Topic 07

University Business

What this course is, and how to use it

This is the seventh course in University Business: **marketing** – the discipline of creating, communicating and delivering value, measured honestly. The tool of the trade is a claim with a clean denominator: every metric in this handout says what it counts over what base, because a marketing number without its denominator is a mood.

The running case is **Nimbus Kitchen's subscription launch**: the founders take the £6.50 frozen meal into a direct-to-consumer weekly bundle. One campaign runs through all seven units – research, segmentation, product, price, channels, promotion and ethics – with every metric recomputable from the campaign sheet at the end.

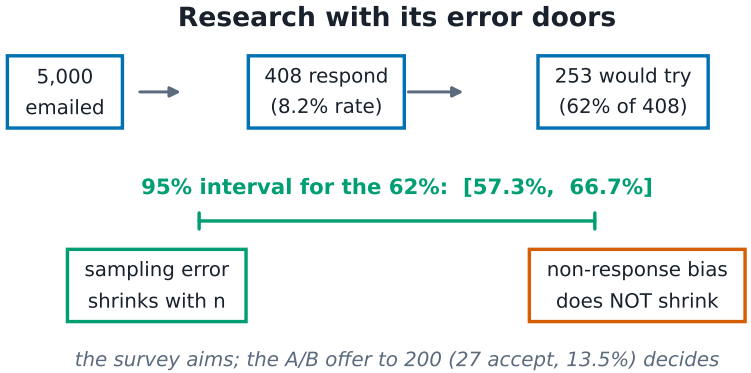
- Units 7.1–7.2 aim the campaign: research that can be wrong, segments that can be acted on.
- Units 7.3–7.5 build the offer: the product levels, the price, the channel economics.
- Units 7.6–7.7 run and judge it: the funnel with honest denominators, and the ethics that keep the numbers worth having.

How to read each unit: the question first, then the machinery, then two escalating worked examples, the mistakes people actually make, and a three-line summary.

Marketing research

Why it matters. The campaign is about to spend £3,600 on the belief that "people want weekly meal bundles". Research exists to make that belief testable – and to kill it cheaply if it is wrong.

The idea, with a picture. A **research question** must be answerable by some data you could actually collect; a hypothesis makes it falsifiable. Nimbus emails 5,000 past customers; 408 respond (an 8.2% **response rate** – already a warning: non-respondents differ), and 62% say they would try a weekly bundle.



the survey aims; the A/B offer to 200 (27 accept, 13.5%) decides

Every number carries its error doors - sampling on the interval, design in the response rate

The definitions. Sampling error shrinks with $\sqrt{n}$; the 95% interval for a proportion is $p \pm 1.96\sqrt{p(1-p)/n}$. **Non-response and selection bias** do NOT shrink – they are design facts. **Exploratory** research generates hypotheses; **descriptive** measures; **causal** needs experiments (A/B tests), not surveys.

The derivation. For $p = 0.62$, $n = 408$: $1.96\sqrt{0.62 \times 0.38/408} = 0.047 \rightarrow$ the interval is **[57.3%, 66.7%]** ✓. Stated honestly: "somewhere between 57 and 67 percent of people like the ones who answered would try it" – two hedges, both earned.

> **Worked example 1 – the question that could not fail.** First draft: "Do customers love our meals?" – no data could refute it. Second draft: "Of past customers, what share would order a 5-meal bundle at £29 within 30 days?" – measurable, falsifiable, and the survey's 62% intent becomes a hypothesis the launch will TEST (intent overstates purchase; unit 7.6's 15% trial rate is the real number). A question you cannot be wrong about is not research.

> **Worked example 2 – escalate to the design's limits.** The 8.2% response rate: who answers a food survey? Regulars – the 62% intent likely overstates the customer base. Two repairs: (a) weight by order frequency if the frame allows; (b) better, a small **A/B test** – offer the bundle to 200 random past customers and watch behaviour. 27 accept (13.5%): the behavioural number lands far below the attitudinal 62%, exactly as intent literature predicts. Survey for direction; experiment for truth.

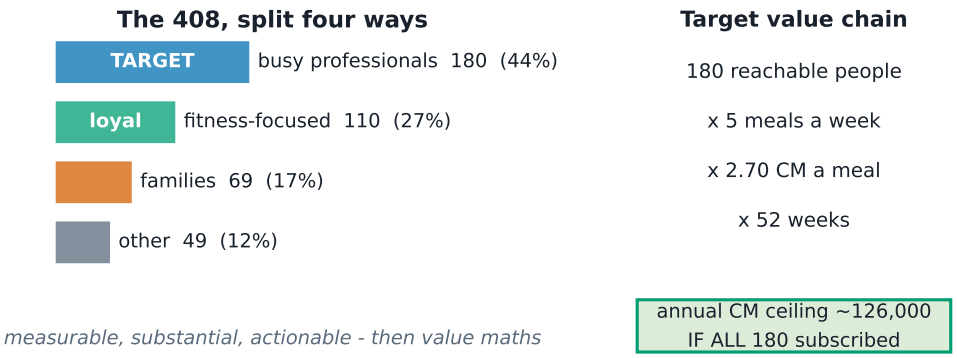
Common mistakes. Quoting 62% without the interval or the response rate. Treating a convenience sample of responders as the customer base. And asking "would you buy?" – intention is not behaviour, and the gap is the most replicated finding in market research.

Summary. Questions must be falsifiable; intervals carry sampling error; design carries bias. Surveys aim, experiments decide.

Segmentation and positioning

Why it matters. "Everyone" is not a target; it is the absence of one. Segmentation splits the market into groups the campaign can actually reach, and positioning chooses the race the product can actually win.

The idea, with a picture. The 408 responses split four ways: busy professionals 180 (44.1%), fitness-focused 110 (27%), families 69 (17%), other 49 (12.0%). A segment is useful only if it is **measurable, substantial, actionable** – the criteria, not the demographics.



Measurable, substantial, actionable - then the value maths decides the target

The definitions. **Segmentation** by behaviour, needs or demographics; **targeting** picks by attractiveness (size × growth × fit); **positioning** is the claim held in the customer's mind against competitors – defined by differences the firm can deliver. A **value proposition** states for whom, what benefit, why believed.

The derivation. Busy professionals: 180 respondents, 27% say they would pay for convenience NOW; the launch target is the reachable slice – call it 180 in the pilot city. Weekly value if every one of the 180 subscribed: 180 × 52 × £2.70 × 5 = £126,360 a year – the full-subscription ceiling, not the forecast (the 27% trial rate and churn do the discounting) ✓. The fitness segment (27%) is smaller but more loyal (64% weekly vs 44%); at scale it may beat the professionals on lifetime value. Segment value = size × conversion × retention × margin – never size alone.

> **Worked example 1 – the actionable test.** "Left-handed bakers" is measurable (rarely substantial). "People who hate cooking" is substantial (rarely actionable – no list, no channel). Busy professionals pass all three: countable (LinkedIn targeting), substantial (44% of responders), actionable (a channel and a message: "evenings back"). The criteria are a sieve, and each failure mode has a name.

> **Worked example 2 – escalate to positioning.** The competitor's frozen meal sits at £6.90 on shelf, positioned on "healthy". Nimbus cannot out-health the health brand; it CAN out-convenience everyone: the weekly bundle, the subscription, the 6 pm delivery slot. Position: "Tuesday evenings, solved" – a specific moment, a credible difference, a reason to believe (5 dinners, one decision). Check with a **perceptual map**: convenience (x) against health (y) puts the competitor high-left; Nimbus claims high-right. A position nobody can attack is a difference nobody can copy quickly.

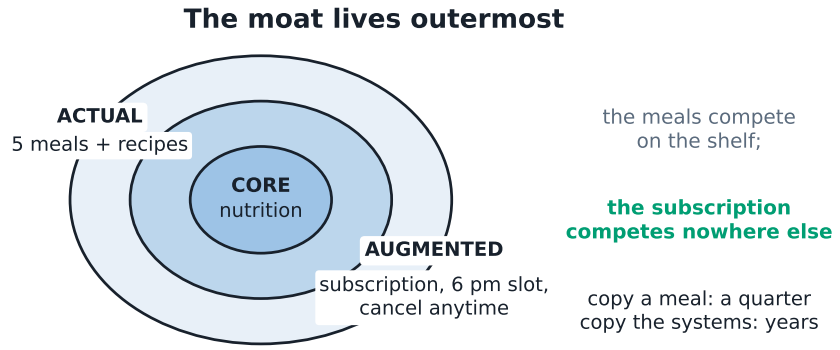
Common mistakes. Segmenting by demographics when the BEHAVIOUR differs. Choosing a segment the firm cannot reach (fails actionable). And positioning on "quality" – everyone says quality; a difference everyone claims is not a position.

Summary. Measurable, substantial, actionable – then value = size × conversion × retention × margin. Position on a deliverable difference, owned against the map.

Product and value

Why it matters. "The product" is more than the meal in the box: the bundle, the subscription, the delivery slot and the cancel-anytime promise are all product. Value is what the customer gets; the price (and unit 7.4) merely captures part of it.

The idea, with a picture. The **three product levels** : core (nutrition), actual (5 frozen meals, recipes, packaging), augmented (subscription flexibility, the 6 pm slot, the app). Nimbus's differentiation lives almost entirely in the augmented layer – the meals compete, the subscription does not.



The meals compete on shelf; the subscription competes nowhere else

The definitions. Core benefit; actual product; augmented product. **Value** = perceived benefits − perceived costs (money, effort, risk). The **value proposition** names the segment, the benefit and the proof. A **brand** is the promise plus the evidence kept.

The derivation. Customer's ledger for the bundle: benefits – 5 dinners solved (£35

of takeaway avoided), one decision a week, zero spoilage; costs – £29, remembering to freeze, novelty risk. Economic value ≈ £35 − £29 = £6 visible, plus time and mental-load benefits the price does not charge for – the surplus a price rise could later capture. Value math is the ceiling of price; the next unit stands on it.

> **Worked example 1 – the augmentation that carries the margin.** The competitor can copy a frozen meal in a quarter. Copying the 6 pm slot, the subscription logistics and the cancel-anytime trust takes years of systems. So the campaign LEADS with augmented features – "skip a week in one tap" is a meal the rival cannot serve. Product policy follows the defensibility, not the recipe.

> **Worked example 2 – escalate to the value test.** "Meal-kit fatigue" is a known churn driver. Nimbus's augmented answer: week-6 menu rotation and a two-minute "change my box" flow – designed against the churn data, not the chef's pride. The test: does the augmentation move retention? The pilot's month-2 retention is 66.4% (142 of 214 – unit 7.6); each +5 points of retention is worth roughly 214 × 0.05 × 5 × £2.70 × remaining-months – product decisions priced in cohort math, not adjectives.

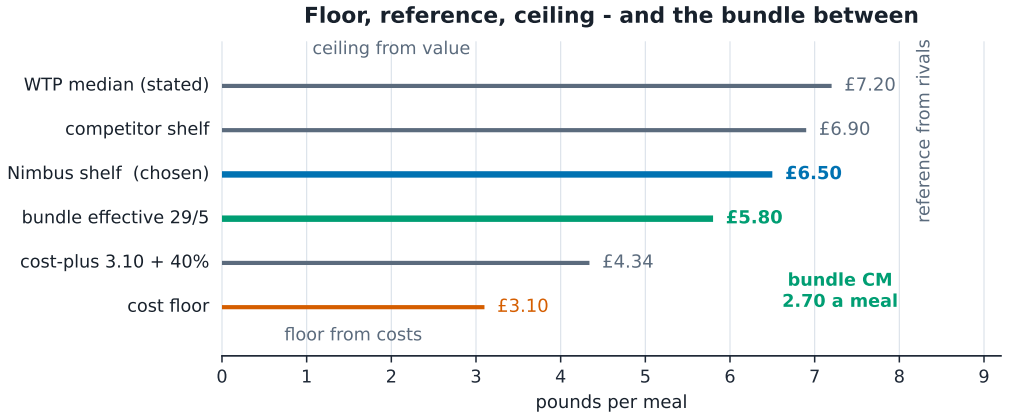
Common mistakes. Competing on the actual product when the augmentation is the moat. Confusing features with value (a benefit the customer does not perceive is a cost). And pricing before counting value – the ceiling comes first.

Summary. Core, actual, augmented – the moat lives outermost. Value = benefits − costs, and it caps the price. Augment against churn, priced in cohort math.

Pricing

Why it matters. Price is the only marketing lever that is pure revenue: a 1% price improvement flows almost untouched to profit. It is also where the cost, the competition and the value all meet – and where the honest marketer states which one they are following.

The idea, with a picture. Nimbus's ladders: cost-plus says £3.10 + 40% = £4.34 (underbids the value); the competitor sits at £6.90; the survey's median willingness-to-pay is £7.20. The shelf price stays £6.50 – but the SUBSCRIPTION bundle prices at £29 for 5 meals (£5.80 a meal, CM £2.70) with a two-meal minimum.



The floor from costs, the ceiling from value, the reference from rivals - and the bundle lives between

The definitions. **Cost-plus** covers cost plus a margin; **competition-based** follows rivals; **value-based** captures a share of the value created. **Penetration pricing** buys volume; **price skimming** skims the eager. Bundles trade price for commitment.

The derivation. Bundle economics: £29/5 = £5.80 effective per meal; CM = 5.80 − 3.10 = £2.70 ✓. The two-meal minimum protects the picking-apart of the bundle (buying 2 of 5 favourite meals at £11.60 undercuts nothing). Shelf £6.50 keeps the anchor high so the bundle reads as the deal: £32.50 à la carte → £29 saves £3.50 and one decision. Price architecture is arithmetic in the customer's head – use theirs, not yours.

> **Worked example 1 – which basis, when.** Cost-plus £4.34 leaves £2.86 of created value uncaptured and signals "cheap" against the £6.90 rival. Matching the rival at £6.90 forfeits the subscription story. Value-based: the bundle's visible surplus is £6 (unit 7.3) – capture part of it (£3.50) in the bundle discount, bank the rest as the customer's reason to subscribe. The honest sentence: "the floor is cost, the ceiling is value, the rival sets the reference – we set the architecture."

> **Worked example 2 – escalate to penetration with a deadline.** Launch month: first 100 subscribers get the 5-meal bundle at £24.80 (£4.96 a meal, CM £1.86). The maths: 100 × 5 × (2.70 − 1.86) = £420 of contribution deliberately spent to seed reviews, referrals and the cohort the funnel (unit 7.6) will judge. Penetration without a reversion date is just a price cut with a story; the sheet says "reverts to £29 in month 2" – and the churn at reversion is itself a measurement.

Common mistakes. Pricing the bundle by averaging shelf prices without the minimum-order protection. Confusing WTP STATED (survey £7.20) with WTP revealed (behaviour – always lower). And discounting forever: promotions are loans

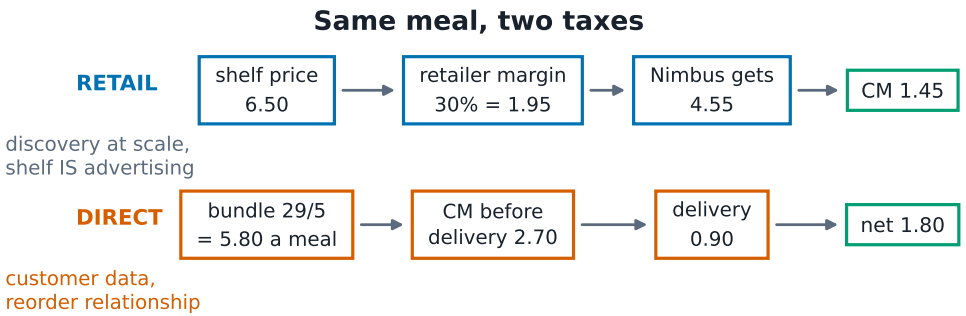
against the brand.

Summary. Floor = cost, ceiling = value, reference = competition. The bundle trades price for commitment at £2.70 CM. Penetration has a reversion date.

Channels

Why it matters. The customer meets the product through a channel, and the channel's margin is a tax on the story: retail takes 30%, direct takes logistics. Channel choice decides which promise survives contact with economics.

The idea, with a picture. Nimbus's two routes: retail (shelf £6.50, retailer pays £4.55 after 30% margin → CM £1.45 a meal) and direct subscription (£5.80 effective, CM £2.70 but with delivery costs). Same meal, different economics.



Same meal, two taxes - the retailer takes margin, the subscription takes logistics

The definitions. A **channel** moves product, information, promotion and money. **Channel margin** is each layer's cut. **Disintermediation** goes direct; **multichannel** runs both. **Conflict** arises when channels compete on price – the classic shelf-vs-subscription war.

The derivation. Retail: £6.50 × 0.70 = £4.55 to Nimbus; CM = 4.55 − 3.10 = **£1.45** ✓. Direct: £2.70 CM minus £0.90 per-box delivery = £1.80 net – richer than retail AND owns the customer data, the reorder relationship and the feedback loop. Retail still earns its place: discovery at scale, zero logistics, and the shelf IS advertising. The channel P&L prices each job the channel does.

> **Worked example 1 – the margin map.** One city, one retailer, 40 shelf slots at £6.50: 40 × £1.45 = £58 a restock of contribution, plus 600 passer-by impressions a week. The subscription: 214 subscribers × 5 × £1.80 = £1,926 a month. Different jobs, different scales – the map decides where each campaign dollar works, and the honest marketer prices the SHELF's advertising value before cutting retail for "low margin".

> **Worked example 2 – escalate to conflict.** The retailer sees the £29 bundle (£5.80/meal) undercutting its £6.50 shelf and threatens delisting. Three answers, honestly priced: (a) differentiate the bundle (5-meal commitment, delivery – a different product, defensible); (b) share the subscription in-store via QR with a finder's fee (£0.30/meal of the CM – still £1.50 net direct); (c) let the channels fight and lose one. Channel conflict is not a pricing error; it is a design decision with a bill. The chosen answer: differentiate, and say so to the retailer with the map on the table.

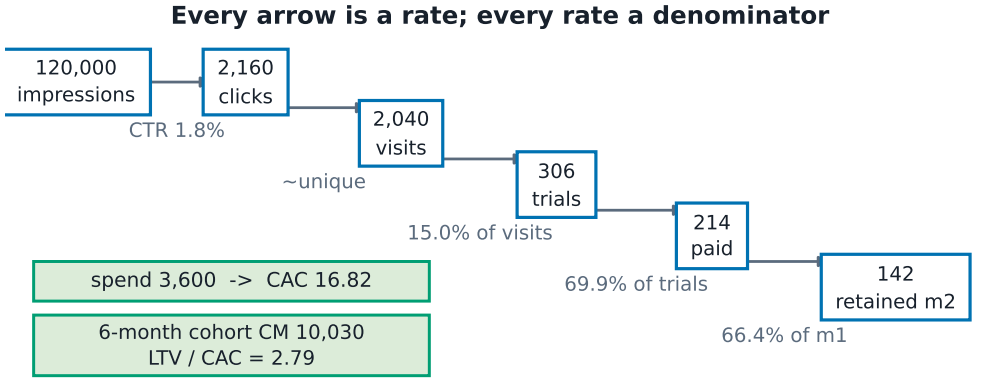
Common mistakes. Comparing channels on margin alone (each does different jobs). Letting direct price undercut the shelf partner without a differentiation story. And forgetting the channel's information flow: direct owns the data that makes unit 7.6's funnel measurable at all.

Summary. Retail: £4.55 in, £1.45 CM, discovery. Direct: £2.70 less logistics, plus the data. Conflict is designed, not stumbled into.

Promotion and digital measurement

Why it matters. £3,600 is about to be spent telling the segment the news. Digital marketing's gift is that everything is measurable; its curse is that everything measurable is not therefore meaningful. Honest denominators are the whole unit.

The idea, with a picture. The month-1 funnel: 120,000 impressions → 1.8% CTR = 2,160 clicks → 2,040 unique landing visits → 306 trial signups (15.0%) → 214 paid conversions (69.9%) → month-2 retention 142 (66.4%). Every arrow is a rate with a denominator, and every rate can be gamed.



Every arrow a rate, every rate a denominator - and CAC meets LTV at the bottom

The definitions. **CTR** = clicks/impressions; **conversion rate** = orders/visits (state which). **CAC** = spend ÷ new customers; **Lifetime value** (LTV) = cohort contribution over life. **Attribution** assigns credit across touches; last-click is the default and the confession. **ROAS** =

ad revenue/spend.

The derivation. CAC = 3,600/214 = **£16.82** ✓. Six-month cohort contribution at 34% first-month churn then 15%: £10,030 total → **LTV/CAC = 2.79** ✓ (month-1 CM £2,889; cohorts 214, 142, 121, 103, 88, 75 subscriber-months × £13.50). The healthy range most operators cite: ≥ 3 to absorb overheads – at 2.79 the campaign is close but not home; the lever is retention, not CTR: +5 points of month-2 retention is worth ~£900 of cohort contribution, which no plausible CTR gain buys.

> **Worked example 1 – the denominator police.** "15% conversion!" – visits or clicks? 306/2,040 unique visits = 15.0% ✓, but 306/2,160 clicks = 14.2%, and against impressions it is 0.26%. Each is true; only one is the campaign's steering number. Same discipline on CAC: £3,600 of PAID spend only – the founder's cousin's graphic design is a cost of the decision and belongs in it. Honest metrics name their numerator and denominator in the same breath.

> **Worked example 2 – escalate to attribution.** A subscriber saw a podcast ad, searched, clicked a retargeting ad, bought. Last-click gives retargeting the credit – so retargeting looks like the best channel and the budget migrates there, every quarter, forever. Two repairs: a holdout test (turn retargeting off for a random city – did sales fall by the claimed amount?) and incrementality thinking: credit a channel only for orders that would NOT have happened without it. The funnel measures the pipes; experiments measure the pressure.

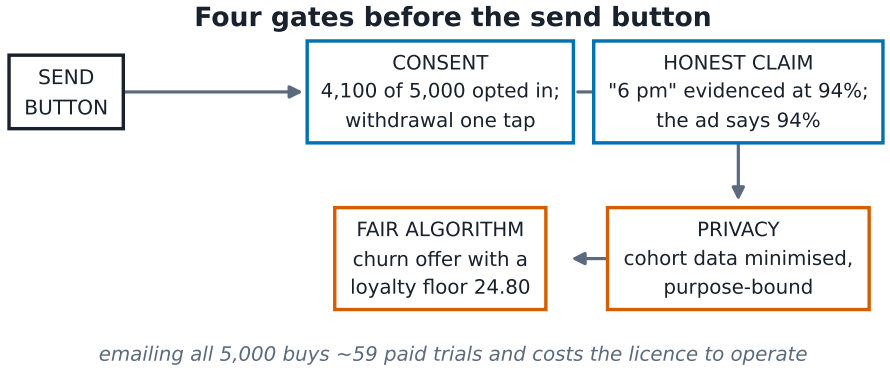
Common mistakes. Optimising CTR into clickbait (the rate rises as the denominator fills with the wrong people). Comparing CAC across periods with different channel mixes. And quoting LTV/CAC from churn assumptions nobody re-tests – the model is a hypothesis with a spreadsheet attached.

Summary. Rates are numerator-over-denominator, stated together. CAC £16.82 vs 6-month LTV/CAC 2.79 – retention is the lever. Attribution defaults lie; holdouts tell the truth.

Marketing ethics

Why it matters. The campaign's final asset is trust: the consent behind the email list, the honesty behind the ad claims, the fairness inside the algorithm. Every earlier number is only worth having if the data and the promises behind it were clean.

The idea, with a picture. The email list: 5,000 addresses, 4,100 with marketing consent. Contacting all 5,000 reaches 18% who never agreed – a small conversion gain bought with the campaign's licence to operate.



Four gates before the send button - each with its evidence line on the sheet

The definitions. Consent : opt-in, informed, withdrawable. **Honest advertising**: claims that are true, fair and substantiable – puffery ("delicious") is legal, measurable claims ("6 pm guaranteed") must be deliverable. **Dark patterns** : design that tricks (the hard cancel, the sneaky basket). **Privacy**: data minimisation and purpose limitation – collect what the purpose needs, use it for that purpose.

The derivation. The arithmetic of the gates: emailing all 5,000 adds 900 × (62% × 15% trial ≈ 84) ≈ 84 trials and ~59 paid (at 69.9%) – call it £1,000 of cohort contribution. The cost: regulator exposure, list poisoning, and the one-word review that outlives the campaign. Trust compounds; the sheet's final line prices it in retention, not fines.

> **Worked example 1 – the substantiable claim.** The ad draft says "always delivered by 6 pm". Operations says 94% by 6 pm. Three options: run it (an unsubstantiated absolute – one 6:15 delivery makes the ad false), soften to "delivered by 6 pm, 94% of the time" (true, checkable, and the number itself signals honesty), or fix operations first. The honest claim is not the cautious one – it is the one the evidence carries.

> **Worked example 2 – escalate to the algorithm.** Churn prediction flags the 72 subscribers likeliest to leave; the retention offer is £10 off. Should it go to (a) the 72 flagged, or (b) everyone? Option (a) is efficient – and punishes the loyal who do not churn, trains customers that threatening to leave pays, and if the model's error falls on the quiet, they never even know they were passed over. Fairness fix: a loyalty floor (never a worse offer than a new customer gets – unit 7.4's reversion priced this), an offer budget open to all at some level, and a human read of the model's misses. Algorithms inherit the objectives they are given; ethics is the objective statement.

Common mistakes. Treating consent as a tick-box (it is withdrawable – and the withdrawal must be as easy as the sign-up). Confusing legal with right. And measuring

only the campaign that ran – the people the gate PROTECTED are in the denominator of trust.

Summary. Consent is the licence; claims must be substantiable. Fairness is designed into offers and algorithms, not audited in afterwards. Trust compounds in the retention line.

The campaign sheet (reproducible)

Month 1, Nimbus Kitchen subscription launch:

- Email: 5,000 sent, 408 responses (8.2%), 62% intent, 95% CI [57.3, 66.7]; A/B offer to 200: 27 accept (13.5%).
- Funnel: 120,000 impressions; CTR 1.8% → 2,160 clicks; 2,040 unique visits; 306 trials (15.0% of visits); 214 paid (69.9%); month-2 retention 142 (66.4%).
- Economics: spend £3,600 → CAC £16.82; bundle £29/5 meals → £5.80 effective, CM £2.70/meal, delivery £0.90 → £1.80 net; 6-month cohort CM £10,030 → LTV/CAC 2.79. Shelf: £6.50, retailer pays £4.55, CM £1.45. Launch offer: first 100 at £24.80 (CM £1.86), reverts month 2.
- Segments of 408: professionals 180 (44.1%), fitness 110 (27%), families 69 (17%), other 49 (12.0%).
- Ethics gates: consent 4,100/5,000; "6 pm" claim evidenced at 94%; retention-offer fairness floor £24.80 for all.

Every figure in this handout recomputes from these values.

exercises.lol exercises.com.cn