SIMULATED MARKET REACTION & DECISION BRIEF

Greenbowl lunch ad (sample)

IDEA TYPE	PACKAGE	VALID ANSWERS	COMPLETED
Advertising idea	Pick the Winner · $49	10,000 of 10,000	October 27, 2026

STRONGEST VERSION IN THIS SIMULATION

Version C

31% would respond

3,078 would respond · 2,810 maybe · 4,112 would ignore it

OVERALL CALL

Mixed: fix weak points, then validate

Other versions: A 26% · B 24%

WHAT IT MEANS

The speed message plus the free drink landed best in this simulation, mostly with office workers who eat at their desks. The biggest gap is clarity: many people didn't know where the salad bar is or what the bowls look like.

RECOMMENDED NEXT MOVE

Fix the clarity gaps in Version C (location, a bowl photo, the offer end date), then run it against one alternative in a small real campaign.

ILLUSTRATIVE EXAMPLE

Illustrative sample. The simulated answers and the AI write-up were scripted for this demonstration and run through the real counting and checking steps. Not a real customer's test and not market research.

CONTENTS

Inside this decision report

01	The short answer	3
02	What was tested and how	4
03	How the simulated audience reacted	5
04	What the simulated people said	6
05	What’s holding your idea back?	8
06	What I would improve first	10
07	Who to focus on first	12
08	Strategic comparison & decision matrix	13
09	Pricing and positioning hypotheses	14
10	Ad targeting guidance	15
11	Your next business decision	17
12	A practical two-week plan	19
13	Before you decide	20
14	What this report can and can’t tell you	21

Simulated results are counted from the AI-simulated answers. AI hypotheses are ideas to test. Nothing in this report is real customer data, market research, a forecast or proof of demand.

Executive decision brief

The short answer

3 versions · price Free drink with any bowl · 10,000 simulated people in the audience in 25 groups

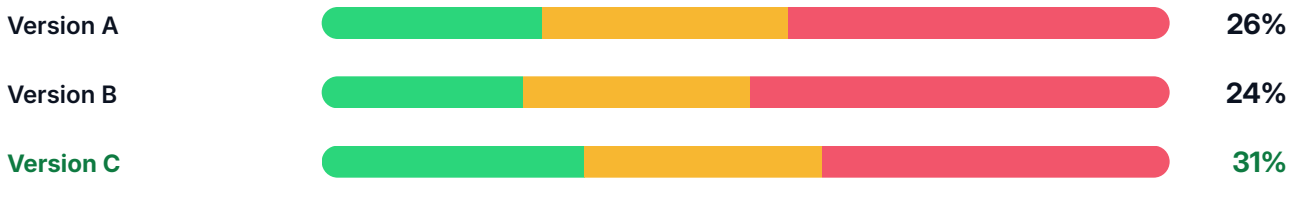
STRONGEST VERSION
31%
Version C · would respond

VALID ANSWERS
10,000
of 10,000 AI-simulated participants

OVERALL CALL
Mixed: fix weak points, then validate

LEAD OVER THE NEXT VERSION
+5
points ahead of Version A

HOW EACH VERSION WAS ANSWERED — SIMULATED RESULT



STRONGEST POSITIVE SIGNAL
Price and value
48%
2,093 of 4,346 yes reasons
Speed (bowls in five minutes) is the benefit people repeated most when they responded.

BIGGEST OBJECTION
Clarity: unclear what it is or does
43%
2,450 of 5,654 not-yes reasons
HIGH
People can't tell where it is or what they get

MOST PROMISING GROUP
Health-focused office workers, 25–27
46%
441 people · overall 31%

RECOMMENDED NEXT MOVE — AI HYPOTHESIS

Fix the clarity gaps in Version C (location, a bowl photo, the offer end date), then run it against one alternative in a small real campaign.

Unclear location was the most common reason to scroll past, and it is the easiest thing to fix.

Test overview

What was tested and how

The idea as you described it, the simulated audience and how the answers were collected.

VERSION A

Instagram ad. Headline: “Your lunch break, upgraded.” A salad bar near downtown offering a free drink with any bowl this week. Button: Order now.

Price: Free drink with any bowl

VERSION B

Instagram ad. Headline: “Fresh bowls in 5 minutes.” A salad bar near downtown. Order ahead and skip the line. Button: Order now.

VERSION C

Instagram ad. Headline: “Lunch that doesn’t slow you down.” A salad bar near downtown with bowls ready in 5 minutes and a free drink this week. Button: Order now.

Price: Free drink with any bowl

ITEM	DETAILS
Idea type	Advertising idea
Audience you described	Office workers 22–50 near downtown Los Angeles
Package	Pick the Winner · 10,000 AI-simulated participants · $49
Valid simulated answers	10,000 of 10,000
Audience groups	25
Versions compared	3
Reasons collected	10,000 (topic given with each answer)
Completed	October 27, 2026 · Test ST-99ZGTR

How the answers were produced

An AI model played each simulated participant: a person from one of the audience groups, with their own age, job, budget and, for newer tests, what they want, what they use today and what worries them. Each one gave an answer (Would respond, Maybe or Would ignore) and a short reason. SellTo.us counted the answers. A second AI model wrote the explanations and recommendations, and automatic checks compared every number and claim with the counted results.

Group sizes and person traits are assumptions about the audience you described, not population statistics. Each test is one simulation run.

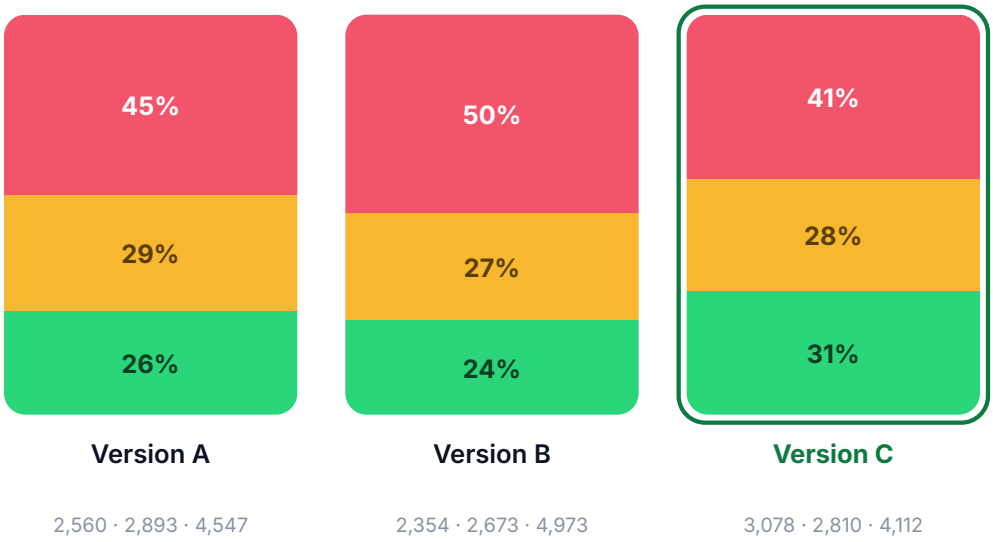
Simulated audience results

How the simulated audience reacted

10,000 valid answers. Every version was shown to the same simulated people.

Answer split

SIMULATED RESULT



Counts under each column: Would respond · Maybe · Would ignore

Said yes to only one version: A 546 · B 429 · C 852 (of 10,000).

Version C and Version A are 5 points apart in this single run.

Results by audience group and version

SIMULATED RESULT

AUDIENCE GROUP	SHARE OF CROWD	VERSION A	VERSION B	VERSION C
Health-focused office workers, 25–27	4% · 441	41%	36%	46%
Health-focused office workers, 28–31	4% · 428	39%	33%	46%
Fitness fans, 26–30	4% · 367	40%	34%	46%
Health-focused office workers, 35–38	4% · 412	38%	31%	45%
Health-focused office workers, 32–34	4% · 418	38%	30%	44%
Health-focused office workers, 39–41	4% · 411	34%	27%	44%
Time-pressed managers, 38–41	4% · 397	25%	39%	43%

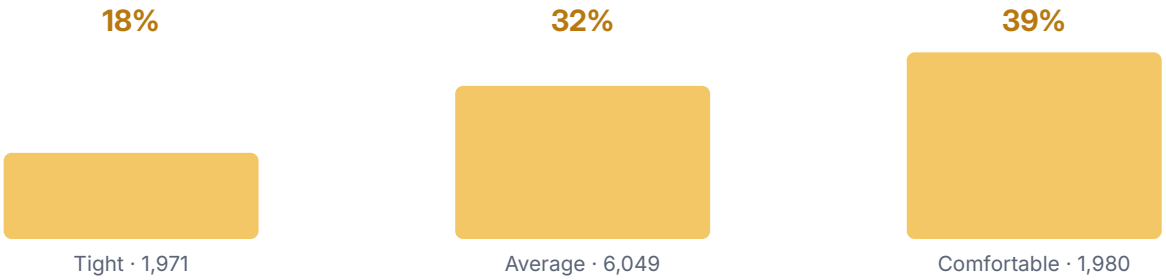
AUDIENCE GROUP	SHARE OF CROWD	VERSION A	VERSION B	VERSION C
Time-pressed managers, 34–37	4% · 403	22%	**43%**	41%
Health-focused office workers, 42–45	4% · 444	32%	27%	**40%**
Time-pressed managers, 42–45	4% · 382	23%	38%	**40%**
Time-pressed managers, 30–33	4% · 410	26%	37%	**39%**
Fitness fans, 22–25	4% · 342	**39%**	33%	38%
Fast-food regulars, 25–28	4% · 421	11%	**12%**	12%
Fast-food regulars, 29–32	4% · 394	9%	9%	**12%**
Fast-food regulars, 33–36	4% · 406	9%	9%	**10%**
Fast-food regulars, 37–40	4% · 397	6%	6%	**7%**

Darker cells: a larger share answered “Would respond”. The highest value in each row is in bold.

Showing the 12 most and 4 least interested groups with enough answers, of 25.

Would respond by budget level · Version C

SIMULATED RESULT



Key findings & evidence

What the simulated people said

Findings computed from the answers, the topics behind yes and no, and reasons in the participants’ own words.

Key findings

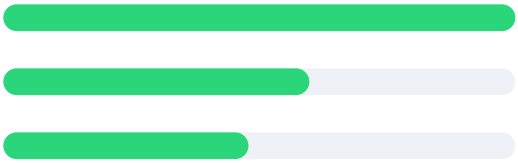
SIMULATED RESULT

- Version C: 31% would respond to it, against 26% for Version A (a 5-point difference in this single simulation run; repeat runs were not done, so how stable the gap is has not been measured).
- Most interested: Health-focused office workers, 25–27 (46%) and Health-focused office workers, 28–31 (46%) for Version C.
- Least interested: Fast-food regulars, 37–40 (7% would respond; 4% of participants).

- Version C and Version A differ in features or details (5 minutes), claims (in 5 minutes) and wording at the same time, so this test can't say which difference moved people in the audience.

Why people said yes

SIMULATED RESULT



One topic per reason. Shares of all 4,346 yes reasons. The reasons for not saying yes are charted in the next section.

In their words

SIMULATED RESULT

> "A free drink this week would get me to tap."
>
> Would respond · A · Age 29 · trainer · Fitness fans, 26–30

> "A free drink with a bowl is a nice push to try it."
>
> Would respond · A · Age 41 · accountant · Time-pressed managers, 38–41

> "Is it delivery or pickup? The ad isn't clear."
>
> Would ignore · A · Age 40 · recruiter · Health-focused office workers, 39–41

> "The ad doesn't say where it is."
>
> Maybe · A · Age 43 · recruiter · Health-focused office workers, 42–45

> "Downtown is big, where exactly is this?"
>
> Maybe · A · Age 38 · nurse · Fitness fans, 36–40

Answers written by AI-simulated participants, copied word for word. They are not real customers.

Main objections & business risks

What's holding your idea back?

Problems ranked by fixed rules from the counted reasons. Severity, counts and affected groups come from the answers; explanations and actions are AI hypotheses.

Top reasons for not saying yes

SIMULATED RESULT

1	Clarity: unclear what it is or does	HIGH	43%	2,450
2	No need, or already covered	HIGH	29%	1,656
3	Look, design and quality	MEDIUM	13%	729
4	Price and cost	MEDIUM	13%	713
5	Trust, proof and reviews		2%	106

One topic per reason, so these are shares of all 5,654 not-yes reasons (from people who said maybe or no). Coloured bars are the problems below; severity follows the rules at the end of this section.

1 HIGH CLARITY: UNCLEAR WHAT IT IS OR DOES

People can't tell where it is or what they get

2,450 of 5,654 not-yes reasons (43%)

WHAT IT MEANS

Many who would scroll past didn't know the location, the menu or whether it's pickup or delivery.

IF IT STAYS UNRESOLVED (HYPOTHESIS)

Even people who like the message may not act if they can't picture the trip or the bowl.

WHAT TO DO

Add the street or a map pin, a photo of a bowl and the word pickup to the ad.

MOST AFFECTED GROUPS

Health-focused office workers, 32–34	69%	124 of 180
Health-focused office workers, 25–27	67%	112 of 166

"Is it delivery or pickup? The ad isn't clear."

2 HIGH NO NEED, OR ALREADY COVERED

Many already have their lunch sorted

1,656 of 5,654 not-yes reasons (29%)

WHAT IT MEANS
People who bring lunch or eat fast food see no reason to switch.

IF IT STAYS UNRESOLVED (HYPOTHESIS)
Ads shown to these groups are likely to be ignored.

WHAT TO DO
Leave fast-food regulars out of the first campaign and focus on people who already eat healthy.

MOST AFFECTED GROUPS

Fast-food regulars, 33–36	61% · 212 of 345	
Fast-food regulars, 29–32	59% · 191 of 324	

"I already have a lunch spot I like."

3 MEDIUM LOOK, DESIGN AND QUALITY

The ad doesn't show the food

729 of 5,654 not-yes reasons (13%)

WHAT IT MEANS
Some people want to see the bowl before they care.

IF IT STAYS UNRESOLVED (HYPOTHESIS)
A text-only creative may look like every other salad ad.

WHAT TO DO
Test a creative with a close-up of a real bowl.

MOST AFFECTED GROUPS

Fitness fans, 36–40	43% · 83 of 192	
Fitness fans, 22–25	41% · 62 of 152	

"A text ad without a food photo doesn't make me hungry."

4 MEDIUM PRICE AND COST

Budget lunch buyers find bowls pricey

713 of 5,654 not-yes reasons (13%)

WHAT IT MEANS
For people who bring lunch, a free drink doesn't change the bowl's price.

IF IT STAYS UNRESOLVED (HYPOTHESIS)
This group may stay out regardless of the message.

WHAT TO DO
Don't target budget lunch buyers until a lower-priced option exists.

CONTINUED

MOST AFFECTED GROUPS

Budget lunch buyers, 27–29	46% · 121 of 264	"Salad bowls downtown are usually too pricey for me."
Budget lunch buyers, 30–32	44% · 134 of 304	

How severity is decided

REPORT RULE

LEVEL	RULE
Critical	The most common topic among not-yes reasons, mentioned in at least 30% of them, while the strongest version gets under 30% positive answers.
High	Mentioned in at least 20% of not-yes reasons, or the most common topic with at least 12%.
Medium	Mentioned in at least 8% of not-yes reasons, or in at least 25% of the not-yes reasons of one audience group with 15 or more of them.
Low	Supported by some not-yes reasons, below the thresholds above.

These are report rules applied to one simulation run, not statistical tests.

Recommended improvements

What I would improve first

Ranked changes to try, with how to test each one with real people. AI hypotheses checked against the counted results, not promises.

1

LOW EFFORT

Add location and pickup details

WHAT TO CHANGE

Put the street name or a map pin and the word pickup in the first line of Version C.

WHY IT MATTERS

Unclear location was the most common reason to scroll past.

HOW TO TEST IT

Run the clearer ad next to the current Version C to the same audience and compare orders.

RISKS OR DEPENDENCIES

Longer text may push the headline out of view on small screens.

EVIDENCE

- Clarity: unclear what it is or does: 2,450 of 5,654 not-yes reasons (43%)
- "Is it delivery or pickup? The ad isn't clear."

2 LOW EFFORT

Show a real bowl in the image

WHAT TO CHANGE
Replace the text-only creative with a close-up photo of a popular bowl.

WHY IT MATTERS
Some people wanted to see the food before acting.

HOW TO TEST IT
Compare the photo creative with the text creative, keeping the same headline and offer.

RISKS OR DEPENDENCIES
The photo must match what people actually get.

EVIDENCE
- Look, design and quality: 729 of 5,654 not-yes reasons (13%)

3 LOW EFFORT

Say when the offer ends

WHAT TO CHANGE
Add the last day of the free drink offer to the ad.

WHY IT MATTERS
People who liked the free drink may wait without a deadline.

HOW TO TEST IT
Compare orders in the days before and after the deadline.

RISKS OR DEPENDENCIES
A deadline only works if you keep it.

EVIDENCE
- Price and value: 2,093 of 4,346 yes reasons (48%)

4 MEDIUM EFFORT

Test the headline separately from the offer

WHAT TO CHANGE
Keep the Version C offer and test two headlines, then keep the headline and test the offer.

WHY IT MATTERS
The versions differ in several ways, so the effect of each part is unknown.

HOW TO TEST IT
Two small campaigns that each change one thing.

RISKS OR DEPENDENCIES
More rounds take more time.

EVIDENCE
- The versions differ in several ways at once

Audience opportunities

Who to focus on first

The groups that responded best in the simulation, and where the opportunities and gaps are.

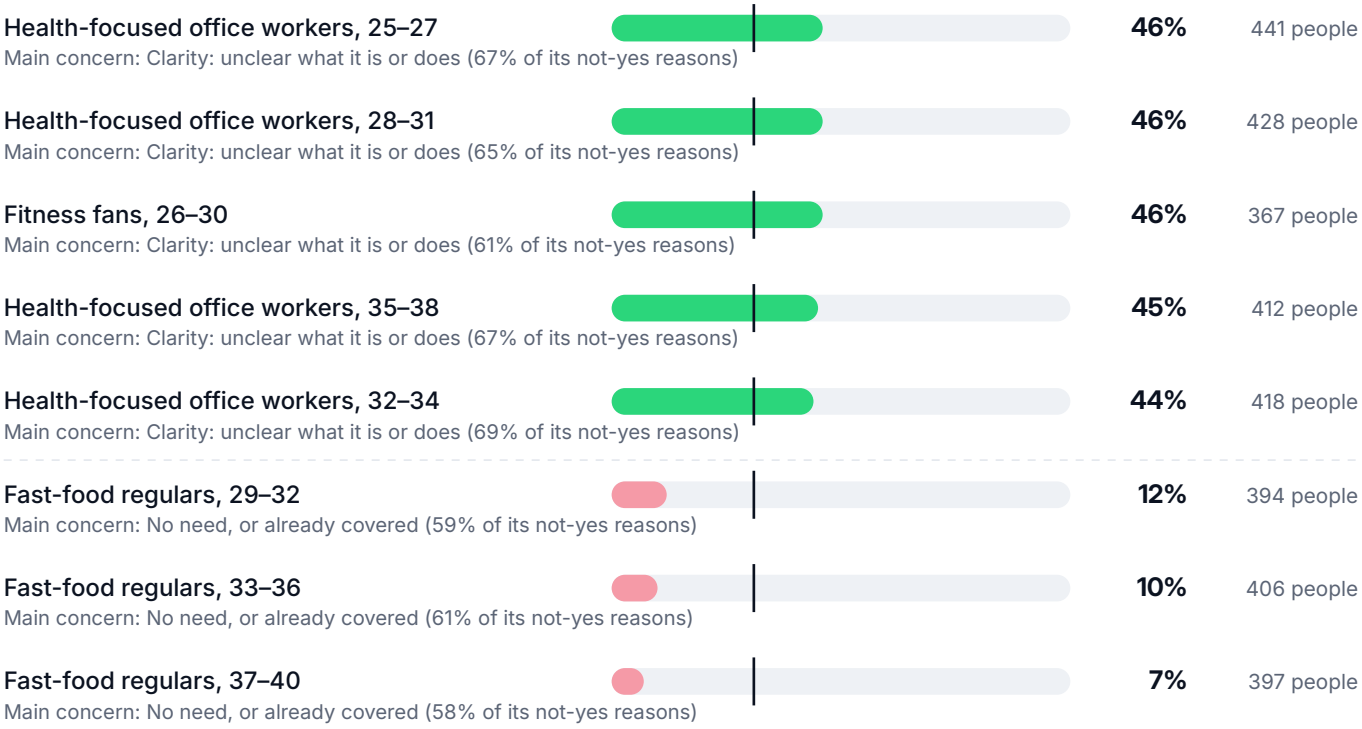
AUDIENCE TO FOCUS ON FIRST

AI HYPOTHESIS

Start with health-focused office workers and time-pressed managers: they responded best to Version C and work within walking distance.

Most and least interested groups

SIMULATED RESULT



Green: most interested groups. Pink, below the dashed line: least interested. Vertical line: all simulated people (31%). Groups with fewer than 100 answers are left out.

Version & image comparisons

Strategic comparison & decision matrix

Every version went to the same 10,000 simulated people in the audience. Counts are from their answers; strengths and weaknesses are AI hypotheses.

RECOMMENDED CANDIDATE

SIMULATED RESULT

Version C

Version C scored 31% against 26% for Version A (5 points apart in this single run).

	VERSION A	VERSION B	VERSION C
Would respond	**26% · 2,560**	**24% · 2,354**	**31% · 3,078**
Maybe	29% · 2,893	27% · 2,673	28% · 2,810
Would ignore it	45% · 4,547	50% · 4,973	41% · 4,112
Said yes only to this one	546	429	852
Strongest group	Health-focused office workers, 25–27 · 41%	Time-pressed managers, 34–37 · 43%	Health-focused office workers, 25–27 · 46%
Does well	The free drink gives a reason to act this week.	The speed promise appeals to managers who eat at their desks.	Combines the speed promise with the free drink.
Weak spot	Doesn’t say anything about speed, which people valued most.	No offer, and budget lunch buyers barely responded.	Still doesn’t say where the salad bar is.
Price or offer	Free drink with any bowl	—	Free drink with any bowl

WHY NO SINGLE CAUSE IS NAMED

Version C and Version A differ in several ways at once (features or details (5 minutes); claims (in 5 minutes); wording), so this test shows which version did better, not which change caused it.

How to read it

AI HYPOTHESIS

Version C scored highest in this simulation. The three ads differ in headline, offer and wording at once, so this test can’t tell which part made the difference.

A cleaner follow-up comparison

AI HYPOTHESIS

Keep the Version C headline and test the free drink against no offer, so only the offer changes.

- Offer test: keep the price, wording and image the same and change only what is included or the features (5 minutes).
- Wording test: only meaningful once what is offered is identical. A same-price comparison of these versions would not isolate the wording, because what is offered also differs.

Assumptions that need a cleaner comparison

AI HYPOTHESIS

- The speed headline matters more than the free drink.

Pricing & positioning

Pricing and positioning hypotheses

Ideas to test with real people. Nothing here is verified market research; no competitor prices are shown unless a dated source is listed.

A CLEARER MESSAGE TO TEST

AI HYPOTHESIS

“Fresh bowls in 5 minutes on Spring Street, free drink this week”

A salad bar near your office: order ahead, pick up in five minutes, and get a free drink this week.

It keeps the speed message and answers where and what in one line.

OFFER HYPOTHESIS TO VALIDATE

AI HYPOTHESIS

Test the free drink against a small discount on the bowl.

Two versions of the ad that differ only in the offer, shown to the same audience.

Possible positioning

AI HYPOTHESIS

Lead with speed for busy office workers.

Alternatives people may choose instead

AI HYPOTHESIS

ALTERNATIVE	WHY SOME PICK IT	HOW YOU COULD STAND APART
Nearby salad chains	Familiar and close to offices.	A local spot with a this-week offer.

AI-suggested alternatives, not market research. Names and details are not verified.

SWOT

AI HYPOTHESIS

STRENGTHS

- A clear speed promise

WEAKNESSES

- No location in the ad

OPPORTUNITIES
- Office lunch groups

THREATS
- Many lunch ads downtown

Advertising & targeting

Ad targeting guidance

Hypotheses for a small real ad test, built from the simulated groups. Not results and not a promise of performance. One example budget is used throughout: about $30 a day for 14 days ($420 in total). Scale it to your means; keep the split.

Show Version C to health-focused office workers and managers near downtown at lunchtime.

Health-focused office workers · 50% of budget (~$15/day)

Simulated: 44% would respond (Version C) · 26% of the crowd
The largest group with strong interest in Version C.

AGES	25–45
WHERE	Downtown Los Angeles, 1 mile radius
PLATFORMS	Instagram Stories
INTERESTS	Healthy eating, Salads, Wellness, Yoga, Meal planning, Nutrition, Fitness, Plant-based food
BEHAVIORS	Engaged shoppers
TIMING TO TEST	11am–1pm weekdays
SHOW	Version C
MESSAGE TO TEST	“Fresh bowls in 5 minutes” Healthy lunch without the wait.

Busy managers · 35% of budget (~$11/day)

Simulated: 39% would respond (Version C) · 20% of the crowd
They eat at their desks and care most about speed.

AGES	30–50
WHERE	Downtown Los Angeles
PLATFORMS	Instagram Feed
INTERESTS	Productivity, Business news, Healthy eating, Coffee, Fitness, Leadership, Commuting, Food delivery
TIMING TO TEST	11am–12pm weekdays
SHOW	Version C
MESSAGE TO TEST	“Lunch in 5 minutes” Lunch that doesn’t slow you down.

Fitness fans · 15% of budget (~$4/day)

Simulated: 39% would respond (Version C) · 15% of the crowd
They look for a quick healthy option around training.

AGES	22–40
WHERE	Downtown Los Angeles
PLATFORMS	Instagram Reels
INTERESTS	Fitness, Healthy eating, Protein, Running, Gyms, Meal prep, Sports nutrition, CrossFit
TIMING TO TEST	12–1pm weekdays
SHOW	Version C
MESSAGE TO TEST	“Post-gym lunch, sorted” Protein-packed bowls, ready when you are.

Google search keywords

AI HYPOTHESIS

salad near me · healthy lunch downtown la · quick lunch downtown · salad bowl pickup · order ahead salad · lunch near me · healthy bowls la · grain bowl

Exclude: recipe · catering

Leave out

AI HYPOTHESIS

- Fast-food regulars (all ages): The least interested group.

Your first two weeks (example ad test)

REPORT RULE

- Days 1–4: run all 3 ad sets at about $30 a day in total (Health-focused office workers $15, Busy managers $11, Fitness fans $4).
- Day 5: compare cost per click and real sign-ups or inquiries; pause the ad sets clearly behind the others, keeping at least two running.
- Days 6–10: move the paused budget to the best one or two ad sets; keep the daily total at $30.
- Days 11–14: retarget site visitors. Total test spend: about $420. Decide on scaling from real sign-ups, inquiries or sales, not from this report.

Lookalike and retargeting

AI HYPOTHESIS

Once real orders start, build a lookalike of customers who ordered.

Interest names differ slightly between Meta, TikTok and Google; pick the closest match. Start small and let real results guide you.

Validation experiments

Your next business decision

What the simulation suggests, what it can't tell you, and the experiments that would let you decide with real evidence.

THE DECISION

AI HYPOTHESIS

Should Version C be the ad to run, and what needs fixing first?

From problem to decision

AI HYPOTHESIS

PROBLEM IDENTIFIED
People can't tell where it is or what they get

→

EXPERIMENT TO RUN
A small real campaign of the clearer Version C against the current one, with pickup orders counted.

→

WHAT TO MEASURE
Pickup orders from each version, not likes

After the real test

WOULD SUPPORT PROCEEDING
The clearer ad brings real pickup orders from office workers.

WOULD SUGGEST REVISING OR STOPPING
People click but don't order, or the clearer ad gets no more orders than the current one.

Decision guidance only: what real people do in the experiment decides which way to go.

WHAT THIS SIMULATION SUGGESTS

- Version C scored highest in this simulation.
- Unclear location was the most common reason to scroll past.

WHAT REMAINS UNKNOWN

- Whether real office workers order after seeing it.
- Which part of Version C matters most.

UNRESOLVED ASSUMPTIONS TO CHECK

AI HYPOTHESIS

- People will walk to the salad bar once they know where it is.
- A food photo would make the ad more appealing than text alone.
- The free drink matters less than speed for busy managers.

If the evidence is mixed

AI HYPOTHESIS

- Promote the salad bar to office lunch groups with a group order offer.
- Hand out flyers with a QR code in nearby buildings at lunchtime.

Validation experiments

AI HYPOTHESIS

1 Clarity fix campaign

GOAL
See whether location and a photo bring more orders

HYPOTHESIS
The clearer Version C brings more pickup orders

WHO TO INVOLVE
Office workers within walking distance

HOW TO SET IT UP
Run two versions of Version C to the same audience at lunchtime: the current one and one with location, photo and pickup details.

WHAT TO ASK OR SHOW
Use an order link or code for each version.

WHAT TO MEASURE
Pickup orders from each version, not likes

ENCOURAGING
The clearer version brings more orders

CONCERNING
Neither version brings orders

LIMITATIONS
Small campaigns can swing from day to day

NEXT STEP
Keep the better version and test the offer next

2 Five-second test

GOAL
Check that people understand the ad

HYPOTHESIS
People can say what and where after five seconds

WHO TO INVOLVE
Five to ten office workers nearby

HOW TO SET IT UP
Show the ad for five seconds on a phone, then ask questions.

WHAT TO ASK OR SHOW
What is this offering, and where would you go?

WHAT TO MEASURE
How many answer both correctly

ENCOURAGING
Most answer both correctly

CONCERNING
Many can’t say where it is

LIMITATIONS
A tiny group; useful for spotting confusion, not for measuring demand

NEXT STEP
Fix whatever people miss before spending on ads

3 Lunchtime flyer test

GOAL
Compare the message offline

HYPOTHESIS
The speed message works in person too

WHO TO INVOLVE
People leaving nearby office buildings at lunch

HOW TO SET IT UP
Hand out flyers with the Version C message and a QR code for one week.

WHAT TO ASK OR SHOW
Scan for the free drink.

WHAT TO MEASURE
QR scans and redeemed drinks

ENCOURAGING
Steady redemptions during the week

CONCERNING
Few scans or redemptions

LIMITATIONS
Depends on location and weather

NEXT STEP
Use what people redeemed to shape the online ad

Small tests show direction; they don't prove demand. Decide what counts as encouraging or concerning before you start, and judge by what people do, not only by what they say.

Practical action roadmap

A practical two-week plan

What to do in the next two weeks to validate the strongest option with real people.

Days 1–3	Days 4–8	Days 9–11	Days 12–14
2 tasks	2 tasks	1 tasks	1 tasks

1

DAYS 1–3

- Make the clearer Version C with location, photo and pickup details
- Run the five-second test with a few office workers

2

DAYS 4–8

- Run the current and the clearer Version C at lunchtime to the same audience
- Track pickup orders from each, not likes

3

DAYS 9–11

- Keep the better version and prepare the offer test

4

DAYS 12–14

- Run Version C with and without the free drink and compare orders

Decision checkpoint: compare what real people did with “Would support proceeding” and “Would suggest revising or stopping” in Your next business decision.

Final decision checklist

Before you decide

Tick these off as you go. They come from the recommendations in this report.

Next-step checklist

- [ ] Add the street or a map pin to Version C
- [ ] Shoot a close-up photo of a popular bowl
- [ ] Set the offer end date
- [ ] Create order links for each ad version
- [ ] Run the five-second test with a few office workers
- [ ] Decide what result would make you keep Version C

Limitations & methodology

What this report can and can't tell you

How the numbers were produced, how the recommendations were made, and what nobody has verified.

What the numbers mean

- Percentages are shares of the simulated participants who answered (10,000), unless a different base is named next to them.
- Topic counts use one topic per reason. Their base is the reasons of people who said yes, or of people who said maybe or no ("not-yes"), as shown.
- Missing answers were not replaced or invented.
- Group sizes and person traits are assumptions about the audience you described, not population statistics.

How the recommendations were made

- The counts, the problem ranking, the severity levels, the most affected groups and the recommended version come from the answers, using fixed report rules.
- An AI model wrote the explanations, improvements and experiments, citing the evidence behind each one. Items without measured evidence are labeled exploratory.
- Automatic checks removed or corrected text with numbers not in the results, forecasts, guarantees, invented quotes or sources, budgets and causes the test can't separate.
- Versions fewer than 3 points apart in one run are reported as too close to call.

Limitations

- Answers come from AI personas built for this test, not from real people in the audience; group sizes are assumptions about the audience you described.
- Each test is one simulation run, so how stable these results are has not been measured.
- The versions differ in several ways at once, so the comparison shows which did better here, not which change caused it.

Automatic report checks

REPORT RULE

CHECK	RESULT
Raw counts add up across versions and groups	OK
Numbers in AI text match the results	OK
No variation or confidence claims (one run, not measured)	OK
No cause-and-effect claims the test can't separate	OK
No new claims recommended	OK
Your own claims marked as unverified where repeated	1 fixed

CHECK	RESULT
Competitor prices, market figures and citations removed	OK
Quotes match their answers (keyword check)	OK
One ad budget; no audience both targeted and excluded	OK
Verdict wording limited to what simulation can show	OK
No forecasts, guarantees or simulated answers presented as real research	OK

Consistency and wording checks on this report. They do not verify that its content is true.

Not verified

- Whether the claims in your description are true
- Competitor names, products, prices and details (no outside research was done)
- Real-world demand, sign-ups, sales or ad performance
- Stability across repeated runs (each test is run once)
- How closely simulated people in the audience match the real ones

AI models: answers claude-haiku-4-5-20251001 · write-up claude-sonnet-5-5. One simulation run.