

AI In Product Management
Prof. Zillur Rahman
Department of Management Studies
Indian Institute of Technology, Roorkee

Lecture - 45
Price Optimization using AI (Part 5)

Welcome to this NPTEL online certification course on Artificial Intelligence in Product. Now we will talk about Module 45, wherein we will understand price optimization using AI for product managers, and we are talking about Part 9, which is price optimization using AI. So, this is the fifth part of this topic, and we are talking about Module 45. So, to give you an overview of this module, we will start with understanding AI-assisted dynamic pricing in detail through the case study of Airbnb. Then we will understand the pitfalls in AI algorithmic pricing through the case study of Uber.

And then we will understand how to overcome these pitfalls. So, let us now talk about the Airbnb dynamic pricing strategy. That is a case study. Airbnb, founded in 2008, turned the idea of renting out extra space in your home into a global phenomenon. With millions of

listings worldwide, managing pricing was a huge challenge. They needed a system to ensure prices were competitive and optimized for both hosts and guests. Enter dynamic pricing, a strategy that adjusts prices in real time based on various factors. So, how does the dynamic pricing system work? Airbnb's dynamic pricing strategy is like a smart assistant that helps hosts set the right price at the right time.

And here's how it works. So, it starts with data collection—historical data. Airbnb collects loads of data on past bookings, prices, and seasonal trends. Then, it also collects current market data, keeping an eye on listings, occupancy rates, and competitors' pricing. External factors—local events, holidays, weather, and even flight data—help predict demand spikes and user behavior.

Look at what users are searching for, booking patterns, and preferences. Next, we will talk about the algorithms and techniques. Machine learning models predict the best prices based on historical and real-time data. Time series analysis techniques like ARIMA help analyze trends over time, such as seasonal demand changes. Regression models understand how factors like location and amenities affect prices.

Reinforcement learning helps the system learn from the outcomes of its pricing decisions and elasticity models. These models gauge how sensitive demand is to price changes. Next comes price optimization. And we will start with smart pricing.

This tool automatically adjusts prices based on demand, availability, and market trends. Then, there are customization options. Hosts can set minimum and maximum prices and adjust prices manually if they prefer. Localized adjustments. Prices are adjusted differently based on location, taking into account local market conditions and competition.

Now, the use case is dynamic pricing in action. Let us imagine a popular beach town during the summer. Here is how Airbnb's dynamic pricing system would work. So, it will start with data collection. Historical data shows that the town sees a surge in bookings from June to August.

Real-time data indicates that many listings are already booked for the upcoming holiday weekend. An upcoming music festival is expected to draw large crowds. Then, it does price prediction. The machine learning models predict high demand for the holiday weekend. The system analyzes how similar events in the past affected bookings and prices.

Then it goes for price adjustments. Smart pricing suggests increasing prices as the holiday weekend approaches. Hosts are notified of the suggested price changes and can choose to accept or adjust them. Prices are adjusted daily based on the latest data, ensuring they remain competitive. The outcome is hosts who follow the pricing suggestions see higher bookings and revenues.

Guests benefit from finding a variety of options at different price points. Now, how does this technical architecture look? So, the first is the data pipeline: data ingestion. Data is pulled from various sources, including internal databases and third-party APIs. Then comes data storage. Data is stored in systems like Apache Cassandra and Amazon S3. Then it goes for data processing. Real-time data streaming is managed with Apache

Kafka, and batch processing is done with Apache Spark. Then comes model training and deployment. So, feature engineering: complete features are created from raw data, such as seasonality, indices, and market trends. Model training. Models are trained on frameworks like TensorFlow and PyTorch using GPU clusters.

Then, model deployment happens. Trained models are deployed using containerized environments like Docker and Kubernetes to ensure scalability and reliability. Continuous learning happens. The system continuously learns from new data, keeping models up to date. Scalability and performance.

Airbnb uses a microservice architecture to handle different parts of the dynamic pricing system. Load balancing and caching strategies help manage high traffic volumes and ensure quick responses. What is the impact on the business for all this? The first is increased revenue. Dynamic pricing helps maximize revenue for hosts by adjusting prices based on demand.

Hosts using smart pricing often see an increase in bookings and overall revenue. Then, it also leads to a competitive advantage. Offering competitive prices helps Airbnb attract more guests, even during peak times. This helps maintain a steady flow of bookings and reduces vacancy rates. Improved user satisfaction.

Guests get fair and competitive pricing while hosts enjoy optimized income from their listings. This balance enhances overall user satisfaction and operational efficiency. Dynamic pricing reduces the need for manual price adjustments by hosts, saving time and reducing errors. It also allows Airbnb to better manage its inventory and predict demand more accurately. What are the challenges and future directions? The first is data privacy. Ensuring the privacy and security of user data is crucial. Airbnb uses robust encryption and anonymization techniques to protect data. Market volatility. Rapid changes in market conditions, such as those caused by the COVID-19 pandemic, require the dynamic pricing system to be highly adaptable and resilient.

Algorithmic fairness. Ensuring that dynamic pricing algorithms do not inadvertently disadvantage certain hosts or guests is essential. Airbnb continuously monitors and updates its models to maintain fairness. Advanced forecasting. Incorporating more advanced forecasting techniques and external data sources can further enhance the accuracy of dynamic pricing models.

Now, what are the takeaways for product managers? The first is data-driven decision-making. Product managers should leverage data to make informed decisions. Airbnb's dynamic pricing system shows the power of using data analytics and machine learning to drive business outcomes. Another is a customer-centric approach.

Balancing the needs of different stakeholders, like hosts and guests, is crucial. Airbnb's dynamic pricing strategy ensures that both parties benefit, leading to higher satisfaction and loyalty. Continuous improvement: The iterative approach to model training and deployment highlights the importance of continuous improvement. Product managers should adopt a similar mindset, always looking for ways to refine and enhance their products.

Scalability and flexibility: Building a scalable and flexible system is vital for handling growth and adapting to changes. Airbnb's microservices architecture and use of a containerized environment ensure that its dynamic pricing system can scale effectively. What are the ethical considerations? Addressing ethical considerations, such as data privacy and algorithmic fairness, is essential. Product managers must ensure that their systems are not only effective but also fair and responsible.

So, what are the pitfalls in AI algorithmic pricing? On June 3, 2017, Blue lights flashed across London Bridge as police cars responded to reports of a terrorist attack. They blazed past thousands of people who were enjoying a Saturday night at restaurants and pubs in the area. Many of those who were out on the streets, sensing danger, attempted to order an Uber and head home to safety.

But for 43 minutes after the first emergency call came in, at 10:07 pm, Uber's dynamic pricing algorithm caused rates in that part of the city to jump more than 200%. The London episode is just one of many troubling examples of Uber's pricing surges during moments of collective anxiety. Similar spikes occurred during a 2016 bombing in New York City, a 2017 taxi driver strike in protest of US anti-immigration policy, and a 2020 Seattle mass shooting. The last of which sent fares soaring as much as 500%.

So this is what it shows: 7 minutes away and it is charging \$85. Uber's algorithmic pricing has consistently sparked criticism from the ride-sharing company's 93 million active users. Even on the night of the London Bridge attack, after Uber manually halted surge pricing near London Bridge, it remained in effect for the surrounding areas of central London for another 50 minutes. An economist might applaud Uber's pricing engine: as demand increases relative to supply, the price of a ride climbs.

For customers, however, the cost of using the service can seem as unpredictable as the spin of a roulette wheel. Uber is not the only company facing this problem. Firms in many industries, including advertising, e-commerce, entertainment, insurance, sports, travel, and utilities have employed dynamic pricing with varying degrees of success. A

classic and well-known example is Coca-Cola, which experimented in the late 1990s with temperature-sensitive vending machines that would increase the price of a beverage on a hot day.

The company quickly abandoned the project in the wake of public outrage. Pricing algorithms are intended to help firms determine optimal prices on a near real-time basis. They use artificial intelligence and machine learning to weigh variables such as supply and demand, competitive pricing, and delivery time. Unfortunately, algorithms occasionally go rogue and come up with figures no one would ever pay, from \$14,000 for a cabinet listed on Wayfair to almost \$24 million for a textbook offered on Amazon.

But such snafus are just one of the risks when companies entrust decision-making to computers. The constant changes in price points send strong signals to customers that need to be properly managed. Yet many organizations fail to appreciate this. They know that prices affect decisions about when and what to purchase, but they overlook the fact that continual ups and downs may trigger unfavorable perceptions of their offerings and, importantly, the company itself. Brands thus need to consider more than simple math when employing algorithmic systems.

These systems can create an uncomfortable tension between easing customers' loyalty between earning customer loyalty and earning money. But implemented correctly, they can maximize revenues while also making customers feel as if they have paid the appropriate amount for a product or service. In the worst case, algorithms turn the already delicate task of asking customers for money into an experience that drives them away.

Brands thus need to consider more than simple math when employing algorithmic systems. Overcoming the Pitfalls: Following are four ways by which the pitfalls in algorithmic pricing can be overcome. One is to determine the appropriate use case and narrative. Second is to designate a pricing algorithm owner. The third is to set and monitor pricing guardrails.

And the fourth is to override the algorithms when necessary. So let us start with the first one: determine an appropriate use case and narrative. In 2020, the Swedish furniture retailer IKEA launched a novel initiative at its Dubai location. For a limited period, the company allowed customers to pay different prices for products according to the time they spent driving to the store. Every item—from a sandwich at the restaurant to a complete bedroom set—had a price expressed in two units: local currency and the time amount.

A family that drove, say, 45 minutes to the IKEA store earned a certain value tied to the distance of their trip. At checkout, the family could show the cashier a Google Maps timeline readout, using a feature of the Google Maps cell phone app that tracks and records all the routes one takes. The cashier would run an algorithm that factored in time spent, distance traveled, and the average hourly wage of a Dubai worker to calculate the monetary value of the ride. The store then offered that value as a form of currency.

The longer the trip, the more time credits the family got, and the less money it needed to fork out. The clear inference shoppers drew from IKEA's program was that the retailer wanted to incentivize them to travel great distances to its stores. Although different customers would pay different prices for the same item, an individual customer might see different prices each time they visited, depending upon where they came from, they nonetheless felt that they had agency in how much they would pay. That contrasted with the helplessness people often experience during price surges. Best of all, because customers' out-of-pocket costs could only decrease in conjunction with distance traveled—as opposed to increasing—As a result of heightened demand, no one ever paid more than the price advertised on the company's website.

So, this is 'buy with your time.' How did it work? In other words, IKEA used distance-based algorithms to reward customers rather than penalize them. It might have lost some immediate revenue. Shoppers who drove far enough would get steep discounts or even obtain some products for free.

But by choosing a proper use case with built-in incentives for people to visit the store, the company probably attracted more remote customers and increased all customers' loyalty and theoretical lifetime value. Models like IKEA's are rare. Companies typically employ dynamic pricing to further their short-term financial goals with less regard for customers' perceptions. Yet the sheer volume and intensity of price changes implemented by algorithms send unequivocal signals to buyers about everything—from a company's mission and values to the quality of its offerings.

These signals can crowd out other efforts to shape the narrative in a brand's relationship with its customers. In the worst case, algorithms turn the already delicate task of asking for money into an experience that drives customers away. This is why firms cannot leave the management of emerging technology to data scientists alone. The path to improvement is not just technical but also organizational and psychological.

As paradoxical as it might sound, a better algorithm might make matters worse. by exploiting circumstances and stirring resentment, as happened with Uber during the London Bridge attack. Overcoming the organizational challenge starts with recognizing that algorithmic pricing is not simply a means to generate prices but to bring supply and demand into balance. It is, in fact, a principle that needs to align with an organization from top to bottom.

When customers have the impression that a firm bases its price solely on supply and demand, the inferences they draw can be harmful. Think of an innovative firm with highly differentiated offerings. When that firm emphasizes supply and demand in its pricing algorithm, it is essentially telling customers that the value of its product is mostly related to availability, not how well it solves customers' problems or performs relative to competitors. Additionally, customers can learn to game the system and time their purchases to coincide with moments when they believe the price is low. This again drives commoditization.

By contrast, IKEA's dynamic pricing model focused on attracting unlikely customers rather than penalizing likely customers due to a lack of supply. The next pitfall is designating a price algorithm owner. In 2019, United Airlines eliminated the mileage tables that frequent flyers relied on to redeem their reward points. It replaced the tables with an algorithmic pricing model, explaining why it was necessary

to align supply and demand, and emphasized how customers could benefit by spending fewer award miles for off-peak flights. The new system did result, however, in higher award prices for high-demand flights. That frequently frustrated reward customers, but the airline communicated all the changes in an easily understandable way and focused its efforts on the specific and presumably loyal customer base. In doing so, it was able to mitigate significant reputational damage. Additionally, because it delegated

management of the new algorithm to the team that supervised the loyalty program, United gave clear ownership of the pricing system to a department highly attuned to the sensitivities of its most steadfast customers. That strategy enabled the airline to monitor and quickly respond to glitches with the algorithm or challenges in customer relationships. It is easy to blame the algorithms themselves when they go haywire, but the root cause of the problem usually lies in other areas. Inadequate organizational attention or a failure to appreciate customer psychology. Most firms have an incomplete understanding of what really happens when they ask customers for money. They focus

too intensively on the numbers, which they view as little more than the passive outcomes of the market forces that shape supply and demand. To use Adam Smith's term, the invisible hand does the work, not the firm itself.

This myopia leads companies to overlook all the other information that prices convey. Even when organizations do recognize the power of this information and its implications, most firms cannot manage it effectively. Because pricing is an organizational orphan, with no clearly defined leadership, responsibility, or accountability. When companies belligerently hand off the heavy lifting of pricing to automation, they cede to the algorithm not only the control of the mass but also the messaging.

While the data scientists, data analysts, and pricing specialists focus on optimizing the numbers, who is making sure that the messages are optimal? The answer in many organizations is no one. A pricing algorithm on its own has two weaknesses. First, it lacks the empathy required to anticipate and understand the behavior and psychological effects that price changes have on customers. Second, it lacks the long-term perspective required to ensure compliance with a corporate strategy or overarching purpose.

By emphasizing only supply and demand fluctuations in real time, the algorithm runs counter to marketing teams' aims for long-term relationships and loyalty. This conflict between long-range thinking and real-time price changes does not merely intensify the clash between earning goodwill and earning money. It also increases the urgency of finding a solution before the brand suffers irreversible damage. If a firm does not manage its price setting and messaging proactively and strategically, it can trigger and even accelerate the commoditization of its offerings by heightening price sensitivity, undermine price-value relationships, and tarnish brand reputation.

But by empowering a team that can plan its initiatives and make in-the-moment decisions about them, the company can pivot quickly when a predicament occurs. The third is to set and monitor pricing guardrails. Think about a typical poor experience at a theme park. Guests have to suffer through long lines for rides, food, and restrooms, plus a lack of personal attention from overwhelmed or undertrained support staff. Such an off-putting experience leaves many customers wondering whether their steep investment in tickets, parking, refreshments, and lodging is even worth it.

Guests would have a more pleasant visit if they encountered shorter lines and wait times and had better interactions with park personnel. To increase customer satisfaction, Walt Disney World in Orlando, Florida, changed its dynamic pricing structure from a manual

to an algorithmic one in 2018. The new program, which raised multi-day ticket prices overall but decreased the price of tickets for off-peak dates, encouraged customers to plan their trips well in advance or book tickets during off-peak periods to take advantage of lower prices. The business program had several merits. First, it shows that dynamic pricing can serve other objectives,

besides increasing revenue or volume. Even if total revenue and overall guest count stay constant over time, the pricing structure makes the flow of customers steadier, which means less volatility in Disney's need for staff and other resources. This can lead to significant cost savings. Second, the customer experience improves dramatically because guests can enjoy more rides, visit more attractions, and better use their time in the parks. Finally, the dynamic pricing program can be explicitly publicized as a commitment to long-term customer satisfaction, despite an overall increase in price.

When Disney World switched to this algorithmic system, it also determined that it would be in its best interest to no longer dynamically price single-day entry to its individual theme parks. Pricing for a single-day ticket across all four properties was set from \$109 to \$129, no matter what time of year a customer chose to visit, regardless of demand. This guardrail limited the amount that Disney could charge for a single-day pass, but it set clear parameters that helped customers anticipate their costs and plan their visits. And by observing how they self-selected their trips, Disney could sharpen its communication about the park experience and design additional service packages to cater to different customer segments.

Other companies can use guardrails in a similar way, not just to protect customers from wild price swings, but also to judge how pricing impacts every area of the organization. When establishing the initial guardrails and continuing to deploy them, firms should encourage information sharing among different lines of business. That is the best way to extract key learnings and use them for the company's benefit. Now, there are three primary areas for close collaboration across functions to glean insights from algorithms: experimentation, monitoring, and strategy.

So, we will start with the first: experimentation. Controlled periodic testing of prices can help a company measure the extent to which customers value a product or service or any of its features and understand the context of when and how they derive that value. Indeed, pricing experimentation can be far more powerful than traditional market research because customers are reacting to actual offerings and making real transactions. Their

responses to price shifts help firms discover what works, what does not, and at what point buyers first make their purchase decisions.

Monitoring firms can develop a new key performance indicator or compare existing indicators to ensure that the frequency and magnitude of price changes are not eroding customers' loyalty or brand reputation. No company wants to be perceived as unfair, manipulative, or greedy. Thus, it is important to take measures to constrain and manage the output of the pricing algorithm and vital to think through the message and its consequences in advance. This enables firms to avoid extreme and free-floating prices by implementing hard floors and ceilings, as Disney did with its fixed single-day pricing. Then comes the strategy.

This is essentially a long-term integrated view of the first two elements. Are the firm's product development, branding, positioning, and pricing all working in harmony or with the least amount of friction to fulfill the company's strategic objectives? The firm must strive to ascertain, directly or indirectly, how customers perceive its mission and purpose and whether its price actions reinforce or harm the reputation it is trying to establish. The messages that customers infer from prices should sync up with the explicit messages that a company communicates through its non-price activities to promote itself and its products.

When firms pay attention to all the various ways that price changes can alter what customers believe and how they behave beyond the immediate buy-or-no-buy decision, they can enhance the customer relationship rather than diminish it—even when they raise prices. Firms can tap into the power of price changes to improve their operations and, at the same time, create a better overall experience for customers. Override the algorithm when necessary—far from the set-it-and-forget-it approach to pricing that was common in the past. Organizations with a dynamic strategy must take a more proactive, creative stance to achieve the desired results.

For Disney, IKEA and United Airlines, the aim was simple. The brands wanted to make it worth the customers while to transact, even under less than ideal circumstances. For Disney, IKEA and United Airlines, the aim was simple. The brands wanted to make it worth the customers' while to transact, even under less than ideal circumstances. They also wanted to benefit from being able to manage how, when and why changes in prices were communicated.

The best pricing algorithm can analyze customers' data and other information to generate optimal prices for any given customer at any given time. but from whose perspective are those prices optimal that question gets at the conflict between earning customer goodwill and earning more money which presents a complicated organizational challenge that should be overseen by a clear owner and managed when necessary sometimes the algorithm might need to be tweaked other time it might need to be temporarily suspended The day after the London Bridge attack, Uber announced that it had refunded the payments of all the rides who had hired a car in the affected area. It also boasted that its drivers had helped tens of thousands of people flee the scene. Both announcements would likely have enhanced the company's reputation had it not just been tarnished by a swift backlash to the price surge.

Although it is difficult to quantify the lasting negative impact of that surge on Uber's relationship with its customers, it is clear that a faster response or a more proactive mechanism for preventing the soaring prices would have benefited the brand and the riders served that evening. All companies should understand what their pricing algorithms are communicating to customers and how best to control that message. To effectively do so, they must develop a proper use case and narrative for implementing algorithm pricing, assigning an owner to monitor price guardrails and empower the owner to a manager or override the automation when necessary. By doing so, companies will be able to optimize dynamic pricing in real time without sacrificing customer loyalty or harming their reputation.

So, to conclude, in this module, we have first understood AI-assisted dynamic pricing in detail through the case study of Airbnb. We have also discussed the pitfalls in AI algorithmic pricing through the case study on Uber. And finally, we have understood how to overcome such pitfalls while using AI algorithmic pricing. These are some of the sources from which the material for this module was taken. Thank you.