

UNIT 3

Customers, Markets, and Marketing

How firms choose their customers, set their prices, and measure what a customer is worth

Image: Benh LIEU SONG, via Wikimedia Commons, CC BY-SA 2.0

By the end of this unit, I should be able to answer these questions:

- How does a company divide a market of millions of people into a handful of segments, and how does it choose which one to serve?
- Why can a 10 per cent price cut lower a company's revenue, and when does it raise it?
- What is a customer worth over the years, and how much may a company spend to win one?
- Why does a subscription business with a 90 per cent yearly retention rate look very different from one with 80 per cent, although the numbers seem close?
- What is a brand, in money, and where does the money a shopper pays go on its way back to the maker?
- Why does a well-designed survey still mislead, and what does a company do instead?

In 2023 the residents of Germany made about 251 million trips with at least one night away from home, which is roughly three trips for every resident. No hotel chain, airline, or tour operator sells to all of these travellers. Each of them chooses a part of this crowd, learns what that part wants, sets a price it will pay, and then tries to keep it coming back. The same choices face a bakery, a software company, and a maker of aspirin. This unit is about the customer: who the customer is, what the customer pays, what a customer is worth to a company over the years, and how a company finds all of this out without fooling itself.

1 The customer and the market

1.1 What marketing is

Most people hear the word marketing and think of advertising. Advertising is a small part of it. The American Marketing Association, the professional body of the field in the United States, defines **marketing** as "the activity, set of institutions, and processes for creating, communicating, delivering, and exchanging offerings that have value for customers, clients, partners, and society at large". In plainer words, marketing is the work of deciding whom a company serves, what it offers them, at what price, through which shops or websites, and how it tells them about it. An advertisement is one way of telling, and it comes late in the process, because the decisions about whom to serve and what to offer have been taken before it.

A **market**, in this unit, is the set of all actual and possible buyers of a kind of product. The market for overnight trips by German residents has about 251 million trips a year in it, counted by Eurostat, the statistical office of the European Union. A market this large is never uniform. A sales representative who flies to Hamburg for one night on a Tuesday wants something different from a family of four spending two weeks on Crete, although both count as trips.

1.2 Segmentation

Dividing a market into groups whose members want similar things and respond to a company's offer in a similar way is called **segmentation**, and each group is a **segment**. Three kinds of information are used to draw the lines:

- **Needs.** What the customer wants to get done. The representative needs a bed near the station for one night and an invoice for the employer; the family needs a beach, two rooms, and a price for fourteen nights. Need-based segments are the ones that matter for the product, because people with different needs buy different things.
- **Behaviour.** How often the customer buys, how much, through which channel, and how loyal they are. A hotel sees in its own records who books ten times a year and who books once.
- **Demographics.** Age, income, household size, education, and where people live. These are the easiest to measure, because statistical offices count them, and they are often the weakest guide, because two people of the same age and income can want completely different things.

In practice a segment is described in all three ways at once, by need, by behaviour, and by demographics, since a company must be able to find the people whose needs it serves. A segment that cannot be reached, because nobody knows who belongs to it, is of no use to anyone.

WORKED CASE Segmenting the German travel market

Eurostat publishes the trips of residents aged 15 and over by purpose, destination, and length. For Germany in 2023 the figures are those of Table 1. Purpose and length divide the 250.6 million trips into four segments, and the destination splits each in two.

Business trips are 34.3 million, or 13.7 per cent of all trips, and 24.1 million of them last one to three nights. Trips for personal reasons are 216.4 million, and the single largest cell of the table is the short personal trip inside Germany, with 78.5 million trips, which is 31 per cent of the whole market. Long personal trips abroad, with 77.2 million trips, are the second largest cell and the one a tour operator with charter flights and beach hotels serves. A chain of hotels in city centres lives from a different combination: the 19.6 million short business trips inside Germany, which fill rooms from Monday to Thursday, and the short personal trips, which fill them at the weekend. The largest segment of the market is therefore short personal trips inside Germany, 78.5 million trips a year, and a company that counts travellers instead of trips, or revenue instead of trips, will draw the lines elsewhere, because a two-week trip abroad brings a travel agent far more revenue than a weekend in Dresden.

Segment	In Germany	Abroad	All Trips
Personal reasons, 1 to 3 nights	78.5	16.4	94.8
Personal reasons, 4 nights or more	44.3	77.2	121.5
Business, 1 to 3 nights	19.6	4.5	24.1
Business, 4 nights or more	6.2	4.0	10.1
All trips	148.6	102.1	250.6

Table 1. Trips with at least one overnight stay made by residents of Germany aged 15 and over in 2023, in millions. Source: Eurostat, dataset tour_dem_tttot, read on 25 November 2026.

CHECK YOUR UNDERSTANDING Question 1

A streaming service finds that 60 per cent of its subscribers are between 18 and 34 years old. Which of the following would make the group aged 18 to 34 a useful segment in the sense of this unit?

(A) Its members respond to prices and to content differently from older subscribers

(B) Its members are the majority of all subscribers of the service

(C) Its members are easy to count from the date of birth on the sign-up form

(D) Its members were the first to subscribe when the service was launched

The solution is at the end of the unit.

1.3 Targeting and positioning

Once the segments are drawn, a company chooses which of them to serve. This choice is called **targeting**. A segment is worth targeting when it is large enough or growing, when its members can be reached with a message and a product, when the company has something to offer that rivals do not, and when the rivals already serving it are not too strong. A small company usually targets one segment, since serving two well costs twice the attention. Large companies serve several segments with separate products, and the five iPhone models Apple introduced in 2025 are an example.

Positioning is the place a product occupies in the customers' minds relative to the rival products, measured on the attributes customers use to compare them. The usual tool is a **positioning map**, a diagram with two such attributes as its axes and every product in the category drawn as a point. Figure 1 shows one for smartphones sold in the United States, drawn with two attributes a shopper sees before anything else, the price and the size of the screen. Three companies appear on it. Apple, the company based in Cupertino, California, that sells the iPhone, places five phones on the map, from the iPhone 16e at 599 dollars to the iPhone 17 Pro Max at 1,199 dollars; its revenue was about 479 billion US dollars in its financial year 2026. Alphabet, the parent company of Google, places three phones, and its revenue was about 403 billion dollars in 2025. Samsung Electronics of South Korea places many more, of which two appear here; its revenue was about 334 trillion won in 2025, which is about 235 billion dollars at the average exchange rate of that year.

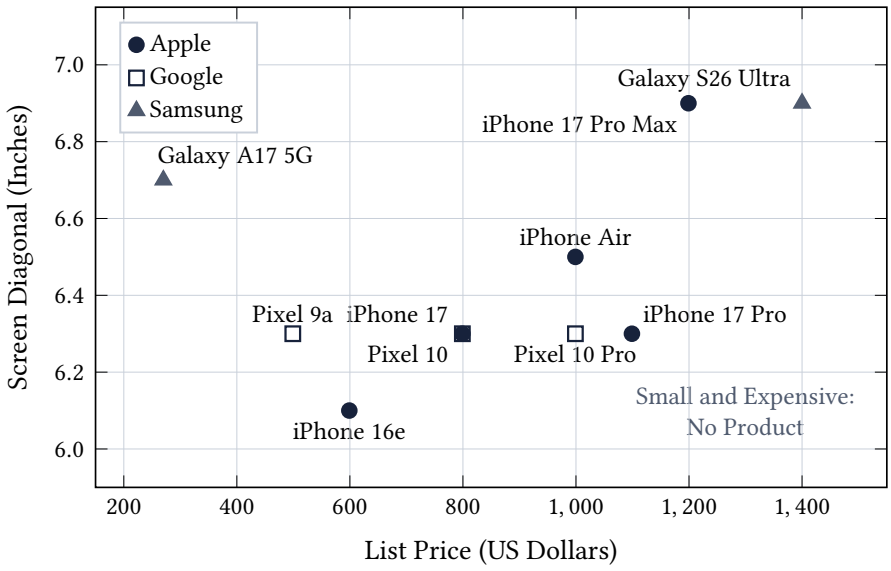

Figure 1. A positioning map of smartphones sold in the United States by list price and screen diagonal. Sources: Apple press releases of 19 February 2025 (iPhone 16e) and 9 September 2025 (iPhone 17, iPhone Air, iPhone 17 Pro, and iPhone 17 Pro Max); the Google Store and Samsung US product pages for the Pixel 9a, Pixel 10, Pixel 10 Pro, Galaxy A17 5G, and Galaxy S26 Ultra, read on 25 November 2026.

The map shows three things at once. The phones form clusters, which are the segments the companies target: a budget cluster below 500 dollars, a middle cluster at 799 dollars where the iPhone 17 and the Pixel 10 sit on exactly the same point, and a premium cluster from 999 dollars upward. The corner for a small and expensive phone is empty; Apple sold a 5.4-inch iPhone 13 mini until 2023 and then stopped, because too few customers wanted that position. And the map also shows the limit of its own two axes. The iPhone 17 Pro costs 300 dollars more than the iPhone 17 and has a screen of the same size, so the extra 300 dollars, which is called a premium, is paid for attributes that are not on the axes, such as the cameras, the materials, and the promise the name carries. Most positioning maps in practice use attributes that customers perceive rather than measure, such as quality or style, and the positions then come from a survey of customers. The two measured attributes were chosen here so that every point can be checked.

EXCURSION ***Marketing Myopia***

In 1960 the Harvard Business Review published an article by Theodore Levitt (1925-2006), who had joined the faculty of Harvard Business School the year before, with the title "Marketing Myopia". Its argument was about segmentation by need. The American railways, Levitt wrote, did not decline because the need to move people and goods declined; that need grew. They declined because they defined their business as the railway business rather than the transportation business, and so they let cars, trucks, and aeroplanes take their customers. Hollywood, in the same way, thought it was in the film business when it was in the entertainment business, and nearly missed television. A company that defines its market by the product it makes sees only the customers it already has; a company that defines it by the need it serves sees the rivals that are coming. The article became one of the most reprinted in the history of the journal, and the question it asks, what business are you really in, is still the first question a marketing plan has to answer.

2 What the customer pays

2.1 Demand and the price elasticity

Of all the decisions in marketing, the price is the only one that brings money in; every other one spends it. When the price of a product rises, fewer people buy it, and the relation between the price and the quantity bought is the **demand curve**. How steep that relation is decides whether a price change pays, and the measure of steepness is the **price elasticity of demand**, the percentage change in the quantity bought divided by the percentage change in the price that caused it.

KEY IDEA **The price elasticity of demand and the revenue rule**

$$\varepsilon = \frac{\Delta Q/Q}{\Delta P/P}$$

Q is the quantity sold and P the price per item; ΔQ and ΔP are the changes. The elasticity is negative, because the quantity falls when the price rises. Demand is called inelastic when $|\varepsilon| < 1$ and elastic when $|\varepsilon| > 1$. A price cut raises revenue $R = PQ$ only when demand is elastic, $|\varepsilon| > 1$; when demand is inelastic, a price cut lowers revenue and a price rise raises it.

For example, an elasticity of -0.79 means that a price rise of 10 per cent lowers the quantity sold by about 7.9 per cent, so the quantity reacts less than the price and demand is inelastic. The comparison between elasticities is always made in absolute value, so an elasticity of -2.6 counts as larger than -0.8.

The revenue rule follows from the definition, and the Case in Detail works it out. Figure 2 shows the same result as a picture. Revenue is price times quantity, so it is the area of the rectangle under the point on the demand curve. A price cut removes a strip along the top of the rectangle, the money lost on every item that would have sold anyway, and adds a strip along the right, the money earned on the new items. Whether the strip on the

right has a larger area than the strip at the top is exactly the question whether $|\varepsilon| > 1$.

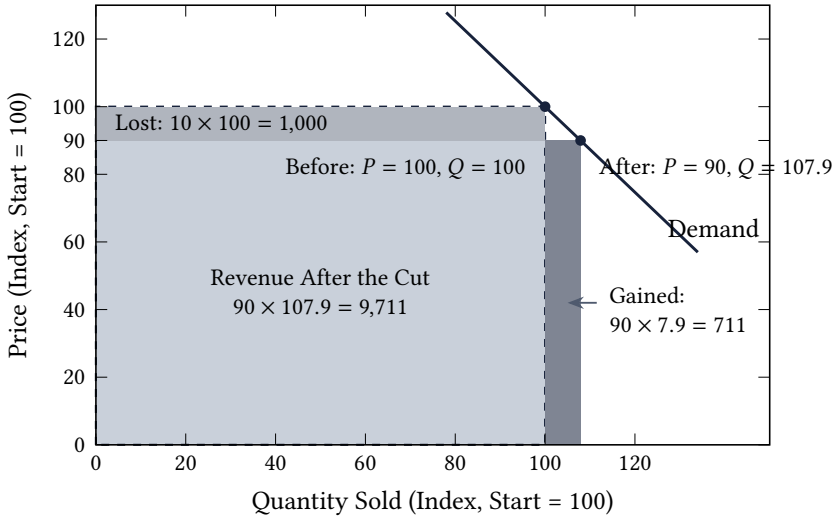

Figure 2. A 10 per cent price cut on a straight demand line with an elasticity of −0.79 at the starting point. Revenue is the area of the rectangle under the point on the demand line: the dashed outline is the revenue before the cut and the shaded area the revenue after it. The dark grey strip at the top is the revenue lost on the items that would have sold at the old price, and the dark strip on the right the revenue gained on the new items. Computed from $Q = 179 - 0.79\,P$.

IN DETAIL Why a price cut raises revenue only when demand is elastic

Let the price change by the fraction p, so that $p = -0.10$ means a 10 per cent cut, and let the quantity change by the fraction εp, which is what the elasticity says. The new revenue is

$$R_1 = P\,(1 + p) \cdot Q\,(1 + \varepsilon p) = PQ\left(1 + p + \varepsilon p + \varepsilon p^2\right).$$

For a small change the last term is small compared with the others, and the change in revenue is approximately

$$R_1 - R_0 \approx PQ\,p\,(1 + \varepsilon).$$

For a price cut p is negative, so the revenue rises when $1 + \varepsilon$ is negative, that is, when $\varepsilon < -1$, which is the same as $|\varepsilon| > 1$. For a price rise the sign of p flips, and the revenue rises when $|\varepsilon| < 1$. The exact product form gives the same answer with a small correction: with $p = -0.10$ and $\varepsilon = -0.79$ the approximation gives a change of $-0.10 \times 0.21 = -2.1$ per cent, and the exact product $0.90 \times 1.079 = 0.971$ gives -2.9 per cent.

Elasticities are measured, and the measured values carry a pattern worth remembering. A review of 160 studies published in 2010 in the American Journal of Public Health by Tatiana Andreyeva, Michael Long, and Kelly Brownell found that the demand for whole categories of food is inelastic, with mean elasticities between −0.27 for eggs and −0.81 for food eaten away from home, and −0.79 for soft drinks. Demand for a single brand inside a category is a different matter. A meta-analysis, which is a study that pools the results of many earlier studies, of 1,851 published brand-level elasticities from 81 studies, by Tammo Bijmolt, Harald van Heerde, and Rik Pieters in the Journal of Marketing Research in 2005, found an average of −2.62. The reason for the difference is **substitution**. When all soft drinks become dearer, people drink a little less of them; when one brand becomes dearer, people switch to the brand on the next shelf. The closer the substitutes, the more elastic the demand, which is why a company with a product that is the only one of its kind can raise its price and a company with an identical rival cannot.

WORKED CASE A 10 per cent price cut at two elasticities

Consider a company that sells 100,000 items a year at 100 US dollars each, so that its revenue is 10 million dollars, and let it cut the price to 90 dollars. If the company owned the whole category and demand had the soft drink elasticity of −0.79 from the review above, the quantity would rise by 7.9 per cent to 107,900 items and the revenue would be $90 \times 107{,}900 = 9.71$ million dollars, which is 2.9 per cent less than before. If the company sold one brand among several and demand had the brand-level average of −2.62, the quantity would rise by 26.2 per cent to 126,200 items and the revenue would be $90 \times 126{,}200 = 11.36$ million dollars, which is 13.6 per cent more.

Revenue is not profit, and the extra items cost money to make. Suppose each item costs 40 dollars in materials and labour, so that the company kept $100 - 40 = 60$ dollars per item before the cut as its contribution, which is what each item contributes towards the company's fixed costs and profit, and which is 6 million dollars a year in all. After the cut it keeps $90 - 40 = 50$ dollars per item. With the brand elasticity the contribution becomes $50 \times 126{,}200 = 6.31$ million dollars, a rise of 5.2 per cent, so the cut still pays, although far less well than the revenue figure suggested. With the category elasticity the contribution falls to $50 \times 107{,}900 = 5.40$ million dollars, a fall of 10 per cent. The cut therefore pays only for the brand with elastic demand, and even there it adds 5 per cent to the contribution, not 13.6. Unit 4 and Unit 5 explain where the cost of 40 dollars comes from and how the contribution reaches the income statement.

CHECK YOUR UNDERSTANDING Question 2

The demand for a product has a price elasticity of exactly −1.0. The company cuts the price by 10 per cent. What happens to its revenue?

(A) It rises by about 10 per cent, because the quantity rises by 10 per cent

(B) It falls by about 10 per cent, because every item earns 10 per cent less

(C) It stays roughly the same, because the extra items make up for the lower price

(D) It rises by about 1 per cent, because the quantity rises slightly more than the price falls

The solution is at the end of the unit.

2.2 Cost-plus and value-based pricing

Two habits of setting a price are found in practice. **Cost-plus pricing** starts from the cost of making one item and adds a percentage, the mark-up. A bakery that pays 1.20 euros for the ingredients and labour of a loaf and adds 50 per cent sells it at 1.80 euros. The method is simple, and it has one weakness, which is that the customer never sees the cost. A loaf is worth to the customer what the customer would pay for it, and that amount has nothing to do with the price of flour. The second habit, **value-based pricing**, starts from the customer instead: it asks what the product is worth to the people in the targeted segment and what the next best alternative costs them, and then sets the price between that value and the company's own cost.

KEY IDEA Value-based pricing

The price of a product is bounded above by the value the customer places on it, which is the most the customer would pay rather than do without it or buy the next best alternative, and bounded below by the cost of providing it. A price above the value sells nothing; a price below the cost loses money on every item. Where in that range the price lands depends on the rivals, on the segment, and on how much of the value the company chooses to leave with the customer.

The two habits give different answers whenever the value differs from the cost by a lot, which is the normal case. A software company whose extra copy costs almost nothing to produce cannot price from cost, and a

pharmacy that sells a branded aspirin for three times the price of an identical store brand, as Section 4 shows, is not pricing from cost either. Cost-plus pricing survives where value is hard to measure and in businesses with many similar items, such as a hardware shop, where the mark-up is really a rule for covering the shop's overheads, the rent and wages that do not depend on any single item. Value-based pricing requires the company to know its customers, which brings the subject back to segmentation.

2.3 Versions, time, and bundles

A single price throws value away. Some customers in a segment would have paid more, and some who would have paid a little less than the price buy nothing. Companies therefore sell the same thing at several prices to different people, a practice economists call **price discrimination** and marketers call differential pricing. It takes three usual forms:

1. **Versions.** The product is offered in a basic and a better version, and customers sort themselves. The iPhone 17 at 799 dollars and the iPhone 17 Pro at 1,099 dollars in Figure 1 are versions; so are a standard and a premium subscription, or a paperback and a hardback.
2. **Time.** The same seat or room is sold at different prices at different times. An airline sells the seats booked three months ahead cheaply to holidaymakers and the last seats dearly to business travellers who decide late; a cinema charges less on a weekday afternoon.
3. **Bundles.** Several products are sold together for less than the sum of their separate prices: a set menu, a software suite, a season ticket. Customers who value one part highly and another little pay for the bundle anyway, and the company sells more of the part they would not have bought alone.

All three work because the segments cannot easily trade with each other. A holidaymaker cannot sell a cheap early ticket to a late-booking business traveller, and a student discount is checked against a student card. Where resale is easy, the prices converge.

Prices also work on the mind, and the evidence for this is real but smaller than popular belief suggests. In field experiments with a mail-order clothing catalogue published in 2003, Eric Anderson and Duncan Simester found that the same item sold more at a price ending in 9 than at a lower round price, so that a dress at 39 dollars outsold the same dress at 34 dollars. Customers also judge a price against a reference, such as the price they paid last time or a crossed-out "was" price next to it, and a high reference makes the same price feel cheap. These effects are measurable, but they are modest, they vary from one product to another, and they do not rescue a product priced above its value.

3 What a customer is worth

A company that sells a subscription, a phone contract, a bank account, or insurance does not sell an item; it starts a relationship and earns a margin for as long as the relationship lasts. Three numbers describe such a relationship. The **customer acquisition cost**, written CAC, is what the company spends on advertising, sales staff, and discounts to win one new customer. The **retention rate** ρ is the share of this year's customers who are still customers next year, and its complement, $1 - \rho$, is the **churn**, the share who leave. The margin m is what one customer brings the company in a year after the direct costs of serving them. The field calls the arithmetic of these three numbers "unit economics"; this unit calls it the economics of one customer.

Put together, the three give the **customer lifetime value**, written CLV, which is the value today of all the margins a customer will bring in future years. Two things make the margins of the later years count for less. The customer may have left, which happens with a probability that grows year by year, and money that arrives in five years is worth less than money today. To discount a future sum means to reduce it to what it is worth today, which is why each year's margin is divided by $(1 + r)^t$, where r is the yearly rate at which a company discounts future money. Unit 6 explains where that rate comes from; here it is taken as 10 per cent a year.

KEY IDEA Customer lifetime value

$$CLV = \frac{m\,\rho}{1 + r - \rho}$$

m is the margin a customer brings in one year, ρ the yearly retention rate (the share of customers who stay from one year to the next), and r the yearly discount rate. The formula counts the margins of the years to come, from the first anniversary onward, and assumes that the margin and the retention rate stay constant for ever.

IN DETAIL The lifetime value as a geometric series

A customer who has just been won pays the margin m at the end of the first year if they are still a customer then, which happens with probability ρ; at the end of the second year with probability ρ^2; and at the end of year t with probability ρ^t. Each payment is divided by $(1 + r)^t$ to bring it to today's value. The expected value of all future margins is therefore

$$CLV = m\,\frac{\rho}{1+r} + m\,\frac{\rho^2}{(1+r)^2} + m\,\frac{\rho^3}{(1+r)^3} + \cdots = m\,(x + x^2 + x^3 + \cdots), \qquad x = \frac{\rho}{1+r}.$$

Call the sum in brackets S. Multiplying it by x shifts every term one place, $xS = x^2 + x^3 + \cdots$, and subtracting gives $S - xS = x$, so $S = x/(1 - x)$. The series adds up to a finite number because $x < 1$. Putting $x = \rho/(1 + r)$ back in,

$$CLV = m\,\frac{\rho/(1+r)}{1 - \rho/(1+r)} = \frac{m\,\rho}{1 + r - \rho}.$$

Without discounting, $r = 0$, the formula gives $m\rho/(1 - \rho)$, and the expected number of years a customer stays is $1/(1 - \rho)$, which is 5 years at 80 per cent retention and 10 years at 90 per cent. Two conventions are in use: this unit counts the first margin at the end of the first year, while many textbooks book it at once, which multiplies every value by $(1 + r)$.

WORKED CASE A subscription customer at 80 and at 90 per cent retention

A streaming service earns a margin of 60 euros a year on each subscriber after the cost of the content they watch and of serving them, and it discounts future money at 10 per cent a year. If 80 per cent of its subscribers renew each year, a new subscriber is worth

$$CLV = \frac{60 \times 0.8}{1 + 0.10 - 0.8} = \frac{48}{0.30} = 160 \text{ euros},$$

which is 2.67 years of margin. If 90 per cent renew, the same subscriber is worth

$$CLV = \frac{60 \times 0.9}{1 + 0.10 - 0.9} = \frac{54}{0.20} = 270 \text{ euros},$$

or 4.5 years of margin. The two retention rates look close, 80 and 90, yet the second customer is worth 270 euros against 160 euros, which is 69 per cent more. The reason is in the denominator: the retention rate is subtracted from $1 + r$, and a change from 0.8 to 0.9 cuts the denominator from 0.30 to 0.20. The same arithmetic sets the ceiling for acquisition. The service may spend at most 160 euros to win a subscriber at 80 per cent retention and 270 euros at 90 per cent before the subscriber loses it money, and in practice it will want to spend much less, because the margin of 60 euros a year and the retention rate are forecasts, and the money for acquisition is spent today.

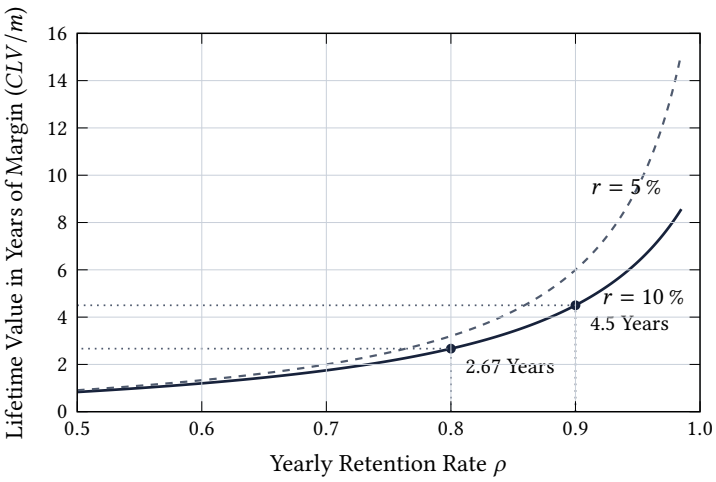

Figure 3. The lifetime value of a customer in years of margin, $CLV/m = \rho/(1 + r - \rho)$, against the yearly retention rate, for discount rates of 10 per cent (solid line) and 5 per cent (dashed line). The two marked points are the cases of the worked case. Above a retention rate of 90 per cent the curve rises steeply, which is why companies with subscriptions watch their churn so closely. Computed from the formula of the unit.

INTERACTIVE The lifetime value of a customer

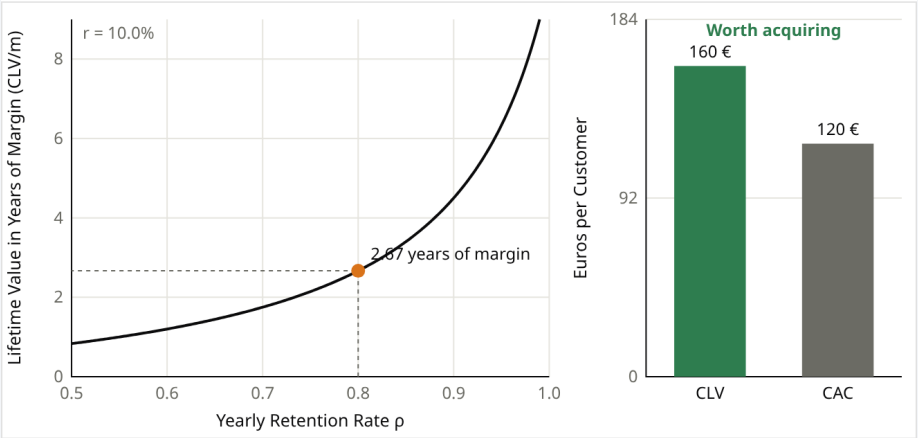

Move the retention rate and watch how many years of margin a customer is worth: the curve is the formula of the unit, \(CLV/m = \rho/(1 + r - \rho)\), for the chosen discount rate, and it bends upward sharply above 90 per cent. The bars on the right compare the lifetime value with the acquisition cost and turn red as soon as winning a customer costs more than the customer brings back.

Click the picture or this link to try it in the online version of this unit.

KEY IDEA The acquisition rule

A company may spend on winning a customer at most what that customer is worth: $CAC < CLV$. A rule of thumb used by investors in subscription businesses asks for a lifetime value of at least three times the acquisition cost, because the lifetime value is a forecast and the acquisition cost is paid in cash today. A company whose acquisition cost exceeds its lifetime value grows poorer with every customer it wins.

Figure 3 makes the shape of the formula visible. Below a retention rate of about 70 per cent the curve is flat and a customer is worth one or two years of margin; above 90 per cent it bends sharply upward, because the denominator $1 + r - \rho$ approaches r, and a customer who never leaves is worth m/r, which is 10 years of margin

at a 10 per cent discount rate. A company with a 95 per cent retention rate therefore has far more to lose from a small rise in churn than a company with 70 per cent, and it is the first kind of company that pays for loyalty programmes, customer service, and the small gifts that keep people from cancelling.

CHECK YOUR UNDERSTANDING Question 3

A software company earns a margin of 50 euros a year per customer, discounts at 10 per cent, and keeps 90 per cent of its customers each year. A better onboarding programme raises retention to 95 per cent. By how much does the lifetime value of a customer change?

(A) It rises by about 5 per cent, as the retention rate did

(B) It rises by about 40 per cent, from 225 to about 317 euros

(C) It doubles, because the expected lifetime doubles from 10 to 20 years

(D) It rises by about 20 per cent, because the denominator shrinks by a fifth

The solution is at the end of the unit.

The formula assumes that every customer has the same retention rate and that it stays the same for ever. Both assumptions are false in every real customer base, and the research spotlight shows what the first one does to a valuation.

RESEARCH SPOTLIGHT

Peter S. Fader *The Wharton School, University of Pennsylvania*

Why the average retention rate undervalues a customer base

A company that wants to value all of its customers at once, for a budget or for the sale of the company, will be tempted to take its overall retention rate and put it into the lifetime value formula of this section. Peter Fader and Bruce Hardie of London Business School showed in a paper published in Marketing Science in 2010 how far wrong that goes. They described an imaginary firm that wins 10,000 customers a year, whose contracts are annual and bring a net cash flow, which is the money left after the costs, of 100 US dollars each, and that has 26,569 customers at the end of 2007. Its overall retention rate is 69.1 per cent, and the textbook formula at a 10 per cent discount rate values the customer base at 4.9 million dollars. The paper counts each year's 100 dollars at the start of the year rather than at its end, so the formula of this unit gives 4.5 million dollars for the same firm; the error found below is the same under both conventions.

The flaw is in the average. Followed cohort by cohort, the retention rate of the customers won in 2003 rises year after year, from 63.3 per cent in the first year to 79.8 per cent in the fourth, although nobody becomes more loyal. The cohort consists of two kinds of people, a third with a retention probability of 0.9 and two thirds with 0.5, and the second kind leaves faster, so the survivors are increasingly the loyal ones. Valuing each cohort with its own future retention gives 7.9 million dollars, so the simple calculation undervalues the base by 38 per cent. With a statistical model of customer heterogeneity, which means that customers differ from one another, the shifted-beta-geometric model, the authors show that the error lies between 25 and 50 per cent in typical settings. The retention rate of Section 3 is the right tool for one customer and the wrong one for a mixed crowd.

Read the paper: Fader and Hardie (2010), Customer-Base Valuation in a Contractual Setting: The Perils of Ignoring Heterogeneity, Marketing Science 29, 85

4 Communicating and delivering value

4.1 The marketing mix

The set of decisions a company makes to bring its offer to the targeted segment is called the **marketing mix**, and the oldest list of those decisions is the one E. Jerome McCarthy (1928-2015) published in 1960 in his textbook *Basic Marketing*.

PRINCIPLES The four Ps of marketing (McCarthy, 1960)

A marketing plan has to decide on four things, each beginning with the letter P in English:

1. **Product**: what is offered, including its quality, its design, its versions, the packaging, and the service around it;
2. **Price**: what the customer pays, including discounts, payment terms, and the structure of versions, times, and bundles of Section 2;
3. **Place**: how the product reaches the customer, which means the channels, the shops, the logistics, and the inventory in them;
4. **Promotion**: how the customer learns about the offer, which covers advertising, personal selling, public relations, and sales promotions.

The four Ps are a checklist, not a theory. The list says what has to be decided and nothing about how, and later authors extended it for services with three more Ps, people, process, and physical evidence. Its value is that a plan which forgets one of the four has a hole in it.

The first two Ps, the product and the price, have had their sections already. The third and the fourth are where most of a marketing budget goes, and they are treated together here, because the people who deliver a product are often the ones who communicate it.

4.2 Channels and their margins

A **channel** is the path a product takes from the company that makes it to the customer who uses it. A company can sell directly, through its own shops, its website, or a sales force, or through intermediaries, which are wholesalers that buy in bulk and resell to shops, and retailers that sell to the public. Every intermediary takes a share of the price for its work, and that share is its **gross margin**, the difference between what it sells an item for and what it paid for it, usually stated as a percentage of its own selling price. The gross margin is not profit. It must pay for the shop's rent, the staff, the goods that are stolen or marked down, and the delivery van, and what is left after that is the profit.

WORKED CASE The path of a 100-dollar jacket and of 100 dollars of groceries

The US Census Bureau, the statistical office of the United States, surveys retailers and wholesalers every year and publishes their gross margins by kind of business. In 2022, clothing stores had a gross margin of 50.6 per cent of their sales, and wholesalers of apparel had one of 43.7 per cent of theirs. Follow a jacket that a shopper buys for 100 US dollars. The shop keeps 50.60 dollars and has paid the wholesaler 49.40 dollars. The wholesaler keeps 43.7 per cent of its own selling price, which is $0.437 \times 49.40 = 21.59$ dollars, and has paid the manufacturer the rest, $49.40 - 21.59 = 27.81$ dollars. The company that designed and made the jacket therefore receives 27.81 dollars of the 100 dollars the shopper paid, and out of that it pays for the cloth, the sewing, the shipping, and its own profit.

Groceries travel more cheaply. Grocery stores had a gross margin of 28.0 per cent in 2022 and grocery wholesalers one of 16.2 per cent, so of 100 dollars spent at the till the store keeps 28.00, the wholesaler keeps $0.162 \times 72.00 = 11.66$ of the 72.00 dollars the store paid it, and the food manufacturer receives

72.00 − 11.66 = 60.34 dollars. Figure 4 draws both paths. The difference is not greed. A clothing shop sells each jacket once, holds it for a season, marks down what is left, and needs a shop assistant per customer; a supermarket turns its inventory over every few weeks and sells thousands of items an hour with a dozen staff. The channel's margin is the price of the work the channel does.

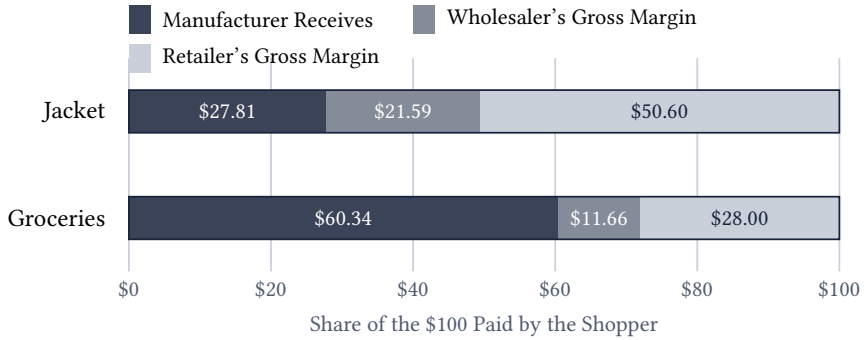

Figure 4. Where 100 US dollars paid by the shopper go, for a jacket bought in a clothing store and for groceries bought in a supermarket, when both pass through a wholesaler. Each step keeps its gross margin as a percentage of its own selling price. Source: US Census Bureau, Annual Retail Trade Survey (clothing stores 50.6 per cent, grocery stores 28.0 per cent of sales) and Annual Wholesale Trade Survey (apparel wholesalers 43.7 per cent, grocery wholesalers 16.2 per cent of sales), 2022 figures benchmarked to the 2022 Economic Census, read on 25 November 2026.

The arithmetic explains why so many companies have tried to sell directly to the customer. The maker of the jacket that receives 27.81 dollars through the channel receives the whole 100 dollars from its own website. It also inherits the channel's work and costs: the warehouse, the parcels, the returns, which in clothing can be a third of what is shipped, and above all the acquisition of customers. A shop on a high street brings in the customers who walk past for the price of its rent, while a website brings nobody until it is advertised. Whether direct sale pays is the question of Section 3 in another form: the acquisition cost of a customer against what that customer is worth. A brand that customers already seek out can afford its own channel; an unknown one usually cannot.

CHECK YOUR UNDERSTANDING Question 4

A retailer works with a gross margin of 40 per cent of its own selling price and sells a lamp for 50 euros. How much did it pay the wholesaler for the lamp?

(A) 20 euros, because the retailer keeps 40 per cent of what it paid

(B) 35.71 euros, because 50 euros divided by 1.40 gives the purchase price

(C) 70 euros, because the retailer adds 40 per cent to its purchase price

(D) 30 euros, because the retailer keeps 40 per cent of its selling price

The solution is at the end of the unit.

EXCURSION *Madam C. J. Walker builds a channel for a segment*

Madam C. J. Walker (1867-1919) was born Sarah Breedlove near Delta, Louisiana, the first child of her family born free, and was orphaned at seven. She worked as a laundress in St Louis for about a dollar a day. Around 1905 she began selling hair-care products for Black women, first as an agent for another company and then, from Denver, under her own name, which she took from her third husband in 1906. The segment she served was large and ignored by the department stores and the mail-order houses of the time, and she reached it with a channel of her own. In 1908 she opened a school in Pittsburgh, Lelia

College, to train “hair culturists”, and in 1910 she moved her headquarters to Indianapolis, where the Madam C. J. Walker Manufacturing Company had a factory, a salon, a laboratory, and a school on one site. Her agents demonstrated the products in customers’ homes and in salons, sold them on commission, and recruited and trained new agents; by 1917 the company claimed to have trained nearly 20,000 women. In the summer of 1917 she gathered 200 of them in Philadelphia for a convention, one of the first national meetings of women in business in the United States, and gave prizes to the agents who had sold the most and to those who had given the most to charity. The agents were the third P and the fourth at once, the distribution and the advertising, and the company reached customers whom no shop was trying to reach. Walker died at her estate in Irvington, New York, on 25 May 1919, and was then described as the wealthiest self-made Black woman in America.



Figure 5. Madam C. J. Walker, photographed by Addison N. Scurlock around 1914. Credit: Addison N. Scurlock, restored by Adam Cuerden, via Wikimedia Commons, public domain.

4.3 Brands

A **brand** is a name, a sign, or a design that identifies the seller of a product and carries a promise about what the buyer will get. The promise is the part of a brand that has a value in money. A buyer who cannot test a product before paying, which is the case for medicines, for a hotel room, and for a used car, runs a risk, and a brand that has kept its promise before lowers that risk and saves the buyer the time of checking. For that service buyers pay a premium, a price above that of the plain product, and the premium, multiplied by the quantity sold, is the brand’s value in money, year after year, less what it costs the company to keep the promise and to keep the name in people’s heads.

WORKED CASE The price of a brand in a pharmacy

Bart Bronnenberg, Jean-Pierre Dubé, Matthew Gentzkow, and Jesse Shapiro opened a paper published in 2015 in the Quarterly Journal of Economics with a price list. At cvs.com in 2013, a package of 100 tablets of Bayer aspirin, 325 milligrams each, cost 6.29 US dollars, and a package of 100 tablets of the CVS store brand, with the same dose, the same directions, and the same active ingredient, cost 1.99 dollars. Bayer is the German pharmaceutical and crop science company with a revenue of about 46.6 billion euros in 2024; CVS Health is the US pharmacy chain with a revenue of about 402 billion dollars in its financial year 2025. The brand premium was 4.30 dollars a package, a price 3.2 times that of the identical store brand.

Who pays the premium tells us what it is for. In the purchase records of a large panel of US households from 2004 to 2011, a group of households whose purchases are recorded over the years, the authors found that pharmacists chose a national brand of headache remedy 9 per cent of the time and the average shopper 26 per cent of the time. The people who knew that the two packages held the same thing bought the cheaper one. Across all packaged consumer goods, American consumers would have spent about 44 billion dollars a year less, 22 per cent of that spending, had they bought the store brand wherever one existed. For a company, the premium is revenue, and the sum of all such premiums in a year is what a

brand earns; a brand worth owning earns more than the advertising and the quality control that maintain it cost. Unit 2 treats the differentiation advantage from which such a premium comes.

5 Measuring marketing

5.1 Market research and its traps

Market research is the collection of the facts a marketing decision needs: how large a segment is, what its members want, what they would pay, and what they did after the last campaign. **Secondary data** already exist and are cheap, such as the Eurostat table of Section 1, the Census Bureau's margins, and a company's own sales records. **Primary data** are collected for the question at hand, by surveys, interviews, observation, and experiments, and the survey is the most used and the most misleading of these.

A well-designed survey, with a clear question put to a random sample of the right people, still misleads in three ways. The people who answer are not the people who were asked, because the ones with strong opinions or spare time reply more often, so the sample sorts itself. People do not know what they will do, so a stated intention to buy is a weak forecast of a purchase. And a price named in a questionnaire costs nothing to accept. A review by John List and Craig Gallet published in 2001 in the journal Environmental and Resource Economics compared experiments in which people first stated what they would pay for a product and then had the chance to buy it with real money, and found that the stated values were on average about three times the real payments, with a wide range from study to study. A survey in which 60 per cent of respondents say they would pay 5 euros a month for a new service is therefore not evidence that 60 per cent will.

What a company does instead is to watch behaviour and, where it can, to run an experiment. Behaviour is recorded in the sales data, in loyalty cards, and on a website, where every click is counted. An experiment offers two versions of a price, a page, or a letter to two randomly chosen groups of customers and compares what each group does; online this is called an **A/B test**, with the two versions as A and B. The random split is what makes it an experiment, because any difference between the groups other than the version is then a matter of chance. How such experiments are designed, how many customers they need, and where they go wrong goes beyond this unit.

5.2 The conversion funnel

Behaviour on a website is read with the **conversion funnel**, the sequence of steps from a visit to a purchase, with the number of people at each step: visitors, those who open a product page, those who put an item in the cart, those who start to pay, and those who pay. The share of people who move from one step to the next is the **conversion rate** of that step, and the share of all visitors who buy is the conversion rate of the shop. The funnel shows where customers are lost, and a company that improves the worst step by a few percentage points sells more without spending a cent more on acquisition.

WORKED CASE Reading the funnel of one online shop

A dataset donated to the machine learning repository of the University of California, Irvine, in 2018 by Can Okan Sakar and Yomi Kastro records 12,330 visits, called sessions in the file, to one online shop over a year, each by a different visitor. Of these sessions, 12,292 opened at least one product page and 1,908 ended in a purchase, so the shop's conversion rate was 1,908/12,330 = 15.5 per cent. The funnel in Figure 6 is drawn from these counts; the cart step is not recorded in the file.

The average hides the segments. New visitors bought in 24.9 per cent of their 1,694 sessions and returning visitors in 13.9 per cent of their 10,551 sessions. Sessions in November, with its sales before Christmas, converted at 25.4 per cent, sessions in May at 10.9 per cent, and the 184 sessions in February at 1.6 per cent. The shop therefore converted 15.5 per cent of its visits into purchases, with a range from 1.6 per

cent in February to 25.4 per cent in November, and a manager who judged a campaign by the overall rate would be judging the calendar. The missing cart step has been measured many times elsewhere. The Baymard Institute, a research company in Copenhagen that studies the usability of online shops, keeps a list of 50 published measurements of the share of carts that are abandoned before payment, from 2006 to 2025, and their average is 70.2 per cent. Roughly seven out of ten carts are left behind.

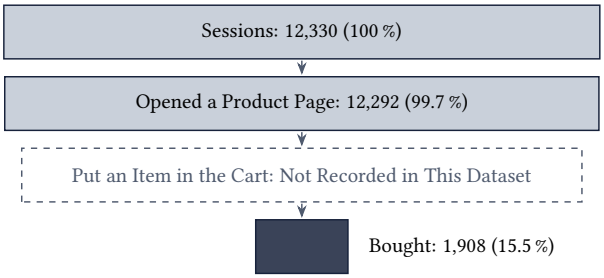

Figure 6. The conversion funnel of one online shop over a year: 12,330 sessions, of which 12,292 opened a product page and 1,908 ended in a purchase; the widths of the bars are proportional to the counts. The cart step is not recorded in this dataset. Source: Online Shoppers Purchasing Intention Dataset (Sakar and Kastro, UCI Machine Learning Repository, 2018, CC BY 4.0), counted from the file.

The funnel also closes the circle of the unit. The visitors at its top were won at an acquisition cost, the buyers at its bottom have a lifetime value, and the ratio of the two decides whether the shop's marketing earns its keep. A conversion rate of 15.5 per cent means that the acquisition cost of a buyer is 6.5 times the cost of bringing one visitor to the site, since $1/0.155 = 6.5$ visitors have to be brought to the site for each buyer, and that number, next to the lifetime value of Section 3, is the one a manager has to know.

KEY POINTS

- Marketing is the work of choosing a segment of a market, learning what it needs, and delivering value to it; segmentation, targeting, and positioning are its first three steps, and a positioning map shows where a product stands against its rivals on the attributes customers compare.
- The price elasticity of demand, $\varepsilon = (\Delta Q/Q)/(\Delta P/P)$, decides whether a price cut pays: revenue rises only when $|\varepsilon| > 1$. Whole categories are usually inelastic (soft drinks -0.79) and single brands elastic (an average of -2.62 across 1,851 estimates), because brands have close substitutes.
- A price is bounded above by the value to the customer and below by the cost; versions, time, and bundles sell the same product at several prices to customers who sort themselves.
- A customer is worth $CLV = m\rho/(1 + r - \rho)$, the discounted sum of future margins; at a 10 per cent discount rate a 90 per cent retention rate is worth 4.5 years of margin and 80 per cent only 2.67 years, and the acquisition cost must stay below this value.
- The four Ps, product, price, place, and promotion, are McCarthy's checklist of 1960. Every step of a channel keeps a gross margin on its own selling price, so the maker of a 100-dollar jacket receives about 28 dollars; a brand is a promise for which buyers pay a premium, 3.2 times the store-brand price in the case of Bayer aspirin.
- Surveys overstate what people will pay by a factor of about three; a company measures behaviour instead, with the conversion funnel and with randomized experiments.

Solutions to the questions in the text: 1-A, 2-C, 3-B, 4-D. The online version of this unit explains each answer.