International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

The New Rules of Enterprise Software Procurement in the AI Era for Cost Optimization

Rajan Seth

Abstract: Enterprise technology procurement has become a material financial lever as organizations navigate complex vendor ecosystems spanning SaaS, cloud infrastructure, and emerging AI workloads. Because software is economically different from physical goods, distribution costs are low relative to upfront development, and pricing is shaped as much by sales motions and channel incentives as by production cost list prices often diverge sharply from realized transaction prices. Sophisticated buyers can translate these dynamics into leverage through credible volume commitments, competitive alternatives, and timing, but many organizations still approach procurement without a systematic framework for converting supplier economics into repeatable bargaining and governance practices. As a result, cost leakage persists after contracts are signed: departments buy in silos, redundant subscriptions accumulate, license utilization goes unmonitored, and cloud consumption and commitments are under-optimized. AI adds another layer of complexity by introducing new unit economics and deployment choices that require disciplined cost models rather than ad hoc experimentation. This article offers a practical, end-to-end framework for sustainable cost optimization across the technology lifecycle: (1) understand how vendors segment and route customers, (2) decode pricing construction and where economics sit across direct and channel paths, (3) negotiate new purchases and renewals using competition, timing, and value exchanges beyond price, and (4) institutionalize post-signature governance across SaaS, cloud, and AI so savings persist without degrading capability.

Keywords: *Enterprise Procurement, Software Licensing, Cost Optimization, Cloud Computing, Negotiation Strategies, Artificial Intelligence Investments*

1. Introduction

Enterprise software procurement now requires cross-disciplinary fluency across pricing and packaging, channel economics, contract terms, finance, and operating governance. Vendor environments have grown more complex: subscription offerings coexist with legacy licensing models, add-ons proliferate, and consumption-based cloud and AI spend introduces volatility that traditional purchasing practices were not designed to manage [1]. A consistent theme connects these categories: vendor pricing is not a simple function of delivery cost. Software economics differ from physical goods; distribution costs are low relative to upfront development, and incremental delivery costs are often limited so pricing is shaped heavily by segmentation, sales incentives, competitive context, switching costs, and channel structures. As a result, list prices frequently diverge from realized transaction prices, creating meaningful but often under-captured opportunities to improve total outcomes across price, terms, and risk.

Despite this, many organizations still lack a systematic way to translate vendor economics into repeatable procurement strategy. Buying decisions are frequently made in departmental silos, which fragments vendor relationships, reduces bargaining power, and allows redundant subscriptions to accumulate. License utilization is often insufficiently monitored during contract periods, and cloud consumption and commitments are rarely governed with the rigor required to prevent spend leakage over time [2]. AI compounds the challenge by introducing new unit economics and deployment tradeoffs that require explicit cost models and governance.

This article addresses that gap by synthesizing practitioner expertise with core procurement principles into an end-to-end framework that spans vendor selection, initial negotiation, renewal strategy, and post-signature optimization. It focuses

Independent Researcher, USA

on how vendors route and price deals across market segments and channels, how quota cycles and partner ecosystems influence negotiation leverage, and how organizations can institutionalize governance across SaaS, cloud, and AI to preserve savings without degrading capability.

2. Understanding the Playing Field: Market Segments

Effective technology procurement begins with understanding how vendors segment customers and route deals. Most enterprise software providers classify buyers into tiers commonly SMB, mid-market, and enterprise based on organizational scale and transaction volume, though thresholds and benefits vary by vendor, product line, and geography [3]. Segmentation matters because it determines (a) the selling motion you face (self-serve vs sales-led), (b) whether intermediaries participate (resellers/partners), and (c) what commercial flexibility exists across price, terms, and support.

SMB / self-serve: Purchasing typically happens through online portals with standardized packaging and payment methods (often credit card). Commercial flexibility is usually limited to published promotions, annual prepay incentives, or formal programs (e.g., education/nonprofit where applicable). The convenience of self-serve often comes at the cost of limited negotiation leverage.

Mid-market: As deal sizes increase, channel partners frequently enter the transaction. Resellers and implementation partners provide sales support and services, and they earn revenue through a margin structure layered on top of vendor pricing. This intermediary layer creates additional negotiation surfaces: buyers can introduce competition across partners, negotiate partner markup, and influence how services and support are bundled into the overall proposal.

Enterprise: At the enterprise tier, direct engagement with vendor sales teams becomes common, commercial terms expand (custom clauses, bundles, multi-year structures), and discounts tend to become more negotiable especially when competition is credible and spend is strategic. Dedicated account management and alternative procurement paths (direct vs partner-led) further increase leverage for sophisticated buyers [4].

Growth and transition windows: Organizations often migrate across tiers as adoption expands from individual or team purchases to departmental rollouts and eventually enterprise-wide usage. These transition points can be valuable negotiation windows: vendors and sales teams are often incentivized to convert growing self-serve accounts into managed relationships and to lock in longer-term commitments. Procurement teams should anticipate these inflection points and use credible growth plans to shape pricing and terms before spend fully scales.

Once you know how the vendor routes your deal (self-serve, partner-led, direct), the next question is how the quoted price is constructed and where economics and profit sit across the chain. That is what creates real leverage.

Segment	Illustrative annual deal size (USD)	Common distribution channel	Illustrative commercial flexibility (not guaranteed)	Common payment method
SMB (Small/Medium Business)	< $10K–$50K	Online / Self-Serve	Limited (published promos, annual prepay, standardized programs)	Credit card
Mid-Market	~$10K–$500K	Channel partners / reseller-led (often sales-assisted)	Moderate (partner competition, bundling of services/support, negotiated terms in some cases)	PO or credit card

Enterprise	> ~$500K	Direct sales and/or strategic partners	High (multi-year structures, custom clauses, bundling, negotiated pricing)	PO, financing, custom terms

Table 1. Illustrative Segmentation Patterns Observed in Practice

3. Decoding Pricing Mechanics

Enterprise technology transactions rarely occur as simple bilateral exchanges. Depending on deal size, geography, and services scope, a purchase may involve multiple parties such as the software publisher, a distributor (in some ecosystems), and a reseller or services partner. Each participant has distinct incentives and margin expectations, and the same product can be routed through different paths [5]. Procurement excellence therefore requires visibility into the pricing architecture behind a quote, because leverage typically sits in specific layers of the value chain, not in the headline discount alone.

3.1 How Channel Economics Shape List Price and Net Price

Software vendors often set list prices above the net prices they expect to realize, in part to support discounting dynamics and to leave room for channel economics. In ecosystems that rely on distributors and resellers, this structure can accommodate the margin expectations of intermediaries while still meeting the vendor's revenue goals.

Illustrative example (for clarity): Assume a software product has a list price of $100. The vendor may sell to a distributor for approximately $50. The distributor sells to a channel partner for $55, capturing $5 in gross spread. The channel partner then sells to the end customer for $70, earning $15 on the transaction. Effective pricing strategies require alignment between price points and value delivery across all participants in the chain [5].

Seller → Buyer	Price (USD)	Notes
Published list price	100	List price shown to customers (anchor)
Vendor → Distributor	50	Illustrative vendor sell-in price
Distributor → Channel partner	55	Distributor gross spread
Channel partner → End customer	70	Partner gross spread

This value chain persists because each participant plays a distinct role and operates under different economics [6]:

Vendors typically retain the largest share of transaction value due to product ownership and licensing rights.

Distributors generally operate on thinner spreads, reflecting a high-volume, low-touch model.

Channel partners earn economics through a combination of resale margin and services (implementation, support, relationship management), with margin varying by role and relationship depth.

For procurement teams, the practical takeaway is that leverage is often distributed across the chain. Understanding which route is being used and where economics sit enables targeted negotiation rather than generic "discount asks." In partner-led deals, competitive pressure among partners can create additional flexibility, particularly when multiple partners are quoting comparable scope.

Transparency into channel economics varies by vendor. Some vendors publish partner program structures and incentives, while others maintain opacity. Procurement teams can improve negotiating positions by researching vendor partner programs, reviewing industry and partner community commentary, and directly asking partners to explain how their commercial proposal is constructed. Partners competing for the relationship will often be more candid about where

they have flexibility especially on services scope, rate cards, and bundling decisions.

3.2 Strategic Margin Compression and "Discount Theater"

In partner-led deals, partner economics are often negotiable. Create leverage by soliciting multiple comparable partner quotes (same scope, term, and support tier) and using them to negotiate partner markup, services rate cards, and bundling. Partners frequently accept lower margins rather than lose the deal especially when they expect follow-on services or expansion.

Distributor economics can also affect outcomes, but are often less visible and not always directly negotiable. For large transactions, consider alternative routing (where feasible) to reduce layers, and insist on quote transparency by separating product, services, and support so pass-through spreads are easier to pressure.

At the vendor level, list prices are often anchored, not the economic truth. This can create "discount theater," where headline discounts look large while the real outcome is determined by net effective price, renewal mechanics, support/services, and commitments. Skilled procurement shifts the discussion from "% off list" to absolute price targets, multi-year economics, and protections that prevent re-basing or unfavorable future increases. Preparation matters: benchmarks, competitive evaluations, and execution-ready alternatives strengthen credibility. Timing can help, but only when you can actually switch.

4. Tactics for Negotiating New Deals (and Renewal Outcomes)

Approaching new technology purchases requires strategic preparation. Leverage rarely comes from price discussion alone; it comes from the full outcome across price, terms, risk allocation, and operational flexibility. The goal is to structure an agreement that is economically efficient and resilient over time, rather than winning a one-time discount that erodes through renewals, add-ons, or restrictive commitments.

4.1 Strategic Positioning

Effective negotiation starts with positioning that reduces vendor uncertainty and increases your strategic value. Smaller organizations should avoid anchoring negotiations solely on today's deal size. Instead, present a credible growth narrative tied to the vendor's usage drivers (users, seats, workloads, business units), including adoption milestones and the conditions under which spend scales. Vendors frequently evaluate accounts on lifetime value potential, not just current transaction value, and interest-based negotiation principles reinforce focusing on underlying incentives rather than fixed positions [7].

Larger organizations can emphasize enterprise-wide adoption plans, standardized rollouts, and executive sponsorship. Multi-level engagement (procurement, IT, security, finance, and business leadership) signals commitment and reduces the perceived risk of churn often improving the vendor's willingness to offer better commercial terms and governance concessions.

Regardless of size, competitive pressure is essential. Parallel evaluations and benchmark quotes create pricing floors and expose alternative paths if negotiation stalls. Avoid arbitrary thresholds (e.g., "we'll never pay more than X% premium"). Instead:

- Define a walk-away price based on quantified switching costs, capability differences, and risk (including implementation and operational overhead), and
- Use competitive quotes to enforce it credibly and consistently.

4.2 Value Exchange Beyond Price

Price is only one negotiation dimension. Vendors also value mechanisms that de-risk future revenue and accelerate market adoption. These can be traded for improved outcomes on pricing, services, and terms when governed properly [8].

Common value exchanges include:

- **Reference potential** (customer references, case studies, quotes, conference participation),
- **Co-marketing** (joint announcements, webinars),
- **Roadmap feedback** and structured product input, and
- **Beta/early access participation** (when it delivers real value and is resourced appropriately).

These should not be open-ended. Treat them like commercial assets: time-bound, contingent on delivery outcomes, and scoped to specific commitments (e.g., "one case study after measurable success," "reference calls limited per quarter"). Buyer–seller interactions increasingly involve collaborative value creation beyond transactional exchanges [8], but that collaboration

should be explicit, governed, and reversible if outcomes are not met.

Operationally, negotiate for concessions that improve time-to-value and reduce risk, such as:

- clearer implementation support and escalation paths,
- solution architecture assistance for initial rollout,
- service credits or structured onboarding, and
- support terms that match your criticality needs (not "premium by default").

4.3 Timing and Sales Incentives

Quota cycles create urgency. Negotiation leverage often increases near fiscal period ends (quarter and year), but only when you have credible alternatives and internal readiness to execute (legal review, stakeholder alignment, implementation plan, budget approval). Use timing to improve terms not to force rushed decisions that lock in long-term switching costs.

When appropriate, timing can also expand flexibility on:

- multi-year structures and ramp schedules,
- payment terms (e.g., upfront payment or milestone-based payment), and
- non-price terms that reduce downstream risk (renewal protections, caps on increases, scope clarity, and termination/portability clauses).

The objective is not "maximum discount." It is a durable agreement: competitive net effective economics, protective renewal mechanics, and operational terms that preserve flexibility as the relationship scales. Sustainable SaaS savings generally come from three repeatable levers (1) consolidation and standardization, (2) utilization governance, and (3) controlled deprovisioning ("scream tests") enabled by the right contract terms and internal operating controls. Procurement should run SaaS cost management as a lifecycle process, not a one-time pricing exercise.

5. Sustainable SaaS Cost Management: Consolidation, Utilization, and Deprovisioning

5.1 Consolidation and Sandardization

Fragmented departmental purchasing weakens bargaining power and creates duplicate spend. Functions often buy similar tools through separate budgets and accounts, while "shadow IT" purchases outside procurement visibility can duplicate enterprise entitlements. Consolidation into a single enterprise agreement can improve leverage and administration if paired with competitive tension and disciplined right-sizing to avoid lock-in or over-commitment [9]. Any claims about governance benefits should be tied to defined oversight practices and measurable outcomes, supported by appropriate research or benchmark data.

How to execute

- **Build visibility:** audit subscriptions across business units, map accounts to owners, and quantify redundancies.
- **Rationalize by capability, not name:** go beyond exact duplicates to overlapping tools (e.g., multiple work management platforms).
- **Standardize with exceptions:** converge on fewer tools while preserving narrowly scoped exceptions for legitimate needs.

This reduces license waste, simplifies training and administration, and improves cross-functional collaboration.

5.2 Utilization Governance ("shelfware" control)

User-based pricing models create waste when licenses are purchased but not actively used. Shelfware accumulates quietly and becomes a recurring leakage. A practical control is a recurring usage cadence (often quarterly) that compares paid entitlements against measured utilization and role-based need, while accounting for exceptions (seasonal users, service accounts, and workflows where login counts don't reflect value). As delivery models evolve, procurement practices must evolve with them [10].

Make usage insights actionable:

- **Contract for flexibility:** include adjustment windows or right-sizing mechanisms during the term.
- **Define "active use" precisely:** use tool-specific metrics and agreed measurement rules.
- **Secure data access:** ensure administrative access to entitlement and utilization data.

Contractual rights create the option to save; operational controls, role-based provisioning, approvals, automation, and owner accountability convert that option into realized savings.

5.3 Controlled Deprovisioning ("scream tests")

For legacy subscriptions with unclear value, controlled deprovisioning clarifies true dependency. A scream test temporarily disables access (or restricts provisioning) to determine

whether a tool is essential, redundant, or unused. Lack of immediate complaints is a signal, not proof: dependencies can be delayed, infrequent, or invisible until a reporting cycle or high-stakes moment.

Run scream tests safely:

- **Rollback plan:** fast re-enable path, clear decision criteria, named owners.
- **Monitoring window:** long enough to capture delayed workflows and periodic processes.
- **Risk tiers and change management:** avoid mission-critical systems unless controls, alignment, and backup options are in place.
- **Document outcomes:** what broke, who escalated, time-to-impact, and what value was demonstrated.

These results improve renewal decisions, refine consolidation priorities, and strengthen future negotiation posture with evidence rather than anecdotes.

6. Saving on Public Cloud Services

Cloud infrastructure is a significant and growing component of enterprise technology spend. Cloud shifts many costs from capital expenditure to operating expenditure, which increases flexibility but also introduces new cost optimization challenges [11]. Effective cloud savings require two parallel tracks: (1) commercial structure (credits, commitments, and negotiated terms) and (2) consumption discipline (engineering and FinOps practices). Commercial negotiation can amplify good consumption hygiene, but it rarely substitutes for it.

6.1 Credits and Commitments

Startups and early-stage organizations should pursue official cloud startup programs where eligible and plan a clear transition from credits to sustainable pricing. Credits can meaningfully reduce early burn, but they are time-bound and should not be treated as a long-term cost model. Venture-backed companies may also unlock incremental leverage through investor relationships or portfolio programs that provide preferred commercial terms.

As credits phase out, organizations typically shift to commitment-based discount structures (provider-specific), such as committed-use instruments and enterprise discount programs. These can reduce effective rates, but only when aligned with stable, forecastable usage. Commitments should be continuously re-validated against workload reality not treated as "set and forget." In some environments, commitment burn-down can also occur through marketplace purchases of third-party software, which can help meet commercial thresholds while diversifying procurement channels.

6.2 Consumption Optimization

The fastest and most durable savings often come from engineering and FinOps practices: rightsizing, scheduling non-production shutdowns, enforcing storage lifecycle policies, and making architecture choices that minimize data movement and unnecessary managed-service sprawl [12]. Benchmarking and modeling usage patterns (not just list prices) helps identify where cost differences actually emerge across services and providers.

On the commercial side, consider whether fixed pricing constructs (where available) provide better predictability than percentage-based discounts, especially when list prices change over time. Re-negotiation opportunities generally improve as total spend grows, but the best outcomes come from pairing commercial structures with measurable consumption controls.

Pay particular attention to "sticky" services with high switching costs (e.g., deeply integrated databases or security tooling): negotiate protections early, before architectural dependencies make reversal difficult.

Service category	Expected discount / effective rate reduction (reference, illustrative)	Primary negotiation / optimization levers
Compute	**30–60%**	Commitments sized to stable baselines; rightsizing; non-prod shutdown schedules
Storage	**25–70%**	Lifecycle tiering; retention policies; data placement; minimizing unnecessary replication

Managed databases / data services	**20–45%**	Lock terms early; cap uplifts; negotiate migration/support; avoid over-provisioning
Security / identity / observability	**20–35%**	Consolidate overlapping tools; enterprise-wide standards; bundle rationalization

Table 3. Rate Card Discounting by Service Category

7. Optimizing AI Investments

Artificial intelligence spending is the newest frontier of enterprise technology cost management. AI introduces new unit economics token-based pricing, GPU-hour consumption, and end-to-end system costs that can create material budget risk without disciplined governance [13]. A procurement-ready AI strategy therefore separates two phases with different cost logic: experimentation (high flexibility; higher unit price is tolerable) and production (unit economics, reliability, and repeatability dominate). Avoiding "bill shock" requires thoughtful deployment model selection and ongoing optimization.

7.1 Deployment model selection

AI deployments generally fall into two approaches, each with distinct cost drivers and organizational requirements. Across both, it is important to distinguish open models (available via API or self-hosted) from closed models (available via API only, with no self-hosted option).

Model provider (API) approach. This functions like a rental model: access is delivered through APIs and costs are driven primarily by token consumption. API-based models are often the fastest path for prototyping, proof-of-concept work, and early product validation because they minimize infrastructure overhead and accelerate time-to-value. This approach is well suited when flexibility and speed matter more than per-unit cost.

Self-hosted approach (cloud or on-premises). This functions more like an ownership model: costs are driven by GPUs, infrastructure, and operational overhead. Self-hosting is an option for open models, and can become attractive when utilization is high, requirements are specialized, latency/throughput constraints matter, or data constraints are strict but only if operational capability exists (in-house or via managed partners). The economics depend on sustained utilization: idle capacity can quickly erase expected savings [14].

Cloud-based self-hosting can improve flexibility through reserved capacity and other discounted compute options, while on-premises deployment can provide maximum control when the organization can reliably operate the full stack. Closed models do not have a self-hosted option, so organizations that require full deployment control must evaluate open alternatives.

The key procurement decision is not "API vs self-hosted" in the abstract it is whether the organization has the usage stability, governance, and operational maturity to convert infrastructure ownership into lower unit cost.

7.2 Optimization strategies

Once a use case moves from experimentation to production, optimization becomes the primary lever for cost control. Model selection and distillation can reduce ongoing operating cost, but outcomes vary widely by workload, quality targets, latency/SLA requirements, utilization, and operational overhead [14]. The correct claim is conditional: open models and self-hosting can reduce unit cost, but only when the full system cost is included compute, engineering labor, observability, security, reliability, and ongoing model operations.

A practical pattern is to validate capability early using high-performing models (often via APIs), then migrate stable workloads to smaller, more efficient models or self-hosted deployments once quality requirements are understood and usage is predictable. Treating inference compute as a procurement surface where portability exists can introduce competitive dynamics across providers, but the economic win depends on end-to-end system design and governance, not compute price alone.

Deployment Model	Cost Driver	Best Suited For	Key Requirement
Model Provider (API)	Tokens	Experimentation, Rapid Prototyping, Standard Use Cases	Operational Expenditure Flexibility
Self-Hosted (Cloud)	GPUs and Infrastructure	High Data Privacy, Massive Scale, Custom Capabilities	CapEx Readiness and High Utilization
Self-Hosted (On-Premises)	GPUs and Infrastructure	Maximum Data Control, Dedicated Resources	Engineering Talent for Stack Management

Table 4. Cost Drivers and Strategic Considerations [11,12]

Conclusion

Enterprise technology procurement is no longer a back-office buying function; it is a repeatable discipline that directly shapes unit economics, risk exposure, and operating flexibility across SaaS, cloud, and AI. The framework in this article highlights a consistent pattern: meaningful savings come from understanding how vendors segment and route deals, decoding the pricing architecture behind quotes, and negotiating outcomes that optimize net effective cost and contractual protections, not headline discounts. But durable results require post-signature governance: consolidating and standardizing overlapping tools, controlling shelfware through utilization reviews and scale-down rights, and applying disciplined FinOps practices to cloud consumption and commitments. As AI spend expands, organizations must also separate experimentation from production and evaluate open versus closed models and API versus self-hosting through full-system cost, reliability, and operational readiness. Teams that institutionalize these practices will reduce cost leakage, improve predictability, and sustain capability as technology portfolios and their underlying economics continue to evolve.

References

[1]Sonja Lehmann and Peter Buxmann, "Pricing Strategies of Software Vendors," Springer Nature Link, [Online]. Available: https://link.springer.com/article/10.1007/s12599-009-0075-y

[2] Kalle Lyytinen and Gregory M Rose, "Disruptive information system innovation: The case of internet computing," Information Systems Journal, 2003. [Online]. Available: https://www.researchgate.net/publication/220356570_Disruptive_information_system_innovation_The_case_of_internet_computing

[3] Chor-Beng Anthony Liew, "Strategic integration of knowledge management and customer relationship management," JOURNAL OF KNOWLEDGE MANAGEMENT, 2008. [Online]. Available: https://www.researchgate.net/profile/Anthony-Liew-2/publication/220363412_Strategic_integration_of_knowledge_management_and_customer_relationship_management/links/09e414fad45b69cf28000000/Strategic-integration-of-knowledge-management-and-customer-relationship-management.pdf

[4] Roland T. Rust et al., "Getting Return on Quality: Revenue Expansion, Cost Reduction, or Both?," Journal of Marketing, 2002. [Online]. Available: https://www.researchgate.net/profile/Christine-Moorman/publication/237326541_Getting_Return_on_Quality_Revenue_Expansion_Cost_Reduction_or_Both/links/563120df08ae13bc6c355da2/Getting-Return-on-Quality-Revenue-Expansion-Cost-Reduction-or-Both.pdf

[5] Andreas Hinterhuber and Stephan Liozu, "Is Innovation in Pricing Your Next Source of Competitive Advantage?," Hinterhurber & Partners, [Online]. Available: https://hinterhuber.com/publications/is-innovation-in-pricing-your-next-source-of-competitive-advantage/ [6] Kangzhou Wang et al., "Cooperation or Competition? Channel Choice for a Remanufacturing Fashion Supply Chain with Government Subsidy," ResearchGate, 2014. [Online]. Available: https://www.researchgate.net/publication/277673683_Cooperation_or_Competition_Channel_Choice_for_a_Remanufacturing_Fashion_Supply_Chain_with_Government_Subsidy

[7] ROGER FISHER et al., "Getting to YES NEGOTIATING AGREEMENT WITHOUT GIVING IN," [Online]. Available: https://files.addictbooks.com/wp-content/uploads/2024/03/Getting-To-Yes-Negotiating-Agreement-Without-Giving-In.pdf

[8] Stephen L. Vargo and Robert Lusch, "Service-Dominant Logic" Continuing the Evolution,"

Journal of the Academy of Marketing Science, 2008. [Online]. Available: https://www.researchgate.net/publication/272566616_Service-Dominant_Logic_Continuing_the_Evolution

[9] Roberto Verdecchia et al., "Building and evaluating a theory of architectural technical debt in software-intensive systems," Journal of Systems and Software, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0164121221000224

[10] Michael Seifert et al., "Hybrid Clouds Arising from Soware as a Service Adoption: Challenges, Solutions, and Future Research Directions," ACM Computing Surveys, 2023. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3570156

[11] Ali Khajeh-Hosseini et al., "The Cloud Adoption Toolkit: Supporting Cloud Adoption Decisions in the Enterprise," arXiv, [Online]. Available: https://arxiv.org/pdf/1008.1900

[12] Cheng-Kai Hsu et al., "Time is money: Effects of time-sensitive incentive on risky driving among on-demand food-delivery motorcyclists in Taiwan," ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925753524001395

[13] Senthil Kumar Jagatheesaperumal et al., "The Duo of Artificial Intelligence and Big Data for Industry 4.0: Review of Applications, Techniques, Challenges, and Future Research Directions," arXiv, 2021. [Online]. Available: https://arxiv.org/pdf/2104.02425

[14] Tom B. Brown et al., "Language Models are Few-Shot Learners," 34th Conference on Neural Information Processing Systems, 2020. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf