

Artificial Intelligence in Microeconomics: Machine Learning for Demand Estimation and Pricing

Author: **Charbel Attieh** ¹

Affiliation:

1. *Business School, Holy Spirit University of Kaslik, P.O. Box 446, Jounieh, Lebanon*
2. *CIRAME Research Center, Holy Spirit University of Kaslik, P.O. Box 446, Jounieh, Lebanon*

Abstract:

This paper looks at how artificial intelligence (AI) is used in microeconomics. It focuses on two tasks: estimating demand and setting prices. The literature review covers three areas. Machine learning (ML) is used to make predictions in economics. Pricing algorithms can change prices in some markets. Generative AI helped workers complete tasks faster in several studies. The paper then uses a computer simulation to test one question: does a firm earn more profit when it uses an ML demand model instead of a standard regression? In the simulation, a firm used one of three models to choose its price. With 2,000 observations, gradient boosting (an ML method) reached 99.0% of the best possible profit. Linear regression reached 96.0%, and log-log regression reached 90.7%. With 100 observations, the gap between methods was smaller and the ML results varied more.

Keywords:

Artificial intelligence, machine learning, demand estimation, algorithmic pricing, microeconomics

Introduction:

Microeconomics studies how consumers and firms make choices. Two common tasks are finding out how buyers react to price changes (demand) and choosing the price that gives a firm the most profit. Both tasks involve predicting what will happen in the future. Machine learning is a type of AI that is built to make predictions (Agrawal et al., 2018).

AI is used in economics in three main ways. First, economists use ML methods to predict outcomes and to work with large datasets (Athey & Imbens, 2019; Mullainathan & Spiess, 2017; Varian, 2014). Second, firms use algorithms to set prices, and researchers have studied how this affects competition (Assad et al., 2024; Calvano et al., 2020). Third, generative AI tools are used for writing, coding and customer service (Brynjolfsson et al., 2023; Korinek, 2023; Noy & Zhang, 2023).

Furthermore, the transition from traditional econometric modeling to machine learning represents a fundamental shift in analytical philosophy for modern firms. Standard regression models, such as linear and log-log specifications, have long been favored by economists for their interpretability and their direct alignment with established theories of elasticity. However, these parametric approaches risk oversimplifying complex, real-world consumer behaviors by forcing data into pre-specified functional forms. In contrast, machine learning algorithms like gradient boosting are designed to capture intricate, non-linear demand dynamics organically. Comparing these approaches clarifies the underlying trade-off between the structural transparency of traditional econometrics and the predictive flexibility of modern AI, helping to identify exactly when advanced computational complexity translates into tangible financial gains.

Many studies describe what AI can do. Fewer studies measure how much data a firm needs before an ML demand model gives better prices than a standard regression. This paper addresses that gap. It has two research questions. RQ1: How is AI used in microeconomic analysis and in firm decisions? RQ2: Under what conditions does an ML demand model lead to a better price than a regression model? The objectives are to summarize the existing research and to test RQ2 with a simulation.

Literature Review:

2.1 Machine learning in economics

Varian (2014) described ML methods such as decision trees and penalized regression as tools for economists who work with large datasets. Mullainathan and Spiess (2017) explained the difference between two tasks: estimating the value of a parameter and predicting an outcome. They described ML as a tool for the second task. Athey and Imbens (2019) reviewed ML methods for economists. They noted that a model that predicts well does not always give a correct estimate of the effect of a price change, so ML methods need to be adjusted when the goal is to find causes.

2.2 AI and prediction

Agrawal et al. (2018) describe AI as a fall in the cost of prediction. In their framework, cheaper prediction increases the value of data and of decisions that depend on predictions. For a firm, a better demand forecast can lead to a price closer to the profit-maximizing price. This is the link tested in this paper.

2.3 Algorithmic pricing

Calvano et al. (2020) built a simulation in which pricing programs that learn by trial and error ended up charging prices above the competitive level, even though the programs did not communicate. Assad et al. (2024) studied retail gasoline stations in Germany. They found that profit margins rose when competing stations in the same market adopted algorithmic pricing. Brown and MacKay (2023) found that online retailers change their prices at different speeds, and that faster algorithmic pricing can lead to higher prices.

2.4 Generative AI at work

Noy and Zhang (2023) ran an experiment in which professionals used a generative AI tool for writing tasks. The group that used the tool finished faster, and its work received higher quality ratings. Brynjolfsson et al. (2023) studied customer-support agents who received an AI assistant. The agents resolved more issues per hour, and the increase was largest for less experienced workers. Korinek (2023) listed ways economists can use large language models in research, such as coding, data work and drafting, and noted that the outputs need to be checked. Horton (2023) tested whether language models can act as simulated economic agents in experiments.

2.5 Advanced ML in Market Prediction and Asset Pricing

While algorithmic pricing alters retail consumer markets, machine learning has also fundamentally transformed price prediction, valuation, and forecasting in broader financial contexts. Recent applications demonstrate AI's capacity to decode complex, non-linear market signals that traditional econometric models often miss. For example, Khalife et al. (2025) applied deep learning models to capture the impact of investor sentiment on S&P 500 prices, highlighting the ability of advanced neural networks to translate unstructured behavioral data into actionable price forecasts. Similarly, Nasser et al. (2025) utilized ensemble machine learning techniques to predict macroeconomic phases through asset returns, showcasing how combining multiple algorithms improves predictive robustness in economic analysis. In the realm of complex valuation, AI is now being actively synthesized with traditional mathematical frameworks; Khalife et al. (2023) successfully integrated artificial intelligence into Monte Carlo simulations to enhance the pricing accuracy of complex financial derivatives under uncertainty. However, as the use of these powerful algorithms expands across business and economic functions, the ethical implications and cultural challenges of AI adoption must also be carefully navigated by organizations (Diab & El Hajj, 2024).

2.6 Research gap

The studies above show that AI is used for prediction and that it can change market outcomes. They do not directly show how the amount of data affects the price a firm chooses when it uses ML or regression. This paper tests that question in a simulation where the true demand curve is known.

Methodology:

3.1 Research design

The study uses a Monte Carlo simulation. In a simulation, the researcher creates the data and knows the true demand curve. This makes it possible to measure exactly how far each method's price is from the best price. The results describe the simulated market. They do not describe a real market.

3.2 Simulated market

One firm sells a product. Demand is $Q = 220(0.6 + 0.8z)\exp(-0.012P^{1.5}) \times \varepsilon$. Here P is price, z is a demand shifter between 0 and 1 (for example, customer traffic or season), and ε is random noise with a standard deviation of 0.15. Marginal cost is \$5. Demand becomes more sensitive to price as price rises, so the demand curve is not a straight line. The firm has past data with n observations. Prices in the past data were random values between \$6 and \$40.

3.3 Estimation methods

The firm uses one of three models to predict demand: (a) linear regression of quantity on price and z ; (b) log-log regression, which assumes a constant price elasticity; and (c) gradient boosting, an ML method (150 trees, depth 3, learning rate 0.08). For each of 21 values of z , the firm picks the price on a \$0.50 grid that gives the highest predicted profit, calculated as $(P - MC) \times \text{predicted } Q$.

3.4 Measures

Each chosen price is checked against the true demand curve. Profit as a percentage of the best possible profit is the first measure. The second measure is the average gap between the chosen price and the best price, in dollars. Each case was repeated 150 times with sample sizes of 100, 500 and 2,000. The table shows the averages. The code was written in Python with NumPy and scikit-learn, using a fixed random seed.

Results:

Table 1 shows the results. The best price is \$18 for every value of z , because z multiplies demand and does not change the best price.

Table 1. *Profit and price error by method (150 repetitions per row)*

Method	Sample size	Profit as % of best possible	Std. dev. (pp)	Average price error (\$)
Linear regression	100	95.9	0.7	3.11
Log-log regression	100	90.0	5.4	4.33
Gradient boosting (ML)	100	97.0	2.8	2.20
Linear regression	500	96.0	0.3	3.10
Log-log regression	500	90.9	2.2	4.30
Gradient boosting (ML)	500	98.2	1.4	1.70
Linear regression	2,000	96.0	0.1	3.08
Log-log regression	2,000	90.7	1.3	4.36
Gradient boosting (ML)	2,000	99.0	0.8	1.23

Note. *pp* = percentage points.

With 2,000 observations, gradient boosting reached 99.0% of the best possible profit. Linear regression reached 96.0% and log-log regression reached 90.7%. The average price error of gradient boosting fell from \$2.20 at 100 observations to \$1.23 at 2,000 observations. The price error of linear regression stayed near \$3.08 at every sample size. Log-log regression had an error of about \$4.36.

With 100 observations, gradient boosting reached 97.0% and linear regression reached 95.9%. The ML results varied more between repetitions (standard deviation of 2.8 percentage points, compared with 0.7 for linear regression). Log-log regression had the lowest profit at every sample size, between 90.0% and 90.9%.




Figure 1. *Profit as a percentage of the best possible profit*

Discussion and Conclusion:

The literature shows three uses of AI in microeconomics: ML for prediction, algorithms for pricing, and generative AI for analytical work. In the simulation, the ML demand model led to higher profit than both regression models when the demand curve was not a straight line and the sample was large. With 2,000 observations, the regression models reached about 91% to 96% of the best possible profit. The gap between methods was smaller with 100 observations.

The study has limits. The data are simulated from one demand function. The market has one seller and no competitors. The ML model was used for prediction only and did not include a method for finding causal effects, which Athey and Imbens (2019) describe as a difficulty with real data. Future studies could use real sales data, add competing firms to study the pricing outcomes described by Calvano et al. (2020) and Assad et al. (2024), and look at effects on consumers such as price discrimination.

Ultimately, these findings carry significant managerial implications for the deployment of algorithmic pricing strategies. The results underscore that the financial superiority of machine learning is highly conditional on data abundance. For firms operating with limited historical data, the marginal benefit of deploying complex algorithms over traditional regression may not justify the associated computational costs, particularly given the higher variance in ML performance observed in small samples. Conversely, in data-rich environments, the capacity of AI to flexibly map complex demand responses unlocks substantial profit opportunities that rigid linear models leave on the table. As algorithmic pricing becomes a standard competitive necessity, a firm's success will increasingly depend not just on adopting AI, but on carefully aligning its analytical sophistication with the depth and quality of its data architecture.

References

- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- Assad, S., Clark, R., Ershov, D., & Xu, L. (2024). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy*, 132(3), 723–771.
- Athey, S., & Imbens, G. W. (2019). Machine learning methods that economists should know about. *Annual Review of Economics*, 11, 685–725.
- Brown, Z. Y., & MacKay, A. (2023). Are online prices higher because of pricing algorithms? *American Economic Journal: Microeconomics*, 15(2), 275–308.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research.
- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297.
- Diab, B., & El Hajj, M. (2024). Ethics in the age of artificial intelligence: unveiling challenges and opportunities in business culture. *Cogent Business & Management*, 11(1). <https://doi.org/10.1080/23311975.2024.2408440>
- Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from Homo Silicus? (NBER Working Paper No. 31122). National Bureau of Economic Research.
- Khalife, D., Yammine, J., El Bazi, T., Zaki, C., & Jabbour Al Maalouf, N. (2025). Analyzing the impact of investor sentiment on S&P 500 prices using deep learning models. *Journal of Financial Reporting and Accounting*, 23(4), 1365-1386.
- Khalife, D., Yammine, J., Rahal, S., & Freiha, S. (2023). Pricing Asian and Barrier Options Using a Combined Heston Model and Monte Carlo Simulation Approach with Artificial Intelligence. *Mathematical Modelling of Engineering Problems*, 10(5), 1690.
- Korinek, A. (2023). Generative AI for economic research: Use cases and implications for economists. *Journal of Economic Literature*, 61(4), 1281–1317.
- Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106.
- Nasser, A., Yammine, J., Harouny, J., & Khalife, D. (2025). Predicting Economic Quads Through Asset Returns Using Ensemble Machine Learning. *IEEE Access*.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–28.