

Scarcity, choice and opportunity cost

Build a clear method for comparing feasible alternatives, separating sunk costs from future consequences, and thinking at the margin. All scenarios are illustrative.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Identify scarcity and feasible alternatives without equating scarcity with poverty.
2. Calculate budget trade-offs and opportunity costs with consistent units.
3. Separate sunk, explicit and implicit costs without double counting.
4. Apply marginal reasoning while checking assumptions, uncertainty and values.

Slide 1: Choices begin with scarcity

Scarcity means that available time, resources and knowledge cannot satisfy every possible use. It does not mean that everything is rare or that everyone is poor. Economics studies choices under these limits. Throughout this lesson, every person, price and situation is illustrative. Start with this question: what must be given up?

Walkthrough

Imagine a completely free afternoon. Even without a cash expense, you cannot spend the same hour sleeping, visiting a friend and learning a skill in different places. Time is scarce because its uses compete. Wealth can relax some constraints, but it does not remove all of them. Economics begins by identifying the decision-maker, the available alternatives and the limits on those alternatives. This course pays particular attention to individual agency and voluntary arrangements. That interest does not settle every factual question: we must still examine evidence, assumptions and the consequences of different choices.

Slide 2: A feasible choice is an available choice

A feasible set contains the options that a person can actually choose, given relevant constraints. A budget, opening hours, physical capacity and existing commitments can all matter. Preferences rank feasible alternatives; they do not make an unavailable alternative available. Separate “I would like it” from “I can choose it now”.

Walkthrough

Suppose a learner wants to take three evening classes. Enough money is not sufficient if all three meet at the same time. Equally, an available timetable is not enough if a required prerequisite is missing. Write down each relevant constraint before recommending a choice. A constraint is binding at a particular choice when that choice uses the entire stated limit; a non-binding constraint has slack. Some models simplify this to money alone. That can be useful, provided we remember that the model deliberately leaves out other limits rather than proving that they do not exist.

Slide 3: Worked example: draw a budget boundary

An illustrative weekly budget is £36. Books cost £9 each and lunches cost £6 each. Let B be books and L be lunches: $9B + 6L$ must be no more than 36. All books gives 4 books; all lunches gives 6 lunches. Two books and three lunches exactly exhaust the budget.

Walkthrough

Step one is to name the quantities and their units: books and lunches per week. Step two is to multiply each quantity by its price. Two books cost 2 times £9, or £18; three lunches cost 3 times £6, also £18. Step three adds those expenditures: £36, so that bundle is feasible and lies on the budget boundary. One book and three lunches cost £27, leaving £9 unspent. Four books and one lunch cost £42, so that bundle is infeasible under the stated budget. We are holding prices fixed and allowing no borrowing.

Slide 4: Worked example: calculate the trade-off

On that budget boundary, one additional book uses £9 that could buy 1.5 lunches. Its opportunity cost is therefore 1.5 lunches when lunch spending is the displaced alternative. With indivisible items, swapping two books for three lunches gives an exact feasible exchange. A price ratio describes the boundary's trade-off, not the learner's preferences.

Walkthrough

Divide the book price by the lunch price: 9 divided by 6 equals 1.5 lunches per book. Check the units: pounds per book divided by pounds per lunch becomes lunches per book. Because half a lunch may not be purchasable, compare whole bundles as well. Starting from two books and three lunches, moving to four books and zero lunches gains two books and sacrifices three lunches. The ratio is still three divided by two. Whether that change is desirable depends on the learner's priorities. Feasibility and desirability are separate questions.

Slide 5: Opportunity cost is the best forgone alternative

Opportunity cost is the value of the next-best feasible alternative sacrificed by a choice. Do not add together mutually exclusive alternatives. If the same hour could earn £18 tutoring or £12 gardening, choosing a third activity for that hour does not forgo £30. It forgoes whichever available alternative the person ranks highest.

Walkthrough

List the alternatives before calculating anything. Tutoring and gardening cannot both occupy the same hour, so adding their earnings invents an opportunity that never existed. Non-cash alternatives also count. Someone might value resting above either paid activity. We cannot discover that ranking from wages alone. When we express benefits in pounds, we are using a stated or estimated money equivalent for this particular person, not claiming that every experience has a publicly observable price. Opportunity costs may differ between people facing the same shop prices because their alternatives and priorities differ.

Slide 6: Keep cash and economic cost separate

An explicit cost is a payment. An implicit cost is the value of an owned resource used instead of its next-best alternative. An afternoon project can use materials and unpaid time. Count each sacrifice once. Include an implicit time cost only when the comparison has not already accounted for the forgone activity.

Walkthrough

Suppose making a shelf costs £14 for materials and takes an afternoon that could otherwise earn £40 after relevant work expenses. Under these illustrative assumptions, the economic cost is £54. Calling the labour unpaid does not make its opportunity cost zero. However, if you compare the project's net outcome directly with the £40 work alternative, do not subtract £40 again on the other side. Consistent comparisons can use total outcomes or incremental differences. Both methods should give the same ranking when they include the same alternatives and avoid counting a sacrifice twice.

Slide 7: Sunk costs belong to the past

A sunk cost has already been incurred and cannot be recovered through the current decision. It can explain how someone reached a situation, but it does not change the relative future consequences of available options. A refundable deposit is not fully sunk. Resale value, cancellation charges and future commitments need separate checks.

Walkthrough

The practical test is not simply whether money has already changed hands. Ask whether choosing differently now would recover any of it. A non-refundable ticket has a sunk purchase price. A ticket that can be resold has a current opportunity cost: using it sacrifices its net resale value. An unpaid cancellation charge may also be an avoidable future cost. Decisions should reflect those differences. Ignoring a genuinely sunk cost does not excuse the original decision or prevent learning from it; it means refusing to let an irreversible past loss distort the present comparison.

Slide 8: Worked example: continue or leave?

An illustrative concert ticket cost £18 and cannot be refunded or resold. Attending now brings enjoyment valued at £25 and requires £8 of travel. Staying home brings enjoyment valued at £14 with no extra cost. Compare future net values: attending gives £17; staying gives £14. Attending has a £3 advantage.

Walkthrough

Step one removes the £18 ticket payment from both current alternatives because neither decision can recover it. Step two calculates attending's future net value: £25 minus £8 equals £17. Step three compares that with the £14 home alternative. The difference is £3 in favour of attending, given these stated valuations. If illness reduced concert enjoyment to £19, attending would instead give £11 and staying home would be preferable. The ticket price remains unchanged in both cases. The example demonstrates a decision rule; it does not prescribe how much anyone should enjoy music or rest.

Slide 9: Think at the margin

Marginal means additional: the benefit or cost of one more unit, hour or step. For a divisible activity, expand while the next increment's benefit exceeds its opportunity cost. For discrete options, compare actual available bundles. A large total benefit does not imply that every additional unit is worth its cost.

Walkthrough

A first glass of water may be extremely valuable, while a tenth glass in the same sitting may not be welcome. Total benefit and the benefit of the next glass answer different questions. Marginal reasoning focuses on the choice that remains open now. It is especially useful when deciding how much of an activity to do. The simple rule needs care when choices come in bundles, require a minimum scale, or affect future opportunities. In those cases, compare complete feasible plans rather than mechanically applying a one-unit rule that ignores connections between decisions.

Slide 10: Worked example: choose the next session

Four optional practice sessions have additional benefits valued at £31, £22, £12 and £5, in that order. Each session has a £15 opportunity cost. Net increments are £16, £7, minus £3 and minus £10. Taking the first two gives a total net benefit of £23. The third reduces that total.

Walkthrough

First calculate each extra session separately: 31 minus 15 is 16; 22 minus 15 is 7; 12 minus 15 is negative 3; and 5 minus 15 is negative 10. Then build feasible totals. Zero sessions gives zero; one gives 16; two give 23; three give 20; and four give 10. Two sessions is best among these options. This conclusion assumes sessions are taken in order, the stated benefits include relevant consequences, and no separate package discount applies. An earlier non-refundable enrolment fee would not change this current choice if it cannot be recovered.

Slide 11: Rules need their assumptions

A useful model states what is held fixed and what is omitted. The session example assumes diminishing additional benefits, constant additional cost and no hidden interaction. Real learning may have thresholds: a final session might unlock a qualification. When increments depend on one another, evaluate the relevant package before applying a marginal rule.

Walkthrough

Imagine that completing all four sessions produces a certificate worth an additional £30 to the learner. That changes the fourth session's consequences and therefore changes the original calculation. The earlier conclusion was not mathematically wrong; its assumptions no longer describe the choice. This is why good economic reasoning records the conditions behind a result. A model can make a complex question manageable, but simplicity is a tool rather than proof of universal validity. Ask which omitted feature could reverse the ranking, then investigate that feature instead of adding detail that cannot affect the decision.

Slide 12: Uncertainty changes what we compare

Many outcomes are uncertain. Expected value weights each possible outcome by its probability, but people may also care about risk, timing and worst cases. Gathering information has a cost and can improve a decision. More information is worthwhile only when its expected usefulness exceeds its relevant cost, including delay.

Walkthrough

Suppose an activity has a one-half chance of delivering a benefit valued at £40 and a one-half chance of delivering zero. Its expected benefit is £20, calculated as one-half times 40 plus one-half times zero. That does not mean everyone would exchange a certain £20 for this gamble. Preferences over risk and the ability to absorb a bad outcome matter. Estimates of probabilities can also be poor. Treat expected value as one transparent component of a decision, state where the probabilities came from, and avoid presenting an uncertain forecast as a guaranteed outcome.

Slide 13: Analysis and values answer different questions

A positive claim concerns what happens or could happen under specified conditions. A normative claim concerns what ought to happen or which values should guide action. Respect for individual choice is a normative commitment. Whether a particular rule expands feasible choices is an empirical or model-based question that still needs evidence.

Walkthrough

A libertarian approach typically gives substantial weight to individual liberty, property rights and voluntary cooperation. We can examine that position openly while distinguishing it from a prediction. For example, saying that adults should generally direct their own lives is a value judgement. Saying that a new restriction raises a particular cost is a causal claim requiring a mechanism and evidence. A person's choice reveals something about their ranking among perceived options, but it does not establish perfect information, freedom from pressure, or effects on other people. Keep those questions visible rather than assuming them away.

Slide 14: A reusable decision checklist

Name the decision and the person making it. Identify feasible alternatives and relevant constraints. Compare future benefits and opportunity costs, counting each sacrifice once. Ignore only costs that are genuinely sunk. Examine the next increment or complete package. State uncertainty, affected people and the values behind your final recommendation.

Walkthrough

Use the checklist on a modest decision before applying it to a complicated policy debate. A clear comparison of two available study plans can teach more than a confident claim about an entire economy with no stated assumptions. Notice where arithmetic ends and judgement begins. The calculations organise implications of the information supplied; they cannot create missing preferences or reliable probabilities. In the exercises, numerical and choice responses can be checked automatically. Short written explanations are explicitly ungraded reflection. Use them to state an assumption or limitation, not to guess a hidden keyword that supposedly proves understanding.

Questions students ask

If a ticket was expensive, should I use it to avoid wasting money?

Only if using it is your best current option. An unrecoverable payment cannot be changed now. Check refunds or resale first, then compare the future consequences of attending and doing something else.

Does choosing something prove that it is good for everyone?

No. A choice may indicate the chooser's ranking among perceived options. It does not by itself establish informed consent, absence of mistakes, effects on others, or a shared moral judgement.

Does economics say I should maximise money?

No. Money can make some comparisons easier, but rest, relationships, learning and other aims can enter a person's preferences. The method asks for a coherent comparison, not a particular life goal.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. A new budget boundary

Illustrative case: you have £84 for kits costing £12 each and meals costing £7 each. No borrowing or other spending is allowed. Give each all-in maximum and the budget-boundary opportunity cost of one kit in meals. Use the price ratio for the last answer, even though purchases are whole items. Round the opportunity cost to two decimal places.

Maximum kits (kits)

Maximum meals (meals)

Opportunity cost of one kit (meals per kit)

2. Reconsider an outing

Illustrative case: an outing ticket cost £27 and cannot be refunded or resold. From now, attending brings enjoyment valued at £24 and costs £9 in travel. Staying home brings enjoyment valued at £12 and costs nothing extra. Calculate attending's future net value and its advantage over staying home. Then select the better current option under these assumptions.

Future net value of attending (£)

Attending minus staying home (£)

Better current option

1. Attend
2. Stay home
3. Equal net values

3. How many practice blocks?

Illustrative case: four optional practice blocks must be taken in order. Their additional benefits are valued at £28, £21, £13 and £7. Each costs £16 including time. There are no package benefits or discounts. Choose the number of blocks that maximises total net benefit and calculate that net benefit relative to taking none.

Best number of blocks (blocks)

Total net benefit (£)

Ungraded reflection: describe one changed assumption that could alter your choice.

4. Find the next-best option

Illustrative case: one afternoon allows exactly one activity. Tutoring has net value £45, repairing your furniture has net value £38, and resting has value £30, all according to your own stated valuations and including relevant costs. If you choose tutoring, what is its opportunity cost and its advantage over the next-best alternative?

Opportunity cost of tutoring (£)

Advantage over the next-best alternative (£)

Ungraded reflection: why should the two forgone alternatives not be added together?

Specialisation and voluntary exchange

Calculate comparative advantage, construct feasible trades and distinguish possible mutual gains from stronger claims about real-world outcomes. All cases are illustrative.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Distinguish absolute advantage from comparative advantage.
2. Calculate opportunity costs and feasible terms of trade.
3. Verify production, transfers and each participant's gains step by step.
4. Explain transaction costs, distributional effects and the limits of a voluntary-exchange model.

Slide 1: Exchange can rearrange possibilities

People can gain from exchange because they value goods differently, face different opportunity costs, or both. Specialisation can also change what a group produces. These are related but distinct sources of gain. Every example in this lesson is illustrative. A voluntary agreement shows expected gains for participants under their information, not guaranteed gains for everyone.

Walkthrough

Consider two people who already own different books. Swapping books may benefit both without producing another book. Now consider two makers who can rearrange their working time before trading. They may increase attainable output as well. We will keep these mechanisms separate. An agreement is voluntary in the relevant sense only when the parties can meaningfully refuse and are not subject to force or fraud. Even then, uncertainty or mistaken beliefs can make an outcome disappointing. Economics can explain the scope for mutual benefit without treating every observed transaction as proof of a desirable outcome.

Slide 2: Absolute advantage asks about productivity

Absolute advantage means producing more output with the same resources, or using fewer resources for the same output. It is measured good by good. Comparative advantage instead asks which producer sacrifices less of another good. A producer can be more productive in every activity yet still benefit from exchange with someone less productive.

Walkthrough

Suppose one person types and edits faster than another. It is tempting to conclude that the faster person should do everything. That ignores their limited time. If their relative edge in typing is much larger than their edge in editing, allocating some editing to the other person can release time for valuable typing. The important comparison is not simply who is better. It is what each person must give up to do the next task. Our numerical examples make that comparison explicit while holding quality, working time and technology fixed.

Slide 3: Worked example A: set up production

Ada and Ben each have eight working hours. Ada makes four loaves or two metres of cloth per hour. Ben makes one loaf or one metre of cloth per hour. Within this illustrative model, each person can split time freely and has constant productivity. Ada has an absolute advantage in both goods.

Walkthrough

Start with the resource constraint: each person has eight hours, not unlimited labour. Ada could make 32 loaves if she used every hour for bread, or 16 metres if she used every hour for cloth. Ben could make eight loaves or eight metres. These endpoints describe feasible production, not consumption preferences. The word “or” matters: producing bread uses hours that cannot simultaneously produce cloth. We assume goods of equal quality, no setup time and no transport cost for now. Those simplifying assumptions let us isolate the opportunity-cost mechanism before adding complications.

Slide 4: Worked example A: compute opportunity costs

For Ada, one loaf costs half a metre of cloth: two divided by four. For Ben, one loaf costs one metre: one divided by one. Ada has comparative advantage in bread. Conversely, one metre costs Ada two loaves and Ben one loaf. Ben has comparative advantage in cloth despite lower absolute productivity.

Walkthrough

Calculate what an hour could do in the alternative activity, then divide by output in the chosen activity. Ada gives up two metres when she spends an hour making four loaves. Two divided by four is half a metre per loaf. Ben gives up one metre to make one loaf. The smaller sacrifice identifies comparative advantage in bread. Reverse the ratios for cloth: Ada sacrifices two loaves per metre, while Ben sacrifices one. In this two-person, two-good constant-productivity example, different opportunity costs create a role for each producer in a more specialised allocation.

Slide 5: Worked example A: establish a baseline

Initially, each person spends four hours on each good. Ada produces sixteen loaves and eight metres. Ben produces four loaves and four metres. Combined production is twenty loaves and twelve metres. With no trade, each consumes their own output. Record this baseline before claiming that a different arrangement benefits either person.

Walkthrough

For Ada, four bread hours multiplied by four loaves per hour gives sixteen loaves. Four cloth hours multiplied by two metres per hour gives eight metres. Ben's corresponding calculations give four loaves and four metres. Add like units only: sixteen plus four gives twenty loaves; eight plus four gives twelve metres. Do not add loaves and metres into an unexplained grand total. We now know the original bundles. To demonstrate a gain without knowing detailed preferences, we can seek an arrangement that gives each person more of at least one good and no less of the other.

Slide 6: Worked example A: reorganise and exchange

Ada switches to six bread hours and two cloth hours, producing twenty-four loaves and four metres. Ben makes cloth for eight hours, producing eight metres. Ada then exchanges six loaves for four of Ben's metres. Final bundles: Ada has eighteen loaves and eight metres; Ben has six loaves and four metres.

Walkthrough

First check production: Ada's six times four gives twenty-four loaves, and two times two gives four metres. Ben's eight cloth hours yield eight metres. Total output is now twenty-four loaves and twelve metres, an increase of four loaves with cloth unchanged. Next subtract what each gives and add what each receives. Ada has 24 minus 6 loaves and 4 plus 4 metres. Ben has six loaves and 8 minus 4 metres. Compared with the baseline, each gains two loaves without sacrificing cloth. We have demonstrated feasible mutual gains under the model's assumptions.

Slide 7: Terms of trade divide the available gain

The exchange rate was four metres for six loaves, or two-thirds of a metre per loaf. That lies between Ada's bread opportunity cost of half a metre and Ben's cost of one metre. Rates strictly between these costs can benefit both at the margin, provided desired quantities, feasibility and exchange costs also permit agreement.

Walkthrough

Ada can make a loaf by sacrificing half a metre. Receiving more than half a metre for it offers an advantage over producing that cloth herself. Ben would sacrifice one metre to make a loaf, so buying one for less than a metre saves resources. The precise agreed rate affects who receives more of the gain. Comparative advantage alone does not determine bargaining power or a unique market price. Nor does it establish how many units each person wants. At an endpoint of the interval, one party may be indifferent before transaction costs rather than strictly better off.

Slide 8: Worked example B: a different technology

Mira produces four app screens or two tests per hour. Sol produces one screen or three tests per hour. Each has six hours. One screen costs Mira half a test and Sol three tests. With three hours on each task initially, Mira produces twelve screens and six tests; Sol produces three screens and nine tests.

Walkthrough

Step one calculates opportunity costs from hourly rates: Mira gives up two tests for four screens, while Sol gives up three tests for one screen. Mira therefore has comparative advantage in screens and Sol in tests. Step two establishes each initial bundle using three hours per activity. Mira makes three times four screens and three times two tests. Sol makes three times one screen and three times three tests. The initial combined output is fifteen screens and fifteen tests. As before, the tasks have comparable quality and constant productivity, and these are illustrative task units rather than claims about actual software work.

Slide 9: Worked example B: verify the gains

Full specialisation gives Mira twenty-four screens and Sol eighteen tests. Mira trades six screens for nine tests. Mira finishes with eighteen screens and nine tests; Sol finishes with six screens and nine tests. Each improves on the baseline. The exchange rate of 1.5 tests per screen lies between their opportunity costs.

Walkthrough

First multiply six working hours by the specialised production rates: Mira produces 24 screens and Sol produces 18 tests. Then account for the transfer. Mira retains 24 minus 6 screens and receives nine tests. Sol receives six screens and retains 18 minus 9 tests. Mira gains six screens and three tests relative to her original bundle. Sol gains three screens with tests unchanged. This example permits full specialisation with gains for both. The bread example used partial specialisation instead. Full specialisation is not a universal recommendation; the useful allocation depends on technologies, preferences and the quantities needed.

Slide 10: Trading uses resources too

Finding partners, comparing quality, negotiating, transporting goods and enforcing agreements all have costs. In the bread example, each participant gained two loaves before exchange costs. If arranging the exchange costs each the equivalent of one loaf, each retains a one-loaf net gain. Larger costs can eliminate the advantage or prevent trade entirely.

Walkthrough

Transaction costs are not merely fees charged by a middleman. They include time spent searching, uncertainty about quality and resources needed to resolve a dispute. A platform or broker may charge money yet reduce the total cost of exchange by saving more elsewhere. Conversely, an apparently free exchange can be expensive in time. In our numerical illustration, subtract the one-loaf equivalent from each person's two-loaf gain. The net gain is one loaf each. This comparison assumes those equivalent costs reflect their own valuations and do not hide uncompensated costs imposed on someone else.

Slide 11: Preferences and alternatives still matter

Lower opportunity cost creates a production opportunity; it does not force anyone to want the resulting goods. Potential partners may disagree on quality, quantity or timing. Their outside alternatives shape acceptable terms. A mutually beneficial reallocation exists only when the final bundles are feasible and preferred after all relevant costs are considered.

Walkthrough

A person who dislikes bread may not value additional loaves in the way our simple illustration assumes. We demonstrated gains using more of one good and no less of another, but even that comparison presumes the extra good is welcome and can be stored, used or passed on without a burden. In broader applications, preferences determine which points on a production boundary are useful. Alternative trading partners also change bargaining positions. Comparative advantage is a powerful explanation of possible cooperation, but it is one component of a complete account of actual exchange.

Slide 12: More total output need not mean everyone gains

An increase in aggregate output does not by itself show that each person benefits. Prices, ownership and adjustment costs affect distribution. A worker may lose a valued occupation even when consumers gain. Possible compensation is different from compensation actually received. Assess the relevant people and transition costs rather than treating a group total as everyone's experience.

Walkthrough

Our worked examples carefully specified final individual bundles. Many public arguments do not: they jump from an estimated aggregate gain to the claim that nobody loses. A more productive arrangement may change wages, asset values, locations and the usefulness of acquired skills. Some losses are temporary; others can persist. Whether institutions should compensate people, and who should pay, introduces normative questions as well as factual ones. A liberty-oriented analysis can place strong weight on voluntary agreements and secure rights while still recognising that distributional outcomes and effects on non-participants require separate examination.

Slide 13: Institutions support or obstruct cooperation

Clear ownership, reliable information and workable contracts can reduce uncertainty about what is exchanged and on what terms. Competition among alternatives can improve options. These are mechanisms to investigate, not guarantees of perfect outcomes. A rights-based judgement about permissible exchange differs from a positive claim about productivity, prices or measured welfare.

Walkthrough

If a seller cannot establish the right to transfer an item, a buyer faces a risk that ownership will later be challenged. If quality cannot be inspected, reputation or a warranty may help. Each arrangement has costs and limits. The appropriate institutional comparison is between realistic alternatives, including their enforcement and information requirements. It is not enough to compare an imperfect market with an imaginary institution that makes no mistakes. Equally, describing an exchange as voluntary does not answer questions about fraud, harms to third parties or the origin of the rights being exchanged.

Slide 14: Audit a claim about gains from trade

Check resources, production rates, units and opportunity costs. Specify the baseline, the new production plan and the exchange. Verify that no goods appear from nowhere. Compare each person's final bundle after transaction costs. State which assumptions support the claimed gains, which people are included, and what the calculation does not establish.

Walkthrough

The strongest elementary demonstration is an accounting exercise followed by a preference comparison. Production must fit available hours, transfers must balance across parties, and final quantities must add up to production after any losses. Only then ask whether each person prefers the outcome. This discipline prevents persuasive stories from outrunning their numbers. In the exercises, you will work with new technologies and exchange terms. A checked numerical response verifies the calculation under the stated assumptions. The ungraded reflections invite you to identify limitations, because accurate arithmetic alone cannot establish that every real-world exchange is beneficial or just.

Questions students ask

Can the less productive person still have something useful to offer?

Yes. Comparative advantage depends on relative opportunity costs rather than absolute output alone. The less productive person can specialise where their relative disadvantage is smaller.

Does comparative advantage require complete specialisation?

No. Desired quantities, increasing opportunity costs, risk, capacity and exchange costs may favour partial specialisation. The production plan must be checked against the actual problem.

Does voluntary trade prove that the agreed price is morally fair?

It can indicate that participants expected to prefer agreement to their perceived alternatives. Moral fairness, adequate information, the legitimacy of rights and effects on outsiders are additional questions.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Compare two makers

Illustrative case: Lea makes six mats or four baskets per hour. Noor makes three mats or five baskets per hour. Calculate each maker's opportunity cost of one mat in baskets, then identify comparative advantage in mats. Assume constant productivity and equal quality. Round numerical answers to two decimal places.

Lea's opportunity cost (baskets per mat)

Noor's opportunity cost (baskets per mat)

Comparative advantage in mats

1. Lea
2. Noor
3. Equal opportunity costs

2. Check a production reallocation

Illustrative case: Emi makes two shirts or four scarves per hour; Ivo makes one shirt or one scarf. Each has ten hours. Initially each spends five hours on each good. A new plan gives Emi four shirt hours and six scarf hours, and Ivo eight shirt hours and two scarf hours. Calculate the new combined output and the increase in each good over the initial plan.

New combined shirts (shirts)

New combined scarves (scarves)

Increase in shirts (shirts)

Increase in scarves (scarves)

3. Account for a complete trade

Illustrative case: Rina makes five pots or one tile per hour; Sam makes two pots or two tiles. Each has six hours. Rina spends all six hours on pots and Sam all six on tiles. Rina then gives Sam nine pots in return for five tiles. There are no losses or exchange costs. Calculate both final consumption bundles.

Rina's final pots (pots)

Rina's final tiles (tiles)

Sam's final pots (pots)

Sam's final tiles (tiles)

Ungraded reflection: what extra information would help you compare this outcome with a no-trade situation?

4. Deduct the costs of exchange

Illustrative case: a buyer values a used camera at £140. The seller is indifferent between keeping it and selling it for £90 before exchange costs. They agree on £110. The buyer incurs £12 of transaction costs and the seller £8. Calculate the buyer's net gain and seller's net gain relative to no trade. Assume the item is accurately described and no one else is affected.

Buyer's net gain (£)

Seller's net gain (£)

Ungraded reflection: explain why a gain for both parties does not determine a unique bargaining price.

Demand, supply and changing prices

Solve simple market schedules, distinguish movements from shifts, and explain what equilibrium can and cannot establish. All markets are illustrative.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Define demand, supply and equilibrium with consistent units and time periods.
2. Distinguish movements along a schedule from shifts of that schedule.
3. Solve linear equilibria and calculate shortages or surpluses.
4. Analyse separate and simultaneous changes while recognising causal and normative limits.

Slide 1: A market model organises plans

Demand and supply describe quantities people plan to buy and sell at different prices over a specified period. They are schedules, not single quantities. An equilibrium price makes those planned quantities equal within the model. All markets and numbers in this lesson are illustrative; actual markets may have search, contracts, queues and bargaining.

Walkthrough

Imagine a weekly market for boxes of identical notebooks. Buyers have different uses for notebooks, and sellers face different costs and capacities. Rather than describe every person separately, we summarise their plans at each possible price. The model asks what quantity would be demanded and supplied if that price prevailed while specified other conditions remained unchanged. It does not assume that every quoted price is actually observed or that all planned trades occur. Define the product, quality, location, period and units before combining different people's plans into market schedules.

Slide 2: Demand relates quantity to the good's own price

A demand schedule gives quantity demanded at each own price, holding relevant other conditions fixed. Our illustrative market has $Q_d = 60 - 2P$, where P is pounds per box and Q_d is boxes per week. At £10, quantity demanded is forty boxes; at £15, it is thirty boxes.

Walkthrough

Read the equation as a conditional plan. Substitute 10 for P: 60 minus twice 10 equals 40 boxes per week. Substitute 15: 60 minus 30 equals 30. The negative coefficient represents a downward-sloping demand relationship over the relevant range. It does not mean people always buy fewer boxes whenever a higher price is observed in the real world; other conditions might have changed too. This linear schedule is intended only for prices that imply meaningful non-negative quantities. Extending it indefinitely to predict negative purchases would misuse the model.

Slide 3: A movement is different from a shift

A change in the good's own price causes a movement along an unchanged demand curve. A change in income, preferences, buyer numbers, expectations or related-good prices can shift demand. A rightward shift means more is demanded at a given own price. Demand is a schedule; quantity demanded is one point on it.

Walkthrough

Suppose the notebook price falls while incomes, preferences and other relevant circumstances stay fixed. The original schedule tells us the new quantity demanded: that is a movement along it. Now suppose a new group of students enters the market. More boxes might be wanted at every notebook price, requiring a different schedule. That is a shift. Income effects depend on the good: demand for a normal good rises with income, while demand for an inferior good can fall. These labels describe responses to income, not a moral or physical judgement about the product's quality.

Slide 4: Supply describes sellers' conditional plans

A supply schedule gives quantity supplied at each own price, holding technology, input prices and other relevant conditions fixed. Here $Q_s = 12 + P$. At £10, sellers plan twenty-two boxes per week; at £15, twenty-seven. The upward slope represents a larger planned quantity at a higher price within the illustrative range.

Walkthrough

Substitute the price into the supply equation just as we did for demand. Twelve plus ten gives twenty-two boxes; twelve plus fifteen gives twenty-seven. The positive relationship can reflect additional production becoming worthwhile as price covers higher marginal costs. It may also involve inventories or entry over the relevant time horizon. A market supply schedule need not be upward sloping in every possible setting. Our simple model assumes a competitive setting in which individual sellers take price as given and can adjust quantity. Later models can examine market power, fixed capacity and strategic pricing.

Slide 5: Supply shifts have different causes


Lower input costs or improved technology can increase supply: more is offered at each price. Higher costs or lost capacity can reduce it. A change in the product's own price is a movement along supply, not a supply shift. Time matters because firms can adjust some inputs quickly and others only slowly.

Walkthrough

Suppose the price of paper falls while notebook prices remain unchanged. Producing boxes may become worthwhile for sellers who previously found it too costly. This is an outward supply shift. By contrast, a rise in the notebook price with unchanged production conditions moves sellers along the existing schedule. Be precise about the time horizon. Tomorrow's supply may depend on stock already printed, whereas next year's supply can include new equipment and firms. The same event may therefore have different short-run and long-run effects. A useful prediction specifies which adjustment opportunities are available.

Slide 6: Worked example: solve the equilibrium

Set the planned quantities equal: $60 - 2P = 12 + P$. Subtract 12 and add $2P$ to get $48 = 3P$. Divide by three: $P = £16$. Substitute into either schedule: $Q = 28$ boxes per week. Check both: $60 - 32 = 28$ and $12 + 16 = 28$.



Equilibrium is where both plans agree: 28 boxes per week at £16 per box.

Walkthrough

Step one writes the equilibrium condition: quantity demanded equals quantity supplied. Step two collects the constant terms on one side and price terms on the other. Starting with $60 - 2P = 12 + P$, subtracting 12 gives $48 - 2P = P$. Adding $2P$ gives $48 = 3P$. Step three divides both sides by three to obtain 16 pounds per box. Step four substitutes this price into both original equations. Both give 28 boxes per week, so the calculation is internally consistent. Equality is a model condition, not a guarantee that any actual price instantly adjusts to it.

Slide 7: Worked example: diagnose a shortage

At a price of £12, $Q_d = 60 - 24 = 36$ boxes, while $Q_s = 12 + 12 = 24$ boxes. Excess demand is twelve boxes per week. This is a shortage at that price: planned purchases exceed planned sales. It is different from scarcity, which exists even when the market clears.

Walkthrough

First evaluate both schedules at the same price and over the same period. Second subtract supplied quantity from demanded quantity: $36 - 24 = 12$. That difference measures unmatched purchase plans, not necessarily a recorded queue of exactly twelve boxes. Actual sales may be limited by available supply and by matching or other frictions. Buyers might compete through higher offers, waiting or search effort. Which adjustment occurs depends on the market's rules. Scarcity remains at the equilibrium too, because producing and buying notebooks still uses resources that have alternative uses even when no excess demand appears in this model.

Slide 8: Worked example: diagnose a surplus


At £20, quantity demanded is twenty boxes and quantity supplied is thirty-two. Excess supply is twelve boxes per week. Sellers' plans exceed buyers' plans at that price. A surplus does not mean the goods have no possible use; it means the specified price and conditions do not reconcile those purchase and sale plans.

Walkthrough

Calculate demand as $60 - 2 \times 20$, which gives 20. Calculate supply as $12 + 20$, which gives 32. Subtract demand from supply: $32 - 20 = 12$. Some sellers may cut prices, reduce output, store goods or seek other buyers. Our static model does not specify the speed or exact process of adjustment. Nor does a planned surplus prove that twelve physical boxes have already been produced and wasted. Supply is a planned quantity at a price; the timing of production and whether unsold inventories exist require additional information.

Slide 9: Worked example: a demand shift

Suppose demand becomes $Q_d = 72 - 2P$ while supply remains $Q_s = 12 + P$. At each price, buyers now want twelve additional boxes. Equilibrium satisfies $72 - 2P = 12 + P$, giving $P = £20$ and $Q = 32$. Demand shifted; the resulting price change causes movement along unchanged supply.



An outward demand shift moves equilibrium from (28, 16) to (32, 20).

Walkthrough

Step one identifies what changed: the demand intercept rises from 60 to 72, while its price coefficient and the entire supply schedule remain unchanged. Step two equates the new demand with the old supply. Rearranging gives $60 = 3P$, so price is 20. Step three substitutes: $12 + 20$ gives 32 boxes. Compared with the original equilibrium, price rises by four pounds and quantity by four boxes. Quantity demanded also changes, but calling this only a movement along demand misses the initial change in conditions. The equilibrium comparison uses two different demand curves.

Slide 10: Worked example: a supply contraction

Return to original demand, $Q_d = 60 - 2P$. Now supply becomes $Q_s = 6 + P$, six fewer boxes at each price. Equating schedules gives $60 - 2P = 6 + P$, so $P = £18$ and $Q = 24$. A contraction of supply raises price and lowers equilibrium quantity in this model.

Walkthrough

Begin from the original baseline rather than silently carrying forward the previous demand shift. The new equality rearranges to $54 = 3P$, which gives a price of 18. At that price, demand is $60 - 36$, or 24 boxes. Supply is $6 + 18$, also 24. Relative to the original price of 16 and quantity of 28, price increases by two pounds and quantity falls by four boxes. The higher price partly offsets the supply contraction through movements along the schedules. It does not undo the underlying loss of supply at each given price.

Slide 11: Several changes can occur together

If demand rises while supply falls, both changes tend to raise price in this model. Their quantity effects oppose one another, so quantity's direction cannot be inferred without relative magnitudes. With $Q_d = 72 - 2P$ and $Q_s = 6 + P$, equilibrium is £22 and twenty-eight boxes, the original quantity at a higher price.

Walkthrough

Combine the two shifted schedules and solve rather than adding verbal predictions mechanically. Seventy-two minus $2P$ equals six plus P becomes $66 = 3P$, so price is 22. Supply then gives six plus 22, or 28 boxes. Demand gives the same result. In this particular numerical case, the quantity effects exactly offset. Other shift sizes could raise or lower quantity. A qualitative diagram often identifies some directions but leaves others ambiguous. Saying that an effect is ambiguous without additional information is a valid analytical conclusion, not a failure to understand the model.

Slide 12: Observations do not identify the cause automatically

A higher observed price and quantity are consistent with a demand increase under an unchanged upward-sloping supply curve. They do not prove that demand was the only thing that changed. Quality, income, costs, seasonality and market boundaries may also differ. A causal explanation needs a justified account of what moved and what stayed fixed.

Walkthrough

Two observed transactions cannot reveal an entire demand curve if demand or supply shifted between them. For example, a festival could increase both the number of buyers and sellers' staffing costs. Observed price and quantity would reflect both mechanisms. To investigate causality, look for evidence about the changes in relevant conditions, a credible comparison group, or a research design that separates influences. Even then, report uncertainty. The phrase "holding other things constant" is a condition of a model comparison; it is not something we can assume simply because a spreadsheet records only price and quantity.

Slide 13: Equilibrium is not a moral verdict

Market clearing means planned demand equals planned supply under specified conditions. It does not establish justice, adequate information, equal opportunity or absence of harm to outsiders. Efficiency claims need additional assumptions about competition, rights, information and external effects. Whether an outcome should be accepted also depends on values beyond this equality.

Walkthrough

A price can clear a market even when some people cannot afford the product, when ownership is disputed, or when production imposes pollution costs on neighbours. Those issues are not disproved by the intersection of two curves. Under particular competitive assumptions, economists can make stronger claims about gains from trade and allocation, but the assumptions must be stated. A liberty-oriented perspective may prioritise voluntary exchange and secure rights; another perspective may weigh distribution differently. Clear analysis separates the mechanism that determines an outcome from the ethical reasons for supporting, changing or rejecting the institutions around it.

Slide 14: Use the model carefully

Define the good, period and units. Write the demand and supply schedules and their valid ranges. Solve for equality and check both sides. For any proposed price, distinguish shortage from surplus. When conditions change, identify the shifted schedule before calculating a new equilibrium. Finally, state assumptions, uncertainty and questions the model leaves unanswered.

Walkthrough

A diagram and an equation can tell the same story if they use consistent axes and definitions. Conventionally, price appears on the vertical axis and quantity on the horizontal axis, even though our equations express quantity as a function of price. You can practise without a graph by checking quantities at several prices. The exercises use new schedules and require explicit numerical answers. Reflection fields are ungraded: they are a place to explain which conditions must remain fixed or why observed correlations are insufficient. Avoid converting a correct algebraic solution into an unsupported claim about a real market.

Questions students ask

If price rises and purchases fall, has demand decreased?

Not necessarily. If the other conditions of demand are unchanged, this is a decrease in quantity demanded along the same schedule. A decrease in demand means a change in the schedule itself.

Is scarcity the same as a shortage?

No. Scarcity concerns competing uses of limited resources. A shortage is excess demand at a specified price and can disappear as conditions change, while scarcity remains.

Must real markets always be at equilibrium?

No. Search frictions, contracts, adjustment delays and other features can prevent or slow market clearing. A static equilibrium is a useful conditional benchmark, not a complete description of every transaction.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Find a new market equilibrium

Illustrative weekly market: $Q_d = 80 - 3P$ and $Q_s = 8 + P$. P is pounds per box; quantities are boxes per week. Solve for the equilibrium price and quantity. Use only the non-negative economically relevant range.

Equilibrium price (£ per box)

Equilibrium quantity (boxes per week)

2. Diagnose two quoted prices

Illustrative weekly market: $Q_d = 90 - 2P$ and $Q_s = 10 + 2P$. At £15 and separately at £25, identify whether there is a shortage or surplus and calculate its positive magnitude. Treat these as plans at each specified price.

At £15

1. Shortage
2. Surplus
3. Balanced

Magnitude at £15 (units per week)

At £25

1. Shortage
2. Surplus
3. Balanced

Magnitude at £25 (units per week)

3. New buyers enter

Illustrative weekly market: demand changes from $Q_d = 100 - 2P$ to $Q_d = 115 - 2P$. Supply remains $Q_s = 10 + P$. Calculate the new equilibrium and identify the initial change.

New equilibrium price (£ per unit)

New equilibrium quantity (units per week)

Initial change

1. Rightward shift of demand
2. Movement along the original demand curve only
3. Leftward shift of supply

4. A loss of available output

Illustrative weekly market: $Q_d = 96 - 3P$ and initial $Q_s = 16 + P$. A cost change reduces quantity supplied by twelve units at every price, leaving the slope unchanged. Calculate the new equilibrium. Then explain, in the ungraded field, why observing a higher price alone would not prove that this cost change caused it.

New equilibrium price (£ per unit)

New equilibrium quantity (units per week)

Ungraded causal explanation

Elasticity, revenue and who bears a tax

Measure proportional responses, connect demand elasticity to revenue, solve a tax-wedge example and separate causal responses from observed correlations. All examples are illustrative.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Calculate signed and absolute midpoint price elasticity with correct bases.
2. Distinguish elasticity from slope and interpret the revenue relationship.
3. Solve a basic per-unit tax model using separate buyer and seller prices.
4. Evaluate the assumptions behind elasticity estimates and causal claims.

Slide 1: Percentage changes make responses comparable

A change of ten units means different things in a market selling twenty units and one selling ten thousand. Percentages express a change relative to a base. Elasticity compares percentage changes in two variables. All data in this lesson are illustrative. Distinguish a percentage change in a rate from a percentage-point change.

Walkthrough

If a quantity rises from 50 to 55, the change is five units and the starting-base percentage increase is five divided by 50, or ten per cent. If an interest rate rises from four per cent to five per cent, it rises by one percentage point but by 25 per cent relative to its initial value. These are different measures. Elasticity uses proportional changes, allowing comparisons across units such as boxes, kilograms or litres. Always identify the relevant base and sign before calculating. A calculator cannot correct an ambiguous definition of the percentage you intend to measure.

Slide 2: Demand elasticity is a conditional response

Price elasticity of demand is the percentage change in quantity demanded divided by the percentage change in own price, along an unchanged demand relationship. Ordinary downward-sloping demand gives a negative signed elasticity. We usually classify its absolute magnitude: above one is elastic, below one is inelastic, and exactly one is unit elastic.

Walkthrough

An elasticity of minus two means that, locally, a one per cent price increase is associated with about a two per cent reduction in quantity demanded, other relevant conditions unchanged. The absolute magnitude is two. Be explicit about whether you report the signed value or its magnitude. Elastic does not mean that quantity is high; it means quantity responds proportionally more than price. An inelastic response is not necessarily zero. Perfectly inelastic demand has zero quantity response in the specified model and range, a much stronger condition than merely having a magnitude below one.

Slide 3: Worked example: use the midpoint method

Price rises from £10 to £12 and quantity falls from 100 to 80 units on one demand curve. Average price is £11; average quantity is 90. Midpoint changes are $2/11 = 18.18\%$ and $-20/90 = -22.22\%$. Signed elasticity is -1.222 approximately; its magnitude exceeds one, so this interval is elastic.

Walkthrough

Step one calculates the changes: price increases by two pounds and quantity decreases by twenty units. Step two calculates the average of each pair of endpoints: eleven pounds and ninety units. Step three divides each change by its own average, giving approximately 0.1818 and negative 0.2222. Step four divides the quantity proportion by the price proportion. The result is negative eleven-ninths, or approximately negative 1.222. Using midpoint bases gives the same magnitude when the direction of comparison is reversed. Do not divide a percentage quantity change by a raw pound change; that would not be an elasticity.

Slide 4: Revenue connects price and quantity

Total revenue equals price multiplied by quantity sold. In the preceding example, revenue changes from £1,000 to £960. The price rise is outweighed by the quantity response. Along unchanged demand, an elastic interval makes price and revenue move in opposite directions; an inelastic interval makes them move in the same direction.

Walkthrough

Calculate revenue independently to check the interpretation: ten pounds times one hundred is one thousand pounds; twelve times eighty is nine hundred and sixty. Revenue falls by forty pounds. For positive endpoints, the midpoint elasticity classification is consistent with this finite-change revenue test. If its magnitude is exactly one, the two endpoint revenues are equal. The relationship assumes the observed quantity change is a movement along unchanged demand and that all units sell at the stated price. Revenue is not profit. Costs, returns, discounts and capacity constraints may change whether a pricing decision benefits the seller.

Slide 5: Worked example: an inelastic price cut

Price falls from £20 to £18 while quantity rises from 50 to 54. Average price is £19 and average quantity is 52. The midpoint elasticity is $(4/52)$ divided by $(-2/19)$, approximately -0.731 . Revenue falls from £1,000 to £972: the percentage quantity gain is too small to offset the price reduction.

Walkthrough

First take the signed changes: quantity rises by four and price falls by two. Divide by the midpoint bases to get about positive 7.692 per cent and negative 10.526 per cent. Dividing those proportions gives negative nineteen-twenty-sixths, approximately negative 0.731. Its magnitude is below one, so demand over this interval is inelastic. Next calculate revenues: 20 times 50 equals 1,000, while 18 times 54 equals 972. The price cut reduces revenue by £28. More customers or more units sold therefore does not necessarily imply more sales revenue, still less more profit.

Slide 6: A straight line can have changing elasticity

Slope compares changes in units; elasticity compares proportional changes. For $Q = 200 - 10P$, the slope is always minus ten units per pound. Point elasticity is $-10P/Q$. At $P = £5$ and $Q = 150$ its magnitude is one-third; at $P = £15$ and $Q = 50$ it is three.

Walkthrough

The same ten-unit response is a larger percentage of a small starting quantity than of a large one. Likewise, a one-pound change is a smaller percentage of a high price than of a low price. That is why a constant slope need not imply constant elasticity. The point formula multiplies the quantity response per pound by price divided by quantity. At five pounds, it gives negative ten times five divided by 150. At fifteen pounds, it gives negative ten times fifteen divided by fifty. Point elasticity describes a small local change; midpoint elasticity summarises an interval between two distinct observations.

Slide 7: Substitution and time influence responsiveness

Demand often responds more when close substitutes are available, when spending represents a larger budget share, or when buyers have time to adjust. These are tendencies to investigate, not universal numerical laws. A narrowly defined product can face more substitutes than a broad category. Essential uses do not make every quantity completely unresponsive.

Walkthrough

A particular café may have many nearby substitutes, while coffee as a broad category has fewer close alternatives for some customers. Buyers may respond to a lasting fuel-price change by changing routes, vehicles or housing, even if tomorrow's commute barely changes. The relevant time horizon therefore matters. Different groups can also have different responses. An average market elasticity can conceal people who have few short-run alternatives. To use an estimate responsibly, check its product definition, population, period, price range and research design. An elasticity borrowed from a different setting may be a poor guide.

Slide 8: Supply has an elasticity too

Price elasticity of supply compares the percentage change in quantity supplied with the percentage change in the price received by sellers, other supply conditions unchanged. It is often positive. Capacity, inventories, input availability and time to expand matter. Demand responsiveness and supply responsiveness jointly determine how a market adjusts to a disturbance.

Walkthrough

A bakery with unused ovens may increase output quickly, while a fully booked venue cannot create more seats for tonight merely because offers rise. Over a longer period, investment or entry can make supply more responsive. These examples concern conditional responses, not shifts caused by new technology or input-price changes. As with demand, distinguish a local elasticity from an estimate over a large interval. A vertical supply curve means fixed quantity in the model's relevant range, whereas a horizontal curve represents an extremely responsive supply at the given price. Neither shape is appropriate for every market.

Slide 9: Other elasticities answer other questions

Income elasticity relates quantity demanded to income. A positive value describes a normal good; a negative value describes an inferior good over the relevant range. Cross-price elasticity relates demand for one good to another good's price. Positive responses can indicate substitutes; negative responses can indicate complements, with relevant conditions held fixed.

Walkthrough

Do not interpret all elasticities as own-price demand elasticity. If bus fares rise and train demand increases, the cross-price response can indicate substitution between travel options. If printer prices rise and demand for ink falls, the goods may be complements. These labels depend on the context and level of aggregation; they are not timeless characteristics of every individual's preferences. Income elasticity also varies across groups and income ranges. The word inferior has no moral meaning here. As always, an observed correlation is not automatically a causal response when incomes, prices and other conditions move together.

Slide 10: A tax introduces a price wedge


For a per-unit tax, distinguish the price buyers pay, P_b , from the net price sellers receive, P_s . With a £6 tax, $P_b = P_s + 6$. Our illustrative schedules are $Q_d = 100 - 2P_b$ and $Q_s = 10 + P_s$. Without the tax, one common price is £30 and quantity is forty.

Walkthrough

Before introducing the tax, set buyer and seller prices equal to P . Equilibrium is $100 \text{ minus } 2P \text{ equals } 10 \text{ plus } P$, giving $90 \text{ equals } 3P$ and $P \text{ equals } 30$. Supply then gives 40 units. With the tax, using a single price in both original schedules would erase the distinction we need to analyse. Buyers respond to the amount they pay, while sellers respond to what they retain per unit after the tax. We assume a competitive market, compliance, unchanged underlying schedules and no separate administrative costs or other policy effects.

Slide 11: Worked example: calculate tax incidence

Substitute $P_b = P_s + 6$ into demand: $100 - 2(P_s + 6) = 10 + P_s$. This simplifies to $78 = 3P_s$, so $P_s = £26$ and $P_b = £32$. Quantity is thirty-six. Compared with the untaxed price, buyers pay £2 more and sellers receive £4 less per unit.



At 36 units buyers pay £32 and sellers receive £26; the vertical wedge is £6.

Walkthrough

Expand the left side carefully: $100 \text{ minus } 2P_s \text{ minus } 12 \text{ equals } 10 \text{ plus } P_s$. Subtract ten and collect price terms to obtain $78 \text{ equals } 3P_s$. Dividing by three gives 26, then adding the six-pound tax gives the buyer price of 32. Check quantities: demand is $100 \text{ minus } 64$, or 36; supply is $10 \text{ plus } 26$, also 36. Tax revenue is six times 36, or £216. The buyer and seller burdens per traded unit sum to the tax: two plus four equals six. This is a particular model calculation, not a general two-to-one incidence rule.

Slide 12: Legal payment and economic burden can differ

Who sends a tax payment to the authority need not determine who bears its economic cost. In the simple competitive model, the less responsive side generally bears more of a small tax's price burden. At the untaxed equilibrium here, supply elasticity is 0.75 and demand's magnitude is 1.5, consistent with sellers bearing more.

Walkthrough

Use the local elasticity formulas at the original price of 30 and quantity of 40. Supply's slope is one, so its elasticity is one times 30 divided by 40, or 0.75. Demand's slope is negative two, giving magnitude two times 30 divided by 40, or 1.5. Sellers are less responsive in this example. This intuition concerns relative ability to adjust quantities. Market power, evasion, price regulation, cross-border options and longer-run changes can alter the result. A positive incidence calculation also does not decide whether a tax is justified; that judgement needs stated goals, rights and other consequences.

Slide 13: A correlation is not a demand experiment

A hot weekend can increase both cold-drink prices and sales. That does not establish an upward-sloping demand curve; demand may have shifted. To estimate a causal price response, isolate a price change from other influences. Randomised prices or credible quasi-experimental variation can help, but each design needs assumptions, appropriate permissions and careful interpretation.

Walkthrough

Imagine comparing a quiet rainy Tuesday with a crowded sunny Saturday. Weather, visitors, opening hours and stock could differ alongside price. Dividing the observed percentage quantity change by the observed price change gives a descriptive ratio, but not necessarily demand elasticity. A controlled comparison might hold other conditions more nearly constant, while a research design may use a source of price variation unrelated to demand changes. Neither method is automatically flawless. Check selection, spillovers, measurement, compliance and the range of prices studied. Report uncertainty rather than presenting a single estimate as a universal behavioural constant.

Slide 14: Calculate, interpret, then qualify

Name the response being measured. Use consistent percentage bases and report whether elasticity is signed or absolute. Check revenue directly when relevant, and separate revenue from profit. For incidence, distinguish buyer and seller prices. Finally, ask whether other conditions were genuinely controlled and whether the estimate applies to the decision at hand.

Walkthrough

A reliable answer has three layers. First comes arithmetic: changes, averages, ratios and units. Second comes interpretation: what the magnitude means for the specified demand or supply relationship. Third comes identification and scope: why the data support a causal statement and where the estimate might fail to transfer. The exercises deliberately include a case where arithmetic is possible but a causal demand interpretation is not established. Written reflections remain ungraded. Use them to explain assumptions and evidence, because checking a number cannot substitute for understanding the model or the way the observations were generated.

Questions students ask

Should I report a negative elasticity or a positive one?

State the convention. Signed own-price demand elasticity is normally negative for downward-sloping demand. Elastic, inelastic and unit-elastic classifications commonly use its absolute magnitude.

Can a price cut increase sales and reduce revenue?

Yes. With inelastic demand over the interval, the proportional quantity increase is too small to offset the proportional price reduction. Calculate price times quantity at both endpoints to check.

Can the party legally paying a tax pass all of it on?

That is not automatic. In the simple competitive model, the division depends on demand and supply responsiveness. Other market institutions and adjustment mechanisms can also matter.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Calculate an interval response

Illustrative unchanged demand curve: price rises from £8 to £10 and quantity falls from 150 to 120 units. Calculate the absolute midpoint elasticity and total revenue at each endpoint. Round elasticity to two decimal places if needed.

Absolute midpoint elasticity

Revenue before (£)

Revenue after (£)

2. Evaluate another price cut

Illustrative unchanged demand curve: price falls from £30 to £27 and quantity rises from 200 to 210. Calculate absolute midpoint elasticity to two decimal places, the signed change in revenue (after minus before), and the elasticity classification.

Absolute midpoint elasticity

Revenue after minus before (£)

Classification

1. Elastic
2. Inelastic
3. Unit elastic

3. Solve a new tax wedge

Illustrative competitive market: $Q_d = 120 - 3P_b$ and $Q_s = 12 + P_s$. A per-unit tax of £8 creates $P_b = P_s + 8$. Calculate the buyer price, seller net price and quantity after the tax. All schedules remain unchanged apart from this wedge.

Buyer price after tax (£ per unit)

Seller net price after tax (£ per unit)

Quantity after tax (units)

Ungraded reflection: what would you compare to calculate each side's burden per traded unit?

4. Separate a ratio from a causal estimate

Illustrative café observations: on a quiet weekday, price is £4 and sales are 80 drinks; during a busy festival, price is £5 and sales are 100. Calculate the signed midpoint percentage-change ratio (quantity change divided by price change). Then decide what the observations alone establish about causal own-price demand elasticity.

Signed midpoint ratio

Supported conclusion

1. Causal demand elasticity must be positive
2. Causal demand elasticity must be negative one
3. Causal demand elasticity is not identified by these observations alone

Ungraded reflection: name one plausible confounding change.

Production, costs and feasible profit

Connect production increments to costs, contribution, operating choices and capacity without confusing cash receipts with economic gains. All scenarios are illustrative.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Calculate total, marginal and average product.
2. Distinguish fixed, variable, avoidable and sunk costs for a stated decision.
3. Reconcile cash surplus, economic contribution and economic surplus.
4. Calculate break-even and check workforce or sales choices against capacity.

Slide 1: Production transforms scarce resources

Production combines labour, equipment, materials, knowledge and time into goods or services. Output is a physical or service quantity; revenue is money received for selling it. More output is not automatically more profit or greater social value. Every business, price and production relationship in this lesson is an illustrative teaching assumption.

Walkthrough

A workshop can produce more chairs yet earn less if additional chairs sell cheaply or require costly overtime. It can also earn money while imposing a cost on people outside the transaction. Keep three questions separate: what can be produced, what the producer gains, and what happens to everyone affected. We begin with the technical relationship between inputs and output, then add prices and opportunity costs. This sequence makes it easier to see which conclusions arise from technology and which depend on markets, alternatives and values.

Slide 2: A production function holds some things fixed


A production function describes attainable output from specified inputs under a given technology. In a short-run example, equipment is fixed while labour can vary. “Short run” means at least one relevant input cannot adjust; it is not a universal number of months. Long-run comparisons allow all relevant inputs to change.

Walkthrough

Imagine a workshop with one cutting machine and different numbers of workers. Adding a worker can change how effectively the machine is used, but adding a second machine is outside this particular short-run choice. A different business may adjust equipment much faster or slower. The distinction is about available adjustment, not a fixed calendar rule. The production relationship also assumes comparable worker hours, skill and material quality. If those change across observations, a higher output cannot automatically be attributed to worker numbers alone.

Slide 3: Worked example: marginal product

With zero through five workers, illustrative daily output is 0, 12, 28, 39, 46 and 50 units. Each additional worker therefore adds 12, 16, 11, 7 and 4 units. These increments are marginal products. Marginal product starts falling after the second worker, although total output continues rising throughout the table.



Total output rises even after marginal product starts falling after worker two. Worker counts are discrete observations; connecting segments are visual guides.

Walkthrough

Subtract each output from the next one. The first worker adds 12 minus zero. The second adds 28 minus 12, or 16. The third adds 39 minus 28, or 11. The fourth adds seven and the fifth adds four. Early specialisation can explain an initial rise in marginal product; crowding around fixed equipment can explain a later decline. These are possible mechanisms, not facts established merely by the table. Diminishing marginal product means smaller additions to output, not necessarily falling total output or a judgement about any worker's individual worth.

Slide 4: Average and marginal product differ

Average product is total output divided by the number of workers. In the same example, three workers produce thirty-nine units, so average product is thirteen units per worker. The third worker's marginal product is eleven. Adding an increment below the previous average pulls the average down; an increment above it raises the average.

Walkthrough

At two workers, output of 28 gives an average of fourteen. The third worker adds eleven units, below fourteen, so the new average falls to thirteen. This is a general property of averages. A new observation below an existing average lowers that average when weights are comparable. Average product and marginal product answer different questions: output per worker across the whole team versus output added by a change in team size. Neither alone determines the profit-maximising workforce, because we also need output prices, labour costs and any other costs caused by expansion.

Slide 5: Classify costs along separate dimensions

A fixed cost does not vary with output over the relevant range; a variable cost does. An avoidable cost changes if the activity is stopped; a sunk cost cannot now be recovered. These distinctions are independent. A future monthly subscription can be fixed with respect to sales yet avoidable if the business does not operate.

Walkthrough

Do not use fixed and sunk as synonyms. Suppose a workshop can cancel next month's room hire at no charge. The hire is fixed over the month's feasible production range, but avoiding it is relevant to whether the workshop opens. A non-refundable design payment made last year is sunk for today's decision. A machine already owned may still have a resale or rental alternative, creating an opportunity cost even if no new cash payment is needed. Classify costs for the specific decision and period rather than attaching permanent labels without context.

Slide 6: Cash surplus is not economic surplus

Cash surplus subtracts the stated cash outlays from receipts. Economic surplus also deducts the value of resources used instead of their best feasible alternatives. An owner's unpaid time can therefore matter. Define the measure and comparison before using "profit". Do not subtract an opportunity cost twice when comparing complete alternative plans.

Walkthrough

If an owner could earn money elsewhere during the same hours, business receipts alone do not show whether operating improves their position. One method subtracts forgone alternative earnings from the business cash surplus. Another adds remaining alternative earnings to business receipts and compares that full allocation with working elsewhere throughout. Both should agree when they include the same costs and time. An accounting allocation may or may not reflect a relevant sacrifice. Ask whether the decision changes cash spending, resource use or the value of an available alternative.

Slide 7: Worked example: contribution per unit

An illustrative product sells for £18. Each unit needs £7 of materials and delivery plus twelve minutes of the owner's time. Alternative work pays £15 per hour. Twelve minutes is 0.2 hours, worth £3. Economic contribution per unit is therefore $£18 - £7 - £3 = £8$ before monthly fixed costs.

Walkthrough

Begin with consistent time units. Divide twelve minutes by sixty to obtain 0.2 hours. Multiply by the available alternative hourly return of fifteen pounds, giving a three-pound opportunity cost per unit. Subtract both the seven-pound cash cost and this time cost from the selling price. Eight pounds remains to cover the costs that do not vary with unit volume and then provide an economic surplus. This contribution assumes each extra sale leaves other sales unchanged, uses the stated resources and fits within capacity. It is not yet the profit on an entire month.

Slide 8: Worked example: build the monthly result

The same business has £90 of avoidable monthly cash costs and four maintenance hours. At £15 per hour, fixed economic costs total £150. Thirty sales give thirty times £8 minus £150, or £90 economic surplus. Cash surplus is £240, but ten total working hours have an opportunity cost of £150.

Walkthrough

Check the result by two routes. The contribution route gives 30 times eight, or £240, less fixed economic costs of £90 plus four times £15, yielding £90. The cash route begins with 30 times £18 revenue, or £540. Subtract £210 variable cash costs and £90 monthly cash costs to get £240. Production takes thirty times 0.2 hours, or six hours, and maintenance takes four more. Ten hours valued at £15 cost £150 in forgone earnings. Subtracting that from cash surplus gives the same £90 economic result.

Slide 9: Break-even depends on the decision

With £8 contribution per unit and £150 avoidable fixed economic costs, monthly economic break-even is $150/8 = 18.75$ units. Nineteen whole sales exceed that threshold under the assumptions. The minimum price for one additional sale after fixed commitments are unavoidable is a different question from the volume needed to justify operating.

Walkthrough

Write economic surplus as eight times quantity minus 150. Set it equal to zero, add 150 to both sides and divide by eight. The result is 18.75, which is a planning boundary rather than a possible whole-unit sale count. At eighteen units, contribution is £144 and the business is six pounds short. At nineteen, contribution is £152 and the surplus is two pounds. This threshold assumes unchanged prices, costs and product mix. If contribution were zero or negative, increasing volume alone would not cover positive fixed costs within this linear model.

Slide 10: Past expenditure does not force future production

Suppose the owner previously paid £300 for a design that cannot be refunded or sold. That payment is sunk for next month's operating decision. It may matter when evaluating the original investment, but it does not change next month's alternatives. Future avoidable costs and current resale opportunities still belong in the comparison.

Walkthrough

A project can have been disappointing overall and still be worth continuing from today, or have done well historically and now be worth stopping. The relevant question is what each available decision changes. Sunk-cost reasoning does not require forgetting the past: past results can inform forecasts and reveal mistakes. It requires separating information from an unrecoverable payment. Also check whether something described as sunk really is unrecoverable. If the design could be licensed to someone else only when unused by this business, that alternative would create a current opportunity cost.

Slide 11: Capacity makes algebra conditional

The owner has twenty hours available. Four maintenance hours leave sixteen production hours. At 0.2 hours per unit, capacity is eighty units. A formula predicting profitable sales beyond eighty is outside the stated feasible set. Expansion would need another assumption about hours, equipment, productivity or outsourcing, with its associated costs.

Walkthrough

Subtract fixed time before dividing by time per unit: twenty minus four is sixteen, and sixteen divided by 0.2 is eighty. At full capacity, an extra order cannot simply consume another otherwise unused twelve minutes. It may displace an existing order, require overtime or require someone else's help. Its opportunity cost must reflect that actual alternative. A constant-contribution formula is useful only while its resource and pricing assumptions hold. Always check a proposed break-even or switching threshold against capacity before presenting it as a feasible business plan.

Slide 12: Worked example: choose a workforce

Return to outputs 0, 12, 28, 39, 46 and 50. Output sells for £12 per unit; each worker costs £100 per day. An already committed £60 daily cost is unavoidable now. Net results for zero through five workers are –£60, –£16, £76, £108, £92 and £40. Three workers is best.

Walkthrough

Multiply each output by twelve, then subtract one hundred pounds per worker and the committed sixty pounds. For three workers, revenue is 39 times twelve, or £468; costs are £300 plus £60, leaving £108. The fourth adds seven units worth £84 but costs £100, reducing the result by £16. The fifth also reduces it. The first three additions increase the result, even though the one-worker total remains negative after the committed cost. This example shows why marginal contributions and complete-plan comparisons work together. The conclusion assumes all listed staffing options are feasible.

Slide 13: Profit guides choices under specific conditions

Economic profit can signal that resources earn more here than in their stated alternatives. It can encourage entry and imitation where feasible. Losses can encourage adjustment or exit. These mechanisms depend on information, competition and institutions; profit is not a complete measure of social contribution or proof that every cost has been borne privately.

Walkthrough

A business may discover a useful way to serve customers, or it may benefit from restricted entry, misleading information or costs imposed on others. The same accounting label cannot distinguish those cases by itself. A liberty-oriented approach values experimentation and voluntary cooperation, while a careful analysis asks how rights, liability and competition shape the available opportunities. Long-run entry can reduce some excess returns, but barriers, innovation and uncertainty can sustain differences. Avoid turning a useful tendency into a promise that all markets instantly reach identical returns or perfectly reward social value.

Slide 14: Audit the production decision

State output units, technology, relevant period and feasible inputs. Calculate marginal and average product separately. Classify costs by what the decision changes, including opportunity costs. Reconcile contribution and complete-plan calculations. Test break-even against capacity, and distinguish a marginal order from opening, continuing or expanding an entire operation.

Walkthrough

A good business calculation should be reproducible by someone who knows only the assumptions and arithmetic. Label pounds per unit, hours per month and total pounds distinctly. Show intermediate values before rounding and explain whether an average represents expected product mix or actual sales. Then test the assumptions most likely to reverse the choice, such as demand, service time or the availability of alternative work. The numerical exercises check specified calculations. Their ungraded reflection fields ask you to identify what the arithmetic leaves unresolved rather than pretending that an automatic marker can fully assess business judgement.

Questions students ask

Does a fixed cost belong in a decision about one extra order?

Only if that order changes it. A cost fixed over the relevant range can still matter when deciding whether to operate for the whole period.

Is unpaid owner labour free?

It has no wage payment, but may sacrifice a valuable alternative. Identify the next-best feasible use rather than automatically valuing it at zero or an unavailable wage.

Does diminishing marginal product mean a worker produces nothing?

No. The additional output can remain positive while becoming smaller. The table describes the whole production arrangement with specified fixed inputs, not an intrinsic measure of a person's worth.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Read a new production table

Illustrative daily output at zero through four workers is 0, 15, 33, 45 and 52 units. Calculate the third worker's marginal product and average product with three workers.

Third worker's marginal product (units per additional worker)

Average product with three workers (units per worker)

2. Build an economic surplus

Illustrative product: price £24, variable cash cost £9, twelve minutes per unit, alternative earnings £20 per hour. Monthly fixed cash costs are £75 and maintenance takes three hours. At twenty sales, calculate contribution per unit and monthly economic surplus. Also calculate the smallest whole number of sales with a non-negative economic surplus, assuming capacity permits it.

Contribution per unit (£ per unit)

Economic surplus at twenty sales (£)

Smallest whole break-even-or-better volume (units)

3. Respect the time budget

Illustrative business: eighteen hours are available each month, routine maintenance uses three hours, and each order needs fifteen minutes. Calculate normal maximum order capacity. Then calculate switching-month capacity if an additional two hours are required for setup. No extra labour is available.

Normal monthly capacity (orders)

Switching-month capacity (orders)

Ungraded reflection: why might a profitable algebraic threshold still be unusable?

4. Compare feasible staffing plans

Illustrative output with zero through five workers is 0, 10, 24, 33, 39 and 42 units. Each unit sells for £14. A worker costs £80 for the day. An additional £50 is already committed and unavoidable for every staffing choice. Calculate the best workforce and its net result after all stated costs.

Best number of workers (workers)

Net result (£)

Ungraded reflection: distinguish the committed cost from the cost of the next worker.

Rights, contracts and institutional choices

Examine how rules support cooperation, why externalities and public goods create challenges, and how to compare realistic institutions while distinguishing efficiency from rights.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Explain property rights, contracts and transaction costs as institutional mechanisms.
2. Identify external effects and public-good properties precisely.
3. Calculate illustrative private and combined consequences without treating monetary sums as complete moral judgements.
4. Compare feasible institutional alternatives using explicit evidence, costs, consent and rights criteria.

Slide 1: Institutions shape the available choices

Institutions are formal and informal rules that structure interaction: ownership arrangements, contracts, customs, organisational rules and enforcement practices. Organisations are groups operating within or changing those rules. Institutions influence information, incentives and transaction costs. Every numerical situation here is illustrative, and the lesson explains economic concepts rather than rules of any particular legal jurisdiction.

Walkthrough

Compare two markets selling the same product. In one, sellers' identities and terms are clear and disputes can be resolved cheaply. In the other, neither party knows whether promises will be honoured. Even with identical technology, the feasible trades and their costs can differ. Institutions help explain that difference. Some rules are written and enforced through courts or organisations; others rely on repeated interaction and reputation. Do not assume that a formal rule is effective merely because it exists on paper. Actual enforcement, expectations and informal practices also shape behaviour.

Slide 2: Property rights specify permitted uses

Property rights describe claims about using, excluding others from, benefiting from and transferring resources. Different rights can be separated or conditional. Clear and secure rights can support planning and investment, but defining boundaries and addressing conflicts can be costly. Economic effects and the moral legitimacy of a particular allocation of rights are separate questions.

Walkthrough

Owning a resource is not always a single unlimited permission. A tenant may have a right to use a workshop without a right to sell the building. A neighbouring owner may have a competing claim about harmful noise. The exact legal details vary, so our concern is the general economic mechanism: predictable claims can reduce uncertainty and make longer-term arrangements possible. They can also determine bargaining positions and who must obtain consent. Calling a right clear does not establish that its origin was just, nor does showing a useful incentive settle every rights-based objection.

Slide 3: Contracts make plans conditional on agreement

A contract records agreed obligations, such as delivery, payment, quality and timing. Better specification can reduce misunderstanding, but writing and enforcing every possible contingency is costly. Incomplete contracts leave some future situations unspecified. Reputation, warranties, repeated dealings and dispute procedures can supplement formal terms, each with its own limitations.

Walkthrough

Suppose a buyer orders a custom desk. Stating the dimensions, material and delivery date narrows the space for disagreement, but unexpected damage or supply disruption can still occur. A detailed agreement costs time to prepare and may not anticipate every relevant event. The parties need a way to interpret terms and respond to surprises. These arrangements influence the attractiveness of exchange. This is not a claim that any informal promise is legally enforceable; enforceability depends on applicable rules and facts. The economic question is how the arrangement changes expected costs, trust and feasible cooperation.

Slide 4: Transaction costs change the gains from cooperation

Search, measurement, negotiation, monitoring and dispute resolution use resources. These transaction costs can prevent a trade whose gross benefits exceed production costs. Intermediaries, standards and organisations may reduce some costs while introducing others. Compare total relevant costs, including time and expected failures, rather than selecting the arrangement with the smallest visible fee.

Walkthrough

A platform charging five pounds may still be cheaper overall if it saves an hour of search worth more than five pounds. Conversely, a zero-fee arrangement may expose participants to expensive disputes or poor information. Some activities are organised inside firms partly because repeated market negotiation would be costly; internal management also creates coordination and incentive costs. Neither markets nor organisations are costless defaults. Identify the task, the realistic alternatives and the margins on which their costs differ. Include uncertainty about quality rather than treating a promised service as identical to a reliably delivered one.

Slide 5: Incentives depend on who receives consequences

A person who expects to keep benefits from improving a resource may have a stronger reason to invest in it. Someone who expects others to bear its costs may take too little care. These are conditional incentive mechanisms, not universal descriptions of human motivation. Norms, preferences, monitoring and the security of claims also matter.

Walkthrough

Consider a shared garden. Participants may invest because they enjoy gardening, care about neighbours or expect future produce. A rule assigning plots might improve incentives in one setting and undermine useful cooperation in another. The relevant comparison depends on which problems are actually present and what arrangements are feasible. A liberty-oriented perspective values personal agency and consent; empirical analysis asks how particular rules affect stewardship and conflict. Avoid assuming that people respond only to money or that every collective arrangement fails. Equally, good intentions alone do not guarantee that a system handles incentives well.

Slide 6: External effects cross the transaction boundary

An externality is an effect on people outside the decision that is not fully reflected in the decision-maker's relevant costs or benefits. Pollution and uncompensated spillover learning are common examples. A price change that affects another seller is not automatically the same kind of externality. Specify the mechanism, affected people and missing incentive.

Walkthrough

A workshop's fumes may harm nearby residents even if buyers and sellers are satisfied with the furniture transaction. The harm falls outside the private comparison unless some institution brings it into the decision. By contrast, a new competitor lowering an incumbent's profits through a lower market price is not, by that fact alone, a technological externality. Economic analysis distinguishes these mechanisms because their implications differ. Estimating an external effect requires evidence about causation and magnitude. Calling something an externality does not by itself identify the best response or resolve whose rights should constrain the activity.

Slide 7: Worked example: a private output decision

An illustrative workshop can make up to five batches. Additional benefits after all private resource costs are £16, £12, £8, £4 and £1. With no other private consequence, all five add positive value to the owner. Total private net benefit is £41. Now suppose each batch also causes £6 of uncompensated harm to neighbours.

Walkthrough

Start with the private decision before adding the external effect. Add the five net increments: sixteen plus twelve plus eight plus four plus one equals forty-one pounds. Because every increment is positive, the owner prefers five batches to fewer under the stated assumptions. The neighbours' harm is not part of that private calculation. We are not assuming the harm is acceptable or that the workshop has a legitimate right to impose it. We are making the incentive gap explicit so that a later institutional comparison can address both consequences and rights.

Slide 8: Worked example: include the external cost

Subtract £6 of harm from each batch's private increment. Combined net increments become £10, £6, £2, minus £2 and minus £5. Three batches maximise their summed net value at £18. Five batches give £41 of private benefit minus £30 of harm, or £11. The fourth and fifth reduce the combined total.

Walkthrough

Compute each affected margin rather than merely comparing average harm with average benefit. The first batch adds sixteen minus six, or ten. The next adds six, the third adds two, the fourth loses two and the fifth loses five. Cumulative totals are ten, sixteen, eighteen, sixteen and eleven. Three batches therefore maximise this particular monetary sum. That conclusion assumes the valuations are meaningful, all relevant effects are included and the options are feasible. It does not prove that monetary gains can ethically cancel someone else's rights violation or that the best enforcement arrangement is known.

Slide 9: Bargaining may help, but conditions matter

If affected parties can identify one another, specify enforceable rights and negotiate cheaply, they may agree on changes that create mutual gains. Bargaining outcomes also depend on initial rights and available alternatives. Many parties, uncertain damage, holdouts, free riding or enforcement problems can block agreement. Potential compensation is different from consent and actual compensation.

Walkthrough

Suppose a noise-reduction measure costs less than the total harm it would prevent. There may be room for an agreement that funds the measure and leaves participants better off. But someone must identify the harm, agree on obligations and verify the outcome. Initial rights influence who can demand payment and who must obtain permission. A stylised result about costless bargaining does not establish that bargaining is costless in practice or that the initial allocation of rights is irrelevant to distribution. Assess the real obstacles before recommending either private negotiation or another institutional response.

Slide 10: Public goods have two technical properties

A public good is non-rival over the relevant range and difficult to exclude non-payers from using. One person's benefit does not substantially reduce another's, and selective access is costly. These properties create a possible free-rider problem. "Public good" does not simply mean beneficial, publicly owned or supplied by government.

Walkthrough

An uncongested local warning signal can benefit several people at once, and excluding a nearby listener may be difficult. By contrast, a publicly operated swimming pool can become crowded and can charge admission. It is not a pure public good merely because a public body runs it. Some goods are partly rival or partly excludable, and technology can change these properties. Free riding arises when someone can benefit without contributing, potentially weakening incentives to fund provision. That is a mechanism to analyse, not proof that voluntary provision never occurs or that compulsory provision always succeeds.

Slide 11: Worked example: funding a shared light

Three neighbours value an illustrative shared light at £18, £15 and £11. Provision costs £36. Total stated benefit is £44, leaving an £8 net gain. Equal £12 shares would exceed the third neighbour's benefit. Contributions of £14, £12 and £10 cover the cost and leave individual gains of £4, £3 and £1.

Walkthrough

First add the benefits: eighteen plus fifteen plus eleven equals forty-four. Subtract the thirty-six-pound cost to obtain eight pounds of possible combined gain. Then test a proposed funding rule. Equal shares cost each person twelve pounds, which the third values more than their eleven-pound benefit, so unanimity is not established. The alternative contributions sum to thirty-six and are each below the respective benefit. This demonstrates a feasible individually beneficial allocation under the supplied valuations. It does not prove people will reveal valuations truthfully, contribute as promised or organise the arrangement without further costs.

Slide 12: Compare realistic institutional alternatives

Possible responses include clearer rights, contracts, voluntary associations, reputation systems, pricing mechanisms, liability, rules and collective provision. Each requires information, administration and enforcement. Compare how alternatives perform on the actual problem, including errors and side effects. An imperfect existing arrangement should not be judged against an imagined alternative that has no costs or mistakes.

Walkthrough

For a pollution problem, a numerical charge may create incentives but require information about damage and reliable enforcement. A performance rule may be easier to verify in some settings but less flexible in others. Bargaining may work well with a few identifiable parties and poorly with dispersed populations. Voluntary associations may use local knowledge but face participation problems. These possibilities are analytical categories rather than legal recommendations. A careful comparison asks who knows what, who decides, who pays, how disputes are resolved and what happens when the institution itself performs badly.

Slide 13: Knowledge and power need scrutiny everywhere

Markets, firms, associations and governments all operate with limited information and imperfect incentives. Competition, transparency, accountability and exit can improve some arrangements, but their availability varies. Private market power and public decision-making power both deserve examination. Institutional analysis should explain mechanisms and evidence rather than assume benevolence or competence from an organisation's label.

Walkthrough

A regulator may lack local cost information; a firm may conceal product quality; an association may be dominated by a few members. These are different problems requiring different safeguards. The relevant question is which feasible arrangement handles the particular information and incentive problems better, under the rights and objectives being considered. A liberty-oriented framework gives strong weight to consent and limits on coercion. It can still investigate actual outcomes rather than treating voluntary language as sufficient evidence that meaningful choice exists. Specify both the evaluative principles and the empirical claims instead of allowing one to stand in for the other.

Slide 14: Keep efficiency and legitimacy distinct

An efficiency calculation asks what consequences follow under stated assumptions and an evaluative measure. A rights-based judgement asks which actions and claims are permissible. Neither is identical to fairness, consent or total money income. A defensible recommendation identifies affected people, relevant rights, feasible alternatives, enforcement costs, uncertainty and the values guiding the decision.

Walkthrough

Return to the workshop example. The monetary calculation identified an output level with a larger combined measured surplus. It did not establish who should hold which rights, whether anyone may be compelled to accept harm, or which institution can implement a better outcome at acceptable cost. A recommendation should say what additional evidence and moral premises it uses. In the exercises, numerical responses check the supplied benefit and cost arithmetic. The ungraded written explanations are where you can discuss missing information and competing principles. They are deliberately not marked by a simplistic keyword rule.

Questions students ask

Does an externality automatically prove that a particular intervention is best?

No. It identifies a potentially missing incentive. A response must still be compared with realistic alternatives, including its information, enforcement, distributional and rights implications.

Can a privately organised group provide a public good?

Yes. Public-good properties describe rivalry and exclusion, not the supplier's ownership. Voluntary provision may work under some conditions, although funding and participation problems can remain.

Does a larger combined monetary gain override someone's rights?

That does not follow from the arithmetic. Whether and how rights constrain permissible actions is a normative question that must be addressed explicitly.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Calculate the gains after coordination costs

Illustrative one-off trade: the buyer values the item at £85 and the seller's opportunity cost is £55. Search, verification and enforcement use resources worth £6, £4 and £5 respectively. Calculate the gross potential gain and the gain after these transaction costs. Do not assume a particular division between the parties.

Gross potential gain (£)

Net potential gain (£)

2. Compare private and combined increments

Illustrative activity: four possible units have private net benefits of £20, £14, £9 and £3. Each unit causes an additional £7 of uncompensated external harm. Units must be taken in order and there are no other effects. Calculate the privately preferred quantity, the quantity maximising summed net value including harm, and that maximum summed value.

Privately preferred quantity (units)

Quantity maximising the stated combined measure (units)

Maximum combined net value (£)

Ungraded reflection: what moral or institutional question does this calculation leave open?

3. Test a shared funding proposal

Illustrative shared service: three people value provision at £21, £17 and £14. Cost is £45. Calculate total stated benefit and combined net benefit. Then ask whether an equal three-way cost split gives every person at least as much benefit as payment. Assume these valuations are accurate and no other effects exist.

Total stated benefit (£)

Combined net benefit (£)

Does every person weakly benefit from the equal split?

1. Yes, all payments are no greater than benefits
2. No, at least one payment exceeds its benefit
3. Cannot compare the supplied values

4. Compare like with like

Illustrative institutional comparison: arrangement A has administration costs £12, expected error costs £25 and delay costs £8. Arrangement B has corresponding costs £20, £10 and £6. Categories do not overlap. Calculate each total and B's cost saving relative to A. These figures alone do not measure rights or consent.

A's total cost (£)

B's total cost (£)

A's cost minus B's cost (£)

Ungraded reflection: identify one additional criterion needed for a defensible recommendation.

Money, time and uncertain outcomes

Distinguish nominal and real values, compound and discount dated amounts, and understand basic risk and bank balance sheets without treating forecasts as guarantees.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Distinguish money, income and wealth, and nominal from real changes.
2. Calculate compound growth, present values and a simple net present value.
3. Interpret expected values while recognising risk attitudes and uncertain probabilities.
4. Describe simplified bank lending and distinguish liquidity from solvency.

Slide 1: Money connects transactions across time

Money commonly serves as a medium of exchange, a unit of account and a store of purchasing power. These functions can be imperfect: purchasing power may change, and acceptance varies. Money is not identical to wealth or income. Every amount and rate in this lesson is illustrative, not a current quotation or personal financial recommendation.

Walkthrough

A person may own a valuable house but have little immediately spendable money. Income is a flow over a period, while wealth is a stock of assets minus liabilities at a date. Money is one part of that balance sheet and helps settle transactions without arranging a direct barter match. Its usefulness depends on acceptance and institutions, not simply physical form. We will focus on interpreting money amounts through time, then introduce credit and bank balance sheets. The arithmetic is general, but actual financial products can have fees, taxes and risks omitted from our teaching examples.

Slide 2: Worked example: nominal versus real purchasing power

An illustrative income rises from £500 to £530 while a relevant price index rises from 100 to 108. Nominal income increases by 6%. Expressed in initial-period purchasing power, the new income is $£530/1.08$, approximately £490.74. Real purchasing power falls by about 1.85%, despite the higher number of pounds.

Walkthrough

Step one calculates the nominal growth factor: 530 divided by 500 equals 1.06. Step two calculates the price-level factor: 108 divided by 100 equals 1.08. Step three divides the income factor by the price factor, giving approximately 0.98148. Subtract one to obtain a real change of negative 1.85185 per cent. Equivalently, deflate the new income by dividing £530 by 1.08, then compare the resulting £490.74 with £500. The result depends on whether the price index represents the purchases relevant to the person or question.

Slide 3: Inflation is a rate of price-level change

Inflation is an increase in a specified general price index over a period. Lower positive inflation means prices rise more slowly, not that the price level falls. Deflation is a decline in that index. Individual relative prices can move differently from the overall index, reflecting changes in demand, supply or quality.

Walkthrough

If an index rises from 100 to 110 and then to 115.5, inflation falls from ten per cent to five per cent while the price level keeps rising. The second increase is calculated from 110, not from the original 100. A single expensive product does not by itself establish general inflation. Index construction also requires choices about baskets, weights and quality adjustment. Different households buy different things, so their experiences can differ from a published average. Be clear about the index, period and comparison before claiming that an income or price has risen in real terms.

Slide 4: Worked example: compound growth

A hypothetical £1,000 balance grows at 5% per year, with interest retained and no fees or taxes. After one year it is £1,050. After two, it is $£1,050 \times 1.05 = £1,102.50$. The general formula is $\text{future value} = \text{present value} \times (1 + \text{rate})^{\text{number of periods}}$.

Walkthrough

First convert five per cent to 0.05. One year's growth multiplies the starting balance by 1.05. In the second year, the same rate applies to the larger balance, including the first year's interest. That produces £52.50 in second-year interest rather than £50. Total interest over the two years is £102.50. Simple interest on the original principal alone would produce £100. Compounding is an arithmetic rule under stated terms, not a promise that any real investment earns a constant return. Negative returns, charges or withdrawals would require a different sequence of balance changes.

Slide 5: Match the rate to the period

A one-per-cent monthly growth rate compounds by multiplying by 1.01 each month. Over twelve months, the factor is 1.01 raised to twelve, approximately 1.126825: effective annual growth of 12.6825%. Multiplying the monthly rate by twelve gives 12%, which omits growth on earlier months' additions. Always check the compounding convention.

Walkthrough

A rate without a period is incomplete information. Five per cent per month and five per cent per year describe very different growth paths. In this example, apply twelve monthly factors rather than one factor of 1.12. Actual quoted financial rates may use different conventions and include or exclude charges, so read their definitions rather than assuming labels are interchangeable. For a sequence of varying period rates, multiply the corresponding growth factors. An arithmetic average of returns generally does not reproduce the final compounded balance. Consistency of periods is essential before comparing alternatives.

Slide 6: Discounting translates future amounts into present equivalents

At a stated discount rate r , a payment received t periods later has present value equal to future payment divided by $(1 + r)$ raised to t . Discounting reverses compounding. At 5% annually, £1,102.50 received in two years has a present value of £1,000 under the same illustrative assumptions.

Walkthrough

Ask what amount today would grow into the future payment at the specified rate. Since £1,000 grows to £1,102.50 over two years at five per cent, dividing the future amount by 1.05 squared returns £1,000. A discount rate represents a comparison across time and may incorporate opportunity cost and other considerations. It must be justified for the problem rather than chosen to force a preferred result. Keep nominal cash flows with a compatible nominal rate and real cash flows with a compatible real rate. Risk and uncertainty require additional care beyond this mechanical conversion.

Slide 7: Worked example: compare a dated project

An illustrative project costs £400 today and returns £220 after one year and £230 after two. At a 10% annual discount rate, present values are £200 and approximately £190.08. Net present value is $£200 + £190.08 - £400 = \text{minus } £9.92$. Under these assumptions, it falls short of the stated alternative return.

Walkthrough

Step one places every payment on a timeline. The £400 outlay occurs now, so it needs no discounting. Step two divides £220 by 1.10, giving £200. Step three divides £230 by 1.10 squared, or 1.21, giving approximately £190.0826. Step four adds the present receipts and subtracts the initial cost. The negative result is approximately £9.9174. Although undiscounted receipts total £450, that total ignores timing. The conclusion assumes the stated receipts occur, the discount rate is an appropriate alternative comparison, and no omitted costs or benefits change the result.

Slide 8: Nominal and real returns must be distinguished

If a hypothetical nominal balance grows by 8% while the relevant price level rises by 5%, the exact real growth is $1.08/1.05 - 1$, approximately 2.86%. Subtracting inflation from the nominal rate gives an approximation of 3%. The exact formula compares growth factors, and actual inflation may differ from what was expected.

Walkthrough

Suppose one hundred pounds becomes one hundred and eight. If a basket that cost one hundred now costs one hundred and five, the final balance buys 108 divided by 105 baskets relative to the starting one. That ratio is approximately 1.02857. Subtracting one gives the real proportional gain. The simple subtraction approximation is often close for modest rates, but it can be inaccurate for large changes. Expected real returns also depend on expected inflation, which need not equal realised inflation. Do not present a nominal guarantee as a guarantee of purchasing power.

Slide 9: Expected value is not a certain outcome

With stated probabilities, expected value is the sum of each outcome multiplied by its probability. An illustrative choice yielding £80 with probability 0.6 and minus £40 with probability 0.4 has expected value £32. Neither possible outcome is £32. Risk attitudes, affordability of losses and confidence in the probabilities still matter.

Walkthrough

Multiply eighty by 0.6 to get forty-eight, and negative forty by 0.4 to get negative sixteen. Adding them gives thirty-two pounds. Repeated comparable choices can make an average meaningful, but a single participant still faces the actual possible outcomes. Risk often refers to settings where a probability model is available; uncertainty can include situations where probabilities themselves are poorly known. The terminology varies, so define it. A precise calculation based on unjustified probabilities can create false confidence. Identify what evidence supports the outcomes and probabilities before treating expected value as a reliable guide.

Slide 10: Several exposures can share the same risk

Spreading resources across different activities can reduce exposure to some independent or imperfectly related risks. It does not remove all risk. If outcomes respond to the same underlying shock, apparent variety may offer little protection. Liquidity needs and the consequences of a bad outcome can matter as much as the average expected return.

Walkthrough

Imagine two shops selling different goods but relying on the same festival visitors. Their sales risks may be closely linked despite the different products. A cancelled festival could hurt both together. To assess diversification, ask about the drivers of outcomes rather than merely count separate names. Some risks can be pooled; others affect many people simultaneously. These are general concepts, not a recommendation to buy particular assets. A decision should also consider when funds are needed, whether commitments can be reversed and whether a plausible loss would prevent meeting essential obligations.

Slide 11: Bank lending changes balance sheets

In a simplified bank loan, lending £1,000 creates a £1,000 loan asset for the bank and a £1,000 deposit liability when the borrower's account is credited. The borrower receives a deposit asset and a matching debt. This does not create £1,000 of net wealth for the borrower or give banks unlimited lending capacity.

Walkthrough

Write both sides of each balance sheet. The bank holds a claim on the borrower and owes the deposit balance to its account holder. The borrower can spend the deposit but must repay the loan under its terms. Payments to another bank can require settlement funds, so liquidity and funding matter. Capital, credit risk, profitability, regulation and demand for loans also constrain lending. Repaying principal reduces the relevant loan and deposit balances in the simplified system; interest payments are a separate flow. The process is more precise than saying banks merely pass a fixed pile of existing deposits from one person to another.

Slide 12: Liquidity differs from solvency

Liquidity concerns meeting payments when due. Solvency concerns whether assets cover liabilities under the relevant valuation. An illustrative bank has assets of £100 and liabilities of £95, leaving equity of £5. An £8 asset loss reduces assets to £92 and equity to minus £3. A solvent institution can nevertheless face a short-term liquidity problem.

Walkthrough

Subtract liabilities from assets to calculate equity: one hundred minus ninety-five equals five. After the asset loss, ninety-two minus ninety-five equals negative three. That is a solvency problem in this simplified balance sheet. A liquidity problem can arise without such a loss if valuable assets cannot be converted into settlement funds in time to meet withdrawals. Selling assets quickly at a discount can then turn liquidity pressure into losses. This distinction explains why payment timing and asset quality both matter. Actual assessment requires detailed valuation and institutional information beyond these teaching numbers.

Slide 13: Credit reallocates purchasing power and obligations

Credit allows spending before the borrower has earned or saved the full amount, while creating future repayment obligations. It can fund productive activity or unsustainable commitments. Effects on output and prices depend on demand, capacity, expectations, institutions and how credit is used. A balance-sheet mechanism alone does not determine a complete macroeconomic outcome.

Walkthrough

A loan financing useful equipment can expand future productive capacity; another loan might finance an asset purchase without directly adding new production. Both create obligations, and neither description guarantees successful repayment. At the economy-wide level, changes in spending interact with available resources, imported goods, pricing behaviour and expectations. Avoid claiming that any one monetary quantity mechanically produces a fixed immediate inflation result under all conditions. A liberty-oriented analysis also asks about contract terms, consent and who bears losses. Those questions complement rather than replace investigation of balance sheets and causal mechanisms.

Slide 14: Put time, purchasing power and uncertainty together

Label each amount as a stock or flow, nominal or real, and present or future. Match rates to periods and compounding conventions. Discount dated amounts consistently, identify the alternative being used, and test uncertain assumptions. For credit, record assets and liabilities separately. An attractive receipt today can carry a substantial obligation tomorrow.

Walkthrough

Use a timeline and a small balance sheet before making a complicated judgement. They expose missing payments, inconsistent periods and confusion between money received and wealth gained. Then ask how the conclusion changes if inflation, timing or outcomes differ. The exercises use hypothetical amounts so that you can practise the method without relying on current market products. Numerical fields check arithmetic under stated assumptions. Ungraded explanations are for assessing the limits of those assumptions, including why a positive expected value or a positive nominal return is not a guarantee of a desirable outcome.

Questions students ask

If inflation falls, are prices back where they started?

Not necessarily. Lower positive inflation means a slower increase in the index. Prices fall only if the relevant index decreases over the stated period.

Why discount a future payment that is larger than today's cost?

Timing matters because resources committed today have alternatives. Discounting expresses dated amounts in comparable present terms using a justified rate and compatible assumptions.

Does receiving a bank loan make the borrower richer?

The initial transaction gives the borrower a spendable asset and a matching liability. What happens to net wealth subsequently depends on how the funds are used, costs and repayment obligations.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Deflate a nominal increase

Illustrative income rises from £800 to £872 while a relevant price index rises from 125 to 135. Calculate the new income in initial-period purchasing power and its percentage change from the initial income. Give the signed percentage change to two decimal places.

New income in initial-period pounds (£)

Real percentage change (%)

2. Match growth and discounting

Illustrative balance: £1,500 grows at 4% annually for three years with all interest retained and no fees or taxes. Calculate the final balance. Separately, calculate the present value of £2,000 received in two years at a 6% annual discount rate. Round money to two decimal places only at the end.

Three-year future balance (£)

Present value of the separate future payment (£)

3. Evaluate a dated hypothetical project

Illustrative project: pay £600 now, receive £350 after one year and £330 after two. Use a 10% annual discount rate. Calculate net present value to two decimal places. Separately, a choice has net outcomes £100 with probability 0.7 and minus £50 with probability 0.3; calculate its expected value.

Project net present value (£)

Separate choice's expected value (£)

Ungraded reflection: why does expected value not establish what a single attempt will deliver?

4. Separate a loan from net wealth

Illustrative bank: assets are £250 and liabilities £232 before a £23 asset loss. Liabilities are unchanged. Calculate equity after the loss. Separately, a borrower receives a £700 deposit and an equal £700 loan obligation with no fees at origination; calculate the immediate change in net wealth from those two entries alone.

Bank equity after the loss (£)

Borrower's immediate net-wealth change (£)

Ungraded reflection: distinguish an inability to make a payment today from liabilities exceeding asset value.

Capstone: Leila's map-shop decisions

Advise the fictional map maker using weighted contributions, opportunity costs, price-response boundaries, switching horizons and capacity. Worked teaching examples use different businesses and numbers.

Written walkthrough. Work through each explanation alongside the PDF slides. The lesson page also includes a captioned video with a generic synthetic narrator. It is not a recording of Craig Wright.

Fictional examples teach reasoning rather than describe a real business. No external academic credit is awarded.

Learning objectives

1. Construct weighted per-order revenue, fees and economic contribution.
2. Reconcile monthly cash and economic surplus with complete time-allocation comparisons.
3. Derive a price-change demand boundary without confusing it with a forecast.
4. Evaluate switching costs over stated horizons and reject infeasible plans.
5. Write an auditable, conditional recommendation and identify evidence that could change it.

Slide 1: Leila's map shop: identify the decisions

Leila makes downloadable walking maps. A single map costs £2 and a three-map bundle £5. Each purchase is one order and one payment. Her forecast is 120 orders: 75% singles and 25% bundles. You will evaluate operating, an additional order, a price change and a service switch. All people, services and figures are fictional.

Walkthrough

This capstone combines the earlier tools rather than introducing a slogan about business success. A positive cash receipt may use time that has a better alternative. A lower transaction charge may accompany a larger monthly bill. A price increase may improve contribution per order while reducing total orders. Treat these as different decisions with different relevant costs. The teaching walkthroughs use other fictional businesses and different numbers, so you can learn the method before independently calculating Leila's case. Her exercise responses are submitted for checking; the teaching slides do not disclose their solutions.

Slide 2: State Leila's time alternative and past commitments

Leila has twenty hours for maps and editing together next month. Editing is available for any of those hours at £18 per hour after its incremental expenses. She values the activities equally apart from financial return. Last month's £240 illustrations and six design hours are unrecoverable. A £12 laptop allocation changes neither spending nor alternative use.

Walkthrough

The editing opportunity supplies a specific, feasible comparison for time spent on maps. Do not replace it with an imagined wage from work she cannot obtain. Equally, do not assume time is free because Leila does not pay herself a wage. The earlier design work can be used with either service and cannot now be recovered. The laptop is retained for editing whichever option she chooses. These facts explain why some accounting entries will not change this decision. A different resale opportunity, personal preference or future benefit would require revisiting the assumptions rather than silently inserting it into the calculation.

Slide 3: Service A: record the complete operating terms

Service A is already configured. It charges £0.20 plus 3% of each order's price and an avoidable £40 monthly operating charge. Delivery and storage cost £0.10 per order. Each order needs three minutes of service time, and monthly maintenance takes four hours. Leila can stop next month without a cancellation charge.

Walkthrough

Write the per-order fee as a fixed amount plus a percentage of that order's price. The word fixed here describes the flat part of a transaction charge; it still varies with the number of orders in a month. Keep it distinct from the monthly charge, which does not vary with orders within the stated range. A bundle is one payment even though it contains three files, so its fixed transaction fee, delivery expense and service time occur once. The monthly charge and maintenance time are avoidable if Leila does not operate next month.

Slide 4: Service B: separate recurring and one-off differences

Service B charges £0.05 plus 5% per order and £55 per operating month. Delivery expense, three-minute service time and four-hour maintenance are unchanged. Switching from A requires £36 cash and two additional hours once. Both services have identical reliability, customer experience and demand in the model. These are assumptions, not claims about actual providers.

Walkthrough

Build two columns before comparing services: recurring monthly consequences and one-off switching consequences. A lower flat transaction charge is only one term. The percentage charge applies to the whole order price, so its importance depends on product mix and prices. The larger monthly fee must also be covered. Switching time displaces available editing just as operating time does, and it uses some of the first month's capacity. Because reliability and demand are stipulated to be identical, do not invent a customer benefit for one provider. Later, question that assumption explicitly as part of sensitivity analysis.

Slide 5: Define measures and rounding before calculating

Calculate fees exactly and aggregate before rounding. There are no taxes, refunds, advertising costs, failed payments or other expenses in this model. Cash surplus subtracts stated cash outlays from revenue. Economic surplus relative to editing also subtracts the value of displaced editing time. Expected fractional orders describe averages, while realised orders are whole counts.

Walkthrough

Keep intermediate calculations to at least four decimal places, and follow each exercise's requested final precision. A weighted average is useful when the product mix is a planning assumption. If actual counts are known, calculate from those counts instead. An expected 75:25 mix need not be realised exactly in a particular month. Do not call the defined cash surplus a complete accounting profit statement, because the model intentionally excludes many real-world items. Our purpose is to compare stated alternatives consistently, not to imply that running an actual business requires no tax, legal or operational investigation.

Slide 6: Worked example: another seller's weighted order

Noor sells illustrative templates for £4 or £9, with a 60:40 order mix. Weighted price is $0.6 \times £4 + 0.4 \times £9 = £6$. A fee of £0.10 plus 2% averages £0.22. Delivery costs £0.18; two minutes valued at £24 per hour cost £0.80. Contribution is therefore £4.80 per order.

Walkthrough

First apply the order weights to prices: £2.40 plus £3.60 equals £6. Because the percentage fee is linear and each order has one flat charge, its weighted average is ten pence plus two per cent of six pounds, or twenty-two pence. Convert two minutes to hours by dividing by sixty, then multiply by twenty-four pounds to obtain eighty pence. Finally subtract the fee, delivery and time cost: six minus 0.22 minus 0.18 minus 0.80 equals 4.80. Monthly costs remain separate. An unweighted price average would incorrectly assume equal numbers of the two order types.

Slide 7: Worked example: reconcile two monthly calculations

Noor expects fifty orders, pays £65 monthly and spends 1.5 maintenance hours. Fixed economic costs are $£65 + £36 = £101$. Economic surplus is $50 \times £4.80 - £101 = £139$. Alternatively, £300 revenue less £11 fees, £9 delivery and £65 monthly charges gives £215 cash surplus; total time costs £76.

Walkthrough

Check the time calculation carefully. Fifty orders at two minutes each require one hundred minutes, or five-thirds of an hour. Add one and a half maintenance hours to obtain nineteen-sixths of an hour. At twenty-four pounds per hour, that is seventy-six pounds. Subtracting £76 from the £215 cash surplus gives £139, matching the contribution method. This reconciliation is a powerful error check: the two routes organise the same consequences differently. If they disagree, look for omitted maintenance, inconsistent time units, an incorrect mix or an opportunity cost counted twice.

Slide 8: Worked example: compare complete time allocations

Suppose Noor has ten hours available for templates and alternative work paying £24 per hour. Alternative work alone gives £240. Operating uses nineteen-sixths of an hour, leaving forty-one-sixths for other work, worth £164. Adding £215 business cash surplus gives £379. The advantage over £240 is £139, matching the incremental method.

Walkthrough

The complete-allocation method includes the earnings actually received under each plan. Noor's remaining work hours earn one hundred and sixty-four pounds, while the template business produces two hundred and fifteen pounds after its stated cash outlays. Together they give three hundred and seventy-nine pounds. Compare that with two hundred and forty from using all ten hours for the alternative. The difference is one hundred and thirty-nine. Do not subtract the seventy-six-pound time opportunity cost again from this difference: the forgone earnings have already been reflected in the reduced hours available for the alternative.

Slide 9: Worked example: an additional order has a different threshold

For Noor, an extra order priced at p contributes $0.98p - £1.08$ before any new fixed costs. Setting this equal to zero gives $p = 1.08/0.98$, approximately £1.10204. With penny prices, £1.11 is the lowest positive-contribution price. This assumes spare capacity, unchanged other sales and no additional setup requirement.

Walkthrough

Derive the expression instead of memorising a threshold. Revenue is p . Subtract the two-per-cent fee to retain $0.98p$. The remaining per-order costs are ten pence of flat fee, eighteen pence of delivery and eighty pence of time, totalling £1.08. At £1.10, the expression is slightly negative; £1.11 makes it positive. This marginal price does not cover the monthly £101 fixed economic costs if charged across the business without enough contribution. Accepting one additional order after commitments are unavoidable and deciding whether to operate for a month are different comparisons.

Slide 10: Worked example: derive a demand boundary

If Noor raises both prices by 20% with mix unchanged, weighted price becomes £7.20 and the fee £0.244. Contribution becomes £5.976. To match the previous £139 surplus, solve $N \times £5.976 - £101 = £139$. The boundary is approximately 40.1606 expected orders, a 19.6787% decline from fifty. This is a condition, not a demand forecast.

Walkthrough

Calculate the new contribution by subtracting £0.244 fees, £0.18 delivery and £0.80 time from £7.20. To preserve the old result, new total contribution must still cover £101 fixed economic costs and leave £139, so it must equal £240. Divide 240 by 5.976 to find the required expected volume. Then subtract that volume from fifty and divide by fifty to measure the allowed proportional decline. The result says how much loss of demand the model can tolerate. It does not estimate how buyers actually respond to the higher prices or identify a profit-maximising price.

Slide 11: Apply the method to Leila's proposed prices

Leila considers a 10% increase to £2.20 and £5.50 on Service A. Keep the same expected mix and other costs. Compare scenarios of 102 and 96 expected orders, and derive the volume that matches her original economic surplus. No demand curve is supplied. Changes in mix, support time or customer exposure could change the comparison.

Walkthrough

Work from weighted contribution rather than assuming that a ten-per-cent price rise creates a ten-per-cent improvement in surplus. Percentage fees change with prices, while other per-order costs remain fixed under the assumptions. The fixed monthly economic costs remain the same. Compute each scenario separately and compare it with the starting case. Then solve for the exact expected order count needed to match that starting result. Fractional expected counts are allowed for this boundary. A recommendation must distinguish this conditional arithmetic from evidence about which demand scenario is likely.

Slide 12: Worked example: compare recurring service costs

A separate seller, Rowan, has £10 orders. Provider X charges £0.30 plus 2%; provider Y charges £0.10 plus 3%. Their fees are £0.50 and £0.40 per order. Y saves £0.10 per order but costs £8 more monthly. Its recurring advantage is therefore $£0.10N - £8$, positive above eighty orders.

Walkthrough

First calculate each provider's entire transaction charge at the same order price. For X, thirty pence plus twenty pence gives fifty pence. For Y, ten pence plus thirty pence gives forty pence. Subtract to obtain ten pence saved per order. Multiplying by monthly volume gives total variable savings, from which the eight-pound extra monthly cost must be deducted. At eighty orders, those effects exactly offset. This recurring comparison deliberately excludes a one-off switch, which belongs in a separate calculation. It also assumes no difference in demand, quality or service time.

Slide 13: Worked example: add switching costs and time

Rowan's switch costs £20 cash and 1.5 hours worth £20 per hour, totalling £50. At 180 monthly orders, recurring advantage is £10. Cumulative advantage after m months is $£10m - £50$: minus £20 after three months, zero after five and plus £10 after six. This undiscounted calculation assumes stable monthly conditions.

Walkthrough

Value switching time using Rowan's stated alternative: one and a half hours times twenty pounds is thirty pounds. Add the twenty-pound payment to get fifty pounds. At 180 orders, ten pence per order saves eighteen pounds before the extra eight-pound monthly charge, leaving ten pounds per month. Multiply that recurring advantage by the horizon and subtract the switch once. Do not subtract it every month or omit it entirely. A longer horizon can make a switch attractive, but only if the recurring assumptions hold and the transition fits available capacity. Discounting would be needed if the chosen comparison required it.

Slide 14: Evaluate Leila's services over stated horizons

Return Leila to the original £2 and £5 prices. Compare B with A at 120 orders, first on recurring costs and then including switching. Derive recurring and first-month volume thresholds. Separately examine 240 orders per month for ten months, with no discounting, and compare the cumulative results after three and ten months.

Walkthrough

Do not mix the higher-price scenario with the service-switching scenario unless explicitly asked to build a new combined case. For this comparison, both services use the original prices, original mix and identical demand. The first-month threshold must include the cash and time cost of switching, while later months should not repeat that one-off cost. The separate 240-order scenario is a planning assumption, not evidence that demand has increased. Your recommendation should state what must remain true over the horizon and whether a small conditional advantage is robust to plausible changes in workload or sales.

Slide 15: A threshold must fit the resource budget

Check Leila's twenty-hour limit before recommending any order threshold. Maintenance uses time even at low volume, and switching consumes additional first-month time. For comparison, Noor's ten-hour budget, 1.5 maintenance hours and two minutes per order permit 255 orders. An extra one-hour setup would reduce that illustrative first-month capacity to 225.

Walkthrough

The capacity method starts by subtracting all relevant fixed time from available time. For Noor, ten minus one and a half leaves eight and a half hours. Multiplying by sixty gives 510 minutes; dividing by two gives 255 orders. With the additional setup hour, remaining minutes fall to 450, giving 225 orders. Apply the same method independently to Leila's own figures. A profitable-looking algebraic threshold outside capacity is not an available choice. Meeting it would require a new assumption about labour, service time or available hours, and that change may alter the cost comparison.

Slide 16: Write a conditional, auditable recommendation

Your brief should cover operating, pricing and switching under their separate scenarios. Support it with checked calculations, units and stated assumptions. Identify two material uncertainties and practical ways to investigate them. Distinguish Leila's private economic return from effects on walkers or neighbourhoods. The written brief is ungraded reflection; numerical checks do not certify the quality of the whole recommendation.

Walkthrough

A strong brief says what should be done under which conditions and what evidence would change that view. Useful tests might measure service minutes by product, track the realised bundle share or compare customer exposure across periods. Avoid a vague instruction to collect more data without specifying the decision it would inform. If the editing work disappears, identify the new feasible alternative rather than mechanically retaining the old wage or declaring time worthless. Your conclusion should show how assumptions connect to numbers and numbers connect to a decision, while acknowledging questions of rights and wider consequences that the private-return model does not settle.

Questions students ask

Should I count the files in a bundle as separate orders?

No. In this case a bundle is one order and one payment. Apply the stated per-order fee, delivery cost and service time once to a bundle.

Why can the best marginal decision differ from the best monthly decision?

An additional order may contribute towards commitments that cannot now be avoided. Opening next month can involve avoidable monthly cash and time costs that the extra-order calculation does not include.

Can I claim the price increase is optimal if one scenario improves the result?

No. A scenario shows a conditional outcome, and a boundary identifies a required response. Neither supplies a full demand curve or proves which price maximises the relevant objective.

Your exercises

Submit your answers on the lesson page at craigwright.net/learning. Numerical and choice checks receive server feedback. Written responses are saved ungraded. No assessment answer key appears in this guide.

1. Leila: build the per-order economics

Use Leila's original prices, 75:25 order mix and Service A terms from the lesson. Calculate the single-order fee, bundle-order fee, weighted order price and weighted economic contribution after fees, delivery and service-time opportunity cost. Keep monthly costs separate. Give fees and contribution to four decimal places and weighted price to two.

Fee on one single-map order (£ per order)

Fee on one bundle order (£ per order)

Weighted order price (£ per order)

Weighted economic contribution (£ per order)

2. Leila: evaluate next month

At 120 orders under the original Service A assumptions, calculate monthly cash surplus, total map-business hours and economic surplus relative to editing. Then calculate the continuous expected-order boundary for zero monthly economic surplus, maintaining the expected mix. Round money to two decimals and the order boundary to four; do not require the expected mix to be exactly realised at every count.

Cash surplus (£)

Total map-business hours (hours)

Economic surplus relative to editing (£)

Zero-economic-surplus expected volume (expected orders)

3. Leila: test the 10% price increase

Keep Service A and the original expected mix. With prices £2.20 and £5.50, calculate weighted economic contribution to five decimal places; economic surplus at 102 and 96 orders to two decimals; the expected order count matching the original surplus to four decimals; and the corresponding maximum percentage decline from 120 to two decimals. All other assumptions stay fixed.

New weighted economic contribution (£ per order)

Economic surplus at 102 orders (£)

Economic surplus at 96 orders (£)

Expected volume matching the original surplus (expected orders)

Maximum decline from the original expected volume (%)

4. Leila: compare services and capacity

Return to the original prices and mix. Calculate B's recurring economic surplus at 120 orders, the expected-volume threshold where its recurring result equals A, and its first-month equality threshold including switching. Give thresholds to four decimals. Calculate A's normal capacity and B's switching-month capacity. Separately, at 240 orders every month with no discounting, calculate B minus A cumulative advantage after three and ten months, including the switch once. Give money to two decimals.

B's recurring surplus at 120 orders (£)

Recurring equality threshold (expected orders)

First-month equality threshold (expected orders)

A's normal monthly capacity (orders)

B's switching-month capacity (orders)

B minus A after three months at 240 orders/month (£)

B minus A after ten months at 240 orders/month (£)

5. Leila: check the framing and write the decision brief

First answer the objective framing check. Its mark applies only to that choice, never to the written brief. Write 400–600 words advising Leila about operating next month, the proposed price increase, and switching services. Keep the 120-order starting case separate from the ten-month 240-order scenario. Support your conclusions with your checked calculations, explain the editing alternative without double counting, address capacity and identify two material assumptions with practical tests. This written brief is explicitly ungraded: review it for economic framing, numerical support, uncertainty, conditional decision quality and clear communication.

Which description correctly frames the economic comparison with editing?

1. Subtract only the cash payments made by the map business.
2. Subtract relevant cash costs and the value of the editing time given up.
3. Subtract cash costs, then add the editing income given up as a benefit.
4. Subtract every past payment, even when it cannot change between the available choices.

Ungraded decision brief