

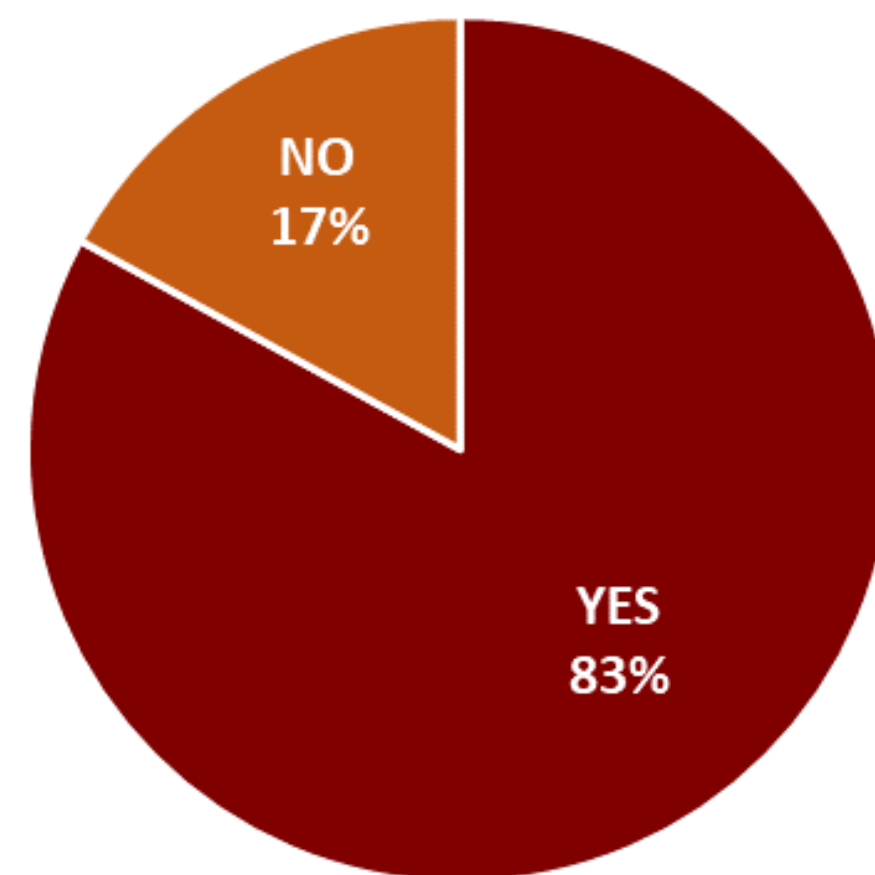
Marketing Analytics in the Age of AI

Jean-Pierre Dubé
MSI
March 2025

Most CEOs value marketing

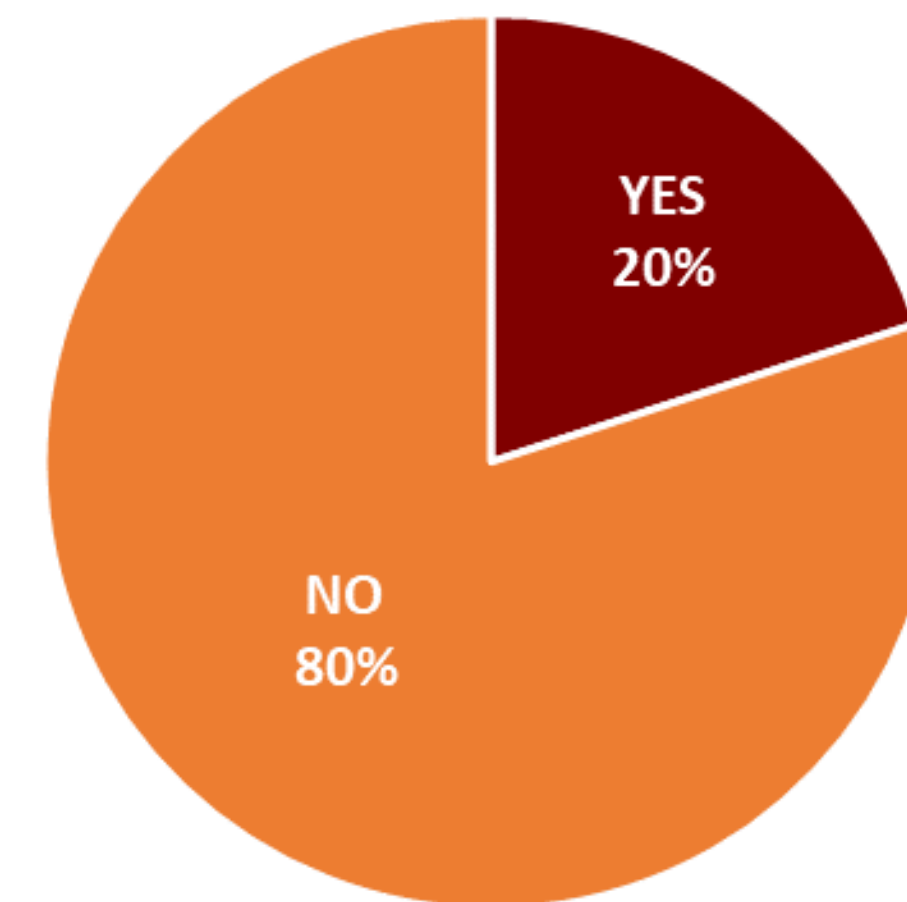
But most don't trust their CMO

CEOs: Does marketing drive growth?



83% of CEOs believe marketing is key to driving growth

CEOs: Do you trust your CMO?



Only 20% of CEOs trust their CMO (90% trust their CFO)

Source: *Harvard Business Review*. ["Why CMOs Never Last,"](#) July-August 2017

Old-School marketing sometimes feels like ...



Trust me. I've
been doing
this for 20
years...

The Chicago-Booth Approach to Marketing



Scientific approach to marketing decisions

- 1 What is your **theory**?
- 2 What is your **evidence**?
 - Which data/method(s)?
- 3 Was it good for business (**accountability/evaluation**)?



Theory



Evidence



Evaluation

Where does AI fit into this **Scientific Approach**?



Ability of machine to perform tasks usually associated with human intelligence.

e.g., learning, reasoning, & making decisions

Currently we really only have **Narrow AI** which is effective at specific tasks.

So what exactly is AI?

“It’s also a problem that “A.I.” has no fixed definition”

“As a computer scientist, I don’t like the term ‘A.I.’ In fact, I think it’s misleading—maybe even a little dangerous.”

“...think of A.I. as a **tool, not a creature”**



Jaron Lanier
Tech Guru & Prime Unifying Scientist at
Microsoft’s Office of the Chief Technology
Officer

Need to set more realistic expectations for business uses

Some definitions

Machine Learning: subfield of AI that uses algorithms to teach computers to learn and improve from data (especially deep learning using vast amounts of **data** to find patterns).

Large Language Models: type of AI using deep learning algorithms to understand, process and generate human language. Trained on **large amounts of text data** (e.g., books, articles, code ...).

Generative AI: broader concept of AI systems that can generate new content, such as text, audio, video and 3D models. Uses ML trained on various data. e.g., ChatGPT built on foundation of an **LLM**.

→ **STILL EXPERIMENTAL!**

Let's look at some marketing applications to flesh these ideas out...

9

Starts with data... and AI as the analyst

example: Price Optimization



Pricing at Ziprecruiter.com

- Academic collaboration with Ziprecruiter.com
- Fastest growing HR company: 75,000 employers per month
- Online matching platform for recruiters and applicants
- High conversion rates and high quality matches
- August 2015: charging \$99 per month to “starter” businesses

... but why \$99?

Pricing at Ziprecruiter.com

- How do you evaluate the \$99 price?
- How do you measure **demand** for a B2B digital service?
- **method:**

We ran a B2B price experiment to measure **demand** and determine the optimal price at Ziprecruiter...

The Ziprecruiter Price Experiment

Method

10 prices tested:

{ \$19, \$39, \$59, \$79, \$99, \$159, \$199, \$249, \$299, \$399 }

All new customers in September 2015

- Customer had to register and advance to paywall
- Each customer assigned randomly to a price
- Track whether customer submitted credit card information and was billed for a month of service

Source: [Dubé and Misra \(2023\)](#)

The Ziprecruiter Price Experiment

Data

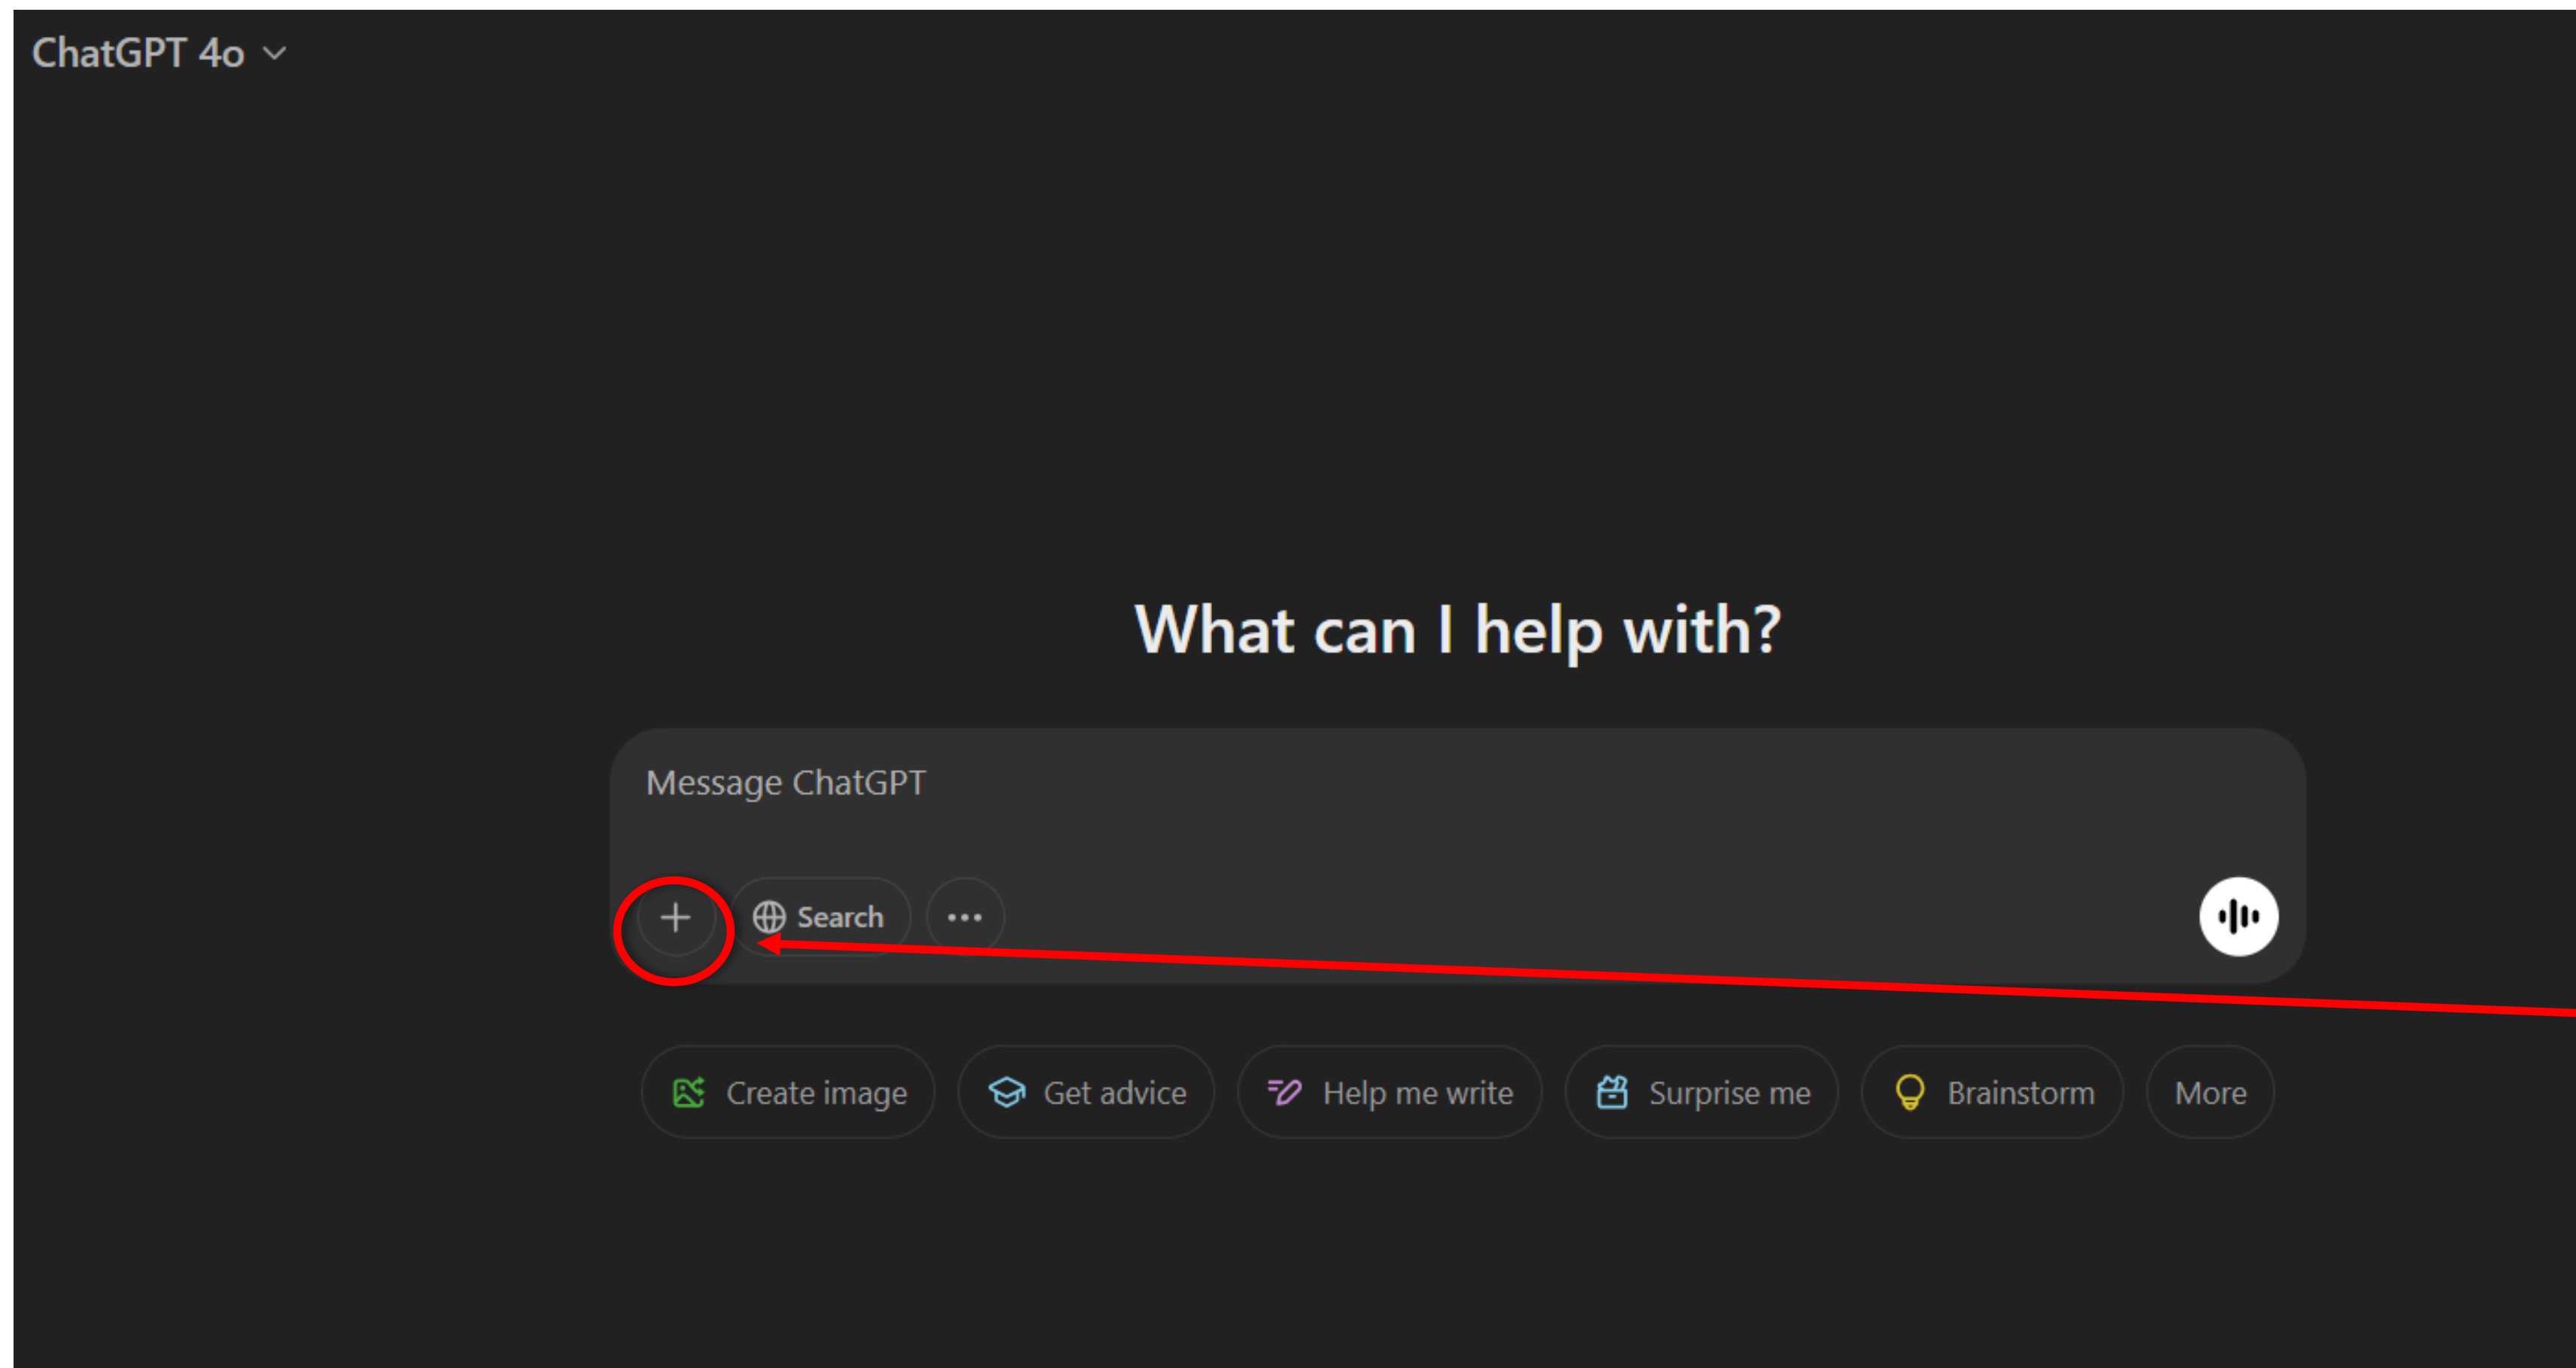
Can AI handle the analytics?

Experimental Cells	Price Charged	# firms	Conversion Rate
Test 1	\$19	286	0.36
Test 2	\$39	260	0.32
Test 3	\$59	214	0.27
Test 4	\$79	246	0.29
Control	\$99	192	0.24
Test 5	\$159	153	0.2
Test 6	\$199	140	0.18
Test 7	\$249	127	0.17
Test 8	\$299	99	0.13
Test 9	\$399	80	0.11

Source: [Dubé and Misra \(2023\)](#)

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...



Upload the data

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...

using the attached data, what is the optimal price to charge

To determine the optimal price to charge using the provided data, I'll follow these steps:

1. Inspect the data to understand its structure.
2. Identify relevant variables (e.g., price, demand, revenue, cost, etc.).
3. Analyze price elasticity of demand if applicable.
4. Find the revenue-maximizing and profit-maximizing price using statistical modeling.

Let me first examine the data and summarize its contents.

The dataset contains the following key columns:

- Experimental Cells: Identifies different pricing tests (e.g., Test 1, Test 2, etc.).
- Price Charged: The price per unit in each test.
- # firms: The number of firms (potential customers).
- Conversion Rate: The proportion of firms that made a purchase.

Next Steps:

To determine the optimal price, I will:

1. Calculate demand (number of customers converted) using $\# \text{ firms} * \text{Conversion Rate}$.
2. Compute total revenue at each price ($\text{Revenue} = \text{Price} * \text{Demand}$).
3. Identify the price that maximizes revenue.

Let me perform this analysis now. [~]

Pricing Analysis				
	Experimental Cells	Price Charged	# firms	Convi
1	Test 1	19	286	0.36
2	Test 2	39	260	0.32
3	Test 3	59	214	0.27
4	Test 4	79	246	0.29
5	Control	99	192	0.24

The optimal price to charge based on revenue maximization is \$79, which results in the highest total revenue of \$5,635.86.

Let me know if you need further analysis, such as profit maximization (if cost data is available) or price elasticity insights. [~]

using the attached data, what is the optimal price to charge

(my prompt)

VERY WRONG!

- AI confused how to model demand
- human can easily reason this is incorrect but AI cannot reason!

The optimal price to charge based on revenue maximization is \$79, which results in the highest total revenue of \$5,635.86.

Let me know if you need further analysis, such as profit maximization (if cost data is available) or price elasticity insights. [~]

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...

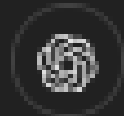
Ok so we need to walk ChatGPT through the micro

- 1. help it figure out demand**
- 2. set price using the “correct” demand**

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...

use the attached data to estimate a logistic regression. Plot the predicted conversion



I'll start by inspecting the contents of the Data.xlsx file to understand its structure and identify the variables that can be used for logistic regression. Then, I'll estimate the logistic regression model and plot the predicted conversion probabilities.

The dataset contains the following columns:

- **Experimental Cells:** Labels for different test groups, including a control group.
- **Price Charged:** The price associated with each experimental cell.
- **# firms:** The number of firms in each experimental cell.
- **Conversion Rate:** The proportion of firms that converted.

I will estimate a logistic regression model using Price Charged as the independent variable and Conversion Rate as the dependent variable. Then, I'll plot the predicted conversion probabilities.

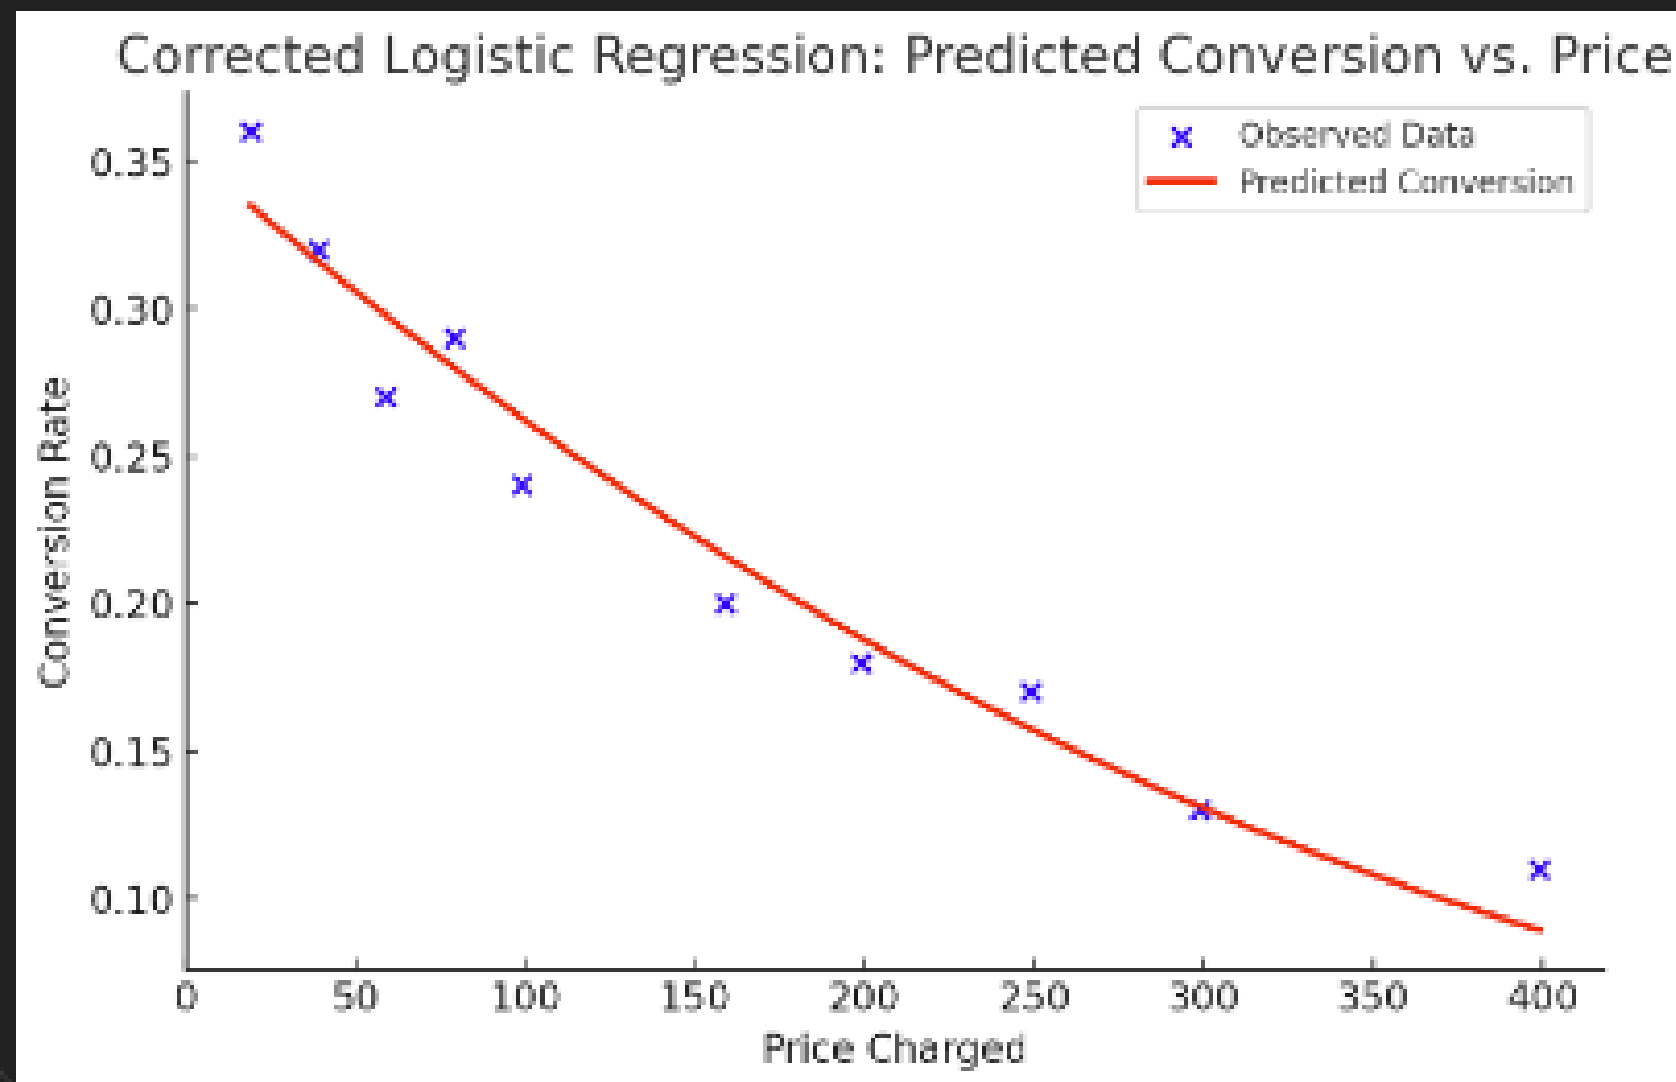
use the attached data to plot a logistic demand curve

(my new prompt)

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...

Corrected Logistic Regression: Predicted Conversion Vs. Price



The corrected plot now accurately represents the logistic regression model. The blue scatter points represent the observed conversion rates for different prices, while the red curve represents the predicted conversion probabilities from the logistic regression model.

This should now be correctly fitted. Let me know if you need further refinements or additional insights! [:-]



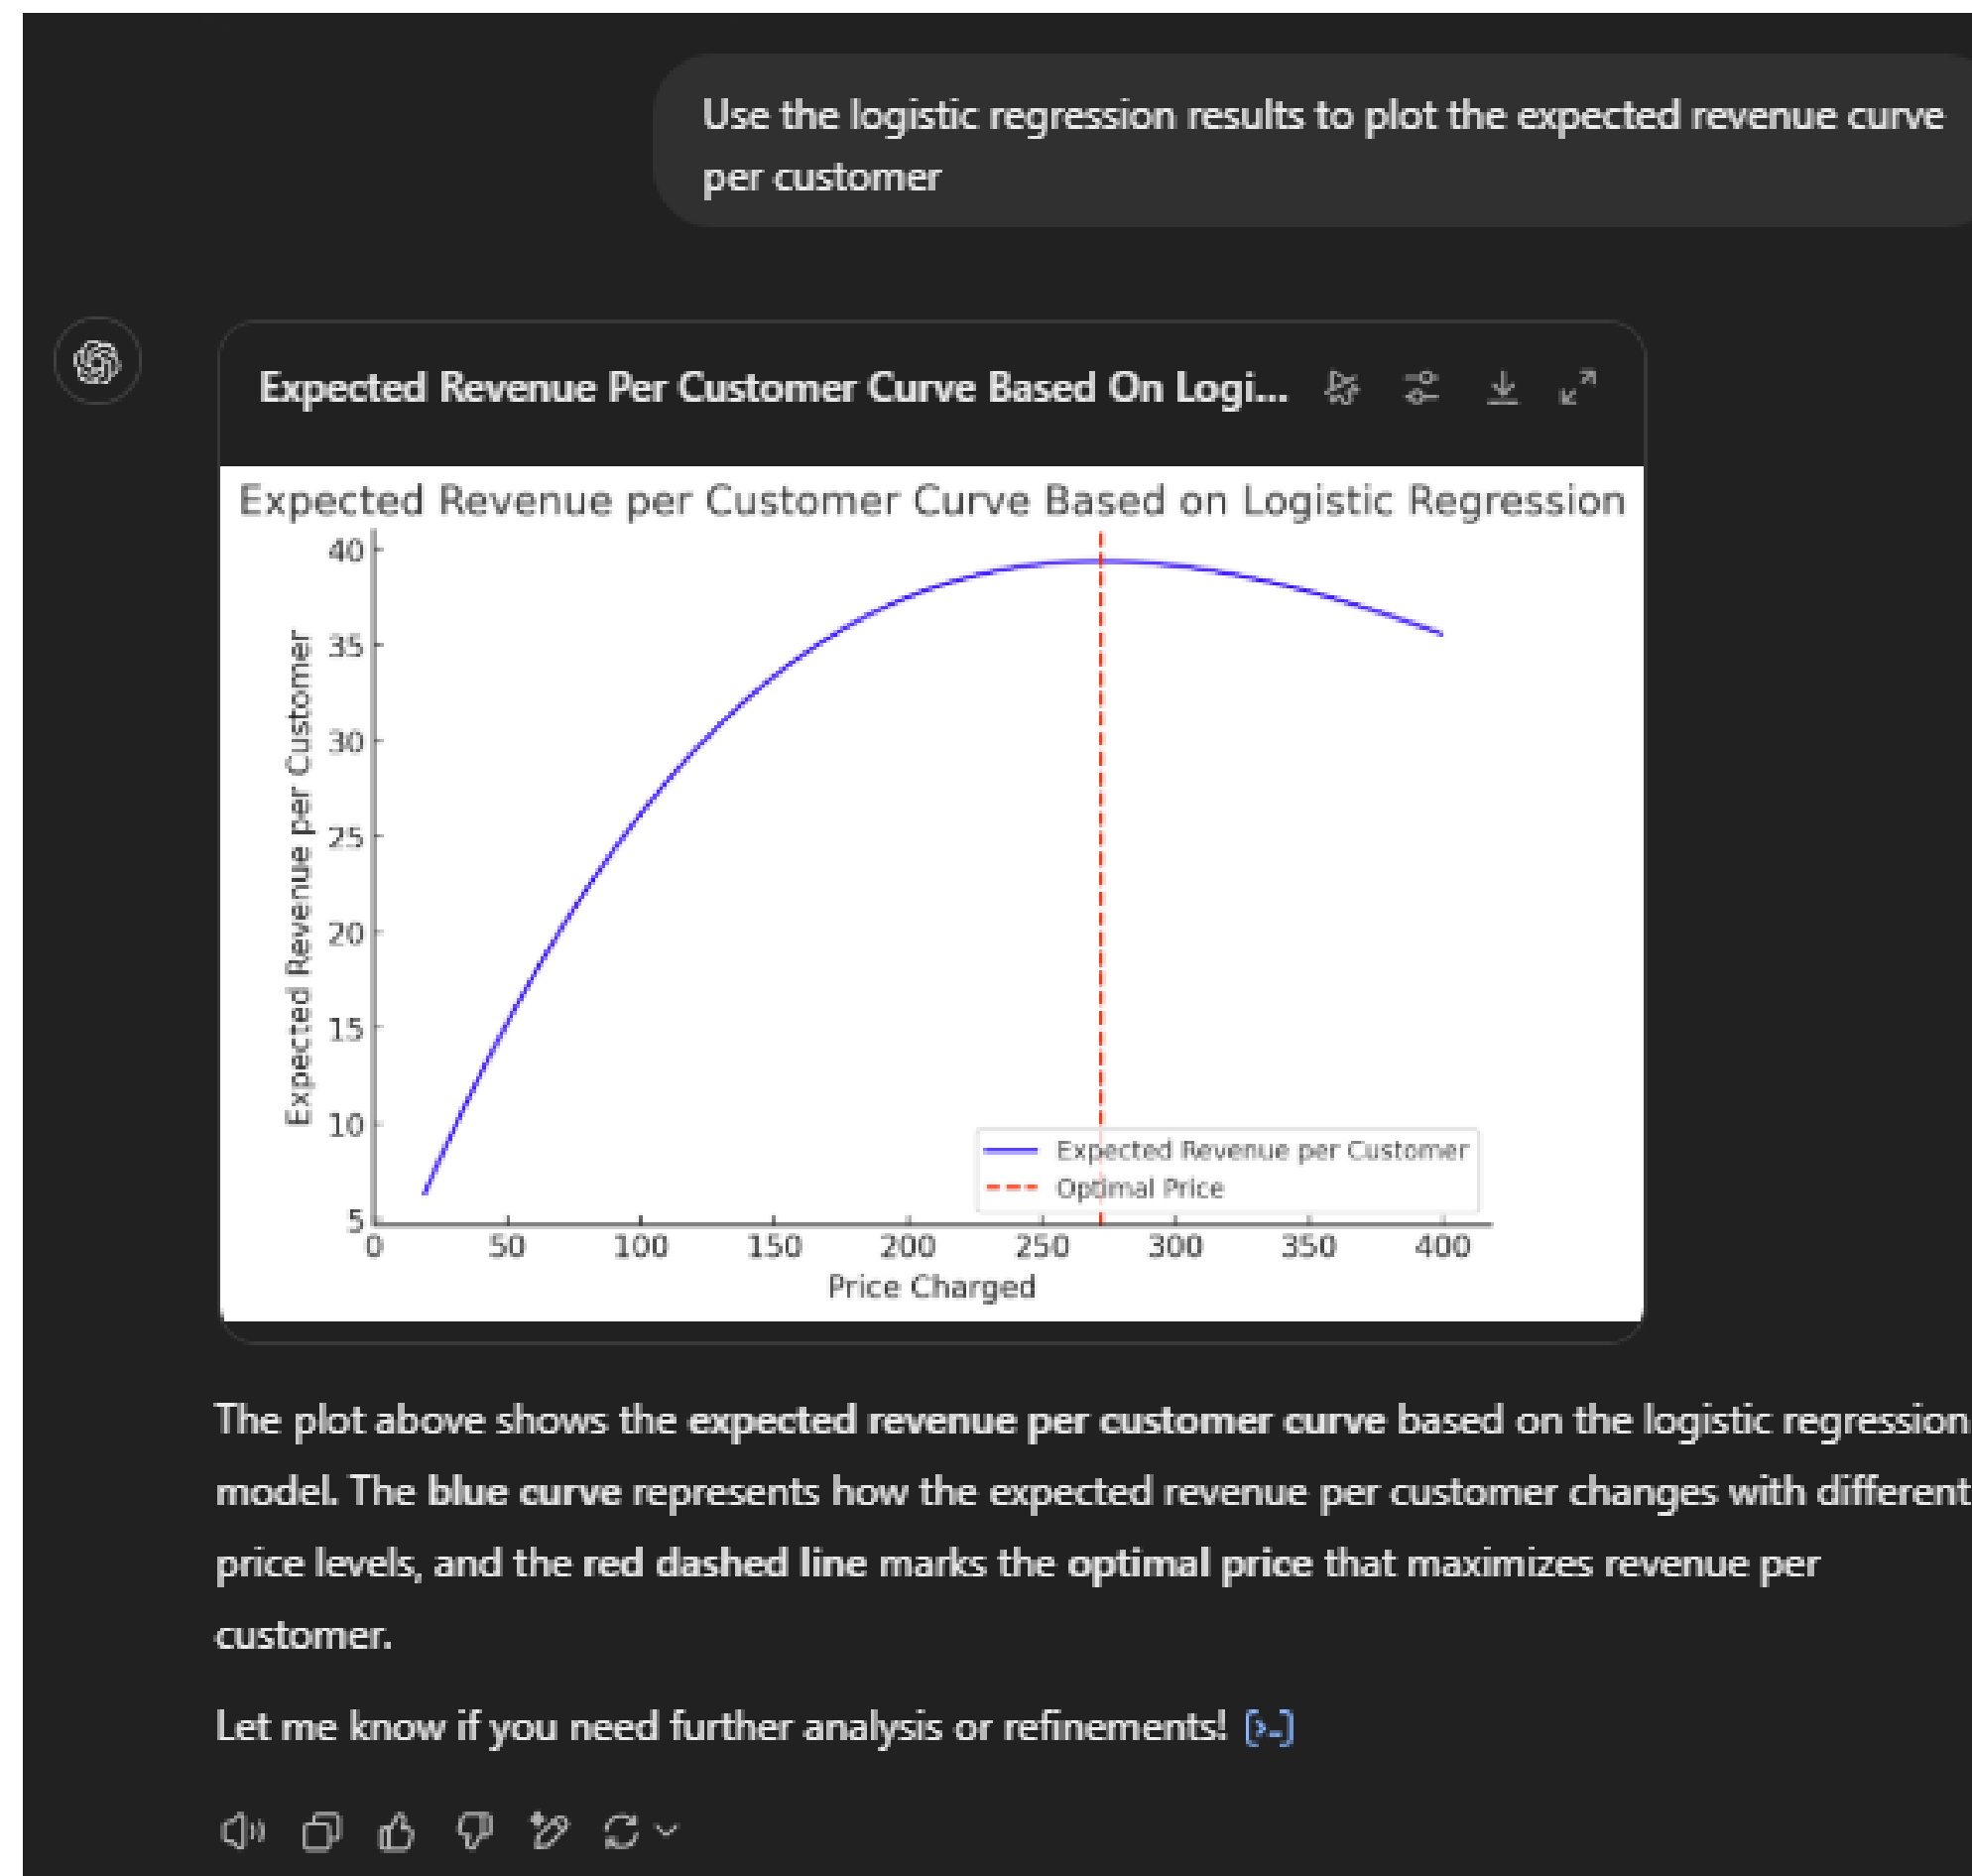
use the attached data to plot a logistic demand curve

(my new prompt)

Demand!!!

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...



Use the logistic regression results to plot the expected revenue curve per customer

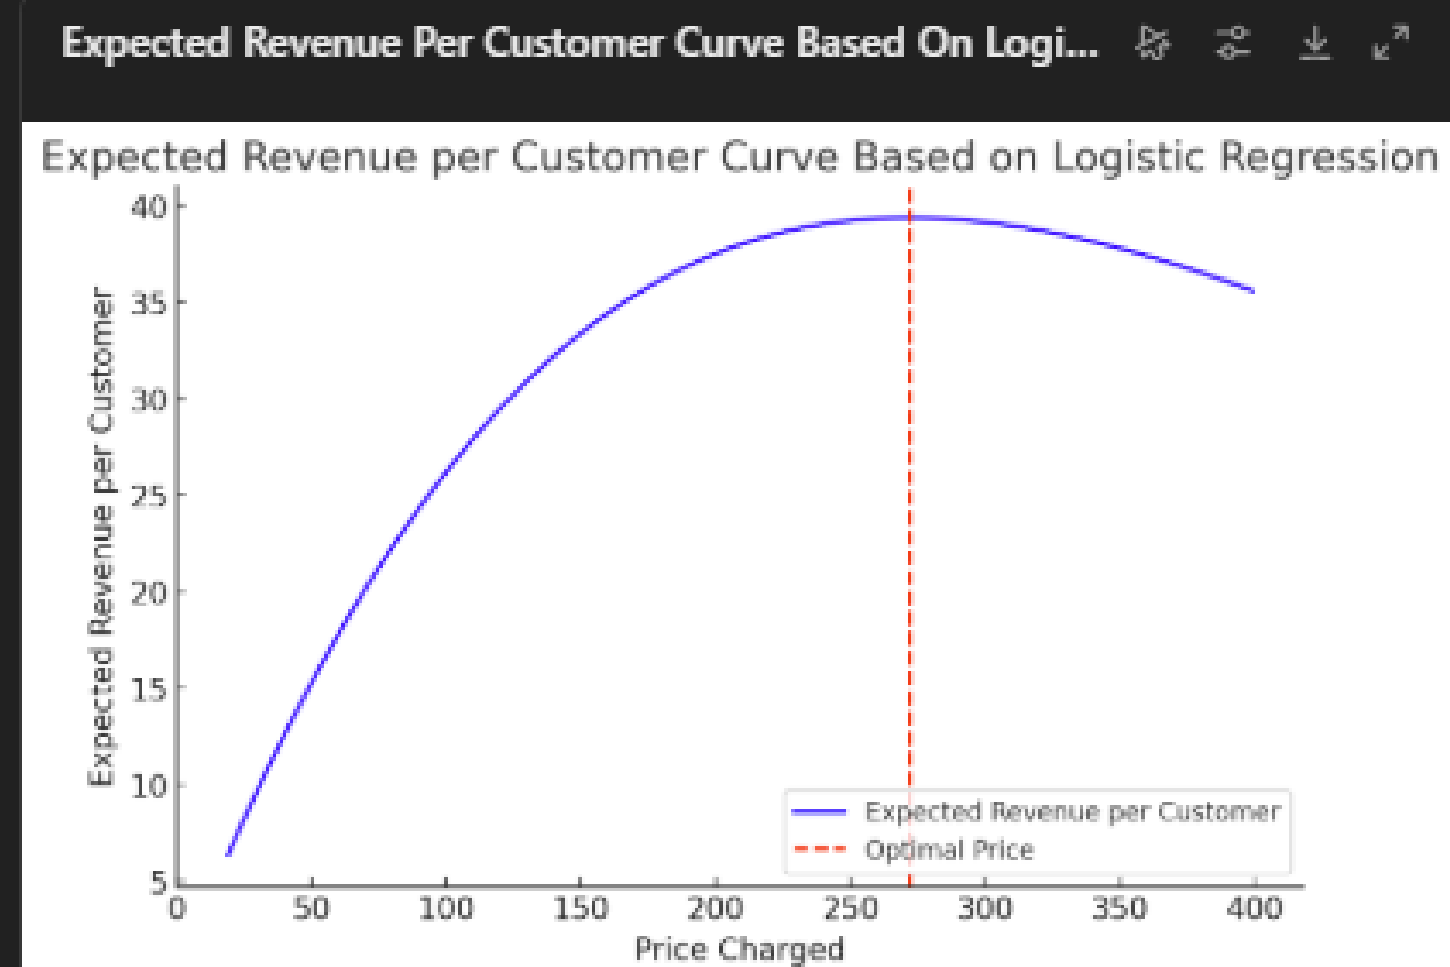
(update the prompt)

what's the optimal price

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...

Use the logistic regression results to plot the expected revenue curve per customer



The plot above shows the expected revenue per customer curve based on the logistic regression model. The blue curve represents how the expected revenue per customer changes with different price levels, and the red dashed line marks the optimal price that maximizes revenue per customer.

Let me know if you need further analysis or refinements! [?]

[?]

Use the logistic regression results to plot the expected revenue curve per customer

what's the optimal price

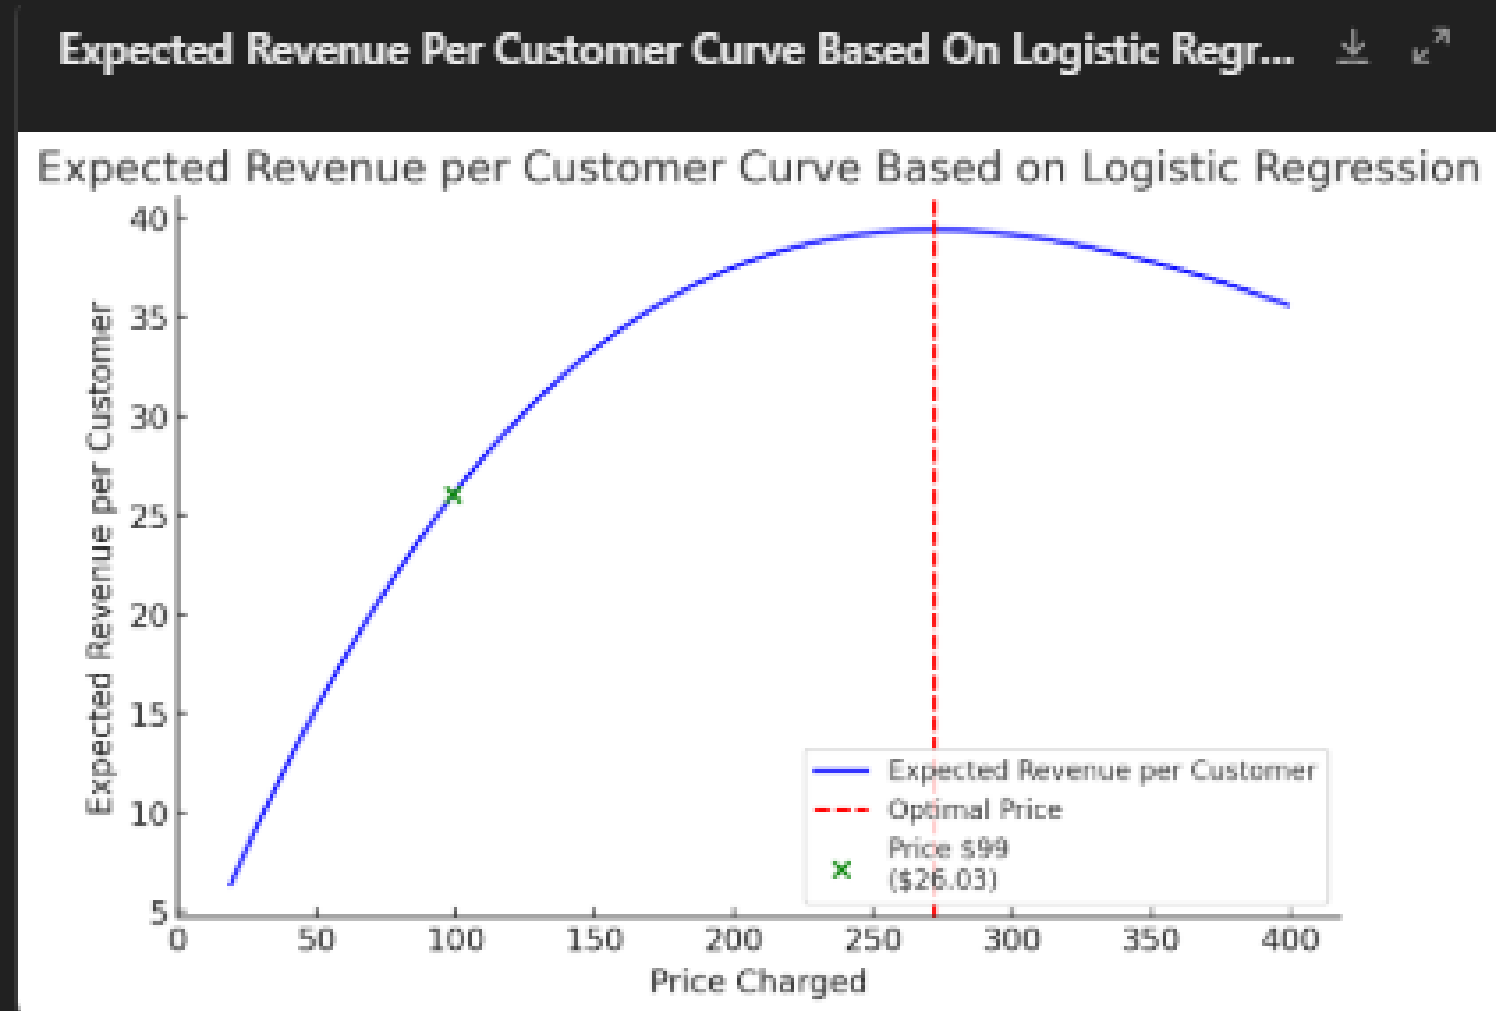
The optimal price per customer that maximizes expected revenue per customer is \$272.33. Let me know if you need further analysis or adjustments! [?]

[?]

The Ziprecruiter Price Experiment

Ask ChatGPT- 4o for help...

what is the expected revenue per customer if the price is \$99? Show this on the graph above



The expected revenue per customer at a price of \$99 is \$26.03. This value is now correctly plotted on the graph with a green marker.

Let me know if you need further clarifications or refinements! [:-]



what is the expected revenue per customer if the price is \$99? Show this on the graph above

The Ziprecruiter Price Experiment

Conclusions using ChatGPT- 4o as the Analyst

Oops ...

under-pricing by over 60% (leaving a lot of money on the table)

- Insight used basic prompt-based AI (Generative AI ... i.e., LLM)
- But
 - required simple experiment to get the **right data**
 - required manager to guide ChatGPT through the microeconomic modeling of demand

Opportunity for AI to generate additional monetization from these data?

23

Let's sprinkle on more Machine Learning...

example: Price Optimization continued

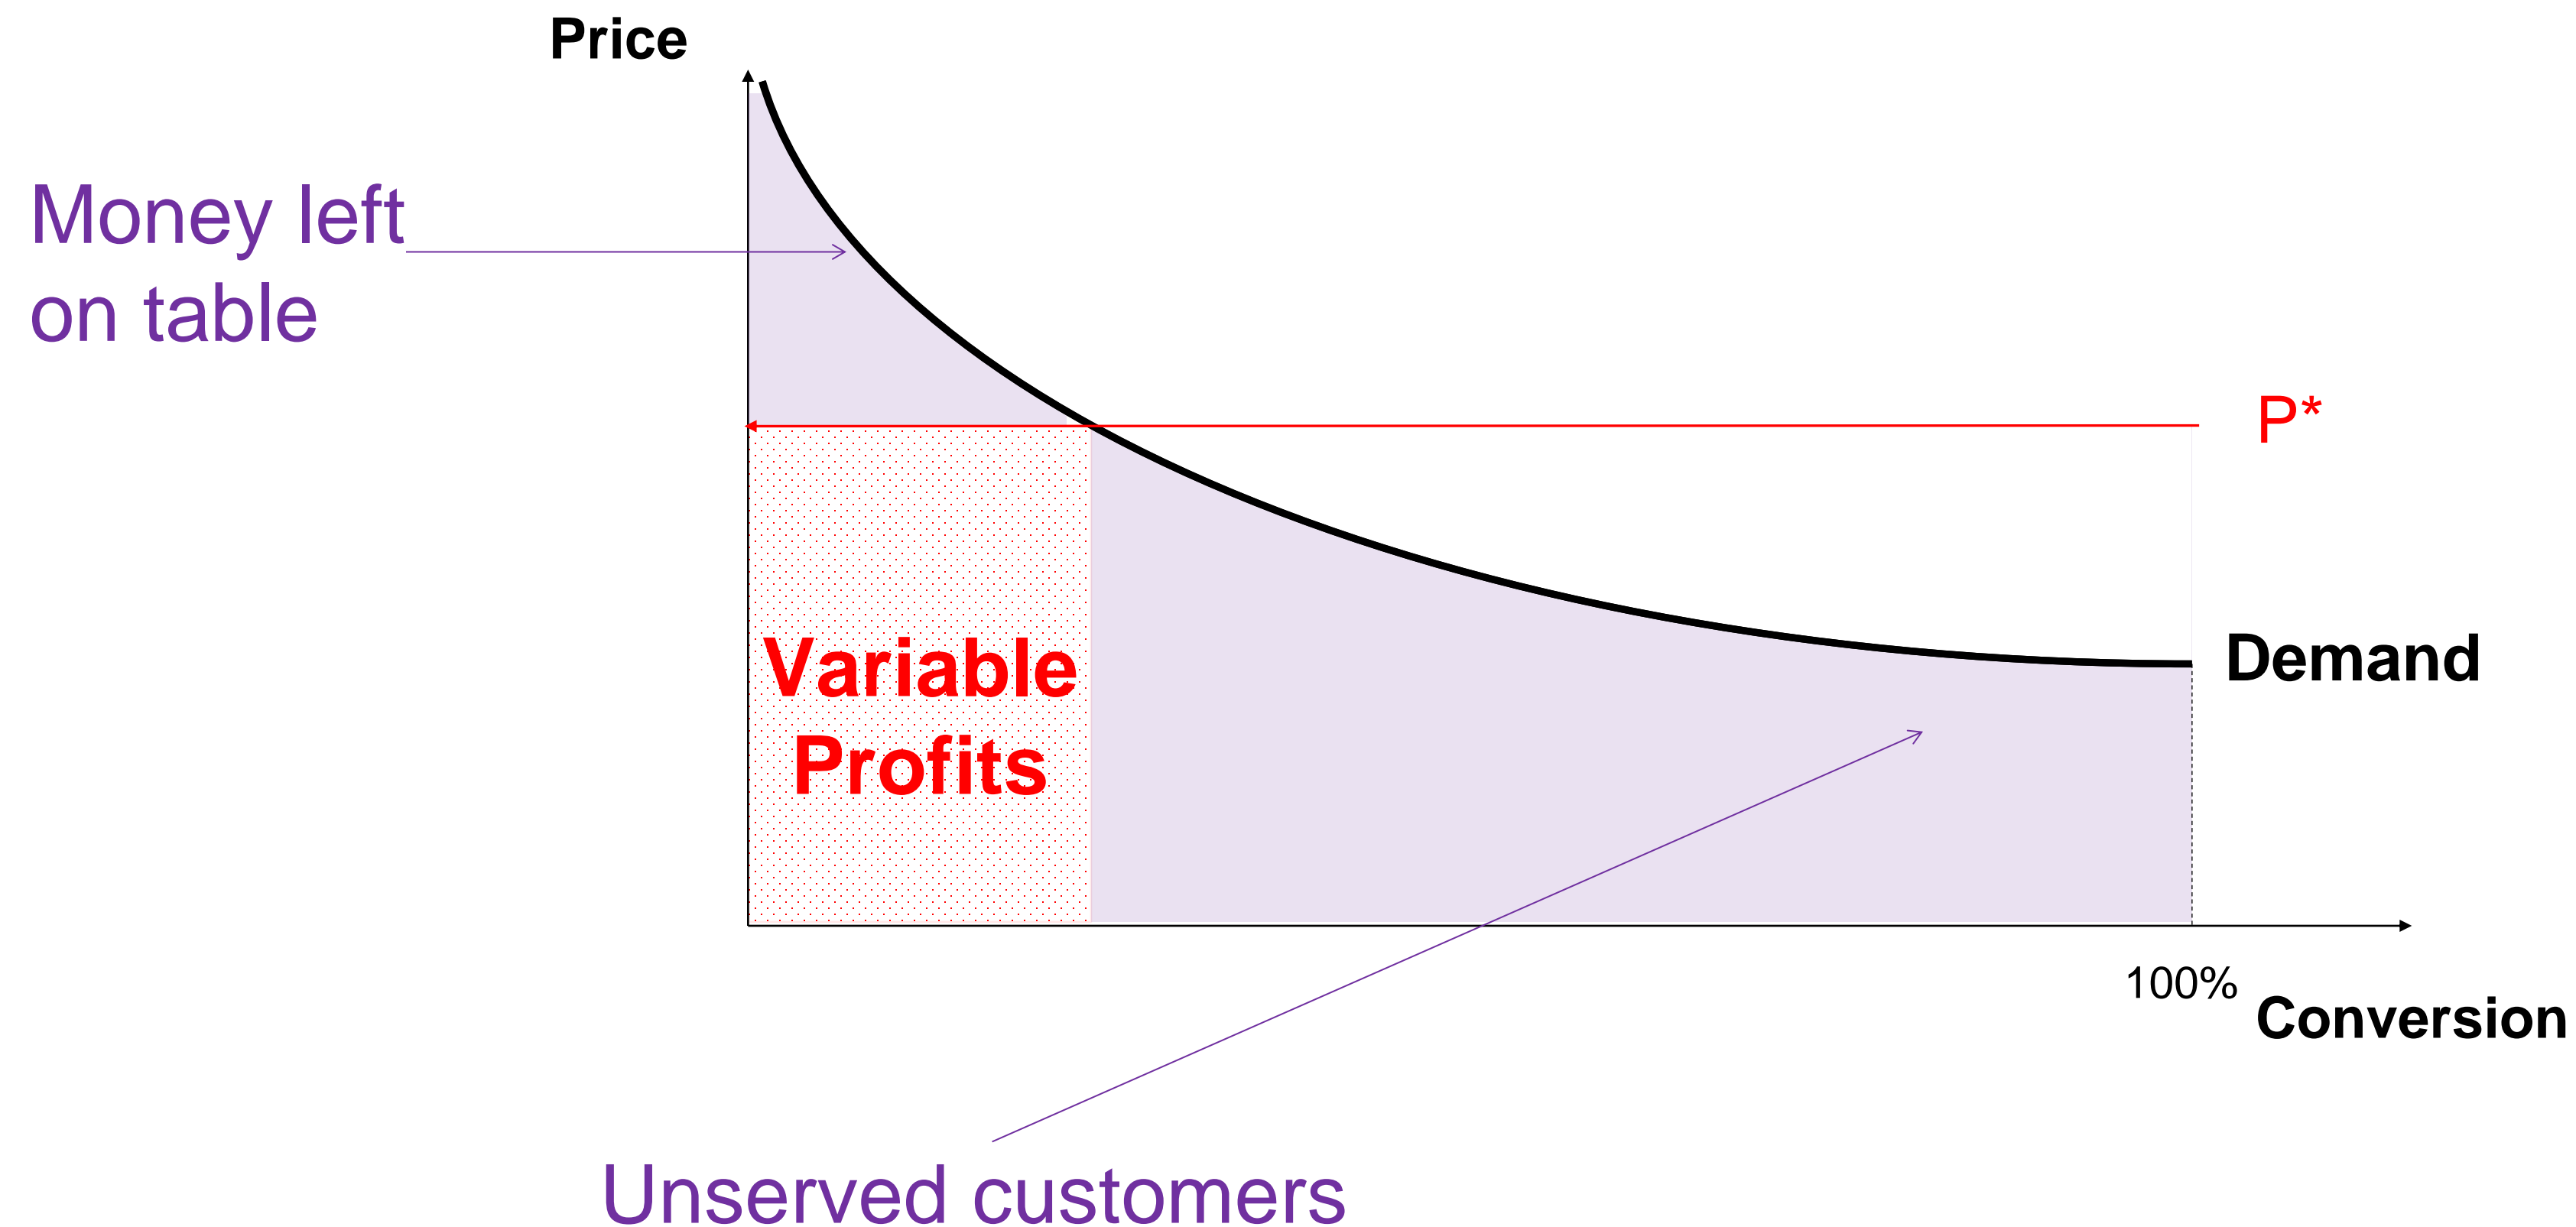


First revisit the underlying microeconomics



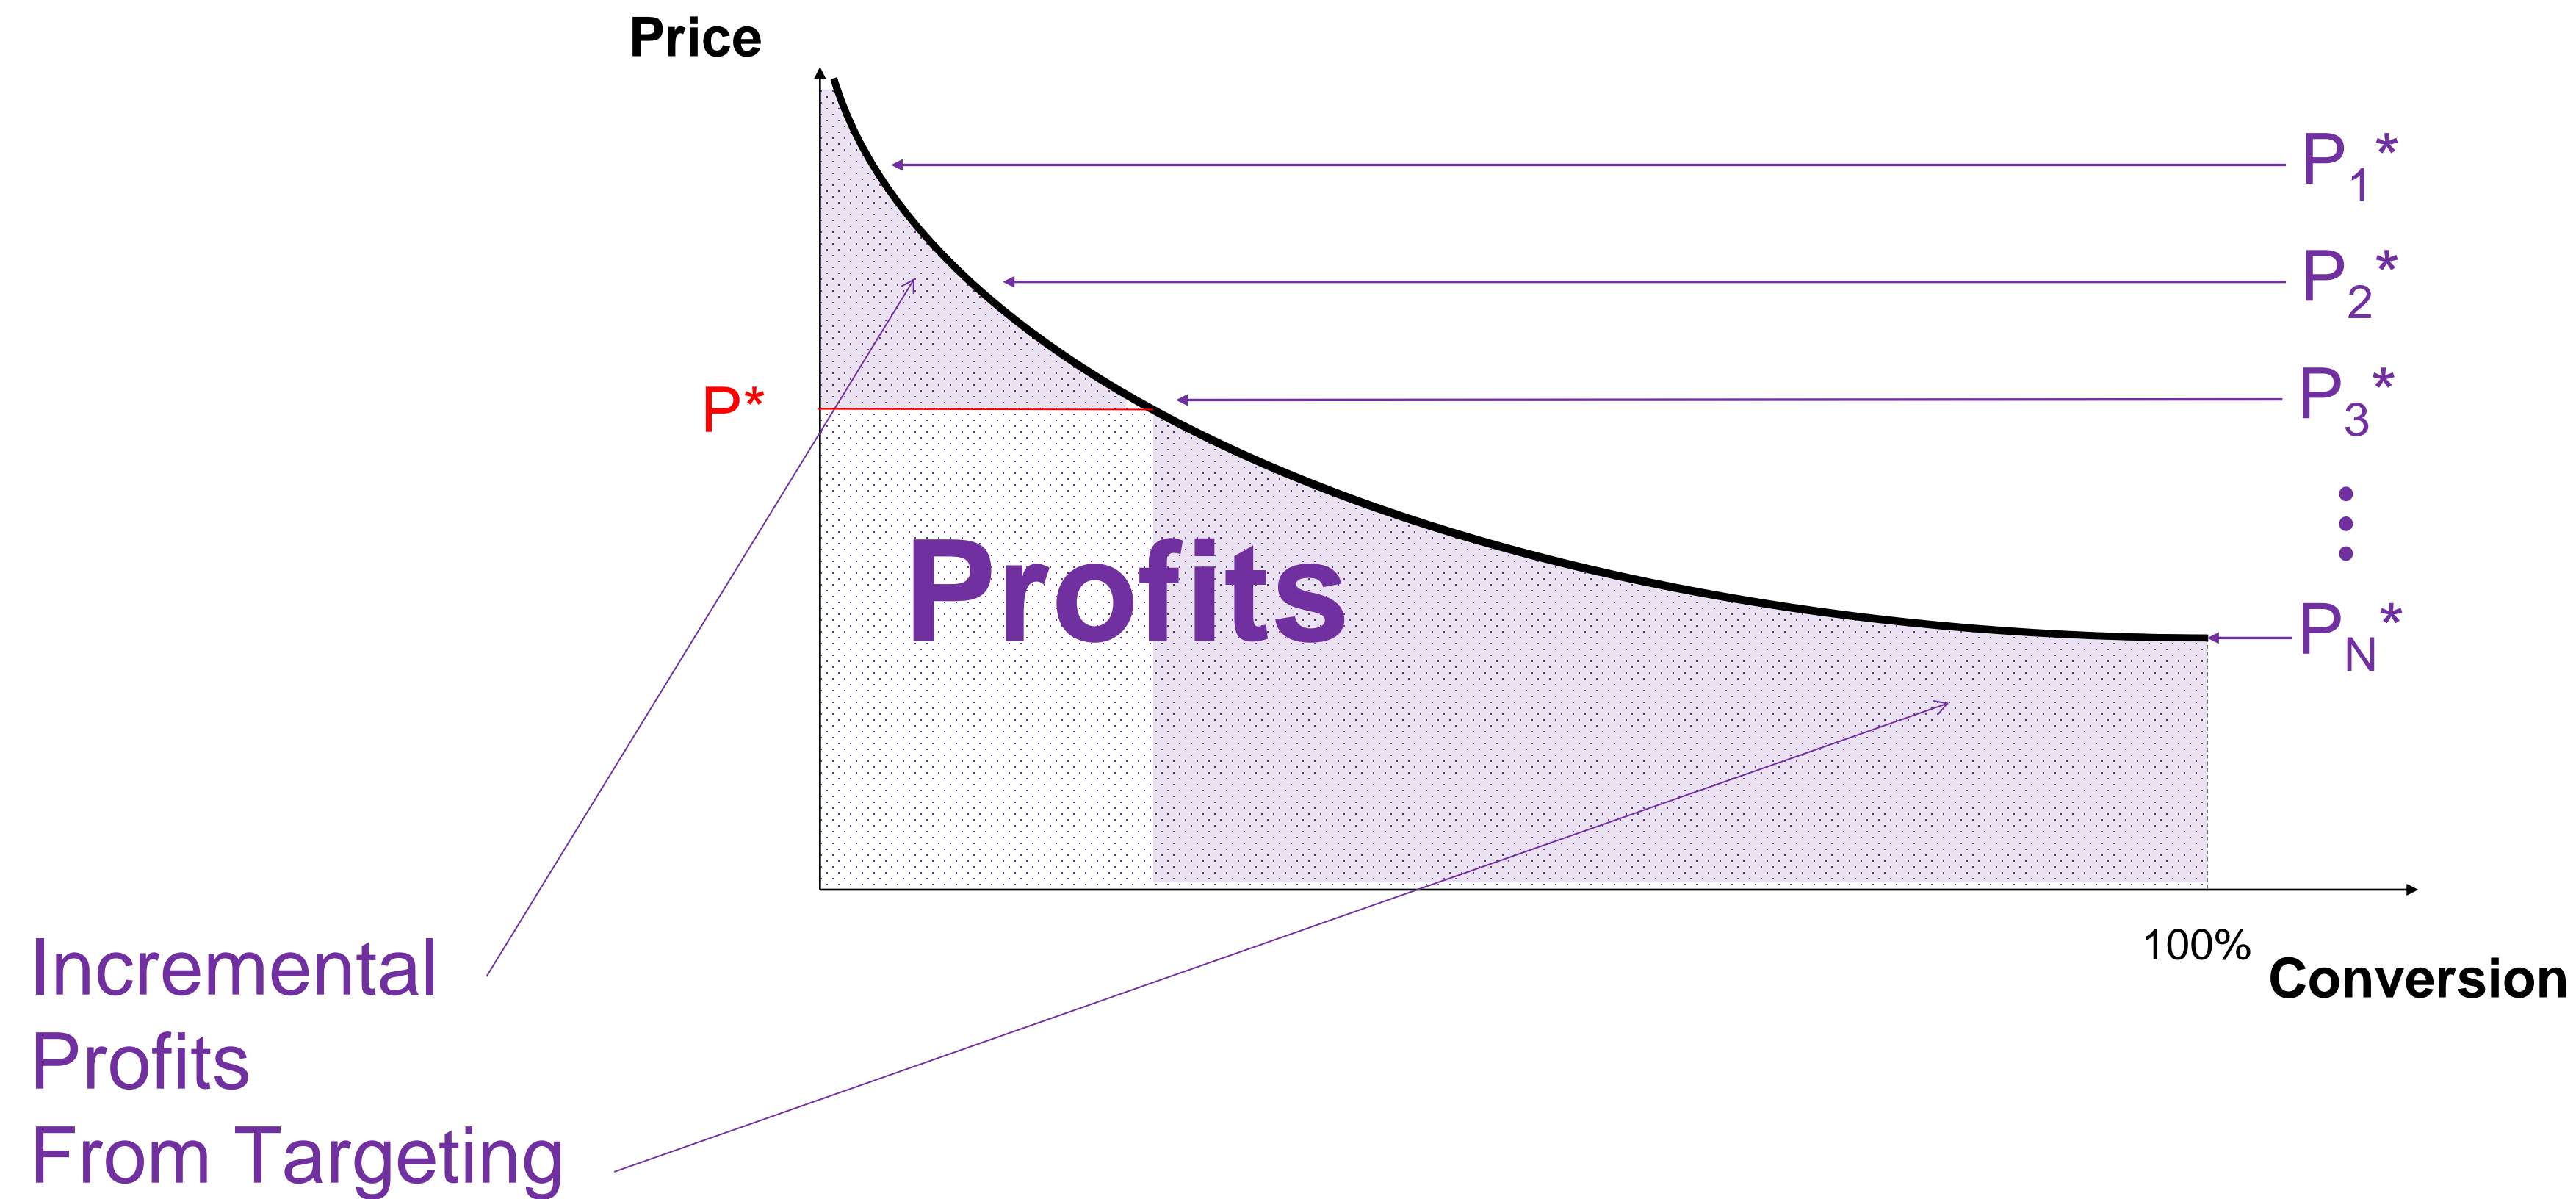
Optimal Uniform Price

Charge customers same price for same product



Personalized Pricing

Suppose we could identify customers/segments



Monetizing Customer Data

Customer data Ziprecruiter observes at registration stage

Company/Job Variable Names
Job state
Company type
Has commissions
Company declared job slots needed
Job total benefits
Employment type
Resume required
Job medical benefit
Job vision benefit
Job life insurance benefit
Job category

Can Ziprecruiter **monetize** these data?

Scalable Personalized Pricing

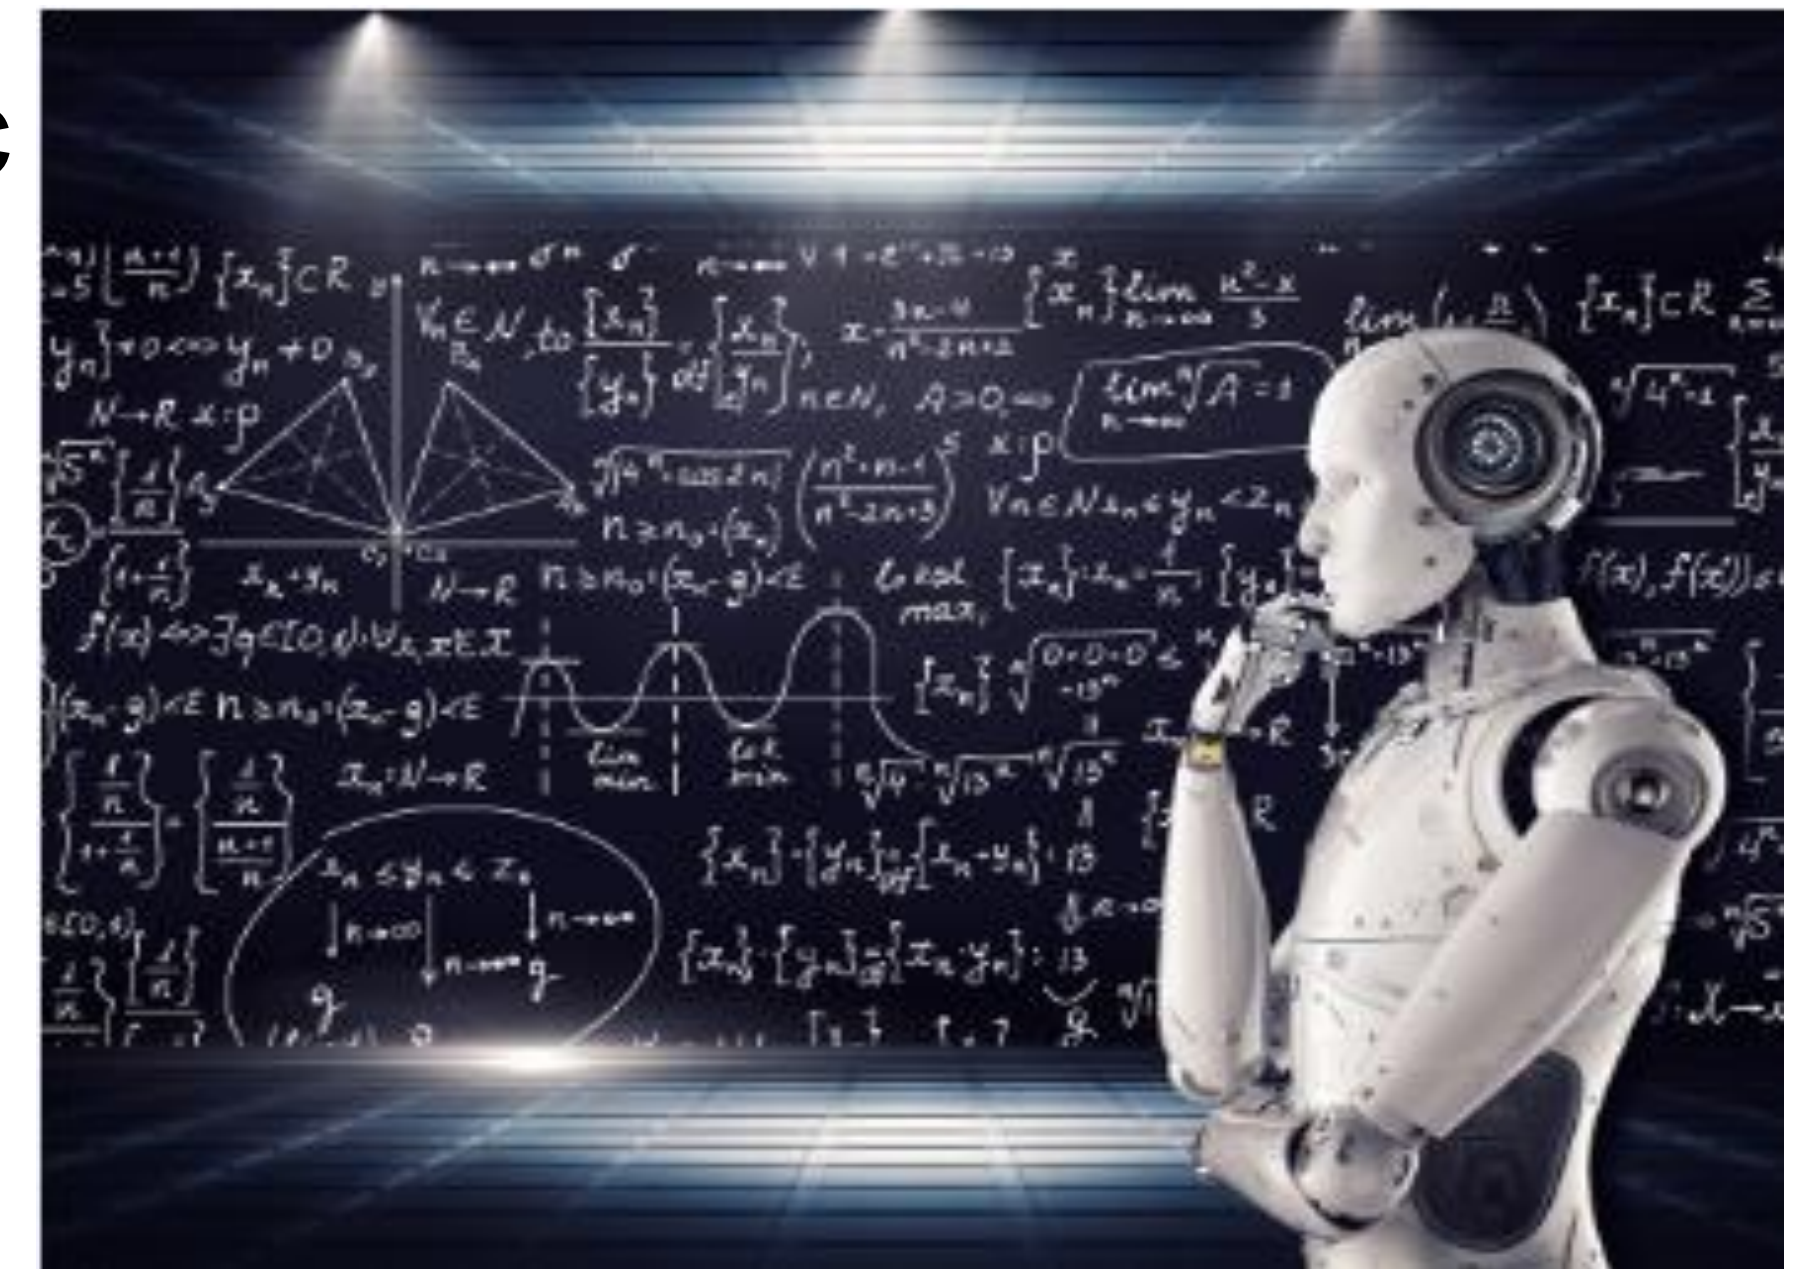
Pricing as a **Machine Learning** problem

Thousands of potentially targetable *features*

- As many features as observations!

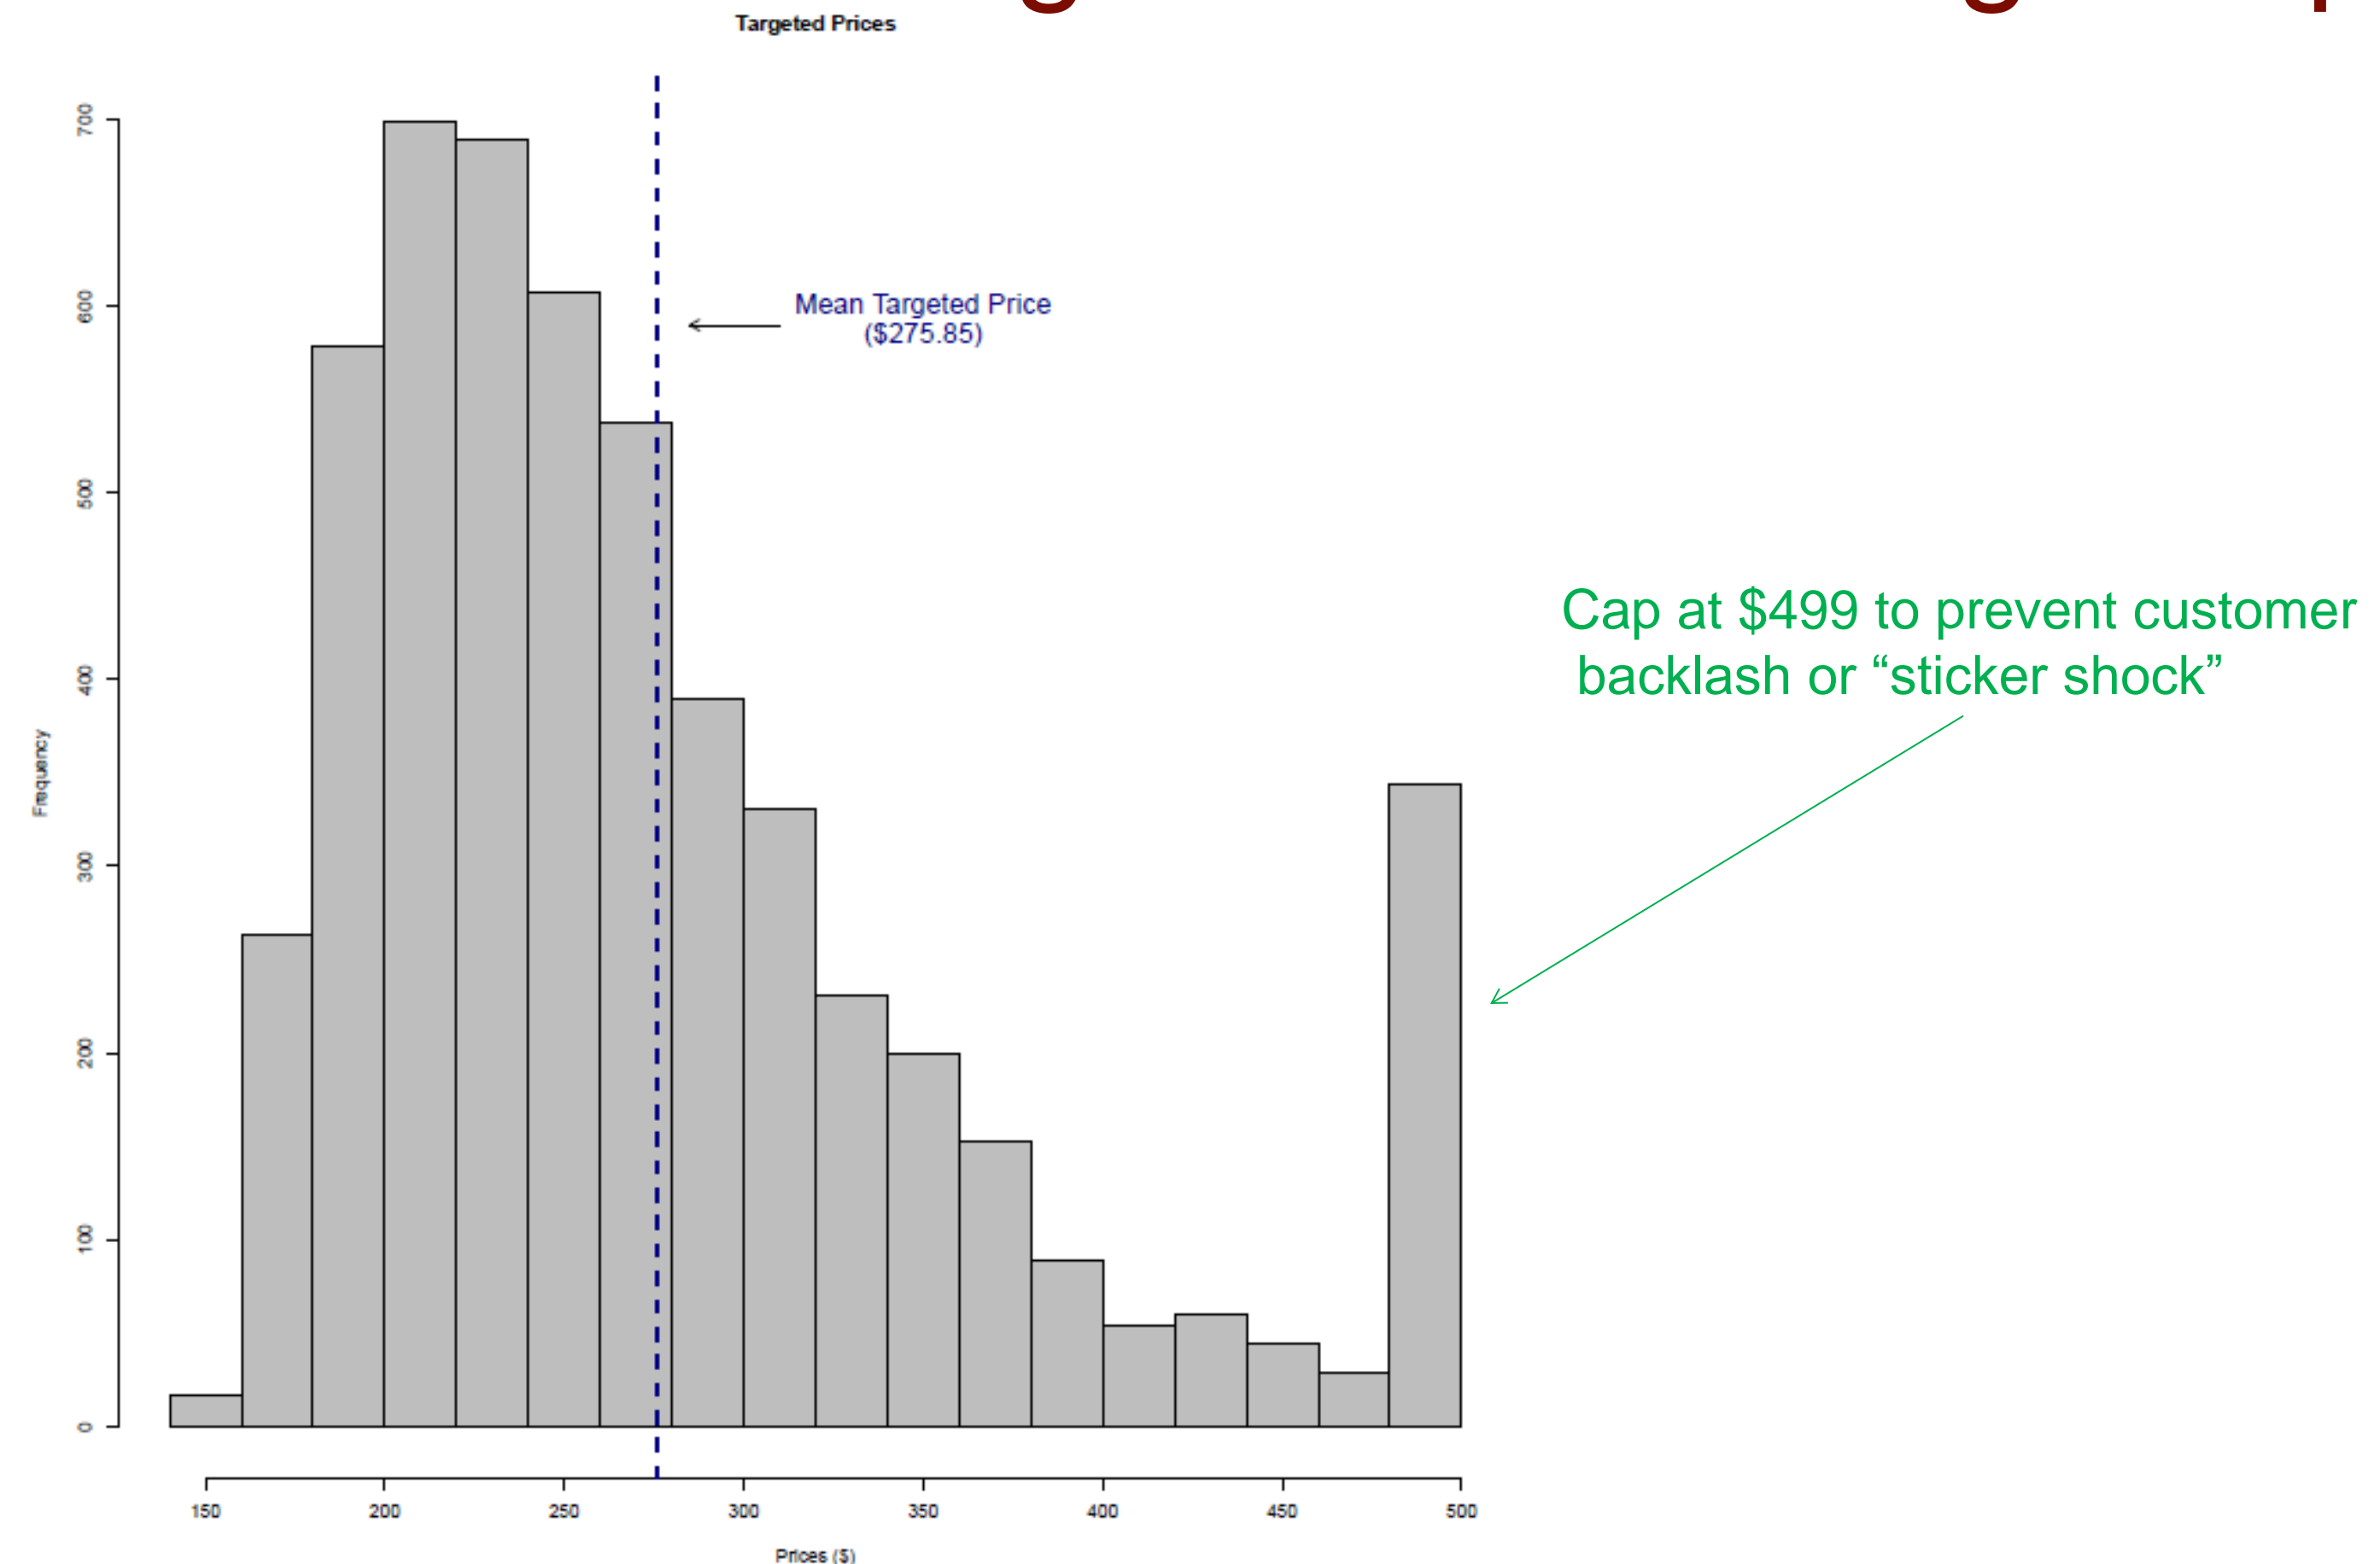
Apply **Deep Learning Algorithm** to our logistic demand

- Combine statistics and machine learning
- Link *willingness-to-pay* to features



Scalable Personalized Pricing

Simulated Personalized Prices for August 2015 training sample



Source: [Dubé and Misra \(2023\)](#)

Distribution of Personalized Prices (N= 7,867)

Scalable Personalized Pricing

Projected Revenues



Pricing	Projected Profit Per Lead
Base (\$99)	\$25
Uniform (\$249)	\$40
Personalized	\$44

Source: [Dubé and Misra \(2023\)](#)

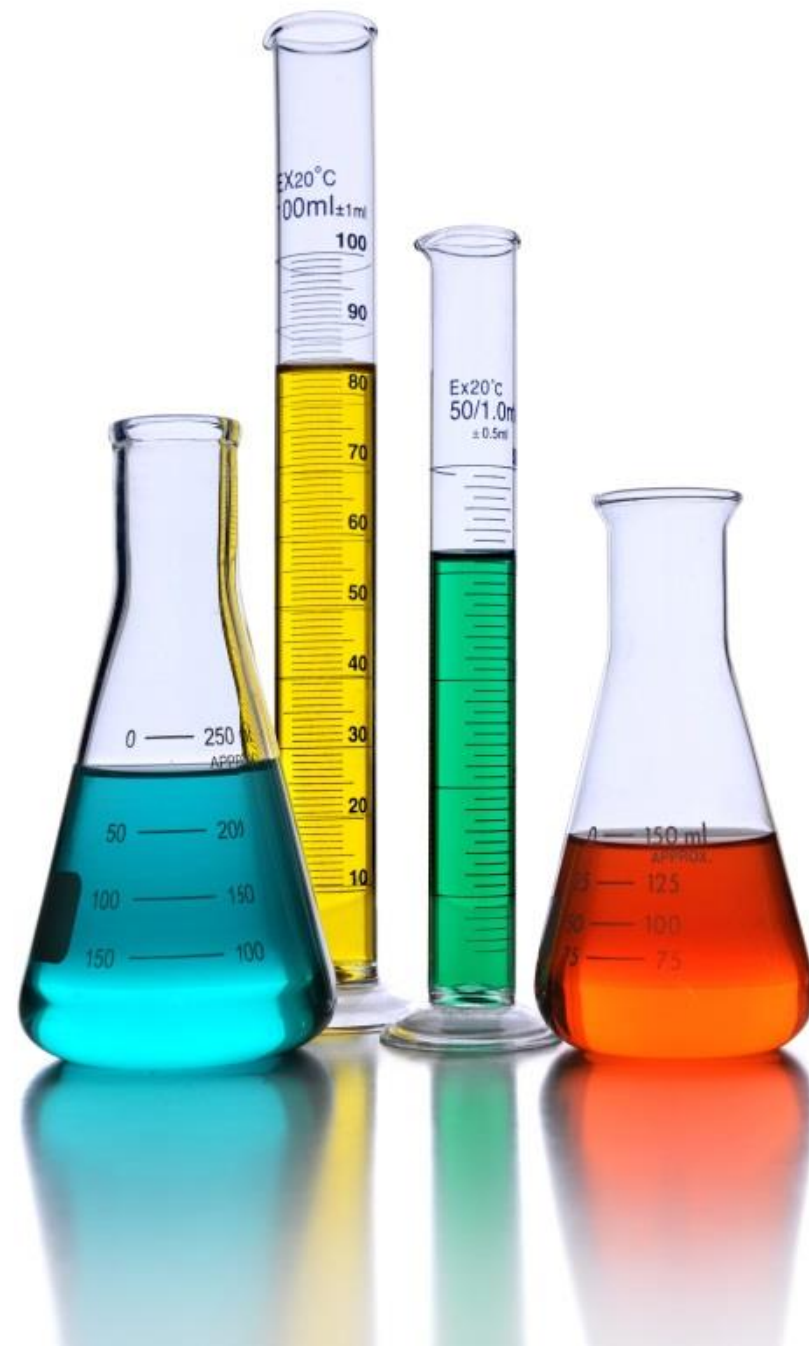
Scalable Personalized Pricing Validation

Nov 2015: Implement a 3 bin experiment

5,315 new “starter” firms

1. Control pricing – \$99 (25%)
2. Uniform pricing – \$249 (25%)
3. Personalized pricing (50%)

Run and evaluate.



Source: [Dubé and Misra \(2023\)](#)

Scalable Personalized Pricing

Actual Profits in Second Experiment

Pricing	Projected Profit Per Lead
Base (\$99)	\$25
Uniform (\$249)	\$40
Personalized	\$44

65%
Rev ↑
83%
Rev ↑

Source: [Dubé and Misra \(2023\)](#)

A Price Experiment at Ziprecruiter.com

ML squeezed more monetizable information out of the data

But again it (currently) requires a human to supply the microeconomic framework for the ML to be effective

34

What about non-numeric content?

Example: An E-Mail Marketing Newsletter



- Over 27,000 subscribers
- 2 daily newsletters featuring novel wines & expert reviews
- 3 professional writers (each costing \$125,000 per year)

Can a GPT capture the “voice” of the writer?
 (“Generative Pre-trained Transformer”)

The Daily Wine Access Newsletter

Training LLM to write a Wine Newsletter

In Dec 2023, ChatGPT off-the-shelf was *terrible!*

- ***Clichés, poor grammar and unprofessional language***
- ***Hallucinations***
- ***Opaque algorithm makes it difficult to decipher sources of errors***

Training LLM to write a Wine Newsletter

3-step training approach

1 Gather training data

- 5 years of writing team's high-revenue newsletters
- same prompts as writers: wine attributes, expert reviews, scores, etc.



2 Create synthetic data using inputs from (1) & GPT-4

- Real-world data insufficient to train LLM
- So repeat (1) many times using synthetic outputs



3 Fine-tune the LLM with real & synthetic data

- uses algorithm Orca/Mistral



Training LLM to write a Wine Newsletter

Experimental Design

Treatment Period: 2 weeks (01/15/24-01/28/24)

Cell Sizes

- No e-mail: 500 (No e-mails sent)
- Human: 9,000 (e-mails written by human writing team -- **Control**)
- AI: 9,000 (e-mails written by AI)
- Hybrid: 9,000 (e-mails written by AI but edited by other human writing team)

Prediction in early 2024 → **hybrid cell** would “win” because it would make writers more productive

Example of a Newsletter

January 15, 2024

WINE ACCESS

Years of Study and Investment Have Paid Off

Making the very finest Amarone requires technical skill, precision equipment, and a fanatical attention to detail that 98% of vintners in the region don't have.

The 2% who do—such as legends Romano dal Forno and Giuseppe Quintarelli—often charge \$400 or more for their efforts.

Riccardo Tedeschi is also one of this region's conscientious elite: His family spent a full decade studying the optimum drying conditions for their grapes, and installing state-of-the-art fans and humidity control. And all their attention to detail and knowhow shines through in the 2019 Marne 180 Amarone, which **earned 97 points—outscoring 1,607 other Italian red wines**—at the 2023 *Decanter* World Wine Awards.

In a premium category, this is a powerful and concentrated Amarone that punches FAR above its weight. You can lay in a six-pack for the same price as one bottle of Dal Forno—and the price drops to just \$45 per bottle when you buy eight. You can't beat that anywhere!



2019 Tedeschi Marne 180 Amarone della Valpolicella Veneto

97 Decanter World Wine Awards

~~\$64~~

\$50 per bottle

Buy Now 22% Off

Shipping included on \$150+

Human

WINE ACCESS

Years of Study and Investment Have Paid Off

97 points from the *Decanter* World Wine Awards. A stunning Amarone that's worth every penny of its \$64 price—and then some. We've got it today for a staggering discount.

Amarone is one of our most popular categories, and this wine would normally be snapped up by collectors looking for a powerful, long-lived red that boasts the best of Italy. But in this topsy-turvy year, those collectors are staying home, which means we've got the keys to a cellar that normally would be locked tight.

That's how we scored this superstar at a price that's unthinkable for a wine of this quality. Today, six bottles cost about what you'd pay for one from a prestigious producer like Dal Forno.



2019 Tedeschi Marne 180 Amarone della Valpolicella Veneto

97 Decanter World Wine Awards

~~\$64~~

\$50 per bottle

Buy Now 22% Off

Shipping included on \$150+

Dark ruby at the center, with crimson highlights toward the rim, the nose is pure hedonism. Richly fruited with purple boysenberries, dark strawberries, and spiced red plums accented with licorice and kirsch liqueur. The mouthfeel is clean and polished, with supple tannins and sweet flower petals from start to finish.

A.I.

WINE ACCESS

Years of Study and Investment Have Paid Off

Amarone is one of our most popular categories, and the 2019 Tedeschi Marne 180 Amarone earned 97 points from the *Decanter* World Wine Awards. It's a stunning bottle that's worth every penny of its \$64 price—and considering today's staggering discount, it won't be long for this world.

For anyone looking for a powerful, long-lived red that boasts the best of Italy, this is your bottle. It boasts unthinkable quality for a wine at this price. **Consider that you can get six bottles today for about what you'd pay for one from a prestigious producer like Dal Forno.**



2019 Tedeschi Marne 180 Amarone della Valpolicella Veneto

97 Decanter World Wine Awards

~~\$64~~

\$50 per bottle

Buy Now 22% Off

Shipping included on \$150+

Dark ruby at the center, with crimson highlights toward the rim, the nose on this bad boy is pure hedonism—richly fruited with purple boysenberries, dark strawberries, and spiced red plums accented with licorice and kirsch liqueur. The palate is clean and polished, with supple tannins and sweet flower petals from start to finish.

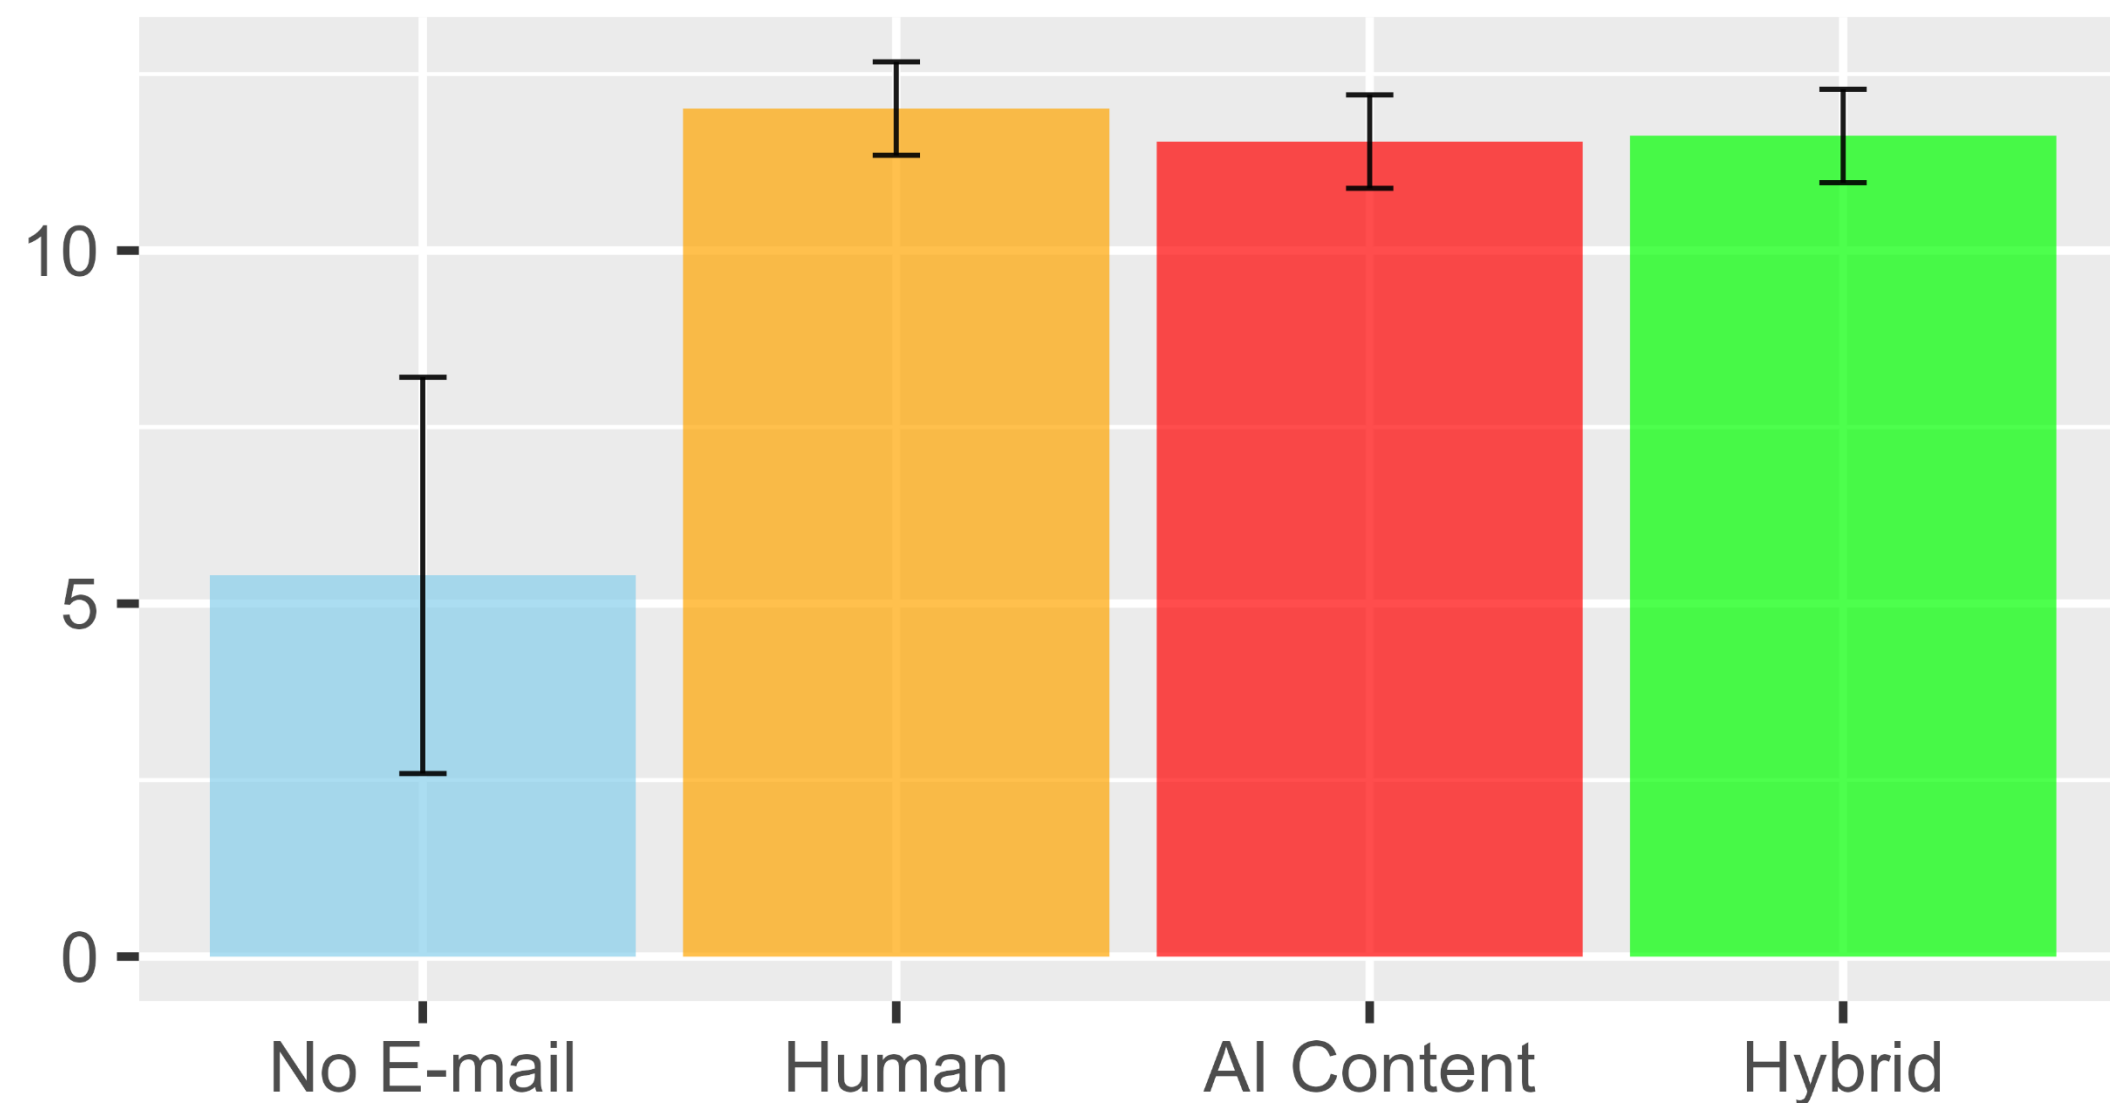
Hybrid

Training LLM to write a Wine Newsletter

Average outcome by cell

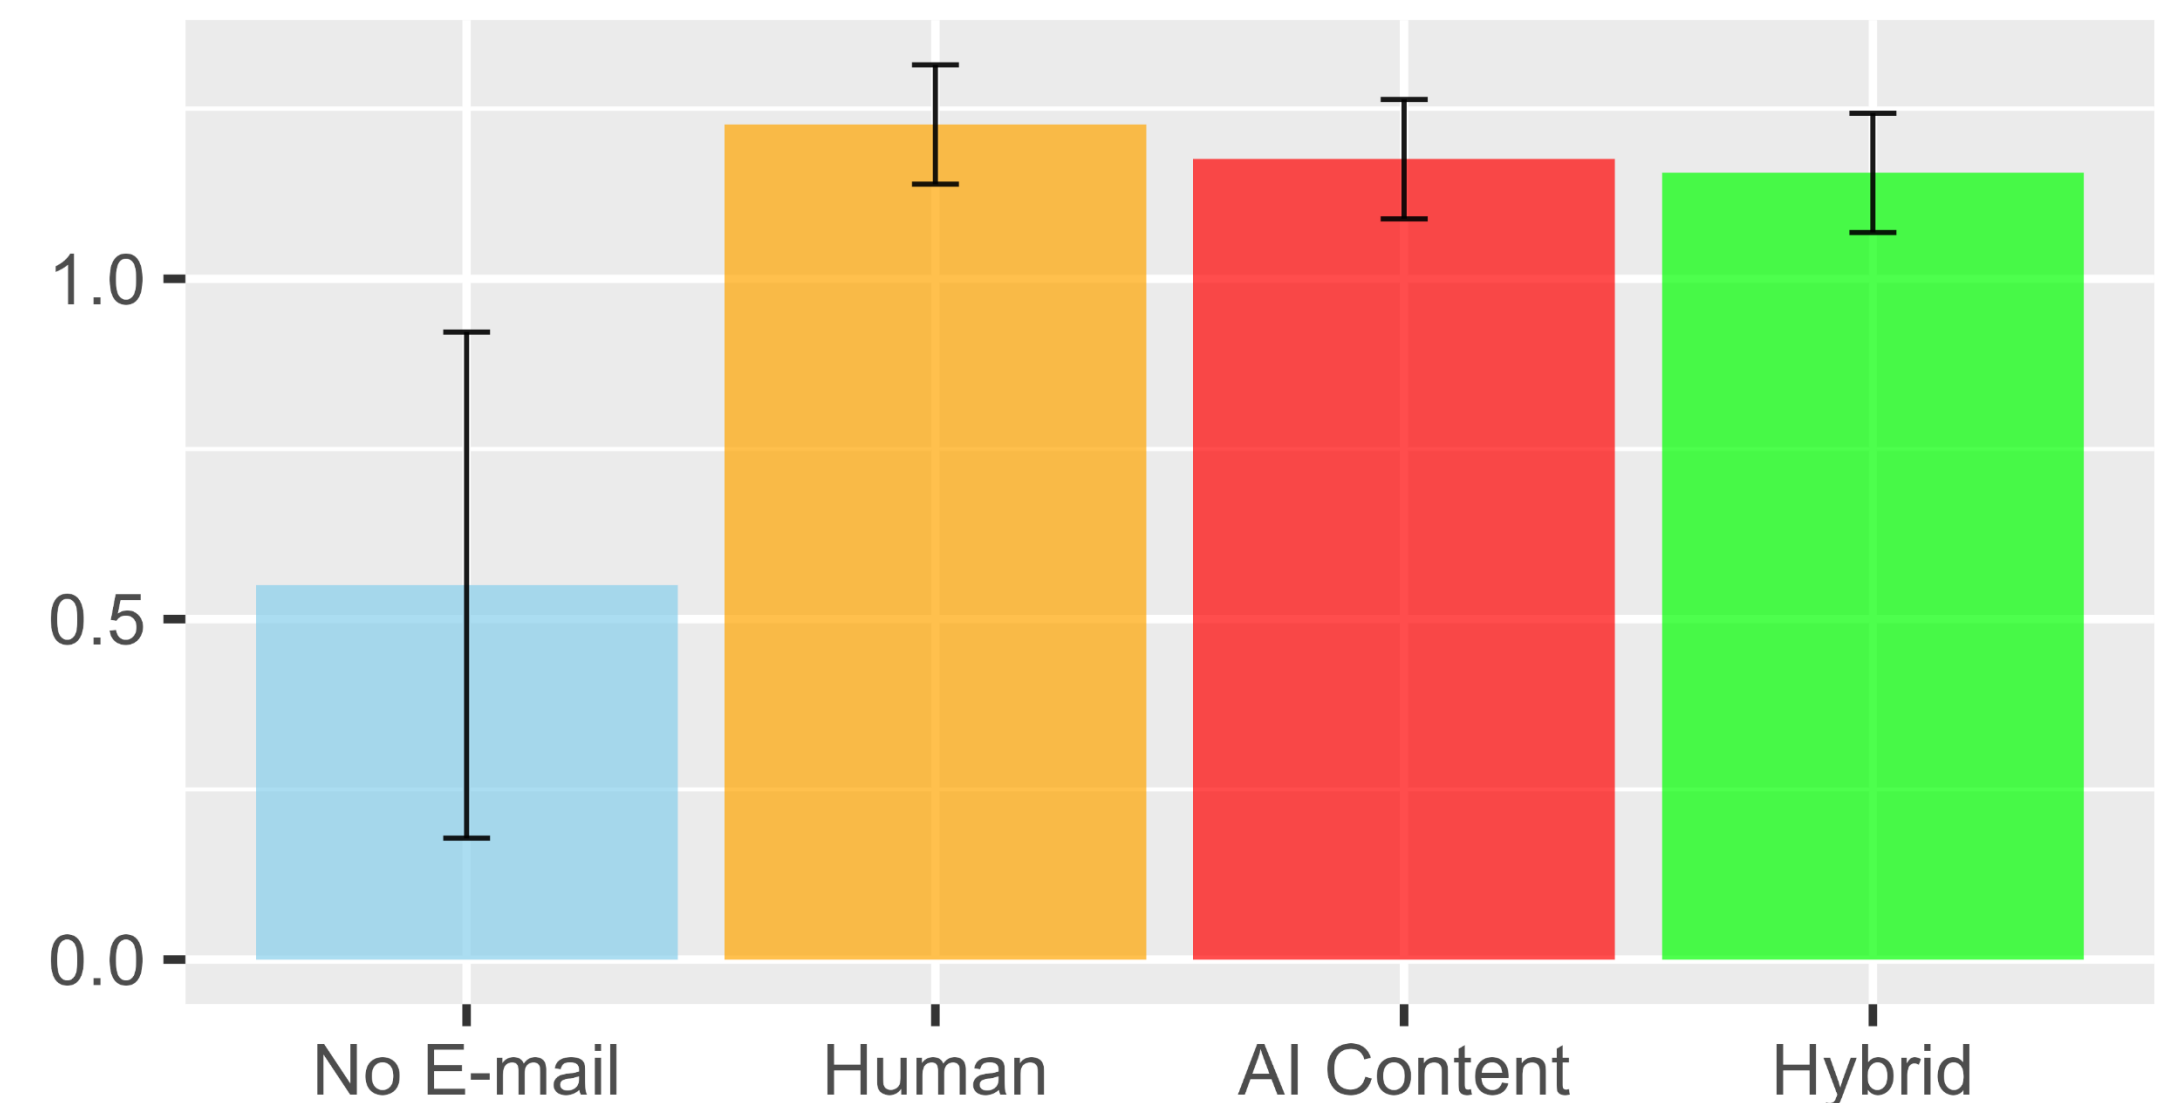
We fail to reject that AI is equivalent to Human

Purchase Probability (%)



E-mail more than doubles purchase incidence

Quantity of Bottles (#)



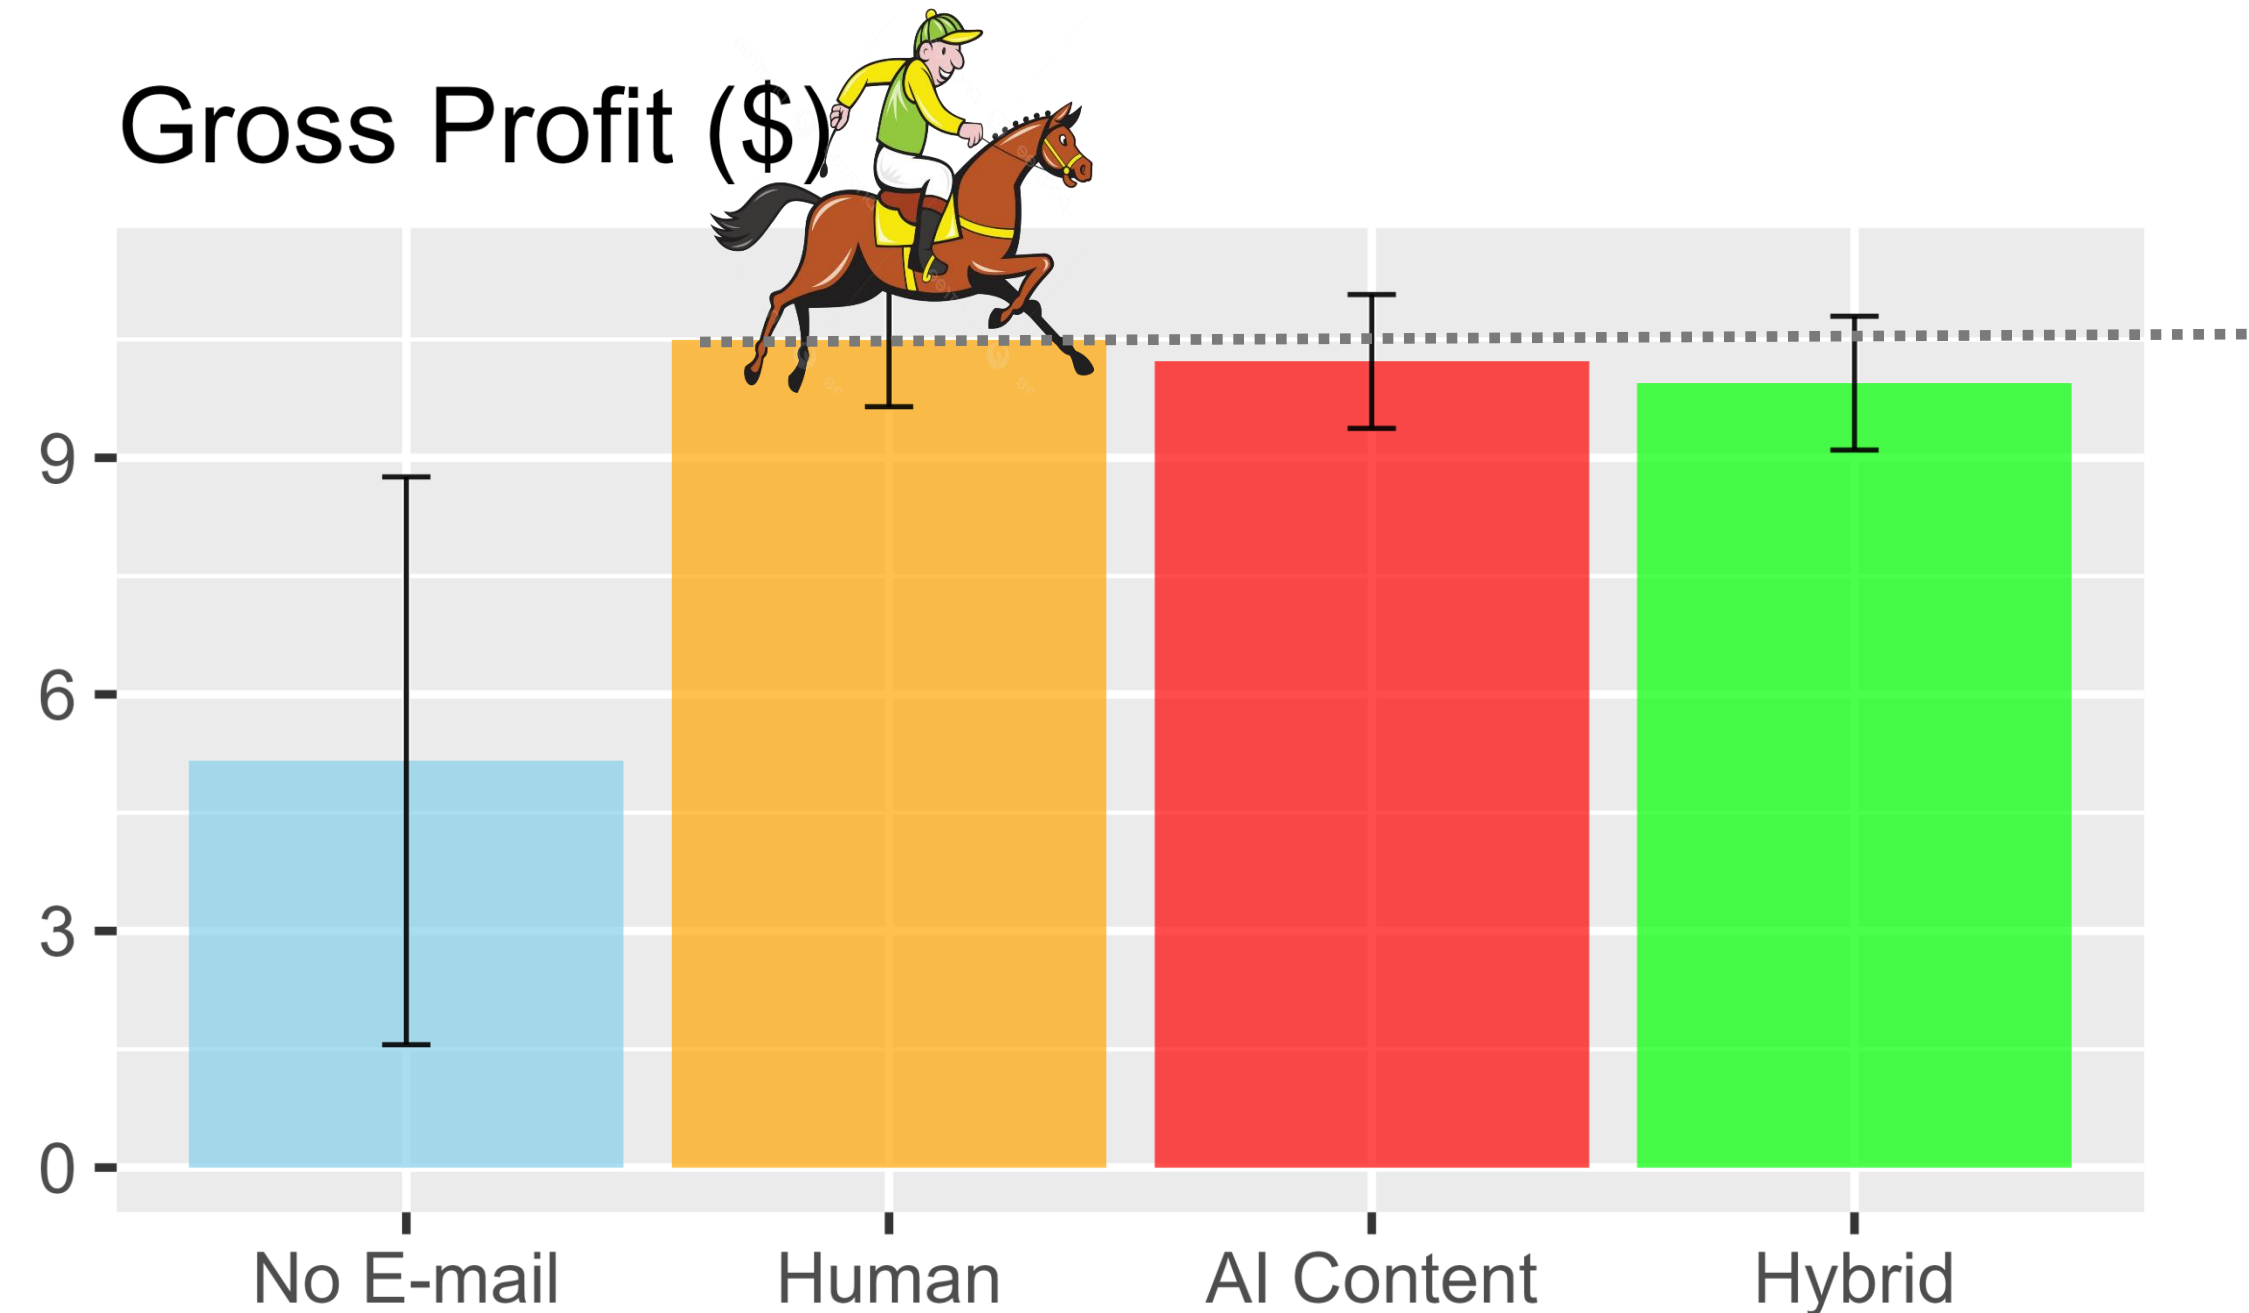
E-mail more than doubles # bottles purchased

Training LLM to write a Wine Newsletter

Average outcome by cell

Which cell should WA select?

Again, we
fail to reject
that AI is
equivalent
to Human



But must also consider avoidable fixed costs


E-mail almost doubles profits!!!

Training LLM to write a Wine Newsletter

Average outcome by cell

- Each writer costs \$125,000 per year!
- AI allows reducing writing team by 3 people

Aggregate Annual Projection



	Control	Human Content	AI Content	Hybrid
Profit	\$3,687,652	\$7,504,844	\$7,308,570	\$7,111,912
Cost	\$ 0 (No writers, no AI license)	\$375,000 (3 writers, no AI license)	\$1,000 (AI license)	\$126,000 (1 write + AI license)
Net Profit	\$3,687,652	\$7,129,844	\$7,307,570	\$6,985,912

This is after one experiment & limited GPT training

Training LLM to write a Wine Newsletter

Adjusting the Prediction

Hybrid performed the worst!

New study by Caplin et al (2024) finds

“AI assistance is more valuable for people who are calibrated, meaning they have accurate beliefs about their own ability.”

Bennett (2024) finds literary community (writers) most critical of AI (legal fights over copyright)

Training LLM to write a Wine Newsletter

Follow-up Study Design

Treatment Period: 2 weeks (09/30/24-10/13/24)

Cell Sizes

- Human: 8,273 (e-mails written by human writing team)
- AI: 8,274 (e-mails written by AI)
- Hybrid: 8,273 (e-mails written by AI but edited by **Marketing Team**)

Training LLM to write a Wine Newsletter

Average outcome by cell

- Each writer costs \$125,000 per year!
- AI allows reducing writing team by 3 people

Aggregate Incremental Annual Projection (Move to Human)		
	AI Content	Hybrid
Incremental Gross Profit	-\$254,015	\$46,424
Incremental Cost	-\$374,000 (AI license)	-\$249,000 (1 write + AI license)
Incremental Net Profit	\$119,985	\$295,424

This is after two experiments & limited GPT training

Training LLM to write a Wine Newsletter

Updating the AI one year later

Revisit using a GTP with prompts

Claude worked the best

Training LLM to write a Wine Newsletter

Follow-up Study Design

Treatment Period: 2 weeks (02/23/25-03/09/25)

Cell Sizes

- Human: 7,987 (e-mails written by human writing team)
- AI: 7,988 (e-mails written by AI (Claude))
- Hybrid: 7,988 (e-mails written by AI but edited by **Marketing Team**)

Training LLM to write a Wine Newsletter

Average outcome by cell

- Each writer costs \$125,000 per year!
- AI allows reducing writing team by 3 people

Aggregate Incremental Annual Projection(relative to Human)		
	AI Content	Hybrid
Incremental Gross Profit	\$278,314	\$315,428
Incremental Cost	-\$374,500 (AI license)	-\$248,500 (1 write + AI license)
Incremental Net Profit	\$651,814.3	\$563,928

This is after two experiments & limited GPT training

Training LLM to write a Wine Newsletter

A pre-trained LLM seems to produce promising e-mail copy

→ But needed historic e-mails (data) to train it properly

GenAI does even better (Claude)

- seems to produce better copy with prompts

Conclusions

Marketing should not be a gut-instinct guessing game

AI creates many new marketing opportunities

“Trust me, I’ve tested the AI application scientifically and ROMI is high!”

But we still need **theory & data** (not just ML) the analysis

We also (currently) still need a human to guide the AI

 **Marketing spending can be more accountable & drive growth!!!**

Thank you.

Jean-Pierre Dubé

James M. Kilts Distinguished Service Professor of Marketing

jdube@chicagobooth.edu

University of Chicago Booth School of Business

