PHANTOM: Pricing Heuristics Against Non-human Transaction Orchestration Mechanisms

Daniel Rösel

Bachelor of Computer Science & Artificial Intelligence

Supervised by:

Alberto Martín Izquierdo

IE University, Madrid, Spain

January 26, 2026

Abstract

With accelerated growth of Lager Language Model agents in e-commerce a novel adversarial dynamic to digital markets emerges. This paper address the vulnerability of dynamic pricing systems to AI intermediaries that decouple the information gather stages from the transaction execution. By conducing reconnaissance isolates sessions, agents circumvent the "Cost of Information" (COI) defined as the accumulated price premium typically thought demand expression estimators. We formally define this phenomenon and derive the Cost of Information Theorem, proving that as the saturation of independent, utility-maximizing agents increases, the platform's ability to sustain a COI converges to zero, rendering standard dynamic pricing mechanisms incentive-incompatible. To respond to this threat we propose a defensive framework which integrates behavioral economics with Adversarially Distributionally Robust Optimization (DRO). We introduce a custom e-commerce research platform built on hybrid Kappa-Lambda architecture, designed to capture and simulate high-fidelity controlled interaction trajectories. We further demonstrate through modeling that human and agent behaviors exhibit distinct transition probability kernels, enabling the construction of discriminative models based on Kullback-Leibler divergence. These behavioral signals serve as inputs for a Distributionally Robust Reinforcement Learning (DR-RL) agent. We formulate the pricing problem as a Stackelberg game where the learner optimizes against an ambiguity set of demand distributions defined by the Wasserstein distance. This approach allows the pricing policy to remain robust against non-stationary contamination without overfitting to deterministic demand curves. The research validates a mechanism for preserving margin integrity and market equilibrium in an agent-mediated economy, while minimizing degradation to the legitimate human user experience (UX).

Keywords: Dynamic Pricing, LLM Agents, Adversarial Machine Learning, E-commerce, Behavioral Detection, Reinforcement Learning

Acknowledgments: Eugene Bykovets, PhD - ETH for helping with problem formulation. This research was supported by the TPU Research Cloud program.

1 Introduction

In this paper we present an exploration and defense against the presence of new commercial entities in digitally powered platforms, preserving market equilibrium in the age of AI. This research establishes the following contributions: definition and formalization of non-human transactors in e-commerce platforms, development of a testing-ground for capturing the behavioral essence of these transactors across a large variety of digital systems, construction of a discriminative model (to prove separability) as a strong learner for downstream mitigation of contamination by non-human entities, translation of such learned separability into existing dynamic pricing machine learning loops, and finally establishment of a high-level KPI-affecting causal effect and cost-saving framework for the future of internet commerce in the presence of such non-human learners.

This research effort touches a large variety of domains, spanning behavioral economics for understanding the rationality of behavior as theorized by the concept of homo economicus, agent-based modeling to translate our learned separability into disjoint dynamic pricing systems, reinforcement learning which serves as the SOTA for price-learners, and dynamic pricing and market equilibrium theory to understand the risks of possible supra-competitive pricing phenomena in cases of adversarial pricing systems driving the market out of equilibrium.

1.1 Motivation and Market Context

The current innovation boom in generative artificial intelligence and its applications to knowledge-based work tasks has brought many competing technologies for browser-use automation, with benchmarks and evaluations (Xia et al., 2025) motivating the development of capabilities focused on commercial research, understanding, and transaction execution (Xie et al., 2024). The “AI Agent” market is forecasted to grow from around USD 5-8 billion in 2025 to USD 42-52 billion by 2030. This surge reflects adoption in e-commerce, customer service, and enterprise automation, where agents handle interactions previously done by humans, raising the question of how these systems should be designed for future robustness as well as how to maintain a competitive edge in the analytical components of e-commerce platforms (MarkNtel Advisors, 2025).

The key stakeholders affected by the threat of increasing agent-driven traffic include online businesses and platform operators (especially in bot-heavy sectors like retail, travel, and financial services), their security, fraud, and engineering teams, end users whose accounts and data are exposed and whose experience degrades, regulators and legal stakeholders responding to breaches and fraud, and the attackers or bot operators driving the automation (Imperva, 2025).

The industry has already seen legal action in cases like Amazon against Perplexity (Ghaffary and Day, 2025), stemming from the difficulty of identifying traffic from hybrid systems like the Commet browser. This paper explores such systems to better understand what the interaction data looks like and what it means for dynamic pricing and recommendation systems downstream. This observed impact indicates a need for prevention of secondary negative effects on the "legacy" systems which power modern revenue sources for many companies. Dynamic pricing algorithms rely on directly translating demand features q to new price assignments $\hat{p}$ across a catalogue of products of size N. This opens opportunities to design a *tabula rasa* of digital market mechanisms that will shape the future of commerce in the age of artificial intelligence.

1.2 Solution Space Overview

Dynamic pricing systems, as presented by Mueller et al. (2019), often deal with sparse low-rank data of demand signals which, combined with contamination from agents, creates complex interactions that impact pricing. To further complicate the problem, certain commercial settings such as the one presented by Amjad and Shah (2017) must address the true demand of products under censored observations. This provides a formulation for handling demand in our case with multiple kinds of commercial mediators: $\hat{q} \leftarrow q_A + q_H$ where q_A represents the distribution of demand generated by agentic mediators and q_H represents that of true human demand, these are two distinct populations with divergent objective functions.

We formally define interaction data as coming from some actor which can either be an agent (A) or human (H). For purposes of this research, an agent is an algorithmic loop with the ability

Algorithm 1: AI Agent's Interaction Loop

Input : Goal G, Platform URL u, LLM $\mathcal{M}$
Output : Task completion result r
Initialize browser instance $\mathcal{B}$ with connection to u
Construct prompt $\pi \leftarrow \textsc{BuildPrompt}(G, u)$
done $\leftarrow$ False
while $\neg$*done* **do**
  Observe current page state s_t from $\mathcal{B}$
  Query $\mathcal{M}$ with (π, s_t) to determine next action $a_t \in$ {click, scroll, fill, navigate}
  Execute a_t on $\mathcal{B}$ to transition to state s_{t+1}
  done $\leftarrow \mathcal{M}.\textsc{JudgeCompletion}(G, s_{t+1})$
Extract final result r from terminal state
return r

to access a web platform and perform actions such as clicks, scrolls, and input field fills. The loop terminates when the internal large language model judges the provided task definition as complete. A detailed breakdown can be found in algorithm 1.

1.3 Research Questions

This work addresses three core research questions:

RQ1 *Separability*: Can agent and human sessions be reliably distinguished from behavioral interaction signals alone, without relying on network-level or device fingerprinting?

RQ2 *Theoretical Impact*: What is the formal relationship between agent contamination levels and the erosion of pricing power in dynamic pricing systems?

RQ3 *Robust Mitigation*: How can pricing policies be constructed to maintain margin integrity under unknown and non-stationary levels of agent contamination?

The previously described goal of separability allows us to formulate a task which entails taking raw interaction data for either actor and creating a composite demand estimate $\hat{q}$. We propose a robust optimization objective defined in our methodology, transforming the pricing problem into a form of Distributionally Robust Optimization (Kuhn et al., 2025) where the learner must guard against adversarial contamination in observed demand distributors. In this setting we must learn to make decision that perform under the assumption of not having a single estimated probability distribution but under an ambiguity set of any distribution, of which we have limited information. In our case as stated is a mixture of distributions with a parameter which is unknown and non-stationary.

2 Literature Review

To better understand all wedges of the current works, we must start by exploring the nature of agents, agentic computer use and web automation, complementing that with economic reasoning and strategic interaction. The final surface to cover, leads us to data-driven dynamic pricing under uncertainty. The key technical risk is not “agents buying things” per se, but agents shaping the behavioral and demand signals that downstream pricing systems consume and depend on. This latter case of agents shopping is currently pending legal action in the case of *Amazon.com Services LLC v. Perplexity AI, Inc* (2026) which is currently being treated as a violation of the Computer Fraud and Abuse Act. The introduction of these mediating actor entities into economic systems, is further creating a threat of false-name bidding (Yokoo et al., 2004), which prior research has explored in a trading context. Other research on pseudonyms in dynamic systems, demonstrate whitewashing in AI agents which can ignore defensive mechanisms by re-entry with different identities (Feldman et al., 2004). Dynamic pricing assumes demand proxies are behaviorally meaningful, while bot detection aims at security and access control. The missing bridge is a principled framework for separating non-human reconnaissance from genuine human demand expression and integrating that separation into pricing heuristics without degrading legitimate user experience (in our research tracked by the user-experience index). This gap, is what our contribution aims to address, particularly for the aforementioned stakeholder groups.

2.1 Agent Taxonomy and Definitions

An agent in the context of artificial intelligence is generally defined by anything that can reason and act upon observations of its environments (collected through some sensory inputs) and carry out actions through effectors. Moreover, a rational agent is an entity that is capable of perceiving the world around them and taking actions to advance specified goals. This definition by Russell and Norvig (2021) is further developed in an economic context by Parkes and Wellman (2015), suggesting AI research attempts to construct a synthetic *homo economicus*, which may also be termed *machina economicus*. A specific class or taxon of this *machina economicus*, the Large Language Model (LLM) agent, is defined as an autonomous system capable of achieving goals and adapting post-training, often without needing explicit code or fundamental model changes (Xia et al., 2025).

We must however acknowledge the current SOTA as presented by OSWORLD simulations by Xie et al. (2024) have demonstrated that multi-modal tasks across desktop and web interaction modes, have a top-performing score of only 12.24% success, whereas humans have a higher 72% success rate; this is linked to the lack of grounding of these agents and their inability of handling unexpected errors. This weakness matters for this research because it clarifies the near-term threat model: practical exploitation does not require a fully competent “computer assistant”, only enough automation to perform high-volume reconnaissance actions (search/filter/open product pages, probe availability/price boundaries) that can contaminate behavioral signals. With the expected growth of these capabilities, this threat only becomes more perilous to revenue management systems.

We model an agent session as producing some events with lower in-session conversion levels relative to humans, this we state in our assumption that $P(\text{purchase}|A) < P(\text{purchase}|H)$ but with a potentially higher volatility in $\hat{q}$, which we observe through the look-to-book metrics in our simulation.

2.2 Economic Agents: From Homo Economicus to Machina Economicus

Existing behavioral economic models tend to be criticized for the assumption of rational behavior, as is embodied in the term of homo economicus. The definition of a machina economicus by Parkes and Wellman (2015) is quite appropriate for our case, particularly because these assumptions of rationality have been argued to be a very adequate reference for AI research by Varian (1995) due to its expected utility maximizing nature. For modeling this behavior, the trajectories of these agents can be formally defined to be partially observable Markov decision processes (Xie et al., 2024). Agents are however not to be confused with web-bots which have previously been known as automated software applications or scrapers which are set with a purpose of carrying out specific tasks on the internet, without a higher level of internal judgement (Imperva, 2025). In our research, we refer to this actor simply as an Agent belonging to the distribution A.

This economic framing also helps separate two related but distinct phenomena of agents as buyers (changing market demand composition), and agents as information gatherers (changing the observed interactions used by pricing/recommendation systems). The thesis focuses on the second, where information acquisition strategically precedes purchase execution. We do not however dismiss the proposed expectation that existing economic systems serving humans, will not be populated by AIs across multiple channels and with various possibly misaligned goals as stated by Parkes and Wellman (2015).

A HAP (HTTP Agent Profile) protocol has been developed as an internet draft by Dhir (2025) in an effort to separate agentic and human internet traffic, however the majority adoption by both the sellers and agent providers would be required for the implementation of such a solution.

2.3 Problem Evidence and Market Impact

The statistical issue of contamination in dynamic pricing systems that observe demand features as a means to update prices has been documented in various previous contexts. The airline industry (which has accounted for 24% of observed disruptions) has seen malicious activity with a measureable impact on skewing key performance indicators by behavior visible in the look-to-book metrics. Excessive reconnaissance traffic inflates search volume without corresponding completed bookings, thereby skewing demand forecasts and disrupting dynamic pricing models. Demand proxies have also been observed to cause significant threat to inventory management by creating artificial scarcity that distorts the demand-supply relationships in the enterprise model. Censored demand as shown by Amjad and Shah (2017) can also be observed in low-bias demand under-estimation caused by a distortion effect coming from non-human traffic data (Imperva, 2025).

When dynamic pricing algorithms operate on highly contaminated or noisy data, the risk grows significantly in creating inaccurate price inferences. The emergent mitigation driven by un-informed reward and regret signals might lead to price suppression for sales continuity which results in harming margins and resulting in a revenue loss. System that poorly fit undesired behavior might result in price gouging, which calls for strong guardrails while preserving targeted business strategy (Mullapudi, 2025).

2.4 Theoretical Foundations: Economic Parallels

Early hints of exploration of prices in a standard English auction explored by Varian (1995) which hints at exploration of prices in a sequential manner, which leads to a marginally different cost to the bidder than the reservation price of the seller. This is a setting in which there is no cost incured by the buyer for their actions or exploring prices in the market. They propose that any agent responsable for the pricing of a good must be imune to dynamic strategies which might extract private information from a market. A key take-away which relates to the Vickery auction mechanism (also called a *direct mechanism*) suggests that not only would defenses against such exploitation be necessary, but the construction of a mechanism in which revelation of the true willingness to pay is the dominant strategy for commerce.

Like in classical revenue-maximizing auctions (Roughgarden, 2013) we assume that the human actor in our system has a private valuation v which we formally draw from intrinsically defined distributions. The important note here is that the agent proxy does not have a mechanism to convey this private information into the demand data which directly impacts the pricing systems.

The key component of this mediation between agents and commercial platforms lays in the transaction costs related to information gathering and negotiation. As proposed by Shahidi et al. (2025) these costs are bound to collapse towards zero (which we demonstrate mathematically), calling for a re-evaluation of the boundaries between firms and markets. As argued by Coase (1937), the market participation and time associated with that participation, is critical part of the Coasean transaction cost logic which includes the discovery or relevant pricing within a given market. This process of price discovery without the presence of AI Agents can be time consuming and resource intensive. To build on top of this work we provide a proof of optimal conditions theorised by Coaes as an extension to AI-mediated markets.

2.5 Landscape of Existing Work

Explorations of the algorithmic collusion by LLMs (Fish et al., 2025) has demonstrated a cross-model tendency of market division with a strong sensitivity to instructions provided in the "system prompt". If a dynamic pricing algorithm which is trained to respond to market signals learns to coordinate with competitor agents (or become manipulated by those agents), the market equilibrium is under threat of destabilization. This is particularly true for Q-learning pricing learners as demonstrated by Calvano et al. (2018).

Our effort to combat contamination stems from research by Hardt et al. (2015) on strategic classification, in conjunction with Liu et al. (2024) who demonstrate a linear regret if contamination is ignored. The strategic classification adversarial effect comes from an effort to manipulate some representative features used in a learning pipeline, which can result in lower prices on loans or lower prices from dynamic pricing algorithms.

To bridge the gap between detection and robust pricing, we look at work in Distributionally Robust Optimization (DRO). As defined by Kuhn et al. (2024), DRO provides a framework for decision-making under ambiguity, where the true data distribution is unknown but lies within a "Wasserstein ball" of a target distribution. In our context, the "ambiguity set" represents the uncertainty introduced by agentic reconnaissance. By optimizing for the worst-case distribution within this set, pricing mechanisms can become resilient to the distributional shifts such as the ones caused by non-human actors, effectively robustifying the revenue function against the contamination described in our problem statement.

In order to create an environment in which prices can be tested against a demand estimate generated by some behavioral model, we take inspiration from the architecture proposed by Ie et al. (2019) in the RecSim platform built for recommendation systems. By modeling the distinct user behavior as POMDPs we can generate faithful interactions which allow us to generalize, past the constraint which is also present in recommendation systems, of rarely having enough experience with individual actor's interactions for good recommendations without generalization. The key inspiration comes from the user choice modeling which we translate to a user transition model for each distinct actor type (agent or human). We further consider the possibility of modeling our quantitative research platform using dynamic Bayesian networks for the sake of tractability within the system. The contribution or RecSim enables researchers to better understand learning algorithms in fixed environments, a gap we identify as needing to be bridged within the space of dynamic pricing.

We also acknowledge the difficulty in similarly affected fields such as authorship, where Ganie (2025) demonstrate the theoretical limits of the distributional divergence between text authored by a human or large language model. Their approach of computing the divergence between two distributions demonstrates purely theoretically that no classifier can outperform random guessing on their particular task. This is yet another factor to take into consideration when exploring the potential mitigation strategies.

The setting of our work is quite complex and covers a wide range of topics, each with its own set of issues that further complicate the task at hand. There is however promise in the field of reinforcement learning and adversarial robustness to combat these problems. We can summarize the characteristics learned from the review of our environment as: (i) non-stationary demand with temporal noise ϵ_t (ii) contaminated behavioral signals from mixed human-agent traffic with unknown mixing ratio α (iii) partial observability where only demand proxies $\hat{q}$ are available, not true demand $d(\cdot)$ (iv) strategic actors capable of feature manipulation to influence pricing outcomes (v) information asymmetry with private valuations v drawn from unknown distributions (vi) session-based interactions modeled as POMDPs with trajectories τ_s (vii) low conversion probability for agents: $P(\text{purchase} \mid A) < P(\text{purchase} \mid H)$ (viii) distributional uncertainty requiring robust optimization within Wasserstein ambiguity sets (ix) potential for adversarial exploitation through false-name bidding and identity whitewashing.

References

Amazon.com Services LLC v. Perplexity AI, Inc (Jan. 2026). en. No. 3:25-cv-09514-MMC.

Amjad, Muhammad J. and Devavrat Shah (Dec. 2017). "Censored Demand Estimation in Retail". In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1.2. Publisher: Association for Computing Machinery (ACM), pp. 1–28. DOI: 10.1145/3154489. URL: https://par.nsf.gov/servlets/purl/10066022 (visited on 11/12/2025).

Calvano, Emilio et al. (2018). "Artificial Intelligence, Algorithmic Pricing and Collusion". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3304991. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3304991.

Coase, R. H. (1937). "The Nature of the Firm". en. In: *Economica* 4.16, pp. 386–405. ISSN: 1468-0335. DOI: 10.1111/j.1468-0335.1937.tb00002.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0335.1937.tb00002.x (visited on 01/20/2026).

Dhir, Sanat (Nov. 2025). *HTTP Agent Profile (HAP): Authenticated and Monetized Agent Traffic on the Web*. Internet Draft draft-dhir-http-agent-profile-00. Num Pages: 13. Internet Engineering Task Force. URL: https://datatracker.ietf.org/doc/draft-dhir-http-agent-profile (visited on 01/20/2026).

Feldman, Michal et al. (2004). "Free-riding and whitewashing in peer-to-peer systems". In: *Proceedings of the ACM SIGCOMM 2004 Workshops*. Association for Computing Machinery, pp. 228–235. ISBN: 1-58113-942-X. DOI: 10.1145/1016527.1016539.

Fish, Sara, Yannai A. Gonczarowski, and Ran I. Shorrer (Sept. 2025). *Algorithmic Collusion by Large Language Models*. en. arXiv:2404.00806 [econ]. DOI: 10.48550/arXiv.2404.00806. URL: http://arxiv.org/abs/2404.00806 (visited on 01/20/2026).

Ganie, Aadil Gani (Sept. 2025). *Uncertainty in Authorship: Why Perfect AI Detection Is Mathematically Impossible*. en. arXiv:2509.11915 [cs]. DOI: 10.48550/arXiv.2509.11915. URL: http://arxiv.org/abs/2509.11915 (visited on 01/05/2026).

Ghaffary, Shirin and Matt Day (Nov. 2025). *Amazon Sues to Stop Perplexity From Using AI Tool to Buy Stuff*. URL: https://www.bloomberg.com/news/articles/2025-11-04/amazon-demands-perplexity-stop-ai-agent-from-making-purchases.

Hardt, Moritz et al. (Nov. 2015). *Strategic Classification*. en. arXiv:1506.06980 [cs]. DOI: 10.48550/arXiv.1506.06980. URL: http://arxiv.org/abs/1506.06980 (visited on 01/20/2026).

Ie, Eugene et al. (Sept. 2019). *RecSim: A Configurable Simulation Platform for Recommender Systems*. arXiv:1909.04847 [cs]. DOI: 10.48550/arXiv.1909.04847. URL: http://arxiv.org/abs/1909.04847 (visited on 12/29/2025).

Imperva (2025). *The Rapid Rise of Bots and the Unseen Risk for Business #2025BADBOTREPORT*. Tech. rep.

Kuhn, Daniel, Soroosh Shafiee, and Wolfram Wiesemann (2025). *Distributionally Robust Optimization*. Tech. rep. arXiv: 2411.02549v3.

Kuhn, Daniel et al. (Nov. 2024). *Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning*. en. arXiv:1908.08729 [stat]. DOI: 10.48550/arXiv.1908.08729. URL: http://arxiv.org/abs/1908.08729 (visited on 12/27/2025).

Liu, Pangpang et al. (June 2024). *Contextual Dynamic Pricing with Strategic Buyers*. en. arXiv:2307.04055 [stat]. DOI: 10.48550/arXiv.2307.04055. URL: http://arxiv.org/abs/2307.04055 (visited on 01/20/2026).

MarkNtel Advisors (2025). *Global AI Agent Market Research Report: Forecast (2026–2032)*. Tech. rep. Noida, Uttar Pradesh, India: MarkNtel Advisors. URL: https://www.marknteladvisors.com/research-library/ai-agent-market.html (visited on 12/12/2025).

Mueller, Jonas W, Vasilis Syrgkanis, and Matt Taddy (2019). "Low-Rank Bandit Methods for High-Dimensional Dynamic Pricing". In: *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, pp. 15442–15452.

Mullapudi, Pavan (2025). *A Reinforcement Learning Approach to Dynamic Pricing*. Tech. rep. Publication Title: International Journal on Science and Technology (IJSAT) IJSAT25049558 Volume: 16 Issue: 4.

Parkes, David C. and Michael P. Wellman (July 2015). "Economic reasoning and artificial intelligence". In: *Science* 349.6245. Publisher: American Association for the Advancement of Science, pp. 267–272. ISSN: 10959203. DOI: 10.1126/science.aaa8403.

Roughgarden, Tim (2013). *CS364A: Algorithmic Game Theory Lecture #5: Revenue-Maximizing Auctions* *. Tech. rep.

Russell, Stuart and Peter Norvig (2021). *Artificial Intelligence A Modern Approach Fourth Edition Global Edition*. ISBN: 978-1-292-40117-1.

Shahidi, Peyman et al. (2025). "The Coasean Singularity? Demand, Supply, and Market Design with AI Agents". en. In.

Varian, Hal R (1995). *Economic Mechanism Design for Computerized Agents*. Tech. rep.

Xia, Boming et al. (Nov. 2025). "Evaluation-Driven Development and Operations of LLM Agents: A Process Model and Reference Architecture". In: arXiv: 2411.13768. URL: http://arxiv.org/abs/2411.13768.

Xie, Tianbao et al. (May 2024). *OSWORLD: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments*. Tech. rep. arXiv: 2404.07972v2. URL: https://os-world.github.io.

Yokoo, Makoto, Yuko Sakurai, and Shigeo Matsubara (2004). "The effect of false-name bids in combinatorial auctions: New fraud in internet auctions". In: *Games and Economic Behavior*

46.1. Publisher: Academic Press Inc., pp. 174–188. ISSN: 08998256. DOI: 10.1016/S0899-8256(03)00045-9.

A Terminology

Agent A An actor of non-human nature, powered by an LLM.

Human H An individual human with some job to be done.