AUTOMOTIVE CASE STUDY / CONTENT AND PRICING

Every Feature Is a Pricing Decision

Vehicle content and price optimization with mixed logit choice models, Monte Carlo simulation and combinatorial search

EXECUTIVE CONCLUSION

Whether a feature is standard, optional or unavailable on each trim, and what each trim and option costs, is one coupled decision. It should be made against a model of what heterogeneous buyers will actually choose, with uncertainty carried all the way into profit.

A vehicle line carries dozens of customer-facing features across several trims, and each feature on each trim can be standard, optional or unavailable. The combinations number in the trillions even for a small line, and every choice moves both what customers value and what the automaker earns. This note describes the method family behind content optimization: a random-utility model of buyers estimated from choice experiments, Monte Carlo simulation of choices and of parameter uncertainty, and a decomposed optimization that searches content while pricing it. A synthetic worked example shows why content, not price alone, is the lever.

S / O / U
for every feature on every trim

Mixed logit
heterogeneous willingness to pay

Monte Carlo
uncertainty carried into profit

Joint
content and price optimized together

Work led by Adam DeJans Jr. and John Brandon Elam in earlier industry roles at a global automaker, not a Bit Bros client engagement. The automaker and program are not named and no financial outcomes are published. The worked example is synthetic.

The planning problem

Long before a model year launches, an automaker must decide its content: which features each trim carries as standard, which are offered as options, which are not offered at all, and at what prices. Engineering, sourcing and manufacturing commit to those choices years ahead. Make a popular feature unavailable on the trim most buyers choose and they defect to a competitor; make an expensive feature standard on the entry trim and margin disappears for the many buyers who did not value it. The decision is combinatorial, it is priced, and it depends on what millions of heterogeneous customers would choose under alternatives they have never seen.

THE DECISION

For each trim and each customer-facing feature, should the feature be standard, optional or unavailable, and what should each trim and option cost, so that expected profit is high and robust while the line remains attractive to the customers it is meant for?

1. The decision space

Notation

Symbol	Meaning
$t \in \mathcal{T},\ f \in \mathcal{F}$	Trims (ordered, base to top) and features
$s \in \{S, O, U\}$	Standard, optional, unavailable
$z_{tfs} \in \{0,1\}$	Feature f has status s on trim t
$p_t,\ q_f$	Trim price; option price
$c_f,\ C_t$	Variable cost of a feature; body and powertrain cost of a trim
$n,\ \beta_{nf},\ \alpha_n$	Buyer; taste for feature f; price sensitivity
$w_{nf} = \beta_{nf}/\alpha_n$	Willingness to pay (WTP)
P_{nt}	Probability buyer n chooses trim t; $t = 0$ is the outside option
$M,\ \kappa$	Market size; complexity cost per optional cell

A combinatorial, priced program

Each cell takes exactly one status, and content must respect the trim ladder and engineering rules:

$$\sum_s z_{tfs} = 1 \qquad \forall t, f \qquad (1)$$

$$z_{tfS} \le z_{t+1,f,S} \qquad \forall t, f \qquad (2)$$

$$z_{t+1,f,U} \le z_{tfU} \qquad \forall t, f \qquad (3)$$

$$z_{tgS} \le z_{tfS} \ \text{ if } g \text{ requires } f \qquad \forall t \qquad (4)$$

$$p_{t+1} \ge p_t + \delta \qquad \forall t \qquad (5)$$

Constraint (2) keeps a feature standard on every trim above one where it is standard; (3) keeps it available above wherever it is available; (4) encodes package dependencies, for example leather requiring heated seats; and (5) orders the price ladder with a minimum step. Incompatibilities ($z_{tgS} + z_{tfS} \le 1$) and engineering exclusions ($z_{tfU} = 1$) are added the same way. The objective is expected profit over the market,

$$\pi(z,p,q) = M\,\mathbb{E}_n\Big[\textstyle\sum_t P_{nt}(z,p,q)\big(m_t + \sum_f \tau_{ntf}(q_f - c_f)\big)\Big] - \kappa \textstyle\sum_{t,f} z_{tfO}, \qquad (6)$$

where $m_t = p_t - C_t - \sum_f c_f z_{tfS}$ is the trim margin and τ_{ntf} indicates that buyer n adds option f to trim t. The complexity term charges for each optional cell, since every option multiplies buildable combinations, part numbers and plant sequencing effort. With $|\mathcal{T}| = 3$ trims and $|\mathcal{F}| = 9$ features there are already $3^{27} \approx 7.6 \times 10^{12}$ status assignments before prices; a production line with over a hundred features is astronomically larger. The program is a mixed-integer nonlinear program whose objective is neither concave nor separable, because choice probabilities couple every cell to every other.

2. A model of what buyers would choose

Random utility with heterogeneous tastes

Buyer n derives from trim t the utility

$$U_{nt} = \gamma_t + \alpha_n\Big(\textstyle\sum_f w_{nf} z_{tfS} + \sum_f (w_{nf} - q_f)^+ z_{tfO} - (p_t - p^{\text{ref}})\Big) + \varepsilon_{nt}, \qquad (7)$$

with the outside option, a competing vehicle at reference price p^{ref} or no purchase, normalized to $U_{n0} = \varepsilon_{n0}$. A buyer who chooses trim t adds option f exactly when $w_{nf} > q_f$, so $\tau_{ntf} = \mathbb{1}\{w_{nf} > q_f\} z_{tfO}$ and the option surplus $(w_{nf} - q_f)^+$ enters the trim's utility. With ε independent Gumbel, the choice probability conditional on tastes is logit,

$$P_{nt} = \frac{e^{V_{nt}}}{1 + \sum_{t'} e^{V_{nt'}}}, \qquad V_{nt} = U_{nt} - \varepsilon_{nt}. \qquad (8)$$

In a *mixed* logit the tastes vary across buyers, $(\beta_n, \alpha_n) \sim g(\cdot \mid \theta)$, typically normal or lognormal with mean μ and covariance Σ. Heterogeneity is what makes content optimization worthwhile: if every buyer valued a sunroof equally, there would be no reason to make it optional anywhere. Mixed logit also removes the independence-of-irrelevant-alternatives restriction of plain logit, so adding a trim draws share from similar trims first.

Estimation, calibration and validation

The taste distribution is estimated from choice-based conjoint experiments, in which respondents repeatedly choose among priced vehicle profiles that vary content systematically. Hierarchical Bayes estimation draws individual-level tastes from a population distribution by Markov chain Monte Carlo; simulated maximum likelihood maximizes $\sum_n \ln \frac{1}{R} \sum_r \prod_k P_{nk}(\beta^{(r)})$ over the population parameters. Either way, WTP for a feature follows as $w_{nf} = \beta_{nf}/\alpha_n$, a ratio whose distribution deserves care because α_n near zero produces implausible tails. Survey choices overstate purchase intent, so the trim constants γ_t are recalibrated until simulated shares of the current line match observed market shares, and the model is validated on held-out choice tasks (hit rate, share error) before it is trusted to rank content plans.

3. Monte Carlo: carrying uncertainty into profit

Simulated choice probabilities

The unconditional probability $\bar{P}_t = \int P_t(\beta) g(\beta \mid \theta)\, d\beta$ has no closed form and is estimated by averaging over R draws of buyer tastes,

$$\hat{P}_t = \frac{1}{R} \sum_{r=1}^{R} P_t\big(\beta^{(r)}\big). \qquad (9)$$

Quasi-Monte Carlo sequences such as scrambled Halton points cover the taste space more evenly than pseudo-random draws, reducing the error from roughly $O(R^{-1/2})$ toward $O(R^{-1})$ for smooth integrands, so far fewer draws reach a given precision.

Parameter uncertainty and risk

Estimated tastes are themselves uncertain. Drawing parameter vectors $\theta^{(k)}$, $k = 1, \ldots, K$, from the posterior (or a bootstrap) and re-evaluating a plan under each turns its profit into a distribution rather than a number. That distribution supports objectives a single estimate cannot: the probability that a plan beats the current line, $\Pr\{\pi(x) > \pi(x^0)\}$, or the conditional value at risk of the uplift at level a,

$$\mathrm{CVaR}_a(\Delta) = \mathbb{E}\big[\Delta \mid \Delta \le \mathrm{VaR}_a(\Delta)\big], \tag{10}$$

the mean uplift in the worst a fraction of plausible markets.

Comparing plans on common random numbers

Optimization compares thousands of candidate plans, and comparing them on independent draws would let simulation noise masquerade as improvement. Evaluating every candidate on the same draws, a sample average approximation with common random numbers, cuts the variance of a difference because

$$\mathrm{Var}(\hat{\pi}_x - \hat{\pi}_y) = \mathrm{Var}\,\hat{\pi}_x + \mathrm{Var}\,\hat{\pi}_y - 2\,\mathrm{Cov}(\hat{\pi}_x, \hat{\pi}_y), \tag{11}$$

and plans that differ in a few cells are strongly positively correlated. The chosen plan should then be re-evaluated on a fresh, larger sample, so the optimizer is not graded on the draws it was tuned to.

4. Optimizing content and price together

Pricing has a clean first-order condition

For fixed content the pricing problem is smooth. In a plain logit where one firm sells all trims to buyers with common α, setting $\partial\pi/\partial p_t = 0$ gives equal absolute markups across trims,

$$p_t - c_t = \frac{1}{\alpha\, P_0}, \tag{12}$$

where P_0 is the outside share: the more buyers the line already wins, the less a price cut can add. With heterogeneous α_n the condition becomes a weighted average over buyers and is solved numerically, but the intuition survives: price is set against the outside option, and content shifts how many buyers the outside option keeps.

A decomposition: search content, price it inside

Because the content space is discrete and huge while prices are continuous, a practical method decomposes the program. An outer search moves through status assignments; for each candidate, an inner problem optimizes trim and option prices with a gradient method, warm-started from the incumbent's prices. Moves that break the ladder or dependency rules are repaired by the smallest change that restores feasibility (making a feature standard on the trims above, for example), so the search never wastes evaluations on infeasible plans. Local search, simulated annealing or a genetic algorithm over the status matrix all fit this frame; the inner pricing step is what turns a content plan into a fair comparison.

CONTENT SEARCH WITH INNER PRICING

fix taste draws $\{\beta^{(r)}\}_{r=1}^{R}$ (common to all candidates)
$z \leftarrow$ current line; $(p, q) \leftarrow \arg\max_{p,q} \hat{\pi}(z, p, q)$
repeat
  for each cell (t, f) and status $s \neq z_{tf}$ **do**
    $z' \leftarrow \textsc{Repair}(z \text{ with } z_{tf} = s)$
    $(p', q') \leftarrow \arg\max \hat{\pi}(z', \cdot)$, warm-started at (p, q)
    if $\hat{\pi}(z', p', q') > \hat{\pi}(z, p, q)$ **then** $(z, p, q) \leftarrow (z', p', q')$
    end if
  end for
until no improving move
re-evaluate (z, p, q) on fresh draws and posterior samples

Profit and customer value are both objectives

Profit is not the only goal. Share supports scale and dealer health, and consumer surplus, the logsum $\mathbb{E}_n\big[\alpha_n^{-1} \ln(1 + \sum_t e^{V_{nt}})\big]$ in money terms, measures how much value the line delivers to buyers. An ε-constraint program, maximize π subject to share $\ge \varepsilon$, or a weighted objective $\pi + \lambda \cdot$share swept over λ, traces the efficient frontier between them, and leadership chooses a point on it knowingly.

5. A worked example

The example is a synthetic three-trim line with nine features, from heated seats to leather, each with a variable cost and a lognormal WTP distribution across buyers, plus lognormal price sensitivity and a competitor at a reference price. Choice probabilities use 2,000 scrambled Halton draws. The baseline is an engineering-led plan: features accumulate up the ladder, most options are unavailable, and prices are cost-plus at 12%, calibrated to a 31% share. The optimized plan comes from the content search above.

Plan	Profit per buyer	Share	Surplus
Engineering-led, cost-plus	\$1,023	31.0%	\$832
Same content, optimized prices	\$1,032	32.3%	\$879
Content and prices optimized	\$1,201	35.2%	\$1,005

Expected profit and logsum consumer surplus per buyer in the market (surplus is meaningful only as a difference). Synthetic.

Re-pricing the baseline content gains under 1%. Re-contenting it gains 17% in expected profit while *raising* share and consumer surplus, because the plan stops forcing buyers to pay for features they do not value. In the optimized plan the base trim becomes an option-rich entry point, letting high-WTP buyers add exactly what they want; the mid trim carries nearly everything as standard with the sunroof optional; and the price ladder widens to \$27.2k, \$32.0k and \$34.6k.

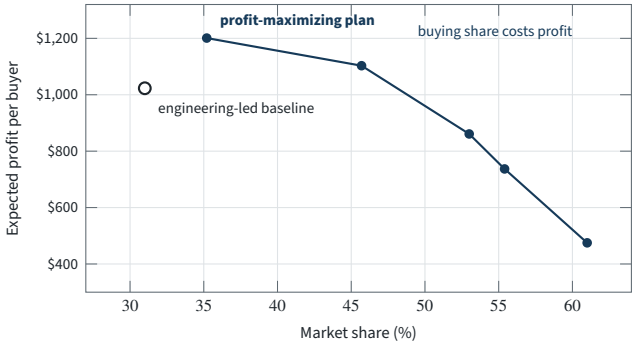


Exhibit 1. Efficient frontier of expected profit against share, each point a content-and-price plan found by the search with a weight on share. The baseline sits inside the frontier: plans exist with more profit and more share. Synthetic.

The frontier in Exhibit 1 also shows the cost of chasing share: moving from 35% to 46% gives up about \$100 per buyer, and beyond 55% profit falls steeply. Robustness comes from Exhibit 2. Re-evaluating both plans under 300 perturbed parameter vectors, with WTP means varied by about 15% and price sensitivity by 10%, the optimized plan beats the baseline in every case: the mean uplift is \$175 per buyer, the 5th percentile \$123, and the CVaR of the worst 10% \$120.

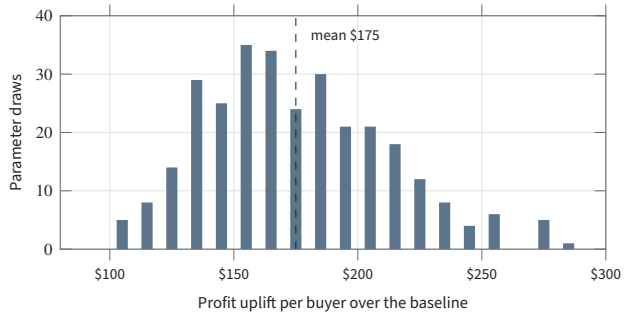


Exhibit 2. Distribution of the optimized plan's profit uplift over 300 posterior-style parameter draws, each evaluated on common random numbers. Every draw is positive. Synthetic.

6. Validation and implementation

Validate the model before the plan

A sound program validates in layers: holdout choice tasks for the choice model, back-casting of past model years (would the model have predicted the shares the market actually delivered after content changes?), and sensitivity of the recommended plan to the taste distribution's tails. A plan whose ranking flips under plausible parameter draws is a finding, not a recommendation.

Put it inside the vehicle development process

Content decisions are made by marketing, engineering, sourcing and finance together, at gates years before launch. The model earns its place by being available

at those gates as a scenario tool: what if a feature's cost rises, a competitor adds it as standard, or a regulation makes it mandatory? Finance should own the cost inputs, marketing the choice model, and engineering the feasibility rules, with every recommended plan reproducible from versioned inputs.

REFERENCE DECISION FLOW

Choice experiments → mixed logit, calibrated → Monte Carlo choices and uncertainty → content search with inner pricing → frontier and plan at the gate

7. What is known about the work

This note describes work led by Adam DeJans Jr. and John Brandon Elam in earlier industry roles at a global automaker, combining statistical choice modeling, probability, Monte Carlo simulation and optimization to inform vehicle content decisions. No financial outcomes are published here. For a sense of what this class of method can achieve at scale, General Motors has published a closely related system in which GM Finance verified $4.4 billion in incremental profit over the average product life cycle since 2018 (Wu-Smith et al., 2023); that figure is GM's and is not attributed to the work described here.

Limitations

Stated-preference data can misstate real purchase behavior, so calibration to market shares is essential and imperfect. Competitor responses are usually held fixed, which is optimistic when rivals react to content moves. Option take-up modeled as independent across features ignores package bundling and dealer stocking. And local search finds a good plan, not a certified optimum; restarts and alternative neighborhoods test how much is left on the table.

8. Lessons that transfer

Content is the lever, price is the fine-tuning

In the worked example, re-pricing fixed content gained under 1% while re-contenting gained 17%. Most of the value lies in what is offered where.

Heterogeneity is the reason to optimize

Options exist because buyers differ. A model that averages tastes away cannot find the plans that let buyers sort themselves.

Report a distribution and a frontier

Leaders need the probability a plan beats today's line and the price of share, not a single expected profit.

A USEFUL EXECUTIVE DIAGNOSTIC

For the features on your entry trim, which are standard because customers value them more than they cost, and which are standard because they always have been?

Further reading

Train, K. E. *Discrete Choice Methods with Simulation*, 2nd ed. Cambridge University Press, 2009.

McFadden, D. and Train, K. Mixed MNL models for discrete response. *Journal of Applied Econometrics* 15(5), 2000.

Rossi, P. E., Allenby, G. M. and McCulloch, R. *Bayesian Statistics and Marketing*. Wiley, 2005.

Wu-Smith, P. et al. General Motors optimizes vehicle content for customer value and profitability. *INFORMS Journal on Applied Analytics* 53(1), 2023.

About the case and Bit Bros

This paper draws on vehicle content work led by **Adam DeJans Jr.** and **John Brandon Elam** in earlier industry roles at a global automaker. It was not a Bit Bros client engagement. The automaker, its program, models, data and results are not disclosed, and the worked example is synthetic. Bit Bros is a boutique applied economics and optimization firm that helps operating teams price their decisions and build the systems that make them repeatable.



Adam DeJans Jr.
Technical Managing Director, Bit Bros LLC
Decision scientist specializing in optimization, simulation and decisions under uncertainty.



John Brandon Elam
Managing Director, Bit Bros LLC
Leads client engagements, working with executives to find the decisions that carry the most value.

Price the feature. Then decide where it goes.