Data-Driven Pricing Optimisation

A Framework for Maximising Profitability in Dynamic Markets

Author: Rakesh Gupta
Senior Product Manager – SaaS & AI Platforms
Date: August 2024

Executive Summary

In markets where products, demand, and customer intent fluctuate rapidly, static pricing models are no longer sufficient. Businesses must align pricing decisions with both internal and external signals to balance conversion and margin.

This paper presents a practical, data-driven framework for pricing optimisation using a combination of historical analysis and adaptive machine learning. The approach aims to maximise total profit — not by lowering or raising prices arbitrarily, but by continuously learning where each product, category, or condition reaches its optimal trade-off between conversion and margin.

The framework is designed for teams managing complex product catalogues or multi-channel environments such as travel, retail, or digital marketplaces. It highlights how organisations can evolve from static markups to self-learning optimisation systems that respond to market dynamics in real time.

1. The Pricing Challenge

Most organisations set prices or markups through rules, competitor benchmarks, or manual judgement. These methods often achieve short-term goals but leave long-term value on the table.

Two opposing objectives drive the problem:
- Higher prices increase margin per sale but reduce conversion.
- Lower prices increase conversion but reduce profit per sale.

The true goal is to maximise overall profit, defined as:

Profit = Margin per transaction × Number of transactions

This objective highlights the need for a balance — the optimal point where price and conversion work together to maximise total revenue contribution.

2. Data Foundations for Pricing Optimisation

Achieving this balance requires reliable, granular data across multiple sources. These typically include:

- Transactional data: Bookings, orders, or sales logs with pricing and margin details.
- Customer behaviour data: Conversions, sessions, and bounce rates from analytics platforms.
- Operational data: Supplier rates, cost inputs, or category hierarchies.
- External context: Market seasonality, competitor benchmarks, currency, or macroeconomic factors.

For modelling purposes, we define two categories of variables:
- Controllable variables: Factors we can adjust — price, channel, supplier, product type, duration, etc.
- Contextual variables: External influences we can't directly control — seasonality, weather, macro events, competitor actions.

By isolating controllable dimensions and treating contextual ones as constants, we can model how pricing decisions affect conversion and profit under stable conditions.

3. Modelling Framework

The model's objective is to maximise:

Total Profit = (Average Transaction Value × Markup%) × (Traffic × Conversion)

Since average transaction value and traffic are relatively constant in short windows, the focus shifts to the relationship between markup percentage and conversion rate — two variables that are inversely correlated.

In essence, the model seeks to find the markup level that maximises the product of markup × conversion.

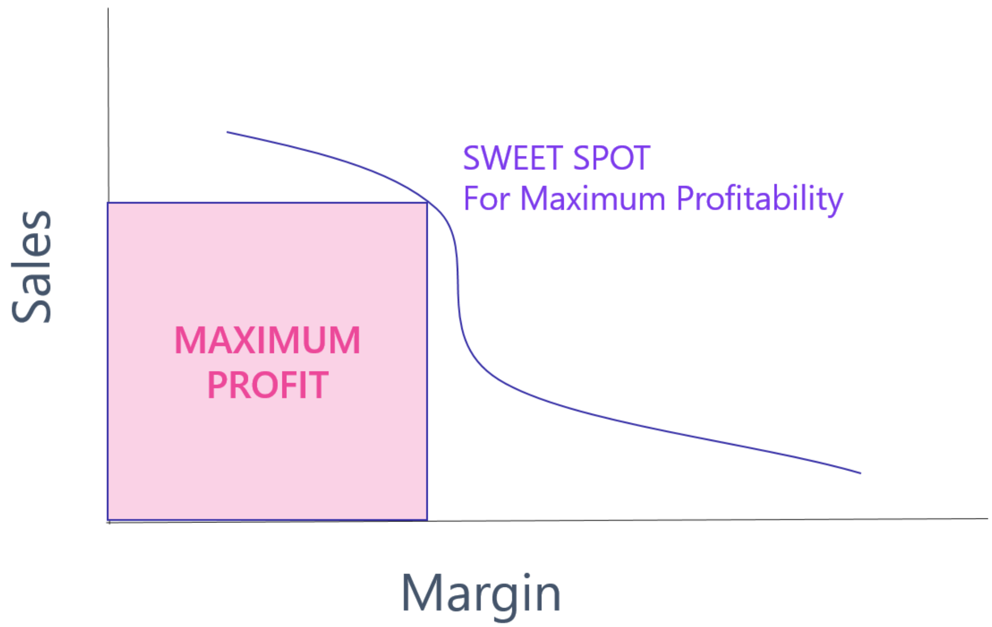

Figure 1. *Revenue curve showing conversion versus markup, with total profit peaking at the optimal midpoint.*

4. Approaches to Pricing Optimisation

Method 1: Historical Data Modelling
This method analyses past data to identify patterns between markup and conversion. By plotting revenue outcomes at different markup levels, we can estimate the optimal markup for a given set of conditions.

This approach is simple, interpretable, and works well in stable markets. However, it requires regular retraining to adapt to external shifts such as seasonality or changing competition.

Method 2: Continuous Learning via Machine Learning
A more dynamic approach uses reinforcement or adaptive learning to search for the optimum in real time. The system adjusts markup percentages incrementally, observes performance, and shifts direction based on outcomes.

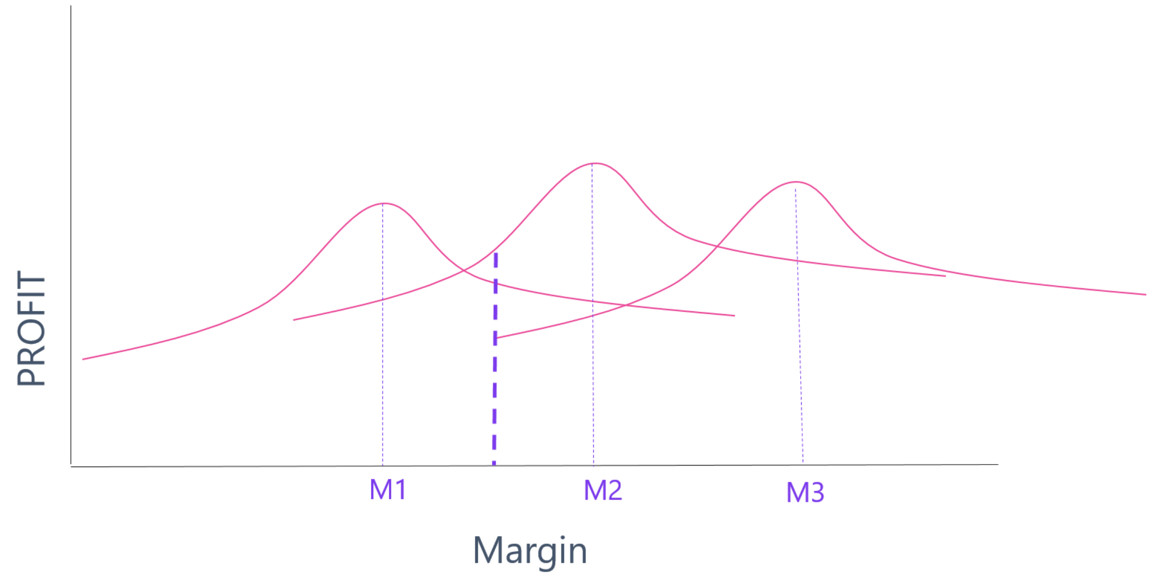

Figure 2. *Illustration of continuous learning in pricing optimisation — showing iterative*

5. Balancing Granularity and Accuracy

The model can be applied across multiple dimensions (conditions) such as channel, supplier, product type, destination, or star rating.

However, increasing the number of dimensions reduces available data per segment. This introduces a trade-off:
- More dimensions → higher precision but lower statistical reliability
- Fewer dimensions → higher reliability but less contextual accuracy

Teams should start with a few high-impact dimensions, validate results, and gradually expand as data density improves.

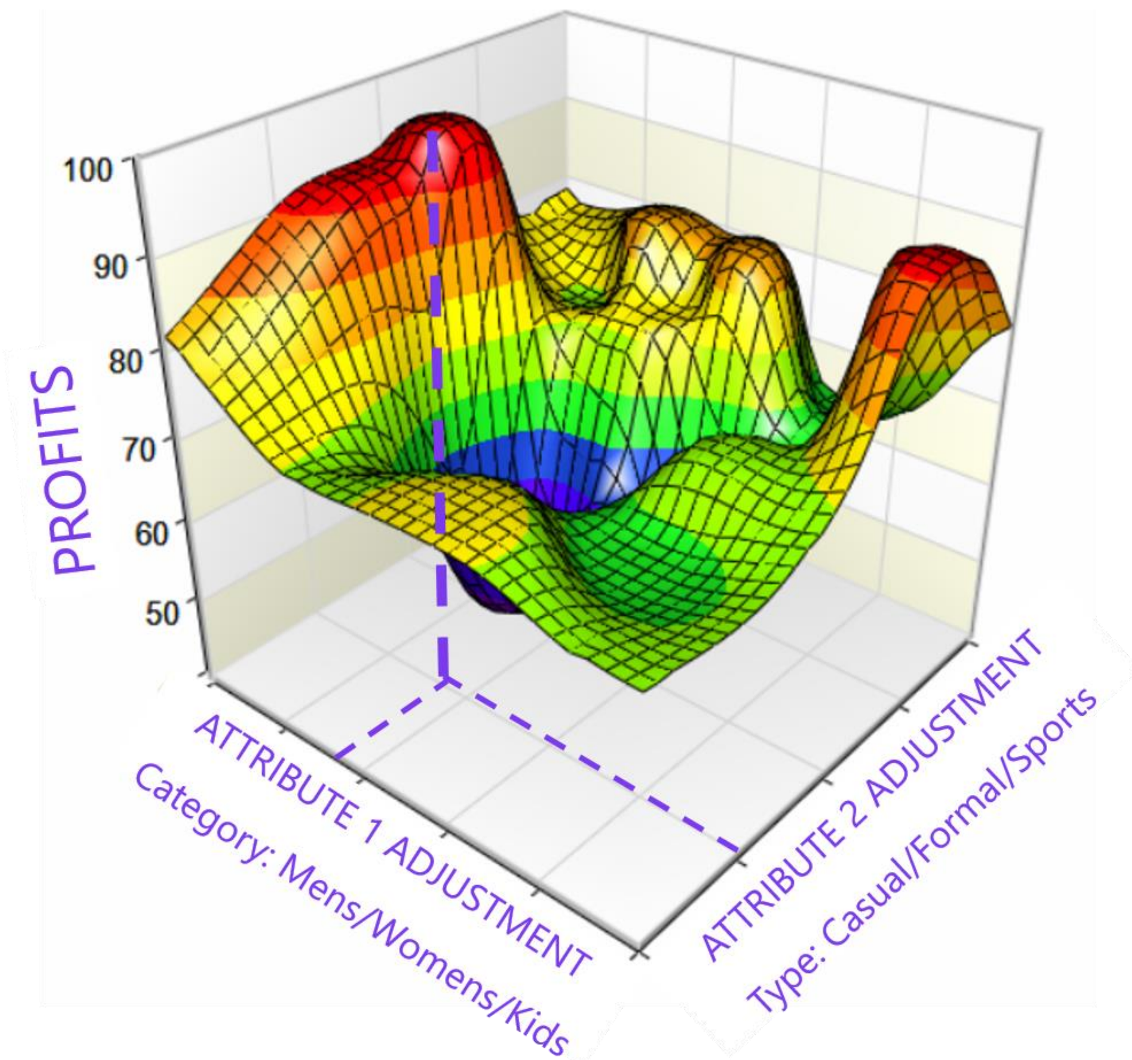

Figure 3. *Multi-dimensional optimisation view showing how margin adjustments vary across Attribute 1 (e.g., category) and Attribute 2 (e.g., style or segment), illustrating how optimal pricing differs by attribute combination.*

6. Implementation Roadmap

1. Define Objectives — Clarify whether the goal is profit maximisation, conversion lift, or competitive parity.
2. Prepare and Integrate Data — Consolidate transactional, behavioural, and contextual data into a unified model-ready dataset.
3. Develop Models — Begin with historical regression analysis to identify baseline elasticity, then progress toward automated or reinforcement models for continuous learning.

4. Run Controlled Experiments — Test recommendations on small segments, monitor results, and refine parameters before full rollout.
5. Integrate and Automate — Connect model outputs to pricing or markup engines with human oversight.
6. Monitor and Refine Continuously — Track key metrics and retrain models as new data and conditions emerge.

7. Conclusion and Key Takeaways

Dynamic pricing is no longer a luxury — it is an operational necessity in competitive marketplaces. The framework presented here transforms pricing from a manual, reactive process into a measurable, self-improving system.

By combining business logic with adaptive data science, organisations can:
- Capture margin uplift without harming conversion.
- Respond faster to market fluctuations.
- Build institutional learning around pricing elasticity.

Key Takeaways:
- Optimal pricing balances conversion and margin, not just one metric.
- Data quality and model granularity determine accuracy.
- Continuous learning models adapt best to changing conditions.
- Success depends on cross-functional collaboration between product, data, and commercial teams.

Author: Rakesh Gupta
Senior Product Manager – SaaS & AI Platforms

https://www.linkedin.com/in/rakeshgupta2/