JNGR 5.0

Journal of Next-Generation Research 5.0

Website: www.jngr5.com Email: editor@jngr5.com

Generic Multi-Agent AI Framework for Weighted Dynamic Corridor Price Optimisation

Walter Kurz[1]

[1]Business Artificial Intelligence, Signum Magnum College, Malta, https://orcid.org/0009-0006-8045-4775.
Contributing authors: walter.kurz@smc.college;

Abstract

The objective of this analysis is to address the challenges encountered by pricing systems in managing real-time market dynamics. This study presents a fundamental theoretical framework with a focus on taxonomy and ontology for a domain-specific multi-agentic artificial intelligence (AI) serving as an internal price advisor to optimise pricing strategies for products and services. The system is designed to function in conjunction with other corporate AI systems and an Enterprise Resource Planning System (ERP). The ERP serves as a high-quality data foundation, and several other internal and external sources can provide essential data with varying quality. Methods: The proposed AI model builds upon the Weighted Dynamic Corridor Price Optimisation framework, which integrates cost-plus and value-based pricing methodologies within a non-linear price corridor bounded by lower and upper thresholds. In the context of supply chain integration, fully-cooperative pricing models can apply Nash equilibrium to enhance supply chain profitability, whilst semi-cooperative models mitigate information asymmetry through the principal-agent theory. The findings from the theoretical analysis of the generic industry- and product-agnostic multi-agentic AI system suggest the system's potential capacity for dynamically computing optimal prices. A generative AI module could facilitate real-time decision-making, enabling sales teams and similar stakeholders to simulate scenarios and refine pricing strategies. In conclusion, the proposed AI system should be capable of delivering adaptive, context-aware, and data-driven recommendations. Depending on its application, the AI system could become very complex, susceptible to errors, and require significant maintenance. Future research should focus on customising the proposed AI system for specific industries and product categories and validating its applicability through empirical research.

Keywords: Dynamic Pricing Optimization, Multi-Agent Systems, AI, Artificial Intelligence, Pricing Strategy Analytics

1 Introduction

Determining the optimal price is vital for maximising profit and maintaining competitiveness. Dynamic pricing based on empirical data involves complex calculations due to variables such as production costs, competitor prices, market demand, and elasticity. The absence of artificial intelligence (AI) tools capable of real-time dynamic pricing presents an opportunity to develop systems that function as negotiation advisors. These systems integrate live data, analyse price impacts interactively, and provide tailored strategies that outperform static algorithms.

1.1 Context Framework

We propose a high-level, generic AI framework that is currently industry- and product-agnostic, aiming to establish a foundation for further empirical research and detailed testing, given the limited research at the intersection of business administration, price calculation, and AI.

The proposed AI model for price optimisation is not an isolated system but a core component of a central corporate AI infrastructure. This infrastructure, embedded within a company's broader Enterprise Information System (EIS), operates as an integral part of the supply chain management process. The AI model interacts with other specialised AI agents (e.g., marketing, compliance, or strategy agents) to enable seamless communication, task delegation, and decision-making.

The central corporate AI system ensures that pricing decisions are aligned with real-time market conditions, corporate objectives, and supply chain dynamics. Figure 1 illustrates the structural positioning and integration of the AI pricing model within the enterprise and supply chain context, indicating its interactions with other AI agents and business entities.

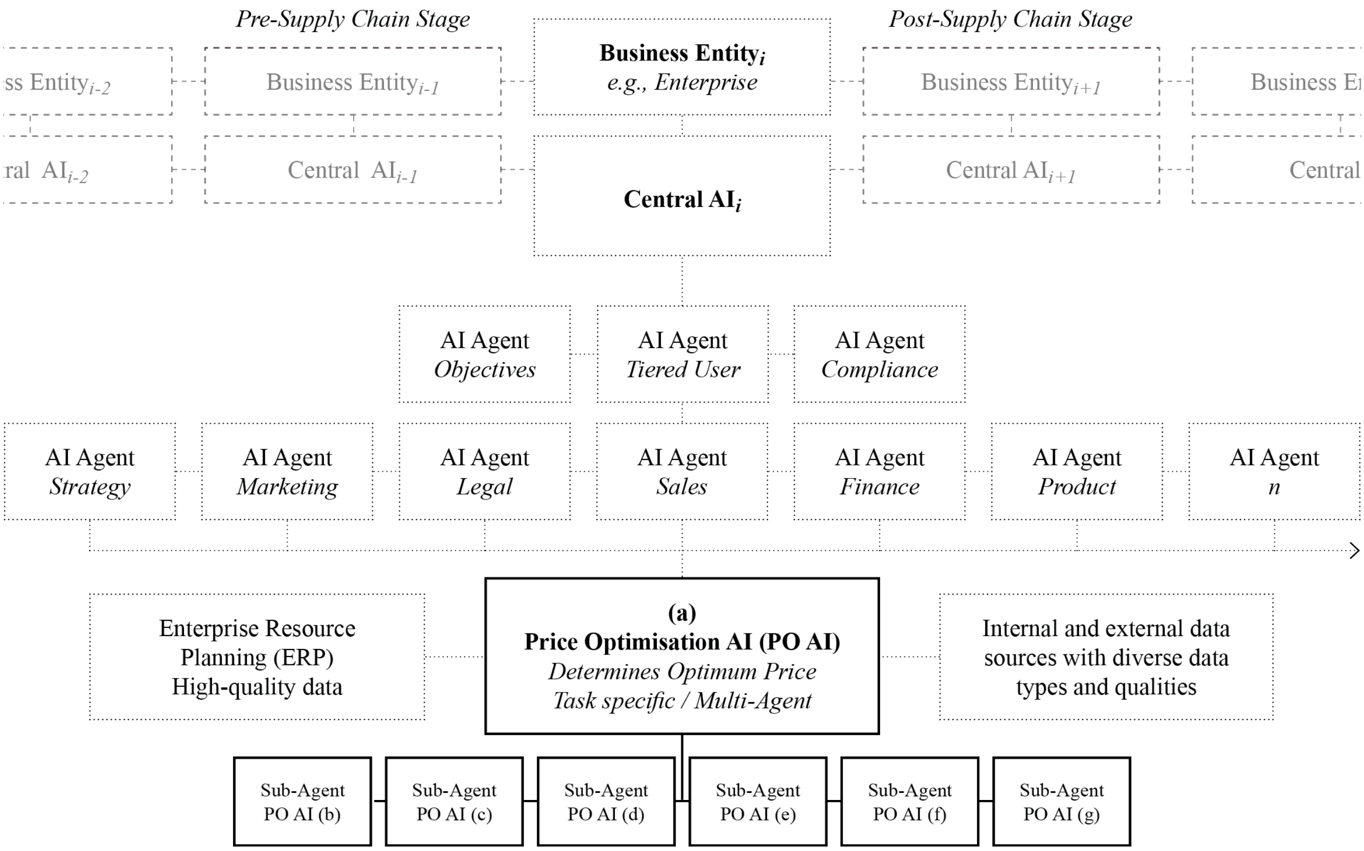

Fig. 1 Positioning of the Price Optimisation AI Agent within the context of the central enterprise AI and a supply chain.

The generic multi-agent AI framework builds on the Weighted Dynamic Corridor Price Optimisation (WDCP) model, proposed by Stromeyer and Kurz in their study [50], which combines cost-plus and value-based approaches within a non-linear pricing corridor.

The lower boundary is determined by production costs, including fixed cost degression and economies of scale, while the upper boundary incorporates market demand elasticity, competitor prices, customer willingness-to-pay, and psychological pricing thresholds. The WDCP framework models two interdependent price-demand functions: internal orientated functions reflecting, for example, cost behaviour based on production levels and market orientated functions analysing, for example, the relationship between price points and customer demand. The proposed AI calculates the optimum price floor based on this framework and processes empirical data from production systems, market analytics, and customer platforms to dynamically determine optimal prices. It includes a generative AI module that enables interactive engagement, allowing, for example, sales teams to simulate pricing scenarios and develop negotiation strategies based on live data and calculated impacts. While not designed to autonomously set or implement prices, the AI provides detailed recommendations, leaving final pricing decisions to human users. For this paper, we assume the deploying company operates a functional Enterprise Resource Planning (ERP) system, ensuring precise processes and high-quality data for price-related values. Conversely, data quality for other price-influencing factors, such as market demand or customer behaviour, may be lower due to variability and less structured collection methods. The WDCP model ultimately aims to maximise profitability, expressed as a function of price [50]:

$$\Pi(P) = Q(P) \times (P - C) + R(Q) \quad (1)$$

This equation defines profit (Π) as a function of price (P), where $Q(P)$ represents the quantity sold as a function of price, and $P - C$ captures the contribution margin, calculated as the difference between the selling price (P) and the unit cost (C). The term $R(Q)$ incorporates volume-based rebates or discounts, allowing strategic adjustments to align with specific business objectives.

1.2 Contribution to the Field

This research provides a contribution to the intersection of artificial intelligence and dynamic pricing by introducing the Weighted Dynamic Corridor Price Optimisation (WDCPO) framework and its integration into a modular multi-agent AI system. In contrast to existing methodologies that often lack adaptability to real-time market dynamics or fail to incorporate both cost- and value-based pricing strategies, the proposed framework addresses these limitations through a structured, dual-boundary approach. The primary contribution lies in the mathematical formalisation of dynamic pricing corridors that account for production cost degression, economies of scale, customer willingness-to-pay, and market elasticity. Through the integration of these variables into a unified model, this framework ensures that pricing decisions are aligned with both profitability and market competitiveness. The study pioneers the application of advanced cooperative and semi-cooperative pricing frameworks within supply chains, applying Nash equilibrium principles and principal-agent theory to address inter-organisational dynamics. The development of a multi-agent AI system contributes to the field by enabling domain-specific agents to process price related complex datasets in real time, providing actionable, context-sensitive pricing recommendations. This work also introduces a generative AI module for interactive negotiation and scenario analysis, offering a significant advancement in the usability of AI systems for pricing optimisation. The modularity and adaptability of the proposed generic framework contribute to the foundational taxonomy and ontology of artificial intelligence in multi-agent systems and its practical applicability across diverse industries, providing a robust foundation for empirical testing and future customisation.

1.3 Related Work

This research builds upon the Weighted Dynamic Corridor Price Optimisation (WDCPO) framework by Stromeyer and Kurz [50], which integrates cost-plus and value-based pricing within a dynamic, non-linear corridor. Previous studies [48, 41] established foundational pricing methods; these lacked adaptability to real-time dynamics and advanced cooperative frameworks. The WDCPO framework addresses these limitations by incorporating cost degression, economies of scale, market elasticity, psychological factors, and advanced game-theoretic models. It applies Nash equilibrium to align profitability across supply chains and principal-agent theory to address information asymmetry. Recent analyses [18, 3] explored artificial intelligence applications in pricing but did not fully integrate these theoretical approaches or utilise real-time multi-agent systems.

2 Methodology

This research employs a formal discussion approach, grounded in a comprehensive review of existing literature and the integration of theoretical frameworks. Contributions from peer-reviewed publications, textbooks, and seminal works are synthesised to construct a conceptual framework for dynamic pricing architecture. Emphasis is placed on formal approaches to ensure mathematical rigour and alignment with contemporary pricing challenges. The literature review serves as the primary method for identifying, analysing, and integrating state-of-the-art research findings. Formal studies on cost-based pricing, value-based pricing, Nash equilibrium, principal-agent theory, non-linear optimisation, and multi-agent systems are examined to combine established methodologies with innovative AI-driven solutions. Conceptual models are adapted and extended to reflect real-world applications and current technological capabilities. The research is inherently theoretical, presenting a structured, generic framework that is relevant to academic discourse and practical business applications.

3 Research

Prior to designing the proposed AI framework, it is important to comprehend the mathematical foundation of the pricing model, as this complex theoretical construct necessitates the collection of extensive data from diverse sources and the application of a comprehensive set of mathematical approaches. The Weighted Dynamic Corridor Price Optimisation model establishes the foundation for the AI's architecture, determines its functionality, and ensures alignment with the model's objectives. While the pricing model relies heavily on mathematical equations, it also incorporates qualitative aspects that require advanced reasoning capabilities. Large Language Models (LLMs) must apply objective-under-constraints frameworks and internal reasoning to interpret and integrate non-quantitative inputs. This ensures that the AI can account for qualitative factors, such as customer perceptions, market dynamics, and strategic considerations, which are vital for generating effective pricing recommendations.

3.1 Specifying the Lower Boundary of the Pricing Model

We focus first on the lower boundary of the Weighted Dynamic Corridor Price Optimisation model, which represents the minimum price a company can charge without jeopardising its financial stability. Key components to consider include fixed costs, which remain constant regardless of production levels, variable costs that change with output, and cost degression due to economies of scale [48].

Additional considerations include contribution margins, regulatory constraints such as price floors, and strategic adjustments like selling below cost for market penetration or key account relationships [38, 14].

These elements collectively define the foundational inputs necessary for constructing a robust lower boundary. Our assumptions regarding the lower boundary are predicated on the requirement to cover both fixed and variable costs while accounting for economies of scale, cost degression, and other strategic and external considerations. Fixed costs, such as rent and salaries, are distributed across the production quantity, reducing the fixed cost per unit as production scales up. Variable costs, including raw materials and energy, decrease with economies of scale through bulk purchasing and operational efficiencies. Additional factors include contribution margins to ensure profitability for each unit sold and external constraints such as legal price floors or minimum wage regulations.

Strategic adjustments, including selling at a loss for key accounts or market penetration, and quality considerations, such as maintaining perceived value thresholds, are essential. Brand reputation plays a role in ensuring the minimum price aligns with customers' expectations of quality and trust in the brand. This encompasses the influence of OEM suppliers, whose reliability and integration into the value chain impact brand image and pricing strategies. Market-specific factors, including competitor pricing, customer price sensitivity, inventory management, and regulatory compliance, are also integrated into the framework. Collectively, these components form a dynamic foundation for determining the lower boundary in the pricing corridor, expressed as:

$$P_{\min} = \frac{FC}{Q} + \sum_{j=1}^{m} (VC_j \cdot u_j) + S + R_{\min} + C_{\text{neg}} + M + B \tag{2}$$

where $P_{\min}$ is the minimum price per unit, FC represents the total fixed costs, Q is the quantity produced, VC_j denotes the variable cost of the j-th factor (e.g., raw materials, energy), and u_j is the usage of the j-th factor per unit. S accounts for strategic slack or adjustments (e.g., for market penetration or key accounts), $R_{\min}$ represents regulatory or legal constraints affecting minimum pricing, and C_{neg} refers to negative contribution margin allowances for strategic purposes. M encompasses market considerations (e.g., perceived value thresholds, competitor benchmarks), and B captures the brand value component, reflecting the minimum price necessary to maintain perceived quality and avoid brand dilution.

To further refine the lower boundary, we incorporate time-dependent and non-linear aspects that reflect the dynamic nature of costs and market conditions. Fixed costs may vary over time due to lease renewals, maintenance expenses, or depreciation schedules, necessitating their modelling as time-dependent components [48].

Variable costs, such as raw materials and energy, fluctuate with supply chain conditions, market trends, and currency exchange rates, which can also be represented dynamically [38].

Over time, production processes often benefit from efficiency gains, following learning curve effects that reduce variable costs through technological advancements and improved resource utilisation [14].

Historical and forecasted data can enhance accuracy by identifying seasonality and long-term trends, while customer demand elasticity, influenced by time-sensitive factors such as brand perception or competing offers, necessitates the integration of dynamic price-quantity relationships [27]. External economic indicators such as inflation and interest rates directly affect both fixed and variable costs, reinforcing the importance of real-time adjustments [16]. These enhancements provide a foundation for a dynamic model that extends the static lower boundary into a time-sensitive and adaptable framework. To capture the dynamic and time-sensitive aspects of the lower boundary, we refine its mathematical representation. Fixed costs, which may vary over time due to periodic changes such as lease renewals or depreciation schedules, are expressed as a time-dependent function:

$$FC(t) = \text{Fixed Costs Over Time} \tag{3}$$

Variable costs are similarly modelled as a time-dependent function to account for fluctuations resulting from supply chain dynamics, market trends, or currency exchange rates:

$$VC_j(t) = \text{Variable Costs Over Time for the } j\text{-th Factor} \tag{4}$$

Incorporating economies of scale and learning curve effects, we express the variable cost per unit as:

$$VC(Q,t) = VC_{\text{base}} \cdot \left(1 - \frac{\beta}{Q(t)^n}\right) \tag{5}$$

Here, $Q(t)$ represents the quantity produced at time t, β is the bulk discount factor, and n captures non-linear cost degression effects. This formulation dynamically adjusts variable costs to reflect improved efficiencies over time. To incorporate historical trends and forecasted data, we introduce a time-series function that refines costs based on seasonal or long-term patterns:

$$P_{\min}(t) = \frac{FC(t)}{Q(t)} + \sum_{j=1}^{m} (VC_j(t) \cdot u_j) + S(t) + R_{\min} + C_{\text{neg}} + M(t) + B \tag{6}$$

This equation integrates time-dependent fixed and variable costs, strategic adjustments over time ($S(t)$), and external economic factors such as inflation, represented within $M(t)$. The inclusion of B ensures the minimum price maintains brand value, aligning with perceived quality and trust. To understand how the lower boundary evolves dynamically over time, we can express it as a differential equation. This approach shows the time sensitivity of cost components and strategic adjustments, which are important for an AI framework to adapt in real time.

$$\begin{aligned} \frac{dP_{\min}}{dt} = {} & \frac{dFC(t)}{dt} \cdot \frac{1}{Q(t)} - \frac{FC(t)}{Q(t)^2} \cdot \frac{dQ(t)}{dt} \\ & + \sum_{j=1}^{m} \left(\frac{dVC_j(t)}{dt} \cdot u_j\right) + \frac{dS(t)}{dt} + \frac{dM(t)}{dt} + \frac{dB(t)}{dt} \end{aligned} \tag{7}$$

Here, $\frac{dP_{\min}}{dt}$ represents the rate of change of the minimum price over time, $\frac{dFC(t)}{dt}$ captures changes in fixed costs over time (e.g., lease or depreciation changes), and $\frac{dVC_j(t)}{dt}$ reflects changes in variable costs (e.g., raw material fluctuations). $\frac{dQ(t)}{dt}$ denotes changes in production volume over time, while $\frac{dS(t)}{dt}$ represents strategic adjustments, such as responding to market shifts. $\frac{dM(t)}{dt}$ accounts for market-related considerations, including inflation or competitive pricing trends, and $\frac{dB(t)}{dt}$ represents brand adjustments, such as shifts in brand perception or quality considerations. The term u_j serves as the proportional weighting factor for the j-th variable cost component, and m is the number of variable cost components. The differential approach provides a dynamic framework to model the evolution of the lower boundary over time by capturing the temporal changes in fixed costs, variable costs, production volume, market adjustments, and brand considerations.

3.2 Specifying the Upper Boundary of the Pricing Model

The Weighted Dynamic Corridor Price Optimisation framework proposes an approach for constructing the upper boundary, taking into account market demand elasticity, customer willingness-to-pay, and psychological pricing thresholds. These factors encompass both quantitative components, such as competitive pricing and product differentiation, and qualitative factors, such as brand perception and consumer psychology.

In contrast to static algorithms, these qualitative aspects necessitate advanced AI systems, particularly large language models (LLMs), to process and integrate non-linear, subjective inputs by interpreting qualitative data alongside quantitative variables. The upper boundary of the pricing corridor is influenced by several key components. Market demand elasticity measures how price changes affect demand, providing an upper constraint to prevent significant decreases in sales volume [48]. Customer willingness-to-pay (WTP) represents the maximum price perceived as fair or acceptable, based on the product's value and market alternatives [41]. Chen et al [10] found that AI-initiated pricing increases consumer repurchase and recommendation behaviours, reduces complaints and switching, is mediated by ethical perceptions, and is negatively influenced by perceived enterprise control. Psychological pricing thresholds, such as round numbers or prestige pricing, influence buyer decisions by applying cognitive biases [14].

Brand perception further impacts the upper boundary, as strong brands associated with trust and quality can justify higher prices [27]. Competitive pricing defines the range within which a product must compete, especially in markets where substitutes are readily available. Product differentiation, through unique features or innovations, allows firms to extend the upper boundary beyond the competition [16]. The position of a product in its market lifecycle also matters, with pricing flexibility expanding for innovative products and narrowing for mature or saturated ones. Dynamic and qualitative market factors, such as customer sentiment or behavioural trends sourced from social media or reviews, require interpretation through advanced AI systems, making these qualitative aspects particularly suitable for LLMs that can synthesise diverse inputs. Emerging studies discuss the importance of real-time adaptability in pricing. Sanchez-Cartas and Katsamakas [47] demonstrate how AI-driven systems can account for platform competition and network effects to optimise pricing strategies. Neubert [43] highlight the role of dynamic pricing systems in identifying trends and patterns across global markets. Gerpott and Berends [18] examine the interdisciplinary nature of pricing decisions, integrating perspectives from economics, marketing, and behavioural science. Building on these components, the upper boundary of the pricing corridor is mathematically expressed to capture market-driven, psychological, and competitive factors that influence the maximum price a company can charge:

$$P_{\max} = \left(a - \alpha \cdot Q^{1/n}\right) \cdot \left(1 + T \cdot V\right) + WTP \tag{8}$$

Here, $P_{\max}$ represents the maximum price per unit, a denotes the maximum potential demand or price ceiling, and α is the demand sensitivity factor, indicating how demand changes with price. The term $Q^{1/n}$ captures non-linear elasticity, where n is the elasticity exponent. T accounts for the technological adjustment factor, reflecting innovation or obsolescence, while V represents the market differentiation factor, which includes brand strength or unique value. WTP stands for willingness-to-pay, which is the maximum price customers perceive as fair or acceptable based on the perceived value of the product.

This equation expresses the maximum price ($P_{\max}$) a company could charge, determined by a combination of market-driven factors such as demand elasticity and technological adjustments, along with qualitative inputs like market differentiation and customer willingness-to-pay (WTP). In practical application, it can be used to dynamically assess pricing strategies, enabling businesses to optimise their price ceilings based on real-time data, customer insights, and competitive conditions.

The term $Q^{1/n}$, referred to as the non-linear elasticity term, models how demand (Q) responds to price changes in a non-linear manner. The parameter n, known as the elasticity exponent, determines the degree of sensitivity of demand to price fluctuations, allowing for greater flexibility in representing market behaviours compared to linear elasticity models. When $n > 1$, the relationship between price and demand is less sensitive (inelastic), indicating that changes in price lead to smaller proportional changes in demand. Conversely, when $n < 1$, demand becomes highly sensitive (elastic), meaning small price changes result in significant shifts in demand (Simon and Fassnacht, 2016). Identifying the elasticity exponent (n) in a market requires analysing historical pricing and sales data.

Advanced statistical methods, such as regression analysis, can estimate how quantity demanded responds to variations in price over time. Experimental pricing strategies, such as A/B testing with different price points, can provide insights into elasticity by observing customer behaviour in controlled environments. For markets with complex demand patterns, machine learning models may be employed to elucidate non-linear relationships between price and demand. This flexibility in capturing non-linear dynamics renders the elasticity term crucial for accurately defining the upper boundary in dynamic pricing models. Willingness-to-pay (WTP) refers to the maximum price a customer perceives as fair or acceptable for a product or service, reflecting its perceived value. This concept is influenced by various interrelated factors. The utility and functionality of a product play a critical role, as consumers demonstrate a higher propensity to pay increased prices for products that meet their needs effectively or offer superior features. Products with unique attributes or advanced customisation options tend to command higher WTP [49]. Brand perception and trust significantly influence WTP. Customers associate strong, reputable brands with quality and reliability, enabling companies with high brand equity to justify premium pricing [30]. Social influences also contribute to WTP, particularly for products that serve as status symbols or align with social acceptance. Luxury goods and exclusive brands, for example, are often purchased not only for their functional value but also for their ability to convey social recognition and prestige [23]. Emotional connections further elevate WTP by creating perceived value beyond a product's functional benefits. Brands that resonate emotionally with consumers, either through compelling narratives or alignment with personal values, can increase perceived worth and therefore WTP [51].

The association between price and quality is another important factor. Consumers often infer higher quality from higher prices, especially in cases where quality is difficult to evaluate before purchase. This association can lead to a higher WTP for products perceived as premium or high-quality offerings [39].

Cultural and psychological factors also shape WTP. For example, in collectivist cultures, social acceptance and group preferences often play a significant role in pricing perceptions, whereas individualistic cultures might emphasise personal utility and unique value [25]. The economic context, including disposable income levels and economic stability, heavily influences WTP. During periods of economic growth, consumers may attribute higher value to aspirational or luxury products, whereas economic downturns tend to lower perceived value due to reduced purchasing power [2].

To account for the dynamic nature of the upper boundary, we introduce time-dependence into the model. This allows for adjustments based on real-time market changes, customer behaviour, and evolving economic and technological factors. The time-dependent upper boundary is expressed as:

$$P_{\max}(t) = \big(a(t) - \alpha(t) \cdot Q(t)^{1/n}\big) \cdot \big(1 + T(t) \cdot V(t)\big) + WTP(t) \tag{9}$$

Here, $P_{\max}(t)$ represents the maximum price per unit at time t, $a(t)$ denotes the maximum potential demand or price ceiling over time, and $\alpha(t)$ is the demand sensitivity factor varying with market dynamics. The term $Q(t)^{1/n}$ reflects the non-linear elasticity term, capturing demand-price sensitivity over time. $T(t)$ accounts for the technological adjustment factor, representing innovation or obsolescence trends, while $V(t)$ represents the market differentiation factor, which evolves with brand strength or unique features over time. $WTP(t)$ stands for willingness-to-pay, dynamically adjusting with perceived value, economic factors, and consumer sentiment.

The dynamic changes in the upper boundary can be captured with a differential equation:

$$\begin{aligned} \frac{dP_{\max}}{dt} = \frac{da(t)}{dt} - \frac{d\alpha(t)}{dt} \cdot Q(t)^{1/n} - \frac{\alpha(t)}{n \cdot Q(t)^{1-1/n}} \cdot \frac{dQ(t)}{dt} \\ + \frac{dT(t)}{dt} \cdot V(t) + T(t) \cdot \frac{dV(t)}{dt} + \frac{dWTP(t)}{dt} \end{aligned} \tag{10}$$

Here, $\frac{dP_{\max}}{dt}$ represents the change in the maximum price over time, $\frac{da(t)}{dt}$ captures the change in potential demand or price ceiling over time, and $\frac{d\alpha(t)}{dt}$ reflects the change in demand sensitivity due to market conditions. The term $\frac{dQ(t)}{dt}$ denotes the change in quantity demanded over time, while $\frac{dT(t)}{dt}$ accounts for technological advancements or obsolescence effects. Additionally, $\frac{dV(t)}{dt}$ represents changes in market differentiation, such as brand value or competitive positioning, and $\frac{dWTP(t)}{dt}$ describes adjustments in willingness-to-pay based on perceived value or economic trends. These equations for the upper boundary are essential for dynamically accounting for both quantitative factors, such as market elasticity and demand, and qualitative factors, such as psychological pricing thresholds and brand perception, which evolve over time. Traditional fixed algorithms lack the capacity to interpret qualitative and non-linear data, rendering advanced AI systems, particularly LLMs, relevant for integrating subjective inputs and real-time market changes into precise pricing strategies.

3.3 Specifying the Optimal Price Function

To determine the optimum price within the defined corridor of upper and lower boundaries, we incorporate key variables including profit margin, boundary weighting, demand elasticity, negotiation flexibility, real-time market data, and psychological factors. Consequently, we balance profitability, competitiveness, and market dynamics by dynamically integrating these variables. We commence with the market and competitive view, which considers the influence of market leaders, low-cost competitors, and other market-specific factors, and is expressed as:

$$P_{\text{competitive}} = W_s \cdot P_{\text{leaders}} + (1 - W_s) \cdot P_{\text{low-end}} + w_t \cdot T_m + w_p \cdot P_{\text{customer}} + w_e \cdot E_m \tag{11}$$

Here, $P_{\text{competitive}}$ represents the competitive price dynamically positioned between market leaders and low-cost competitors. W_s is the strategy weight ($0 \leq W_s \leq 1$), determined by the company: $W_s = 1$ means fully aligned with market leader pricing, while $W_s = 0$ means fully aligned with low-cost competitors. P_{leaders} denotes the average price of the top two market leaders, and $P_{\text{low-end}}$ represents the average price of low-cost competitors (value players). The term w_t is the weight assigned to market trends, and T_m is the market trend index capturing seasonal demand or macroeconomic shifts. w_p represents the weight assigned to customer preferences, while P_{customer} reflects customer-driven price insights derived from historical data or segmentation. Lastly, w_e is the weight assigned to economic and purchasing power adjustments, and E_m is the market-specific economic index reflecting factors such as purchasing power parity (PPP), local currency value, and cost of living.

This equation calculates $P_{\text{competitive}}$, a dynamic competitive price positioned between market leaders and low-cost competitors, while incorporating market trends, customer preferences, and economic adjustments to reflect both local and global pricing dynamics. The inclusion of E_m enables companies to adapt their pricing strategies for international markets by accounting for factors such as purchasing power parity, local currency value, and cost of living, ensuring alignment with regional market conditions. Real-time market data captures dynamic external factors that influence pricing decisions, such as exchange-traded prices, supply chain conditions, demand fluctuations, and currency exchange rates. These components are important to ensure that pricing strategies remain responsive to market shifts, align with customer purchasing power, and reflect competitive and economic realities. This dynamic adjustment factor can be expressed in the following equation:

$$\begin{aligned} D_m = P_{\text{competitive}} + w_e \cdot P_{\text{exchange}} + w_s \cdot S_c \quad &+ w_d \cdot D_f + w_{fx} \cdot Fx + w_i \cdot E_i \\ + w_c \cdot C_p + w_{sm} \cdot S_m + w_r \cdot R_p \end{aligned} \tag{12}$$

Here, D_m represents the real-time market data adjustment factor, while $P_{\text{competitive}}$ denotes the dynamic price positioned between market leaders and low-cost competitors. P_{exchange} reflects live prices from commodity exchanges or other regulated pricing sources, and S_c captures supply chain data, including costs or disruptions. The term D_f represents demand fluctuations derived from sales systems or forecasts, and F_x accounts for currency exchange rates impacting international pricing. E_i includes economic indicators such as inflation or GDP trends, whereas C_p highlights competitor promotions or tactical discounts. S_m captures social media sentiment or consumer feedback, and R_p reflects regulatory or political changes affecting pricing. The equation uses weighted components to combine multiple diverse real-time data sources into a single adjustment factor, D_m. Each data source, such as competitor prices, exchange rates, or demand fluctuations, contributes differently to pricing decisions based on the product, market context, or business goals. The weights, represented by w_x, determine the relative influence of each component on D_m, enabling flexibility and adaptability. For instance, in the case of commodities with fixed exchange prices, the weight assigned to P_{exchange} (e.g., w_e) may dominate, while the weight for competitor promotions (w_c) might be negligible. Conversely, in consumer goods markets where competition is more intense, $P_{\text{competitive}}$ and C_p may carry greater weights, reflecting their relevance to tactical pricing. The output of D_m is not a specific price but an adjustment factor that modifies the final price calculation to account for real-time external market conditions. This ensures the pricing strategy dynamically aligns with competitive, economic, and customer-related factors. For example, if competitor promotions are prominent, D_m may result in a downward adjustment to maintain competitiveness. Alternatively, a rise in exchange rates may cause D_m to increase the price for international markets, ensuring profitability is preserved despite cost fluctuations.

This dynamic adjustment ensures the pricing remains contextually relevant and strategically aligned. While it is straightforward to define these variables theoretically, their real-world application is highly complex, requiring data from diverse and often unstructured sources. This is where AI, particularly advanced systems like large language models (LLMs), excels by synthesising information from disparate APIs, databases, and real-time inputs. Unlike static algorithms, which are limited to predefined rules, AI dynamically processes and integrates qualitative and quantitative data, enabling context-aware adjustments that traditional methods cannot achieve.

To construct a comprehensive pricing equation that incorporates all six variables, we combine real-time market data (D_m), profit margin (M_p), boundary weighting ($W_{\min}, W_{\max}$), demand elasticity adjustment (E_d), negotiation flexibility (K_n), and psychological factors (P_s). We also integrate a strategic slack variable (S_k) to account for key account adjustments, allowing flexibility in pricing for strategic clients where a long-term partnership is prioritised over immediate profit margins. This ensures the equation is adaptable to different scenarios, optimisable for objectives like profit maximisation, and flexible under constraints such as upper ($P_{\max}$) and lower ($P_{\min}$) boundaries.

The optimal price can be expressed as a dynamic function that integrates all relevant variables while remaining constrained by the upper ($P_{\max}$) and lower ($P_{\min}$) boundaries:

$$P_{\text{optimal}} = W_{\min} \cdot P_{\min} + W_{\max} \cdot P_{\max} + M_p + D_m + E_d + P_s + K_n + S_k \tag{13}$$

Here, P_{optimal} represents the calculated optimal price, while $P_{\min}$ and $P_{\max}$ denote the lower and upper price boundaries, respectively. $W_{\min}$ and $W_{\max}$ are the boundary weighting factors, balancing cost recovery and value capture. M_p reflects the minimum profit margin or desired markup, and D_m is the real-time market data adjustment factor, incorporating competitive prices, market trends, and economic conditions. The term E_d accounts for demand elasticity adjustment, optimising volume and revenue based on price sensitivity, while P_s captures psychological factors, such as pricing thresholds or perceived value effects. K_n represents negotiation flexibility, allowing for discounts or strategic adjustments, and S_k provides strategic slack for key accounts, enabling price adjustments for long-term partnerships or high-priority clients.
This equation is designed to be used as an objective function for profit maximisation while ensuring constraints from $P_{\min}$ and $P_{\max}$ are respected. The flexibility introduced by S_k, K_n, and D_m makes it adaptable

to real-time conditions and strategic priorities, constituting a feasible framework for AI-driven pricing optimisation. To introduce time sensitivity, we can model the optimal price (P_{optimal}) as a time-dependent function:

$$\begin{aligned} P_{\text{optimal}}(t) &= W_{\min}(t) \cdot P_{\min}(t) + W_{\max}(t) \cdot P_{\max}(t) \\ &\quad + M_p(t) + D_m(t) + E_d(t) + P_s(t) + K_n(t) + S_k(t) \end{aligned} \tag{14}$$

To capture how $P_{\text{optimal}}(t)$ evolves over time, we can calculate its time derivative:

$$\begin{aligned} \frac{dP_{\text{optimal}}(t)}{dt} &= \frac{d}{dt}\Big(W_{\min}(t) \cdot P_{\min}(t) + W_{\max}(t) \cdot P_{\max}(t) \\ &\quad + M_p(t) + D_m(t) + E_d(t) + P_s(t) + K_n(t) + S_k(t)\Big) \end{aligned} \tag{15}$$

Here, $\frac{dP_{\text{optimal}}(t)}{dt}$ represents the changes in the optimal price over time, while $\frac{dW_{\min}(t)}{dt}$ and $\frac{dW_{\max}(t)}{dt}$ denote the changes in weighting factors for the lower and upper boundaries, respectively. $\frac{dM_p(t)}{dt}$ captures changes in cost structures, such as raw material costs or supply chain delays, and $\frac{dD_m(t)}{dt}$ reflects real-time market fluctuations, including competitor promotions and exchange rates. The term $\frac{dE_d(t)}{dt}$ accounts for demand elasticity shifts due to seasonal or economic factors, while $\frac{dP_s(t)}{dt}$ represents changes in psychological factors influencing perceived value. $\frac{dK_n(t)}{dt}$ highlights strategic decisions evolving over time, such as renegotiation of key account terms, and $\frac{dS_k(t)}{dt}$ addresses adjustments for strategic slack variables over time. This time-dependent formulation of $P_{\text{optimal}}(t)$ is important for capturing dynamic changes in pricing influenced by evolving cost structures ($P_{\min}(t)$), market conditions ($D_m(t)$), and strategic adjustments ($S_k(t)$). Integrating these variables, the model ensures that pricing decisions remain adaptive to real-time fluctuations, providing businesses with a framework to optimise profitability under shifting economic and market constraints.

3.4 Specifying the Profit Maximisation Objective

Having established the lower boundary, upper boundary, and the methodology for determining the optimal price (P_{optimal}) within the pricing corridor, the next step is to formalise the profit maximisation objective. This is important for ensuring that the calculated price aligns with profitability goals while adhering to the constraints defined by the pricing corridor. The model adapts to real-time conditions and strategic business priorities by integrating dynamic market and cost factors.
To maximise total profit (Π), the objective function is expressed as:

$$\begin{aligned} \max_{P_{\text{optimal}}} \quad & \Pi(P_{\text{optimal}}) = Q(P_{\text{optimal}}) \cdot (P_{\text{optimal}} - C(Q)) - R(Q) \\ \text{subject to:} \quad & P_{\min} \leq P_{\text{optimal}} \leq P_{\max} \end{aligned} \tag{16}$$

Here, $\Pi(P_{\text{optimal}})$ represents the total profit as a function of the optimal price, while $Q(P_{\text{optimal}})$ denotes the sales quantity as a function of the price. P_{optimal} is the calculated optimal price within the pricing corridor. The term $C(Q)$ represents the cost function, capturing total costs as a function of quantity produced, and $R(Q)$ accounts for rebates or discounts applied based on volume.

This formulation provides the mathematical foundation for deriving the optimal price that maximises profit, ensuring that all constraints and real-world complexities, such as demand elasticity and cost dynamics, are effectively integrated into the decision-making process. In practice, the artificial intelligence system would iteratively solve this optimisation problem by extracting real-time data, recalculating the price, and providing actionable recommendations to decision-makers. This approach would yield a dynamically updated "optimal price" based on current conditions and strategic objectives. When implemented effectively, this framework renders the pricing process both data-driven and strategic, facilitating informed decision-making while maximising profitability.

3.5 Specifying the AI-Driven Fully Cooperative Price Equilibrium Framework

In a fully cooperative setting where participants in a supply chain or business collaboration are willing to freely exchange relevant information, we propose a novel application of the Nash equilibrium—an aspect of game theory—to optimise pricing strategies within cooperative supply chains. Traditionally, the Nash equilibrium has been utilised to model strategic interactions where no participant can improve their outcome unilaterally [42, 19, 40]. Applying this concept to real-time supply chain management has been limited due to the complexity and volume of data involved. For human decision-makers, processing such data and simultaneously optimising outcomes across multiple stakeholders is not feasible. With advancements in AI, particularly in large language models (LLMs) combined with mathematical optimisation techniques, it becomes possible to dynamically compute Nash equilibrium solutions in real-world supply chain scenarios. This framework utilises the AI's capacity to handle vast and diverse datasets, integrating real-time market data, production costs, customer preferences, and competitive factors across multiple parties in the supply chain. The AI can identify equilibrium prices that maximise collective profitability while balancing individual constraints, such as minimum profit margins and strategic priorities. Incorporating fairness constraints into dynamic pricing strategies, Bérczi et al [6] explore how dynamic pricing strategies in unit-demand markets can maintain global envy-freeness across various time perspectives, developing algorithms that achieve envy-free optimal dynamic prices for social welfare maximisation, though highlighting the complexity of such tasks for revenue maximisation. These approaches align with the framework's objective of maximising collective profitability while respecting individual constraints. The approach redefines supply chain optimisation by allowing entities to cooperate more effectively, even in highly complex environments. Unlike traditional static models, which focus solely on localised decision-making, AI-driven frameworks enable real-time adjustments, ensuring that all parties benefit from shared optimisation goals. A generic form of a Nash equilibrium can be expressed as:

$$U_i(P_i^*, P_{-i}^*) \geq U_i(P_i, P_{-i}^*), \quad \forall P_i, \forall i \tag{17}$$

Here, U_i represents the utility or payoff function of participant i, while P_i^* denotes the strategy (e.g., price or decision) of participant i at equilibrium. P_{-i}^* captures the strategies of all other participants except i at equilibrium, and P_i represents any alternative strategy for participant i. The inequality ensures that no participant can improve their utility by unilaterally changing their strategy at equilibrium. We propose the AI-Driven Fully Cooperative Price Equilibrium Framework (AI-FCPEF), wherein two or more AI systems (or a single integrated AI) negotiate and optimise prices within a supply chain to identify a Nash equilibrium that maximises collective profitability. In contrast to traditional scenarios, where companies independently optimise prices using adversarial strategies, this framework facilitates cooperative negotiation to achieve individual price optima while maintaining a supply-chain-wide strategic optimum. The integration of real-time data, strategic priorities, and profitability constraints enables the AI systems to establish pricing that balances mutual benefits and supports long-term collaboration. This approach redefines Nash equilibrium applications, shifting from individual profit maximisation to cooperative optimisation, leveraging AI's capacity to process vast datasets and adapt to dynamic conditions in real time.

The Generic AI-Driven Fully Cooperative Price Equilibrium Framework (AI-FCPEF) can be expressed as:

$$\max_{P_1^*, P_2^*} \Big[U_1(P_1^*, P_2^*) + U_2(P_1^*, P_2^*)\Big] \quad \text{subject to: } U_i(P_i^*, P_{-i}^*) \geq U_i(P_i, P_{-i}^*) \quad \forall P_i, \forall i \tag{18}$$

Here, P_1^* and P_2^* represent the optimal prices for company 1 and company 2, respectively. $U_i(P_i^*, P_{-i}^*)$ denotes the utility or profit function of company i when using the optimal price P_i^* and considering the pricing strategy of the other company P_{-i}^*. The term P_{-i}^* captures the pricing strategy of the other company in relation to company i at equilibrium. Finally, $U_i(P_i, P_{-i}^*)$ represents the utility or profit function of company i when using an alternative price P_i, while the other company maintains its equilibrium strategy P_{-i}^*.

The inequality ensures that no company can unilaterally improve its profit by altering its price while the other company maintains its optimal pricing strategy. To explicitly incorporate the profit functions of both companies (16) into the cooperative Nash equilibrium, the equation becomes:

$$\begin{aligned}
\max_{P_1^*, P_2^*} \Big[& Q_1(P_1^*) \cdot (P_1^* - C_1(Q_1)) - R_1(Q_1) + W_{\min,1} \cdot P_{\min,1} + W_{\max,1} \cdot P_{\max,1} \\
& + D_{m,1} + S_{k,1} \\
& + Q_2(P_2^*) \cdot (P_2^* - C_2(Q_2)) - R_2(Q_2) + W_{\min,2} \cdot P_{\min,2} + W_{\max,2} \cdot P_{\max,2} \\
& + D_{m,2} + S_{k,2} \Big] \\
\text{subject to:} \quad & P_{\min,1} \leq P_1^* \leq P_{\max,1}, \quad P_{\min,2} \leq P_2^* \leq P_{\max,2}
\end{aligned} \tag{19}$$

Here, P_1^* and P_2^* are the optimised prices for company 1 and company 2, respectively. $Q_1(P_1^*)$ and $Q_2(P_2^*)$ represent the sales quantities as functions of their respective prices, while $C_1(Q_1)$ and $C_2(Q_2)$ are the cost functions for each company based on their sales quantities. $R_1(Q_1)$ and $R_2(Q_2)$ denote rebates or discounts applied based on volume for each company. The terms $W_{\min,1}$, $W_{\max,1}$, $W_{\min,2}$, and $W_{\max,2}$ are the weighting factors for upper and lower boundaries, and $P_{\min,1}$, $P_{\max,1}$, $P_{\min,2}$, $P_{\max,2}$ denote the lower and upper price boundaries for each company. $D_{m,1}$ and $D_{m,2}$ are the real-time market data adjustments, and $S_{k,1}$ and $S_{k,2}$ represent strategic slack variables for key accounts or priority clients.

This equation combines the previously developed individual profit functions (16) into a single cooperative framework. The AI system can optimise both (or more) functions simultaneously while ensuring that individual constraints are maintained. This approach can be extended across the entire supply chain, enabling all participating companies to collaboratively optimise their pricing strategies through the AI-Driven Fully Cooperative Price Equilibrium Framework (AI-FCPEF).Treating the supply chain as a unified entity, the framework dynamically aligns individual prices with strategic objectives, such as increasing market share or achieving market leadership. Companies may also adjust margins at specific points in the supply chain to support collective advantages, such as entering new markets or strengthening competitive positioning.

3.6 Specifying the AI-Driven Semi-Cooperative Price Equilibrium Framework

In a semi-cooperative setting, participants engage in business transactions while withholding critical proprietary information. This necessitates an alternative approach where optimisation occurs under partial information constraints. Principal-agent theory offers a suitable foundation for addressing such scenarios. It models interactions where one party (the principal) delegates tasks to another (the agent) while managing information asymmetry and accounting for risk aversion [15, 32].

In the context of semi-cooperative supply chains, this theory can guide the development of pricing mechanisms by aligning incentives between parties through structured contracts. These contracts adaptively balance individual utilities and risks, allowing for effective optimisation despite limited information sharing [26]. An AI system specialised in semi-cooperative frameworks can estimate missing information and reason under uncertainty to address information gaps. It applies principal-agent theory models to compute pricing strategies that align incentives and approach the optimum price in this constrained setting. The principal-agent problem is typically modelled as [32]:

$$\max_{a} \mathbb{E}[U_p(a,s)] \quad \text{subject to:} \quad \mathbb{E}[U_a(a,s)] \geq U_a^{\min} \tag{20}$$

Here, a represents the agent's action or decision variable, and s denotes the state of nature or external factors influencing outcomes. U_p is the utility function of the principal, which depends on the agent's action and the state of nature, while U_a represents the utility function of the agent, dependent on their action and the state of nature. The term $U_a^{\min}$ reflects the agent's minimum acceptable utility, encapsulating their participation constraint. Finally, $\mathbb{E}$ is the expectation operator, accounting for uncertainty in the state of nature.
To transition this framework into the *AI-Driven Semi-Cooperative Price Equilibrium Framework (AI-SCPEF)*, the optimisation problem incorporates specific variables and constraints relevant to semi-cooperative supply chains. The principal (buyer) aims to maximise utility, defined by the value obtained from purchasing goods or services from the agent (seller) at an optimised price, $P_{\text{optimal, principal}}$. This price must remain within a dynamic corridor defined by a lower boundary, $P_{\text{min, agent}}$, and an upper boundary, $P_{\text{max, agent}}$, reflecting the agent's cost structures, competitive dynamics, and market conditions.

$$U_p = V(Q) - Q \cdot P_{\text{optimal, principal}} \tag{21}$$

where $V(Q)$ represents the value derived by the principal from the quantity purchased, and $Q \cdot P_{\text{optimal, principal}}$ reflects the expenditure on the transaction. The agent's utility function captures their profit and is defined as:

$$U_a = Q(P_{\text{optimal, agent}}) \cdot \big(P_{\text{optimal, agent}} - C(Q)\big) - R(Q) + S_k \tag{22}$$

where $P_{\text{optimal, agent}}$ is the price the seller receives, $C(Q)$ represents production costs, $R(Q)$ accounts for volume-based rebates or discounts, and S_k introduces strategic slack for priority clients. To ensure alignment and participation, the agent's price, $P_{\text{optimal, agent}}$, must also satisfy the buyer's constraints. The optimisation problem for the entire system, incorporating principal-agent interactions, is formalised as:

$$\begin{gathered} \max_{P_{\text{optimal}}} \ \mathbb{E}\Big[V(Q) - Q \cdot P_{\text{optimal, principal}}\Big] \\ \text{subject to: } \mathbb{E}\Big[Q(P_{\text{optimal, agent}}) \cdot \big(P_{\text{optimal, agent}} - C(Q)\big) - R(Q) + S_k\Big] \geq U_a^{\min} \end{gathered} \tag{23}$$

Equation (23) integrates real-time market adjustments (D_m), demand elasticity (E_d), and strategic factors (S_k) to calculate an optimised price across the supply chain. The AI system inherently processes this data, reducing the need for assumptions and enabling precise optimisation. Within a fully integrated multi-agent system, this framework creates a chain of optimised prices across all levels of the supply chain. Each transaction, from one company to another, aligns individual objectives with collective goals, ensuring that the principal-agent relationships are both efficient and strategically aligned.

3.7 Specifying the AI Architecture for Dynamic Pricing Using Multi-Agent Systems

The Weighted Dynamic Corridor Price Optimisation (WDCPO) model we propose would benefit from implementation through a dedicated AI system that is owned, trained, and fully accessible by the company. This recommendation does not pertain to hosting decisions, such as whether the system resides on in-house servers or in the cloud, but rather to ensuring complete ownership and control of the AI.

Such ownership addresses the complexity of the model, the sensitive nature of the data involved, and the need for compliance with data protection regulations and legal frameworks. It also ensures that intellectual property and trade secrets remain secure, which would otherwise be challenging to guarantee with external AI solutions. While this is one possible approach, decision-makers should carefully evaluate the trade-offs between external and internal implementations.

The WDCPO model requires extensive, continuous real-time API calls to process dynamic data efficiently. Centralised AI systems, particularly single-instance models, face significant performance limitations due to synchronous API execution, where each call blocks further processing until a response is received. Sequential execution introduces delays and restricts scalability, rendering such systems unsuitable for handling the volume and frequency of real-time data required for dynamic pricing. While asynchronous methods, such as those presented in AsyncLM [20], reduce latency by overlapping API interactions, their implementation within monolithic architectures remains limited. These systems require retraining the entire model for updates, resulting in high time and operational costs.

Distributed frameworks such as Pathways [4] demonstrate the effectiveness of asynchronous distributed dataflow and task scheduling across thousands of accelerators, utilising techniques such as sharded dataflow graphs and asynchronous gang-scheduling to substantially enhance the management and execution of complex machine learning operations. Asseman et al [3] apply game theory to optimise dynamic pricing in a multi-agent blockchain environment, addressing a game with imperfect information and harmonising diverse incentives to enhance economic outcomes. Their model facilitates revenue optimisation through real-time identification of consumer budgets, underpinning adaptive, data-driven pricing strategies within a blockchain protocol. When applied to complex, multi-functional pricing models, single-instance AI structures lack the modularity needed to integrate such methods effectively.

We propose a multi-agent AI system that incorporates an objective-under-constraints framework for each agent. This configuration allows specialised agents to operate independently with clear and focused objectives and constraints. A centralised AI attempting to handle diverse objective-under-constraints frameworks simultaneously would likely encounter inefficiencies and conflicts due to the complexity of managing such varied tasks within a single model. A multi-agent system resolves these issues by distributing tasks among agents, enabling concurrent execution and improving scalability and efficiency. The modular design further simplifies updates, as only individual agents need retraining rather than the entire system. This approach is a key recommendation for meeting the demands of our dynamic pricing model.

We identify two approaches for designing the multi-agent AI system for dynamic pricing: *domain-specific agents* and *function-specific agents*. Domain-specific agents focus on particular areas such as monitoring customer behaviour, tracking market developments, or analysing competitor activities. These agents align with human intuition and workflows, facilitating comprehension, maintenance, and fine-tuning. For instance, a domain-specific agent monitoring competitors would oversee tasks such as tracking product launches, pricing strategies, and branding activities, enabling human resources to concentrate on clear, domain-specific outcomes.

Conversely, function-specific agents are designed for individual computational tasks, such as regression analysis, time-series forecasting, and optimisation routines. While this approach supports highly specialised operations, it disperses functions across workflows, increasing complexity, coordination challenges, and the risk of errors. Fine-tuning and updating such systems become challenging as fragmented tasks lack a unified structure. A domain-specific agent system offers several advantages. It is inherently scalable, as new domains can be added with minimal impact on existing agents. It simplifies error isolation, as agents operate within well-defined boundaries, facilitating issue identification and resolution. It also integrates domain knowledge, allowing industry-specific insights and human expertise to guide agent development.

Domain-specific agents enable parallel training and fine-tuning, avoiding dependency conflicts. The equations underlying the Weighted Dynamic Corridor Price Optimisation model, particularly its dynamic upper and lower boundaries and interdependent price-demand functions, are highly complex and designed generically. In real-world applications, they require extensive modifications to align with an organisation's

strategy, product types, and contextual components. A domain-specific agent system is particularly suited for this purpose, as different industries and product categories demand unique analytical priorities.
For example, monitoring competitors in consumer goods differs significantly from evaluating pricing for capital goods. Over time, organisations can further optimise results by incorporating or refining agents to address evolving business requirements. Distributing tasks such as cost analysis, demand forecasting, and competitor tracking to specialised agents mitigates the risk of errors while enhancing efficiency. These agents independently process their results, which are subsequently aggregated by a central AI to calculate the optimal price range.

We also propose that each agent operates within a domain-specific objective-under-constraints framework. This approach would not be feasible in a function-specific system, where fragmented tasks lack alignment with domain-level goals, further complicating management and optimisation. Recent research supports our rationale for implementing domain-specific agents in a multi-agent AI system. Calvaresi et al [7] elucidate the challenges associated with achieving real-time performance in multi-agent systems, attributing these difficulties to the limitations of conventional agent schedulers, communication middleware, and negotiation protocols.

These components are not inherently designed to meet the stringent timing constraints necessitated by safety-critical applications in sectors such as healthcare and automotive industries. Similarly, Condurache et al [13] propose a framework for dynamic multi-agent systems, emphasising the importance of modular design for scalability, error isolation, and flexibility, all of which align closely with the benefits of domain-specific agents. Korbel and Tichý [31] explore dynamic pricing models using multi-agent reinforcement learning and demonstrate that agents specialising in specific domains improve the flexibility and efficiency of pricing strategies.

Lu et al [35] further illustrate how multi-agent reinforcement learning can address dynamic pricing challenges, such as traffic congestion, by assigning clear objectives to agents that adapt to evolving conditions. Their research supports the notion that domain-specific agents simplify task optimisation and improve long-term adaptability. Research by Asseman et al [3] explores distributed agent systems, concluding that fragmenting highly specialised tasks across workflows increases the risk of errors and coordination challenges, further strengthening the case for domain-specific agents.

3.8 Constructing the Generic Domain-Specific Multi-Agent AI System

We propose the following domain-specific agents for the multi-agent AI system: the *Price Optimisation Agent* (a) as the central coordinating AI, the *Market Analysis Agent* (b), the *Cost Tracking Agent* (c), the *Time-Series Price Analysis Agent* (d), the *Corporate Pricing Strategy Agent* (e), and the *Price Policy Effect Agent* (f). The *Corporate Strategy Agent* (g) ensures alignment between the *Price Optimisation Agent* and broader strategic priorities, such as market expansion, innovation strategies, and long-term profitability goals.

The *Market Analysis Agent* (b) consists of sub-agents, including the *Competitor Price/Product Observer Agent* (b1), the *Customer and Client Behaviour Analyst Agent* (b2), and the *Replacement Product Analyst Agent* (b3). The *Cost Tracking Agent* (c) has sub-agents, including the *Internal Costs Analyst Agent* (c1), the *Contract Analyst Agent* (c2), and the *Supply Chain Analyst Agent* (c3).

In practical applications, the specific configuration of the system will be contingent upon the organisation's requirements and the nature of the products or services being priced. The internal architecture of each agent is dependent on the specific task it is designed to perform. It is reasonable to posit that every agent is, or at minimum should be, equipped with a comprehensive set of state-of-the-art components, including artificial intelligence algorithms, machine learning models, analytical methods, and generative AI capabilities to fulfil its assigned role effectively. It is proposed that these agents are not merely background systems delivering isolated functions. When organisations invest substantial resources into establishing such a system, the integration of generative AI agents can provide additional utility by enabling human-in-the-loop interaction. Users can interact with individual agents separately to gain specific insights into the domain knowledge accumulated by that agent.

3.9 Specifying the Agent-Specific Objective-Under-Constraints Framework

The Price Optimisation Agent (a) applies equations (14), (15), (16) and others to calculate the optimum price or price corridor based on the Weighted Dynamic Corridor Price Optimisation model. While the equations provide the mathematical foundation for determining the optimum price, the agent itself operates within a specific objective-under-constraints framework that governs how these equations are applied. The overarching objective of the Price Optimisation Agent (a) is to calculate a price that aligns with the company's strategic goals as overseen by the Corporate Strategy Agent (g). This ensures that the output remains consistent with priorities such as profitability targets, market positioning, and innovation strategies. A second critical

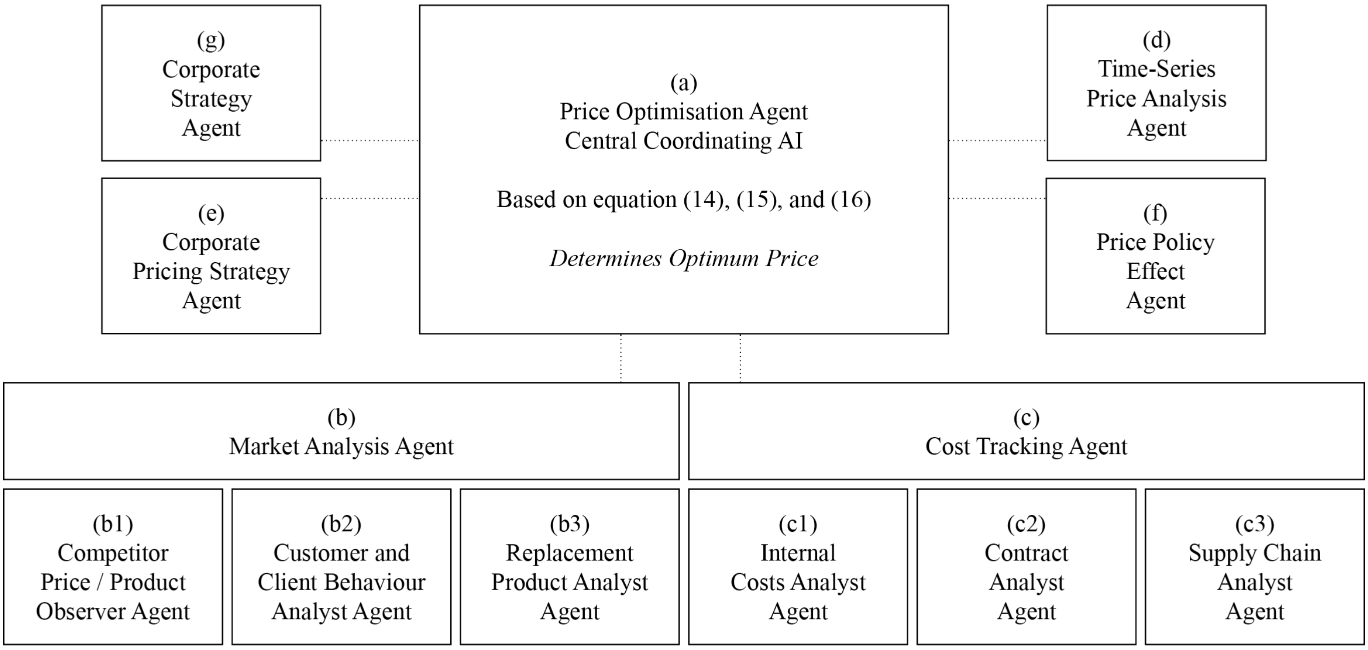

Fig. 2 Generic Domain-Specific Multi-Agent AI System for Price Optimisation

constraint arises from the applied and predefined pricing strategies set by the Corporate Pricing Strategy Agent (e). The price must comply with the chosen pricing strategy, which reflects the company's overall approach to market positioning. If the company applies an international pricing model diversified across different markets, such as the European, Asian, or other geographically distinct regions, the Corporate Pricing Strategy Agent (e) ensures that the Price Optimisation Agent (a) adheres to the appropriate strategy for each market, accounting for regional economic, competitive, and customer-specific conditions. The Price Optimisation Agent (a) must also operate within legal and regulatory frameworks. The calculated price must respect external constraints, such as industry regulations and competitive rules, ensuring compliance and avoiding illegal anti-competitive behaviours, such as predatory pricing. Furthermore, the optimum price or price corridor must remain bounded within the dynamic upper and lower boundaries implied by the equations, ensuring both cost coverage and market feasibility. The generic objective-under-constraints framework for the Price Optimisation Agent (a), including the three generally applicable constraints, can be expressed as follows:

$$O_{(a)}: \quad \underset{P_{\text{optimal}}}{\arg\max} \; \{P_{\text{optimal}}\} \tag{24}$$

$$\begin{aligned} \text{subject to:} \quad & C_1 : P_{\text{optimal}} \text{ aligns with strategic objectives set by agent (h),} \\ & C_2 : P_{\text{optimal}} \text{ complies with pricing strategies defined by agent (e),} \\ & C_3 : P_{\text{optimal}} \text{ adheres to legal and competitive constraints,} \\ & C_4 : P_{\min} \leq P_{\text{optimal}} \leq P_{\max}, \\ & C_n : \text{Other constraints specific to real-world applications.} \end{aligned}$$

The objective defines the core task of the agent, while the constraints ensure that the price calculation aligns with strategic, applied, and regulatory requirements.

The group of *market-oriented agents* under the *Market Analysis Agent* (b): The system is tasked with continuously monitoring and analysing market developments across various domains. These components are designed to anticipate the requirements of the Price Optimisation Agent (a) and ensure that the data and insights they provide are structured, relevant, and readily accessible for seamless integration into the price calculation process, for example, equation (14). All components operate independently and proactively, tracking relevant data streams and aggregating information even without direct requests from the Price Optimisation Agent (a). Their primary function is to extract and process data from a variety of sources to capture changes in competitor pricing, customer behaviour, and product availability. To ensure the captured data reflects developments over time, it is proposed that all gathered information is stored in a vector database with an associated timestamp. The timestamp is essential to enable the components to identify trends, analyse temporal sequences of events, and understand the evolution of market dynamics. Without this functionality, the data would become a static archive rather than a structured, development-oriented resource. While the individual sub-components (b1), (b2), and (b3) focus on their specific domains, the overarching Market Analysis Agent (b) coordinates their outputs and maintains the vector database, integrating the domain-specific data into a coherent and accessible structure for further analysis.

The Competitor Price/Product Observer Agent (b1) tracks and maintains accurate, time-stamped data on competitor prices and product availability to ensure real-time visibility of external market conditions. The framework can be expressed as follows: The Competitor Price/Product Observer Agent ($O_{(b1)}$) is tasked with the continuous tracking and aggregation of time-stamped data on competitor pricing and product availability. This component ensures real-time visibility into external market conditions, supporting data-driven optimisation by the Price Optimisation Agent ($O_{(a)}$). The agent utilises advanced data pipelines to retrieve and validate information, maintaining compliance with legal and ethical guidelines while preserving data accuracy and temporal coherence. The optimisation framework for $O_{(b1)}$ can be expressed as:

$$O_{(b1)}: \quad \underset{\mathcal{D}_{\text{comp}},\mathcal{T}}{\text{Track and optimise}} \; \{\mathcal{D}_{\text{comp}}(t) \,|\, t \in [t_0, T]\} \tag{25}$$

$$\begin{aligned} \text{subject to:} \quad & C_1 : \mathcal{D}_{\text{comp}} \text{ must originate from verified and approved sources,} \\ & C_2 : \mathcal{T}_{\text{obs}}(t) \in \mathbb{R}^+ \text{ to ensure all data is time-stamped for trend analysis,} \\ & C_3 : \mathbb{E}[\text{Acc}(\mathcal{D}_{\text{comp}})] \geq \epsilon \text{ to guarantee data accuracy and completeness,} \\ & C_4 : \mathcal{D}_{\text{comp}} \text{ complies with legal standards for competitive monitoring,} \\ & C_n : \text{Additional constraints tailored to specific business applications.} \end{aligned}$$

Here, $\mathcal{D}_{\text{comp}}(t)$ represents the dataset of competitor prices and product availability as a function of time t, while $\mathcal{T}_{\text{obs}}(t)$ denotes the observation timestamp associated with each data point, ensuring temporal consistency. The term $\text{Acc}(\mathcal{D}_{\text{comp}})$ refers to the accuracy measure of the collected data, defined as the proportion of error-free entries, with ϵ representing the minimum acceptable threshold for data accuracy. Finally, $[t_0, T]$ defines the time window over which competitor data is tracked and analysed.

The Competitor Price/Product Observer Agent (b1) has the additional objective of identifying the most relevant competitors and their competitive products or services. We propose the agent generates a reasoning-based ranking of competitors, assessing their relevance and competitive impact using a "Competitive Relevance Factor". This factor reflects the degree of competitive pressure posed by each competitor.

The agent tracks competitor positioning, pricing strategies, price developments, discounts, and, where possible, derives assumptions about client bases, customer segments, and other relevant indicators. In an international context, the agent monitors the markets in which competitors operate, ensuring that market-specific competitors are identified. This allows the Price Optimisation Agent (a) to determine the relevant competitive landscape for any specific country or market where a price needs to be calculated. All data collection adheres to legal boundaries and ethical standards, creating comprehensive competitor profiles and ensuring compliance.

Since we have already proposed that every agent integrates a generative AI aspect, the Competitor Price/Product Observer Agent (b1) allows humans to interact directly and retrieve specific insights, also outside of active price calculations. This capability extends the agent's objective to include focused tasks, where humans in the loop can instruct the agent to monitor specific competitors, markets, or products and services. These tasks can be executed as immediate requests or as semi-permanent assignments, enforcing the agent to maintain ongoing observation and analysis over specified focus areas while ensuring alignment with its core function.

The Customer and Client Behaviour Analyst Agent (b2) operates also within an objective-under-constraints framework, focusing on monitoring and analysing customer and client behaviour in alignment with the company's overarching strategy. This agent receives specific instructions and constraints from the Corporate Strategy Agent (g) and the Corporate Pricing Strategy Agent (e), which define the targeted customer and client segments based on sociodemographic, behavioural, and strategic considerations.

The agent's core task mirrors the general functionality of the Competitor Price/Product Observer Agent (b1), such as continuous tracking, building a time-stamped vector database, and providing relevant insights. The focus is to analysing purchasing patterns, demand behaviours, client preferences, and segment-specific price sensitivities. The agent identifies trends, evaluates behavioural shifts, and builds comprehensive profiles for defined customer and client segments. These profiles include assumptions about purchasing power, demand elasticity, and response to pricing strategies, all within the defined constraints. The framework can be expressed as follows:

$$O_{(b2)}: \quad \underset{\mathcal{WTP},\mathcal{B}}{\text{Determine and analyse}} \left\{\mathcal{WTP}(t), \mathcal{B}(t) \,\middle|\, t \in [t_0, T]\right\} \tag{26}$$

$$\begin{aligned}\text{subject to:}\quad & C_1 : \text{Target segments and focus areas are defined by agent (h)},\\ & C_2 : \text{Pricing strategy alignment is ensured as defined by agent (e)},\\ & C_3 : \text{Collected behavioural data must include timestamps},\\ & C_4 : \text{Data collection adheres to legal, ethical, and privacy regulations},\\ & C_n : \text{Other constraints specific to real-world applications}.\end{aligned}$$

Here, $\mathcal{WTP}(t)$ represents the willingness to pay across target segments as a function of time t, while $\mathcal{B}(t)$ denotes customer and client behaviour data, including purchasing patterns, demand elasticity, and response to pricing strategies. The variable $t \in [t_0, T]$ defines the time window for data collection and analysis. C_1 ensures that the target segments and focus areas are defined by the Corporate Strategy Agent (h), while C_2 guarantees alignment with pricing strategies as specified by the Corporate Pricing Strategy Agent (e). C_3 enforces the inclusion of timestamps to track behavioural trends, and C_4 mandates compliance with legal, ethical, and privacy regulations. Finally, C_n accounts for additional constraints relevant to specific real-world scenarios.

The Customer and Client Behaviour Analyst Agent (b2), like agent (b1), operates proactively and independently, gathering relevant insights even in the absence of direct requests. This agent is also responsible for determining the upper limit of the price corridor, which is heavily influenced by clients' willingness to pay. It specialises in analysing targeted customer segments and processes data through surveys, quick-response studies, and similar quantitative methods. The agent must also process qualitative inputs, such as expert assessments or structured methods like Delphi rounds, to refine insights on willingness to pay. These analyses are conducted across a matrix of customer segments, ensuring the willingness-to-pay data is accurately aligned with the defined sociodemographic and behavioural profiles. This agent is also tasked with integrating value-based pricing research by identifying price-influencing criteria for each customer segment. When multiple variables affect the willingness to pay, we propose the agent applies Analytical Hierarchy Process (AHP) methods, as introduced by Saaty (1980), to prioritise and weight these factors systematically, ensuring a structured evaluation of pricing drivers. The generic equation for AHP in this context can be expressed as:

$$\text{AHP-WTP}_{(b2)} : W_j = \frac{\sum_{i=1}^{n} (C_{ij} \cdot P_{ij})}{\sum_{j=1}^{m} \left(\sum_{i=1}^{n} C_{ij} \cdot P_{ij}\right)} \tag{27}$$

Here, W_j represents the weight of criterion j for a specific customer segment, while C_{ij} denotes the comparison value or score for criterion j with respect to factor i. The term P_{ij} captures the priority or importance assigned to factor i within criterion j. Additionally, n is the number of influencing factors within each criterion, and m represents the number of criteria considered. The resulting weights W_j provide a structured and prioritised evaluation of price-influencing criteria for each customer segment. These weights serve as a key input for determining willingness to pay (WTP), ensuring that the pricing calculation accounts for the most relevant factors and reflects the value perceived by customers.

The Replacement Product Analyst Agent (b3) operates with a specialised focus on identifying potential threats or opportunities related to replacement products or services. This agent analyses the market to detect alternatives that could directly or indirectly replace the company's offerings, including solutions that solve the same customer problem more effectively or efficiently. It also explores opportunities for the company's own products or services to enter new industries or replace existing solutions in adjacent markets. Given the complexity and scope of this task, the agent requires a holistic understanding of product functionality and customer needs. To refine its analysis and ensure precision, humans in the loop can interact with the agent, instructing it to focus on specific products, industries, or solution spaces. Insights generated by this agent are presented to humans for evaluation, and approved findings can be forwarded to the Competitor Price/Product Observer Agent (b1) for further processing and integration. The framework can be expressed as follows:

$$O_{(b3)} : \quad \underset{\mathcal{R},\mathcal{O}}{\text{Analyse and identify}} \left\{ \mathcal{R}(t), \mathcal{O}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{28}$$

$$\begin{aligned}\text{subject to:}\quad & C_1 : \text{Focus areas and priorities refined through human interaction},\\ & C_2 : \text{Analyses include comparisons to identify alternative solutions},\\ & C_3 : \text{Findings are presented to humans for approval},\\ & C_4 : \text{Data collection and evaluation adhere to legal frameworks},\end{aligned}$$

$$C_n : \text{Other constraints specific to real-world applications.}$$

Here, $\mathcal{R}(t)$ represents the identification of replacement threats over time t, while $\mathcal{O}(t)$ captures opportunities for the company's products or services to replace existing solutions in new or adjacent markets. The variable $t \in [t_0, T]$ defines the time window for analysis. Constraint C_1 allows humans to refine focus areas and priorities through direct interaction with the agent, and C_2 ensures that functional comparisons are included in the analysis to identify alternative solutions. Constraint C_3 mandates that all findings are presented to humans for evaluation and approval before being processed further. C_4 enforces compliance with legal and ethical guidelines for data collection and evaluation, while C_n accounts for additional constraints relevant to specific real-world scenarios.

The Replacement Product Analyst Agent (b3) can play an important role in optimising pricing strategies by identifying replacement threats or opportunities for the company's products or services. If a competitor introduces an upcoming product or service that could replace the company's offerings, this agent provides early insights, enabling the company to adjust pricing strategies to remain competitive. For example, the company may lower prices to counter the threat or extend the product lifecycle under competitive pressure. The agent can detect functional alternatives that solve customer problems more effectively or efficiently, which may impact perceived value and demand elasticity.

The group of agents under the Cost Tracking Agent (c) is responsible for continuously monitoring and analysing all cost-related components that influence the lower boundary of the price corridor. These agents systematically track internal costs, handled by the Internal Costs Analyst Agent (c1), external contract terms, managed by the Contract Analyst Agent (c2), and supply chain dynamics, monitored by the Supply Chain Analyst Agent (c3), ensuring the data provided is accurate, structured, and reflects real-time developments. The Cost Tracking Agent (c) coordinates its sub-agents and consolidates their outputs into a coherent framework to support the Price Optimisation Agent (a) in determining cost-based pricing inputs. We propose also here that all gathered data is processed, time-stamped, and integrated into a central cost database to ensure traceability and enable trend analysis. This allows for continuous calculating and monitoring of cost degression, economies of scale, and external cost drivers that influence the lower price boundary.

The Internal Costs Analyst Agent (c1) is tasked with monitoring and analysing all internal cost components necessary for determining the lower boundary of the price corridor. This agent systematically differentiates between fixed costs and variable costs, ensuring data is gathered from controlling departments, accounting systems, bookkeeping processes, human-in-the-loop inputs, and other relevant sources. It processes key cost structures, such as unit cost calculations, project-based costing, and other cost allocation methods relevant to business operations. For fixed costs, the agent monitors expenses that remain constant regardless of production levels, such as rent, salaries, and depreciation schedules. For variable costs, it evaluates expenses that fluctuate with output, including raw materials, energy consumption, and production-related logistics.

The agent must account for economies of scale with a clear data foundation, as these are often contractually defined or based on predetermined agreements, such as bulk purchase discounts or tiered pricing. This requires close interaction with the Contract Analyst Agent (c2) to assess supplier agreements and negotiated cost reductions, as well as the Supply Chain Analyst Agent (c3) to evaluate dynamic supply chain costs and alternative sourcing options. The agent incorporates historical data, forecasts, and trends to reflect changes in internal costs, such as inflation, resource price volatility, and operational efficiencies. Integrating these inputs, including human-in-the-loop assessments, the agent ensures that the lower price boundary reflects both current cost realities and opportunities for cost optimisation. This comprehensive approach allows for the accurate calculation of cost-based inputs, accounting for fixed and variable dynamics, economies of scale, and alternative sourcing possibilities. The objective-under-constraints framework of the Internal Costs Analyst Agent (c1) can be expressed as:

$$O_{(c1)} : \quad \underset{\mathcal{C}_{\text{fixed}}, \mathcal{C}_{\text{variable}}}{\text{Monitor and analyse}} \left\{ \mathcal{C}_{\text{fixed}}(t), \mathcal{C}_{\text{variable}}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{29}$$

$$\begin{aligned} \text{subject to:} \quad & C_1 : \text{Fixed and variable costs are accurately differentiated,} \\ & C_2 : \text{Data is sourced from controlling, accounting, and reliable inputs,} \\ & C_3 : \text{Economies of scale are integrated using contractual insights,} \\ & C_4 : \text{Cost trends and forecasts, including inflation, are considered,} \\ & C_n : \text{Other constraints specific to real-world applications.} \end{aligned}$$

Here, $\mathcal{C}_{\text{fixed}}(t)$ and $\mathcal{C}_{\text{variable}}(t)$ represent fixed and variable internal costs over time t. Constraint C_1 ensures accurate cost categorisation, while C_2 mandates reliable data sources, including controlling and accounting.

C_3 incorporates economies of scale based on contractual agreements or data-driven models, and C_4 accounts for cost trends, inflation, and operational efficiencies. C_n addresses additional constraints relevant to specific applications. The Internal Costs Analyst Agent (c1) ensures the accurate integration of cost components into the optimisation process as defined by equations (14), (15), and (16).

This agent provides critical inputs to the calculation of $C(Q)$, which represents the unit cost as a function of production levels. Fixed costs (FC) and variable costs (VC) are dynamically adjusted based on production volume (Q) and economies of scale. These inputs feed into $P_{\min}$, ensuring that cost structures align with real-time data and operational conditions. The agent accurately represents cost degression effects and bulk purchase discounts, allowing the optimisation framework to reflect realistic and achievable pricing boundaries. These precise calculations play a pivotal role in determining the contribution margin, enabling pricing strategies to remain cost-covering and competitive while maximising profitability.

The Contract Analyst Agent (c2) focuses on analysing contract terms, particularly within supply chain management, to provide cost-related insights. Unlike the mathematically oriented architecture of the Internal Costs Analyst Agent (c1), this agent relies on Large Language Model (LLM) capabilities to reason and interpret the often complex and context-dependent nature of legal contracts. It must understand specific terms within the framework of international pricing strategies, including diversification and market-specific agreements.

The agent interprets clauses such as supplier pricing tiers, volume discounts, and delivery conditions to identify terms relevant to cost calculations. It works in coordination with the Internal Costs Analyst Agent (c1) and the Price Optimisation Agent (a) to contextualise contractual data for accurate integration into pricing decisions. The objective-under-constraints framework of the Contract Analyst Agent (c2) can be expressed as:

$$O_{(c2)}: \quad \underset{\mathcal{T}_{\text{contract}}}{\text{Analyse and interpret}} \left\{ \mathcal{T}_{\text{contract}}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{30}$$

$$\begin{aligned} \text{subject to:} \quad & C_1 : \text{Contracts are interpreted for international pricing strategies,} \\ & C_2 : \text{Key terms are prioritised,} \\ & C_3 : \text{Complex clauses are resolved via LLM capabilities,} \\ & C_4 : \text{Data aligns with inputs for agents (c1) and (a),} \\ & C_n : \text{Other real-world constraints.} \end{aligned}$$

Here, $\mathcal{T}_{\text{contract}}(t)$ represents the time-dependent analysis of contract terms, ensuring that key factors such as volume discounts and delivery conditions are prioritised (C_2). C_1 ensures the context of international pricing strategies is considered, and C_3 employs LLM capabilities to interpret complex clauses. C_4 mandates alignment of contractual data with the inputs required by *Cost Tracking Agent* (c1) and *Price Optimisation Agent* (*a*). C_n accommodates additional constraints for real-world scenarios. A possible human-in-the-loop integration for the Contract Analyst Agent (c2) could involve collaboration with legal departments or contract specialists, who can supply necessary documents or validate interpretations. This agent must be capable of analysing various document formats, such as PDFs or other structured and unstructured legal texts. Given its focus on sensitive contract-level information, the agent must operate within a highly secure environment to prevent accidental exposure of confidential details. Oversight from the compliance department, or a designated compliance agent, is essential to ensure adherence to confidentiality guidelines and prevent unauthorised dissemination of sensitive information. As an alternative, we propose that instead of directly sharing contractual data with other agents, the Contract Analyst Agent (c2) could reverse the information flow. In this approach, other agents query the Contract Analyst Agent (c2) for specific data points, such as agreed-upon prices, volume discounts, or delivery terms, without receiving the full contract details. The Contract Analyst Agent (c2) validates these requests, approving or rejecting queries based on compliance rules, and provides only the necessary values for integration into the pricing process. This design ensures the system maintains confidentiality while enabling seamless collaboration across agents. The Supply Chain Analyst Agent (c3) monitors the supply chain, tracking current suppliers and identifying potential alternatives by building detailed profiles and ratings for each supplier. This agent continuously evaluates supplier pricing, pricing models, and relevant conditions, providing real-time insights into opportunities for cost reductions. For example, if a supplier with a rating of 87 lowers the price for specific parts, the agent could identify the potential cost savings and recommend switching suppliers for a particular production series or more broadly.

In addition to cost monitoring, the Supply Chain Analyst Agent (c3) supports calculations related to Nash equilibrium, as expressed in (19), and principal-agent pricing models, as formalised in (20). These frameworks address information asymmetry and optimise interactions within the supply chain. Given that supply chain monitoring often involves implicit information and context rather than explicitly stated data, the agent requires advanced reasoning and contextual understanding capabilities to integrate these elements effectively into the pricing framework. This ensures it can process and interpret complex data to provide actionable recommendations that align with cost optimisation and strategic objectives. The objective-under-constraints framework of the Supply Chain Analyst Agent (c3) can be expressed as:

$$O_{(c3)}: \quad \underset{\mathcal{S}}{\text{Monitor and evaluate}} \left\{ \mathcal{S}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{31}$$

$$\begin{aligned} \text{subject to:} \quad & C_1 : \text{Supplier profiles and ratings are continuously updated,} \\ & C_2 : \text{Pricing changes and conditions are monitored in real time,} \\ & C_3 : \text{Recommendations align with goals and cost-saving strategies,} \\ & C_4 : \text{Implicit context is interpreted using advanced reasoning,} \\ & C_5 : \text{New supplier opportunities are identified,} \\ & C_6 : \text{Outputs support Nash and principal-agent calculations,} \\ & C_7 : \text{Compliance with legal and regulatory standards is ensured,} \\ & C_n : \text{Other real-world constraints.} \end{aligned}$$

Here, $\mathcal{S}(t)$ represents the supply chain data monitored over time t, including supplier profiles, ratings, pricing changes, and conditions. Constraint C_1 ensures that supplier profiles and evaluations are continuously updated, while C_2 mandates real-time monitoring of pricing changes and supplier conditions. C_3 aligns recommendations with production goals and cost-saving strategies, and C_4 employs advanced reasoning to interpret implicit context and unstructured data. C_5 identifies opportunities for new suppliers, and C_6 supports calculations for Nash equilibrium and principal-agent models to address information asymmetry. C_7 ensures all outputs comply with ethical, legal, and regulatory standards, and C_n accounts for additional real-world constraints.

The Supply Chain Analyst Agent (c3) could benefit from a strong human-in-the-loop component to enhance its performance. Humans can provide specific instructions, such as identifying key suppliers to monitor more closely or prioritising potential new supply partners and stakeholders. Ahn et al. [1] propose a method within their Generative Probabilistic Planning (GPP) framework that combines attention-based graph neural networks (GNNs), offline deep reinforcement learning (Offline RL), and policy simulations to dynamically optimise supply chain actions, utilising historical data to make informed decisions under uncertainty and adapt to changing objectives such as maximising profits or service levels. Quan and Liu (2024) introduce InvAgent, a LLM-based zero-shot multi-agent system for inventory management that significantly enhances supply chain resilience and efficiency through adaptive decision-making, superior explainability with CoT integration, and dynamic responses to fluctuating demands, as validated by extensive evaluations.

To improve the efficiency of agent (c3), we also propose the implementation of an AI-readable, supply chain-wide, central, high-quality, data-driven marketplace that is accessible to all participants and AI agents within the network. This marketplace would allow the agents to access and evaluate profiles, ratings, pricing models, and other critical supplier information in a structured format. The agents could also analyse additional metrics such as Environmental, Social, and Governance (ESG) ratings, financial credibility, and product specialisation to identify high-rated suppliers or strategic opportunities. The agents could then recommend specific actions to humans, such as initiating contact with potential partners. The same marketplace could also support the agents in identifying potential clients and customers within the supply chain, creating opportunities on both the supply and demand sides. The Historical Pricing Strategy Observer Agent (d) focuses on building and maintaining a structured, time-stamped archive of pricing data. This agent monitors the development of the company's historical prices and tracks competitor price trends over time. The inclusion of timestamps is essential for constructing time series data, enabling the agent to analyse temporal price evolutions, detect patterns, and identify long-term pricing trends. A robust time series database allows for the application of forecasting methods to anticipate future price movements or anomalies. The agent deploys advanced analytical and AI algorithms to derive actionable insights. For example, regression analysis can predict price developments based on historical trends and influencing factors, while seasonal decomposition techniques identify periodic patterns within pricing data. Other techniques, such as exponential smoothing and AutoRegressive Integrated Moving Average (ARIMA) models, can support accurate forecasting of pricing trajectories.

These insights assist the Price Optimisation Agent (a) in understanding past behaviours, evaluating strategic adjustments, and predicting optimal price corridors under evolving market conditions. The objective-under-constraints framework of the Historical Pricing Strategy Observer Agent (d) can be expressed as:

$$O_{(d)}: \quad \underset{\mathcal{P}_{\text{time}}}{\text{Store and analyse}} \left\{ \mathcal{P}_{\text{time}}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{32}$$

$$\begin{aligned} \text{subject to:} \quad & C_1: \text{Data must include timestamps for time series construction,} \\ & C_2: \text{Ensure accuracy and consistency of archived pricing data,} \\ & C_3: \text{Collected data must comply with privacy and legal regulations,} \\ & C_4: \text{Integration with external sources for competitor pricing is permitted,} \\ & C_n: \text{Other constraints specific to real-world applications.} \end{aligned}$$

Here, $\mathcal{P}_{\text{time}}(t)$ represents time-stamped pricing data as a function of time t. Constraint C_1 ensures that timestamps are included for constructing time series models, and C_2 mandates accuracy and consistency in archived data. C_3 enforces compliance with privacy and legal standards, while C_4 permits integration with external sources for competitor pricing. C_n accounts for additional real-world constraints.

The Historical Pricing Strategy Observer Agent (d) cooperates closely with other agents by exchanging critical data and insights. It works in coordination with the Competitor Price/Product Observer Agent (b1) to gather competitor pricing trends and align findings with real-time observations. The Price Optimisation Agent (a) utilises this agent's time-stamped data to calibrate its pricing calculations, particularly when incorporating historical performance into optimisation models. The agent supports the Customer and Client Behaviour Analyst Agent (b2) in assessing price elasticity based on past behaviours and identifying segment-specific responses to price changes.

The Corporate Pricing Strategy Agent (e) is responsible for managing and analysing the company's overarching pricing strategies. This agent evaluates, updates, and ensures alignment of pricing strategies with corporate objectives, such as profitability targets, market positioning, and competitive dynamics. It maintains a structured repository of applied pricing strategies, including cost-plus, value-based, and psychological pricing methods, to guide the Price Optimisation Agent (a) and ensure that all pricing decisions align with the company's strategic goals. This agent incorporates data from other agents, such as the Historical Pricing Strategy Observer Agent (d), to assess the effectiveness of past strategies and adjust current ones accordingly. The Corporate Pricing Strategy Agent (e) monitors market trends, competitor approaches, and customer behaviour to refine and adapt pricing strategies in response to dynamic business environments. The objective-under-constraints framework of the Corporate Pricing Strategy Agent (e) can be expressed as:

$$O_{(e)}: \quad \underset{\mathcal{S}_{\text{pricing}}}{\text{Evaluate and adapt}} \left\{ \mathcal{S}_{\text{pricing}}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{33}$$

$$\begin{aligned} \text{subject to:} \quad & C_1: \text{Pricing strategies align with corporate goals,} \\ & C_2: \text{Strategies comply with legal and ethical standards,} \\ & C_3: \text{Dynamic inputs from agents (d), (b2), and (b3) are incorporated,} \\ & C_4: \text{Adjustments to strategies are time-stamped for traceability,} \\ & C_n: \text{Other constraints specific to real-world applications.} \end{aligned}$$

Here, $\mathcal{S}_{\text{pricing}}(t)$ represents pricing strategies as a function of time t. Constraint C_1 ensures alignment with corporate goals such as profitability and market positioning, while C_2 mandates compliance with legal and ethical standards. C_3 incorporates dynamic inputs from Time-Stamped Pricing Data Agent (d), Customer Behaviour Analyst Agent (b2), and Replacement Product Analyst Agent (b3). C_4 requires adjustments to strategies to be time-stamped for traceability, and C_n accounts for additional real-world constraints. This agent collaborates closely with the Price Optimisation Agent (a) to ensure alignment between strategic goals and operational pricing decisions. It also works with the Historical Pricing Strategy Observer Agent (d) to evaluate the outcomes of past strategies, integrating insights into future adjustments.

The Price Policy Effect Agent (f) evaluates the historical effects of implemented pricing strategies across different markets, identifying how past price adjustments influenced sales performance. It focuses on analysing the relationship between price changes and corresponding sales volumes, commonly referred to as price-demand elasticity.

The agent operates with a strong mathematical focus, modelling and optimising the nonlinear price-demand function to determine how price adjustments affect sales quantities, customer retention, and acquisition rates. Regression analysis supports forecasting demand under various pricing scenarios while accounting for additional factors such as marketing activities, competitor pricing, and seasonal trends.

Evaluating this relationships identifies optimal price points that maximise overall profitability while accounting for contribution margins and production costs. The agent generates models to predict demand responses under various pricing scenarios, ensuring alignment with the company's strategic goals.

It assesses how price policy changes impact total revenue, profitability, and unit economics, creating a foundation for dynamic price optimisation. The objective-under-constraints framework of the Price Policy Effect Agent (f) can be expressed as:

$$O_{(f)}: \quad \underset{\mathcal{P}_{\text{opt}},\mathcal{Q}}{\text{Optimise}} \left\{ \mathcal{P}_{\text{opt}}(t), \mathcal{Q}(t) \,\middle|\, t \in [t_0, T] \right\} \tag{34}$$

$$\begin{aligned} \text{subject to:} \quad & C_1 : \text{Price-demand relationship uses historical and real-time data,} \\ & C_2 : \text{Optimisations consider contribution margins,} \\ & C_3 : \text{Nonlinear price elasticity is included,} \\ & C_4 : \text{Models align with goals set by agent (g),} \\ & C_n : \text{Other real-world constraints.} \end{aligned}$$

Here, $\mathcal{P}_{\text{opt}}(t)$ represents the optimised price over time t, while $\mathcal{Q}(t)$ denotes the demand function. Constraint C_1 ensures that the price-demand relationship is modelled using both historical and real-time data, and C_2 incorporates contribution margins and unit profitability. C_3 requires the analysis to include nonlinear price elasticity effects, and C_4 ensures alignment with profitability goals set by *Corporate Strategy Agent* (g). C_n accounts for additional constraints relevant to real-world scenarios. To further support the agent's ability to analyse price-demand relationships and optimise price policy effects, we propose incorporating regression analysis for forecasting demand based on pricing decisions and other influencing factors which can be expressed as: To further support the agent's ability to analyse price-demand relationships and optimise price policy effects, we propose incorporating regression analysis for forecasting demand based on pricing decisions and other influencing factors, which can be expressed as:

$$Q = \beta_0 + \beta_1 P + \beta_2 P^2 + \sum_{i=1}^{n} \gamma_i X_i + \varepsilon \tag{35}$$

Here, Q represents demand or quantity sold (dependent variable), and P is the price of the product (independent variable). The term P^2 captures nonlinear price effects, such as diminishing or increasing elasticity. X_i includes additional explanatory variables, such as competitor prices, marketing spend, or seasonal effects. β_0 is the intercept term, while β_1 and β_2 are the coefficients for price and price squared, respectively. γ_i represents the coefficients for other influencing variables X_i, and ε is the error term accounting for unobserved factors.

This regression model enables the agent to quantify the impact of price changes on sales volume, identify non-linear price elasticity effects, and incorporate additional factors such as marketing activities, competitor pricing, or seasonal trends. The coefficients derived from the regression provide actionable insights that the agent can use to recommend adjustments to optimise demand and maximise profitability. The agent collaborates closely with the Price Optimisation Agent (a), providing critical input for dynamic pricing decisions.

It also integrates with the Time-Series Price Analysis Agent (d) to incorporate historical trends and forecasts into its analysis.

The Corporate Strategy Agent (g) ensures that the calculated prices align with the company's overarching strategic objectives. This agent acts as a reference framework, holding the strategic guidelines that dictate the pricing approach, such as premium positioning, cost leadership, or market penetration strategies. It communicates directly with the Price Optimisation Agent (a), validating that the chosen prices reflect the company's strategic intent.

For example, in a premium strategy, prices must remain high to preserve perceived value and exclusivity, while in a cost leadership strategy, lower prices must reflect optimised production costs and higher sales volumes. Agent (g) can instruct Agent (a) to adjust price calculations where misalignments occur, ensuring coherence between pricing outcomes and strategic goals.

3.10 Relevance of AI Techniques Across Multi-Agent System

The integration of reinforcement learning into multi-agent systems has significantly enhanced real-time price optimisation. Agents are capable of dynamically adjusting prices in response to market fluctuations, thereby improving adaptability and profitability [31]. These technologies facilitate real-time pricing adjustments based on demand, competition, and customer behaviour, rendering pricing strategies more responsive and data-driven. Recent research examines the significance of uncertainty modelling in enhancing the reliability and adaptability of multi-agent systems.

Lockwood and Si [34] discuss methodologies to quantify and manage uncertainty in reinforcement learning, emphasising its critical role in ensuring reliable decision-making in dynamic environments. Similarly, Lütjens et al [37] propose incorporating model uncertainty estimates into reinforcement learning, enhancing safety by mitigating overconfident decisions in unfamiliar scenarios.

These approaches are particularly pertinent for AI-driven pricing strategies, which must navigate uncertain and unpredictable market conditions. Further advancements in uncertainty modelling focus on balancing optimism and pessimism in risk-aware reinforcement learning. Vlastelica et al [53] address both epistemic (knowledge-based) and aleatoric (inherent) uncertainties, improving decision-making robustness in safety-critical applications such as dynamic pricing systems.

Charpentier et al [9] articulate the challenges in achieving real-time performance for agents in reinforcement learning due to the inherent difficulties in accurately predicting and managing aleatoric and epistemic uncertainties within dynamic environments. Jin et al [29] present TIME-LLM, a framework that effectively repurposes large language models for time series forecasting by transforming input data into text prototypes and using "Prompt-as-Prefix" to enhance reasoning, demonstrating superior performance in few-shot and zero-shot scenarios without altering the original model architecture. Becker and Neumann [5] examine the effects of overestimated aleatoric uncertainty in deep state space models, noting it serves as implicit regularisation that improves robustness but can compromise performance in tasks demanding accurate uncertainty assessments, such as environments with occlusions or heterogeneous sensor inputs.
Table 1 presents the subjective relevance scores of selected artificial intelligence techniques across multiple dimensions, representing their applicability in a multi-agent system context. The table highlights the techniques deemed most critical for discussion, based on their significant roles in enhancing the performance, reliability, and decision-making capabilities within complex AI-driven applications.

Table 1 Subjective relevance scores of AI techniques across multi-agent system components.

Technique	**(a)**	**(b)**	**(b1)**	**(b2)**	**(b3)**	**(c)**	**(c1)**	**(c2)**	**(c3)**	**(e)**	**(f)**	**(g)**
Machine Learning	9	8	7	8	6	9	7	8	9	7	9	8
Reinforcement Learning	7	6	5	6	5	7	5	6	6	5	7	6
Large Language Models	8	8	7	8	7	6	9	7	7	6	8	7
Time Series	7	7	7	8	6	7	6	7	9	6	8	7
Optimisation	10	9	7	8	7	9	6	8	8	7	9	8
Data Retrieval	8	8	8	9	8	9	9	9	9	7	8	7
Knowledge Graphs	7	7	6	7	6	8	7	8	7	6	8	7
Uncertainty Modelling	7	6	5	6	6	7	6	6	7	5	6	6
Graph RAG	6	8	7	8	6	8	8	7	8	6	7	6

The table reflects a deliberate evaluation of AI techniques based on their roles in enhancing the multi-agent system's ability to perform dynamic pricing. Machine Learning achieves high scores across components due to its predictive accuracy and adaptability in complex environments. Reinforcement Learning supports iterative bargaining strategies, making it crucial for negotiation scenarios. Large Language Models excel in processing natural language inputs, enabling effective communication with sales teams and customers. Time Series methods are integral for forecasting demand and adjusting pricing strategies dynamically. Optimisation techniques provide the foundation for determining efficient price corridors and thresholds. Data Retrieval ensures the integration of real-time data, supporting informed decision-making. Knowledge Graphs facilitate reasoning over complex relationships among data entities. Uncertainty Modelling enhances the system's robustness by managing quality and mitigating risks. Graph RAG adds contextual understanding through retrieval-augmented generation, proving useful in refining negotiation strategies and generating context-aware suggestions.

In practical applications, the methodologies and associated evaluations may vary significantly based on the organisation's specific context, including the product, market, and business model, and the resultant artificial intelligence model configuration determined by these factors.

To augment the training efficiency of the AI models, particularly those tasked with processing natural language inputs or outputs, we propose the adoption of a self-distillation method inspired by Zhang et al [54]. This approach utilises the internal capabilities of large language models to refine and optimise their own performance without extensive retraining on new datasets. Self-distillation, as detailed in their study, involves the model generating its own training data by creating outputs that are then utilised as new inputs in a repetitive cycle.

Implementing self-distillation could significantly enhance the responsiveness and adaptability of our AI systems, particularly in scenarios requiring dynamic interaction with sales teams and real-time negotiation. Through the efficient processing of complex inputs, including customer queries, market updates, and pricing adjustments, the artificial intelligence system enhances its capacity to support strategic decision-making. This method will be particularly beneficial in the 'Real-Time Data Processing' module of our AI framework, ensuring that the system not only learns more effectively from limited data but also adapts more rapidly to changes in pricing contexts, thereby enhancing overall system performance and reliability in dynamic pricing environments.

3.11 Generic High-Level Overview of Agent Communication and Task Delegation

The Price Optimisation Agent (a) functions as the primary coordinator, facilitating communication and allocating tasks to other agents within the system. Communication is structured through standardised data formats, such as JSON or XML, and shared ontologies to ensure consistency in data interpretation and efficient collaboration.

Predefined APIs enable the Price Optimisation Agent (a) to issue specific queries, such as requesting competitor pricing data from the Competitor Price/Product Observer Agent (b1) or cost updates from the Internal Costs Analyst Agent (c1). Pricing tasks are deconstructed into subtasks and assigned to agents based on their roles and expertise.
A task-oriented communication framework within multi-agent systems, advanced by He [24], utilises deep reinforcement learning to optimise the relevance and efficiency of information flow. This framework addresses architectural and practical challenges while highlighting the necessity for future research in semantic theory and system design to enhance multi-agent collaborative decision-making. Multi-agent planning techniques, including algorithms such as the Contract Net Protocol, dynamically allocate tasks and resolve inter-agent dependencies, effectively managing the coordination and cooperation among agents to optimise the overall system performance by distributing responsibilities based on agent capabilities and current workload [36].

Dynamic scheduling adjusts task priorities in real time, accounting for factors such as supply chain disruptions reported by the Supply Chain Analyst Agent (c3). Data is exchanged via vector databases, which store and manage time-stamped records for trend analysis, and knowledge graphs, which map relationships among agents, pricing variables, and market dynamics. Advanced coordination methods enhance collaboration. Reinforcement learning optimises task delegation by analysing historical interactions among agents and improving future efficiency.

Game-theoretic approaches, such as Nash equilibrium models, balance competing objectives like cost minimisation and competitive pricing strategies. Discrepancies between agents are resolved through negotiation algorithms, while consensus mechanisms confirm shared decisions, such as validating pricing boundaries proposed by the Corporate Pricing Strategy Agent (e). This workflow presents a generic model for multi-agent system communication and task delegation. In practice, specific workflows and configurations may vary depending on the company's product, market, and operational context. To formalise the communication and data aggregation process of the Price Optimisation Agent (a) with other agents, the following equation models the interaction framework:

$$R_a = \sum_{i=1}^{n} w_i \cdot F_i(Q_i, D_i) \tag{36}$$

Here, R_a represents the aggregated response or output of the Price Optimisation Agent (a). n is the total number of interacting agents, such as $b1$, $c1$, and $c3$. w_i denotes the weight or priority assigned to Agent i, while F_i is the function of Agent i, which depends on its query (Q_i) and data state (D_i). Q_i is the query issued by Agent (a) to Agent i, and D_i represents the data or state of Agent i.

This equation allows for an adaptive weighting mechanism (w_i), enabling the Price Optimisation Agent (a) to prioritise responses dynamically based on real-time market conditions or specific pricing scenarios. For instance, reinforcement learning can be employed to fine-tune w_i over time by analysing historical task performance and agent reliability, with inputs from the Price Policy Effect Agent (f), which evaluates the outcomes of past pricing strategies. Modelling and optimising nonlinear price-demand relationships, the Price Policy Effect Agent (f) provides actionable insights that enable the Price Optimisation Agent (a) to adjust its data aggregation priorities and improve decision-making accuracy dynamically.

The aggregated response (R_a) derived through this interaction framework serves as the foundational input for key pricing equations in the Weighted Dynamic Corridor Price Optimisation model, including the dynamic boundaries ($P_{\min}$ and $P_{\max}$, including the dynamic boundaries ($P_{\min}$ and $P_{\max}$, equations (3), (4), (5), (6), (7), (8), and (9)) and the optimisation of the profit-maximising price (P_{optimal}, equation (13)).

4 Discussion

The AI supported Weighted Dynamic Corridor Price Optimisation (WDCPO) model marks an improvement in the field of dynamic pricing techniques. This approach enables swift and adaptable price adjustments by establishing a pricing range with upper and lower boundaries that consider production costs, consumer price tolerance, and market sensitivity. The corridor pricing strategy enables dynamic modifications based on various market-influencing factors. Unlike traditional static methods, this framework ensures profitability whilst maintaining a competitive edge in fluctuating markets. The model's foundation lies in a multi-agent artificial intelligence system, which introduces a structured and modular approach. Each agent functions within well-defined parameters and limitations, contributing to a cohesive pricing strategy. For example, the Price Optimisation Agent amalgamates inputs from cost analysis, market surveillance, and corporate strategy to determine optimal pricing solutions. This design simplifies the complex task of aligning pricing decisions with strategic goals while preserving flexibility.

Benefits: The proposed system offers advantages, particularly in supply chain management and dynamic pricing optimisation. Its capacity to integrate real-time data from multiple sources ensures that pricing strategies remain responsive to fluctuations in production costs, market demand, and competitive conditions. Advanced artificial intelligence techniques, such as reinforcement learning and large language models (LLMs), enable the system to process and interpret data effectively, providing actionable insights for decision-makers. The introduction of cooperative and semi-cooperative pricing frameworks enhances the system's versatility. The Fully Cooperative Price Equilibrium Framework, based on Nash equilibrium principles, allows supply chain participants to align on pricing strategies that maximise collective profitability. In semi-cooperative scenarios, the system applies principal-agent theory to address information asymmetry, ensuring that pricing decisions remain efficient and mutually beneficial. Another key advantage lies in the modular design of the system. Each agent operates independently, with tasks clearly delineated and aligned with overall strategic objectives. This structure enhances scalability and flexibility, allowing the system to adapt to diverse market conditions and operational contexts. Moreover, the inclusion of a centralised, AI-readable supply chain marketplace enables seamless collaboration between agents, improving cost management and decision-making efficiency.

Challenges: Despite its advantages, the system encounters significant challenges related to implementation and data dependency. The development and maintenance of such a multi-agent system necessitates considerable expertise in artificial intelligence development, access to high-quality data, and substantial computational resources. For smaller organisations, these requirements may present barriers to adoption, particularly in the absence of modular, pre-trained artificial intelligence systems tailored to specific industries. Data quality constitutes another critical challenge. The framework relies on consistent and accurate data inputs from diverse sources, including production systems, market analytics, and customer platforms. Fragmented or incomplete datasets can undermine the accuracy of pricing calculations and forecasts, thereby reducing the system's effectiveness in real-world scenarios. Industries with sparse or unstructured data may encounter difficulties in fully leveraging the framework's capabilities. Real-time processing and coordination between agents also impose high computational demands. In complex supply chain environments, ensuring seamless communication and task delegation among agents requires advanced infrastructure and well-designed protocols. In the absence of these, the system's efficiency and accuracy may be compromised.

Legal and Ethical Considerations: The framework proposed in this paper integrates ethical and legal considerations directly into the constraint structures of the domain-specific agents, ensuring these factors are foundational rather than supplementary. Each agent operates within an objective-under-constraints framework where compliance with legal standards and ethical guidelines is explicitly embedded. These constraints are designed to accommodate the specific requirements of the agent's function, such as adherence to competition laws in market analysis or data protection regulations in customer behaviour modelling. This integration assures that the proposed system operates within a legally sound and ethically foundation, irrespective of the agent's domain or task. The system's incorporation of these considerations within its constraint framework enables it to be both scalable and adaptable, whilst reducing the likelihood of breaches. The modular architecture of the system enables the integration of ethical considerations into agent-specific constraints. For instance, agents tasked with competitor analysis must observe antitrust regulations to avoid breaching competition laws. Similarly, agents analysing customer behaviour must adhere to GDPR and similar laws to protect individual privacy and prevent discrimination [11].

The implementation must address several legal and ethical considerations, particularly regarding data privacy and artificial intelligence regulation. Data privacy laws, such as the General Data Protection Regulation (GDPR; European Commission, 2016), require organisations to ensure transparent, secure, and lawful processing of personal data [11]. GDPR specifically mandates that AI systems processing personal data adhere to principles of data minimisation, purpose limitation, and accountability. Similarly, the California Consumer Privacy Act (CCPA; California Legislature, 2018) grants residents the right to access, delete, and opt out of the sale of their personal data, impacting how AI models collect and use consumer information [33]. In Canada, the Personal Information Protection and Electronic Documents Act (PIPEDA; Government of Canada, 2000) governs data collection and obligates organisations to obtain meaningful consent and safeguard personal information [8]. Brazil's Lei Geral de Proteção de Dados (LGPD; Presidência da República, 2018) and South Africa's Protection of Personal Information Act (POPIA; South Africa Government, 2013) reflect similar principles, adapted to local contexts [46, 22]. Singapore's Personal Data Protection Act (PDPA; Singapore Government, 2012) underscores this global trend toward harmonising data privacy standards [21]. Artificial intelligence frameworks are equally crucial for structuring the proposed system. The European Union's Artificial Intelligence Act (AI Act; European Commission, 2021) establishes harmonised rules for trustworthy AI by categorising systems based on risk levels and enforcing accountability measures for high-risk applications [12]. Additionally, UNESCO's Recommendation on the Ethics of Artificial Intelligence (2021) promotes transparency, fairness, and human oversight [52]. In the United States, the Blueprint for an AI Bill of Rights (White House, 2022) outlines principles for safe and equitable AI [28], further reflected in Canada's Directive on Automated Decision-Making (Government of Canada, 2019), which sets clear standards for transparency and bias mitigation in automated processes. Compliance also extends to laws governing international markets. China's Personal Information Protection Law (PIPL) [45] and Saudi Arabia's Personal Data Protection Law (PDPL, 2021) impose stringent localisation and consent requirements. AI-specific guidelines, such as the OECD AI Principles (2019) [44] and the G20 AI Principles (2019) [17], provide additional guidance for supporting human-centric and ethical AI development across jurisdictions.

Future Directions: Advancing the proposed system necessitates addressing its limitations while enhancing its usability. One immediate priority is the development of modular, pre-trained artificial intelligence agents that can be customised for specific industries and scales of operation. Such agents would lower barriers to entry for smaller businesses, reducing deployment times and resource requirements. The integration of human-in-the-loop functionality presents another area for improvement. Allowing human oversight and interaction with individual agents can enhance the system's flexibility and address potential misalignments. Research into workflows that optimise human-artificial intelligence collaboration, including user-friendly interfaces and transparency in decision logic, would render the system more accessible to non-technical stakeholders. Scalability and interoperability should also be prioritised. Extending the system to support global markets with diverse regulatory and economic conditions requires robust algorithms capable of handling cross-border complexities. Exploring advanced techniques, such as graph neural networks and hybrid reinforcement learning models, may further optimise decision-making processes in highly variable environments. The incorporation of emerging technologies, such as the Internet of Things (IoT) and blockchain, could enhance data reliability and granularity. IoT devices can provide real-time inputs on production and logistics, while blockchain can validate supply chain transactions, ensuring data integrity.

Conclusion: The WDCPO framework and its associated multi-agent AI system present an innovative method for dynamic pricing. This model employs sophisticated AI techniques and concepts from game theory to ensure adaptability, profitability, and competitive alignment across diverse market scenarios. The system's modular design and collaborative structures address key challenges in supply chain management, enabling participants to improve pricing strategies collectively or individually. Although the framework is theoretically robust, its intricacy introduces potential obstacles. As the system expands, it may become more susceptible to errors and require considerable upkeep to maintain consistency and dependability. Recognising and resolving inefficiencies could necessitate substantial effort, particularly in troubleshooting, agent coordination, and process refinement. As a result, additional empirical testing is essential to assess its effectiveness in real-world applications. Enhancing the system will necessitate further research into modular agent development, human-AI collaboration, and scalability.

5 Limitation

This study focuses on the theoretical development of the Weighted Dynamic Corridor Price Optimisation (WDCPO) model and its supporting multi-agent AI system. While the conceptual framework is robust, its validation relies on further empirical testing. The absence of real-world implementation and field data introduces uncertainties regarding its scalability and performance in diverse industries and markets.

The model's dependency on high-quality, structured data is another consideration. Industries with fragmented or limited datasets may encounter challenges in fully realising the system's potential. Furthermore, the modular design and multi-agent architecture necessitate advanced computational infrastructure and expertise, which may require tailored adaptations for implementation in specific organisational contexts.

The study does not address all potential use cases exhaustively. Instead, it focuses on establishing a foundational framework that can be adapted to varying needs. This inherent adaptability, while advantageous, means the system's real-world customisation and performance remain subject to practical testing.

6 Future Research

Future research should prioritise empirical validation of the proposed framework through controlled experiments and pilot studies. Testing the system in specific industries, such as retail, manufacturing, or logistics, would provide valuable insights into its effectiveness and areas for refinement. Metrics such as profitability optimisation, adaptability to market changes, and agent collaboration efficiency should guide these evaluations. Further research is also required to explore how human involvement, or human-in-the-loop (HITL) systems, can be effectively integrated into the model. This includes defining the specific roles and functions humans should perform within the system, such as overseeing critical decisions, intervening in exceptional cases, or fine-tuning agent outputs.

A particular focus on the psychological aspects of HITL systems is essential, as the ease of human interaction and cognitive alignment with AI systems directly impacts usability, trust, and error reduction. There is a need to study how the structure and architecture of multi-agent systems can remain comprehensible to human operators. This involves balancing the use of domain-specific agents with the application of diverse algorithms and AI techniques. Research should assess the optimal level of human understandability in these systems, ensuring that domain-specific agents are sufficiently transparent and maintainable. Each agent's objectives, data sources, and decision logic must be designed to facilitate human oversight and intervention, reducing errors and improving adaptability in dynamic scenarios.

The development of modular, pre-trained AI agents tailored to specific domains remains an essential next step. These agents could be fine-tuned with industry-specific data, reducing deployment time and resource requirements. Creating accessible, scalable versions of the system would make it suitable for organisations of varying sizes and operational scales. Designing workflows and tools to visualise agent interactions and decision-making logic will help bridge the gap between technical complexity and managerial usability. A focus on ensuring the system remains human-understandable at a specific level will mitigate risks of misalignment and facilitate adoption by non-technical stakeholders. Expanding the framework's capabilities to global, multi-market environments is another key area. Research into algorithms for handling cross-border pricing strategies and regulatory compliance would increase the system's relevance across diverse economic contexts. Emerging technologies, such as IoT for real-time data collection and blockchain for data integrity, should also be investigated to enhance the system's accuracy and reliability.

References

[1] Ahn H, Olivar S, Mehta H, et al (2024) Generative probabilistic planning for optimizing supply chain networks. arXiv preprint URL https://arxiv.org/abs/2404.07511

[2] Angrist JD, Pischke J (2021) Mastering 'Metrics: The Path from Cause to Effect. Princeton University Press, URL https://press.princeton.edu/books/paperback/9780691152844/mastering-metrics

[3] Asseman A, Kornuta T, Patel A, et al (2022) Multi-agent dynamic pricing in a blockchain protocol using gaussian bandits. arXiv preprint URL https://arxiv.org/abs/2212.07942

[4] Barham P, Chowdhery A, Dean J, et al (2022) Pathways: Asynchronous distributed dataflow for ml. arXiv preprint URL https://arxiv.org/abs/2203.12533

[5] Becker P, Neumann G (2022) On uncertainty in deep state space models for model-based reinforcement learning. arXiv preprint URL https://arxiv.org/abs/2210.09256

[6] Bérczi K, Codazzi L, Golak J, et al (2023) Envy-free dynamic pricing schemes. arXiv preprint URL https://arxiv.org/abs/2301.01529

[7] Calvaresi D, Dicente Cid Y, Marinoni M, et al (2021) Real-time multi-agent systems: Rationality, formal model, and empirical results. Autonomous Agents and Multi-Agent Systems 35:Article 12. URL https://link.springer.com/article/10.1007/s10458-020-09492-5

[8] of Canada G (2000) Personal information protection and electronic documents act (pipeda). SC 2000, c 5 URL https://laws-lois.justice.gc.ca/eng/acts/P-8.6/

[9] Charpentier B, Senanayake R, Kochenderfer M, et al (2022) Disentangling epistemic and aleatoric uncertainty in reinforcement learning. arXiv preprint URL https://arxiv.org/abs/2206.01558

[10] Chen J, Zhang Y, Wu Y (2024) Consumer psychological and behavioral responses to ai-initiated pricing. BMC Psychology URL https://bmcpsychology.biomedcentral.com/articles/10.1186/s40359-024-01928-x

[11] Commission E (2016) General data protection regulation (eu) 2016/679. Official Journal of the European Union URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679

[12] Commission E (2021) Artificial intelligence act. COM/2021/206 final URL https://eur-lex.europa.eu/
[13] Condurache R, De Masellis R, Goranko V (2019) Dynamic multi-agent systems: Conceptual framework, automata-based modelling and verification. In: PRIMA 2019: Principles and Practice of Multi-Agent Systems. Springer, p 106–122, URL https://link.springer.com/chapter/10.1007/978-3-030-33792-6_7
[14] Diller H, Beinert M, Ivens B, et al (2020) Pricing: Prinzipien und Prozesse der betrieblichen Preispolitik. Kohlhammer, Stuttgart, Germany, URL https://doi.org/10.17433/978-3-17-034893-6
[15] Eisenhardt KM (1989) Agency theory: An assessment and review. Academy of Management Review 14(1):57–74. URL https://www.jstor.org/stable/258191
[16] Frohmann F (2023) Digitales Pricing: Strategische Preisbildung mit dem 3-Level Ansatz. Springer, Wiesbaden, Germany, URL https://link.springer.com/book/10.1007/978-3-658-37941-4
[17] G20 (2019) G20 ai principles. G20 Insights URL https://www.g20-insights.org/
[18] Gerpott TJ, Berends J (2022) Competitive pricing in online markets: An interdisciplinary framework. Journal of Business Economics URL https://link.springer.com/article/10.1057/s41272-022-00390-x
[19] Gibbons R (1992) Game Theory for Applied Economists. Princeton University Press, URL https://doi.org/10.2307/j.ctvcmxrzd
[20] Gim I, Lee S, Zhong L (2024) Asynchronous llm function calling. arXiv preprint URL https://arxiv.org/abs/2412.07017
[21] Government S (2012) Personal data protection act (pdpa). Singapore Personal Data Protection Commission URL https://www.pdpc.gov.sg/
[22] Government SA (2013) Protection of personal information act (popia). Act 4 of 2013 URL https://www.justice.gov.za/
[23] Han YJ, Nunes JC, Drèze X (2010) Signaling status with luxury goods: The role of brand prominence. Journal of Marketing 74(4):15–30. URL https://psycnet.apa.org/record/2010-13693-002
[24] He W (2022) Task-oriented communication using deep reinforcement learning. URL https://arxiv.org/abs/2208.10165
[25] Hofstede G (2001) Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations. Sage Publications, URL https://journals.sagepub.com/doi/10.1177/0312896202027001 05
[26] Holmström B (1979) Moral hazard and observability. The Bell Journal of Economics 10(1). URL https://www.jstor.org/stable/3003320
[27] Homburg C, Totzek D (2011) Preismanagement auf Business-to-Business Märkten. Gabler, Wiesbaden, Germany, URL https://link.springer.com/book/9783834915597
[28] House W (2022) Blueprint for an ai bill of rights. United States Government URL https://www.whitehouse.gov/
[29] Jin M, Wang S, Ma L, et al (2024) Time-llm: Time series forecasting by reprogramming large language models. arXiv preprint URL https://arxiv.org/pdf/2310.01728
[30] Keller KL (2021) Strategic Brand Management: Building, Measuring, and Managing Brand Equity. Pearson Education, URL https://www.amazon.com/Strategic-Brand-Management-Kevin-Keller/dp/0131888595
[31] Korbel J, Tichý T (2019) Dynamic pricing of product clusters: A multi-agent reinforcement learning approach. ResearchGate URL https://www.researchgate.net/publication/333164223_Dynamic_Pricing_of_Product_Clusters_a_Multi-Agent_Reinforcement_Learning_Approach
[32] Laffont JJ, Martimort D (2002) The Theory of Incentives: The Principal-Agent Model. Princeton University Press, URL https://press.princeton.edu/books/paperback/9780691091846/the-theory-of-incentives
[33] Legislature C (2018) California consumer privacy act (ccpa). California Civil Code § 1798100 et seq URL https://leginfo.legislature.ca.gov/
[34] Lockwood O, Si M (2022) A review of uncertainty for deep reinforcement learning. arXiv preprint URL https://arxiv.org/abs/2208.09052
[35] Lu J, Hong C, Wang R (2024) Magt-toll: A multi-agent reinforcement learning approach to dynamic traffic congestion pricing. PLOS ONE 19(11):e0313828. URL https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0313828
[36] Lu K, Zhou Q, Li R, et al (2022) Rethinking modern communication from semantic coding to semantic communication. IEEE Wireless Communications pp 1–13. URL https://doi.org/10.1109/MWC.001.2000108
[37] Lütjens B, Everett M, How JP (2019) Safe reinforcement learning with model uncertainty estimates. arXiv preprint URL https://arxiv.org/abs/1810.08700
[38] Milz M, Höfelmeier M (2023) Professionelles Pricing in turbulenten Zeiten. Haufe, Freiburg, Germany, URL https://shop.haufe.de/prod/professionelles-pricing-in-turbulenten-zeiten

[39] Monroe KB, Krishnan R (1985) The effect of price on subjective product evaluations. Perceived Quality pp 209–232. URL https://www.scirp.org/reference/referencespapers?referenceid=2296992
[40] Myerson RB (1991) Game Theory: Analysis of Conflict. Harvard University Press, URL https://doi.org/10.2307/j.ctvjsf522
[41] Nagle TT, Müller G (2018) The strategy and tactics of pricing: A guide to growing more profitably. Routledge, London, England, URL https://doi.org/10.4324/9781315185309
[42] Nash JF (1950) Equilibrium points in n-person games. Proceedings of the National Academy of Sciences 36(1):48–49. URL https://doi.org/10.1073/pnas.36.1.48
[43] Neubert M (2022) Dynamic pricing strategies: A systematic literature review. SSRN Electronic Journal URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4611545
[44] OECD (2019) Oecd ai principles. Organisation for Economic Co-operation and Development URL https://www.oecd.org/going-digital/ai/principles/
[45] People's Republic of China (2021) Personal information protection law (pipl). URL https://www.chinalawtranslate.com/en/pipl/, legal document
[46] da República P (2018) Lei geral de proteção de dados (lgpd). Law No 13,709/2018 URL https://www.planalto.gov.br/
[47] Sanchez-Cartas JM, Katsamakas E (2024) Performance of ai pricing algorithms under platform competition. Electronic Markets URL https://link.springer.com/article/10.1007/s10660-024-09821-w
[48] Simon H, Fassnacht M (2016) Preismanagement. Gabler, Wiesbaden, Germany, URL https://link.springer.com/book/10.1007/978-3-658-45056-4
[49] Steenkamp JEM (2022) Global Brand Strategy: Unlocking Branding Potential Across Countries, Cultures and Markets. Springer, URL https://link.springer.com/book/10.1057/978-1-349-94994-6
[50] Stromeyer K, Kurz W (2024) Weighted dynamic corridor price optimization. Business Management Studies 10(3). URL http://redfame.com/journal/index.php/bms/index
[51] Thompson CJ, Rindfleisch A, Arsel Z (2006) Emotional branding and the strategic value of the doppelgÃ¤nger brand image. Journal of Marketing 70(1):50–64. URL https://www.jstor.org/stable/30162072
[52] UNESCO (2021) Recommendation on the ethics of artificial intelligence. UNESCO Official Recommendations URL https://unesdoc.unesco.org/
[53] Vlastelica M, Blaes S, Pineri C, et al (2023) Mind the uncertainty: Risk-aware and actively exploring model-based reinforcement learning. arXiv preprint URL https://arxiv.org/abs/2309.05582
[54] Zhang Y, Wang Y, Liu Z, et al (2024) Enhancing multilingual capabilities of large language models through self-distillation from resource-rich languages. Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) pp 11189–11204. URL https://doi.org/10.18653/v1/2024.acl-long.603