. These are a very strong indicator of how healthy a business is in given point in time
- Market and Investor Sentiments : what are the ongoing sentiments about a given Business in market. Are their customers happy? Are there any news articles indicating customer dissent and risk factors due to the Business's long term decisions? All these factors are also important to consider when assessing the WTP of a business at given point in time
- Multi-modal Margin : Margin would usually depend on factors outside numeric series alone. It depends on macro economic growth conditions, regional decisions, willingness to pay for your customers etc. We will use multi-modal forecasting where projected margin depends on multiple factors
- Deal Pricing Simulation : Given the factors that influence pricing, what will the metrics like margin, transaction volume, revenue etc. look like based on different scenarios generated around pricing and determines floor/ ceiling prices

All above are individual ML/AI projects and form a critical consolidation of factors that will guide the deal negotiation further

2.2. Approach –Fetching the Required Data

First step would be to gather the metrics and insights from all individual projects and then pass them through the MCP server which will consolidate the inputs using LLMs, and then pass the information to another LLM which will act as the negotiator agent. In subsequent sections we will discuss about each project in details. Below Figure 1 represents a high level overview of all the projects.

- Current Pricing Information : this project is specifically taking into account what price points have been agreed upon and the details or earlier rounds of negotiation. We are aiming to get the data around what prices were offered for each product (along with information like bundles or subscriptions offered) and what was accepted or rejected. This information will be used to generate the prompt for the agent that understands this given customer's price perception and strategies for negotiation that have been successful historically.

  The pricing information agreed upon and currently being charged can be called in using API connections to existing processes in the company. If not, then we will need to create interfaces or dashboards where current pricing information could be showcased for consumption

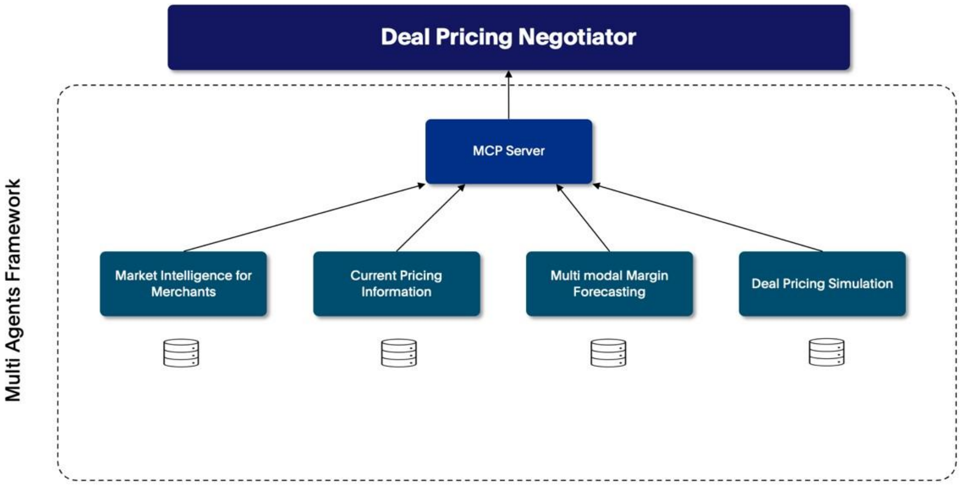

Figure 1

- Market Intelligence :Current Financial Health assessment leverages AI-driven analysis to evaluate a business's real-time financial performance across multiple dimensions. This component utilizes advanced financial statement analysis techniques powered by LLMs to process and interpret complex financial documents, extracting key performance indicators and identifying trends that may impact negotiation strategies.

  The system employs sophisticated prompt engineering techniques specifically designed for financial applications, enabling the LLM to perform comprehensive financial analysis including ratio analysis, cash flow assessment, and profitability evaluation. By integrating real-time financial data through secure API connections, the system can access current balance sheets, income statements, and cash flow statements to provide up-to-date financial health scores.

  Key metrics analyzed include debt-to-equity ratios, current liquidity positions, revenue growth trends, operating margin stability, and EBITDA performance compared to industry benchmarks. The AI system also incorporates external economic indicators and industry-specific factors to contextualize the financial health assessment within broader market conditions

This also includes Market and Investor Sentiments. Market sentiment analysis represents a critical component that leverages natural language processing and sentiment analysis algorithms to evaluate public perception and market mood surrounding target businesses. This multi-source sentiment analysis engine processes data from news articles, social media platforms, analyst reports, investor communications, and regulatory filings to provide comprehensive sentiment intelligence.

The sentiment analysis system employs advanced NLP techniques including transformer-based models specifically fine-tuned for financial text analysis. The system continuously monitors sentiment trends across multiple timeframes - real-time, short-term (weekly/monthly), and longterm (quarterly/yearly) - to identify sentiment shifts that could impact negotiation positioning.

Integration with financial news APIs, social media monitoring tools, and market data feeds enables real-time sentiment tracking. The system assigns sentiment scores on multiple dimensions including overall market perception, investor confidence, customer satisfaction, and regulatory sentiment.

All above data is consolidated to create an integrated score which represents customer performance on external sources of information and can be a proxy for any textual data coming out of financial and sentiment analysis

- Multi-modal Margin Forecasting : Multi-modal margin forecasting represents the core innovation of our approach, combining numerical time series data with textual information, market sentiment, and external economic indicators to predict future margin performance under various pricing scenarios. The forecasting system integrates multiple data modalities through sophisticated attention mechanisms that allow the model to focus on the most relevant features within each data type while fostering cross-modal understanding. The basic concept behind this methodology is that Margins do not depend on only historical trends but are influenced by multiple other factors. Our approach for multi modal margin forecast takes below trends into account:
- Historical profitability trends: margin, revenue and cost attributes taken together to create a time series data on overall profitability for the given customer
- Industry dependent seasonality : since the customers can come from various industries which themselves have a seasonality cycle, we need to account for industry specific trends
- Profitability across similar looking business customers : also taking into account how margin trends look like for similar customers, this gives an understanding of whether the given customer has similar behavior as others in their cohort and identify anomalies

Feature-level attention layers enable the model to identify the most relevant predictive features within each modality, while temporal attention mechanisms capture time-dependent relationships across different data sources.

The forecasting system generates margin predictions with confidence intervals, scenario-based projections, and sensitivity analysis to understand how different factors influence margin outcomes. This enables final negotiator agent to understand the range of possible margin scenarios and identify optimal pricing strategies

- Deal Pricing Simulation : utilizes AI-driven scenario modeling to test various pricing strategies and predict their impact on key business metrics including margin, transaction volume, revenue, and customer willingness to pay. This component employs LLM based prompting to generate comprehensive pricing scenario analyses. This also takes inputs from other projects described above to detect for changes in margins based on increase or decrease in current pricing offered which eventually depends on individual merchant score on external data sources.

These scenarios will provide a Floor and Ceiling pricing for the given negotiation, that will be used by the MCP server LLM to guide the user in negotiation process. It will be trained to be within 30 percentile to 70 percentile range for any discounts offered, crossing over 50 percentile only in hard negotiation cases.

3. Technical Architecture and Implementation

3.1. Model Context Protocol Implementation

The Model Context Protocol (MCP) serves as the backbone of our system architecture, providing a standardized interface for connecting our LLM-powered negotiation agent with multiple specialized data sources and AI models.

Our MCP implementation follows the established client-server architecture where the central negotiation agent acts as the MCP host, managing multiple MCP clients that connect to specialized MCP servers. Each MCP server is designed to handle specific data types and analysis functions, ensuring modularity and scalability. The architecture consists of four primary Agents:

i. Pricing Intelligence Agent: Manages historical pricing data and negotiation outcomes
ii. Financial Health and Sentiment Agent: Processes real-time financial statements and health metrics
iii. Margin Forecasting Agent: Executes multi-modal forecasting models
iv. Pricing Simulation Agent: Runs scenario simulations and optimization algorithms

Each MCP server exposes standardized tools, resources, and prompts that the central negotiation agent can access through the MCP protocol. This design enables seamless integration of new data sources and analytical capabilities without requiring changes to the core negotiation logic

3.2. Architecture Design

The technical architecture implements a sophisticated multi-agent LLM framework that leverages the Model Context Protocol (MCP) to orchestrate specialized AI agents for automated deal pricing negotiation. The system follows a hierarchical multi-agent architecture with clear separation of concerns, enabling modular development, specialization, and controlled communication between agents. Refer to below Figure 2 for detailed architecture design.

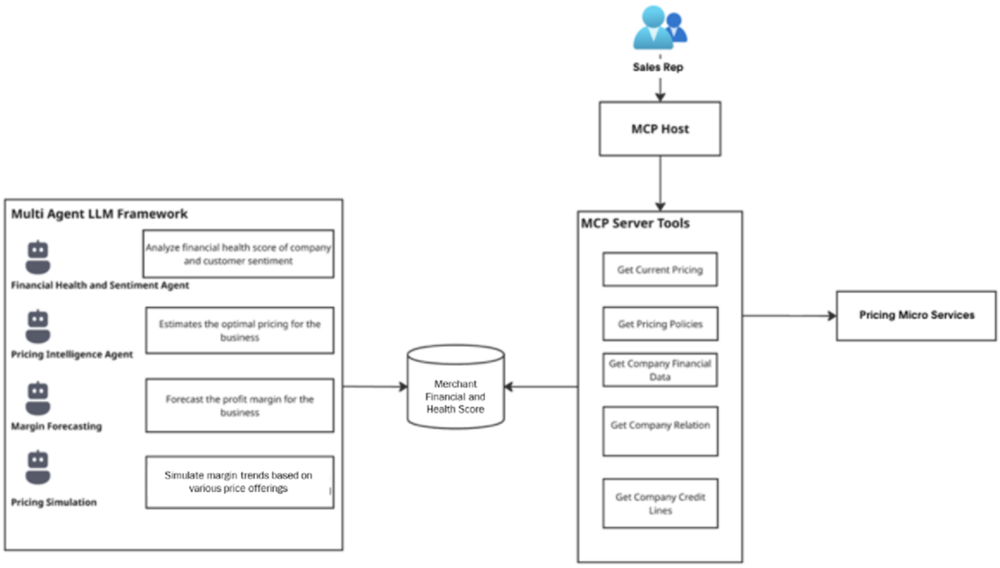

Figure 2

The MCP Host serves as the central orchestration layer implementing a client-host-server architecture where the host manages multiple client instances, each maintaining a 1:1 relationship with specialized MCP servers

Specialized Agents Architecture:

- Financial Health and Sentiment Agent

  Purpose: Analyzes comprehensive financial health metrics and market sentiment to determine customer negotiation capacity and willingness to pay.

  Tool Interfaces Exposed via MCP:

  `get_company_financial_data`: Retrieves real-time financial statements, cash flow, and key ratios
  `get_company_relation`: Historical engagement metrics and payment behavior analysis
  `get_company_credit_lines`: Credit capacity and utilization analysis for risk assessment
  `analyze_market_sentiment`: Natural language processing of news, social media, and analyst reports

- Pricing Intelligence Agent

  Purpose: Estimates optimal pricing strategies based on competitive analysis, market positioning, and historical negotiation patterns.

  The agent employs reinforcement learning algorithms trained on historical negotiation outcomes to predict likely customer responses to different pricing strategies. The system models

- *Price Elasticity: Customer sensitivity to price changes across different segments*
- *Competitve Dynamics: Real-time competitor pricing adjustments and market responses○Temporal Factors: Seasonal trends, market cycles, and economic indicators affecting pricing acceptance*

• Margin Forecasting Agent

Purpose: Provides multi-modal forecasting of profit margins under various pricing scenarios and market conditions.

The system implements joint trend-seasonal decomposition with feature-wise augmentation methods that have demonstrated improvement over single-modal approaches.

The architecture processes:

- *Structured Data: Historical margin data, transaction volumes, cost structures*
- *Unstructured Data: Market reports, regulatory announcements, economic indicators*
- *External Signals: Commodity prices, interest rates, currency fluctuations, industry benchmarks*

• Pricing Simulation Agent

Purpose: Simulates comprehensive business impact scenarios based on different pricing strategies and policy configurations **Simulation engine architecture**:

Using Monte Carlo framework for scenario simulations processing pricing tiers, discount policies, and market responses, using the Rule-based pricing policy engine system for modeling different pricing policies and approval workflows, and impact on Multi-dimensional analysis of pricing changes on revenue, volume, customer satisfaction, and competitive position

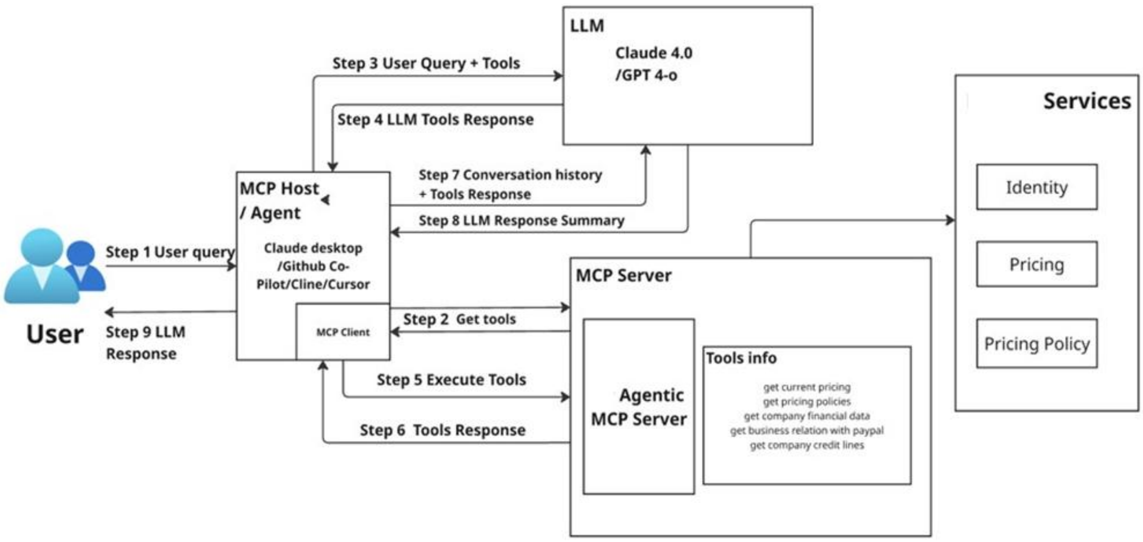

Figure 3

3.3. Detailed MCP Workflow Design

The diagram Figure 3 represents a comprehensive 9-step workflow architecture that demonstrates the complete lifecycle of an LLM request with tool execution through the Model Context Protocol (MCP). This technical implementation follows the client-server-host architecture pattern where each component has clearly defined responsibilities and communication protocols

The workflow begins when a user submits a natural language request through the client interface. This could be initiated through various LLM-enabled applications and a front end design over them.

The system accepts natural language requests such as:

- "What's the optimal pricing strategy for Enterprise Customer XYZ?"
- "Analyze customer financial health to come up with discount offerings"
- "Generate negotiation recommendations for this renewal opportunity"

The system automatically identifies which business intelligence tools and data sources are needed to fulfill the request. This eliminates the manual process of determining which systems to consult and ensures comprehensive analysis without human oversight gaps. It combines the sales representative's query with available business intelligence tools, creating a comprehensive analysis framework that considers all relevant factors simultaneously rather than requiring sequential manual research.

Below is the workflow design:

User Interaction Layer

- Step 1: Sales representatives submit pricing queries through integrated tools (Claude Desktop, GitHub Copilot, Cline/Cursor)Step 9: Comprehensive negotiation recommendations delivered back to user interface

AI Processing Core

- Step 3: MCP Host enhances user query with available tool capabilities, sends to LLM (Claude
4.0/GPT-4o)
- Step 4: LLM analyzes request and generates structured tool execution commands
- Step 8: LLM synthesizes all data into actionable negotiation strategies

Tool Discovery & Execution

- Step 2: MCP Server discovers available analysis tools from Agentic MCP Server
- Step 5: Parallel execution of five specialized tools
- Step 6: Tool responses consolidated into unified intelligence report

Context Integration

- Step 7: Historical conversation context combined with current analysis results for complete situational awareness

4. EXPERIMENTATION AND METRICS

While actual experimentation and empirical validation are planned for future work, the anticipated performance improvements are informed by extensive industry benchmarks from recent studies in automated negotiation, multi-agent pricing, and AI-driven margin forecasting.

We will implement end-to-end integration on enterprise datasets and other input parameters to this model to measure these metrics, especially adoption. This will include:

- Controlled A/B testing comparing negotiation outcomes with and without LLM+MCP guidance
- Real-world pilot deployments to observe margin, closure velocity, and win rate impacts in live negotiation settings
- Monitoring and feedback by Sales teams who would be using this system on the ground

5. IMPLEMENTATION CHALLENGES AND SOLUTIONS

MCP Server Configuration and Integration Complexity: The implementation of multiple specialized MCP servers presents significant configuration challenges, particularly in enterprise environments with existing legacy systems. Common issues include JSON configuration errors, dependency conflicts between different Python environments, and networking complications when integrating with existing firewall configurations.

Solution Framework: Implementation teams should establish robust configuration management protocols using infrastructure-as-code approaches. Utilizing containerized deployment with Docker and Kubernetes helps isolate dependencies and ensures consistent environments across development, staging, and production systems. Comprehensive logging mechanisms using Python’s built-in logging modules provide visibility into server operations and facilitate rapid debugging.

Data Quality and Integration Challenges: Enterprise deployments face significant data quality issues when integrating multiple heterogeneous data sources for financial analysis, sentiment monitoring, and pricing intelligence. Inconsistent data formats, API rate limiting from external services, and data synchronization delays can severely impact system performance and recommendation accuracy.

Solution Approach: Implementing robust data validation pipelines with automated quality checks ensures data integrity before processing. Circuit breaker patterns prevent cascading failures when external APIs experience issues, while intelligent caching strategies with Redis clusters reduce dependency on real-time external calls. Data standardization layers normalize information from different sources into consistent formats for processing.

Security and Compliance Complexities: MCP implementations introduce unique security challenges, including potential command injection vulnerabilities and the need for secure credential management across multiple service integrations. Enterprise environments require strict compliance with regulations like SOX, PCI DSS, and GDPR while maintaining system performance.

Solution Architecture: Implementation of comprehensive security frameworks including end-to-end encryption for all MCP communications, OAuth-based authentication with role-based access controls, and regular security audits of all MCP server components. Automated compliance monitoring ensures adherence to regulatory requirements without manual oversight overhead

6. PERFORMANCE EVALUATION METRICS

Quantitative Performance Indicators

Primary success indicators focus on measurable improvements in negotiation results and business outcomes:

- Margin Improvement Rate: Target 25-40% improvement in average deal margins compared to baseline manual negotiation processes
- Deal Closure Velocity: Reduction in negotiation cycle time from initial quote to signed agreement, with targets of 40-60% improvement in time-to-closure
- Win Rate Enhancement: Increase in successful deal closure rates, targeting 20-35% improvement in conversion from proposal to signed contract
- Revenue Per Negotiation: Measurement of total contract value secured per negotiation session, adjusted for deal complexity and customer segment

Business Impact Assessment Metrics

ROI-focused metrics demonstrate direct business value creation:

- Cost Reduction Metrics: Measurement of reduced manual analysis time, decreased requirement for multiple department consultations, and elimination of external consulting costs
- Revenue Enhancement: Tracking of incremental revenue generated through improved negotiation outcomes and faster deal closure cycles
- Customer Retention Impact: Analysis of relationship quality improvements and reduced customer churn resulting from more informed, fair pricing strategies

Long-term Strategic Impact Measurement

Enterprise-level metrics demonstrate sustainable competitive advantage:

- Market Share Growth: Correlation between improved negotiation capabilities and market position improvements
- Customer Lifetime Value Enhancement: Long-term customer value improvements resulting from more strategic, relationship-focused negotiations
- Organizational Learning Acceleration: Measurement of knowledge transfer and capability development across the sales organization

7. CONCLUSION

The implementation of LLM and MCP-based automated deal pricing negotiation systems represents a huge advancement in sales technology and business process optimization. This detailed analysis demonstrates how sophisticated AI orchestration through standardized protocols can deliver substantial improvements in negotiation outcomes while maintaining enterprise-grade security and compliance requirements.

The multi-modal approach to margin forecasting and pricing scenario simulation provides organizations with unprecedented analytical capabilities, enabling data-driven negotiation strategies that consistently outperform traditional manual approaches. The successful integration of specialized AI agents through MCP architecture proves that complex business intelligence can be democratized across sales organizations without requiring extensive technical expertise.

While implementation challenges around system integration, organizational adoption, and performance optimization require careful planning and execution, the documented solutions and best practices provide clear pathways to successful deployment. The future evolution toward fully autonomous negotiation capabilities and advanced predictive analytics promises even greater competitive advantages for early adopters.

The performance evaluation framework establishes measurable success criteria that demonstrate both immediate operational improvements and long-term strategic value creation. Organizations implementing these systems can expect significant returns on investment through improved margins, faster deal closure, and enhanced customer relationship management.

This research contributes to the broader understanding of enterprise AI implementation while providing practical guidance for organizations seeking to leverage advanced AI capabilities for competitive advantage in an increasingly data-driven business environment.

REFERENCES

[1] Anthropic. (2024). "Introducing the Model Context Protocol." Technical Report. https://www.anthropic.com/news/model-context-protocol

[2] Model Context Protocol Specification. (2025). "Architecture Overview." https://modelcontextprotocol.io/docs/concepts/architecture

[3] Microsoft. (2024). "LLM evaluation metrics and best practices." https://learn.microsoft.com/enus/ai/playbook/technology-guidance/generative-ai/working-with-llms/evaluation/list-of-eval-metrics

[4] OPTIMIZING PRICING FRAMEWORK USING REASONING LLMS AND AGENTS FOR ENHANCED DEAL NEGOTIATIONS, Chirag Soni and Swati Shah, IJAIA, Vol 16, No 2 March 2025 https://aircconline.com/ijaia/V16N2/16225ijaia02.pdf

[5] OpenAI Platform. (2025). "Building MCP servers for ChatGPT and API integrations." OpenAI API Documentation. https://platform.openai.com/docs/mcp

[6] Renting, B. M., et al. (2024). "Towards General Negotiation Strategies with End-to-End Reinforcement Learning." arXiv preprint arXiv:2406.15096. https://arxiv.org/pdf/2406.15096.pdf

[7] Springer. (2025). "A systematic mapping study on automated negotiation for software engineering." Automated Software Engineering, 32(1). https://link.springer.com/article/10.1007/s10515-025-00515-7

[8] SSRN. (2025). "Machine Learning Algorithms for Dynamic Pricing Optimization in Retail." SSRN Electronic Journal. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5228515