

ADDALY · THE AI CLASSROOM

Making It Say Something

Clean work that communicates nothing is the commonest self-taught failure, and it is fixable.

8 modules · 64 lessons · 26 figures · 8 case studies · 52,031 words · about 9 hours

Dr Mustafa Mun
Author and editor

ABOUT THIS BOOK

Making It Say Something, published by Addaly, 2026. Author and editor: **Dr Mustafa Mun**.

The manuscript was drafted with the assistance of a large language model and was edited and published under his name. That is stated plainly here rather than left to be discovered, because a reader is entitled to know how a book they are trusting was made.

This book is the printed form of the course of the same name at addaly.com/learn/visual-storytelling. The website is the living version and is corrected first; where the two differ, the website is right.

Free to read, free to download, free to print, and free to teach from. There is no paid edition and there never will be. If you paid for this, somebody has sold you something that was given away.

Every diagram in it is drawn from data by code, not generated as a picture, so the numbers on the page are the numbers in the argument. Answers to every exercise, test and examination question are at the back, with a note on why each wrong option attracts people — those notes are the teaching, not the marking.

Contents

1. Before you open the software

Turn a vague request into a written brief naming one specific viewer, their belief before and after, the seven constraints the work must survive, and an agreed measure of success — and defend a design decision by pointing at that document rather than at taste

1. Nice work. What is it for?
2. Nobody is going to hand you a real brief
3. "Young urban professionals" cannot be persuaded
4. You cannot move someone without knowing where they are standing
5. The constraints were true all along; you just found out late
6. Six questions, and what to do when they cannot answer
7. Reference, and the line you actually cross
8. Agree the measure before you make the thing
9. Case study: The donor poster that answered the wrong objection
10. Module test

2. Making one message survive

Rank a design into three levels using at most two emphasis levers per element, measure its contrast against a stated ratio, prove the hierarchy with a blur, greyscale, flip and distance test, and delete the element you like most when it competes with the message

1. A design that says four things says none
2. There are only five ways to make something look important
3. The eye reports differences, never values
4. You see far less at once than you think
5. Every element you add taxes the ones already there
6. The empty part is doing the most work
7. Four cheap tests that fail a design before an audience does
8. The bit you are proudest of is usually the bit in the way
9. Case study: Six departments, one A3, and the surgeon's quotation
10. Module test

3. The shape of a story

Build a short piece as a chain of state changes joined by but and therefore, allocate its seconds to a stated structure, name the protagonist and the cost they carry, replace generic claims with checkable detail, and land one ask the piece has earned

1. "And then" is the weakest word you own
2. The first second is a promise, and you have to pay it
3. If nothing changes, it is not a story
4. Where the ninety seconds actually go
5. One checkable detail outperforms ten adjectives
6. Nothing is at stake, so nobody is watching
7. The person who changes is the one it is about
8. Endings, and the four asks that cancel each other out
9. Case study: The founder's forty years, and the customer's Thursday
10. Module test

4. Getting the material

Gather the material a planned piece needs — run an interview that yields self-contained answers, distinguish images that carry meaning from images that decorate, obtain consent whose scope matches the intended use, protect a subject against jigsaw identification, clear found material against its actual licence, present a number with its denominator, and audit every beat against the asset that will carry it before building begins

1. Draw ugly boxes now or pay for it later
2. Most interview footage is unusable, and the question is why
3. Half your footage is wallpaper
4. Consent is a scope, not a signature
5. Anonymity fails by combination, not by name
6. Free to download is not free to use
7. A number is a claim with the argument removed
8. Find the missing shot now, not in the edit
9. Case study: The detail that made the story mean something, and identified her
10. Module test

5. Words that carry the picture

Write the words around a visual piece so they do the work the pictures cannot — a headline that earns the first line, an opening that survives truncation, narration that can be followed without rewinding, subtitles inside real reading-speed limits, a caption that adds rather than restates, and a draft cut by a third without losing an idea

1. A headline is not a summary
2. Everything you write has a fold in it
3. A listener cannot go back
4. Most of your audience has the sound off
5. The caption decides what the photograph means
6. Your first draft contains the route, not the destination
7. You cannot un-know the thing you are explaining
8. Most people reading your English did not grow up in it
9. Case study: The subject line that doubled opens and halved the next email
10. Module test

6. Type, colour and image as meaning

Choose type, colour and image for what they say rather than for taste — justify a typeface by its associations and its script coverage, set a measure and line-height that can actually be read, keep meaning out of hue alone, meet a stated contrast ratio, and say what a photograph, an illustration and a diagram each assert that the others cannot

1. Your brand is not the logo, it is what you repeat
2. A typeface means where you have seen it before
3. Point size is the wrong number to argue about
4. Red is not danger everywhere
5. One in twelve men cannot see the difference you just used
6. Every photograph is mostly the part you cut off
7. The photograph that fits anywhere belongs nowhere
8. Three pictures, three different promises
9. Case study: The red and green flags on four thousand lab reports
10. Module test

7. The surface and the system carrying it

Deliver a piece into the conditions it will actually meet — survive the crop chain and interface overlays, set a cover frame and link preview deliberately, carry meaning with the sound off and on a phone speaker, mock a design up inside a real feed, distinguish the platform's objective from your own, and read a retention curve to name the specific failure it records

1. The same design fails in five different ways
2. The thumbnail and the title are one thing
3. Something is always sitting on top of your picture
4. The still is what recruits the play
5. It has to work silent and reward being unmuted
6. Nothing is ever seen on a white wall
7. The system is optimising something, and it is not your work
8. Which number answers which question
9. Case study: The carpentry reel with a button on the joint
10. Module test

8. Judgement, and the work that accumulates

Take a piece from private judgement into public responsibility — run a five-second and an explain-back test on strangers, separate a client's observation from their prescription, defend a decision from the agreed brief and concede what is not load-bearing, identify manipulation by whether it survives being explained, represent people you are not without using them as decoration, disclose AI use, and construct the feedback that deliberate practice needs

1. Three tests, no budget, twenty minutes
2. "Make the logo bigger" is a symptom, not an instruction
3. Argue from the document, not from taste
4. If explaining it would stop it working, it is manipulation
5. Somebody is in this piece, and it is not you
6. What the tool can be trusted with, and what it cannot
7. Take apart the thing that worked on you
8. Fifteen pieces with a point of view beat fifty with a filter
9. Case study: The countdown that was nearly real
10. Module test

Making It Say Something — final examination

The paper that opens the next course. Answers to everything follow it.

MODULE 1

Before you open the software

Most work that communicates nothing was doomed before the first file was created, because nobody wrote down who it was for or what it had to change. This block builds the document that every later decision is checked against: the one person who will see it, what they believe before they see it, what they should believe afterwards, the constraints that were true all along, and the measure everyone agreed to in advance.

BY THE END OF THIS MODULE YOU CAN

Turn a vague request into a written brief naming one specific viewer, their belief before and after, the seven constraints the work must survive, and an agreed measure of success — and defend a design decision by pointing at that document rather than at taste

1 · 8 min

Nice work. What is it for?

THE ONE THING TO KEEP

Before you open any software, write one sentence naming who sees this and what should be different in their head afterwards.

The most common failure is not ugliness

The piece that gets no response is usually not ugly. It is a well-spaced poster with a good photograph, a competent typeface and nothing to say. The colours

agree with each other. The alignment is honest. And a person walks past it without one thought changing in their head.

This happens to self-taught people more than to people who came through a studio, and there is a reason. Craft is visible in other people's work. You can see the kerning, the grid, the grade on the photograph, and you can copy it. Effect is invisible. You cannot look at a poster and see the eleven people who called the number. So you practise the part you can see, get good at it, and stay stuck.

The fix is not more craft. It is a habit of asking one question before you open any software.

The sentence that closes the gap

Write this before you design, in your own words, in your own language:

After seeing this, [a specific person] should [think, feel or do] [a specific thing], because [the reason they would believe it].

Most of the value is in refusing to be vague at each bracket.

Weak: *raise awareness of our café*. Awareness is not a thought a human has. "A specific person" is not "everyone".

Stronger: *Someone walking to the bus stop at 8am who currently buys instant coffee from the petrol station should think this place makes proper coffee and is open right now, because they can see the machine and the word OPEN from six metres away.*

That second sentence has already made design decisions for you. It tells you the piece is read at six metres, by someone moving, in the morning. It tells you the machine matters more than the logo. It tells you OPEN is the largest word. You have not chosen a font yet and the layout is half decided.

What each sentence has already decided for you

Raise awareness

- Names no particular viewer
- Names no belief to change
- Sets no size and no distance
- Decides no order of elements
- Cannot be tested on anybody

The bus-stop sentence

- One person, walking, at eight
- Buys instant from the petrol station
- Read at six metres, moving
- OPEN is the largest word
- Testable in three seconds

Both sentences take about the same time to write. The second one has made four design decisions before a typeface is chosen, which is why it is worth refusing to be vague at every bracket.

Three tests you can run this afternoon

The three-second test. Show the piece to a real person for three seconds, take it away, and ask what they remember. Do not ask if they like it — people are polite, and liking is not the variable. If the answer is "it looked nice", you have craft with no message. If the answer is "coffee, and something about being open", it works.

The removal test. Cover the most beautiful element with your hand. Does the message survive? If yes, that element is decoration. Decoration is allowed — but you should know which parts of your piece are load-bearing and which are furniture. Now cover the message instead. If the piece still feels finished, you never had one.

The distance test. Stand back three metres from the screen, or shrink the image to the size it will actually be seen. A flyer held in the hand, a post at 400 pixels wide in a feed, a thumbnail at 150 pixels. Almost every beginner design is composed at the size of the artboard and consumed at a fifth of it.

Where this advice is wrong

Not everything carries a message in words. A record sleeve, a textile pattern, a title sequence that is pure mood — these communicate a feeling and nothing else, and that is a complete job. The failure is never "no sentence". It is no intention: making a choice because it looked nice in someone else's work, with no answer to why it belongs in yours.

And the opposite failure is just as real. A sharp message rendered badly still fails, because people infer competence from surface. A flyer with a brilliant

line set in stretched Arial with a drop shadow tells the reader this business does not check its work, and they generalise from the flyer to the plumbing. Both halves are required. If your craft is the weak half, design foundations is the other side of this course and they are meant to be taken together.

There is a third failure worth naming because it looks like success: the piece that communicates perfectly, to nobody. Technically clean, message clear, aimed at an audience that does not exist or does not care. That is not a design problem. That is the next lesson.

Do this today

Open the last thing you finished. Do not change it. Write its sentence after the fact — who it was for, what should have changed in their head, and why they would have believed it. If you can write that sentence easily, you have been doing this instinctively and the rest of the course will sharpen it. If you cannot, you have found the reason the work did not land, and it was never the typeface.

BEFORE YOU MOVE ON

A designer shows a friend a café flyer for three seconds. The friend remembers a nice photograph of coffee but not the name of the place or that it opens at 7am. What does this most likely reveal?

- A. The photograph is too high in quality relative to the rest of the piece, so it should be replaced with something plainer.
- B. Three seconds is too short a window for anyone to take in a flyer, so the test result carries no information.
- C. Nothing in the piece was made more prominent than the photograph, so the only thing available to carry away was the photograph.
- D. The friend is not the target audience, so their reaction should be discounted until a real customer sees it.

2 · 9 min

Nobody is going to hand you a real brief

THE ONE THING TO KEEP

A request names an object; a brief names who sees it, what they believe now, and what should be different afterwards — write it down and get it confirmed.

What you will actually be given

You will be given a sentence. "We need a poster for the launch." "Something for Instagram, modern, make it pop." "A logo — you know, professional but friendly."

None of that is a brief. It is a request for an object. A brief tells you what the object has to do. The gap between the two is where most unpaid revision rounds live: you make the thing that was asked for, the client says it is not right, and neither of you can say why, because nobody ever wrote down what right meant.

Extracting the brief is a skill, it takes about fifteen minutes, and it is the single highest-paid quarter hour in the job.

The seven questions

Ask these out loud, in a call or in person if you can, and write the answers down where the client can see them.

1. **Who exactly sees this?** Push past "everyone aged 18 to 45". Push until you get a person: "shop owners on this road who already know us but think we are expensive."
2. **What do they currently believe?** This is the question people skip and it is the most useful one. You are not writing on a blank surface. You are editing something already in their head.

3. **What should they believe or do afterwards?** One thing. If you are given four, ask which one they would keep if the other three were deleted.
4. **Where will they see it?** On a phone at a bus stop. Printed A5 and handed over at a counter. Projected in a hall. This determines size, contrast, and how many words are possible, and it is often the answer that kills the client's original idea.
5. **How will you know it worked?** Phone calls, walk-ins, sign-ups, a full room. Even an unmeasurable answer is useful, because it exposes disagreement between two people on the client's side who thought they agreed.
6. **What must not change?** Legal text, a registered mark, a colour the founder will not give up, a photograph of a person who has given permission for that photograph and no other.
7. **Who signs it off, and when?** If the answer is "the team", find out whether one person can say yes. A piece that needs four yeses is a different job and should be priced as one.

Add money and rounds to the same document. "Two rounds of revision included, further rounds charged at X" prevents a specific argument, and it prevents it in writing.

Write it back

Send the answers back in your own words, in a shared document, and ask them to confirm. This is not bureaucracy. It converts a vague conversation into a thing you can point at three weeks later when someone says they wanted it blue.

A brief that fits on one screen:

FOR: Shop owners on Kandivali market road, 30–55,
 who know us but assume we are expensive.
 NOW: They think a service contract costs 8,000 a year.
 AFTER: They think it starts at 2,400, and they call.
 WHERE: A5 handed over the counter. Also one 4:5 post.
 SUCCESS: Twenty calls in three weeks.
 FIXED: Logo, the phone number, the GST line.
 SIGN-OFF: Priya only. Draft Thursday, final next Wednesday.
 ROUNDS: Two included.

When the client genuinely does not know

Sometimes there is no hidden brief to extract. The client has not thought about it, and interrogating them harder makes them defensive and makes you look like an obstacle. When you hit that, stop asking and start proposing.

Show two directions — not two layouts, two *premises*. Write one sentence above each: "Direction A says we are the cheapest. Direction B says we are the ones who turn up." Keep them visually rough so nobody argues about the shade of green. People are far better at reacting than at specifying, and a wrong direction rejected in five minutes has taught you more than an hour of questions.

The honest limitation: this only works if the two directions are genuinely different. Two versions of the same idea with different corner radii teach you nothing, and clients can feel when they are being offered a fake choice.

Briefing yourself

Personal work needs this more, not less, because there is nobody to argue with. The questions change slightly: who is this for, what do they already assume, what is the one thing I want to be different, and how will I know I got there. Written down, in a file, before you start.

You will fail this test constantly. Everyone does. But an unfinished personal project usually died because there was no way to tell whether it was finished, and a sentence written at the start is how you tell.

Do this today

Take a job you are working on now, even a small one, and answer the seven questions in a plain text file. Where you cannot answer, write "unknown" instead of guessing — then send those specific unknowns to the client as questions. Three unknowns sent on day one costs you one email. Three unknowns discovered on day nine costs you the weekend.

BEFORE YOU MOVE ON

A client asks for a poster and cannot answer any question about who it is for or what should change. Repeated questioning is making them irritated. What is the most productive next move?

- A. Show two rough directions built on visibly different premises, each with a one-line statement of what it claims, and let the client react.
- B. Proceed with your own best judgement and present a single finished piece, since the client clearly wants to be led.
- C. Send the seven questions in writing so the client has time to think properly before answering.
- D. Present two polished layouts of the same concept so the client can pick the execution they prefer.

3 · 9 min

"Young urban professionals" cannot be persuaded

THE ONE THING TO KEEP

Replace the demographic with one written portrait naming the device, the posture, the emotional state and the past disappointment, because those four facts decide real design questions and age never does.

A demographic answers no question you will actually ask

The brief says the audience is women aged 25 to 40 in tier-one cities with disposable income. Now you are three hours into the work and you have two options for the headline. One is dry and specific. One is warm and general. Ask the demographic which to use. It has nothing to say. It never did.

The decisions you make while designing are almost all of this shape: **would this person stop, and would they believe it?** A category cannot stop or believe. A person can. So the useful unit is not a segment, it is one imagined viewer detailed enough that you can predict their behaviour and be wrong in a way you notice.

Write the portrait, in about eighty words

Not a marketing persona with a stock photo and a fictional surname. A paragraph you could read to the client and have them say "yes, that is who walks in".

Sunita runs a two-chair salon in Kalyan. She looks at her phone between clients, sound off, standing up, one hand. She has been burned by a supplier who promised same-day delivery and took nine days. She does not believe advertised prices. She has about four seconds before the next client sits down.

Everything in that paragraph is load-bearing. Sound off kills your voiceover-only opening. One hand and standing up kills your six-line body copy. Burned by a supplier tells you the objection your piece has to answer before it can sell anything. Four seconds sets the length of your first beat.

Three facts that decide more than age does

- **The device and posture.** Phone in one hand on a moving train is a different medium from a laptop at a desk. It changes minimum type size, shot length, whether sound exists, and whether anything with fine detail can survive at all.
- **The state they are in when they meet the work.** Bored, hurried, anxious, sceptical, or actively shopping. Somebody comparing two quotes will read a wall of text. Somebody scrolling at midnight will not read anything.
- **What has already disappointed them.** Almost every audience has been let down by somebody in your category. The specific shape of that disappointment is the objection your work has to survive, and it is usually the most useful sentence in the whole brief.

Notice that none of these are age, income or city.

Where the detail comes from

You are not inventing it. You are collecting it, cheaply.

1. **Ask the client who called last week.** Not "who is your customer" — they will give you the demographic. Ask them to describe the last three people who actually bought, by name and by what they said on the phone. Clients can always do this, and it is gold.
2. **Read the reviews of your competitors,** especially the two-star and three-star ones. Five-star reviews say "great service". Three-star reviews say exactly what people expected and did not get.
3. **Watch one real person use the thing,** for ten minutes, without helping them. Ten minutes of watching beats an hour of speculation. You will

see them scroll past the thing everyone agreed was the most important element.

If you can do only one, do the first. It costs one question.

The honest limitation

A single portrait sharpens your decisions and narrows your sight at the same time, and this is a real cost, not a caveat. You will make work that fits Sunita and fails the customer who is nothing like her. That failure has a mechanism: every decision you make is now tuned to one set of constraints — sound off, four seconds, deep distrust of prices — and a viewer who has none of those constraints gets a piece shaped by problems they do not have.

Two defences, both cheap. Write **two** portraits when the audience genuinely splits — the buyer and the person who has to approve the purchase are often different people with opposite fears. And name, in one line, who this work is explicitly not for. "Not for people who already trust us" is a decision, and writing it down stops the client from later asking why the piece spends so long establishing credibility.

The test before you start

Read the portrait, then look at your first idea and ask one question out loud: *would she stop?* If the answer needs a paragraph of explanation, the answer is no.

Keep the portrait visible while you work. Print it, or paste it into the corner of the canvas on a layer you hide before export. In GIMP, Krita, Inkscape, Figma or Photoshop this costs nothing and it changes what you notice. The reason to keep it on screen rather than in your head is not sentiment — it is that by hour three you will have started designing for yourself, and only a written sentence argues back.

BEFORE YOU MOVE ON

A designer replaces "men 18–34" with a portrait: "Arjun watches on a phone at 1am in bed, sound off so he doesn't wake his brother, and has been scammed once by an online course." Why does this change the work, mechanically?

- A. It makes the work more emotionally resonant, because a specific named person is more relatable to a viewer than a demographic category is.
 - B. It gives each design choice a predicate you can check: sound off rules out a voiceover opening.
 - C. It narrows the audience, so the piece will reach fewer people but convert a higher proportion of them.
 - D. It replaces guesswork with data, since a portrait is drawn from research rather than assumption.
-

4 · 9 min

You cannot move someone without knowing where they are standing

THE ONE THING TO KEEP

Write the viewer's belief before and after in their own words, then check that every element shortens the distance between those two sentences — and diagnose a stalled piece by asking whether it failed at attention, comprehension or belief.

Every piece of work is a delta

A brief that says "we want people to know we are reliable" describes a destination and no starting point. Depending on where the viewer actually stands, the same destination needs four completely different pieces.

- If they have never heard of you, reliability is meaningless noise; you first have to exist.

- If they think you are expensive, a reliability claim reads as a justification for the price and hardens the objection.
- If they were let down by you once, a reliability claim reads as a lie, and the piece has to acknowledge the failure before it can claim anything.
- If they already trust you and are just slow to act, reliability is a wasted beat; what they need is a deadline.

So write two sentences, not one.

Before: *She thinks a salon-supply delivery will take a week and probably arrive wrong.*

After: *She thinks this particular supplier gets it right the same day, and that the claim is checkable.*

The work is the distance between those two sentences. Everything that does not shorten that distance is decoration.

The commonest failure: stating the after

Most weak work simply announces the after state. "Same-day delivery. Guaranteed." Printed large, set well, beautifully spaced. It fails because it treats a belief as something you can install by declaring it, and beliefs do not work that way. The viewer already holds a belief; yours arrives as a competing claim from an interested party, and the existing one has the advantage of experience behind it.

What actually shifts a belief is evidence the viewer can check without trusting you first. A photograph of the dispatch board at 4pm with the day's dockets on it. The line "we missed 11 of 2,400 deliveries last quarter — here is what we did about it". A named customer with a phone number. None of these are more persuasive because they are cleverer; they are more persuasive because they cost you something to publish and a false claim could be caught.

Belief has three components, and they fail differently

When a piece does not move anybody, the stall is usually in one of three places, and the fix for each is different:

1. **Attention.** They never took it in. Nothing about your argument is the problem. Fix the first beat, the size, the placement, the contrast.
2. **Comprehension.** They saw it and did not understand what was on offer. Fix the language, the hierarchy, the number of ideas competing.
3. **Belief.** They saw it, understood it, and did not believe it. No amount of better layout fixes this. You need evidence, a guarantee, a third party or an admission.

Teams spend most of their time on the second, because it is the most enjoyable to work on. In my experience the third is where the traffic is being lost, and it is the only one that requires going back to the client and asking for something they may not want to give.

How to find the before state cheaply

- **The objection log.** Ask whoever answers the phone what people say just before they decide not to buy. Sales staff can recite this list instantly; it is almost never in the brief.
- **Search suggestions.** Type the category into a search box and read the autocompletions: "is X safe", "X vs Y price", "X refund". Those are unedited before-states, at scale, free.
- **The review floor.** Two-star reviews of competitors, again. "Good product but the sizing runs small" is a before-state you can defuse in one line of copy.

Write the sentence the viewer would say

In their words, not yours. Not "perceives the category as unreliable" but "last time this took nine days and nobody answered the phone". The reason is practical: you can answer a sentence, and you cannot answer an abstraction.

Almost every good headline in this kind of work is a direct reply to a sentence somebody actually said.

What this cannot do

A before-and-after pair makes the job tractable; it does not make it achievable in one piece. If the gap is large — from "never heard of you" to "bought" — no single flyer or reel closes it, and pretending otherwise is how a piece ends up carrying five messages and landing none. When the distance is too big, split it: name the one intermediate belief you can actually install this time, and say plainly which piece of work does the next step.

It also cannot tell you whether the after state is true. If the client's deliveries really do take nine days, your job is not to build a more persuasive picture of same-day delivery. It is to tell them that the piece will work once and then cost them more than it earned, because the fastest way to destroy a claim is to make a lot of people check it.

BEFORE YOU MOVE ON

A landing page is seen by plenty of people, is perfectly clear about what is offered, and converts almost nobody. Which diagnosis does the before/after framing point to, and why?

- A. Comprehension — clarity of the offer and clarity of the value are different things, so the body copy needs rewriting until the benefit is unmistakable.
- B. Attention — high traffic does not mean people actually looked, so the hierarchy needs strengthening.
- C. Belief — it arrives and is understood, so what is missing is checkable evidence.
- D. The after state is wrong and should be made more ambitious so the page asks for a bigger commitment.

5 · 8 min

The constraints were true all along; you just found out late

THE ONE THING TO KEEP

Write the seven constraints — format, viewing situation, production route, both deadlines, split budget, existing assets and the approver — before designing, because each one is a filter that rejects whole classes of idea in seconds rather than days.

The rework you did not have to do

A designer spends two days on a poster built around a full-bleed photograph with white type reversed out of the dark half. It is good. Then the printer says the job is one-colour on kraft stock, because that is what the budget bought, and the whole thing dies. Nobody hid this. Nobody asked.

Constraints are not obstacles that appear later. They are facts about the job that existed before you started, and the only variable is whether you knew them on day one or day three. Getting them written down is fifteen minutes of work that routinely saves two days.

The seven to write down

1. **Final format and quantity.** A3 poster, 200 copies. Or 1080×1920 vertical video, under 60 seconds. "Some social stuff" is not a format.
2. **Where it will physically be.** Pasted at eye level in a corridor. Held in the hand. Seen from a car at 40km/h. This decides your smallest legible element more than any other fact.
3. **The production route and its limits.** Screen print in two colours, office laser printer, a phone screen, a large-format vinyl plotter. Each has a list of things it cannot do — the laser printer cannot do full bleed, the plotter cannot do fine gradients, the phone cannot do 6pt type.

4. **Deadline, and the real one.** Ask for both: the date they told you, and the date after which it is worthless. A Diwali flyer has a hard second date; an about-page redesign usually does not.
5. **Budget, split.** Your fee is not the budget. Print, paid media, a photographer, a licence for a typeface or a track — each of these either exists or does not, and a piece designed around a licensed typeface with no budget for one is a piece that will be rebuilt in a substitute font at midnight.
6. **What already exists and must be used.** Logo files, a colour, a tagline the founder will not part with, previous photography. Ask for the actual files early; "we have a logo" and "we have a vector logo" are different situations, and finding out on the day of print that the only copy is a 400px JPEG from a 2014 website is a specific and avoidable misery.
7. **Who approves it, and how many of them there are.** One person is a client. Five people is a committee, and a committee needs a different presentation strategy and a different time estimate. Get the name.

Constraints make decisions, and that is the point

There is a real pleasure in an unconstrained brief, and it is usually followed by paralysis. Two colours and kraft stock is not a lesser version of the full-colour job. It is a different problem with a smaller decision space, and smaller decision spaces are faster to search. Some of the most memorable design in the world was made inside brutal limits — one ink, one plate, one week.

The practical move is to write the constraint list at the top of the working file and treat each item as a filter. An idea that fails a constraint is not a good idea you cannot afford; it is an idea for a different job.

The seven constraints, and what each one rejects

Format and quantity	A3, 200 copies. Some social stuff is not a format.
Where it physically is	Sets the smallest legible element
Production route	One ink kills the full-bleed photograph
Both deadlines	The stated date, and the date it is worthless
Budget, split	Your fee is not the print or the licence
What already exists	A vector logo and a 400px JPEG differ
Who approves it	One person is a client. Five is a committee.

Each line is a filter. An idea that fails one is not a good idea you cannot afford, it is an idea for a different job, and it costs seconds to reject here and two days to reject later.

The free path

Every constraint above can be worked inside on free software, and the tool is not usually what is missing.

- Two-colour print separations, spot colours and bleed: **Scribus** does layout and PDF/X export, **Inkscape** does vector artwork and can output spot-colour-safe files. The paid counterparts are InDesign and Illustrator, which job ads will name.
- Photo work bound to a small palette: **GIMP** and **Krita** both do indexed colour, duotone by hand, and posterisation. **Photopea** runs in a browser on a low-spec laptop and opens PSD files.
- Video at a fixed length and aspect ratio: **DaVinci Resolve** (free tier), **Kdenlive** or **Shotcut**; **CapCut** on a phone.

What you cannot get free is a licensed commercial typeface or a stock photo licence. That is a budget line, and it belongs on the list above so somebody decides before you design around it.

The limitation nobody mentions

Writing the constraints down does not make them stable. Clients change quantity, dates move, and the founder discovers a new favourite colour. The list is not a contract; it is a record of what was true when you started, and its

value is in making a change visible as a change. When item 1 moves from 200 copies to 5,000, screen printing stops being the cheap option and the whole design premise is up for review — and you can say that in one sentence because you wrote the premise down.

Send the list back to the client in an email with the words "here is what I have understood". Half the time you will get one correction that saves the job.

BEFORE YOU MOVE ON

A client raises the print run from 200 to 5,000 copies. Why does a written constraint list turn this from a small update into a flagged design review?

- A. Because a larger run means a larger budget, so the design can now afford more colours.
- B. Because the route chosen for 200 copies is no longer cheapest at 5,000, and the design premise rested on that route.
- C. Because 5,000 copies will be seen by a much wider range of people, so the message needs to be broadened and made less specific to survive.
- D. Because more copies means a longer lead time, so the deadline must move.

6 · 9 min

Six questions, and what to do when they cannot answer

THE ONE THING TO KEEP

Ask six specific questions — including what has been tried before, what will never be approved, and who else must say yes — and when the client cannot answer, write a draft brief for them to correct rather than pressing for invention.

The meeting where you ask nothing

The client says what they want. You say yes. Everyone leaves happy and you build the wrong thing. This happens because asking questions feels like admitting you do not already know, and because the client is busy and you do not want to be difficult.

It is worth being slightly difficult for about twelve minutes. Here are six questions, in an order that works, with what each is actually for.

The six

1. "Who do you want to see this, and can you describe the last person like that who actually bought?"

The second half is the working part. The first half gets you a demographic; the second gets you a human being with a name and a sentence they said. If they cannot describe one, that itself is the finding: the business does not know who buys, and your poster will not fix that.

2. "What do you think they believe about you right now?"

Asked flatly, this produces a defensive answer. Asked as "what is the thing people get wrong about you", it produces the truth, often with feeling. The complaint you get here is usually the whole brief.

3. "What should happen after somebody sees it?"

Push for an action, not a state. "Feel confident about us" is unfalsifiable. "Call the number", "come to the shop on Saturday", "stop worrying enough to read the price list" can all be checked.

4. "What has been tried before, and what happened?"

The most under-asked question in the trade. There is almost always a previous attempt. Whether it worked matters less than why they think it did not, because that sentence contains the standard you are being judged against.

5. "Is there anything you already know you will not approve?"

This saves your favourite concept. Every client has a forbidden list: a colour the competitor uses, a word the founder hates, a style that reminds them of a failed campaign. They will not volunteer it, and they will absolutely reject work over it.

6. "Who else has to say yes?"

Asked last because by then they are comfortable. The answer changes your timeline, the number of rounds, and how you present. A concept that survives one enthusiastic owner dies in a committee unless you build the case for it into the presentation.

When they cannot answer

Often they genuinely do not know, and pressing harder produces invention rather than information. Three moves, in order of cost:

- **Offer two wrong answers.** "Is it more that people think you are expensive, or that they have never heard of you?" People who cannot generate an answer can almost always choose between two, and the correction — "neither, it's that they think we only do weddings" — is better than either option you offered.
- **Ask for artefacts instead of opinions.** The last three enquiry emails. A photo of the shop. The price list. Facts are easier to hand over than judgements, and you can read the judgement out of them yourself.

- **Propose and let them react.** Write the brief as you understand it, in six lines, and send it. People who cannot author a brief are excellent at correcting one. This is the single most reliable technique in the list.

Write it down in front of them

Type or write the answers during the conversation, and read the summary back before you leave. It feels awkward and it is the difference between a shared understanding and two private ones. The phrase to use is: "so the job is to get salon owners who think delivery takes a week to try one order — and the piece has to work on a phone with no sound. Have I got that right?"

Then send it in writing the same day. Not as a contract — as a record. When the fourth revision round asks for something that contradicts answer 3, you have a calm, non-accusatory thing to point at.

The honest limit

This process finds what the client knows and can say. It does not find what is true. A client can sincerely tell you their customers value craftsmanship when the customers are in fact buying on price and nothing else, and the discovery meeting will not catch that — every answer came from the same source, so there is nothing to disagree with.

The only correction is a second source, and it need not be expensive: ten minutes of competitor reviews, one conversation with the person who answers the phone, or three sentences from an actual customer. When the second source contradicts the client, you do not win the argument by asserting it. You show them the three sentences. Evidence they can read themselves changes minds; your professional opinion, offered against their lived experience of their own business, rarely does.

BEFORE YOU MOVE ON

A client cannot articulate what customers believe about them. Why is drafting a six-line brief for them to correct more effective than asking the question again more clearly?

- A. It shows professionalism, which makes the client more willing to open up in future meetings.
 - B. It moves them from authoring, which they cannot do, to correcting, which they can.
 - C. It commits them contractually to the brief, so later changes become billable.
 - D. It gives you a defensible position if the work is rejected, since you can prove they approved the direction.
-

7 · 9 min

Reference, and the line you actually cross

THE ONE THING TO KEEP

Take structure, mechanism and constraint from reference and never surface — test it by asking whether someone who knows the original would recognise it in yours — and write under each board item what you are taking, because an unlabelled board drifts into copying by itself.

Everybody works from reference; the question is what you take

There is a persistent idea that looking at other people's work is cheating, held mostly by beginners, and a matching idea that anything on a mood board is fair game, held mostly by people who have not been sued. Both are wrong, and the useful distinction is not legal at first — it is structural.

What you can take freely and usefully:

- **Structure.** Where the emphasis sits. That the headline is one word and the explanation is six. That the piece withholds the product until the last third.
- **Mechanism.** *Why* it works. The poster is legible at speed because there are exactly three tones and no mid-grey. The reel holds because the first shot poses a question the second shot refuses to answer.
- **Constraint.** That it was made in two colours, or entirely in one continuous take, or with no type larger than 12pt.

What you cannot take:

- **Surface.** The specific illustration, the photograph, the exact colour-and-typeface combination that makes it recognisably theirs, the particular joke.

The practical test: could a viewer who knows the original recognise it in yours? If yes, you took surface. Taking structure produces work that looks like nothing in your reference pile; taking surface produces a copy with new words.

Build the board to answer a question

A mood board of forty beautiful things is a procrastination device. A useful board is six to ten items with a note under each saying what you are taking from it. Write the note; without it the board becomes an aesthetic mood and you will unconsciously take surface.

— *Kraft-stock gig poster: taking the three-tone limit, not the woodcut style.*

— *Rival's reel: taking the on-screen-text rhythm — one line per 1.5s. Not the music.*

— *Old railway timetable: taking the dense-information-can-be-calm idea.*

Include at least two items from outside your field. Everything in a designer's board coming from other design is how whole categories start to look identical — a mechanism worth naming: each maker samples the current consensus,

adds a small variation, and publishes it back, so the consensus tightens with every cycle. The escape is reference from outside the loop: signage, packaging in a grocery shop, a textile, a scientific diagram, a film from a country whose advertising you have never seen.

Where the law actually is, and where it is not

Copyright protects expression, not ideas. A layout concept, a colour scheme and a general style are not usually protected; a specific photograph, illustration, piece of text, typeface software file or piece of music is. That much is broadly stable across most jurisdictions.

Beyond that, be careful, because the details genuinely differ by country: how much of a work counts as substantial, whether parody or pastiche is defended, whether trade dress protects an overall look, how long a musical phrase can be. These vary, they are decided case by case, and this course cannot give you legal advice. What it can tell you is where to be cautious: **anything recognisable belonging to somebody with a legal department** — a brand's trade dress, a character, a famous photograph, a typeface's actual font file — is the wrong place to be adventurous on a client job.

AI-generated reference deserves its own warning. A model trained on existing work can reproduce a distinctive style closely enough that the origin is recognisable, and whether that is infringement is under active dispute in several countries right now, with cases running in both directions. Do not treat any confident answer you read on this — including a confident answer that it is all fine — as settled, because it is not.

The practical habits

- **Keep the source with the image.** A board of unattributed JPEGs becomes, six weeks later, a board where you cannot tell what you may use. Save the URL.
- **Never put a reference in a client presentation without saying it is a reference.** Clients remember images, not captions, and a borrowed image shown as an idea becomes an expectation.

- **Move one step away on purpose.** After the first version, put the board away and rebuild the piece from your notes only. What survives without the images in front of you is yours.
- **Free sources worth knowing:** public-domain archives such as the Internet Archive and many national library collections, Wikimedia Commons with the licence read properly, and openly licensed typefaces from Google Fonts. Read the licence rather than the summary, especially for anything marked non-commercial.

The uncomfortable part

Work that takes only structure is harder and looks less finished at the first showing, because surface is what makes something look resolved. You will be tempted, on deadline, to reach for the thing that already works. Everybody is. The reason to resist is not moral purity — it is that a piece assembled from other people's surfaces has no argument of its own, and the moment a client asks why it looks like that, you have nothing to say.

BEFORE YOU MOVE ON

Two designers use the same award-winning poster as reference. One produces work that looks nothing like it; the other produces something recognisably derivative. What most likely separated them?

- A. The first designer has more experience, so they naturally transform influences further.
- B. The first took the constraint and the emphasis pattern; the second took the typeface, palette and imagery.
- C. The second designer used more references, so the influences blended less thoroughly.
- D. The first designer changed the colours and the typeface, which is the standard test for whether borrowing has gone far enough to count as original.

8 • 8 min

Agree the measure before you make the thing

THE ONE THING TO KEEP

Agree one outcome measure and one diagnostic measure before production, build the attribution mechanism into the artwork itself, and state up front how noisy the result will be — because a limitation named in advance is a position and the same sentence afterwards is an excuse.

The argument you will have in four weeks

The work goes out. Two weeks later the client says it did not work. You ask how they know. They say it feels quiet. You say the engagement was up. Neither of you is lying and neither of you can win, because nobody agreed in advance what the piece was supposed to do.

This argument is entirely preventable with one line written before you start.

Three kinds of measure, and which one is honest

Vanity measures — views, likes, impressions, reach. They are easy to get and they move for reasons unconnected to your work: the algorithm, the hour of posting, one share by somebody with a large following. They are not useless; they tell you whether anybody saw the thing at all, which is the first of the three failure stages. But a piece with 50,000 views and no enquiries has failed, and a measure that calls it a success is the wrong measure.

Proxy measures — saves, watch-through rate, time on page, scans of a QR code, clicks to the site. These sit closer to intent. Watch-through especially: for a short video, the proportion still watching at three seconds and at the end tells you where attention broke, which is diagnostic rather than merely score-keeping.

Outcome measures — enquiries, walk-ins, bookings, sales, applications. This is what the client actually wants and the hardest to attribute, because a flyer, a reel, a friend's recommendation and the weather all landed in the same week.

The honest answer is to agree **one outcome measure and one diagnostic measure**, and to say plainly how confident attribution can be.


Make attribution possible while you are still designing

This is the bit that has to happen before production, not after, and it is almost always skipped.

- A different phone number on the poster than on the website. Costs nothing on most systems, and it is the cleanest attribution available to a small business.
- A specific phrase to mention: "say KALYAN when you call".
- A separate short link or QR code per surface, so the flyer and the reel are distinguishable.
- A one-line question from whoever answers the phone: "where did you see us?" — imperfect, people misremember, but it beats nothing and it is free.

If none of these are possible, say so now: "we will not be able to tell this apart from everything else running that week, so we will judge it on enquiries during the period and accept that as rough". A stated limitation before the fact is a professional position. The same sentence after the fact is an excuse.

Enquiries in six ordinary weeks, no campaign running



Nothing was published in any of these weeks. The ordinary spread is about five, so a rise from twelve to fifteen after a flyer is inside the noise, and a campaign judged on it is being judged on the weather.

Write the number down, roughly

Not because the forecast will be right — it will not — but because it turns a vague hope into something checkable. "Last month, twelve enquiries. If this works, we would expect somewhere around twenty, not two hundred." Now everybody's expectations are in the same room before any money is spent.

This is also the moment to say if the target is unreachable by design work. If they need forty enquiries and the poster will be seen by roughly four hundred people, arithmetic says the campaign is too small regardless of how good the poster is. That conversation is much easier before you have made it.

The sample-size problem, in plain terms

Small businesses run small campaigns, and small numbers are noisy. The difference between eleven enquiries and fifteen is not evidence that a design worked; a normal quiet week and a normal busy week differ by about that much. You need either a larger volume or a longer period before a difference means anything, and the honest thing is to say which of the two you are relying on.

Beware, too, of judging a change by the first day. A piece often does best in its first hours and then settles; comparing day one of the new work against a weekly average of the old is comparing two different things and will usually flatter the new.

What the measure cannot tell you

A number tells you whether something happened. It does not tell you why, and it cannot distinguish a good piece badly distributed from a bad piece well distributed. Two additional habits close some of that gap:

1. **Keep the diagnostic measure.** A three-second drop-off says the opening failed. A high watch-through with no clicks says the opening worked and the ask did not.
2. **Collect three sentences of what people actually said.** "I didn't realise you delivered" is worth more than a percentage, because it names the

belief that moved, and belief is what you were working on in the first place.

Write the measure into the brief, one line: *"Success is enquiries mentioning KALYAN over four weeks; we expect around twenty, and we will read three-second drop-off to see whether the opening is doing its job."* That line ends the four-week argument before it starts.

BEFORE YOU MOVE ON

A reel gets 80,000 views, a 4% watch-through to the end, and no enquiries. What does this pattern most usefully indicate?

- A. The piece succeeded at reach and failed at conversion, so the call to action at the end needs to be stronger.
- B. The attribution was broken, so enquiries probably did arrive but were not counted.
- C. Attention was bought but not held, so the failure sits early in the piece.
- D. The audience targeting was wrong, since the right audience would have watched to the end.

CASE STUDY · MODULE 1

The donor poster that answered the wrong objection

A district blood bank in Nagpur, preparing a fortnight of college drives after two years of falling repeat donation.

The brief arrived as one sentence from the director: a poster to raise awareness of blood donation. There was a budget for 2,500 A3 sheets, a week of visits to four colleges, and a target of 900 units.

The designer asked the seven questions and got through five of them. Question two stopped everything. What do they currently believe about donating here? The director said people did not understand that donation was safe. The communications officer said people were busy. Neither of them had asked anybody.

The second source cost one conversation

So she went to the counter and asked the two attendants what people say just before they change their mind and leave. They could recite it without pausing, because they hear it every day: how long is this going to take.

The register supported them. The bank's own arrival and discharge stamps put the median donor in the building for one hour and fifty minutes. Almost all of that was waiting for the medical officer, who screened between eleven and one because that was when he was free between theatre lists.

The before state was not ignorance about safety. It was a specific memory held by people who had already donated once: I lost an afternoon. Awareness posters are written for people who have never heard of you. These people had heard of the bank and had decided.

The decision, and what it cost on each side

The honest line the brief now pointed at was a promise about time. Booked slot, forty minutes, out by ten past. That line could not be printed unless the

bank changed how it worked on the two drive days: the medical officer screening from nine, a second counter staffed, and donors given a slot rather than a queue number.

That was two extra staff shifts and an argument with the pathology head, who said moving the screening earlier would collapse the sample-processing queue behind it and that the medical officer had a theatre list nobody could move.

The alternative was cheap and safe. Print a poster that said donate blood, save a life, over a photograph of a smiling donor. Nobody in the building would object to it. The designer could also predict its result, because it was approximately last year's poster and last year's result was the reason she had been called.

And there was a third option that was worse than both, which was the tempting one: print the forty-minute promise without making the change. That version would work once. The fastest way to destroy a claim is to get a lot of people to check it, and 900 people checking a false forty minutes would cost the bank the donors it still had.

Agreeing the measure before printing

Before anything went to press they wrote one line into the brief. Success was repeat donors, counted as arrivals whose donor identification number already existed in the register, over four weeks. The diagnostic measure was the gap between the arrival stamp and the discharge stamp, which the bank was already recording and had never looked at. A separate helpline number was printed on the poster and nowhere else, so a call could be attributed.

The designer also wrote the sentence she expected to need later: this fortnight also contains a college festival and a newspaper appeal, so we will not be able to separate the poster from either, and we are accepting that as rough. A limitation named before the fact is a professional position; the identical sentence afterwards is an excuse, and the difference is entirely in the date at the top of the document.

The forecast went into the same line, deliberately rough. The bank had been seeing about 40 repeat donors a fortnight. If the change worked, expect some-

thing in the region of 60 to 80, not 300. Writing that down cost a minute and put everybody's expectations in one room before any money moved, which is also where the conversation about arithmetic happened: 2,500 posters reaching roughly 9,000 students at four colleges cannot produce 900 units on any conversion rate anybody has ever measured, so the drive needed the staff canteen and the two mills as well.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

The pathology head was brought into the room rather than told afterwards, and turned out to have a workable objection and a workable fix: screening could start at nine on two named days if a second technician handled samples those mornings. The bank ran one drive with the change and the forty-minute promise, and one without it, on the safe poster.

The drives produced 71 and 44 repeat donors against a normal fortnightly figure of about 40. Two drives is not evidence and the designer said so in the report. The diagnostic measure was not noisy: median time in the building was 46 minutes on the changed days and 108 on the others. That number is what the bank acted on, because it is a fact about the bank rather than about one poster.

The helpline attributed 19 calls. Three of them said, unprompted, that they had come back because of the time.

The director's brief named a destination and no starting point. Why does that make a poster hard to design rather than merely vague?

Every piece of work is the distance between what the viewer believes now and what they should believe afterwards, and the same destination needs completely different work depending on where they are standing. Reliability aimed at someone who

has never heard of you is noise; aimed at someone who was let down once it reads as a lie. Without the before state the designer cannot tell which of four different pieces is required, so the choices get made on taste.

Why did she ask the counter attendants rather than pressing the director for a better answer?

A discovery meeting finds what the client knows and can say, and every answer in it comes from the same source, so there is nothing to disagree with. The objection log held by whoever answers the phone or faces the queue is a second source, it costs one conversation, and it records what people said at the moment they decided not to act, which is exactly the before state the brief needed.

Printing the forty-minute promise without changing the process was the cheapest option available. What is the argument against it that does not rely on ethics?

Specificity is what makes a claim believed, and the same specificity is what makes it checkable. A vague claim that fails disappoints one person quietly; a precise claim that fails is contradicted by nine hundred people on one afternoon, and each of them now discounts everything else the bank says. The campaign would have spent an asset the bank cannot rebuy in order to buy one fortnight.

MODULE 1 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 A supplier's flyer says **SAME-DAY DELIVERY. GUARANTEED**, set large and spaced well. Customers who were let down by a previous supplier ignore it. What is the mechanism of the failure?
 - A. The guarantee is a legal commitment, so readers discount it as boiler-plate they cannot enforce.
 - B. The typeface is not distinctive enough for the claim to be remembered a day later.
 - C. It announces the after state, and a belief arrives as a rival claim from an interested party.
 - D. Beliefs shift only after a message has been repeated enough times for a reader to recognise it.

- 2 A client genuinely cannot say what their customers believe about them now. Pressing the question harder is producing invention rather than information. Which move is the most reliable?
 - A. Ask them to write a full creative brief document before any design work begins.
 - B. Decide the message yourself and confirm it once there is something to look at.
 - C. Keep asking, in different words, until an answer that can be used emerges.
 - D. Write the brief as you understand it in six lines and send it for correction.

- 3 A brief describes the audience as women aged 25 to 40 in tier-one cities. Three hours in, you have two headlines: one dry and specific, one warm and general. Why does the demographic not settle it?
 - A. Demographics are usually guessed at by the client and are therefore too inaccurate to design against.
 - B. Design decisions are about whether one person would stop and believe, and a category can do neither.
 - C. Age bands that wide are useless; a five-year band would have answered the question.
 - D. Demographic data is expensive to obtain and is rarely current enough to trust on a live job.

- 4 Two days go into a poster built on a full-bleed photograph with white type reversed out of the dark half. The print budget bought one ink on kraft stock. Which constraint would have rejected the idea on day one?
 - A. The production route, because one ink on an uncoated stock cannot hold that image.
 - B. The budget, because a photographic poster always costs more than a typographic one.
 - C. Format and quantity, because 200 copies is too few to justify a photograph.
 - D. Where it will physically be, because a corridor is too dark for a dark poster.

- 5 You have built a piece from a reference board. Which test tells you whether you took surface rather than structure?
 - A. Whether you changed more than thirty per cent of the elements you borrowed.
 - B. Whether the reference carries a licence that permits commercial reuse and adaptation.
 - C. Whether the original is old enough for its copyright to have expired where you publish.
 - D. Whether a viewer who knows the original would recognise it in yours.

- 6 A shop had 11 enquiries the week before a flyer and 15 the week after. The client calls it a success. Why is that not yet evidence?
- A. Enquiries are a vanity measure and should never be used as an outcome measure.
 - B. Fifteen is too small a number for any statistical test to be run on it.
 - C. A normal quiet week and a normal busy week already differ by about that much.
 - D. A flyer's effect cannot appear until at least three months of distribution have passed.

MODULE 2

Making one message survive

A viewer gives you a glance, and a glance sees contrast, position and size — not your carefully written third line. This block is the mechanics of that glance: the five levers that create emphasis, why contrast rather than colour is the only thing the eye actually reports, how much each added element costs the rest, what space is doing when it looks like nothing, and the four cheap tests that break a design before an audience does.

BY THE END OF THIS MODULE YOU CAN

Rank a design into three levels using at most two emphasis levers per element, measure its contrast against a stated ratio, prove the hierarchy with a blur, greyscale, flip and distance test, and delete the element you like most when it competes with the message

9 · 8 min

A design that says four things says none

THE ONE THING TO KEEP

Rank every element into three levels and check by blurring the design — if more than one thing survives the blur, nothing does.

How a piece ends up saying nothing

Nobody sets out to make a flyer with six messages. It arrives one addition at a time. The client wants the QR code. Then all three social handles, because the

intern set up TikTok. Then the discount. Then a photograph of the shopfront so people recognise it. Then, at the last minute, the founder's quote.

Every one of those additions is reasonable on its own. Together they produce a surface where nothing is bigger than anything else, and here is the mechanism that matters: the eye needs a single entry point. It finds one by looking for the strongest contrast — in size, in weight, in colour, in isolation. When six elements compete for that, the eye has no reason to prefer any of them, scan order becomes arbitrary, and an arbitrary scan is uncomfortable. The reader resolves the discomfort by leaving.

This is not a taste judgement. You can watch it happen with the three-second test from the first lesson.

Three levels, and only three

The workable structure is almost always the same:

- **One thing** that is unmissable in half a second.
- **Two or three things** that are found in the next two seconds.
- **Everything else**, which is there for the person who has already decided and is now looking for the address.

The discipline is not "remove things". It is "rank things". The GST line and the opening hours can stay — they just go in level three, at 9pt, and they stop fighting.

Three levels, and what belongs in each

Level one	One thing, unmissable in half a second
Level two	Two or three things, found in the next two seconds
Level three	Address, hours, GST line, at nine point

The discipline is ranking, not removing. The GST line and the opening hours can stay; they move to level three and stop competing for the entry point.

Making a level-one element is not about size alone. Contrast comes from difference against its neighbours: one heavy word among light ones, one saturated colour on a muted page, one element with air around it while the rest is

packed. Isolation is the most underrated of these and costs nothing but restraint.

The blur test, done properly

Squinting works. Actually blurring works better, because it is repeatable.

In Photopea — free, runs in a browser tab, opens PSD files — flatten a copy and apply Filter, Blur, Gaussian Blur at around 15 to 20 pixels on a 1000-pixel-wide design. GIMP has the same thing under Filters, Blur, Gaussian Blur. On a phone, take a screenshot and pinch it down to the size of an app icon.

Now look. Whatever is still identifiable at that blur is your level one. If three blobs survive equally, you have three level ones, which is zero. If nothing survives, the piece has no entry point at all and every element is politely the same weight — the most common look of competent, ineffective work.

Do this before you show anyone anything. It takes forty seconds and it catches the problem the client will describe as "make it pop", which almost always means "I cannot find where to start".

Where one message is the wrong rule

A menu is not a one-message design. Neither is a timetable, a spec sheet, a form, or a slide of exam dates. These carry many facts by definition and forcing a single hero onto them makes them slower to use, which is the only thing they are judged on.

For those, the rule becomes one *structure* rather than one message: consistent groups, a clear order, and hierarchy applied inside each group rather than across the whole page. Someone scanning a menu for the vegetarian section is running a search, not receiving a message, and your job is to make the search short.

The other honest exception is a piece with genuine sequence — a brochure, a scrolling page, a video. There you get one message *per moment*. The spread you are looking at right now has one job; the next spread has the next. That is not a loophole. Applied to each screen separately it is the same rule.

The hardest part

The hardest part of this lesson is not technical. It is telling a client that the QR code, the three handles and the founder's quote cannot all be prominent. Do not argue about taste, because taste is a matter of opinion and theirs is as valid as yours. Argue about ranking: "If a person only reads one thing here, which one earns you the money?" That question is answerable, it is theirs to answer, and once they answer it the layout follows.

Sometimes they will say all four. Then make the version they asked for and one alternative, run the three-second test on both with two people from their staff, and show the results. Evidence beats argument and it costs you twenty minutes.

Do this today

Take your busiest recent design. Blur it. Write down what survives. Then re-write the piece so exactly one thing survives — not by deleting the rest, but by demoting it. Compare the two at a glance on your phone. In most cases the demoted version reads as *more* professional, not emptier, which is the part that surprises people.

BEFORE YOU MOVE ON

A client keeps saying a flyer needs to "pop" but rejects every brighter colour you try. A 20-pixel Gaussian blur of the flyer shows four separate elements still clearly readable. What is the most useful reading of this?

- A. The colour palette is too low in contrast overall, so the whole piece needs a stronger accent colour rather than a change in layout.
- B. The blur test is only meaningful for digital surfaces, since a printed flyer is held close and read in full.
- C. The client's feedback is subjective and should be met with two options and a preference vote rather than a diagnosis.
- D. Four elements are competing to be the entry point, so the eye has no reason to start anywhere; the fix is demoting three of them, not intensifying any.

10 · 9 min

There are only five ways to make something look important

THE ONE THING TO KEEP

Emphasis comes from five levers — size, weight, contrast, position and isolation — and each element should use at most two, because every lever spent consumes contrast the lower levels need to stay distinguishable.

Emphasis is comparative, and there is a short list of comparisons

Nothing on a page is important. Things are only more important than the things next to them, and there are five ways to arrange that. Every hierarchy you have ever admired is some combination of these.

1. **Size.** The biggest thing is read first. Crude, reliable, and the one most people reach for.
2. **Weight and density.** Bold against regular; a solid block against an outline; a dense cluster against a sparse one.
3. **Contrast against the ground.** Dark on light, light on dark, one saturated colour among muted ones. This is the strongest lever at a distance and the one people under-use.
4. **Position.** Top and left in left-to-right reading cultures; the optical centre, which sits slightly above the geometric centre; the end of a sequence.
Where something sits is a claim about when it should be read.
5. **Isolation.** A thing alone in space outranks a thing in a crowd, regardless of its size. This is the lever that costs nothing but nerve.

The rule that does most of the work: use two

One lever is usually not enough to be seen. Three or more is shouting, and shouting has a mechanism behind its failure. Each lever consumes some of the design's available contrast — the total difference the viewer can perceive between elements. Spend it all on the headline and the second-level items have nothing left to distinguish them from the third, so the hierarchy collapses from three levels into two: the shouting thing, and everything else.

So: big **and** isolated. Or bold **and** high-contrast. Not big and bold and red and centred and boxed, which is what a flyer looks like when five people each added their own emphasis.

Make the gaps obvious


A hierarchy fails more often from insufficient difference than from the wrong choice. If your headline is 28pt and your subhead is 24pt, the viewer does not perceive two levels; they perceive one level with an inconsistency in it. The difference has to be unmistakable at a glance.

A practical starting scale, which you should then break on purpose:

level 1	48pt bold	
level 2	24pt bold	(half)
level 3	16pt regular	(two thirds)
support	12pt regular	

The ratios matter more than the numbers. Roughly 1.5× to 2× between adjacent levels reads as deliberate; 1.1× reads as a mistake. Type scales built on a fixed ratio — 1.25, 1.333, 1.5 — exist for exactly this reason, and every one of Inkscape, Scribus, Figma, GIMP and Photoshop will let you type the resulting numbers straight in.

A starting type scale, and the gaps that read as deliberate



The ratios matter more than the numbers. Forty-eight over twenty-four is read as two levels; forty-eight over forty-four is read as one level with a mistake in it.

Isolation is the cheap lever nobody pulls

If you cannot make the important thing bigger, because the format is small, make everything else smaller and move it away. A 16pt line alone in the top third of a poster outranks a 40pt line wedged between a logo and a QR code. The eye finds edges and isolated figures fast, before it has resolved any detail, so an isolated element is registered as a distinct object while a crowded one is still part of a texture.

This is also the lever clients resist, because empty space looks like unused space and unused space looks like something you did not do. The counter-argument to have ready: the empty area is not empty, it is the reason the important line is readable from across the room.

What the levers cannot do

They establish rank. They cannot establish meaning. A perfectly ranked design in which the largest element is the wrong message is a well-built failure, and this is much commoner than a badly ranked one. Before you apply any of this, the sentence from the brief has to be settled, because hierarchy is an amplifier and it amplifies whatever you put at level one.

They also have a cultural boundary worth knowing. Position as a lever assumes a reading direction. In a left-to-right script, top-left is the strong entry; in an Urdu or Arabic layout the strong entry is top-right, and mirroring a layout without rethinking it produces something that reads backwards to the people it is for. Size, weight, contrast and isolation carry across scripts; position does not carry for free.

The check

Name your three levels out loud before you open the file: what is read first, second, third. Then, after you have built it, hand it to somebody for two seconds and ask what they saw. If they name your level-two element, you have spent your contrast in the wrong place — and the fix is almost never to make level one bigger. It is to make levels two and three quieter.

BEFORE YOU MOVE ON

A poster's headline is set large, bold, red, centred and inside a box. Viewers report seeing the headline and then nothing else registering. What is the mechanism?

- A. Five levers on one element exhaust the perceivable contrast, so levels two and three collapse together.
- B. Red is an unusually aggressive colour that dominates any composition it appears in, so the palette rather than the emphasis is the real problem here.
- C. The headline is too large in relation to the format, and reducing its size would restore the balance.
- D. Centring removes the reading entry point, so the eye has nowhere to go after the headline.

11 · 9 min

The eye reports differences, never values

THE ONE THING TO KEEP

The visual system reports differences rather than values, so legibility is a measurable relationship — at least 4.5:1 for body text — and a hierarchy built on hue alone vanishes in greyscale, in one-colour print and for anyone with colour vision deficiency.

Your visual system does not measure brightness

A patch of grey on a dark background looks light. The same patch on a white background looks dark. Nothing about the patch changed. This is not an illusion in the sense of a trick — it is how the system works. Retinal cells respond strongly to edges and differences and weakly to steady, uniform values, so what reaches your awareness is already a map of contrasts rather than a map of brightnesses.

Every practical consequence in this lesson follows from that one fact. A colour has no properties on its own. It only has a relationship to what surrounds it, and design is the business of arranging those relationships.

Contrast is measurable, and there is a free way to check it

For text legibility, the widely used figure comes from the Web Content Accessibility Guidelines: a contrast ratio of at least **4.5:1** between text and its background for normal-size text, and **3:1** for large text — large meaning about 18pt regular or 14pt bold and upward. The ratio runs from 1:1 (identical) to 21:1 (pure black on pure white).

Those numbers were derived for screens and for people with moderately reduced vision, and they are a floor rather than a target. They are worth knowing because they turn an argument about taste into a number anybody can

check. Free checkers exist in every browser's developer tools, in the Colour Contrast Analyser, and inside Inkscape and Figma via extensions.

The usual failures, all of which pass a designer's eye on a good monitor in a dark room and fail on a phone in sunlight:

- Light grey body text on white — often around 2.5:1.
- White text on a photograph, where the ratio varies across the image and is fine over the dark half and invisible over the sky.
- Two mid-tone brand colours placed on each other because they are "the brand colours".

Contrast is not only tonal

Tone is the strongest axis, but there are others, and mixing them is what makes work feel composed rather than merely loud:

- **Scale** — one large shape among small ones.
- **Shape** — one curve in a field of rectangles.
- **Texture** — a photograph against flat colour.
- **Direction** — one diagonal in a grid of horizontals.
- **Density** — a crowded block beside an open one.

A design with strong tonal contrast and nothing else tends to look like a warning sign. A design with five kinds of contrast at once looks chaotic. Two is usually right, which is the same arithmetic as the emphasis levers, for the same reason.

The test that settles it: remove the colour

Convert the work to greyscale. It takes one keystroke in GIMP (Image → Mode → Grayscale), Krita, Photopea or Photoshop, and there is a greyscale filter in browser developer tools. If the hierarchy survives, your contrast is doing real work. If everything turns into the same mid-grey soup, you built the hierarchy out of hue alone.

Hue alone is the specific trap. Red and green of identical lightness are strikingly different in colour and almost identical in tone. Around one in twelve men and one in two hundred women has some form of colour vision deficiency, most commonly in exactly that red-green axis, so a chart or a status indicator distinguished only by hue is unreadable to a meaningful share of any audience. It also disappears entirely in a one-colour photocopy, which is how a surprising number of flyers end up being reproduced.

The fix is never to avoid colour. It is to make sure colour is the second signal and tone, shape or position is the first: a red line that is also thicker, a green tick that is also a tick.

Why it fails in the room it was not designed in

You are looking at a bright, calibrated screen at arm's length in controlled light. The viewer is looking at a cheap phone at 30% brightness in daylight, or at a print where the ink has spread slightly into uncoated paper and thin type has thickened and closed up. Every one of those conditions **reduces** effective contrast; none increases it.

So design with headroom. If 4.5:1 is the floor, land nearer 7:1 for anything that has to work outdoors, on cheap paper, or at a distance. The cost of extra contrast is almost always aesthetic restraint. The cost of insufficient contrast is that the piece does not function at all for part of your audience, and you never hear about it, because people who cannot read something do not send feedback — they move on.

BEFORE YOU MOVE ON

A dashboard marks healthy items green and failing ones red, at identical lightness. It tests well in the studio. What is the mechanism of its failure in use?

- A. Hue is the only signal, and hue is the channel that colour vision deficiency degrades.
 - B. Red and green carry opposite cultural meanings in some markets, so the semantics do not transfer internationally.
 - C. Saturated colours fatigue the eye over long viewing sessions, so users stop distinguishing the states.
 - D. The dashboard needs a legend, since colour coding is only meaningful once its key has been learned.
-

12 · 9 min

You see far less at once than you think

THE ONE THING TO KEEP

Detailed vision covers only about two degrees, so viewers navigate by coarse peripheral cues — contrast, shape and especially motion — which is why a blur simulates what pulls the next glance and why the F-pattern describes pages with no hierarchy rather than prescribing one.

The sharp part of your vision is about the size of a thumbnail

Detailed vision comes from the fovea, a small central region of the retina, and it covers roughly two degrees of your visual field — about the width of your thumbnail held at arm's length. Everything outside that is progressively blurrier. You do not notice, because the eye moves three to four times a second in rapid jumps called saccades, and your brain assembles a stable impression from the samples.

Two consequences that matter for everything you make:

- **A viewer never sees your whole design in detail.** They see a sequence of thumbnail-sized samples, and a low-resolution impression of everything else.
- **The periphery still works, but it only reports coarse information:** contrast, motion, gross shape, edges. It cannot report which typeface you used or what the small print says. So the periphery decides *where to look next*, and it decides on the basis of the coarse stuff.

This is the whole justification for the blur test. Blurring a design is a crude simulation of peripheral vision, and what survives the blur is what will pull the next saccade.

Entry, path, exit

A composition should be able to answer three questions:

1. **Where does the eye land first?** Usually the highest-contrast or largest element, modified by reading habit. If two elements compete, the viewer's entry point is effectively random, which means you have no control over anything that follows.
2. **Where does it go next?** Movement between elements is driven by alignment, direction, grouping and implied lines — a face looking towards the headline, an arrow, a strong edge, a row of items that reads as a sequence.
3. **Where does it leave?** Something has to be last, and it should be the thing you want them to act on. If the last stop is the logo and the action is the phone number, you have arranged for people to finish in the wrong place.

Draw these three as arrows on a printout. It is a two-minute exercise and it exposes the layout where everything is beautifully arranged and nothing leads anywhere.

The F-pattern is not a law

You will be told that people read screens in an F shape. That came from eye-tracking studies of text-heavy pages — long, largely unformatted prose with

little visual hierarchy. Under those conditions the F shape is a real and repeatable finding.

What it describes is behaviour in the absence of design. Give people a strong hierarchy, a photograph, a pull quote or a clear entry point, and the pattern changes completely, because contrast overrides habit. Treating the F as a template — jamming everything important along the top and left — is applying a finding outside the conditions that produced it, and it produces exactly the unformatted page the F describes.

The honest version: in the absence of anything stronger, reading direction sets a weak default. Your job is to supply something stronger.

Motion outranks almost everything

Peripheral vision is disproportionately sensitive to movement — a legacy of needing to notice things approaching from the side. In practice, one moving element in a still layout captures the eye regardless of size, contrast or position, and will keep recapturing it.

This is why an auto-playing video beside body text means the body text does not get read, and why a blinking cursor animation next to a headline steals the headline's job. Use motion once, for the thing you want looked at, or not at all. And provide a way to stop it: for some people, including those with vestibular disorders, continuous motion causes genuine discomfort rather than mild annoyance, and most platforms expose a reduced-motion preference you can respect.

What this does not license

Knowing how gaze works tells you how to direct attention. It does not tell you whether you have earned it, and the same mechanisms that make a public-health poster legible make a manipulative advertisement effective. A course on attention mechanics that stops here has taught half a subject; the second half — where directing attention turns into extracting it — has its own lesson later in this course.

One practical caution as well. Gaze is not comprehension. Eye-tracking shows where people looked, not what they understood or believed, and a design that reliably pulls the eye to a headline nobody finds credible has solved the easiest of the three problems. Attention is a precondition. It is not the result.

BEFORE YOU MOVE ON

A landing page follows the F-pattern advice by stacking everything important along the top edge and the left column. Attention data shows people scanning in an F. Why is this weak evidence that the approach worked?

- A. Eye-tracking measures fixation rather than comprehension, so the data cannot show whether the content was understood at all.
- B. The F was observed on pages with no hierarchy, so supplying none just recreates those conditions.
- C. The sample was probably too small, since scan-path studies need many participants before a pattern is reliable.
- D. Reading direction varies between scripts, so the F only applies to left-to-right audiences.

13 · 8 min

Every element you add taxes the ones already there

THE ONE THING TO KEEP

Elements cost $n(n-1)/2$ relationships plus the viewer's decision time, so subtract before rearranging and group what cannot be removed — density is only a problem when it is unstructured.

Addition is not free, and the arithmetic is worse than linear


A layout with three elements has three relationships to manage: A to B, A to C, B to C. Add a fourth and you have six. A fifth gives ten, a sixth fifteen. The general figure is $n(n-1)/2$, and it is the reason a design with eight elements is not slightly harder than one with four — it is four times harder to resolve, because every new object has to be positioned, sized and weighted against everything already present.

This is not an abstraction about beauty. It is the direct explanation for why committee flyers look the way they do. Each person adds one item that is individually reasonable — a QR code, a second phone number, an award badge, a line about parking — and the total number of relationships grows quadratically while the design time stays the same. Nobody made a bad decision. The decisions were made one at a time.

Attention is also a fixed budget

The viewer arrives with a few seconds. Whatever they spend deciding what to look at is not spent understanding anything. Two elements of equal weight cost more than the sum of their parts, because before reading either the viewer has to resolve which one comes first, and that resolution is pure overhead.

Relationships to manage as elements are added



Eight elements are not twice as hard as four. Each addition is individually reasonable and the total grows as $n \times n - 1$, over two, while the design time stays the same.

So the cost of an element is paid twice: once in the composition's complexity, and once in the viewer's decision time.

Ask three questions before anything goes on the page

1. **What job does it do that nothing else does?** If two elements do the same job — a headline and a strapline that say the same thing differently — one is redundant and should go, not be made smaller.
2. **What happens if it is absent?** If the answer is "nothing in particular", that is the answer. If the answer is "somebody in the client's office will be annoyed", that is a political cost, and it should be named as one rather than dressed up as a design requirement.
3. **What does it displace?** Space is finite. The award badge does not arrive in empty space; it arrives by shrinking the headline or by eating the margin that made the headline readable.

Subtract before you rearrange

When a layout is not working, the reflex is to move things. Try removal first, because it is faster and more diagnostic. Hide one element — a layer toggle in GIMP, Krita, Inkscape, Scribus or Figma, one click, fully reversible — and look again. Go through every element in turn. You will usually find two or three that the design is better without, and at least one that you had assumed was load-bearing and is not.

This is also the most effective way to have the conversation with a client. "Here is the version with the award badge and here is the version without it" is a comparison they can judge. "The award badge weakens the hierarchy" is an opinion they have no way to evaluate.

Grouping buys you back some of the cost

You cannot always delete. The opening hours, the address and the phone number all have to be there. But three items placed close together with consistent alignment are perceived as **one** object, not three, because the visual system groups by proximity and similarity before you are aware of it.

That is the practical escape route. Five loose items become two groups; two groups are one relationship rather than ten. It costs you nothing except discipline about spacing: the space *between* groups has to be clearly larger than the space *within* a group, or the grouping does not read and you have five items again, now with inconsistent gaps.

The limit of the principle

"Less is more" is not always true, and applied without thought it produces work that is stylish and useless — a poster so reduced that it no longer says when the event is. The principle is not minimalism. It is that every element must earn its cost, and some elements are worth a great deal.

Dense work can be excellent. A railway timetable, a technical diagram, a product spec sheet: these carry enormous amounts of information and can be entirely calm, because the density is organised into a repeating structure the eye learns once and then applies. The thing that fails is not density. It is *unstructured* density — many items, each competing individually, with no system that lets the viewer stop evaluating them one by one.

BEFORE YOU MOVE ON

A flyer has eleven separate items and feels chaotic. The designer groups the address, hours and phone number tightly together and increases the space around that block. Why does this help more than shrinking each item would?

- A. Smaller items are harder to read, so shrinking would have traded chaos for illegibility without solving anything.
 - B. Contact details are low-priority information, so consolidating them signals their rank to the viewer.
 - C. Increasing surrounding space raises the design's overall contrast, which makes the remaining elements easier to distinguish.
 - D. Proximity makes the three read as one object, cutting the count of competing relationships.
-

14 • 8 min

The empty part is doing the most work

THE ONE THING TO KEEP

Gaps are grouping instructions: keep space within a group smaller than the space around it, work in multiples of one base unit, and when space is unaffordable switch to small consistent gaps rather than abandoning the system.

Space is not what is left over

Beginners fill space because space looks like an unfinished job. Clients ask for it to be filled because they are paying for something and emptiness does not look like something. Both are reading the space as absence, when in fact it is the mechanism by which everything else becomes readable.

The visual system groups by proximity automatically and pre-consciously. Things near each other are one thing; things far apart are separate things. That means every gap you set is a statement about what belongs with what —

and if all your gaps are the same, you have stated that everything belongs with everything, which is the same as stating nothing.

The single rule that fixes most layouts

The space within a group must be smaller than the space around it.

That is nearly the whole of it. A caption sitting 20px below its image and 20px above the next headline belongs to neither. Move it to 8px from its image and 40px from the headline and it snaps into place, and the page suddenly looks designed rather than typed. Nothing else changed.

The most common real-world version of this failure: a form where each label sits exactly halfway between its own field and the field above, so the eye has to work out the pairing from meaning rather than seeing it. That is a small tax charged to every single user of the form, forever.

Use one spacing unit

Pick a base — 8px on screen is a common choice, or a millimetre grid in print — and use multiples of it. 8, 16, 24, 32, 48. The reason is not tidiness. It is that arbitrary spacing produces near-misses, and a 14px gap next to a 16px gap reads as an error rather than as a distinction, because the difference is too small to be intentional but large enough to be noticed.

Write the unit into the one-page brand system, alongside the type scale. In six weeks you will not remember it, and the drift between the version you made in January and the one you make in March is a large part of what makes a body of work look amateur.

Margins are not decoration

The outer margin has a specific job: it stops the composition from fighting the edge of the page. Type running close to a trimmed edge looks accidental, and in print it may actually be cut off, because commercial trimming has a tolerance of a millimetre or two in either direction. This is why bleed exists — the artwork extends past the trim line so that a slight misalignment shows more artwork rather than a white sliver.

Practical figures worth carrying: for most print jobs, 3mm of bleed beyond the trim and keep anything important at least 5mm inside it. Scribus and Inkscape both let you set bleed on export to PDF, free, as do InDesign and Illustrator.

On screen the equivalent is the safe area. A vertical video at 1080×1920 loses the top and bottom to the platform's own interface — the account name, the caption, the buttons down the right side — and the exact amounts change whenever the platform updates. Keep essential type in the middle 60% vertically and away from the right edge, and check on a real phone rather than in the editor.

Active space and the thing clients will actually accept

Space stops feeling like waste when it has a shape. A block of text set against a wide, deliberate column of white on one side reads as composed; the same text centred with even gaps everywhere reads as under-filled. The difference is that the first is asymmetric and therefore obviously chosen.

When asked to fill the space, the useful reply is not a defence of minimalism. It is a comparison: show the version with the gap and the version with something in it, printed at actual size, and let them look from two metres away. The argument is usually over in ten seconds, because at two metres the filled version has lost its headline and the empty one has not.

Where the principle runs out

Generous space is a luxury paid for in square centimetres, and some jobs genuinely cannot afford it. A classified advertisement, a product label carrying legally required text, a ballot paper, a chart with forty labelled points: in all of these the information is fixed and the space is not negotiable.

The move then is not to give up on space but to switch scale — small, consistent, ruthlessly applied gaps, with the hierarchy carried by weight and rule lines rather than by air. Dense typography done properly is a genuine craft, and it is the reason a well-set timetable is calmer than a half-empty poster with four competing headlines.

BEFORE YOU MOVE ON

In a form, each label sits exactly midway between its own field and the field above it. Users are slower and make more pairing errors. Why?

- A. The labels are too small relative to the fields, so they read as secondary information and get skipped.
 - B. Forms should place labels inside the fields as placeholder text, which removes the pairing problem entirely.
 - C. The form lacks visual separators between rows, so the eye has no structural cue about where one entry ends and the next begins.
 - D. Equal gaps give the grouping system nothing to act on, so the pairing has to be worked out from meaning instead of being seen.
-

15 · 8 min

Four cheap tests that fail a design before an audience does

THE ONE THING TO KEEP

Degrade your own vision on purpose — blur, greyscale, flip, and view at true size and distance — because after two hours you read your layout from memory, and these four tests plus a three-second stranger restore something like a first glance.


You cannot see your own work after two hours

This is not a character flaw. You have spent two hours learning where everything is, so your eye goes straight to each element because it already knows the route. A viewer has none of that, and the gap between your reading and theirs widens with every minute you spend.

The four checks below are all ways of degrading your own vision deliberately so that you see something closer to a first glance. Each takes under a minute

and costs nothing.

Four checks that degrade your own vision on purpose



The first four cost under a minute each and test transmission only. The fifth needs another person and is the only one that tests meaning.

1. The blur

Apply a Gaussian blur of roughly 1–2% of the design's longest edge — about 10–20px on a 1000px canvas. Free everywhere: GIMP (Filters → Blur → Gaussian), Krita, Photopea, or simply stand back and squint, which is the same test with worse repeatability.

What it tells you: what survives peripheral vision, which is what decides where the eye goes first. If two shapes survive with equal strength, your entry point is unresolved. If nothing survives, the design has no first level at all.

What it cannot tell you: anything about whether the surviving message is the right one. A blur validates rank, never content.

2. Greyscale

One menu item in every image editor, and a filter in browser developer tools.

What it tells you: whether your hierarchy is built on tone or on hue alone. Hierarchies built on hue collapse here, and they also collapse in photocopies, in cheap newsprint, and for viewers with colour vision deficiency.

3. The flip

Mirror the design horizontally. In Inkscape it is the H key; in GIMP, Image → Transform → Flip Horizontally.

What it tells you: balance and awkward alignment, freed from your familiarity. Because you cannot read the text any more, you see the composition as

shape. A layout that looked settled often turns out to be leaning, and edges you thought were aligned are visibly not. Painters have used a mirror for this for centuries, for exactly this reason.

Note: do not conclude that the mirrored version is better and ship it. Reading direction matters, and a flipped layout reverses the entry point.

4. Actual size, actual distance

The one people skip, and the one that catches the most expensive mistakes. Print the poster at true size and pin it where it will hang, then walk to where people will stand. Put the reel on a phone and hold it at normal viewing distance. Open the email on the oldest device in the house.

A useful rule of thumb from signage practice: legibility requires roughly 25mm of cap height for every 10 metres of viewing distance, with generous contrast. A shop sign read from across a road needs letters far larger than anyone estimates on a laptop.

What it tells you: whether the smallest necessary element is actually legible in the real viewing condition, which no amount of zooming on a monitor will reveal.

The fifth check, which needs another person

Show it to somebody for three seconds — count them out loud — then take it away and ask two questions: *what was it about* and *what did it want you to do*. Do not explain first. Do not defend. Write down what they say verbatim.

This is the only check that tests meaning rather than form, and it is uncomfortable in proportion to how useful it is. Ask two or three people, because one person's confusion may be idiosyncratic and two people's confusion almost never is.

What the checks cannot do

All five are tests of transmission. They tell you whether the message arrives, and they say nothing about whether the message is worth sending or whether

the viewer will believe it. A piece can pass every check and still fail commercially, because the claim was not credible or the offer was not wanted.

There is also a trap in over-testing. A design optimised solely to survive a blur tends towards one enormous word on a flat ground — maximally legible, frequently empty. The checks are a floor to clear, not a score to maximise, and the right response to passing all of them is not satisfaction but the question the checks cannot answer: does this say the thing the brief asked for?

BEFORE YOU MOVE ON

A designer's poster passes the blur test with one dominant shape, passes greyscale, and is legible at true distance. Three strangers still cannot say what it is offering. What does this combination indicate?

- A. Transmission works and meaning does not, so the fault is in what the dominant element says.
- B. The blur test was applied at too weak a radius, so the apparent dominance was an artefact of insufficient degradation.
- C. The strangers were not the target audience, so their confusion says little about how the real viewer will read it.
- D. The design needs a stronger call to action placed near the dominant element so the offer is unmissable.

16 · 8 min

The bit you are proudest of is usually the bit in the way

THE ONE THING TO KEEP

The element you invested most in competes hardest for the first glance, so judge it against the written message rather than against its cost — and demote it to level three before you consider deleting it.

Why it is always that element

There is a reliable pattern in this work. The element you spent longest on — the clever type treatment, the illustration you drew by hand, the transition that took an afternoon — is disproportionately likely to be the one damaging the piece. Two mechanisms drive it, and neither is about talent.

The first is effort. You do not assess the element on its merits; you assess it partly on what it cost you, which is a fact about your afternoon and not about the design. Economists call this the sunk cost fallacy and it operates just as strongly on hours as on money.

The second is that ambitious elements are, by construction, high in contrast and complexity, which is precisely what makes them compete for the first glance. A flourish that is merely pleasant sits quietly at level three. A flourish that is genuinely impressive fights the headline for level one — and wins, because you made it well.

The two-question test

When you suspect a favourite is in the way:

1. **Does it carry the message, or decorate it?** A type treatment that makes the key word mean more is carrying. A type treatment that makes the key word look expensive is decorating. Decoration is not forbidden —

it is what makes work feel like it came from somewhere — but decoration at level one is a hierarchy error.

2. **Would you have added it if somebody else had built the rest?**

This strips out the effort. If you look at a stranger's layout and would not add your flourish to it, it is not solving a problem in the layout.

Do it reversibly, which removes the fear

The resistance to cutting is mostly fear of loss, and version control removes the fear entirely.

- Save the current file with a clear name — `poster-v4-with-flourish.xcf` — before you touch anything.
- Hide the layer rather than deleting it. One click in GIMP, Krita, Inkscape, Scribus, Figma or Photoshop, instantly reversible.
- Keep a `kill/` folder and move cut pieces into it. Nothing is destroyed, and parts cut from one job are frequently the seed of the next.

Then look at the version without it, ideally after a gap. Not five seconds later, when your memory is still supplying the missing element — the next morning, or after lunch. If you cannot tell what is missing, it was not doing anything.

What to do with the good idea you cut

A cut element is rarely bad. It is usually good and misplaced. Three destinations, in order of usefulness:

- **Demote it.** The elaborate treatment moves from the headline to a small detail, where it adds texture without competing. This saves the idea most often.
- **Move it to another piece in the series.** If the campaign has a poster, a reel and a flyer, an idea that overloads the poster may be exactly right as the whole of the reel.
- **Bank it.** Keep the folder. Ideas have a long shelf life and a genuinely good one will fit a brief within the year.

The counterweight: do not cut everything distinctive

This lesson is about elements that compete with the message. It is not a licence to sand a piece down until nothing is left but a centred sans-serif headline on white, which is the failure mode of designers who have learned restraint and nothing else. Work with no personality communicates perfectly and is forgotten immediately, and being forgotten is also a failure — often the more expensive one, since it produces no complaints and therefore no feedback.

The distinction is positional, not moral. A distinctive element at level three gives the work a voice. The same element at level one takes the message's place. Most of the time the correct move is not deletion but demotion.

The habit that makes this bearable

Decide what the piece is saying *before* you make the thing you will fall in love with. If the one-sentence message is written down and agreed in advance, cutting an element becomes an ordinary comparison against a stated standard rather than a judgement on your taste. The brief is what makes the decision impersonal, and impersonal decisions are the only ones you can make about your own work with any reliability.

BEFORE YOU MOVE ON

A designer cuts an elaborate hand-lettered headline treatment and the poster immediately reads better. Why is demotion usually a better first move than deletion?

- A. Deletion wastes the hours already invested, so demotion recovers the value of the work that was done.
- B. Clients notice when elaborate work disappears, so keeping some of it visible avoids an argument about the fee.
- C. The problem is rank, not quality: at level three it keeps the voice and frees level one.
- D. Every design needs one decorative element to avoid looking generic, so something must occupy that role.

CASE STUDY · MODULE 2

Six departments, one A3, and the surgeon's quotation

A charitable eye hospital in Madurai, producing a flyer for free cataract screening camps in eleven villages.

The flyer had been assembled the way flyers usually are, which is one reasonable request at a time. The hospital logo. The trust's logo. The funding partner's logo, because the grant agreement required acknowledgement. A photograph of the senior surgeon. A quotation from him about the right to sight. The dates. The eleven village names. A QR code. Two phone numbers, because one was the ambulance line. The word free, in red, because somebody in the meeting said red meant important.

Every item had a person behind it who could explain why it mattered. Nobody had looked at the sheet as a whole in a month.

Forty seconds of blurring

The communications assistant flattened a copy in a free browser editor and applied a twenty-pixel blur. Nothing survived. The sheet came back as an even field with three darker rectangles in it, and the darkest was the photograph of the surgeon's face, which at that blur was the only identifiable object on the page.

That is the diagnosis stated precisely. The eye needs one entry point and finds it by looking for the strongest local contrast. There were four candidates at the same weight, so scan order was arbitrary, and an arbitrary scan is uncomfortable enough that people resolve it by leaving.

The decision, and the cost on each side

The fix was not deletion but ranking. In the proposed version, the two things that survived a blur were the dates and the words free eye check, taking the top third. The eleven village names went to level three at nine point in a single column. The three logos went to a row along the bottom at the size the grant

agreement actually specified, which nobody had read and which was smaller than what had been used. And the surgeon's photograph and quotation came off the flyer entirely.

The surgeon was also the hospital's largest individual donor and the reason the camp had a theatre list at all.

So there were two real costs. Demoting him risked a relationship worth considerably more than one camp, and the administrator would be the person in that conversation. Keeping him risked eleven villages of people who could not find the date, which was the only fact the flyer existed to transmit.

The administrator did two things instead of choosing in the abstract. He had twenty copies of each version printed, and he took them to the outpatient corridor and ran a three-second test on staff and waiting relatives: look at this, then tell me what it was about and what it wanted you to do.

He also prepared a third option, which is the move that usually saves a cut element. The surgeon's photograph and quotation were not thrown away; they were moved to a separate piece, an A5 card handed to every person screened, where the quotation is read by somebody who has already decided to come and has time.

Why the sheet could not simply be made louder

The review meeting had produced the usual instruction, which was to make the dates bigger. That was not going to work on its own, and the reason is the arithmetic of contrast. Emphasis is comparative: there is a fixed amount of perceivable difference on one surface, and every element that spends some of it leaves less for the rest. The dates had already been enlarged twice. What they needed was not more size but fewer competitors, which meant the other items coming down.

There was a second cost nobody had costed. The sheet had grown to eleven separate objects. Three elements carry three relationships to manage; eleven carry fifty-five, and each of them has to be positioned and weighted against all the others. Nobody had made a bad decision. The decisions had simply been made one at a time, by different people, over a month.

Grouping bought some of it back. The eleven village names, set in one column at one size with one alignment, are perceived as a single object rather than eleven, so long as the space around the block is clearly larger than the space inside it.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

On the ranked version, fourteen of twenty people named the date and eleven said free eye check. On the committee version, three of twenty named the date, and six said something about a doctor.

The administrator showed the surgeon both sets of results before proposing anything, which changed the conversation from a matter of standing into a matter of arithmetic. The surgeon asked one question, which was whether the card handed at the camp would carry the quotation in full rather than cut, and then agreed.

The red was replaced by a heavier weight of the same ink, because the flyer was being printed in one colour on newsprint at two of the three print runs and the red had been separating to a mid-grey that was slightly lighter than the body text.

Attendance across the eleven villages rose from about 640 the previous year to 1,180. An auto-rickshaw announcement also ran for the first time that year, so the flyer cannot be separated from it, and the report said so.

Four elements were competing for the first glance. Why is that perceptually the same as having no first-level element at all?

A level is a relationship, not a size. The eye selects an entry point by finding the strongest local contrast, so when four elements carry the same weight there is no reason to prefer any of them and the scan order becomes arbitrary. Four level-one

elements and zero level-one elements produce the same experience for the viewer, which is why the blur test reports both as a failure.

The administrator described the surgeon's photograph as a political cost rather than a design requirement. Why does naming it that way help?

An element is worth keeping if it does a job nothing else does, and if the honest answer to what happens when it is absent is that a particular person will be annoyed, that is a real cost but it belongs to a different ledger. Dressing it as a design requirement means it gets defended with design arguments that are not true, and the layout is then built around a claim nobody can examine. Named as politics, it can be solved politically, which is what the separate card did.

Why did printing twenty of each version settle an argument that discussion had not?

The claim in dispute was about what a stranger takes from the sheet in three seconds, and that is measurable on strangers and not knowable by anyone who has already read it. Twenty sheets and a corridor converted a contest between two opinions, in which the surgeon's carried more weight and reasonably so, into a comparison against a standard both parties had agreed mattered.

MODULE 2 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 You blur a design at two per cent of its longest edge and three shapes survive with equal strength. What have you found?
 - A. A balanced composition that needs no further structural work.
 - B. Three level-one elements, which for a viewer is the same as none.
 - C. A blur radius set too large for a canvas of that width.
 - D. A palette problem, since competing hues survive a blur more strongly than tone.

- 2 A headline is set big, bold, red, centred and boxed. Three levels of hierarchy were planned and only two are visible. What happened?
 - A. Each lever spent contrast the lower levels needed, so they stopped being distinguishable.
 - B. Five levers at once is beyond what a viewer is able to perceive within a single half-second glance.
 - C. The levers interfere with one another and begin to cancel out above three.
 - D. Red is a saturated hue and saturated hues flatten the tonal range of a page.

- 3 A status panel shows working in green and broken in red, at the same lightness. Which description names the mechanism of its failure?
 - A. Saturated colours always measure below the contrast floor for body text.
 - B. Red and green are culturally loaded and imply loss and gain to many readers.
 - C. Presses cannot reproduce red and green accurately on uncoated stock.
 - D. The information is carried in hue alone, so it disappears in tone.

- 4 You are told that people read screens in an F shape and should place important elements along the top and left. What is wrong with using that as a template?
- A. The finding was measured on printed pages and does not transfer to screens.
 - B. It applies only to right-to-left scripts, where the pattern is mirrored.
 - C. It describes reading where there is no hierarchy, and contrast overrides habit.
 - D. The result has failed to replicate and is now regarded as a statistical artefact.
- 5 A layout is not working and your reflex is to move things around. Why try hiding elements one at a time first?
- A. It is faster and more diagnostic, and a layer toggle makes it fully reversible.
 - B. Clients approve fewer elements more readily, so removal is easier to get signed off.
 - C. Moving an element breaks its alignment and introduces errors elsewhere on the page.
 - D. Rearranging requires a grid, and most layouts have not had one set up properly.
- 6 A caption sits 20px below its photograph and 20px above the next heading. What is the failure, and what rule repairs it?
- A. Both gaps are too tight; captions need at least 40px of separation on either side.
 - B. The caption is on the wrong side; captions belong above the image they describe.
 - C. The gaps are not multiples of eight, so the spacing reads as arbitrary rather than as something chosen.
 - D. Equal gaps say it belongs to neither, so space inside a group must be smaller than around it.

MODULE 3

The shape of a story

A sequence of true statements is not a story, and most brand films are exactly that. This block covers the minimum unit — a change of state — and then the structure built on it: how a ninety-second piece actually divides, why one concrete detail buys more belief than ten adjectives, why a piece with nothing at stake cannot hold anyone, who the story is allowed to be about, and how to end without either trailing off or stacking four requests.

BY THE END OF THIS MODULE YOU CAN

Build a short piece as a chain of state changes joined by but and therefore, allocate its seconds to a stated structure, name the protagonist and the cost they carry, replace generic claims with checkable detail, and land one ask the piece has earned

17 · 9 min

"And then" is the weakest word you own

THE ONE THING TO KEEP

Between any two beats, if the honest joining word is "and then", the second beat is deletable — rewrite until it is "but" or "therefore".

Why a list of true facts is boring

Here is a real thirty-second script from a plumbing business: "We do pipe repairs. And then we also do drain cleaning. And then we service boilers. And

then we have twelve years of experience. Call us."

Every sentence is true. Nothing is wrong with it. It is unwatchable, and the reason is structural, not stylistic. Each beat is joined to the last with *and then*, which means each beat is optional. Delete any one and nothing breaks. A viewer feels that — they cannot say why, but they can feel that leaving early costs them nothing, so they leave.

The writers of South Park made this rule famous and it is the most portable piece of story craft there is: between any two beats you should be able to say **but** or **therefore**, and almost never *and then*.

Rewritten: "Your kitchen floods at eleven at night. **But** every plumber's phone goes to voicemail. **Therefore** we answer at eleven at night." Same business, same facts, and now beat two makes beat three necessary. You cannot delete the middle.

The same plumbing business, two chains

Joined by **and then**

- We do pipe repairs
- And then we do drain cleaning
- And then we service boilers
- And then we have twelve years
- Any beat can be deleted

Joined by **but** and **therefore**

- Your kitchen floods at eleven
- BUT every phone goes to voicemail
- THEREFORE we answer at eleven
- Same facts, fewer of them
- The middle cannot be cut

A viewer cannot name the joining word and can feel it. When every beat is optional, leaving early costs nothing, so they leave.

Setup, tension, turn

The smallest usable shape has three parts, and it fits a thirty-second reel and a four-panel brochure equally.

- **Setup** — a situation the audience recognises as theirs. Not your company history. Their Tuesday.
- **Tension** — something is wrong, missing, or in the way. This is the beat beginners skip, because it feels negative and clients dislike it. It is the only beat doing work.

- **Turn** — something changes. Not necessarily a solution. A reframe, a reveal, a consequence.

A brochure: the cover carries the tension, the inside spread carries the turn, the back carries the practical detail nobody reads first. A case study: the situation, the complication, the decision you made and why, what changed. A single poster gets a compressed version — the tension in the image, the turn in the line.

Tension you invent does not count

The failure mode of this lesson is manufactured drama. A reel that opens with "Most people are doing this **WRONG**" about a thing nobody was worried about, or a nonprofit brochure that invents a crisis to justify itself. Audiences are extremely good at detecting stakes that are not real, and the detection is instant.

Real tension is something the audience already feels before you say it. Your job is to name it accurately, in their words, not to generate it. If you cannot find one, that is information: possibly the product genuinely solves a problem nobody has, and that is a business finding worth reporting rather than a writing problem to paper over.

Not everything is a story

A price list is not a story. A safety notice, a train timetable, an ingredients panel and a set of assembly instructions are not stories, and dramatising them makes them worse — slower to scan, and slightly untrustworthy, because the reader senses the persuasion and starts to wonder what is being hidden.

Use this test: does the audience arrive already wanting the information? If yes, get out of the way and organise it. Structure is for the audience who has not yet decided to care.

There is also a limit to the three-beat shape itself. It is a starting scaffold, not a law of nature; documentary, observational and comedic work often runs on rhythm, contrast or accumulation instead. Learn the scaffold first because it

fixes the specific failure most beginners have, then break it deliberately rather than by accident.

Where the beats go in time

For a thirty-second reel, roughly: setup in the first three to five seconds, tension by about eight, turn somewhere in the middle, and the last three seconds carrying whatever you want remembered. That last part matters more than people think — the final frame is what a viewer holds while deciding whether to follow you, and most beginner reels end on a hard cut mid-sentence.

Write the three beats as three plain sentences before you write any script, choose any shot, or open an editor. If the three sentences are not interesting when read aloud in a flat voice, no amount of cutting, music or colour grading will rescue them. Music is very good at making a bad structure feel briefly acceptable in the edit and then failing on someone else's phone with the sound off.

Do this today

Take any piece of your own work with more than one part — a reel, a three-panel leaflet, a portfolio page. Write out its beats as sentences. Put the actual joining word between them. Count how many are *and then*. Then rewrite so that at least one becomes *but* and one becomes *therefore*, even if it means cutting a fact you liked. The version with fewer facts and better joints performs better almost every time, and you can test that claim yourself within a week.

BEFORE YOU MOVE ON

A charity's thirty-second film lists four programmes it runs, each introduced with "we also". Viewers drop off steadily throughout. Which change addresses the actual cause?

- A. Add music and faster cuts so the pace carries the viewer through the list of programmes.
 - B. Restructure so one programme answers a problem the previous beat established, making the later beats consequences rather than additions.
 - C. Move the strongest programme to the opening so the most compelling material is seen before the drop-off.
 - D. Cut the film to fifteen seconds so the list ends before the typical drop-off point.
-

18 · 9 min

The first second is a promise, and you have to pay it

THE ONE THING TO KEEP

The opening has one job — answer "what is this and is it for me" — and it is only craft rather than bait if the rest of the piece pays what it promised.

What the opening actually has to do

The usual advice is "grab attention", which is useless because it names a goal rather than a job. The job is narrower. In the first second — or the first glance at a static piece — a viewer answers two questions almost without thinking: *what is this*, and *is it for me*. If either answer is missing or ambiguous, they leave, and on a feed leaving costs them nothing.

So the opening is not a trick. It is an answer delivered early. Everything else people call a hook is a way of delivering that answer faster.

On short-form video platforms the drop-off is steepest in the first few seconds — a large share of the audience you lose, you lose before the sentence you were saying finishes. Rather than chase a specific number that changes with every platform update, treat the shape as the fact: attention is not spent evenly, it is spent almost entirely at the front.

Concrete openings that work

- **Open on the end state.** Show the finished cabinet, then go back to the plank. The viewer now knows what they are being paid for staying.
- **Open mid-action.** Not "hello everyone, today we are going to" — the hands already moving, the pan already hot.
- **Put the question on screen as text.** Most feed video autoplays muted. If your opening line is spoken and not written, for a large fraction of viewers it did not happen.
- **Make the first frame deliberate.** It is a thumbnail whether you chose it or not. A frame of your face mid-blink is a choice you made by not making it.
- **For static work,** the same rule applies to the largest element. It has to answer "what is this" without the reader taking a step closer.

Captions are the cheapest single improvement available. CapCut generates them automatically and is free on phone and desktop. If you would rather stay in open-source tools, DaVinci Resolve's free tier has automatic subtitle generation, and Kdenlive and Shotcut both accept an SRT file. Whatever you use, check the words — automatic captions are consistently wrong on names, prices and non-English words, which are exactly the words that matter in a local business video.

The ethical edge, stated plainly

Engineering attention is a real skill and pretending otherwise is dishonest — every decision about order and emphasis is attention engineering, including the ones in this course. The line is not between influencing and not influencing. It is narrower and more useful: **does the piece pay what the opening promised?**

A hook that promises a bamboo house and delivers a bamboo house is craft. A hook that promises a bamboo house and delivers eight minutes of a man talking about his channel is a withdrawal from an account you cannot refill. The mechanism is worth knowing because it is not moral disapproval — it is arithmetic. Bait produces a spike in the first metric and a decline in the second: high clicks, short watch time, unfollows. Platforms have moved steadily toward weighting the second, and audiences remember the last three seconds long after they have forgotten the first.

The opening, and whether the piece pays it

	The piece pays it	The piece does not
Strong opening	Craft. Clicks and watch time spike.	Bait. Clicks spike, watch time falls, unfollows.
Weak opening	A good piece almost drives.	Nothing drives, so nothing is measured.

The line is not between influencing and not influencing. Every choice about order and emphasis is attention engineering. The line is whether the piece pays what the opening promised.

Three places the edge is clearly crossed, whatever the metrics say:

- **A promise the piece does not keep** — an arrow pointing at something that is not in the video, a shocked face at a non-event, "you won't believe" resolving into nothing.
- **Manufactured fear or urgency** — a countdown that does not exist, a health scare, a fake deadline.
- **Aiming this at people who cannot easily disengage** — children, someone in distress, someone at two in the morning. The techniques work best on exactly the people who are least able to refuse them, and knowing that is part of knowing the craft.

If your opening uses generated footage of something that looks like a real event, label it visibly, in the frame, at the start. Several countries now require exactly that for synthetic media, India's IT Rules and China's labelling rules among them, and the rules are still moving; making things with AI covers

where that stands. Beyond the law, a viewer who discovers the opening was fabricated does not carefully separate that from everything else you make.

The check that keeps you honest

Write your opening as a sentence, and then write the sentence the viewer would say afterwards if you delivered. If those two sentences do not match, you have a bait problem, and no amount of retention optimisation is going to fix a broken promise — it will only distribute it more widely.

Do this today

Take your last video and watch only the first two seconds, muted, on a phone, at arm's length, as if it appeared between two other videos. Then ask someone else to do the same and tell you what they think it is about. If they cannot say, the problem is not the algorithm. Fix the first two seconds and republish; it is the cheapest edit in this entire course.

BEFORE YOU MOVE ON

A creator's videos get high click-through but very short average watch time, and follower growth has stalled despite rising views. What is the most likely mechanism?

- A. The openings promise something the videos do not deliver, so clicks rise while watch time and trust fall.
- B. The videos are too long for the platform, so viewers leave before the payoff arrives.
- C. The thumbnails are stronger than the production quality of the videos, so the fix is investing in better camera and lighting.
- D. The platform is throttling the channel, so the fix is posting more frequently to recover reach.

19 · 8 min

If nothing changes, it is not a story

THE ONE THING TO KEEP

A story is a state, an event, and a different state; if the start and end states match you have a description, which must earn attention through usefulness or beauty rather than suspense.

The minimum story is two states and a reason

Here is a true, well-shot, competently edited ninety-second film: the workshop in the morning light, hands at the lathe, a slow pan across finished pieces, the owner saying they have always cared about quality, the logo. Everybody involved is proud of it. Nobody who watches it can tell you anything about it an hour later.

It has no change in it. At the end, everything is exactly as it was at the start — the workshop existed, the owner cared, the pieces were good. That is a description, and descriptions are perfectly legitimate objects. They are just not stories, and they do not behave like stories in a viewer's attention.

The minimum unit of a story is: **a state, a thing that happens, a different state**. Three parts. Remove any one and it stops working.

*The shop had made the same sign for forty years. **Then** the road was widened and the shopfront lost two metres. **So** everything they knew about sign-making at that size became useless in a week.*

That is a story in three sentences, and you already want to know what they did.

The test: name the two states

Before you script anything, write two lines:

```
state at the start: ...
state at the end:   ...
```

If you cannot fill them in with genuinely different content, you have a description. That is not automatically fatal — a product demonstration, a menu, a price list and a location film are all descriptions and all useful — but you should know which you are making, because a description has to earn attention a different way: through usefulness, beauty or brevity, never through suspense.

When the two states are the same and the piece is supposed to be a story, the usual cause is that the change happened to somebody who was left out. The customer's state changed. The film was about the workshop.

Change can be small

People hear "story" and reach for drama — the near-bankruptcy, the flood, the family tragedy. Most work does not have access to that and does not need it. A change of understanding is a change:

- She thought the delivery would take a week. It took four hours. Now she orders on Fridays.
- He assumed the expensive blade was a con. He counted the re-sharpenings. He was wrong by a factor of three.
- The team believed the form was fine because nobody complained. Then they watched somebody fill it in.

Each of these is small and each is a genuine change of state, which is why they hold. The size of the change sets the emotional scale; the presence of the change decides whether there is anything there at all.

Chain the units

A longer piece is several of these joined end to end, where each new state creates the next problem. This is where **but** and **therefore** do their work: they are the joints that make one state cause the next. A piece that can only be joined with "and then" is a list of events, and a list can be reordered without

loss, which is the diagnostic — if you can swap beats two and three and nothing breaks, they are not causally linked and one of them is probably deletable.

The counting exercise

Take any piece you are working on and write one line per beat. Then, beside each, write the state it leaves the viewer in. If two adjacent beats leave the same state, the second one is doing nothing and you have found your cut. This routinely removes a quarter of a rough edit, and it removes it on evidence rather than on feel.

Where this does not apply

Not every piece should be a story, and forcing narrative onto material that does not have it produces the worst genre in the trade: the manufactured journey, where a company that simply started trading in 2019 is given a hero's arc and a moment of doubt that never happened. Viewers detect this reliably, because invented adversity has no specific detail in it — the struggle is always described in general terms while the success is described in concrete ones, and that asymmetry is visible.

If there is no change, say the true thing plainly and make it useful or beautiful instead. A clear, honest description of what you do, well made, outperforms a fictional arc every time — and the fictional arc carries a risk the description does not, which is that somebody who was there remembers it differently.

BEFORE YOU MOVE ON

A ninety-second film shows a workshop, the owner explaining their commitment to quality, and finished products. It is well shot and nobody remembers it. What is structurally missing?

- A. Emotional music and tighter pacing, which would give the material the momentum it currently lacks.
 - B. A clear call to action at the end, without which viewers have no reason to retain what they saw.
 - C. Nothing in it changes state, so there is a description where the viewer expected an event.
 - D. The owner is not a compelling on-screen presence, so the material needs a professional narrator instead.
-

20 • 9 min

Where the ninety seconds actually go

THE ONE THING TO KEEP

Give the turn the largest share of the running time and the setup the smallest, check the script against 140–160 words per minute before recording, and test the structure as seven plain sentences while changes are still free.

Most short films spend half their length on setup

Count it on any amateur brand film. Forty seconds of establishing: the exterior, the sign, the door opening, somebody walking in, the owner sitting down and being introduced. The actual subject arrives at 0:45 and there are fifteen seconds left for it.

This happens because the maker experiences the setup as necessary context and the viewer experiences it as a delay. Both are right about their own experience. The viewer wins, because they have the option of leaving.

A workable division for ninety seconds

This is a starting point to break deliberately, not a template.


- **0:00–0:03 — the promise.** What this is and why it concerns me. Usually the most surprising true statement you have, or the strongest image. Not the logo, not the establishing shot.
- **0:03–0:20 — the situation and the problem.** Who, and what is wrong. Seventeen seconds is enough for both if you stop describing and start stating.
- **0:20–0:60 — the turn.** Something is attempted; it goes wrong or it costs something; a second attempt works. This is the majority of the piece and it is where amateur films are thinnest.
- **0:60–0:80 — the new state.** What is different now, shown rather than claimed.
- **0:80–0:90 — one ask.** One. Named, specific, achievable in the next ten seconds.

The useful discipline here is that the turn gets the largest share. If your turn is eight seconds and your setup is forty, you have the proportions of a film that will be watched by your client and nobody else.

Two numbers that make scripting concrete

Speech runs at roughly 140 to 160 words per minute in natural, unhurried delivery. So ninety seconds is about 210 to 240 words of voiceover — roughly one side of A4 at normal type size. Write your script, count the words, and you know whether it fits before you record anything. Almost every over-long edit was over-long in the document.

Where ninety seconds go, and where amateurs put them



The turn takes the largest share and is where amateur films are thinnest. The last bar is the common failure: forty-five seconds of establishing, and the subject arriving with fifteen seconds left.

Shot length: a comfortable average for this kind of piece is two to four seconds, which gives 25 to 40 shots in ninety seconds. If your shot list has twelve items, expect a slow film; if it has ninety, expect one nobody can follow. This is a rough guide rather than a rule — a single unbroken shot can carry a whole piece, and does, in the hands of somebody who planned it that way.

The same shape compresses and stretches

At **fifteen seconds** you get roughly: two seconds of promise, five of problem, five of turn, three of new state and ask combined. There is no room for a separate ask, so the ask has to be inside the new state — the phone number on screen while the result is shown.

At **six minutes** the shape does not change; the middle subdivides. The turn becomes three or four linked state changes, each with its own small reversal. What does not happen is that the setup grows proportionally. A six-minute film can still put its problem on screen in the first twenty seconds, and the good ones do.

Write the beats before the shots

One line per beat, in plain sentences, before any camera or software is involved:

- 1 Sign was made the same way for 40 years.
- 2 BUT the road widening cut the shopfront by two metres.
- 3 THEREFORE the old layout no longer fits and trade drops.
- 4 They try shrinking the same design. It looks cheap.
- 5 BUT the apprentice suggests splitting it across two panels.
- 6 THEREFORE the sign now works and reads from further away.
- 7 ASK: bring us your awkward shopfront.

Seven lines. You can rearrange, delete and test this in two minutes; you cannot do that with a rough cut. Every hour spent here saves several in the edit, and the beats document is also the thing you show the client before shooting, when a change is still free.

The honest limitation

This structure suits persuasion and explanation, which is most commercial work. It is not universal. Observational documentary, a mood piece, a piece of music promotion or an experimental short may deliberately withhold structure, and imposing a turn at 0:20 on material that is about stillness ruins it.

And a well-proportioned film about nothing is still about nothing. Structure allocates attention; it does not manufacture a reason for the attention to be given. If the beats document reads as boring in seven plain sentences, no amount of grading, music or camera movement will fix it — and knowing that on day one rather than in the edit is the entire point of writing it down first.

BEFORE YOU MOVE ON

A team's ninety-second film has forty seconds of establishing shots and eight seconds on the problem being solved. Why is re-editing to a faster pace unlikely to fix it?

- A. Faster cutting looks amateurish in corporate work, so the change would introduce a new problem while solving the old one.
 - B. The proportions are wrong rather than the tempo, so trimming the setup evenly leaves the turn just as thin as before.
 - C. Viewers leave in the first three seconds, so anything after that point is irrelevant to retention.
 - D. The film needs additional footage, and no amount of editing can add material that was never shot in the first place.
-

21 • 8 min

One checkable detail outperforms ten adjectives

THE ONE THING TO KEEP

A specific detail is believed because it could be contradicted, so replace adjectives with checkable nouns and numbers — but only true ones, because precision makes a false claim worse rather than merely wrong.

Why a detail is believed and a claim is not

Compare two sentences about the same business.

We are committed to quality and customer satisfaction.

We re-cut the sign at eleven at night because the spacing on the second word was two millimetres out and it would have annoyed us every time we drove past.

The second is believed and the first is not, and the reason is not that it is more eloquent. It is that the second sentence **could be false in a checkable way**. Two millimetres, second word, eleven at night: each of those is a fact somebody could contradict. A claim about commitment cannot be contradicted, so it carries no information and the reader's discount rate on it is total.

This is the most reliable mechanism in the whole subject. Specificity implies a witness. Generality implies a copywriter.

Detail also does the work of description

A second effect, separate from belief: concrete detail builds a picture, and a picture is retained where an abstraction is not. "A small salon" occupies no space in memory. "Two chairs, a mirror with a chipped corner, and a radio that only gets one station" is a place you can see. You did not spend more words on the adjectives; you spent them on nouns.

The practical instruction is almost mechanical: **cut the adjective, add the noun**. Not "high-quality ingredients" but "the coriander is cut at four in the morning". Not "experienced team" but "between them, ninety-one years of it".

Where to find the details

They are never in the client's prepared answers, because prepared answers are already generalised. They come from:

- **Asking about a specific occasion, not the general case.** Not "how do you ensure quality" but "tell me about the last time something went wrong". Specific questions produce specific answers; that is the whole trick.
- **The numbers nobody thought were interesting.** How many, how long, how far, how often. "We deliver quickly" becomes "the average is 4 hours 20 minutes, and we know because we log every one".
- **Watching rather than asking.** The gesture, the tool with tape on the handle, the handwritten list on the wall. People do not report their own habits accurately, but they perform them in front of you.

- **What it cost.** The sacrifice, the mistake, the thing they had to stop doing. This overlaps with stakes and is dealt with properly in the next lesson.

In pictures, the same rule

The generic photograph — smiling people around a laptop, a handshake, an abstract cityscape — is the visual equivalent of "committed to quality". It could illustrate any business in any country, which means it identifies none. A badly lit photograph of the actual workshop, with the actual mess in it, communicates more because it could only be that place.

This is also why stock imagery used carelessly damages credibility rather than being neutral. Viewers increasingly recognise stock, and recognising it tells them the specific thing was not available or not thought worth showing.

The two real limits

Specificity narrows. "Two millimetres of spacing on a shopfront sign" speaks precisely to people who care about signage and may exclude a broader audience. That is usually an acceptable trade, because a narrow group who believe you are worth more than a wide group who do not — but it is a trade and you should make it knowingly. The escape is to pair one specific with one general: the detail earns the belief, the general statement makes it applicable.

Detail must be true. This is not a stylistic technique to be applied to invented facts. If the average delivery time is not 4 hours 20 minutes, do not write it — not for ethical reasons alone, but because the specificity that makes a true detail persuasive makes a false one catastrophic. A vague claim that fails disappoints. A precise claim that fails is a lie with a number attached, and it is the kind of thing customers photograph and post.

When the client has no interesting details, that is the finding. Either nobody has asked them the right questions, which you can fix in ten minutes, or the business genuinely does what everybody else does in the same way — in which case the honest job is to find the one thing that is true and say it plainly, rather than dressing an ordinary business in a borrowed story.

BEFORE YOU MOVE ON

Why does "the average delivery is 4 hours 20 minutes, and we log every one" persuade more than "we deliver quickly"?

- A. Numbers feel objective, and audiences trust quantitative claims more than qualitative ones regardless of their source.
 - B. It is falsifiable, so making it implies somebody could check it and the speaker accepted that risk.
 - C. It is more memorable, and memorable claims are more likely to be acted on later when the need arises.
 - D. It contains more information, so the reader can make a better-informed comparison against competitors.
-

22 · 8 min

Nothing is at stake, so nobody is watching

THE ONE THING TO KEEP

Attention attaches to unresolved questions, so name what was risked, given up or went wrong — a real admission makes the surrounding claims believable, while invented adversity is detectable by its vagueness against a concrete resolution.

Attention follows uncertainty

A piece in which everything goes well is a piece with no open question, and an attention system that allocates on the basis of unresolved questions has nothing to attach to. This is why the testimonial that says "it was great, I'd recommend them" is inert while "I nearly cancelled in the second week" makes you lean in. The second one has something unresolved in it.

Stakes are simply: **what could have gone wrong, and what it would have cost.** If the answer is nothing, there is no story, only a report.

Three forms of cost you can actually show

1. **What was risked.** The money put in, the job left, the order accepted that was larger than anything they had done. Risk is only legible if the scale is stated — "a big order" is nothing, "four times anything we had made before" is a stake.
2. **What was given up.** Every real choice excludes something. The workshop that will not take rush jobs is refusing revenue, and saying so is far more persuasive than claiming to care about craft, because it has a price attached.
3. **What went wrong.** The batch that failed, the delivery that missed, the design that had to be rebuilt. This is the most uncomfortable and the most valuable.

Admitting a fault buys belief, and there is a reason

When a party with an obvious interest says something against that interest, the statement carries weight the rest of their claims do not, because the usual explanation — that they said it because it suited them — is unavailable. Legal systems have a version of this in the treatment of statements against interest, and audiences apply it informally all the time.

In practice: "we missed eleven of 2,400 deliveries last quarter" makes the other 2,389 believable in a way that "we are reliable" never will. The admission is doing work for the claim, not against it.

The limits matter. This works when the admitted fault is real, small enough to survive, and paired with what was done about it. It fails when the admission is cosmetic — "our only weakness is that we care too much" is recognised instantly as a job-interview answer and costs credibility rather than buying it. It also fails when the fault is the central thing the customer is buying; a courier admitting that deliveries are often late is not being charmingly candid, they are describing a broken service.

Stakes in a fifteen-second piece

There is time. A stake can be one line of on-screen text or one shot.

- A hand hesitating before a cut that cannot be undone.
- "Third attempt."
- The clock in shot showing 11:40pm.

Each of these poses a question — will it work, why a third attempt, why is anyone working at that hour — and a question is all you need. The answer can arrive two seconds later.

Do not manufacture it

The failure mode is invented adversity, and it is detectable by a specific asymmetry: the struggle is described in general terms while the resolution is described in concrete ones. "There were difficult times" followed by "and now we produce 400 units a week" is the signature of a story whose first half was written afterwards.

If there genuinely were no stakes — the business started, it worked, nothing dramatic occurred — then say that. It is unusual and slightly disarming, which is itself interesting. What you must not do is borrow a hardship narrative, particularly one drawn from somebody else's real circumstances, because the people who recognise it will recognise it precisely.

The question to ask the client

Not "what challenges have you faced", which produces a rehearsed answer. Ask instead: **"what is the thing that nearly went wrong that you still think about?"**

The phrasing works because it asks for a memory rather than a position, and because "still think about" gives permission for the answer to be small. What comes back is usually specific, usually unflattering in a minor way, and almost always the best material in the interview. It is also the point at which you have to exercise judgement: some of what people tell you in that moment is not for publication, and asking before you use it is both a courtesy and the thing that keeps them willing to tell you anything next time.

BEFORE YOU MOVE ON

A delivery firm's advertisement states "we missed 11 of 2,400 deliveries last quarter, and here is what changed". Why does this strengthen rather than weaken the claim of reliability?

- A. Transparency is fashionable, so audiences currently reward brands that publish negative figures.
 - B. The failure rate is low, so stating it is effectively a boast disguised as a confession.
 - C. A statement against the speaker's own interest removes the usual discount, so the surrounding figures inherit credibility.
 - D. It pre-empts negative reviews, so complaints that appear later have less impact on prospective customers.
-

23 · 8 min

The person who changes is the one it is about

THE ONE THING TO KEEP

The protagonist is whoever is different at the end — test a script by deleting the brand name and seeing whether a story remains, and put the brand at the turning point rather than in the leading role.

A simple test you can run on any script

Read the piece and ask: **who is different at the end?** That person is the protagonist, whatever the brief says the film is about.

Run it on the average company film and the answer is the company. It was founded, it grew, it invested, it now serves customers nationwide. The customer appears as evidence of the company's progress rather than as somebody whose life changed. That is a company history, and company histories are interesting to roughly the number of people who work there.

Cross out the name

A second test, faster and crueller. Delete every instance of the brand name from the script and read what is left. If what remains is a story — somebody had a problem, tried something, ended up different — you have a story that a brand appears in. If what remains is nothing, you had an advertisement with anecdotes decorating it.

This is not an argument for hiding the brand. It is a diagnostic for whether anything is being carried other than the brand.

When the brand should be the protagonist

Sometimes it genuinely should, and the blanket advice to "make the customer the hero" is repeated far past its useful range. The brand is the right protagonist when:

- The change being described actually happened to the company and is interesting — a real pivot, a real failure, a decision that cost something.
- The audience is investors, staff or partners, for whom the company's trajectory is the subject.
- The founder's own experience is the product, which is true for a great deal of craft and service work.

The failure is not brand-as-protagonist. It is **nobody** as protagonist: a film in which no one changes, narrated in the passive voice, where achievements accumulate without anyone having done them.

Give the protagonist a want and an obstacle

A protagonist without a want is a subject of a portrait. The want does not need to be grand — to get the order out by Friday, to stop being embarrassed by the shopfront, to understand why the form keeps being filled in wrong. What matters is that it is stated early and that something is in the way, because the gap between want and obstacle is the space the piece occupies.

The commonest scripting error here is a want that is already satisfied. "She wanted a reliable supplier, and she found one" has no gap. "She had a supplier

she did not trust and a client meeting on Thursday" has one.

Where the brand goes instead

If the customer is the protagonist, the brand's structural role is the thing that makes the turn possible. Not the hero, not the narrator, not the moral. In practice this means:

- The brand appears at or just after the turning point, because that is where it is relevant.
- It is shown doing something specific rather than being described. The blade cutting, the van arriving, the person on the phone at 7am.
- It gets less screen time than the client would like, and the case for that is made at the beats stage rather than in the edit, when it becomes a fight.

Two honest complications

Consent. If a real customer is the protagonist, they are a person with a life that continues after your film. They need to know what it will be used for, where it will run and for how long, and they need to be able to say no afterwards. A signed release is normal practice; it is also not a substitute for actually asking. This matters most when the story involves difficulty — somebody's business nearly failing, an illness, a family situation — because the footage that helps your piece may be the footage they regret most in two years.

One protagonist. A film with four customers has four small stories and no arc. If the brief demands several voices, the honest structure is usually not a story at all but a set of statements, presented as such. Do not glue four testimonials together and call it a narrative; the joins will be "and then", and the viewer will feel every one.

BEFORE YOU MOVE ON

A script survives the brand-name deletion test but still feels flat: a customer is shown, described sympathetically, and ends the film as she began. What is the diagnosis?

- A. The brand is too absent, so the piece has no anchor and the viewer has nothing to attach the sympathy to.
 - B. There is a subject but no protagonist, because a protagonist is defined by a change that this character does not undergo.
 - C. The customer is not sufficiently representative of the target audience, so viewers do not identify with her.
 - D. The piece needs a narrator to supply the meaning that the visuals are not currently carrying on their own.
-

24 • 8 min

Endings, and the four asks that cancel each other out

THE ONE THING TO KEEP

An ending must resolve, land the meaning and ask once — four simultaneous asks force a choice at the moment attention is lowest, and an ask larger than the piece earned reframes everything before it as a pitch.

An ending does three jobs, and most do one

The last few seconds have to resolve the thing that was opened, land what it means, and say what happens next. Most endings do the third, badly, and skip the first two — which is why so many otherwise decent pieces finish with a logo, a website, a phone number, a QR code and three social handles arriving simultaneously over music that fades out.

Take them in order.

Resolve. Whatever question the piece opened has to close. If you started with a shopfront that no longer fits, the ending shows the sign that does. Shown, not stated — a claim at the end is the weakest position for a claim, because the viewer has just spent ninety seconds forming their own view.

Land the meaning. One line, or one image, that says what this was about beyond the events. This is the part that gets cut for time and is the reason some endings feel like they stopped rather than finished. It need not be profound. "Two panels. Same sign. Reads from the bus stop now."

Ask. One thing. Which brings us to the real problem.

Why four asks produce fewer actions than one

A viewer at the end of a piece has a small amount of willingness to act. Present one option and they either take it or they do not. Present four and they must first choose, and choosing costs effort at exactly the moment their attention is lowest. A meaningful share resolve the effort by doing nothing, which is always available and always free.

There is a second mechanism, specific to screens: the four asks physically compete. The QR code, the URL, the handle and the phone number each take space and contrast from the others, so all four are smaller and less legible than one would have been, and at the moment where the viewer has least patience for squinting.

So: **one ask, at the size the whole ending can afford.** Everything else goes in the caption, the description, the bio — places that are still there after the piece has ended.

Make the ask small enough to be true

The ask has to match what the piece earned. A ninety-second film introducing an unfamiliar business has not earned "book a consultation"; it may well have earned "save this" or "call and ask what it would cost". An ask that overreaches is not merely ignored — it retroactively reframes the piece as a sales pitch, which devalues everything before it.

A useful check: would a reasonable person who liked the piece do this within ten seconds, on the device they are holding, without standing up? If not, the ask is too large or the wrong channel.

Two mechanics worth knowing

The loop. On platforms that replay automatically, an ending that connects back to the opening makes the second viewing feel intentional rather than accidental, and watch-through figures rise. This is a real effect and a modest one; it is not a reason to contort a piece that has a proper ending.

Holding the last frame. Give the final card two to three seconds of stillness. Editors consistently cut this short because on the tenth viewing it feels interminable — a reliable illusion, since you already know what it says and the viewer is reading it for the first time. Three seconds is roughly the time to read and act on a short line. In DaVinci Resolve, Kdenlive, Shotcut or CapCut this is one drag of the clip's edge.

Endings that do not ask for anything

Not every piece should end with a request. A film whose job is to make somebody think differently about a subject can end on the thought, and appending "follow for more" to it cheapens the whole thing — a sudden switch from talking to selling, in the last two seconds, which viewers register as a reveal of what the piece was really for.

When there is no honest ask, end on the resolution and the meaning, and let the work stand. The brand can be present without a demand attached. This is a harder case to make to a client than a conversion-shaped ending, and the argument that works is about the measure agreed at the start: if success was defined as a shift in belief, then a request to buy is not the ending that was commissioned.

BEFORE YOU MOVE ON

A reel ends with a QR code, a website, a phone number and two social handles. Response rates are lower than for an earlier version with only the phone number. What is the likeliest mechanism?

- A. QR codes require a second device or an app, so including one adds friction that suppresses every other option on screen.
- B. Choosing costs effort at the point of lowest attention, and the four options also shrink each other's legibility.
- C. Multiple contact routes signal a disorganised business, which reduces trust just before the moment of action.
- D. The earlier version performed better because phone calls convert more reliably than digital routes for this audience.

CASE STUDY · MODULE 3

The founder's forty years, and the customer's Thursday

A two-bench bamboo furniture workshop outside Agartala, commissioning a ninety-second film for a state handicrafts exhibition.

The founder's brief was a list. The workshop began in 1978. Three generations have worked in it. It won a state award in 2009. A finishing machine was bought in 2011. It makes chairs, stools, screens, trays and lamp bases.

Written out as beats, every join between those sentences is and then. That is the diagnostic: each beat can be deleted without breaking the next one, and a viewer who senses that leaving costs them nothing will leave.

Two beat lists, seven sentences each

The film-maker wrote both versions as plain sentences before touching a camera, because at that stage a change costs thirty seconds.

The company version, honestly written out, had no change in it anywhere. At the end of ninety seconds the workshop existed, the family cared, the furniture was good, exactly as at the start. That is a description. Descriptions are legitimate objects and they have to earn attention through usefulness or beauty, never through suspense, and a description of a workshop set to music earns neither.

The second version came from a question asked in the interview: what nearly went wrong that you still think about. An architect in Guwahati had ordered sixty chairs for a cafe, on a five-week deadline, and had refused the first delivery because the joint finishing varied visibly across the run. The workshop remade twenty-two chairs at its own cost and changed how it batches finishing.

That is a state, an event, and a different state, and it joins with but and therefore at every seam.

The decision, and the cost on both sides

The founder's son objected to putting the rejected batch in the film, and his objection was not vanity. The exhibition audience included a state procurement officer and at least one competing unit. A recorded admission that a delivery had been refused could be read at a later tender as a quality record rather than as candour.

The counter-argument is the one about statements against interest. A claim made by an interested party carries a discount, and the discount disappears when the statement damages the speaker, because the usual explanation for why they said it is unavailable. Twenty-two remade chairs is what makes every other claim in the film believable.

Both of those are true at the same time, and the choice had a price either way: silence buys a film nobody remembers, and candour buys a sentence that may be quoted back in a procurement room.

There was a second decision underneath it. The architect was a real person whose business continues after the film. She had to be asked what the piece was for, where it would run and for how long, and told she could ask for it to come down.

Who the film was about

The test that settled the structure took one minute. Delete every instance of the workshop's name from each script and read what is left.

From the company version, nothing was left. Founded, grew, invested, won, makes: remove the name and there is no sentence with a person in it. From the second version, a story remained in full, because somebody wanted sixty chairs in five weeks, was let down, and ended up with a supplier she now uses for everything.

That is not an argument for hiding the workshop. It is a diagnostic for whether anything is being carried other than the workshop. The brand still appears, at the turning point, where it is relevant, and it is shown doing something specific rather than being described: the second batch going onto the bench, the new finishing jig, the number chalked on the wall.

Where the ninety seconds went

The first assembly spent forty-four seconds arriving: the road, the gate, the sign, the founder sitting down. The turn arrived at fifty and had twenty seconds left.

Rewritten to the beat list, the promise took three seconds, the situation and the problem seventeen, the turn forty, the new state twenty and the ask ten. The maker counted the narration at 231 words before recording anything, which at about 150 words a minute is where ninety seconds actually lands, and that count is why nothing had to be cut in a hurry at the end.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

They made the version with the rejection in it and did three things to limit the risk the son had named. The cafe was not named. The number was stated rather than softened, because twenty-two is specific and a batch is not, and the architect asked for the number herself when she saw the cut. And the sentence about what changed afterwards was placed immediately after the admission rather than at the end, so the admission is never on screen alone.

The asks were cut from three to one. The first cut ended with a website, a phone number, a handle and an invitation to visit the stall, all arriving together over music. The final version ends on the remade chair, one line about the joint, and one ask: bring us a drawing.

The procurement officer did ask about the refused delivery at the stall. The answer given at the stall was the same as the answer in the film, which is the only reason it was an easy conversation. Three enquiries over the exhibition mentioned the film unprompted, against a target nobody had set, which the report recorded as a failure of the brief rather than of the film.

The founder's list was entirely true. Why does a chain of true facts joined by 'and then' fail to hold attention?

Because each beat is optional. If beat three does not depend on beat two, the viewer can leave after beat two at no cost, and they can feel that even though they could not name it. But and therefore are the joints that make one beat cause the next, and a piece whose beats are causally linked cannot be reordered or cut without visible damage, which is what keeps someone watching.

How does admitting a refused delivery make the rest of the film's claims more believable rather than less?

An interested party's claims carry a discount because the obvious explanation is that they said it because it suited them. A statement that damages the speaker removes that explanation, so it is weighted differently, and the credit transfers to the claims around it. The limits matter: it works when the fault is real, survivable, and paired with what was done about it, and it fails when it is cosmetic or when the fault is the thing the customer is buying.

The first cut ended with four asks. Why does that produce fewer actions than one ask?

A viewer at the end of a piece has a small amount of willingness to act. One option is taken or not; four options require a choice first, and choosing costs effort at the moment attention is lowest, so a meaningful share resolve it by doing nothing. There is a second, physical mechanism: four asks share the same space and contrast, so each is smaller and less legible than one would have been.

MODULE 3 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 You have written a script's beats as plain sentences. What tells you that a beat can be deleted without the piece breaking?
 - A. It contains no number, name or other checkable detail.
 - B. It repeats a fact that was already given in the opening, so the viewer has heard it once.
 - C. The honest joining word between it and the beat before is and then.
 - D. It runs longer than four seconds in the planned edit.
- 2 Beyond the vague instruction to grab attention, what job does the first second of a piece actually have to do?
 - A. Surprise the viewer so they cannot predict what is coming next.
 - B. Answer what this is and whether it is for me.
 - C. Introduce the brand so the viewer knows who is speaking to them.
 - D. State the strongest claim the piece is going to make later.
- 3 A well-shot ninety-second film opens on a workshop in morning light and ends with the owner saying they have always cared about quality. Nothing is different at the end. What have you got?
 - A. A story whose turn is too weak, which a stronger and more decisive ending would repair on its own.
 - B. A testimonial, which needs a named customer to carry it.
 - C. An advertisement, which needs a clearer call to action.
 - D. A description, which must earn attention through usefulness or beauty, never suspense.

- 4 Your narration script for a planned ninety-second film runs to 400 words. What does the count tell you before you record anything?
- A. This is roughly a two-and-a-half-minute film, and nobody has decided that yet.
 - B. The delivery will have to run faster than 200 words a minute to fit.
 - C. About a third of the words are padding and should be removed by the mechanical passes.
 - D. Nothing useful; narration length is set in the edit rather than on the page.
- 5 A courier writes that it missed 11 of 2,400 deliveries last quarter. Why does that make its other claims more believable rather than less?
- A. Precise figures are inherently more persuasive to readers than round ones.
 - B. It demonstrates that careful records are kept, which implies a competently run operation.
 - C. A statement against the speaker's own interest cannot be explained by it suiting them.
 - D. Admitting a fault produces sympathy, and sympathy lowers the reader's scrutiny.
- 6 You want to know whether a script is a story that a brand appears in, or an advertisement with anecdotes on it. Which test answers that fastest?
- A. Remove the music and see whether the piece still holds without it.
 - B. Delete every instance of the brand name and see whether a story remains.
 - C. Cut the first shot and see whether the opening still answers its two questions.
 - D. Remove the customer and see whether the brand still reads clearly on its own.

MODULE 4

Getting the material

Interviews, footage, photographs, archive and numbers are the raw material a structure is built from. This module covers how to ask questions that produce usable answers, which images carry meaning and which are wallpaper, what consent actually has to cover, how anonymity fails, how licences work, and how to audit what you have before you start building.

BY THE END OF THIS MODULE YOU CAN

Gather the material a planned piece needs — run an interview that yields self-contained answers, distinguish images that carry meaning from images that decorate, obtain consent whose scope matches the intended use, protect a subject against jigsaw identification, clear found material against its actual licence, present a number with its denominator, and audit every beat against the asset that will carry it before building begins

25 · 8 min

Draw ugly boxes now or pay for it later

THE ONE THING TO KEEP

Plan in throwaway frames while changes still cost thirty seconds, and get consent in writing before the camera comes out, not after.

The step everyone skips

Six boxes on a sheet of A4. Stick figures. Arrows for camera moves. Handwriting you would not show anyone. Twenty minutes, and nobody ever sees it.

That is a storyboard, and skipping it is the most expensive habit in this course, because the cost does not appear until later and it always appears as something else — a reshoot, a client argument, a piece that never quite works and nobody can say why.

The reason is arithmetic. A change on paper costs thirty seconds. The same change after the shoot costs a reshoot. After the edit it costs a re-edit and possibly a reshoot you can no longer get, because the location has closed, the light has changed, or the person in shot three has left the company. Every hour of planning is buying options that expire.

The same change, priced by when you notice it

On paper	●	Thirty seconds and a new box
On location	●	One more setup before you pack up
That evening	●	A phone call and a second visit
In the edit	●	A rewrite, or a reshoot nobody is paying for
Weeks later	●	The light has changed and the shop has closed

Nothing about the change gets harder. What expires is the option to make it, which is what an hour of planning actually buys.

What the boxes catch

Things you cannot see in your head, and can see immediately in six ugly frames:

- You have no shot of the actual product. You have four shots of atmosphere.
- Shots three and five are the same angle at the same distance, so the cut between them will look like a mistake.
- The story needs an aerial shot of a rooftop you do not have access to.
- Everything is a medium shot. Nothing is close enough to feel and nothing is wide enough to establish.
- You have twelve seconds of setup for an eight-second idea.

Under each box write one line of what happens and one line of what is said or on screen. That is the whole method.

The shot list

The storyboard is the picture; the shot list is the plan. Five columns, on paper or in any text file:

- Shot number
- Size — wide, medium, close, extreme close
- What happens
- Sound or on-screen text
- Roughly how many seconds

A thirty-second reel usually runs to six to ten shots. If your list has three, it will feel static. If it has twenty-five, it will feel frantic and you will not have time to light any of them.

The shot list is also the thing that tells you whether the shoot is possible. Add up the seconds; add up the locations; look at how many setups need the same light. Most one-day shoots that overrun do so because nobody counted the setups.

Thumbnails, the same tool for static work

A poster gets the identical treatment under a different name. Twelve small boxes on one sheet, twenty minutes, before opening any software. Not decorative — just where the mass sits, where the big word goes, where the eye enters.

The mechanism is worth understanding, because "just start in the software" feels faster. Once you have a layout on screen with real type and a real photograph, changing it is expensive: you have invested effort, undoing feels like loss, and you start nudging instead of exploring. Twelve throwaway boxes cost nothing to abandon, so you actually consider the option where the image bleeds off the top, which you would never have tried after forty minutes in Photopea.

Doing it free, including on a phone

- **Paper.** Still the fastest. A pencil and a sheet folded into six.
- **Slides.** Google Slides or LibreOffice Impress, one slide per shot, a text box underneath. Reorderable, shareable with a client, and works on a phone.
- **Photoboard.** Go to the actual location with your phone and take one still per planned shot. This is better than drawing for anything shot in a real place, because it tells you the truth about the light and the background clutter.
- **Krita** or any free sketching app if you would rather draw. None of this needs to look good.

Client-facing tip: show a rough board rather than a finished draft. Clients give braver feedback on something visibly unfinished, and reserve their politeness for things that look done. A polished draft gets "lovely" and then gets rejected at the end.

Two things planning is really for

Consent, before the camera comes out

The moment your board includes a real person — an employee, a customer, a child, a bystander in the background — permission belongs at the planning

stage, in writing, before the camera comes out. This is not paperwork for its own sake. Before the shoot you have a conversation; after the shoot you have a piece you cannot use unless they agree, and your leverage is zero. That applies with more force to anyone whose likeness or voice you intend to reuse, alter or synthesise later.

A one-paragraph release, signed and photographed on a phone, is enough for most small jobs. What counts as valid varies by country and by whether children are involved, so a lawyer where you live is the person who answers that for your case. The habit — ask first, in writing, name what it will be used for — is the same everywhere.

Knowing when not to board

Over-boarding kills observational work. If you are filming a market, a wedding or a real conversation, a locked twelve-frame board makes you chase your drawings instead of watching what is in front of you, and the best moment of the day happens while you are setting up shot seven.

For that kind of work, board the *structure* — the beats from the previous lesson, and a short list of shots you must not leave without — and keep the frames open. Planning what has to exist is different from planning what it looks like.

Do this today

Take a video you plan to make and draw six boxes for it, badly, in under fifteen minutes. Then count how many of the six are the same shot size. If more than three are, you have found something that would have shown up only in the edit.

BEFORE YOU MOVE ON

A designer skips thumbnail sketching and starts a poster directly in Photopea. Four hours later the layout is competent but they cannot see any alternative to it. What best explains the narrowed thinking?

- A. The software's default artboard and grid settings pushed the composition toward a conventional layout from the start.
 - B. Four hours is not long enough to explore alternatives in detail, so the fix is more time in the same file.
 - C. Once real type and images are placed, abandoning the layout feels like losing invested work, so exploration turns into nudging what already exists.
 - D. Working on screen at full artboard size hides how the poster will look small, which is what limits the range of layouts considered.
-

26 • 9 min

Most interview footage is unusable, and the question is why

THE ONE THING TO KEEP

Your question is absent from the finished piece, so ask for self-contained answers and for specific moments rather than summaries — and check every number, because memory reconstructs rather than replays.

The problem is structural, not personal

You record forty minutes with somebody interesting and find, at the edit, that you can use ninety seconds of it. This is normal, and it is almost never because the person was dull. It is because of one mechanical fact: **your question is not in the finished piece, and their answer has to stand without it.**

Ask "how long have you been doing this?" and you get "about eleven years". Eleven years of what? In the edit, that clip is orphaned. You now need a voice-

over or a caption to repair it, which costs you time and costs the piece a moment of somebody's actual voice.

So the first rule is not about rapport. It is: **ask in a way that forces a self-contained answer.**

The instruction that fixes most of it

Before you start, say this, plainly:

"I won't be in the finished thing, so if you can fold my question into your answer, that helps. Instead of 'eleven years', something like 'I've been repairing sewing machines for eleven years'."

Most people get it immediately and do it for the rest of the session. The ones who forget can be nudged once: "sorry, can you give me that one as a full sentence?" You are not correcting them, you are explaining a constraint of the format, and people are generally happy to help once they know what the machine needs.

Questions that produce material, and questions that produce positions

A position is a rehearsed answer — a thing the person has said before and will say the same way again. "What makes you different from your competitors?" produces a position every time.

A memory is specific, unrehearsed, and usually better. You get memories by asking for **a moment rather than a summary**:

- Not "is the work difficult?" but "when did you last get something wrong?"
- Not "what are your values?" but "what have you turned down?"
- Not "how do customers find you?" but "who was your first customer, and how did they hear?"
- Not "why did you start?" but "what were you doing the week before you started?"

The grammar matters. "When did you last..." asks for an event and gets a date, a place and a person. "Do you..." asks for a yes.

Four mechanics that cost nothing

1. **Wait three seconds after they stop.** Most people treat a silence as their responsibility and fill it, and the filling is frequently the best part — the qualification, the second thought, the thing they had decided not to say. You have to be genuinely quiet, and visibly not reaching for your notes, or they read the pause as the interview ending.
2. **Ask the same thing twice, differently, twenty minutes apart.** The second answer is shorter and better, because they have now heard themselves say it once.
3. **Record thirty seconds of the room saying nothing.** This is room tone, and it is what you lay under an edit so that the gaps between sentences do not sound like the audio has dropped out. Forget it and every cut in the interview will click.
4. **End with "is there anything I should have asked?"** It is the single highest-yield question in the set, because it hands over the framing, and people who have been answering for half an hour usually have one thing they have been waiting for.

The free kit is genuinely adequate

A phone's voice-memo app records a better interview than a camera's built-in microphone, because the phone can sit close to the speaker while the camera sits far away and sound falls off fast with distance. Put the phone on the table, on a folded cloth so it does not pick up knocks, and sync it later. Audacity edits it for nothing; DaVinci Resolve's free version will align it to picture. A clip-on lavalier microphone costs about the price of a meal and is the cheapest large improvement available to anyone.

The limitation you must design around

An interview is testimony, not record. Human memory reconstructs rather than replays: each retelling is assembled from fragments plus expectation,

which is why people give dates, quantities and sequences with complete confidence and get them wrong. This is not dishonesty and it is not unusual.

The practical consequence is a rule: **anything checkable gets checked.** If they say the business started in 2011, find the registration. If they say four hundred units a week, ask what the figure is counted from. You are not doubting them; you are protecting them, because the number that goes out under their name with your edit around it is the one a stranger will argue with.

BEFORE YOU MOVE ON

Why does asking an interviewee to fold your question into their answer matter more than getting a warm, relaxed performance out of them?

- A. Warmth is subjective while sentence structure is objective, so it is the only part of an interview a beginner can control.
- B. Longer answers give the editor more choices, and more choices always produce a better cut.
- C. The question is cut out, so an answer that depends on it is orphaned and has to be repaired with narration or a caption.
- D. Interviewees who speak in full sentences are perceived as more credible by audiences, which raises the persuasive power of the clip.

27 · 8 min

Half your footage is wallpaper

THE ONE THING TO KEEP

Every image should illustrate, prove or bridge — cover the picture and then the words, and if understanding is unchanged either way the image is carrying nothing, while an image that contradicts the words quietly overrides them.

Three jobs an image can do, and one it cannot

When a picture appears over somebody speaking, or beside a paragraph of text, it is doing one of three jobs:

1. **Illustrative** — it shows the thing being described. She says "the machine jams on thick denim", you see thick denim jamming a machine.
2. **Evidential** — it proves something the words only claim. He says "we make four hundred a week", you see the rack with four hundred on it.
3. **Continuity** — it covers a join. You have cut two sentences together that were three minutes apart, and the cutaway hides the jump.

The fourth category is **wallpaper**: a slow pan across a workbench, a shot of hands doing something unrelated, a stock photograph of a team smiling at a laptop. It is doing none of the three jobs. Most people's material is largely wallpaper, and it feels productive to shoot because it is easy and looks nice.

The test

Cover the image. Read or listen to the words alone. Then cover the words and look at the image alone. **If your understanding is the same in both cases as it is with both together, the image is carrying nothing.**

This is a workable version of a real finding about how people take in two channels at once: pairing a picture with words helps recall when the picture adds information the words do not have, and helps very little when the picture

merely restates them. Identical is not reinforcement. Identical is one channel running twice while the viewer's attention pays for two.

The shot that contradicts is worse than none

There is an asymmetry worth knowing: when picture and voice disagree, most viewers believe the picture and rewrite the voice to fit. Say "we do everything by hand" over a shot that happens to include an industrial machine in the background and you have not told a small inconsistency, you have told the audience that the claim is untrue. They will not be able to name what bothered them afterwards; they will simply trust the piece less.

This is why "we'll find something to put there" is a dangerous habit. A cutaway grabbed for its prettiness carries whatever it happens to contain, including the thing in the corner you did not notice.

Shooting so you have coverage rather than wallpaper

A reliable habit: for every idea the piece has to carry, get three shots — **wide, detail, hands**.

- The **wide** establishes where you are and how big it is.
- The **detail** is the thing itself, close enough to be unambiguous.
- The **hands** shot is a person doing the action, which supplies motion and tells you a human is involved.

Three shots per idea, ten seconds each, is thirty seconds of shooting that gives an editor real choices. An hour of drifting pans gives them none.

Photographers and documentary crews have taught a version of this for decades as a five-shot sequence — hands, face, wide, over-the-shoulder, and one unusual angle — and the reason it survives is that it produces cuts that work rather than footage that exists.

Doing it for nothing

A phone shoots all of this. Lock the exposure before you start a shot, because a phone left on automatic will hunt for brightness in the middle of your take

and the change is very hard to hide. Hold each shot for ten seconds even when it feels absurd; editors need handles at both ends. Shoot detail shots from about a hand's width further back than feels right, because you cannot uncrop but you can always crop in.

For still work the same rule applies with different names. An image in a document, a slide or a post either carries information, proves a claim, or bridges a structural gap — or it is decoration, and decoration in an information-carrying piece is a cost paid in attention and load time.

The honest limitation

Not every piece needs to be maximally efficient. A pause, a piece of texture, a shot of the street outside, can buy the audience a moment to absorb what just happened, and a relentlessly information-dense edit is exhausting to watch. The distinction is intent: a breathing shot placed deliberately after a heavy idea is doing a job, and you can say what the job is. A breathing shot placed because the edit felt thin is wallpaper with a better explanation.

BEFORE YOU MOVE ON

A video says "every piece is cut by hand" over a workshop shot that includes a laser cutter in the background. What is the mechanism by which this damages the piece?

- A. Viewers resolve a conflict between picture and voice in favour of the picture, so the spoken claim is heard as untrue.
- B. The background detail distracts the eye from the subject, so less of the sentence is absorbed.
- C. Competitors could screenshot the frame and use it to accuse the business of false advertising.
- D. Mixing handmade and machine imagery muddles the brand's positioning across the wider campaign, so later pieces inherit the confusion this one introduced.

28 · 9 min

Consent is a scope, not a signature

THE ONE THING TO KEEP

A release fails on scope rather than signature, so state in plain words what the material is for, where it will appear, for how long and that it may be edited — and agree a withdrawal route before anyone needs it.

What a release is actually for

A release is a record that a person agreed to a specific use of their image, voice or words. The signature is the least interesting part of it. The part that matters, and the part that fails, is the **scope**: what the material will be used for, where it will appear, for how long, and whether it can be edited.

The complaints that arrive months later are almost never from someone who denies agreeing. They are from someone who agreed to one thing and found another. They said yes to a piece about the workshop and turned up in a recruitment advertisement. They said yes to a website and found themselves on a billboard. The signature was real; the scope was not what they understood.

So write the scope in words the person actually reads, and keep it narrow enough to be true.

The plain-language minimum

A usable release, in ordinary sentences, names six things:

1. **Who** is being recorded and who is doing the recording.
2. **What** is being recorded — an interview, photographs, video in the workshop.
3. **What it is for** — "a short film about the business, and stills from it".
4. **Where it may appear** — "the company website, its social accounts, and printed material". If you might sell it on, say so; if you will not, say that.

5. **How long** — a stated period, or "until you ask us to stop".
6. **That it may be edited**, because it will be, and somebody who expected their whole answer will be upset by a five-second cut of it.

This fits on one side of paper. It does not need legal language, and dressing it in legal language tends to make it worse, because a person who does not understand a document has not really consented to it even when they have signed it.

Recording consent when paper is not practical

On a street, at an event, or with somebody who is about to change their mind if you produce a form, film the consent instead. Point the camera at them and ask them to say their name, today's date, what they are taking part in, where it will be published, and that they understand they can ask you to take it down.

This is not a superior legal instrument. It is a better practical one, because it lives in the same file as the material, it is timestamped by the recording itself, and it cannot be separated from the footage it covers the way a loose form in a drawer can.

Children, workplaces and the person who is not the point

- **Under-18s:** a parent or guardian agrees, and in most places the young person should be asked as well, because an unwilling child on camera is a problem regardless of what a guardian signed.
- **Employees:** consent given to an employer is under pressure by construction. The person who says no is saying no to their manager. Give them a route to decline that does not require them to decline to your face — a form collected later, an opt-out email address.
- **Bystanders:** people in the background of a shot who are recognisable are a real category. Repositioning the camera is nearly always easier than clearing them.

The law genuinely differs, and this is not legal advice

What you may film and publish without permission varies sharply by country, and the differences are not small details. Some jurisdictions treat a person's image as a personal right that survives being in a public place; others treat public space as broadly open to photography. Data-protection regimes — the EU's GDPR, India's Digital Personal Data Protection Act, and the several dozen others — can treat an identifiable face as personal data, which brings duties about purpose and retention that a release does not by itself satisfy. Rules for commercial use differ again from rules for journalism and for art.

None of this is settled into a single global rule, and it moves. Find out what applies where you are and where you publish, and for anything with money or risk attached, ask a lawyer in that country. What is portable is the practice, not the law: **narrow scope, plain words, written down, and a route to withdraw.**

Decide the withdrawal policy before anyone asks

Someone will email in three years and ask to be removed. Perhaps they left the company badly, perhaps the piece resurfaced somewhere unpleasant, perhaps they simply changed.

Decide now what you do, and tell them at the time of filming. A defensible policy for a small operation: material comes down from anything you control within a stated number of days, you will not chase copies you do not control, and you will not re-use them in anything new. Saying this at the start makes people far more willing to take part, because the thing they are actually afraid of is permanence.

BEFORE YOU MOVE ON

A signed release grants use "in all media, worldwide, in perpetuity". Why is this more likely to generate a dispute than a narrower one?

- A. It covers uses the person never pictured, so the first unexpected appearance feels like a breach even though the paper allows it.
 - B. Broad wording is unenforceable in most jurisdictions, so a court would strike it out and leave you with nothing.
 - C. Perpetual terms require additional payment in most professional agreements, so the signer will later claim they were underpaid.
 - D. Wide grants attract more scrutiny from platforms, which increasingly require narrow, purpose-limited consent before hosting footage of identifiable people.
-

29 · 9 min

Anonymity fails by combination, not by name

THE ONE THING TO KEEP

Anonymity fails through the combination of details rather than the name, so count how many people each detail still fits, prefer subtractive protection over blur or pitch-shift, strip metadata, and tell the person plainly that identification can be made harder but never impossible.

The failure mode nobody anticipates

Somebody agrees to tell you something difficult — about a debt, an illness, an employer, a family — on the condition that they are not identifiable. You change the name, blur the face, and publish.

They are identified within a week, and not by anyone clever. They are identified because the piece mentioned a town of forty thousand people, a trade with perhaps nine practitioners in it, a daughter at a named school and an ac-

cident "about two years ago". Each fact on its own identifies nobody. Together they identify exactly one person, to anybody who already knows them a little.

This is **jigsaw identification**, and it is the way anonymity actually fails. Removing the name does very little work, because the name was never the identifier. The combination was.


Count the population each detail leaves

The practical method is to go through the piece detail by detail and ask, for each one, roughly how many people it still fits. Then multiply down.

- "A city in the north" — millions.
- "A town of forty thousand" — tens of thousands.
- "...who repairs violins" — perhaps a handful.
- "...whose daughter is at St Anne's" — one.

The moment the number reaches single figures you are publishing an identity. The fix is to make details less specific in ways that do not damage the truth of the story: "a small town", "a skilled trade", "a school-age child". If the specificity was load-bearing — if the story only means something because it is *violins* — then you have a decision to make with the person, not a detail to quietly adjust.

How many people each detail still fits



Not one of these details is the name, and together they publish an identity. Anonymity fails by combination, so the count to watch is the population each addition leaves standing.

Blur, pixelation and voice masking are weaker than they look

- **Pixelation and mosaic** reduce an image to a coarse grid. That grid is a lossy but structured transformation, and where the original is drawn from a small set of possibilities — a known group of people, a limited set of characters — the mapping can often be run backwards well enough to nar-

row the candidates. Pixelated text has been recovered this way repeatedly. Treat pixelation as obfuscation, not as deletion.

- **Moving blur** leaks across frames. A tracked blur that slips for two frames on a turn of the head is two clean frames in a file anyone can step through.
- **Pitch-shifted voice** is a reversible operation. Shifting it back is a menu item in free software. If a voice must be protected, have somebody else read the words, or use a text-to-speech voice and say that you have.

The robust versions are all subtractive: **crop the face out of frame, replace the audio, remove the detail from the shot rather than covering it**. What is not in the file cannot be recovered from the file.

The metadata you forget

A photograph from a phone usually carries EXIF data including the exact coordinates where it was taken, the device, and the timestamp. Publishing a "location withheld" image with its GPS intact is a common and complete failure.

`exiftool` is free and removes it in one command:

```
exiftool -all= -overwrite_original photo.jpg
```

Most social platforms strip EXIF on upload; your own website, a PDF, an email attachment and a file-sharing link generally do not. Check rather than assume, because the platforms that strip it are protecting their users, not yours.

Documents have the same problem in a different place. A PDF exported from a word processor carries the author name and often the file path, which is frequently a person's full name and their employer's network.

Show them before you publish

The single most useful step is also the one people skip: show the person the finished piece and ask whether they feel exposed. They know their own life. They will spot the detail that identifies them to their brother-in-law and that means nothing to you.

This costs you a day and it changes what you are doing from a promise you hope to keep into one you have checked.

Say what you cannot promise

Here is the limitation, and it should be said to the person out loud before they agree to anything. You can make identification harder. You cannot make it impossible, and you cannot make it revocable. Once a file is public it can be copied, archived, indexed and searched by tools that did not exist when you published it, and the risk of re-identification only ever goes up as more data about everyone accumulates elsewhere.

A person deciding whether to speak is entitled to that sentence, not a reassurance. Some of them will decide not to, and that is the system working.

BEFORE YOU MOVE ON

Why is cropping a face out of frame a stronger protection than pixelating it in place?

- A. A crop also removes the background, which is where most identifying detail actually sits.
- B. Pixelation signals to viewers that something is being hidden, which draws attention and invites people to try to identify the subject.
- C. Pixelation is a structured, lossy transform that can often be narrowed back to candidates, whereas cropped pixels are simply absent from the file.
- D. Video codecs reconstruct pixelated regions across successive frames during compression, so detail that was deliberately destroyed gradually reappears in the encoded file.

30 · 9 min

Free to download is not free to use

THE ONE THING TO KEEP

A licence answers whether you may use it commercially, must credit it, may change it, and what it obliges your own work to become — and ShareAlike and NonCommercial are the two clauses that reach furthest into work you thought was yours.

The four questions a licence answers

Every piece of material you did not make — a photograph, a piece of music, a clip, a font, a map, a chart — comes with a licence, even when nobody shows it to you. A licence answers four questions, and if you cannot answer all four you do not have permission, you have a download.

1. **May I use it commercially?**
2. **Must I credit it, and in what exact words?**
3. **May I change it?**
4. **What does my use oblige me to do to my own work?**

That fourth one catches people, and it is worth taking seriously before anything else.

Creative Commons, by mechanism

Creative Commons licences are built from four clauses that combine.

- **BY** — attribution required. Nearly all of them include it. The credit is not a courtesy; it is the condition on which the permission rests, and omitting it means you are using the work without a licence.
- **SA** (ShareAlike) — a copyleft clause. If you make something derived from it, your derivative must carry the same licence. That is the clause that reaches into your own work. A ShareAlike photograph in a commercial

brochure can oblige you to license the brochure the same way, which is very unlikely to be what your client agreed to.

- **NC** (NonCommercial) — no commercial use. This is the most dangerous clause in practice, because "commercial" is not defined by whether you were paid. It is about the nature of the use. A free video on a business's own channel is promoting a business; whether that counts is genuinely arguable, and the argument happens after publication.
- **ND** (NoDerivatives) — no modification. Cropping, colour-grading, cutting a song shorter, laying a title over it: all modifications.

CCo is different in kind: a deliberate waiver as close to public domain as the law in a given country allows. It is the only one with no ongoing obligation.

Public domain is a date, and the date depends where you are

Copyright expires. When it expires is jurisdictional: life of the author plus 70 years in the EU, the UK and the US for most modern works, plus 60 in India, with a tangle of exceptions for anonymous works, sound recordings, films and anything published before various cut-off dates. A work in the public domain in one country can be in copyright in another, and the internet delivers your page to both.

The practical consequence for a small operator is that "it's from 1890, it's fine" is usually right for text and often right for images, and that a 1940s photograph is a research task rather than an assumption.

The free stock sites, and their actual terms

Unsplash, Pexels and Pixabay are genuinely usable and genuinely free, and none of them is CCo any more, whatever people tell you. They run their own licences, which broadly allow commercial use without attribution but forbid selling the image itself unchanged and forbid using an identifiable person in a way that suggests endorsement or shows them in a bad light. That last restriction is the one that bites: a stock face on a poster about debt is a problem regardless of the licence.

For music, the free paths are real: the YouTube Audio Library, the Free Music Archive, ccMixter, and the large catalogue of attribution-licensed tracks by composers such as Kevin MacLeod, where the required credit line is published verbatim and must be copied exactly. For fonts, Google Fonts is licensed for commercial use including embedding; a font downloaded from a font-sharing site is very often a commercial font being distributed without permission, and your client is the one who gets the letter.

Keep a credits file from the first asset

One plain-text file per project, one line per asset: what it is, where it came from with the full URL, the licence name, the date you downloaded it, and the exact attribution string. This takes fifteen seconds at the time and is close to unrecoverable eighteen months later when somebody asks.

```
hero.jpg | unsplash.com/photos/abc123 | Unsplash Licence |  
2026-03-04 | no credit required  
intro.mp3 | freemusicarchive.org/... | CC BY 4.0 | 2026-03-04 |  
"Nocturne" by A. Vaz, CC BY 4.0
```

Fair use and fair dealing are not a plan

You will be told that short excerpts are fine. What is actually true is that different countries have different, non-equivalent exceptions. The United States has fair use, a flexible four-factor test decided case by case. The United Kingdom, India and many Commonwealth systems have fair dealing, which applies only to enumerated purposes such as criticism, review, reporting and research — a promotional video is not on the list. There is no international rule, the tests are decided after the fact rather than before, and platform take-down systems operate on their own logic regardless of either.

The disagreement is real and unresolved, and nothing here is legal advice. The working position for somebody without a legal budget: use material you have cleared, and treat any excerpt you cannot clear as a risk you are choosing knowingly rather than a right you have.

BEFORE YOU MOVE ON

A designer uses a CC BY-SA photograph inside a paid client brochure, crediting the photographer correctly. What is the specific problem?

- A. Attribution licences forbid paid work unless the photographer is separately compensated for the commercial context.
 - B. Cropping or colour-matching the photo to the brochure counts as a derivative, which BY-SA does not permit.
 - C. Creative Commons licences cannot be applied to photographs that contain identifiable people at all, so the licence on this particular image is void from the moment it was offered.
 - D. The ShareAlike clause can oblige the derivative work to carry the same licence, which the client almost certainly did not agree to.
-

31 • 9 min

A number is a claim with the argument removed

THE ONE THING TO KEEP

Every number carries hidden choices of base, window and definition, so state the absolute alongside the relative, name the window, and label the axis — because no chart is neutral and the honest move is making your choices visible rather than claiming not to have made any.

Numbers persuade because they look like they have stopped being opinions

"Sales grew 50%" reads as a fact. It is a fact about a ratio between two numbers, both of which somebody chose, over a window somebody chose, measured by a definition somebody chose. All three choices are arguments, and all three are invisible in the sentence.

This matters in two directions. You will be handed numbers by clients and you have to make them into something honest. You will also be tempted, every time, by the version that sounds better.

Always give the denominator

The commonest failure is a percentage without a base. "A 50% increase in customers" is two customers becoming three. "Doubled our reach" is 400 people becoming 800. Neither is dishonest until it is used to imply scale it does not have.

The repair is a habit rather than a judgement: **state the absolute and the relative in the same sentence.**

The same three numbers, with and without a base

Relative alone	Absolute in the same breath
• Sales grew 50 per cent • We doubled our reach • Complaints fell by a third • Nothing here can be checked	• From 2 customers to 3 • From 400 people to 800 • From 9 complaints to 6 • Both readings see the same pair

Neither column is dishonest until the left one is used to imply a scale it does not have. Stating both is a habit rather than a judgement, which is why it survives a deadline.

"From 2 in every 1,000 to 3 in every 1,000 — a rise of 50%."

That sentence is harder to misread in either direction. A reader who wants the alarming version and a reader who wants the reassuring version both have to look at the same two numbers. Medical and risk communication has converged on this construction for exactly that reason: relative change alone systematically inflates how large an effect feels, and absolute change alone can bury one that matters.

Three windows, three different truths

Ask anybody for a growth figure and they will give you the window that flatters. This is rarely deliberate; the flattering window is simply the one they remember.

- Measured from the worst month, everything is a recovery.
- Measured from the launch, everything is growth.
- Measured year on year, seasonal businesses stop lying to themselves.

The working rule is to **state the window in the same breath as the number**, and to prefer a window chosen for a reason you can say out loud: a full year because the business is seasonal, since the new premises because that is the change being discussed. If the number only works in one particular window, you have found something out about the number.

The chart is an argument too

A line chart makes three editorial decisions before anyone reads it: where the vertical axis starts, how wide the time span is, and what is left off.

- **A truncated vertical axis** — starting at 94 rather than 0 — magnifies small changes into cliffs. There are legitimate uses, mainly where zero is meaningless, such as body temperature or a stock index. There is no legitimate use on a bar chart, because a bar encodes quantity by length, and a bar cut off at the bottom is drawing the wrong length.
- **Squeezing the horizontal axis** steepens every line. Two charts of identical data at different widths tell different stories, and neither is false.
- **What is missing** is the hardest to see. A chart of surviving businesses is a chart with the failures deleted, and the shape it draws is the shape of survival, not of outcome.

Check the definition before you draw anything

"Customers" can mean accounts, people, households, or orders. "Active users" depends entirely on a period somebody picked. "Response time" can be a mean, which one bad week ruins, or a median, which hides the bad week entirely.

Ask two questions of every number you are given: **counted how, and from what?** A surprising proportion of the time the person supplying the figure does not know, and finding out changes it. That is not an obstacle to the work;

it is the most useful hour you will spend on it, because the number that survives the question is the one you can put on a wall.

The free path is complete

You do not need paid software for any of this. LibreOffice Calc and Google Sheets do the arithmetic. Datawrapper's free tier produces clean, accessible, embeddable charts and defaults to sensible axes. For anything programmatic, Python with pandas and matplotlib, or R, cost nothing and are what most published analysis is actually made with. If you are on a phone, Google Sheets and Datawrapper both work in a mobile browser.

Where this stops being technique

There is a line between choosing a clear presentation and choosing a flattering one, and it is not always obvious where you are standing. A usable test: **would you be comfortable if the audience saw the other defensible presentation next to yours?** If the answer is no, you have not designed a chart, you have built an argument that depends on the alternative being hidden.

The honest limitation is that no presentation is neutral. Every chart has an axis, a window and an omission, and a person who says their version is simply the data is either inexperienced or selling something. What you can do is make your choices visible — label the axis, state the window, say what is excluded — so a reader can disagree with the framing rather than only with the conclusion.

BEFORE YOU MOVE ON

Why is a truncated vertical axis defensible on a line chart of a stock index but not on a bar chart of quarterly sales?

- A. Stock indices change too fast for a zero baseline to stay useful, while sales figures are stable enough to fit one.
 - B. A bar encodes quantity by its length, so cutting the baseline draws a length that no longer matches the value.
 - C. Bar charts are read by general audiences while index charts are read by specialists who can correct for the axis.
 - D. Financial conventions permit truncation, and business reporting conventions do not, so the difference is one of professional norms.
-

32 · 8 min

Find the missing shot now, not in the edit

THE ONE THING TO KEEP

List every beat against the specific file that carries it, so that a missing asset becomes a shopping list and an impossible beat becomes a rewrite — and then mark what you will not use, because selection cost grows with the number of pairs you could compare, not with the number of takes.

The gap that costs a day

You have your structure — beats joined by *but* and *therefore*, an opening that makes a promise, an ending the piece has earned. You have four hours of material. You start building, and at beat five you discover there is no shot of the thing beat five is about.

The cost of that discovery depends entirely on when it happens. Noticed while you are still on location: thirty seconds. Noticed the same evening: a phone

call. Noticed three weeks later, after the person has gone back to another city: a rewrite, a compromise, or a reshoot nobody is paying for.

The audit is the fifteen minutes that moves that discovery earlier.

The audit itself

Write your beats down the page. Next to each one, write the specific asset that carries it — not a category, a file.

```
1  promise      - wide of the workshop at 7am, clip 0042
2  the problem  - Anil: "the machine jams on thick denim",
00:14:22
3  but          - MISSING. need denim actually jamming
4  therefore    - hands re-threading, clips 0071-0074
5  the change   - Anil: "we take the jobs we used to turn
down", 00:31:05
6  cost         - MISSING. what did the machine cost him
7  ask          - end card, not made yet
```

Every line is either a file reference with a timecode or the word MISSING. "I think I have something for that" is MISSING. "There's definitely a shot of the workshop" is MISSING until you have found it and written the number down.

What this produces is a short, specific list of things to obtain, which is a very different object from an anxious sense that the material might not be enough.

Two kinds of gap, and they are not equally serious

- **An asset gap** is a beat with no material. It is recoverable: film it, photograph it, find it, or write around it.
- **A structure gap** is a beat that nothing could carry, because the thing it claims did not happen. Beat six asks what the machine cost him, and the honest answer is that it was a routine purchase with no drama in it. No amount of footage repairs that; the structure was wishful.

The audit is the first moment these two look different. Before it, both feel like "I need more material". After it, one of them is a shopping list and the other is

a rewrite, and confusing them is how people end up shooting for two more days to prop up a beat that should have been deleted.

Shoot the structure plus the repair material

An edit needs more than the structure strictly requires, for mechanical reasons rather than artistic ones:

- **Cutaways to hide joins.** Every time you cut two pieces of the same interview together, you need something to cut to, or the subject's head jumps.
- **Handles.** Ten extra seconds at the start and end of each shot, so a transition has somewhere to happen.
- **One alternative for anything fragile.** If the whole piece turns on one sentence, get that sentence twice.

A reasonable planning figure for a first piece is that you will use about a fifth of what you gather, and that the fifth you use will not be the fifth you expected.

The same audit for still work

Nothing here is specific to film. A four-page brochure has beats; each one needs an image, a number or a quote, and every one of those is either in hand or is not. A slide deck built the same way stops producing the familiar disaster where slide nine has a title and nothing underneath it. A pitch is a structure with evidence attached to each turn, and the evidence is either found or invented, and the audit is where you find out which.

Where the audit turns into hoarding

The honest limitation is that gathering more material has a rising cost that is easy to miss, and the cost falls on the edit rather than on the shoot.

Selecting between takes is comparison work, and comparison work does not scale with the number of takes; it scales roughly with the number of pairs you might compare. Three takes of a line is three comparisons. Ten takes is forty-

five. The shoot feels more productive with every additional take, and the edit gets disproportionately slower, which is why experienced editors ask for fewer options rather than more.

So the audit has a second half, done after gathering: go back through and mark what you are **not** going to use. Material that is not in the list is not in the project. It can live in a folder called `unused` and never be opened again. That folder is doing real work — it is the reason you can find the shot you need in ninety seconds rather than forty minutes.

BEFORE YOU MOVE ON

Why does gathering ten takes of a key line rather than three slow the edit by more than three times?

- A. Longer shoots produce tired performances, so later takes are usually worse and have to be discarded individually.
- B. Choosing means comparing takes against each other, and the number of possible comparisons grows far faster than the number of takes.
- C. Storage and transcoding time rise with the amount of footage, which delays every subsequent step.
- D. Editors lose confidence when given many similar options and begin re-cutting decisions they had already settled, so the delay compounds through every later section of the piece.

CASE STUDY · MODULE 4

The detail that made the story mean something, and identified her

A small legal-aid organisation in Muzaffarpur, making a four-minute film about informal lending, with one woman willing to speak on the record.

She agreed on one condition, stated plainly and more than once: she must not be identifiable. The organisation did what most people do. It changed her name, shot her in silhouette, and shifted the pitch of her voice.

Then somebody went through the rough cut detail by detail and counted, for each one, roughly how many people it still fits.

- The town: tens of thousands.
- A woman running a tailoring unit with two machines: a few hundred.
- Whose husband works at the sugar mill, which was named in a wide shot: perhaps a dozen.
- Whose daughter is in class nine: one, to anybody in two wards.

Not one of those details is her name. Together they publish an identity, and this is how anonymity actually fails. The name was never the identifier.

Three of the protections were weaker than they looked

The pitch shift is a reversible operation and undoing it is a menu item in free software. The silhouette held for most of the interview and slipped for two frames when she turned her head, and two clean frames in a file anyone can step through is a face. And the photographer's stills, shot on a phone and sent for the accompanying article, still carried their EXIF block, which contains the coordinates of the room.

The robust protections are all subtractive. Crop the face out of the frame rather than covering it. Replace the audio rather than processing it. Remove the detail from the shot. What is not in the file cannot be recovered from the file.

The decision, and what it cost either way

Here is where it stopped being technical. The whole point of the film was a mechanism: the interest on her loan was being repaid in stitched garments rather than in cash, at a rate she had never been shown in figures. Nothing else in four hours of material carried that. It was the finding.

And the tailoring detail is what made her identifiable. Take it out and the film becomes a general piece about debt that a viewer has seen before, of no use to the people it was made for. Leave it in and the protection is decorative.

A third option was considered and rejected: rebuild the sequence as illustration. That removes the identification entirely and also removes the evidential claim, which in a film about a practice people deny exists is most of the value.

The organisation had also, at the planning stage, agreed a withdrawal policy and told her about it, which is the cheap decision that turned out to matter most later.

The release that would not have covered it

The form she had signed was the organisation's standard one, and it named a documentary film. It did not name the accompanying article, the social cut, the conference screening, or the fact that her words would be edited into four minutes out of four hours.

Nobody in that situation denies signing. The complaints that arrive months later come from people who agreed to one thing and found another, so the part of a release that fails is the scope rather than the signature. The form was rewritten in ordinary sentences before publication and she was asked to agree to it again: what was recorded, what it is for, where it may appear, for how long, that it will be cut, and that she may ask for it to stop.

The audit that found the hole in the structure

The beats were listed down a page with the specific file that carries each one beside them, with a timecode, or the word MISSING. Two lines came back MISSING.

One was an asset gap: there was no shot of a ledger, and one could be filmed. The other was a structure gap. Beat six claimed a moment when she realised the arithmetic, and there was no such moment, because she had known for two years and had no alternative. No amount of filming repairs that. The beat was wishful, and it was deleted rather than propped up, which is the distinction the audit exists to make visible.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

They showed her the cut before publishing and asked whether she felt exposed. She found a detail nobody on the team had: the mill hooter is audible in the room tone, and it places both the location and the time of day.

The final version replaced the pitch-shifted voice with a colleague reading her words, and said so on screen. It kept the payment-in-garments mechanism, because that practice is not unique to her and is the thing the film exists to show. It removed the two machines, the mill, the daughter's class and the room tone, and replaced them with a garment trade and a general reference to a school-age child. The stills went through exiftool -all= before upload.

She was also told, before she agreed to any of it, the sentence she was entitled to hear: identification can be made harder and it cannot be made impossible or revocable, because a public file can be copied, archived and searched by tools that do not exist yet.

Four weeks after publication she asked for one line to be cut. It was cut from everything the organisation controlled within two days, because the policy had been agreed at the planning stage and nobody had to invent it under pressure.

The name was changed, the face was hidden and the voice was altered. Why was she still identifiable?

Anonymity fails by combination rather than by name. Each detail narrows the population it could describe, and the narrowing multiplies: a town of tens of thousands, a specific trade, a named employer and a child's school year intersect at one household. The useful method is to walk the piece detail by detail and ask how many people each one still fits, and to treat single figures as publishing an identity.

Why are cropping and re-recording described as stronger protections than blurring and pitch-shifting?

Blur, pixelation and pitch shift are transformations applied to information that remains in the file, and transformations can often be run backwards or can leak: a tracked blur slips for two frames, a pitch shift is undone from a menu, and pixelated text has been recovered repeatedly. Cropping and replacing the audio are subtractive, and what is not in the file cannot be recovered from it.

Agreeing a withdrawal policy in advance cost nothing at the time. What did it buy?

Two things. Before recording, it made her far more willing to take part, because the fear people actually hold about speaking is permanence rather than publication. Afterwards, it meant the request to cut a line was handled as a routine step within an agreed period, by people who were not deciding under pressure what they owed her, which is when organisations make decisions they later defend badly.

MODULE 4 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 You can show a client six rough boxes on a sheet of A4 or a polished first draft. Why does the rough version get better feedback?
 - A. A rough board costs less to make, so less is lost when it is rejected.
 - B. Clients cannot judge typography, so leaving it out avoids irrelevant comments.
 - C. The rough version conceals the weaknesses that a finished draft would expose to a client far too early.
 - D. People give braver feedback on something visibly unfinished and save politeness for finished things.

- 2 You ask how long somebody has been doing the work and they say about eleven years. In the edit the clip is unusable. Why?
 - A. Eleven years is not a specific enough detail for a viewer to believe.
 - B. Spoken numbers are hard for a listener to retain without a caption.
 - C. Your question is absent from the finished piece, so the answer has to stand alone.
 - D. The clip is too short to cut cleanly against the material on either side of it in the edit.

- 3 A voice says we do everything by hand over a cutaway that happens to include an industrial machine in the background. What does the viewer take away?
 - A. The picture is believed and the claim is quietly rewritten as untrue.
 - B. The inconsistency is noticed and raised, usually in the comments.
 - C. The two channels reinforce each other, and recall of the claim improves.
 - D. The shot is doing continuity work, so a small mismatch is acceptable.

- 4 Somebody who signed a release complains months later that they are in a recruitment advertisement. They do not deny signing. Which part of the release failed?
- A. The signature, because it was collected informally on a phone photograph.
 - B. The scope: what it is for, where it appears, for how long, and that it may be edited.
 - C. The language, because the document was not drafted by a lawyer.
 - D. The date, because permission of this kind expires after a fixed statutory period in most countries.
- 5 You want a Creative Commons photograph for a client's commercial brochure. Which clause can reach into the brochure itself and change what it must become?
- A. BY, because the required credit must appear on every copy that is distributed.
 - B. ND, because laying a title over the image counts as a modification.
 - C. NC, because a brochure promoting a business is a commercial use.
 - D. SA, because a derivative work must carry the same licence as the original.
- 6 A truncated vertical axis is sometimes defensible on a line chart and never on a bar chart. What is the reason for the difference?
- A. A bar encodes quantity by length, so a cut-off bar draws the wrong length.
 - B. Zero is always meaningful in the kind of data that bar charts are used for.
 - C. Bar charts are read by comparing the tops of the bars, not their full extent.
 - D. Charting software cannot label a truncated axis correctly on a bar chart.

MODULE 5

Words that carry the picture

Almost every visual piece contains words — a headline, a caption, a subtitle, a voiceover, the line before the fold. They are usually written last and badly. This module covers what a headline is for, where text is cut off, how writing for the ear differs from writing for the eye, how to caption properly, and how to cut your own work.

BY THE END OF THIS MODULE YOU CAN

Write the words around a visual piece so they do the work the pictures cannot — a headline that earns the first line, an opening that survives truncation, narration that can be followed without rewinding, subtitles inside real reading-speed limits, a caption that adds rather than restates, and a draft cut by a third without losing an idea

33 • 8 min

A headline is not a summary

THE ONE THING TO KEEP

A headline's only job is to get the next line read, so favour photographable nouns, put the most specific word first because text is always cut from the right, and write twenty before keeping one — but never promise something the piece does not pay.

What it is competing against

A headline is not read in isolation. It is read in a feed, an inbox, a search page or a poster site, alongside a dozen other things making the same bid, by some-

body who has already decided to skip most of them. It has one job, and the job is not to describe the piece.

Its job is to get the next line read. That is the whole of it. The headline that describes the piece accurately and is not read has failed completely; the headline that gets the first line read has succeeded even if it left things out, provided the piece then pays what the headline implied.

Once you accept that, a lot of common headline advice reorganises itself. "Be clear" is not a value in itself, it is a tactic that usually works because confusion is a reason to move on. "Be intriguing" is the same tactic from the other side and fails when the intrigue is unresolvable.

Specific nouns do the heavy lifting

Compare:

Improving Operational Efficiency in Small Manufacturing

The £90 tool that cut our reject rate in half

The second is not better because it is punchier. It is better because it contains three things a reader can picture — a price, a tool, a defect — and abstractions cannot be pictured. "Operational efficiency" is a category that contains thousands of things, so the reader has to do work to find out whether any of them concern them, and they will not do that work.

The practical test: underline every noun in your headline and ask whether you could photograph it. If none of them can be photographed, you have written a category, not a headline.

The two failures at the edges

Too vague is the commoner one, and it usually comes from a fear of being wrong. "A new approach to storytelling" commits to nothing and therefore attracts nobody.

Too clever is the rarer and more painful one. A pun that requires reading the piece to understand is a locked door with the key inside. Wordplay works

when it is readable both ways — funny if you get it, plain if you do not.

Between them sits the version that is specific and slightly incomplete: it names the thing and withholds one piece, and the withheld piece is genuinely in the article.

Front-load, because the end may not arrive

Text is cut off in more places than you expect, and always from the right. Search results truncate. Email subject lines truncate, more aggressively on a phone than on a desktop. Link previews truncate. Poster text gets read at an angle by somebody walking.

So the distinguishing word goes at the front, not at the end.

- "How we cut our reject rate in half using a £90 tool" → truncated, reads "How we cut our reject rate in..."
- "The £90 tool that halved our reject rate" → truncated, still reads "The £90 tool that halved..."

The second survives because its most specific element is in the first four words.

Write twenty, keep one

This is the part people skip and it is the part that works. Headlines are cheap to produce and expensive to get right, which is exactly the shape of problem that rewards volume. Set a timer for ten minutes and write twenty, deliberately including bad ones — the literal one, the question one, the number one, the too-clever one, the flatly boring one. The good one is usually somewhere around fifteen, after you have exhausted the obvious moves.

Do not do this at the start. Write the piece first. A headline written before the piece describes the piece you intended; a headline written after describes the one you made, and those are different documents.

The limitation, which is about incentives

Everything above optimises for getting read, and getting read is a measurable thing, which means it can quietly become the objective. This goes wrong in a specific way: the headline formats that test best in a single measurement are the ones that create an unresolved question, and an unresolved question can be created without having an answer. Once a piece is judged only by whether the headline was clicked, the cheapest way to improve is to promise more than the piece contains.

The cost does not appear in the click figure. It appears later, in whether anybody reads the second piece, and by the time it shows up in the numbers you have trained an audience to discount you. The discipline is simple to state and hard to hold: **the headline may withhold, but it may not promise something the piece does not deliver.**

BEFORE YOU MOVE ON

Why does "The £90 tool that halved our reject rate" survive truncation better than "How we halved our reject rate using a £90 tool", when both contain the same facts?

- A. Starting with a definite article rather than a question word signals a statement, which search engines and inboxes rank more prominently in limited space.
- B. Truncation always removes text from the right, so the most distinguishing element has to sit in the opening words to survive.
- C. Shorter headlines are less likely to be truncated at all, and the first version has fewer characters.
- D. Readers scan headlines for concrete nouns and skip function words, so any sentence containing a price will be noticed regardless of where the price sits.

Everything you write has a fold in it

THE ONE THING TO KEEP

Every surface truncates and offers a more link, and the reader decides at that cut — so make the pre-fold text a complete statement rather than a tease, delete the throat-clearing, and check the cut on a real phone rather than in a preview.

The fold is a decision point, not a design detail

A "fold" was originally the crease in a folded newspaper: what appeared above it was what somebody saw on the stand, and what appeared below it was what they saw if they bought the paper. The word survived into web design, got declared dead several times, and then quietly became more important than ever, because nearly every surface now truncates text and offers a **more** link.

The important thing is not the exact position. It is that at the fold, a reader makes a decision, and the decision is cheap to answer with *no*. Whatever is above it has to make continuing worth a tap.

Where the folds actually are

These numbers move — platforms change them without announcement, and they differ between the web, the apps and the various phone sizes — so treat them as orders of magnitude rather than specifications, and check your own posts on your own phone.

- An **Instagram caption** shows roughly the first 125 characters before *more*.
- An **email subject line** is often cut around 40 characters on a phone and around 60 on a desktop client, and the preview text after it is a second, separate fold that most senders waste on "View this email in your browser".

- A **search result description** tends to be cut somewhere near 155–160 characters, and search engines frequently rewrite it anyway.
- A **link preview card** shows a title of a few words and a description of one or two lines.
- A **long page** has a soft fold at the bottom of the first screen, which on a phone is perhaps 300 words less than you think.

Write the first sentence as though it is the whole thing

The reliable move is to make the pre-fold text **complete rather than teasing**. A tease says: there is something here, tap to find out. A complete opening says: here is the thing, and there is more.

Tease: "You won't believe what we found in the old machine shop..."

Complete: "We found a 1954 lathe under a tarpaulin, and it still runs."

The second gives away the entire premise and gets more people through the fold, which looks backwards until you consider what each one asks of the reader. The tease asks them to spend a tap on trust. The complete opening asks them to spend a tap on interest they already have. Only one of those is a thing a stranger has.

The throat-clearing problem

Most people's first two sentences are them getting started. "So, we've been thinking a lot recently about the way small workshops handle..." That is warm-up, and it lands exactly where the space is most expensive.

A mechanical fix that works on nearly every draft: **delete your first paragraph and see whether anything is missing**. Usually the piece starts better without it, because the second paragraph is where you stopped introducing and started saying something.

The email preview line is free space almost nobody uses

Every mail client shows a snippet after the subject, taken from the first text in the message unless you set a preheader. Most senders therefore advertise their unsubscribe link or a browser fallback. Setting that line deliberately gives you a second headline of roughly equal length, at no cost, and the two can work as a pair — the subject naming the thing, the preview adding the detail that makes it matter.

Test it the only way that is real

Open your own post on your own phone, in the app, not in the editing preview. The preview lies, because it is rendering at a different width with a different font and no interface furniture over it. This takes ninety seconds and catches the specific, embarrassing failure where the cut lands in the middle of the one word the sentence depended on.

Where the advice breaks down

Writing for the fold is optimisation against a constraint you do not control, and it can flatten your work into sameness. Every post opening with a complete, declarative, slightly punchy sentence is recognisable as a style, and once an audience recognises a style as a technique, the technique stops working on them.

There is also a category of piece for which this is simply wrong. A long essay, a piece for people who already follow you, a piece whose value is in accumulation rather than in a single claim — these can open slowly, and their readers arrived already having decided. The fold matters in proportion to how much of your audience is strangers. If you are writing for people who came looking for you, you are allowed a first paragraph.

BEFORE YOU MOVE ON

Why does a complete opening line usually get more people past the fold than a teasing one, despite the tease creating curiosity?

- A. A tease asks a stranger to spend a tap on trust, while a complete opening asks them to spend it on interest they can already feel.
 - B. Platforms rank teasing language lower because it resembles engagement bait, so fewer people see the post at all.
 - C. Curiosity fades within a few seconds, so by the time the reader reaches the tap the tease has stopped working.
 - D. Complete sentences contain more searchable keywords, which makes the post easier to find and therefore delivers it to an audience that is already more interested in the subject.
-

35 • 9 min

A listener cannot go back

THE ONE THING TO KEEP

A listener cannot backtrack, so finish each grammatical thought before starting the next, round numbers for the ear and put the exact figure on screen, and read the script aloud to find the sentence where you run out of breath.

One difference explains all the others

A reader controls the pace. They can stop, re-read the clause they lost, look at the previous sentence, skim ahead. A listener can do none of these things without an act of will — finding the scrub bar, dragging it back, missing the next bit while they do it. In practice they simply lose the sentence and keep going, slightly less attached than before.

Everything below follows from that one asymmetry. Writing for the ear is writing that can be parsed **in one pass, in order, with no backtracking**.

The nested clause is the main offender

"The lathe, which had been sitting under a tarpaulin in the back of the shop since the previous owner retired in 1998, still runs."

On paper this is fine. Spoken, the listener holds "the lathe" in memory for nineteen words while a subordinate clause runs, and by the time "still runs" arrives they have lost the subject and have to reconstruct it. They do reconstruct it, mostly, and it costs them the first half of your next sentence.

Unpacked:

"The lathe still runs. It had been under a tarpaulin at the back of the shop since 1998, when the previous owner retired."

Same facts, same length, no held-open structure. The rule that generates this: **finish each grammatical thought before starting the next one.** Subject, verb, object, stop.

Numbers have to be rounded to be heard

"Thirty-one point four per cent" is five distinct pieces of information arriving in two seconds. Almost nobody retains it. "Nearly a third" is one piece, retained by everybody, and in a spoken piece it is more accurate in the only sense that matters — it is what the listener will end up believing.


Put the precise figure on screen if precision matters. The ear gets the shape, the eye gets the decimal.

The same applies to lists. Three items is the practical limit in speech. Beyond that a listener stops tracking and waits for the end.

Read it aloud, and time it

This is the whole quality-control process and it is not optional. Read the script aloud, standing up, at the pace you would actually deliver it. Three things happen:

Words of narration a running time can actually hold



At about 150 words a minute, a ninety-second film is roughly 225 words. A 400-word script is not a ninety-second film that needs tightening; it is a two-and-a-half-minute film you have not decided to make.

1. You run out of breath somewhere, and that is where the sentence is too long. Mark it and break it.
2. You stumble over a phrase, and a phrase you stumble over while reading your own writing is one a listener will stumble over too. Almost always it is a consonant cluster or an accidental rhyme.
3. You find out the actual duration. Clear narration runs at roughly **140 to 160 words a minute** — slower than conversation, because a listener has no facial cues and no ability to ask. A ninety-second piece is therefore about 210 to 240 words of narration, and that is far less than people expect. If your script is 400 words you are not making a ninety-second film, you are making a two-and-a-half-minute one, and you should decide that deliberately.

Write the contractions in

People say "it's", "we'd", "didn't". A script that says "it is", "we would", "did not" will either be read stiffly or silently contracted by the reader, and the second means your timing is wrong. Write what will be said.

Similarly, write the pauses. A full stop is a pause; a paragraph break is a longer one. If you need a specific beat, put it on its own line. Narrators follow layout more than punctuation.

Homophones and the things only the ear hears

Spoken text has failure modes that never occur in print. "Their", "there" and "they're" are identical. So are "our" and "hour" in many accents. "A grade A" is ambiguous aloud. Accidental rhymes and alliteration that look inert on the page become conspicuous when spoken, and a listener who notices the sound has stopped following the sense.

The free test for all of this is a text-to-speech read-back — every phone and operating system has one built in, and hearing your script in a flat synthetic voice is unusually good at exposing the sentences that only work because you know what they mean.

The limitation: ear-writing is not always better writing

This style — short sentences, front-loaded subjects, rounded numbers — is genuinely worse for some purposes. It cannot hold a qualified argument, because qualification is exactly what subordinate clauses are for. A legal position, a technical caveat, a claim that is true only under three conditions: these become misleading when flattened, because the flattening removes the conditions.

When you need precision that will not survive the ear, the answer is not to write a more complicated script. It is to put the precise version on screen or in the description and let the narration carry the shape. Two channels, each doing what it is good at, rather than one channel doing both badly.

BEFORE YOU MOVE ON

A script contains "The lathe, which had sat under a tarpaulin since the previous owner retired in 1998, still runs." Why does this fail when spoken but not when read?

- A. The listener must hold the subject open across a long clause, and cannot re-read to recover it when the verb finally arrives.
 - B. Dates spoken aloud take longer to process than dates on a page, so the year interrupts the sentence's rhythm.
 - C. Commas are inaudible, so the listener cannot tell where the subordinate clause begins and ends.
 - D. Spoken delivery is slower than reading, so a long sentence simply takes more time and the listener becomes impatient before the end arrives.
-

36 • 10 min

Most of your audience has the sound off

THE ONE THING TO KEEP

Captions are limited by reading speed rather than screen space — two lines, around 17 characters a second, broken at clause boundaries — and because a viewer reading text cannot study the picture at the same time, captioning is a structural decision rather than a final step.

Two different things with one name

- **Subtitles** assume the viewer can hear but does not understand — they carry the dialogue, often translated.
- **Captions** assume the viewer cannot hear — they carry the dialogue *and* the sound that matters: [machine starts] , [laughs] , the speaker's name when it is not obvious who is talking.

The distinction is not pedantry. A video whose meaning turns on a sound — a crack, a door, a change in a machine's note — is unintelligible with subtitles and intelligible with captions, and for a deaf viewer that is the difference between a piece that includes them and one that does not.

Then there is the third group, which is most people: viewers with working ears watching in a queue, in bed, in an office, on a train, with the sound off because the platform started it muted and turning it on is a decision. Captions serve all three at once, which is why they have gone from an accessibility duty to the default.

The constraint is reading speed, not screen space

The mistake beginners make is fitting the words on screen. The real limit is whether they can be read in the time they are there, while the viewer is also watching the picture.

The broadcast conventions exist because of this:

- **Two lines maximum.** Three lines covers too much picture and forces the eye to travel.
- **Around 37 to 42 characters per line.** Wider than that and the eye has to sweep, which costs a fixation.
- **Roughly 17 characters per second**, or about 160 to 180 words a minute, as a ceiling for comfortable reading. Faster is legible and exhausting.
- **A minimum of about one second on screen**, even for one word, because anything shorter registers as a flash rather than as text.

When the speech is faster than the reading limit, you edit the caption. Condensing speech into readable text is a real skill and it is not cheating; a verbatim caption that cannot be read has served nobody.

The caption limits, and what sets each one

Two lines maximum	Three covers too much picture
37 to 42 characters a line	Wider costs the eye a fixation
About 17 characters a second	Faster is legible and exhausting
One second minimum on screen	Shorter registers as a flash
Break at clause boundaries	Reading grabs phrase-sized chunks

None of these limits is about fitting the words on screen. They are all about whether the words can be read in the time they are there, while the viewer is also watching the picture.

Break lines where the sense breaks

This is the single largest quality difference between amateur and professional captions, and it costs nothing.

Bad:

```
the machine jams on thick denim which is
why we used to turn the work away
```

Better:

```
the machine jams on thick denim
which is why we used to turn the work away
```

Break at clause boundaries — before a conjunction, before a preposition, between a subject and a long predicate. Never split an article from its noun or a name across two lines. The reason is mechanical: reading proceeds by grabbing phrase-sized chunks, and a break in the middle of a chunk forces a re-read of both halves.

The file, and why a sidecar beats burning in

Subtitles live in a plain-text file, usually `.srt`, which is readable and editable in any text editor:

```

1
00:00:04,120 --> 00:00:06,800
the machine jams on thick denim

2
00:00:06,800 --> 00:00:09,400
which is why we used to turn the work away

```

A **sidecar** file like this can be switched off, translated, indexed by search engines and read by screen readers. **Burned-in** captions are painted into the picture and cannot be any of those things — but they are the only kind that survives being downloaded and re-uploaded, which is why social platforms are full of them.

The defensible answer for most people is both: burn in for the vertical social cut, supply an `.srt` everywhere that accepts one.

Automatic captions, and exactly where they fail

Automatic speech recognition is now good enough that starting from it and correcting is faster than typing from scratch. It fails in a predictable pattern, and knowing the pattern makes the correction pass quick:

- **Proper nouns** — names of people, places, businesses, products. The model has never seen your client's surname.
- **Numbers and units**, particularly currencies and measurements.
- **Accents outside the training distribution**, which in practice means most speakers outside a narrow band of British and American English. Published error rates differ by accent and dialect by a factor of several, and the direction of that difference is consistent enough to plan around.
- **Overlapping speech**, which it renders as a plausible single sentence that nobody said.

The last one is the dangerous failure, because it is fluent. It does not look like an error.

Free tools for the whole job: YouTube's automatic captions with its editor, Whisper if you can run it locally, Subtitle Edit on Windows and Aegisub cross-

platform for timing, CapCut and DaVinci Resolve's free version for burning in. None of this needs a subscription.

The honest limitation

Captions change what your piece is. A viewer reading captions is spending a share of their attention on text, and cannot look at the picture at the same time — the eye cannot be in two places. This means a caption over a shot that requires close looking is a genuine conflict, not a solved problem.

The consequence for the work is upstream: shots that need to be *studied* should not coincide with dense speech. That is an editing decision made at the structure stage, and it is one of the places where accessibility stops being a checklist item added at the end and becomes something that shapes the piece.

BEFORE YOU MOVE ON

Why does breaking a caption line after "which is" rather than after "denim" measurably slow reading?

- A. The second line starts with a lower-case word, which readers briefly interpret as a continuation error.
- B. It puts more characters on the first line, pushing it past the comfortable line-length limit.
- C. Reading proceeds in phrase-sized chunks, so a break inside a chunk forces both halves to be re-processed.
- D. Eye movement between lines is slower than movement within a line, so any two-line caption should be avoided where possible.

37 · 8 min

The caption decides what the photograph means

THE ONE THING TO KEEP

A caption either fixes which of an image's possible meanings applies or adds a fact the image cannot contain, and the second is what earns its space — while alt text is a separate job describing the image for someone who cannot see it.

A photograph is ambiguous until something fixes it

A picture of a man standing in an empty workshop at night can be a portrait of failure, of dedication, of solitude, or of a man who has locked himself out. The image contains all of these. What the reader takes from it is decided by the words next to it.

This gives the caption two genuinely different jobs, and conflating them is why most captions are wasted.

1. **Fixing the meaning.** The image has several possible readings and the caption selects one. "Anil, closing up after a 14-hour day" closes off failure and solitude and leaves dedication.
2. **Adding what the image cannot contain.** A photograph cannot show a date, a name, a price, a consequence or a cause. "The same lathe, four months after the workshop flooded" adds a fact that no arrangement of pixels could have carried.

The second is where captions earn their space. The first is necessary but cheap.

The wasted caption restates

A photograph of a lathe. Caption: "A lathe."

This is extremely common and it is a real cost, not merely a missed opportunity. The reader's eye goes to the caption — captions are among the most-read text on any page, because they are short and attached to something interesting — and it returns nothing. You have spent your most reliable piece of attention on a label.

The test is the same one used for images: **cover the picture, read the caption; cover the caption, look at the picture**. If the caption tells you nothing new, rewrite it.

Alt text is a different job again

Alt text is not a caption. A caption is for everybody. Alt text is read aloud by a screen reader to somebody who cannot see the image at all, and its job is to convey what a sighted reader gets from looking.

The differences in practice:

- **Describe, do not editorialise.** "Man in a workshop at night, standing beside a lathe, coat still on" — not "a moving portrait of dedication".
- **Do not begin with "image of" or "photo of".** The screen reader has already announced that it is an image; you are making the listener sit through it twice.
- **Length follows function.** A decorative image takes empty alt text (`alt=""`), which tells the screen reader to skip it — far better than making somebody listen to a description of a background texture. An image carrying information needs as much text as the information takes, which for a chart can be a paragraph.
- **If the caption already describes it,** the alt text can be short or empty, because otherwise the listener hears the same thing twice.

Every major platform now has an alt-text field, usually two taps away, and almost nobody fills it. It costs fifteen seconds.

Numbers, names and dates belong in the caption

A reliable habit for documentary and commercial work alike: every caption contains at least one **checkable** thing — a name, a date, a place, a quantity. This does two things. It forces the caption past restatement, and it puts the specifics somewhere they survive, because a caption is what gets quoted, indexed and remembered after the image itself has been separated from the article.

The order people read in

On a page with a headline, an image, a caption and body text, readers do not go top to bottom. A common path is image, then caption, then headline, then the first line of body text, then out. That is one reason captions are worth more than their size suggests: for a meaningful share of your audience, the caption is not a supporting detail, it is the second thing they read and sometimes the last.

Write them as though that is true, because for some readers it is.

The limitation worth naming

Fixing an image's meaning with words is a power that can be used dishonestly, and it does not require lying. A true caption placed against a true photograph can produce a false impression, because the caption directs the reader towards one reading and away from another that is equally supported by the picture.

A protest photograph cropped to a confrontation, captioned truthfully with the date and place, will read as a riot; the same photograph captioned with the size of the crowd will read as a large peaceful gathering with an incident in it. Nothing in either caption is false.

There is no technique that resolves this, because there is no neutral caption. What you can do is ask the question directly: **what would somebody who was there say about this pairing?** If the answer is that they would feel misrepresented, you have your answer, regardless of whether each individual word could be defended.

BEFORE YOU MOVE ON

Why is empty alt text (`alt=""`) on a purely decorative background image better practice than describing it?

- A. Screen readers ignore images that have no meaningful content anyway, so the attribute value makes no practical difference.
 - B. It tells the screen reader to skip the image, sparing the listener a description that carries no information.
 - C. Empty alt text keeps the page's accessibility score high, since automated checkers penalise long descriptions.
 - D. Search engines ignore decorative images, so any text placed there is wasted effort that could be spent on the caption.
-

38 · 8 min

Your first draft contains the route, not the destination

THE ONE THING TO KEEP

A first draft records the thinking rather than the conclusion, so a third can usually go — delete the opening paragraph, run the mechanical passes, and cut against a stated limit — but stop before you remove the redundancy that lets a distracted reader catch up.

Why drafts are long

Writing is how most people work out what they think. The consequence is that a first draft is a record of the thinking — the approach, the false start, the qualification you added while deciding, the restatement that was you checking you still agreed with yourself.

None of that is useful to a reader, who wants the conclusion and the reasons, not the journey to them. This is why cutting your own work is so productive

and so uncomfortable: the material you are removing is not bad, it is just yours rather than theirs.

A workable target for a first cut is **a third gone with nothing lost**. People disbelieve this until they try it.

Delete the first paragraph

Start here, because it works almost every time. Read your opening paragraph, delete it, and read what is now the first sentence. In most drafts the piece starts better, because paragraph one was throat-clearing — context nobody asked for, a scene-set, a definition, a *what I want to talk about today*.

The same is true of the last paragraph surprisingly often, which is usually the writer summarising what they just said because they are reluctant to stop.

Six mechanical passes

These are things you can do without judgement, which is what makes them possible on your own work.

1. **Nominalisations back into verbs.** "We made a decision to undertake a review of" → "we decided to review". Look for words ending in *-tion*, *-ment*, *-ance*; each one is usually a verb wearing a coat.
2. **Adverbs before strong verbs.** "Completely destroyed", "absolutely essential", "very unique". If the verb is right the adverb is padding; if the adverb is doing the work, the verb is wrong.
3. **"There is" and "it is" openings.** "There is a problem with the timing" → "the timing is wrong". This one alone cuts several per cent from most drafts.
4. **Qualifier stacks.** "Somewhat quite often", "may possibly be able to". Keep one hedge where you genuinely mean it; delete the rest, because a stack of hedges reads as a lack of confidence rather than as precision.
5. **Repeated ideas in different clothes.** Search your draft for two sentences that make the same point. Keep the more concrete one.

6. **The word "that"**. Delete every instance and put back only those whose absence broke the sentence. Usually about half.

Read the last sentence first

A trick for finding structural flab rather than word-level flab: read your paragraphs from the bottom of the piece upwards, one at a time. Out of sequence, you cannot lean on the momentum of your own argument, and a paragraph that only exists because of the one before it becomes obvious — it has no content of its own to stand on.

Cut against a real limit

Abstract editing drifts. Editing against a number does not. Set the limit before you start: 240 words for a ninety-second narration, 125 characters before the fold, 40 characters on a subject line, a single side of A4.

A hard limit changes the question from "is this good?", which has no stopping point, to "what earns its place in 240?", which does. It is also the mechanism behind the old observation about not having time to write a short letter: brevity is not the absence of effort, it is the product of a second pass that a deadline usually eats.

Read it aloud, again

The same test as for scripts finds different things in prose. Sentences you cannot say in one breath are too long. Sentences where you speed up are usually padding — your own voice skipping ahead to the part that matters. Mark those and cut to the part you sped up towards.

Where cutting goes wrong

There is a failure at the far end, and it is real. Compression removes redundancy, and redundancy is partly what makes text comprehensible: the restatement that lets a reader who missed a clause catch up, the example that makes an abstraction concrete, the transition that says how this paragraph relates to the last one. Strip all of it and you get prose that is technically shorter and ma-

terially harder, where every sentence is load-bearing and missing one means missing everything.

This matters more for a reader who is tired, distracted, reading on a phone on a bus, or reading in their second or third language. Dense is not the same as clear.

The distinction to hold is between **repetition that helps** and **repetition that is you hearing yourself**. An example that makes an idea concrete earns its place. A second statement of an idea already stated does not. The test is whether a reader who understood the first instance gains anything from the second.

BEFORE YOU MOVE ON

A writer cuts every example and transition from a draft, leaving only the load-bearing claims. Why does the result often become harder to read despite being shorter?

- A. Removing redundancy leaves no slack, so a reader who loses one sentence has no way to recover the thread.
- B. Shorter paragraphs increase the number of visual breaks, which interrupts reading rhythm and reduces comprehension.
- C. Examples carry the emotional content, and readers disengage from text that does not move them.
- D. Compressed prose uses lower-frequency vocabulary because the writer reaches for precise single words, which slows recognition.

39 · 8 min

You cannot un-know the thing you are explaining

THE ONE THING TO KEEP

Knowing something changes how the signal seems from the inside, so you cannot audit your own work for clarity — write down what the reader already knows, define in passing on first use, and show it to one person who genuinely does not know.

The tapping experiment

In a study run at Stanford at the end of the 1980s, people were asked to tap out the rhythm of a well-known song on a table while a listener tried to name it. Before each attempt the tapper predicted how likely the listener was to get it. Tappers predicted around **50 per cent**. The actual figure was about **2.5 per cent** — three correct out of a hundred and twenty.

The tappers were not stupid or arrogant. They could hear the song while they tapped. It was playing in their heads, complete with melody and words, and they genuinely could not experience the tapping as the listener experienced it: a series of unrelated knocks.

This is the curse of knowledge, and the important part is that **it is not a failure of effort**. Knowing something changes what the signal sounds like from the inside, and no amount of trying harder restores the earlier state.

How it shows up in visual work

- A diagram that makes sense because you know what the boxes are.
- A video that opens on the detail, because the context is obvious to you.
- An abbreviation used on first appearance.

- A before-and-after where the difference is invisible to anybody who does not know what to look at.
- A chart whose axis label is a field term.
- A piece that answers the question the client asked you rather than the question the audience has.

The last one is the most expensive, and the hardest to see, because it is the whole brief being pitched at the wrong altitude.

Three fixes that work, in increasing cost

Write down what the reader already knows. Three bullet points, literally on the page, before you start: *they know what a lathe is; they do not know what a tolerance is; they have never heard of this business.* This converts an unanswerable question into a list you can check the draft against. It is crude and it catches most of the damage.

Define on first use, once, in passing. Not a glossary box, not a paragraph of background — a clause. "The tolerance, which is how much error the part is allowed, was a tenth of a millimetre." A reader who knew skims past it in half a second; a reader who did not is now in the piece. The cost of including it is nearly zero and the cost of omitting it is the whole paragraph.

Show it to one person who does not know. Not a colleague, not another designer. Somebody from outside, given the piece with no preamble, then asked to say back what it was about. The gap between what they say and what you intended is a measurement, and it is the only one available. One person catches most of it; three catches nearly all.

The opposite failure exists, and it is worse socially

Over-explaining is also a failure, and it is more damaging to the relationship with the reader than under-explaining, because it carries an implication about them. A piece that defines terms its audience uses daily reads as condescension, and condescension makes people stop reading in a way that mild confusion does not — a confused reader might persist, an insulted one will not.

The resolution is not to aim for the middle. It is to be specific about who you are writing for, which is the brief again: **one named reader, with a stated starting position.** "Somebody who runs a small workshop and has never commissioned a film" tells you to define *storyboard* and not to define *lathe*. Without that sentence you are guessing in both directions at once.

The honest limitation

This cannot be fully fixed, only managed, and the reason is the mechanism itself. You cannot audit your own writing for the curse of knowledge, because the audit is performed by the same mind that has the knowledge. Everything you check will look clear to you, and it will look clear precisely because of what you know.

This is why the third fix — an actual person who does not know — is different in kind from the first two rather than merely better. The first two are you predicting somebody else's state, which is the thing you are demonstrably bad at. The third is measurement.

It also means the problem gets worse over a project, not better. On day one you nearly remember not knowing. By week six you are as far inside it as your client is, which is exactly why the piece you show them will make sense to both of you and to nobody else.

BEFORE YOU MOVE ON

Why is showing a draft to one uninformed outsider fundamentally different from carefully re-reading it yourself, rather than just more thorough?

- A. Outsiders read more slowly and therefore notice omissions that a fast re-read skips over.
- B. Writers become attached to their own phrasing, so emotional investment prevents them from making cuts an outsider would make.
- C. A self-review is performed by the mind that holds the knowledge, so the very gaps you are looking for are invisible to it.
- D. A single outside reader provides a sample, and any sample is more reliable evidence than an expert judgement about a general audience.

40 · 9 min

Most people reading your English did not grow up in it

THE ONE THING TO KEEP

The cost for a second-language reader sits in phrasal verbs, idiom, stacked negation and elegant variation rather than in complex ideas — so use the same word for the same thing, prefer single-word verbs, and treat readability scores as a smoke alarm rather than a target.

The arithmetic

There are roughly four hundred million first-language English speakers and somewhere between one and one and a half billion people who use English as a second, third or working language. For anything published on the open internet — and certainly for anything published from or about South Asia, Africa, or most of Europe — the majority of your English-reading audience learned it in a classroom.

They are not less intelligent readers and frequently they have better formal grammar than native speakers. What differs is specific and predictable: **idiom, phrasal verbs, low-frequency vocabulary, and culture-bound reference.**

Phrasal verbs are the single biggest cost

A phrasal verb is a verb plus a particle whose meaning cannot be derived from the parts: *put off, get by, take on, run into, come across, work out, turn down, hold up.*

English has thousands of them, they are extremely common in casual writing, and they are individually unguessable. *Turn down* means reject, reduce, or fold, depending. A reader who knows *turn* and knows *down* has no route to

any of these. Worse, they are not obviously hard — the words are all short and common, so the sentence looks simple and simply will not parse.

Phrasal verbs, and the single words that replace them

Unguessable from the parts

- turn down the order
- the machine gave out
- work out the cost
- come up against a problem
- put off the decision

Survives a second language

- refuse the order
- the machine failed
- calculate the cost
- meet a problem
- delay the decision

Every word on the left is short and common, which is why the sentence looks simple and will not parse. The difficulty is in the language, not in the idea, so simplifying the argument is the wrong lever.

The fix is not to avoid them entirely, which produces stilted text. It is to prefer the single-word verb **where the phrasal one is doing the load-bearing work of the sentence**:

- "we had to turn down the order" → "we had to refuse the order"
- "the machine gave out" → "the machine failed"
- "work out the cost" → "calculate the cost"

Four more things that cost disproportionately

1. **Stacked negation.** "Not unlike", "it is not uncommon for this to fail to work". Each negation must be resolved, and they compound. Say the positive.
2. **Culture-bound reference.** Cricket, baseball, a 1970s advertisement, a politician, a sitcom. These are invisible to the writer because they feel like general knowledge. If a reference is load-bearing, make the sentence work without it.
3. **Elegant variation.** Native writers are taught not to repeat a word, so a piece calls the same object a *lathe*, a *machine*, a *tool* and *the unit*. To a second-language reader this reads as four things. **Use the same word for the same thing, every time.** It is worse prose and much better communication.

4. **Numbers written as words.** "Two thousand four hundred" is harder than "2,400" for almost everyone and much harder across languages with different grouping conventions. Use digits. Also state currencies explicitly — ₹, £, \$ alone is ambiguous across several countries.

What does not help

Readability formulas. Flesch–Kincaid and its relatives count syllables per word and words per sentence. That is all they count. They therefore reward chopping a clear sentence into three fragments and punish using the correct technical term because it has four syllables. A text can score as "easy" while being incomprehensible, and a well-structured explanatory paragraph can score as difficult while being the clearest thing on the page.

Use them as a rough smoke alarm for drafts that have drifted into 60-word sentences. Do not use them as a target, because optimising a proxy is how you end up with text that is measurably simple and functionally useless.

Simplifying the ideas. This is the thing people do instead, and it is the wrong lever. The difficulty is in the language, not the content. A second-language reader can follow a complex argument perfectly well; they cannot follow *come up against*.

Translation, machine and otherwise

If you are publishing across languages, machine translation is now good enough to be useful and not good enough to be trusted with anything that matters. Its characteristic failures are exactly the categories above — idiom, ambiguity, and terms with a specific meaning in your field — plus one more: it produces fluent, confident text when it is wrong, so errors do not look like errors.

The free, practical path: machine-translate, then have one fluent speaker read it for sense, not for grammar. That one pass catches the reversals of meaning, which are the failures that actually damage you. Keep a short list of your key terms with their agreed translations, so the same thing does not get three names across a site.

The limitation

Everything here optimises for comprehension, and comprehension is not the only thing writing does. Idiom, rhythm and reference are also how a piece establishes that it comes from somewhere and belongs to someone. Text stripped of all of it reads like an instruction manual, and an instruction manual persuades nobody.

The judgement is about what each piece is for. Safety information, instructions, prices, anything where being misunderstood has a cost: strip it. A personal essay, a piece whose value is voice, work aimed at a community that shares the references: keep the idiom and accept that some readers will work harder. What you should not do is make that choice by accident, which is what happens when you never think about who is reading.

BEFORE YOU MOVE ON

Why does calling the same object a "lathe", then "the machine", then "the unit" cost a second-language reader more than repeating one word would?

- A. Each synonym is drawn from a lower-frequency band than the one before it, so the vocabulary difficulty accumulates steadily across the paragraph until the reader stalls.
- B. Machine translation renders each synonym differently, so translated versions of the text become inconsistent.
- C. Repetition aids memory, so varied wording makes the object harder to recall later in the piece.
- D. Without a shared assumption that variation is stylistic, the reader has to work out whether these are one object or several.

CASE STUDY · MODULE 5

The subject line that doubled opens and halved the next email

A coaching centre in Kota with 6,400 parents on an email list, announcing a scholarship test.

Last year's subject line was: your child may be missing out. It opened at 41 per cent against a list average of 22, and the counsellor who wrote it wanted the same shape again.

The objection to it was not taste. It was the next email. That one opened at 9 per cent, against the list's normal 19, and 84 people unsubscribed within two days. The cost of the first line did not appear in the first line's number. It appeared a fortnight later, in a place nobody was looking, and the centre had recorded both figures without connecting them.

What the fold was doing to the honest version

The replacement subject was written to be complete rather than teasing: scholarship test, 14 December: who can sit it and what it covers.

Opened on a phone, most clients cut the subject somewhere around forty characters. That left scholarship test, 14 December: who can, which still carries the two facts that matter, because the specific words were at the front. The version the counsellor drafted first, how your child can compete for one of sixty scholarship places, cut to how your child can compete for one of, which carries nothing.

The preview line after the subject was being spent on view this email in your browser. That is a second headline of roughly the same length, available at no cost, and it was set to carry the fee waiver and the deadline.

The body opened with two sentences of throat-clearing about the centre's commitment to accessible education. They were deleted and nothing was missing, which is the usual result.

The decision, and its cost on both sides

The honest subject was going to open lower this week. That was predictable and it was stated in writing before sending, along with the two numbers that would be used to judge it: opens on this email, and opens on the next two.

The cost of saying that out loud is that the admissions head is judged on this week, not on the following month, and an open rate that drops from 41 to something lower is a visible fact with the counsellor's name on it. The cost of not saying it is that the list is spent a little further, and a list is the one route to these parents that nobody else can switch off.

There was a second, quieter piece of work on the body. Most of these parents read English as a second or third language. The draft used test, exam, paper and assessment for the same event, which reads as four things to somebody parsing carefully; it used phrasal verbs in the load-bearing sentences, including put off and sit for; and it wrote numbers as words. The whole draft was then cut against a stated limit of 260 words, from 410.

Cutting against a number rather than against a feeling

The limit did most of the work. Asking whether a paragraph is good has no stopping point; asking what earns its place in 260 words does. The first paragraph went and nothing was missing. There is a and it is openings were rewritten. Three nominalisations became verbs. Two sentences made the same point in different clothes, and the more concrete one was kept.

One cut was put back deliberately. The draft had explained twice, in different words, that the test is free to sit. A second statement of an idea already stated is usually the writer hearing themselves, but this one was doing a different job: a parent reading on a phone between shifts who missed the first mention catches the second, and the fee is the single fact most likely to stop somebody registering. Compression removes redundancy, and some redundancy is what lets a distracted reader catch up.

The test that found the rest

The counsellor read the final version aloud, standing up. She ran out of breath in the second paragraph, which is where the sentence was too long, and stumbled on a phrase that turned out to be an accidental rhyme. Both were fixed in four minutes. Nobody had needed to have an opinion about either.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

The honest subject opened at 27 per cent: below last year's 41 and above the list's ordinary 22. The following email opened at 24. Eleven people unsubscribed.

Registrations came to 212 against 178 the previous year. The centre also sent an SMS for the first time that season, so the two cannot be separated, and the report said so rather than claiming the email.

The argument that persuaded the admissions head was not the registration figure, which is contaminated. It was the pair of open rates, presented together, next to last year's pair. One number went down and one went up, and the second one is the list itself.

The body changes were the part nobody discussed and the part with the clearest evidence: replies asking what the test covered fell from 61 the previous year to 9, which is a direct measure of how many people could not follow the first version.

The teasing subject line opened at nearly twice the rate. Why was it still the more expensive option?

Because the cost is not paid in the metric being optimised. A headline that creates an unresolved question can do so without having an answer, and once it is judged only by whether it was opened, the cheapest improvement is to promise more than

the message contains. The damage appears later, as a lower open rate on the next message and as unsubscribes, by which time the audience has learned to discount everything the sender writes.

Why does front-loading the specific words matter more than the total length of a subject line?

Because text is cut from the right, in more places than people expect, and the cut position varies by client and by device. A line whose distinguishing element sits in the first few words survives every truncation; the same information placed at the end disappears at the exact moment a reader is deciding. The fix is to reorder rather than to shorten.

Calling the same event a test, an exam, a paper and an assessment is normal English style. Why was it changed?

Elegant variation is taught to native writers as good prose and it costs a second-language reader directly: four names read as four things, and the reader spends effort reconstructing that they are one. Using the same word for the same thing every time is worse prose and much better communication, and the reply count showed the size of the effect.

MODULE 5 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 A headline is read in a feed, beside a dozen other bids for the same attention. What is its job?
 - A. To get the next line read.
 - B. To summarise the piece accurately for the reader.
 - C. To contain the term people will be searching for.
 - D. To state the strongest benefit the piece offers.

- 2 Two openings sit above the more link. One teases and one gives the premise away. The complete one gets more people through the fold. Why?
 - A. Complete sentences are indexed more favourably by search engines, which raises how often a post is shown.
 - B. Teases are detected by ranking systems and demoted as low-quality content.
 - C. A tease uses more characters, so less of the message fits before the cut.
 - D. A tease asks a stranger for a tap on trust; a complete opening asks for interest they have.

- 3 The lathe, which had been under a tarpaulin since the previous owner retired in 1998, still runs. On paper this reads well. Why does it fail spoken?
- A. Listeners find long sentences boring and stop paying attention partway through.
 - B. The subject is held open for nineteen words, and reconstructing it costs the next sentence.
 - C. Subordinate clauses cannot be punctuated clearly enough in a script for a narrator to deliver them.
 - D. A narrator runs out of breath, so the take has to be recorded again.
- 4 A caption is on screen for 1.4 seconds and contains 40 characters. What sets the limit you are working against?
- A. The width of the safe area in a vertical video frame.
 - B. The platform's maximum character count for a single subtitle cue.
 - C. Reading speed, about 17 characters a second, while also watching the picture.
 - D. The speaker's delivery rate, which a caption is obliged to match word for word throughout.
- 5 A photograph of a lathe is captioned A lathe. What would make the caption earn its space?
- A. Naming what is in the picture, so that no reader is left uncertain about it.
 - B. Adding a fact the image cannot contain, such as a date, a name or a consequence.
 - C. Repeating the headline, so the central message is reinforced on the page.
 - D. Describing the picture for a reader who is reaching it through a screen reader instead.

- 6 You have read your draft four times looking for anything a newcomer would not follow, and found nothing. Why is that weak evidence that it is clear?
- A. You are too close to the deadline for your judgement to be reliable.
 - B. Writers are trained to check grammar rather than meaning whenever they reread their own drafts.
 - C. Readability formulas measure the wrong thing, so no check is available to you.
 - D. The audit is run by the same mind that holds the knowledge, so everything looks clear.

MODULE 6

Type, colour and image as meaning

Typefaces, colours and pictures carry meaning before anybody reads a word, and most of that meaning is learned rather than intrinsic. This module covers what a typeface actually says and why, what makes type readable at small sizes, how colour meaning changes across cultures while legibility does not, designing for eyes unlike your own, and choosing between a photograph, an illustration and a diagram.

BY THE END OF THIS MODULE YOU CAN

Choose type, colour and image for what they say rather than for taste — justify a typeface by its associations and its script coverage, set a measure and line-height that can actually be read, keep meaning out of hue alone, meet a stated contrast ratio, and say what a photograph, an illustration and a diagram each assert that the others cannot

41 · 10 min

Your brand is not the logo, it is what you repeat

THE ONE THING TO KEEP

Write your type scale, colour roles, spacing unit and voice on one page with actual values, because in six weeks you will not remember them and the drift is what looks amateur.

Why one-person operations look temporary

Two businesses post the same number of times a week. One looks like a company. The other looks like a person trying things. The difference is almost never the logo. It is that the first one's posts share a typeface, a spacing rhythm and a way of talking, and the second one's do not.

Recognition is built by repetition. A reader who has seen your last four posts identifies the fifth before reading a word, and that identification is worth more than any single piece being better. It is also the cheapest advantage available to someone with no budget, because consistency costs nothing except a decision made once and written down.

The reason people fail at it is not laziness. It is memory. In six weeks you genuinely will not remember whether the orange was #E4572E or #E0522A, whether headings were 32 or 34, whether you used sentence case. So you re-decide every time, slightly differently, and the drift is invisible per-post and obvious across twenty.

The four parts of a small real system

Type. Two typefaces at most, and one is usually better. Then a scale — a fixed set of sizes, not whatever looks right. Pick a ratio and stick to it. At 1.25: 16, 20, 25, 31, 39, 49. At 1.5, which suits bolder work: 16, 24, 36, 54. You round

these to whole numbers and you never use a size that is not on the list. Google Fonts is free, licensed for commercial use, and downloadable for offline work.

A one-page system, and the part people leave out

Type	Two faces at most, and a fixed list of sizes
Colour by role	page, ink, muted, line, accent
Spacing	One base unit and its multiples
Voice	Three words you are, three you are not
NEVER	The mistakes you have already made once

Roles answer the question that actually stalls you, which is not which colours you own but when the orange is allowed. Accent means one element on a screen.

Colour by role, not by name. This is the change that makes a palette usable. Instead of "our orange and our navy", define five roles:

- **page** — the background
- **ink** — body text
- **muted** — secondary text, captions
- **line** — rules, borders, dividers
- **accent** — exactly one thing per screen

Roles answer the question that actually stalls you, which is not "which colours do we own" but "when do I use the orange". Accent means one element. They also make a rebrand a five-value edit instead of a sixty-value hunt.

Spacing. One base unit and multiples of it: 8 gives you 8, 16, 24, 32, 48, 64. The reason this matters more than it sounds is that the eye detects near-misses. A 22-pixel gap next to a 24-pixel gap does not read as variety, it reads as slightly wrong, and the reader downgrades the whole piece without being able to say why.

Voice. Three words you are, three you are not, and one real sentence rewritten to show the difference. "Direct, warm, specific. Not corporate, not jokey, not salesy." Then: *before* — "We are delighted to announce the launch of our new service offering." *after* — "We fix boilers on Sundays now."

Write it on one page

TYPE	Headings: Inter Bold. Body: Inter Regular. Sizes: 16 / 20 / 25 / 31 / 39. Nothing else.
COLOUR	page #FBF8F3 ink #1B1B1B muted #6E6A63 line #DED8CE accent #E4572E (one per screen)
SPACING	Base 8. Use 8 / 16 / 24 / 32 / 48.
VOICE	Direct, warm, specific. Not corporate, not jokey, not salesy.
NEVER	Drop shadows. Stretched type. Two accents. Stock photos of people shaking hands.

The NEVER list is the part people leave out and the part that saves you. It is where you record the mistakes you have already made once.

Keep it where you work — a pinned note, a text file in the project folder, a page in the app you already use. If checking it takes more than five seconds you will stop checking it.

Free tools that hold a system properly: **Figma**'s free tier supports styles and variables and runs in a browser. **Penpot** is open source, does the same job and can be self-hosted. Canva's free tier will let you save colours inside a design, but the full brand-kit feature has generally sat behind the paid tier, so check before you build a workflow on it. Honestly, a plain text file with the hex codes in it beats an unused tool, and on a phone a note pinned to the top of your notes app is a complete implementation.

The limitation worth knowing

A system written before you have made anything is fiction. You will invent rules for situations you have never met and then break them all within a month, which teaches you to ignore the document.

Make five real pieces first. Then lay them side by side, look for what you actually did rather than what you meant to do, and write down the parts that worked. A system derived from your own output is one you will obey, because it is a description before it is a rule.

The second limitation: consistency compounds only if the thing being repeated is worth repeating. Applied to a bad type pairing for a year, this lesson produces a recognisably bad brand with excellent recall. Get the individual decisions right first — that is design foundations — then lock them.

Do this today

Open your last five posts side by side. Write down every distinct type size and every distinct colour you find. If the list is longer than eight items, you have your answer, and building the one-page document above will take you about twenty minutes.

BEFORE YOU MOVE ON

A solo designer defines a palette as "our teal, our cream and our charcoal" but the work still looks inconsistent across posts. What is the most likely cause?

- A. Three colours is too few to build a recognisable identity, so the palette needs expanding with tints and shades.
- B. The colours are named as possessions rather than assigned to roles, so each post re-decides which colour is the background, the text and the single accent.
- C. The colours were chosen without checking contrast ratios, so they render differently across screens and appear inconsistent.
- D. The palette is fine but is not stored in a tool that enforces it, so the fix is moving the work into Figma or Penpot.

42 · 9 min

A typeface means where you have seen it before

THE ONE THING TO KEEP

A typeface means what the reader has been trained by exposure to expect, so associations are local and changeable — and before committing to one for multilingual work, check its script coverage, because a missing glyph renders as tofu.

Association, not essence

There is nothing in the shape of a letter that is inherently trustworthy, modern or friendly. What a typeface communicates is the accumulated weight of where the reader has encountered it: in newspapers, on legal documents, on cheap flyers, on a particular decade's advertising, on their school textbooks.

This matters because it makes the meaning **learned, local and changeable**, not a property you can look up. A typeface that reads as institutional authority in one country reads as nothing in particular in another, and as a specific unpleasant government in a third.

So the useful question is never "what does this font mean?" It is "where has my particular reader seen this before?"

The broad families and their usual freight

- **Transitional and old-style serifs** — Times, Garamond, Georgia and their free equivalents. Centuries of books and newspapers. They read as established, printed, considered. They are also what most people's eyes were trained on for long text.
- **Geometric sans** — Futura and its descendants, built from circles and straight lines in the 1920s. Modernist, deliberate, slightly cold. Because they were designed as an argument about modernity, they still carry it.

- **Neo-grotesque sans** — Helvetica, Arial, Inter. Deliberately neutral, which after sixty years of corporate use has become its own meaning: corporate neutrality is not the absence of a voice.
- **Slab serifs** — heavy rectangular serifs. Nineteenth-century posters, railways, industry. Sturdy and loud.
- **Scripts and handwriting faces** — personal, informal, handmade. They also read slowly, and set at small sizes or in capitals they stop being readable at all.
- **Display faces** — designed for one size and one job, usually large. A display face used for body text is the commonest visible amateur mistake.

The Comic Sans story is more useful than the joke

Comic Sans was drawn in 1994 for speech balloons in a Microsoft interface, where a child's cartoon dog explained the computer. Times New Roman in a cartoon balloon looked absurd; the new face fixed it. It was then shipped with Windows, became one of a handful of fonts everybody had, and ended up on funeral notices, hospital signage and government letters — contexts it was never for.

The backlash is not really about the letterforms. It is about the mismatch between a face built for a cartoon dog and the seriousness of where it appeared, which is precisely the point of this lesson.

There is a second part worth knowing. Its irregular, differentiated letterforms — where **b**, **d**, **p** and **q** are genuinely distinct shapes rather than reflections of one another — are the reason some dyslexic readers report finding it easier, and why several accessibility guides list it. The evidence for that is mixed and the effect is not large, but it is a real reason to not sneer reflexively at a typeface in a context you are not in.

The constraint nobody mentions: script coverage

If your work will ever carry Devanagari, Bengali, Tamil, Telugu, Arabic, Thai or any of a hundred others, the elegant Latin typeface you chose almost certainly has no glyphs for them. The reader gets **tofu** — the empty rectangle ☐

that a renderer substitutes for a character it has no shape for — or an ugly automatic fallback that does not match anything else on the page.

This is why the Noto family exists: a project explicitly aimed at covering every script in Unicode, with the name meaning *no more tofu*. It is free, open-licensed, and its sub-families are designed to work together, so Noto Sans Devanagari next to Noto Sans looks like one decision rather than two.

Before you commit to a typeface for anything multilingual, type a sentence of each script you need into the specimen and look. Thirty seconds now, versus a rebuild later.

Doing it free, properly

Google Fonts is licensed for commercial use, downloadable for offline and print work, and covers most scripts through Noto. The Open Font Library and Fontshare add more. Two rules that save real money and real trouble:

1. **A font downloaded from a font-sharing site is usually a commercial font being distributed without permission.** The person who gets the letter is your client, not the site.
2. **Check the licence for embedding** if the work becomes a PDF, an app or a video, because print and web permissions are sometimes separate.

The limitation

Typeface connotation is real but weak, and the published research on it is thin and inconsistent. What is reliably true is that a badly chosen face *damages* a piece — by contradicting it, by being unreadable, by looking like a default. What is much less supported is the idea that the right face carries a specific emotional payload into the reader on its own.

Treat type as something that can lose you the argument and rarely wins it. Get it out of the way of the message, make it readable, make it consistent, and spend your effort on what the piece actually says.

BEFORE YOU MOVE ON

A designer picks a refined Latin serif for a campaign that will later run in Hindi and Tamil. What is the most likely mechanical failure?

- A. Serif shapes derived from Latin calligraphy distort the proportions of Indic letterforms when scaled.
 - B. Serifs reduce legibility in scripts with a headline stroke, so the campaign will read as unprofessional in those markets.
 - C. Commercial font licences usually exclude non-Latin use, so the campaign would need a separate licence for each additional script region.
 - D. The face has no glyphs for those scripts, so text renders as empty boxes or falls back to an unrelated font.
-

43 · 9 min

Point size is the wrong number to argue about

THE ONE THING TO KEEP

Readability is set by x-height, measure and line height rather than by point size — 45 to 75 characters a line, 1.4 to 1.6 line height — and capitals cost reading speed because they flatten the word shapes readers use as cues.

Why two 12pt fonts look completely different sizes

Point size measures the body the letters sit on, a leftover from metal type, not the letters themselves. What the eye actually responds to is **x-height** — the height of a lowercase *x*, and therefore of the bulk of any lowercase text.

Two faces at the same point size can differ in x-height by a third. This is why a 12pt setting in one typeface is comfortable and a 12pt setting in another is squinting, and why "make it 14pt" is not a specification. The specification is: set a paragraph, print or view it at the real size and distance, and look.

A high x-height helps at small sizes and on screens. It costs something at large sizes, where the more varied proportions of a low-x-height face give a page more texture.


The measure: 45 to 75 characters

Line length is the setting that most reliably separates readable body text from unreadable body text, and almost nobody adjusts it.

The working range is **45 to 75 characters per line, including spaces**, with about 66 as a comfortable middle for continuous prose, and 35 to 45 for a narrow column.

The reason is the return sweep. At the end of each line the eye jumps back and down to the start of the next. That jump is not precise; it lands approximately and corrects. The longer the line, the further the jump and the more often it lands on the wrong line — so a reader loses their place, re-reads, and gradually gives up without knowing why. Very short lines fail from the other side: too many sweeps, and the phrase-sized chunks reading depends on get broken across lines.

Line length, and what the return sweep does



The jump back to the start of the next line is approximate and corrects itself. A longer line means a longer jump, so the reader loses their place, re-reads, and gives up without knowing why.

You can count it. Type $\times$ sixty-six times at your body size and set the column to that width.

Line height, and the relationship it has with measure

Body text usually wants line height of about **1.4 to 1.6 times the font size**. Below 1.3 the lines interfere, and the return sweep has nothing to aim at. Above about 1.8 the lines stop reading as a block and start reading as separate items.

The two settings are coupled: a longer measure needs more line height, because a longer return sweep needs a larger target. This is why pasting a comfortable setting into a wider column breaks it.

All-capitals costs reading speed, and here is why

Setting a paragraph in capitals slows reading measurably — the figures in the literature vary, commonly landing in the region of ten to twenty per cent for continuous text.

The mechanism is word shape. Lowercase words have distinctive outlines produced by ascenders and descenders: `yellow` has a low start, two tall letters and a descender. Capitals remove all of it and make every word a rectangle of the same height, so the reader loses a cue they were using and has to work letter by letter.

This does not make capitals wrong. It makes them expensive, so they are for things that are short and want to be conspicuous rather than read fluently: a two-word label, a heading of four words. A paragraph of capitals is a paragraph most people will skip.

Floors worth knowing

- **16px** is the practical minimum for web body text, and it is also the threshold below which mobile Safari zooms a form field automatically when it is tapped — a real, visible bug caused by a type size.
- **9 to 11pt** is normal body text in print. A flyer read at arm's length and a poster read from three metres are different problems entirely; type size follows viewing distance, not habit.

- **Justified text on a narrow measure** produces rivers of white space, because the software has too few word spaces to distribute. Either widen the measure, enable hyphenation, or set it ragged-right. Ragged-right is not a compromise; most books set ragged-right would be fine.

Contrast is part of readability

Light grey text on white is a style that has cost more reading than any typeface choice. Type carries information in thin strokes, and thin strokes lose contrast first — which is why a light weight at low contrast disappears in bright daylight on a phone, a condition a large share of your audience is in.

The next lesson deals with contrast ratios properly. The point here is that it belongs to typography, not to colour: choosing a hairline weight is choosing a contrast problem.

The limitation

Nearly all of the legibility research behind these numbers was done on Latin script, on screens or paper in controlled light, with young adults who had normal vision. The findings are sound within that frame and should be held loosely outside it.

Other writing systems have different determinants — Devanagari hangs from a headline stroke, Arabic is cursive and connected with contextual letterforms, Chinese characters are read by a different process at a different density — and rules derived from the Latin alphabet transfer unevenly at best. Older eyes, low-vision readers, poor phone screens in sunlight and cheap print all shift the thresholds upward. When in doubt, go bigger and higher-contrast than the numbers require, because the cost of overshooting is aesthetic and the cost of undershooting is that nobody reads it.

BEFORE YOU MOVE ON

A column of text is widened from 65 characters to 110 without changing the font size. Why does it become harder to read?

- A. Wider columns hold more words per line, so the reader takes in more information before each pause and tires faster.
 - B. Long lines reduce the number of paragraph breaks visible on screen at once, removing the visual landmarks readers rely on to keep track of how far through a passage they have got.
 - C. The return sweep at the end of each line is longer and lands less accurately, so the eye more often arrives on the wrong line.
 - D. A wider measure lowers the effective x-height relative to line length, making letters seem smaller than they are.
-

44 • 9 min

Red is not danger everywhere

THE ONE THING TO KEEP

Colour meaning is learned by cultural pairing and varies accordingly, while contrast, attention and colour vision are physiological and do not — so never carry information in hue alone, and check the design in greyscale.

The same wavelength, different meanings

Colour association is learned by repeated pairing inside a culture, so it varies exactly as much as cultures do.

- **White** is weddings and purity in much of Europe and North America, and mourning at funerals across much of East Asia and parts of South Asia.
- **Red** is danger, debt and stop in Western finance and signage, and luck, celebration and marriage across China and India. A financial chart in red

says *loss* in London and something closer to *good* in Shanghai, which is why some Chinese exchanges use red for rising prices.

- **Green** carries religious significance across the Islamic world, and is money in the United States because of the currency.
- **Purple** is mourning in Brazil and Thailand, and royalty in Britain, for the mundane reason that Tyrian purple dye was ruinously expensive.

None of this comes from the light. It comes from what the reader has repeatedly seen a colour attached to. That is good news, because it means you can find out rather than guess — look at what the actual audience's own institutions, packaging and festivals do.

What is not cultural

A few colour effects are about the eye and hold broadly:

- **Saturated colours attract attention more than desaturated ones**, which is why one accent works and five do not.
- **Pure blue is a poor colour for small text.** The eye focuses different wavelengths at slightly different distances — chromatic aberration — and the cones sensitive to short wavelengths are sparse and nearly absent from the very centre of the fovea. Small saturated blue type on a bright background reads as slightly soft no matter how good the screen is.
- **Pure red on pure blue** vibrates unpleasantly, because the two focus at opposite ends of that error and the eye cannot settle on both.
- **Around one in twelve men** of northern European descent has some form of red-green colour vision deficiency, which the next lesson takes up properly.

These are constraints rather than meanings, and they do not vary by market.

The rule that survives every culture

Never carry meaning in hue alone.

A status dot that is green for working and red for broken conveys nothing to a substantial minority of your readers, nothing in greyscale print, and nothing

on a badly calibrated projector. Add a second channel every time: a tick and a cross, a filled and hollow shape, a word, a position.

The cheap check: view the design in greyscale. Every phone and operating system has a greyscale accessibility setting, and any image editor will desaturate a screenshot. If the information survives, the colour was decorating; if it disappears, the colour was load-bearing and needs a partner.

Colour psychology claims are weaker than the marketing pages say

You will read that blue increases trust, that orange drives purchases, that red creates urgency. These claims trace back to a small number of studies with small samples, inconsistent replication, and effect sizes nowhere near what is claimed. The much-repeated "red button outperformed green" result is a single test on one site with one context and one set of surrounding colours, from which the general law does not follow.

What holds up better, and is unexciting: **contrast against the surroundings** predicts noticing far more reliably than hue. A red button on a red-heavy page is invisible; a green one on the same page is conspicuous, and the reverse holds on a green page. The effect people attribute to the colour is mostly the effect of the difference.

A palette you can actually work with

Define colour by role, not by name, and keep the accent scarce:

- **page** — background
- **ink** — body text
- **muted** — secondary text
- **line** — borders and rules
- **accent** — exactly one thing per screen

Then check each pairing for contrast before you fall in love with it. Free tools: the WebAIM contrast checker in a browser, the Colour Contrast Analyser as a

desktop app, GIMP or Krita for desaturating a mock-up, and any browser's developer tools for inspecting a live page.

The limitation, stated honestly

Cultural colour associations are generalisations about very large groups of people, and inside any of those groups they vary by age, religion, region, class and context. Red at a wedding and red on a warning label are different reds to the same person. A designer applying a country-level colour rule to an individual audience is making the same mistake as a designer applying an American rule everywhere; the granularity changed, the guessing did not.

The defensible version is narrow: **know that the associations differ so you stop assuming yours are universal, then look at what your actual audience's own materials do**, and reserve your confidence for the physiological constraints and the greyscale test, which hold everywhere.

BEFORE YOU MOVE ON

An A/B test finds a red call-to-action button outperforms a green one on a retail page. Why is "red converts better" the wrong lesson to take away?

- A. Single tests lack statistical power, so the difference is probably noise rather than a real effect.
- B. Colour preference varies by culture, so a result from one market cannot transfer to another.
- C. Conversion depends mainly on the wording of the button label, so the colour comparison isolated the wrong variable entirely.
- D. What was measured was the button's contrast against that particular page, not a property of red.

45 · 9 min

One in twelve men cannot see the difference you just used

THE ONE THING TO KEEP

Colour vision deficiency, poor screens and sunlight make the impaired case the ordinary case, so measure contrast rather than eyeballing luminance, never encode meaning in hue alone, and honour reduced-motion — while knowing the 4.5:1 threshold is a contested floor rather than a guarantee.

Who you are actually designing for

Roughly **8 per cent of men and 0.5 per cent of women** of northern European descent have some form of colour vision deficiency, with lower but non-trivial rates in other populations. The commonest form is deuteranomaly, a shifted green sensitivity, which makes reds, greens, browns and oranges collapse towards one another.

A much larger group has something else going on: presbyopia after forty, cataracts, a cheap screen, a projector with a dying lamp, bright sunlight on a phone, or simply a moving train. In aggregate, the share of your audience not seeing your design under ideal conditions approaches all of them.

This is not a special case to handle at the end. It is the ordinary case.

Contrast ratios, and what the number means

The contrast ratio compares the **relative luminance** of two colours — a weighted measure of how much light they emit or reflect, weighted because the eye is far more sensitive to green than to blue. The scale runs from 1:1 (identical) to 21:1 (pure black on pure white).


The widely used thresholds:

- **4.5:1** for body text.

- **3:1** for large text, meaning about 18pt, or 14pt bold.
- **3:1** for interface components and meaningful graphics — the edge of a form field, an icon that carries information.
- **7:1** is the enhanced level, for text that has to work in poor conditions.

The practical consequence that catches people: **luminance is not lightness as you perceive it**. A mid-yellow looks bright but has high luminance, so yellow on white fails badly while the same-looking blue passes easily. You cannot eyeball this; you have to measure it, and the measurement is free.

Contrast ratios, floors and a common failure



Luminance is not lightness as you perceive it, which is why a mid-yellow that looks bright fails on white while a blue that looks similar passes. You cannot eyeball this, and measuring it is free.

The three checks that take two minutes

1. **Greyscale.** Turn the display to greyscale, or desaturate a screenshot. Anything that becomes indistinguishable was carrying information in hue alone.
2. **Contrast.** Run every text-and-background pair through a checker — WebAIM's web page, the Colour Contrast Analyser desktop app, or the contrast readout built into Chrome and Firefox developer tools.
3. **Sunlight.** Take the phone outside. Light grey text on white, thin weights and low-contrast placeholders fail here first and most.

Motion, and the people it makes ill

Autoplaying parallax, large sliding transitions and looping background video cause genuine symptoms — nausea, dizziness, disorientation — for people with vestibular disorders. This is not a preference.

Operating systems expose a *reduce motion* setting, and the web exposes it to CSS:

```
@media (prefers-reduced-motion: reduce) {  
  * { animation: none !important; transition: none !important;  
  }  
}
```

Six lines, honoured by every modern browser. For video, the equivalent courtesy is not autoplaying anything with rapid flashing, and a rate above three flashes per second is the threshold associated with photosensitive seizures.

Text that can be resized, and other small mercies

- Do not disable pinch zoom. `user-scalable=no` in a viewport tag is a line of code that removes a disabled person's ability to read your page.
- Set sizes in relative units so a reader's own larger default is respected.
- Give a form field a real label, not a placeholder that vanishes the moment somebody starts typing and takes the question with it.
- Give every informative image alt text, and every decorative one empty alt text.
- Make a focus outline visible. Removing it because it looked untidy leaves keyboard users navigating a page with no cursor.

The contrast standard itself is contested

Here is the unsettled part, and it is worth knowing because you will meet designers on both sides.

The current contrast formula is known to model some cases badly. It handles dark mode poorly — it can pass a pairing that is genuinely hard to read on a dark background and fail one that is comfortable — and it ignores font weight and size almost entirely, when a hairline weight at 4.5:1 is plainly harder than a bold at the same ratio. A proposed replacement, APCA, models perceived lightness difference and does take weight into account. It has been in draft for years and is not the adopted standard.

So the position to hold: **the 4.5:1 threshold is the accepted baseline and the one you will be measured against, and it is a floor rather than a guarantee.** Meet it, then look at the thing on a real screen in real light, because a pairing that passes and still looks faint to you is faint.

Why this is not charity

The strongest argument is not moral. Every one of these decisions improves the design for people with no impairment at all — in sunlight, on a bad screen, tired, in a hurry, on a train. Captions serve the hearing majority. High contrast serves everybody outdoors. A visible focus outline serves anybody who prefers a keyboard. Designing for the edge of the range is how you get a design that works in the middle of it.

BEFORE YOU MOVE ON

A designer checks a yellow-on-white label by eye, judges it bright and clear, and ships it. Why does it fail a contrast measurement?

- A. Contrast ratios are calculated for print reflectance, so screen colours routinely score lower than they appear.
- B. Saturated colours are excluded from the contrast formula, which only compares neutral tones.
- C. Yellow is the hue most affected by red-green colour vision deficiency, so the standard applies an additional penalty to yellow pairings in order to protect those readers from unreadable text.
- D. The ratio is computed from relative luminance, and yellow has high luminance, so it sits close to white on the scale despite looking vivid.

46 · 9 min

Every photograph is mostly the part you cut off

THE ONE THING TO KEEP

A photograph's edges, height, lens and chosen instant are all claims made before any editing, so cropping can mislead more than retouching does — the useful test is whether the person who was there would call it a fair account.

The frame is the first claim

A photograph is defined as much by its edges as by its contents. Everything outside the rectangle has been excluded by somebody, and the viewer has no way of knowing what it was.

This is why the familiar defence — *we did not alter the image* — is weaker than it sounds. You do not need to alter an image to make it say something that is not true. You need only to choose where the edges go, and choosing where the edges go is not optional; there is no photograph without it.

A crowd photographed tightly around a confrontation is a riot. The same crowd photographed wide is forty people and an argument. Neither image has been manipulated. Both are accurate. They make opposite claims.

Four decisions the camera makes for you

1. **Height.** A camera below eye level looks up at the subject, which is the geometry of a smaller person looking at a larger one, and reads as imposing or authoritative. Above eye level reverses it and reads as diminished or observed. Level is the neutral one, and neutral is also a choice.
2. **Distance and lens.** A wide lens used close up exaggerates what is nearest — a nose, a hand, a machine — because the distance ratio between near and far parts of the subject is extreme. It reads as aggressive, comic or distorted. A long lens from far away compresses depth, flattens features, and

reads as observational or surveillant. The face has not changed; the relationship between the camera and the face has.

3. **The moment.** Speech makes unflattering shapes. Any politician can be photographed mid-vowel looking deranged, and any calm person can be caught blinking. A fifteen-frame difference is the difference between competent and shifty, and no editing is involved.
4. **What is in focus.** Throwing a background out of focus removes information from the frame as effectively as cropping it. The context is present, illegible, and deniable.

Where the industry draws its line, and why the line is odd

News agencies have explicit rules. Reuters, the Associated Press and their peers forbid adding or removing content — no cloning out a distracting object, no compositing — while allowing cropping, and allowing tonal adjustment within what a darkroom could have done. Photographers have been dismissed for cloning out a shadow.

The rule is defensible as a bright line, because "did you add or remove pixels" is checkable and "is this fair" is not. But it is worth seeing clearly what it permits: **a crop can mislead far more comprehensively than a cloned-out shadow**, and it is entirely within the rules. The line is drawn where it is because it can be enforced, not because it tracks the harm.

This disagreement is live. Documentary photographers, news editors and academics argue about it continuously, and there is no settled resolution to hand you.

The working test for commercial and documentary work alike

You will rarely be deciding whether to clone something out. You will constantly be deciding how tight to crop, which frame to pick, and where to stand.

The question that does useful work: **would the person in the photograph, or the person who was standing there, recognise this as what happened?** Not *would they like it* — nobody likes their own photograph — but *would they say this is a fair account of the moment*.

It is not a rule and it does not decide hard cases. It does catch the ordinary ones, which are the ones you will actually face: the crop that makes a small workshop look derelict, the low angle that makes a friendly person look severe, the frame chosen because it was dramatic rather than because it was representative.

Generated and heavily edited images

An image that is wholly or partly generated carries a further problem: it depicts something that never occurred, in a register the viewer reads as photographic.

The duties here differ by country and keep changing — India's rules for synthetic media, China's labelling requirements and the European transparency provisions all point towards visible marking, on different timetables and with different scope, and none of this is legal advice. Provenance standards such as C2PA attach a signed record of how an image was made, and are being adopted unevenly; metadata also does not survive a screenshot.

The design consequence is the durable part: **a visible label in the frame is the only disclosure that survives being copied.** Everything attached to the file is lost the first time somebody screenshots it.

What this does not mean

None of this makes photography dishonest or unusable. Every representation selects — a sentence selects, a diagram selects, a film selects. The point is not to achieve neutrality, which is unavailable, but to stop believing you have it by default.

A photographer who knows the frame is an argument can make the argument deliberately, defend it when asked, and notice when they are making one they did not intend. A photographer who believes the camera simply records is making arguments constantly and cannot see any of them.

BEFORE YOU MOVE ON

Why is a news agency rule that permits cropping but forbids cloning out an object an imperfect match for the harm it aims to prevent?

- A. Removing one object changes little, while choosing the edges can exclude the entire context and reverse what the image asserts.
 - B. Cropping is reversible from the original file, so a crop can always be audited later while a clone cannot.
 - C. Cloning is easier to detect with software, so the rule targets what is enforceable rather than what is most misleading.
 - D. Readers already expect photographs to be cropped for layout reasons, so they discount framing automatically and can only be deceived by edits made at the level of individual pixels.
-

47 • 8 min

The photograph that fits anywhere belongs nowhere

THE ONE THING TO KEEP

A stock image is built to fit many contexts and therefore carries information about none, so readers skip it — replace it with a mediocre photograph of the actual thing, a diagram, or nothing at all, and reserve stock for abstraction and for protecting a real person.

The mechanism, in one sentence

A stock image is produced to be sold many times into many contexts. To be sellable into many contexts it must be specific to none. The result is an image carrying almost no information about your particular subject, which is why readers' eyes pass over it without stopping.

This is the same phenomenon as the advertising banner that nobody sees. Readers learn, quickly and without noticing, which regions of a page reliably contain nothing they need, and stop allocating attention to those regions. A generic smiling team around a laptop occupies the shape of a photograph and delivers the informational content of a decoration.

So the cost is not that the picture is a bit bland. The cost is that a slot which could have carried evidence carried nothing, and that the reader has been trained slightly further towards skipping your images.

The second cost is recognition

A small number of images are enormously over-used, and a reader who has seen one before has been told something you did not intend to tell them: that nobody here had a camera, or an afternoon, or a real example.

This matters most for the businesses most likely to reach for stock. A large corporation using a generic photograph reads as impersonal. A two-person workshop using a generic photograph of a different, larger, foreign workshop reads as a workshop that would rather not show you itself.

What to use instead, all of it free

1. **Your own phone.** A genuinely mediocre photograph of the actual thing beats an excellent photograph of a different thing, because the actual thing carries information. Daylight from a window, the object on a plain surface, take twenty, keep one.
2. **The process, not the product.** Half-finished work, the workbench, the hands, the failed attempt on the shelf. These are the images clients never think to take and the ones readers stop at.
3. **Screenshots and documents.** For anything digital, the actual screen is evidential and specific. Crop and annotate it.
4. **Diagrams.** If what you need is to show a relationship rather than an object, draw it. Inkscape, Excalidraw and diagrams.net are free.
5. **Typographic covers.** A large, well-set word on a plain field is honest and looks deliberate, where a wrong photograph looks like a mistake. It is

also the cheapest thing in this list.

6. **Stock chosen oddly.** If you must use a library, search past the obvious term. Specific, slightly strange, slightly imperfect images work; the first page of results for *teamwork* does not.

When stock is genuinely the right answer

There are real cases, and pretending otherwise is not useful:

- **A background or texture** doing a compositional job, with no informational claim attached.
- **Deliberate abstraction** — an article about a general phenomenon where a specific photograph would falsely imply a specific case.
- **Protecting somebody.** An article about debt, illness or abuse often should not carry a photograph of a real person, and a model-released stock image is the honest way to avoid it.

In the last case the licence conditions matter. The major free libraries — Unsplash, Pexels, Pixabay — all forbid using an identifiable person in a way that implies endorsement or shows them in a bad light. A stock face beside a headline about fraud is a licence breach and a much larger human problem.

Generated images, honestly

Generating an image is now the obvious substitute, and it solves the specificity problem partially and creates others.

What it is genuinely good for: abstraction, texture, illustration in a consistent style, and situations where the alternative is a wrong photograph. What it is bad for: anything that will be read as evidence, because it is evidence of nothing.

Three things to hold at once. **Disclosure duties differ by country and are moving** — several jurisdictions now require visible marking of synthetic media, with different scope and timetables, and none of this is legal advice. **The copyright position is unresolved:** whether training on scraped images is lawful is being litigated in several countries with different statutory

starting points, and whether a generated image can be owned at all has been answered differently in different places. Anyone telling you this is settled is telling you about one jurisdiction, at one moment, or is selling something.

The artefacts are still readable at large sizes — hands, text on signage, repeated faces in a crowd, jewellery and teeth — so check at full size before you ship.

The test to apply before you place any image

Ask what the reader learns from this picture that they would not learn from the caption alone. If the answer is nothing, you do not have an image problem. You have an empty slot that a photograph is politely covering, and the better fix is usually to remove it and let the text breathe.

BEFORE YOU MOVE ON

Why does a generic stock photograph cost a page more than simply leaving the space empty would?

- A. Readers learn to skip regions that reliably hold nothing, so an uninformative image trains them to ignore your images generally.
- B. Stock images increase page weight and load time, which pushes readers away before the content appears.
- C. Empty space is read as deliberate restraint, while a photograph creates an expectation of quality the rest of the page must then meet.
- D. Licence terms on free stock libraries restrict commercial reuse, so the image carries legal risk the empty space does not.

48 · 8 min

Three pictures, three different promises

THE ONE THING TO KEEP

A photograph claims this existed, an illustration claims this is the kind of thing, and a diagram claims these are the parts and their relations — choose by what the reader must do with the picture, and say what a diagram simplified, because boxes assert boundaries the world may not have.

Each form asserts something different

Before choosing how a picture should look, decide what it has to claim. The three main options make genuinely different promises to the reader, and picking the wrong one either overclaims or fails to carry the load.

- A **photograph** asserts *this existed, and it looked like this*. It is evidential. It is also stuck with the surface — a photograph of a gearbox shows a metal box.
- An **illustration** asserts *this is the kind of thing being discussed*. It is explicitly a representation, so it can be generic without lying, and it can omit. A cutaway drawing shows the inside of the gearbox, which no photograph of an assembled gearbox can.
- A **diagram** asserts *these are the parts and this is how they relate*. It abandons appearance entirely for structure. A London Underground map is geographically wrong and is the correct picture, because the question passengers have is about connections, not distances.

A **chart** is a fourth form, asserting quantity and comparison, and it is covered separately.

Four pictures, four different claims

Photograph	This existed, and it looked like this
Illustration	This is the kind of thing being discussed
Diagram	These are the parts and how they relate
Chart	These are the quantities, and their comparison

Choose by what the reader has to do with the picture. An article on how a sewing machine forms a stitch needs a diagram, and a beautiful photograph of the machine will not substitute, because the stitch happens inside and in stages.

The choosing question

Not "what would look good here" but: **what does the reader need to do with this?**

- Believe it happened → photograph.
- Recognise the type of thing → illustration.
- Follow a process or a relationship → diagram.
- Compare amounts → chart.
- Find their way → map, which is a diagram with geography in it to the extent the task needs.

This resolves most arguments quickly. An article explaining how a sewing machine forms a stitch needs a diagram, and a beautiful photograph of the machine will not substitute, because the stitch happens inside and in stages. An article about a person needs a photograph, and an illustration of a generic person would drain it of the one thing it had.

Illustration solves a problem photographs cannot

When the subject is sensitive, an illustration lets you show a situation without exposing a person. A drawn figure is understood as a type rather than an individual, so nobody is identified, nobody consented under pressure, and nobody's face is attached to a topic for the rest of their life.

It also allows deliberate omission. A drawing can leave out the clutter that a photograph must include, which is the same act as cropping but visible and

honest — nobody thinks a drawing was the whole room.

Diagrams lie by boundary

The limitation of diagrams is specific and worth knowing: a box-and-arrow drawing asserts that the world has boxes and arrows in it.

Draw a system as five boxes and you have claimed there are five distinct components with clean interfaces. Very often the real thing is a continuum, or the boxes overlap, or one of them is actually three things that nobody has separated. The diagram is now clearer than the truth, and the reader will remember the boxes.

This is not an argument against diagrams — a simplification you can name is worth more than a fog. It is an argument for saying what you simplified. One line under the drawing does it: "simplified; stages 3 and 4 overlap in practice."

The free path is complete

- **Inkscape** — vector drawing, illustrations, cutaways, print-ready SVG and PDF.
- **Krita** — painting and raster illustration, good with a tablet or a mouse.
- **diagrams.net** and **Excalidraw** — box-and-arrow diagrams in a browser, including on a phone.
- **Mermaid** — diagrams written as text, which version-control well and regenerate when the content changes.
- **GIMP** or **Photopea** — raster editing and photo preparation; Photopea runs in a browser and opens PSD files.
- **Graphviz** — for a diagram whose layout should be computed rather than arranged by hand.

None of this needs a subscription and most of it runs on a low-end laptop.

Consistency matters more than quality

The practical failure in most small projects is not that the illustrations are poor. It is that there are three styles in one document: a photograph, a clip-art

icon and a hand-drawn diagram, each from a different source.

A reader reads that as three separate voices and downgrades the whole piece. Pick one drawing style — same line weight, same palette from your role-based colours, same level of detail — and apply it to everything you draw. A consistently mediocre set of diagrams looks considered. An inconsistent set of good ones looks assembled.

BEFORE YOU MOVE ON

An article explains how a sewing machine forms a lock stitch. Why will a high-quality photograph of the machine fail where a diagram succeeds?

- A. Photographs of machinery carry irrelevant surface detail, which distracts the reader from the explanation in the text.
- B. Readers trust diagrams more than photographs for technical subjects, because a diagram signals that the author understands the mechanism.
- C. Diagrams can be reproduced at any size without losing detail, while photographs degrade when printed small.
- D. The stitch forms inside the machine and across several stages, and a photograph can only show the outside at one instant.

CASE STUDY · MODULE 6

The red and green flags on four thousand lab reports

A three-branch diagnostic laboratory in Hubballi, reviewing its printed blood-test report after a complaint that never reached anyone.

The report was clean and well set. Out-of-range values were printed in red; in-range values in green. That was the whole signalling system, and it carried information in hue alone.

It came to light by accident. A phlebotomist's father, a retired schoolteacher with a red-green deficiency, asked his son which of the numbers were the bad ones, because he could not tell. Nobody had ever complained, and the reason is worth stating: people who cannot read a document do not send feedback about it. They ask somebody, or they guess, or they take it to a doctor and say nothing.

The size of it, counted rather than estimated

The lab prints roughly 4,000 reports a month and about 55 per cent go to men. Taking the commonly cited figure of around 8 per cent of men with some form of red-green deficiency, that is on the order of 170 men a month receiving a document whose only signal they cannot read.

That is the smaller number. The larger one is everybody else, because the report is not only read in colour.

- Insurance claims are photocopied at the branch, in one colour. Every flag vanishes.
- Reports are photographed and sent on messaging apps, often at low screen brightness in daylight.
- The lab's own PDF is opened by doctors on phones with a warm night filter running.

A screenshot of the PDF converted to greyscale lost 100 per cent of the flags. That single test, which took under a minute, was what moved the conversation

from opinion to fact.

The contrast measurement was the second surprise. The red used, at the body size on the report, measured 3.9 to 1 against white, below the widely used 4.5 floor for body text. The green measured 1.9 to 1, which is near the bottom of the scale. Both had been chosen by eye on a good monitor.

The decision, and what it cost on each side

Changing the report was not a five-minute edit. The template sits inside the lab's accredited document control, so a revision means a change request, re-approval, and a new version number across three branches. There was pre-printed stationery to replace. Fourteen staff read those colours daily and would have to be retrained. And the senior pathologist liked the scheme, on the reasonable ground that referring doctors recognised it at a glance.

Not changing it cost nothing visible and nothing measurable. That was the honest description of the alternative, and it was the strongest argument for acting, because a failure that generates no complaints generates no pressure either.

The proposal deliberately did not remove the colour. Colour was demoted from the only channel to the second one.

The typography was part of the same problem

The review picked up two more things while the template was open, both of which belong to readability rather than to decoration.

The patient's name and the reference ranges were set in a light weight, chosen because it looked modern. Type carries its information in thin strokes, and thin strokes lose contrast first, so the light weight was a contrast decision disguised as a style decision, and it failed on a photocopy before anything else did.

The comments field at the foot of the report ran the full width of the page, about 110 characters a line. The return sweep at the end of each line is approximate, and the longer the line the more often the eye lands on the wrong one.

Set to two columns of about 60 characters, the same text is read without anybody losing their place, and the page got shorter rather than longer.

None of this needed paid software. The template was rebuilt in a free layout tool, the contrast was measured in a free browser checker, and the greyscale test was a setting on a phone.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

The flags kept their colours and gained two more channels: an arrow glyph pointing up or down beside the value, and the word **HIGH** or **LOW** set in bold. The reference range was printed beside every result rather than in a legend at the foot of the page, so a reader who trusts neither channel can compare the numbers themselves.

The red was darkened until the flag text measured 7.1 to 1 against the paper, which also survived the branch photocopier. The greyscale test was repeated on the new PDF: every flag remained legible.

The cost was one week of document control, about forty thousand rupees of stationery, and one afternoon of training.

Six months later the lab could not show a change in any number it routinely collects, and the report to the partners said so. The evidence it did have was the test itself: a photocopy of the new report, handed to the same retired schoolteacher, who read it without asking anyone.

A marketing agency had separately proposed a blue header on the grounds that blue increases trust. That claim was not adopted, on the grounds that the published support for colour-psychology effects of that kind is thin and inconsistent, while the contrast measurement is not.

Nobody had complained in years. Why is the absence of complaints weak evidence that the report was working?

Because the failure removes the route by which it would be reported. A reader who cannot distinguish the flags usually assumes they are being slow rather than that the document is at fault, and asks somebody or leaves. The same is true of everyone reading a photocopy or a dim screen. A design problem that is invisible to the people who could fix it accumulates instead of surfacing.

The fix kept the red and green. Why is adding a second channel the right move rather than replacing the colours?

The failure was not the hue itself but the fact that meaning rested on hue alone, which disappears for viewers with colour vision deficiency, in one-colour reproduction, and on poor screens. Colour is an effective and fast second signal when something else carries the meaning first, so an arrow and a word restore the information while keeping the recognition the referring doctors already had.

The colours had been chosen by eye on a good monitor and looked fine. Why did they still fail a contrast check?

Contrast ratio compares relative luminance, which is a weighted measure of emitted or reflected light, and it does not track perceived lightness. A saturated colour can look bright and vivid and still sit close to white in luminance, which is why a mid-yellow fails badly on white while a similar-looking blue passes. Eyeballing cannot report the number, and measuring it costs nothing.

MODULE 6 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 Two businesses post equally often. One looks like a company and one looks like a person trying things. The logos are equally good. What produces the difference?
 - A. The second one posts at inconsistent times of day, so recognition never has a chance to accumulate at all.
 - B. The second one uses free software, whose output reads as generic to a viewer.
 - C. The second re-decides type, colour and spacing from memory, and the drift shows across twenty posts.
 - D. The second one uses more than two typefaces within a single piece of work.
- 2 A typeface reads as institutional authority to one audience and as nothing in particular to another. What decides what a typeface means?
 - A. Where that particular reader has repeatedly seen it before.
 - B. The geometry of the letterforms, which carries an inherent mood.
 - C. The decade in which the face was originally drawn and released.
 - D. The classification family the face belongs to within the standard system.
- 3 Two typefaces set at 12pt look noticeably different in size. What explains it?
 - A. Screen rendering hints differ between the two font files, which changes the apparent weight of the strokes.
 - B. The software is applying line height differently to each of them.
 - C. One has a wider set width, and wider letters are perceived as larger.
 - D. Point size measures the body the letters sit on, and x-height differs between faces.

- 4 You are designing for an audience whose colour associations you do not share. Which claim holds across markets rather than varying with them?
 - A. Red signals danger, so it should be reserved for warnings and errors.
 - B. Saturated colours attract attention more than desaturated ones do.
 - C. Blue increases trust, which makes it suitable for financial work.
 - D. White reads as celebratory and suits weddings and formal announcements.

- 5 A photographer defends a contested image by saying the file was not altered. Why is that weaker than it sounds?
 - A. Most images are altered in camera before the file is ever written to the card.
 - B. Metadata records every edit, so the claim is checkable and usually turns out to be false.
 - C. The edges were chosen, and a crop can mislead more comprehensively than a retouch.
 - D. Agency rules forbid tonal adjustment as well as cloning, so the defence is irrelevant.

- 6 A generic photograph of a team around a laptop sits at the top of an article and readers' eyes pass straight over it. What is the mechanism?
 - A. It is built to sell into many contexts, so it carries information about none.
 - B. It is usually supplied at a lower resolution than an original photograph would be.
 - C. It carries a watermark that readers have learned to associate with low-quality pages.
 - D. It shows people, and readers have grown tired of photographs containing people.

MODULE 7

The surface and the system carrying it

A piece is never seen in the conditions it was made in. It appears at a particular size, in a particular shape, with the sound off, next to other things, ranked by a system whose objective is not yours. This module covers safe areas and crop chains, the still that recruits the play, designing for silence, the feed as context, what recommender systems actually optimise, and which numbers answer which question.

BY THE END OF THIS MODULE YOU CAN

Deliver a piece into the conditions it will actually meet — survive the crop chain and interface overlays, set a cover frame and link preview deliberately, carry meaning with the sound off and on a phone speaker, mock a design up inside a real feed, distinguish the platform's objective from your own, and read a retention curve to name the specific failure it records

The same design fails in five different ways

THE ONE THING TO KEEP

Design for the physical size and distance the piece will actually be seen at, because every surface has a different smallest legible element and resizing tools do not change what the piece says.

One artwork, five surfaces, five separate problems

A layout that works on your 24-inch monitor is not a design yet. It is a design *for a 24-inch monitor viewed at 60 centimetres*. Change the surface and it can fail completely without a single element changing.

The unifying mechanism is simple and worth holding on to: the smallest legible element is set by physical size divided by viewing distance. Everything below is that one rule applied to five very different situations.

The five surfaces

The vertical video, 9:16

Export at 1080 by 1920. The trap is the platform interface, which sits *on top of* your video: captions, the username, the sound label, a column of buttons down the right side, and the caption sheet along the bottom. Roughly, keep anything that must be read out of the top eighth and the bottom fifth, and away from the right edge. These margins differ per platform and change with app updates, so download the current safe-area template rather than trusting any number, including that one.

Two more things people learn the hard way: the first frame is a thumbnail, and most viewing is muted, so anything spoken must also be written. Video editing covers the cutting side.

The feed image

On Instagram, a 4:5 portrait occupies more vertical space in the feed than a square, which is free attention. Text must survive the feed's own compression — thin type at small sizes turns to mush after re-encoding, so err heavier than feels necessary. Assume it will be seen at about 400 pixels wide on a phone, not at the size of your artboard.

The thumbnail at 150 pixels

You upload 1280 by 720. It is displayed, in the places that matter, at somewhere between about 120 and 250 pixels wide. Every decision follows from that: three or four words at most, one clear subject, and enough contrast that it separates from both light and dark interface backgrounds.

The test is not optional and takes ten seconds: after exporting, zoom the image down to about 12 percent, or send it to your own phone and look at it at arm's length. If you cannot tell what it is, nobody scrolling can either.

The printed flyer

Three things reliably go wrong.

Colour. Screens are RGB and mix light; presses are CMYK and mix ink. Vivid blues, bright greens and neon oranges exist on screen and have no CMYK equivalent, so they come back duller and slightly shifted, and the client thinks the printer made a mistake. Convert and inspect before you send.

Resolution. Print wants around 300 dpi. A5 at 300 dpi is 1748 by 2480 pixels; A4 is 2480 by 3508. An image that looked fine on screen at 800 pixels wide is a mess on paper.

Bleed. Anything running to the edge must extend past it, usually 3mm, because the guillotine is not perfect. Miss it and you get a white sliver down one side of the whole run.

The free path here is real but specific. **Scribus** is the free, open-source layout tool that handles CMYK and PDF/X export properly, which is what a commercial printer wants. **Inkscape** is the free vector tool for the artwork itself. GIMP and Photopea are excellent at many things and weak at CMYK work-

flows, so do not discover that at the print shop. If you would rather buy once than subscribe, Affinity's one-off purchase is the usual answer. Vector and print goes into this properly.

The slide behind a speaker

The room is bright, the projector is older than you hoped, and the back row is twenty metres away. Practical consequences: body text no smaller than about 28 to 30 point, very few words, one idea per slide, and high contrast between text and background.

Do not assume a dark slide is safer. On a weak projector in a lit room, black backgrounds wash out to grey and light text loses most of its contrast; in a dark room, a full-white slide is painful. There is no rule that survives every room, which is why the only reliable step is to put one slide up in the actual room before the talk. And nobody reads a paragraph while listening to a person — the slide supports the speaker or competes with them.

The tempting shortcut, and what it misses

Every tool now offers one-click resizing — Canva's resize feature, Adobe Express, various apps. These genuinely save time on mechanical reflow, and for simple layouts they are fine.

What they cannot do is rethink. They will shrink a seven-word headline to fit a thumbnail rather than tell you that seven words is four too many for that surface. They will keep a caption inside the frame while placing it exactly where the platform's own buttons will sit. Resizing moves elements; designing for a surface changes which elements exist.

Do this today

Take one piece you have already published and view it at the true size on the true surface: the reel on a phone with the app open over it, the flyer printed on ordinary paper, the thumbnail at 12 percent zoom. Write down the first thing that becomes unreadable. That is the piece's actual failure, and it will not be the thing you spent the most time on.

BEFORE YOU MOVE ON

A designer exports a flyer with a vivid electric-blue background from Photopea and sends the RGB file to a commercial printer. The delivered flyers are a duller, slightly purple blue with a thin white strip down one edge. What happened?

- A. The file was exported at screen resolution, so the press interpolated the colour and misaligned the trim.
 - B. The printer used the wrong paper stock, which absorbed more ink than expected and shifted the hue.
 - C. The blue lay outside the CMYK gamut so it converted to the nearest printable ink mix, and the artwork had no bleed, so the guillotine cut inside the edge of the colour.
 - D. Photopea cannot export print-ready files at all, so any file from it will reproduce incorrectly on a press.
-

50 • 9 min

The thumbnail and the title are one thing

THE ONE THING TO KEEP

Title and thumbnail are two channels of one message — make them complementary rather than redundant, and only promise what the piece actually pays.

The mistake that wastes half your space

A video titled "I built a house out of bamboo" with a thumbnail reading BAMBOO HOUSE over a picture of a bamboo house. Nothing wrong with any single part. Together they say one thing twice and use up the only two pieces of real estate you get.

Treat them as one unit with two channels. The title carries the specific — the number, the place, the claim. The thumbnail carries what a sentence cannot: the state of a thing, a face, a before and after, a scale you have to see. When they complement each other, the pair produces a question the viewer wants closed. When they duplicate, they produce a statement, and a statement is easy to walk past.

Worked example. Title: "I built a house out of bamboo for ₹80,000."

Thumbnail: the finished house, one small human figure at the door for scale, no text at all. The title supplies the number; the image supplies the disbelief; the gap between them is the reason to click.

This is not a video-only idea. A book cover and its title work the same way. So do a poster's image and its headline, and an email subject line and its first sentence.

The craft, unsentimentally

- **Three or four words maximum** on the image, if any. At 150 pixels wide, a sentence is texture.
- **One subject.** Two competing subjects at thumbnail size is one blur.
- **Faces work**, and there is a mechanism behind it: humans identify faces and read expressions at very low resolution, far below the resolution needed to read text. A face carrying a real reaction survives shrinking better than almost anything else.
- **Contrast against the interface**, in both light and dark mode. A near-white background disappears into a light theme; a black one disappears into a dark theme.
- **Front-load the title.** On phones, titles truncate — often somewhere around forty to sixty characters depending on where they appear — so the specific noun goes first. "Bamboo house for ₹80,000 — the full build" survives truncation. "How I spent six months learning to..." does not.

Test it in context rather than in isolation. Screenshot a real results page in your subject, open it in Photopea or GIMP, and paste your thumbnail into the grid at the same size as the others. You will see immediately whether yours is

findable or whether it is the fourth dark-blue rectangle in a row. YouTube has also added built-in thumbnail testing on some channels; if you have it, real data beats your opinion and mine.

Clickbait, defined usefully

The distinction is not "exciting versus boring", and treating it that way leads people to make deliberately dull work and call it integrity.

The distinction is whether the piece pays the promise. "I spent thirty days learning to weld" is a strong hook and completely honest if you spent thirty days learning to weld. An arrow pointing at something that does not appear in the video is not, whatever its click-through rate.

The failure has a specific cost, and it is not reputational vagueness. It is that you train an audience to distrust your framing, so your next genuinely surprising piece is read as more of the same, and the one asset you cannot buy — being believed — is the one you spent.

Two things that are your responsibility

Generated thumbnails. If the thumbnail shows a scene that never happened and reads as photographic, that is two problems at once: a promise the video cannot pay, and in a growing number of countries a labelling duty. India's IT Rules for synthetic media and China's labelling rules both require visible marking, the EU's transparency provisions point the same way, and the timetables keep moving — making things with AI sets out where that stands. The rules differ by country and a lawyer where you live is the person who answers it for your case. The design consequence is simple though: a visible label in the frame is the only part that survives a screenshot.

Other people's faces. A recognisable person on your thumbnail, particularly in a reaction or an implied-story crop, is a claim about them. Get permission, and do not fabricate an expression they never made.

Where click-through is a bad master

Optimising purely for click-through selects for the audience that clicks, which is not the audience that stays. Recommendation systems have shifted steadily toward watch time and satisfaction signals, so a high-CTR, low-retention pattern is now actively penalised on most platforms rather than merely being unpleasant.

The practical version: read click-through and retention together, always. A change that raises one and drops the other has not improved anything; it has moved the failure later in the funnel where it is harder to see.

Do this today

Take one published piece and write down the title and describe the thumbnail in a sentence. If both contain the same information, rewrite one of them so it carries what the other cannot. Then look at the pair at actual size on a phone. Most of the time this single edit is worth more than a week of production polish.

BEFORE YOU MOVE ON

A creator changes a thumbnail and click-through rises sharply, but average view duration falls and subscriber growth slows. They are considering keeping the new thumbnail because clicks are the metric that grew. What is the strongest argument against?

- A. The new thumbnail sets an expectation the video does not meet, so it converts strangers into disappointed viewers and moves the failure further down the funnel.
- B. Click-through is a vanity metric and should be ignored in favour of watch time alone.
- C. The drop in duration is probably seasonal noise, so the thumbnail should be kept and re-evaluated after a month.
- D. Thumbnails should never be changed after publishing because it resets the platform's learning about the video.

51 · 9 min

Something is always sitting on top of your picture

THE ONE THING TO KEEP

The visible frame is smaller than the exported one and its boundaries change without notice, so compose for the intersection of every crop you will need, keep text out of the bottom fifth, and archive a wide master so a new shape is an export rather than a reshoot.

The frame you export is not the frame anyone sees

You export a 1080 by 1920 vertical video with your subject beautifully centred. On the platform, a username, a caption of three lines, a follow button, a music credit and a column of icons are drawn over it. The bottom quarter and the right-hand strip are gone. Your subject's chin is behind a share button.

Nothing has gone wrong technically. The file is correct. The **safe area** — the part of the frame that is reliably visible — is much smaller than the frame, and it was never in your editing preview.

The numbers, held loosely

As a working starting point for full-screen vertical video on the current crop of short-video apps: keep anything that matters out of roughly the **bottom 20 per cent** and the **right-hand 10 to 15 per cent**, and leave a margin at the top for the status bar and any header.

These figures move. Platforms change their interfaces without announcement, they differ between iOS and Android, and they differ between phone sizes — a tall phone with a notch and a short one with a home button crop differently from the same file. Treat published specifications as approximate and verify the only way that works: **post it privately, open it on a phone, screen-**

shot it, and look. Two minutes, and it catches the failure that no amount of reading catches.

Title safe is an old idea that is still right

Broadcast television has had safe areas for seventy years. Cathode-ray televisions overscanned — they deliberately drew the picture slightly larger than the visible screen so that ageing components never revealed a black edge — and the amount varied by set. Producers therefore worked to two nested boxes: **action safe** at about 93 per cent of the frame and **title safe** at about 90 per cent, with text inside the inner one.

CRTs are gone and the convention survived, because the underlying problem did not: the frame is still cropped by something you do not control, it is just an app's chrome now rather than a television's bezel. Every editor — DaVinci Resolve, Kdenlive, Shotcut, CapCut, Premiere — can draw these guides, and turning them on is a single setting.

The crop chain is the harder problem

A single shot commonly has to survive several different crops of itself:

- **9:16** full-screen vertical, 1080 × 1920.
- **4:5** in a feed, 1080 × 1350 — the tallest shape most feeds allow, and therefore the one that occupies the most screen.
- **1:1** square, still used for grids and some previews.
- **16:9** landscape for a website, a slide or a television.
- A **thumbnail** crop, often square and often taken from the centre.

Automatic reframing tools will do this for you and will do it badly, because they track a face or a saliency estimate and have no idea which part of the frame carries the meaning.

The robust approach is to compose for the **intersection**: the region that survives every crop you will need is a tall-ish central box, and anything load-bearing goes in it. Shoot slightly wider than the tightest crop so there is material to

give away, and decide the crops before you shoot rather than discovering them at export.

The crop chain, and the box that survives it

Shapes one shot must carry

- 9:16 full screen, 1080 by 1920
- 4:5 in a feed, 1080 by 1350
- 1:1 square for grids
- 16:9 for a site or a slide
- A square thumbnail from the centre

What survives all five

- A tall box near the centre
- Nothing in the bottom fifth
- Nothing in the right-hand strip
- Text kept in the middle third
- A wide master in the archive

Automatic reframing tracks a face or a saliency estimate and has no idea which part of the frame carries the meaning. Composing for the intersection is what makes a new shape an export rather than a reshoot.

Text is the first casualty

Text placed near an edge is the thing that gets cut, and cut text looks like incompetence rather than like a platform quirk. Two habits:

1. **Keep on-screen text in the middle third vertically**, not at the bottom where captions and interface live.
2. **Set it larger than you think**. A phone held at arm's length on a bus is not a monitor at 60 centimetres. If you cannot read it on your own phone at a glance, it is too small, and your eyes know the words already.

Doing it free

DaVinci Resolve's free version, Kdenlive and Shotcut all handle multiple aspect ratios, safe-area guides and reframing on a modest laptop. CapCut runs on a phone and does vertical work competently. For stills, GIMP, Krita and the browser-based Photopea all let you set up guides once and save them as a template — which is the real time-saver, because the setup is the cost and you only pay it once.

The limitation

Everything above is a moving target that you do not control and cannot appeal. A platform can change its interface overnight and turn a year of correctly

composed work into work with a button on its subject. There is no specification you can hold them to.

The defence is structural rather than technical: **keep a master at the widest sensible framing**, so a new crop is an export rather than a reshoot. Archive the project file, not only the finished video. The people who get hurt by a layout change are the ones who only kept the delivered file.

BEFORE YOU MOVE ON

Why does automatic reframing from 16:9 to 9:16 so often cut the wrong part of a shot?

- A. It tracks faces or visual salience, which is not the same as knowing which part of the frame carries the meaning.
- B. Reframing algorithms sample only the first frame, so they cannot follow a subject that moves during the shot.
- C. Converting aspect ratios resamples the image, and the resulting softness makes the subject harder to identify.
- D. Platforms apply their own secondary crop after upload, so the reframing tool's output is cropped a second time before anyone sees it.

52 · 8 min

The still is what recruits the play

THE ONE THING TO KEEP

Every moving piece is recruited by a still — cover, poster, link preview, paused first frame — so set each one deliberately, shoot a still rather than grabbing a frame, and get the preview right before the link is shared because the card is cached.

Every moving piece is first a stationary one

Before anybody watches your video, they see a single frame of it, in at least four places:

- the **cover** in a grid or profile,
- the **poster** in a player before it starts,
- the **link preview** when somebody shares the URL,
- the **first frame** as it scrolls into view, paused or muted.

In each case the still is doing the recruiting. A piece with a good opening thirty seconds and a black first frame has a good opening thirty seconds that nobody reaches.

The default is usually the worst available frame

Left alone, most systems take frame zero, or a frame chosen by an automatic heuristic. Frame zero is very often the darkest frame in the piece, because editors fade in. Automatic selection tends to pick a sharp, well-exposed frame, which is not the same as a meaningful one — a crisp shot of an empty workbench beats a slightly soft shot of somebody's face by every criterion the machine has and by none that matter.

So set it explicitly, everywhere it can be set:

- Upload a **custom cover** wherever the platform offers one.
- Set `poster` on an HTML video element, so the player shows something before it loads.
- Set `og:image` for link previews, at about **1200 × 630**, which is the size most cards crop to.

Link previews are cached, which is why yours is still wrong

The first time a URL is shared, the platform fetches the page, reads its `og:` tags and stores the result. Change the image afterwards and shares keep showing the old one, sometimes for a very long time, because nothing re-fetches on your schedule.

The practical consequence: **get the preview right before the link goes anywhere**. Most platforms offer a debug or card-validator tool that forces a re-fetch and shows what they actually read, and using it before you publish costs a minute. Using it afterwards costs an apology.

```
<meta property="og:title"          content="The £90 tool that
halved our reject rate">
<meta property="og:description" content="A workshop, a jig, and
four months of measurements.">
<meta property="og:image"
content="https://example.com/cover-1200x630.jpg">
```

What makes a good recruiting still

It is not the prettiest frame. It is the frame that makes the piece's promise legible at the size it will appear.

- **One subject.** Two competing subjects at cover size is a blur.
- **A face with a real expression**, if the piece has one, because faces survive shrinking better than nearly anything else — people read expressions at resolutions far below what reading text requires.
- **Legible at a fingernail.** Shrink it to 150 pixels wide and look. Everything that disappears was decoration.

- **Contrast against both light and dark interface**, since you do not know which the viewer is using.
- **Not a frame from the fade-in.** Pull from a point where the piece is running.

Design a frame rather than finding one

The better habit is to plan for a frame rather than hunt for one. While shooting, take a deliberate still — the composition arranged, the subject placed, shot as a photograph. It costs thirty seconds on location and produces a cover that no frame grab will match, because video frames are compromises made for motion: a shutter speed chosen for movement, a moment chosen for what came after it.

For animated and graphic work the equivalent is to build the cover as its own artboard rather than exporting frame one.

The line that matters

A cover is a promise. The same discipline applies here as to a headline: it may withhold, and it may frame, and it may not depict something the piece does not contain.

The specific failure is a cover showing a moment that never happens — a reaction, a result, a dramatic instant assembled or generated for the cover alone. It works once. It is also, for synthetic imagery, increasingly a labelling question, and the duties differ by country and keep moving, so that is a matter for a lawyer where you are rather than a rule this page can give you.

The limitation

Optimising the still is optimising the front door, and it has a ceiling. A well-made cover on a piece that does not pay it produces a click, a few seconds, and an audience slightly less likely to click the next one. Platforms measure this directly — a high click-through paired with an immediate exit is a pattern their ranking systems are built to detect and demote.

So the cover is worth real effort and cannot rescue anything. It determines who arrives. The first second and the structure determine whether arriving was worth it.

BEFORE YOU MOVE ON

A site owner replaces a page's `og:image` after the link has already been shared widely, but old previews keep appearing. Why?

- A. Platforms fetch and store the tags the first time a URL is shared and do not re-fetch on the publisher's schedule.
 - B. Older shares embed a copy of the image at the moment of posting, so each existing post holds its own static copy.
 - C. Image files are served with long cache headers, so viewers' browsers continue showing the previously downloaded file.
 - D. Search engines index the original image and supply it to social platforms, which prefer an indexed version over a live fetch.
-

53 · 9 min

It has to work silent and reward being unmuted

THE ONE THING TO KEEP

Autoplay is muted by policy, so sound-on is an earned state — the silent version must carry the whole meaning and the sound must add something different — and a phone speaker reproduces almost nothing below about 300 Hz, so any effect you built from low frequencies is not there.

Why muted is the default, mechanically

Browsers and apps block audio that starts without a user action. This is not a stylistic preference of platform designers; it is a policy adopted after autoplay-

ing sound became the most complained-about behaviour on the web, and it is now enforced at the level of the browser's media engine. Audio requires a gesture — a tap, a click, an unmute.

The consequence is structural. **Sound-on is an earned state, not a starting condition.** Your piece begins silent, for everybody, every time, and the sound arrives only if something in the silent version persuaded someone to reach for it.

On top of that, a large share of viewing happens where sound is socially impossible: an office, a class, a shared room, a bus, next to a sleeping child.

The double build

The design rule follows directly and it is not "put captions on it". It is:

1. **The silent version must carry the whole meaning.** Not a hint of it — the whole thing. If the piece is unintelligible without audio, it is unintelligible.
2. **The sound version must add something,** or being unmuted was a waste of the viewer's decision.

Most work fails one or the other. A narration-driven film with subtitles bolted on fails the first: it is readable rather than watchable, because the meaning lives in words the viewer is now having to read while also looking at a picture. A silent-first piece built entirely from on-screen text fails the second: unmuting adds a music bed and nothing else, and the viewer has learned there is no reason to do it again.


The pieces that work put **different information in each channel.** The picture and on-screen text carry the argument. The sound carries texture, presence, rhythm, and the moments that only exist as sound — the machine starting, the room's noise stopping, somebody laughing mid-sentence.

The phone speaker throws away your bass

Here is the specific technical fact that ruins more mixes than any other. A phone's built-in speaker is a few millimetres across, and a driver that small

cannot move enough air to reproduce low frequencies. In practice, output falls away sharply below roughly 500 Hz to 1 kHz, and there is effectively nothing under about 300 Hz.

What a phone speaker returns of what you mixed



A driver a few millimetres across cannot move enough air to carry low frequencies. The rumble, the heartbeat and the warmth in a close male voice are not quietly reduced on a phone; they are absent.

So:

- A music bed whose weight is in the bass is, on a phone, a thin scratchy version of itself.
- A male voice recorded warm and close, with the low end that makes it sound rich on headphones, loses that body entirely and can end up sounding weaker than a thinner recording would have.
- A rumble, a heartbeat, a low drone: inaudible. The dramatic effect you mixed is not there.

The fix is not clever processing. It is **checking on a phone speaker**, held at arm's length, in a room with other noise in it. Free tools do the rest: Audacity for a high-pass filter to clear out energy nobody will hear anyway, and any editor's loudness meter to keep the level consistent.

Loudness, and why your piece is quieter than the next one

Platforms normalise loudness, measured in LUFS, so that one video is not dramatically louder than another. Targets vary and change, commonly landing somewhere around -14 LUFS for streaming. A piece mastered much louder gets turned down, and everything you crushed to make it loud stays crushed — you lose the dynamics and gain nothing.

A piece mastered much quieter gets turned up, which is usually fine, but a piece with a quiet voice and loud music gets normalised as a whole and the

voice stays buried.

The practical rule: **mix the voice to sit clearly above everything else, and leave the overall loudness alone.** The platform will handle the rest, and the one thing it cannot fix is a balance problem inside your mix.

The first second, silent

The opening has to answer *what is this and is it for me* with the sound off. That means motion, a face, a legible few words, or a striking frame — something that works as a picture.

A common and avoidable failure is opening on a person starting to speak. Silent, that is a stranger moving their mouth: no information, and the viewer leaves before the sentence that would have hooked them finishes.

The limitation

Designing for silence pushes work towards on-screen text, and on-screen text competes with the picture for the same eyes. A viewer reading cannot simultaneously study the image, and a piece that solves silence by covering itself in text has traded one failure for another.

The resolution is a matter of sequencing rather than volume: put text where the picture is simple, and let the picture carry the moments that need looking at. That is an editing decision made at the structure stage, which is why sound-off design belongs at the planning table and not at the export dialogue.

BEFORE YOU MOVE ON

A film-maker mixes a warm, bass-heavy voiceover that sounds rich on headphones. Why does it sound weak on a phone?

- A. Platform loudness normalisation reduces the overall level, which removes the perceived weight of the low frequencies.
 - B. Compressed audio codecs discard low frequencies first because they carry less perceptual information.
 - C. Phone speakers are tuned to emphasise speech frequencies, so a deliberately deep voice sits outside the range they boost and is pushed behind the music.
 - D. A speaker a few millimetres across cannot move enough air to reproduce low frequencies, so the warmth is simply not produced.
-

54 · 8 min

Nothing is ever seen on a white wall

THE ONE THING TO KEEP

Conspicuousness is a relationship with whatever is beside your work, not a property of it, so mock designs up inside a real screenshot of the feed in both light and dark themes — and remember that adjacency you do not control changes the tone of anything not self-contained.

You design in a gallery and publish into a street

You build a design on a clean canvas, centred, at full size, with nothing around it. Your client approves it the same way, in a PDF, full screen.

It will be seen at a third of that size, between a photograph of someone's lunch and a news headline about a disaster, on a phone held at an angle in changing light, for about a second and a half, by somebody scrolling with their thumb.

The design was never evaluated in the conditions it has to work in. This is the single commonest reason good-looking work underperforms, and the fix costs ten minutes.

Contrast is against neighbours, not against your canvas

A design's conspicuousness is not a property it has. It is a relationship between it and whatever is beside it.

A muted, tasteful, desaturated piece is striking in a feed full of saturated clutter and invisible in a feed full of muted, tasteful work — which is exactly what happens when a style becomes popular. Whatever everybody is doing is, by that fact, no longer distinctive.

This is why chasing the current look is self-defeating in a way that is easy to describe and hard to resist. The look that stood out did so because it was different from its surroundings; adopting it moves you into the surroundings.

Mock it up in situ

The method is embarrassingly simple and almost nobody does it.

1. Screenshot a real feed on a real phone — yours, or better, one belonging to somebody in the audience.
2. Paste your design into the screenshot at the right size, in GIMP, Krita, Photopea or any editor.
3. Look at it for one second and then look away.

What survives that one second is what your design communicates. Everything else is for you and the client.

Do it twice: once in the platform's light theme and once in dark. A near-white background vanishes into a light interface; a near-black one vanishes into a dark one, and a design with a transparent or matching edge stops having a boundary at all. This is why a visible border, or a background tone a few steps away from the interface's, is worth more than it looks.

Adjacency changes meaning, and you do not control it

What appears next to your piece alters how it reads, and you have no say in it.

A cheerful advertisement directly under a report of a disaster reads as callous, and neither party did anything wrong — the ranking system put them together. A joke beside a piece about the same subject played straight reads as mockery. The classic version of this is an airline advertisement beside an aviation accident story, which is why media buyers have had adjacency rules for decades.

You cannot prevent it. What you can do is notice that tone which depends entirely on surrounding context is fragile, and that a piece whose meaning is self-contained travels better. This is the same argument as writing self-contained interview answers, one level up.

Scale, and the one-second test

A feed is read at a fraction of your working size. Run the blur test from earlier in the course, but run it at the real dimensions: shrink the design to the width it will occupy on a phone — often 300 to 400 pixels — and ask what survives.

If the answer is "my headline", good. If the answer is "a colour and a vague shape", then a colour and a vague shape is your entire message, and you should choose which colour and which shape deliberately rather than by accident.

The limitation

Designing to stand out in a feed optimises for one moment — the scroll-past — and that moment is not the whole relationship. Work that wins the second is frequently work that is loud, high-contrast and slightly aggressive, and a body of work made entirely of that is exhausting and forgettable in aggregate.

There is also a large class of work where none of this applies. A piece somebody sought out, a page they arrived at deliberately, a document sent directly, a printed thing held in the hand: these are read in conditions closer to the ones you designed in, and optimising them for a feed makes them worse.

The question to ask before you apply any of this is **how did this person get here**. Strangers scrolling and people who came looking are different audiences in different conditions, and a piece cannot be maximally suited to both.

BEFORE YOU MOVE ON

A studio adopts the muted, desaturated style that has been standing out in its market. Why does the effect usually not transfer?

- A. Muted palettes have lower contrast, so they lose out to saturated work regardless of what surrounds them.
- B. Audiences associate a style with the studio that established it, so later adopters read as derivative and are discounted.
- C. The style stood out by differing from its surroundings, and adopting it makes it part of those surroundings.
- D. Platform compression handles desaturated images poorly, introducing banding that makes copied versions look cheaper than the original.

55 • 9 min

The system is optimising something, and it is not your work

THE ONE THING TO KEEP

Ranking systems maximise a proxy for a proxy, so optimising hard against a metric drifts the work away from what the metric stood for — optimise form, never the promise, treat algorithm folklore as survivorship, and own a route to your audience that nobody can switch off.

What a ranking system is actually doing

A recommender does not assess whether a piece is good. It predicts, for a given viewer, how likely they are to produce some measurable behaviour —

watch to the end, watch again, react, comment, share, not scroll away — and ranks by that prediction.

That measurable behaviour is a **proxy** for the thing the platform wants, which is time and return visits, which is itself a proxy for advertising revenue. There are several layers of substitution between the number being maximised and anything a person would recognise as value, and each layer is a place where the two can come apart.

Goodhart's law, and why it is not cynicism

The economist Charles Goodhart observed that a measure used as a target stops being a good measure. The mechanism is mundane: any proxy correlates with the real thing across the range of behaviour that existed when it was chosen, and optimising hard against the proxy pushes you outside that range, into the region where the correlation breaks.

Concretely:

- **Completion rate** correlates with satisfying pieces, and also with pieces that are short, or looped so the counter never completes, or structured to withhold the answer until the final second.
- **Comments** correlate with pieces people care about, and also with pieces people are angry at. An argument in the replies is indistinguishable, from a metric's point of view, from a conversation.
- **Watch time** correlates with good, and also with slow, padded and artificially delayed.

None of these optimisations require anybody to be dishonest. They happen by selection: the creators who accidentally do them get distributed more, so more people see and copy them, and within a year the format has drifted.

The division that keeps you honest

A workable line, and one you can actually hold under pressure:

Optimise the form. Never optimise the promise.

Form is everything about delivery that does not change what you said: the aspect ratio, the safe areas, the opening frame, the caption, the pacing, cutting a slack ten seconds, putting the specific noun first in the title. All of this makes it easier for the right person to find and follow you, and none of it costs the audience anything.

The promise is what the piece claims it will give. Inflating it — a cover showing something that does not happen, a title asking a question the piece never answers, an opening implying a revelation you do not have — buys a click by spending trust, and trust is the thing you cannot buy back.

The reason to hold this line is not moral, although it is also that. It is that the proxy is measured and trust is not, so the damage is invisible in exactly the place you would look for it.

Nobody knows the algorithm, including the people saying they do

Ranking functions are not published, they change frequently, they are personalised per viewer, and large parts of them are learned rather than written. What circulates as algorithm knowledge is folklore assembled from small samples by people who saw one thing work once.

This matters because the folklore is generated by survivorship. You hear the account of somebody whose video did well and who has an explanation for it. You do not hear from the thousands who did the same thing and did not, because they did not get a platform to explain it from. Any rule derived this way is untestable by you, because you cannot run the counterfactual.

What you can rely on are the durable, stated things: the piece has to work in the shape and conditions the surface imposes, the opening has to pay, and a viewer who leaves immediately after arriving is a signal every system on the market is designed to notice.

Distribution you do not own is a structural risk

A platform can change its ranking, its interface, its terms or its mind overnight, without notice, without explanation and without appeal. Accounts

are suspended by automated systems with no usable route to a human. This is not a rare event; it is a normal one.

The defence is unglamorous and it works: **own something**. An email list, a website, a phone number, a client relationship — anything where reaching your audience does not require somebody else's permission. Use the platforms for reach, and move the people who care onto something you control. A newsletter of two hundred people you can reach directly is worth more than twenty thousand followers on a service that can turn them off.

The honest limitation

None of this is an argument for ignoring distribution. Work nobody sees helps nobody, and the platforms are where the audience is, and refusing to learn how they present your work is a choice to be worse at your job.

The position that holds up: learn the mechanics thoroughly, meet the format's real constraints, and keep the promise fixed. Then judge yourself against something the platform is not measuring — whether anybody came back, whether the client got the outcome in the brief, whether the piece did the thing you wrote down at the start.

BEFORE YOU MOVE ON

Why does optimising hard for completion rate tend to degrade work, even when no one involved is being deliberately manipulative?

- A. Completion rate is measured per viewer, so it rewards pieces that appeal to a narrow audience rather than a broad one.
- B. Platforms weight completion above all other signals, so creators neglect every other aspect of quality in response.
- C. The measure correlates with satisfying work only across ordinary cases, and pushing hard against it moves you into cases where it does not.
- D. Completion is easily inflated simply by making pieces shorter, so the metric systematically drives every creator towards the shortest possible format regardless of subject.

Which number answers which question

THE ONE THING TO KEEP

Each stage of a piece has its own number and the retention curve is the only one that points at a place in the file, but small samples are mostly noise — so pick the success number before publishing, look for patterns across many pieces, and remember that the outcomes that matter most are frequently the ones nothing counts.

Most reporting measures the wrong stage

A piece passes through stages, and each stage has its own number. Reporting one number for the whole thing tells you nothing you can act on.

- **Impressions** — how many times it was shown. This measures distribution, not your work. It is the number most often reported and the least informative.
- **Click-through, or the first three seconds retained** — whether the promise attracted the right person. This measures the cover, the title and the opening frame.
- **Retention through the middle** — whether the structure held. This measures the piece itself.
- **Completion** — whether the ending was worth arriving at.
- **Saves and shares** — whether it had value beyond the moment. A share is somebody spending their own reputation on you, which is why it is the most expensive action a viewer can take and the most meaningful.
- **Comments** — ambiguous by construction. Agreement, confusion and hostility all produce comments and none of them look different in the count.
- **The outcome in the brief** — enquiries, bookings, sign-ups, the thing you wrote down before you started. This is the only one that answers the

question the work existed for.

Match the number to the stage. If the opening is the problem, completion will not tell you; if the ending is the problem, impressions certainly will not.

Reading a retention curve

The retention graph is the most useful diagnostic any platform gives away, and its shape is readable.

- **A cliff in the first few seconds** means the promise and the piece do not match — either the cover attracted the wrong people, or the opening did not answer *what is this and is it for me*. A very steep cliff with a high click-through is the signature of an inflated promise.
- **A steady gentle decline** is normal and healthy. People leave continuously; you cannot fix this and should not try.
- **A sharp dip at a specific timestamp** is the useful one. Something happened at 0:34, and you can go and watch 0:34. It is usually a slow passage, a moment where the piece restates something, a graphic nobody could read, or a claim that lost people's trust.
- **A rise** means rewatching. Either a loop, or something people went back for — a number on screen, a detail, a joke. A deliberate rise is a strong signal; an accidental one usually means something was unreadable the first time.


This one curve, looked at honestly, is worth more than every other number on the page, because it points at a place in the file rather than at a mood.

Your sample is almost certainly too small

Here is the part that saves people from a great deal of wasted effort.

A video with forty views has a completion rate that is nearly meaningless. The difference between 38 per cent and 45 per cent on samples that size is ordinary variation; you can watch the same piece behave differently on two days for reasons that have nothing to do with the piece. Acting on it — changing your style, abandoning a format — is responding to noise.

Two retention curves from the same running time



A cliff in the first few seconds beside a high click-through is the signature of a cover the piece did not pay. A dip at one timestamp is the useful shape, because you can go and watch that second.

What a small operator can do reliably:

1. **Look at patterns across many pieces**, not at one. Twenty pieces where the ones with a face in the cover consistently outperform is a signal. One that did well is an anecdote.
2. **Look at the curve rather than the totals**. A dip at a specific second is diagnostic even with few viewers, because you go and look at the thing rather than believing the percentage.
3. **Prefer the outcome number**, which is usually small but real. Three enquiries is three enquiries.

One number, chosen in advance

The brief already asked what success looks like. The discipline at this end is to hold to it: **pick the one number before publishing and write it down**, along with what result would make you change what you do.

Choosing afterwards guarantees you will find something that went up, because with a dozen metrics moving at once, something always did. That is not evidence. It is the same mistake as choosing the window that flatters a growth figure, arrived at from the other direction.

What these numbers cannot see

The honest limitation is about what is not counted, and it is large.

A piece watched in full by one person who then commissioned you registers as one view. A piece forwarded in a private message — where a great deal of actual sharing happens — is usually invisible to any dashboard you can see. A piece that changed somebody's mind and prompted no action is unmeasured by construction. Meanwhile a piece with excellent numbers can have produced nothing at all.

So the numbers are a diagnostic instrument for specific questions, not a verdict on the work. They are very good at telling you **where** in a piece something went wrong. They are poor at telling you whether the piece was worth making, and treating them as though they answered the second question is how people end up with a body of work that performs well and says nothing.

BEFORE YOU MOVE ON

A retention graph shows a sharp dip at 0:34 and an otherwise gentle decline. Why is this more useful than a low overall completion figure?

- A. Timestamped data is collected more accurately than aggregate figures, so the dip is more likely to be a real effect.
- B. Overall completion is affected by viewers who never intended to finish, so it understates how well the piece performed.
- C. A localised dip indicates a technical fault such as a stall or an encoding error, which is easier to fix than a content problem.
- D. A dip at a known second points at a specific passage you can go and watch, while a completion figure names no place.

CASE STUDY · MODULE 7

The carpentry reel with a button on the joint

A two-person furniture workshop in Coimbatore with 19,000 followers, posting build videos and receiving almost no enquiries.

The footage was good. Daylight from a workshop door, a steady phone on a clamp, a real bench. Exported at 1080 by 1920, correctly. And the workshop's best-performing reel produced two enquiries in a month.

Four separate failures were stacked on one file, and none of them was visible in the editor.

What the platform was drawing over the work

The compositions placed the hands and the joint in the lower third, which is where a person naturally puts the subject of a bench shot. On the platform, that strip carries the caption sheet, the account name and the audio label, and the right edge carries a column of buttons. The exported frame was correct and the visible frame was smaller, and nothing in the editing preview had ever shown the difference.

The cover was the second failure. The reel faded in from black, and the platform's default is frame zero, so the still recruiting every play was almost black.

The third: the explanation of the joint existed only in the spoken audio, and most plays were muted, because autoplay is muted by policy rather than by preference. Silent, the opening was a stranger's hands moving.

The fourth was inaudible by construction. The music bed had been chosen for its weight, which sat between 60 and 80 hertz. A phone speaker a few millimetres across returns almost nothing below about 300 hertz, so on the device nearly everyone used, the bed was a thin scratch and the drama built from it was not there.

What the numbers said, read as stages rather than as a score

Click-through was high. The thumbnail read HIDDEN JOINT over a tight crop, and people tapped it. Retention fell from 100 to 38 per cent within three seconds and then ran almost flat.

That pair has one reading. The cover attracted people and the opening did not pay it. A high click-through beside an immediate exit is the pattern every ranking system on the market is built to detect, so the two numbers were working against each other.

The decision, and what each option cost

Moving the work into the middle third was the change that mattered and there was no cheap version of it.

Re-shooting the build was not available; the piece was finished and sold. Cropping into the existing 1080-wide footage would reframe it at the price of resolution, and the grain of the wood was the thing the reel was about, so a softened crop that then went through the platform's re-encode would lose exactly the detail the viewer came for.

The third option was to change how the next build was shot: record wide at the phone's highest resolution, treat that as a master, and export crops from it. That cost storage the workshop did not have much of, and a noticeably slower edit on a five-year-old laptop.

Underneath all of it sat a structural problem nobody had raised. Nineteen thousand followers existed on a service that could change its interface, its ranking or its mind overnight, and the workshop had no route to those people that it owned.

The line they drew through the fixes

Every change on the list was a change to form: the aspect ratio, the safe area, the cover frame, the size of the on-screen text, the filter on the music, the specific noun at the front of the title. None of them altered what the piece claimed.

One proposed change was not form. Somebody suggested a thumbnail showing the finished cabinet from a build they had not filmed, on the grounds that it was the best photograph they owned. That is the promise rather than the delivery, and it is where a click is bought by spending the one asset that cannot be bought back. It was the only item on the list that was dropped.

The folklore was dropped as well. The workshop had been posting at a particular hour because a video had told them the algorithm preferred it. Ranking functions are unpublished, personalised and frequently changed, and what circulates as knowledge about them is assembled from the accounts of people whose piece happened to do well. The thousands whose identical experiment failed do not get a platform to explain it from, so the rule cannot be tested by anyone repeating it.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

They posted one reel privately, opened it on a tall phone and a short one, screenshotted both, and drew the safe box by hand over the frame. It was smaller than any published figure they had read.

From the next build onwards they shot a wide master and cropped from it. On-screen text moved to the middle third and roughly doubled in size. The cover was shot as a deliberate still during the build rather than pulled from the timeline. The spoken explanation was also written on screen. A high-pass filter at 120 hertz cleared out energy nobody would hear, and the voice was mixed to sit clearly above the bed.

Retention at three seconds went from 38 to 61 per cent across the following six reels. Six reels is a pattern rather than proof, and the report said so. Click-through fell, because the thumbnail stopped promising more than the piece contained.

Enquiries went from about two a month to about nine. Most of them did not come from the reels directly. They came from a WhatsApp number in the profile and a list of 140 people who had asked to be told when a batch was ready, which is the only route the workshop owns and the only one nobody can switch off.

The export was technically correct at 1080 by 1920. Why did the composition still fail?

The frame that is exported is not the frame anyone sees. The platform draws its own interface over the video, so the reliably visible region is smaller than the file and its boundaries change with app updates and phone shapes. Composing for the exported rectangle puts the subject where the caption sheet and the button column will be, and the only way to know the real boundary is to post it privately and look at it on a phone.

Click-through was high and retention collapsed at three seconds. What does that pair diagnose, and what does it rule out?

It says the cover and title attracted people and the opening did not answer what this is and whether it is for them. It rules out a structure problem in the body of the piece, because almost nobody reached the body. It is also the specific signature of an inflated promise, which most ranking systems treat as a negative signal rather than as a high-performing hook.

Why did a workshop with 19,000 followers treat a list of 140 people as the more valuable asset?

Reach on a platform is held at somebody else's discretion: ranking, interface and terms can change without notice, and accounts are suspended by automated systems with no usable appeal. A direct route to people who have already asked to hear from you is small, durable and unmediated, and the enquiry count showed it was also where the work was actually coming from.

MODULE 7 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 A vertical reel, a feed image, a 150-pixel thumbnail, a printed flyer and a projected slide all fail in different ways. What single rule connects them?
 - A. Export at the highest resolution each platform will accept from you.
 - B. Contrast requirements are identical everywhere, so one exported file can serve all five surfaces.
 - C. Each surface needs a typeface chosen to suit how that surface renders type.
 - D. The smallest legible element is set by physical size divided by viewing distance.

- 2 A video is titled I built a house out of bamboo, and its thumbnail reads BAMBOO HOUSE over a picture of a bamboo house. Nothing is technically wrong. What is being wasted?
 - A. Text at thumbnail size is illegible anyway, so any words placed on the image are wasted.
 - B. Two channels carry one message, so the only two pieces of space say the same thing.
 - C. Repeating the same words across title and image is treated as keyword stuffing.
 - D. The title should never describe anything that the thumbnail already shows.

- 3 You have the published safe-area figures for a short-video app. Why still post the file privately and look at it on a phone?
- A. Boundaries change with app updates, with the operating system and with phone shape.
 - B. Published figures are deliberately approximate to discourage automated posting tools.
 - C. Editing previews render at a different colour profile from the finished app.
 - D. Safe areas are determined by the video codec and vary between export settings.
- 4 You replaced the og:image on a page a week ago and shared links still show the old picture. What is happening?
- A. The replacement image exceeds the size limit and is being silently rejected.
 - B. og:image is ignored by most platforms unless og:title is set alongside it.
 - C. The card was fetched and cached on first share, and nothing re-fetches on your schedule.
 - D. Preview cards are built from the first image found in the page body rather than from the tag.
- 5 A narration-led film has subtitles added at the end. Which half of the double build does it fail, and why?
- A. The sound version, because narration adds nothing beyond what the text already says.
 - B. Neither; subtitles satisfy both halves of the requirement at once.
 - C. Both, because subtitles are not captions and carry none of the non-speech sound the piece depends on.
 - D. The silent version, because the meaning now lives in text the viewer must read while watching.

- 6 Retention falls from 100 to 38 per cent within three seconds and then runs flat, beside an unusually high click-through. What does that pair record?
- A. A distribution problem: the system showed the piece to the wrong audience.
 - B. A promise the piece did not pay.
 - C. A pacing problem somewhere in the middle of the piece.
 - D. A sound problem, since almost every play begins muted.

MODULE 8

Judgement, and the work that accumulates

The last skills are the ones that decide whether the others get used: testing on people who are not you, translating feedback into problems, defending a decision from the brief, knowing where persuasion becomes manipulation, representing people you are not, using AI tools honestly, and building a practice that improves rather than repeats.

BY THE END OF THIS MODULE YOU CAN

Take a piece from private judgement into public responsibility — run a five-second and an explain-back test on strangers, separate a client's observation from their prescription, defend a decision from the agreed brief and concede what is not load-bearing, identify manipulation by whether it survives being explained, represent people you are not without using them as decoration, disclose AI use, and construct the feedback that deliberate practice needs

57 · 9 min

Three tests, no budget, twenty minutes

THE ONE THING TO KEEP

Comprehension can only be measured on somebody who does not already know, so run a five-second test, an explain-back and a first-click on five strangers — hand it over in silence and never ask whether they like it.

You cannot test it on yourself, for a reason you already know

You know what the design says, so you will always be able to find the message in it. The curse of knowledge applies to pictures exactly as it applies to sentences. The only instrument available is somebody who does not already know.

The good news is that the tests that work are cheap, fast, and need no software.

The five-second test

Show the design for five seconds. Take it away. Ask two questions:

1. **What was that about?**
2. **What did it want you to do?**

That is the whole test, and it directly measures the thing hierarchy exists for. If three people give three different answers, you do not have a hierarchy problem to investigate — you have the answer already.

Five seconds is generous. A poster gets less, a feed gets a second and a half. Use five because it produces an answer rather than a shrug, and remember you are grading against a harder standard than the one you tested at.

The explain-back test

For anything longer — a video, a page, a brochure — hand it over, say nothing, and afterwards ask them to explain it to you as though you had not seen it.

The gap between their account and your intention is the measurement. Three specific gaps are common:

- They get the subject and miss the point. The piece is clear and says nothing.
- They repeat your words back without content. They have absorbed the phrasing and not the claim.
- They fixate on a detail. Something in the piece is louder than the message, and now you know which thing.

The first-click test

For anything with an action in it: **where would you tap to do X?**

This one is worth knowing because the first click is unusually predictive. People who begin a task by going the right way finish it far more often than people who begin wrongly, and recovery from a wrong first move is much rarer than designers expect. If the first click is wrong, nothing downstream saves it.

Twenty minutes, five strangers, three tests



Never ask whether they like it. Liking is answered socially, and people say yes kindly about a design that told them nothing.

How to run it for nothing

- A family member who is not in your field. The requirement is not expertise, it is ignorance of this particular project.
- A neighbour, a shopkeeper, the person next to you in a queue. Ten seconds of somebody's time is not a big ask and most people are pleased to be

asked.

- A group chat: post the image, ask the two questions, do not explain first.
- Public critique communities, which are free and mostly generous, provided you show the brief so people are critiquing against the right target.

Five people is a good first batch. You will read that five testers find 85 per cent of problems, and that figure comes from a model assuming every problem is independent and roughly equally likely to be noticed. Real problems are not equally likely, and distinct groups of users hit different ones, so the honest version is: **five is enough to find the obvious failures, and not enough to conclude there are none.**

The two things that ruin a test

Explaining the design first. The moment you say "so this is about the workshop's new hours", you have installed the answer in their head and the test is over. Hand it over in silence. This is much harder than it sounds and it is the single commonest way people waste a tester.

Asking whether they like it. Liking is not a measurement and people answer it socially — they will say yes to be kind, and they will say it about a design that told them nothing. Ask what it said, what they would do, what they noticed first. Never ask what they think of it.

Watch the hands and the face

What people do is more reliable than what they report. Somebody who pauses, squints, turns the page back, or scrolls past and returns has told you something they will not say out loud, often because they have already blamed themselves for not understanding.

That self-blame is worth knowing about: when a design fails, most people assume they were being slow. They will not report the failure, and it is still a failure.

The limitation

These tests measure comprehension, and comprehension is not the whole job. They will not tell you whether the piece is moving, whether it will be remembered next week, whether it suits the client's standing, or whether it will work in a market you have not sampled. Five friends in one city are not your audience.

They are also blunt about degree. They reliably distinguish *clear* from *unclear* and are poor at distinguishing *good* from *better*, which means they are for catching failures rather than for choosing between two strong options. For that you need the brief, a judgement, and a willingness to decide.

BEFORE YOU MOVE ON

Why does explaining the design before a five-second test invalidate it, rather than merely helping the tester get started?

- A. The explanation takes time, so the tester has effectively had longer than five seconds to absorb the design.
- B. Testers become reluctant to contradict the designer once they know what was intended, so they under-report problems.
- C. Verbal explanation engages a different channel from visual comprehension, so the two interfere and produce a confused response.
- D. It installs the answer you were trying to measure, so their response reports your words rather than the design's.

58 · 9 min

"Make the logo bigger" is a symptom, not an instruction

THE ONE THING TO KEEP

People are reliable about what they noticed and unreliable about why, so never argue with an observation and never simply execute a prescription — ask what they were doing when they noticed it, batch the feedback, and spend your credibility only on what the message depends on.

Three different things arrive wearing the same clothes

When somebody responds to your work, what they say is one of three things, and telling them apart is most of the skill.

- **Observation** — what they noticed or failed to notice. "I didn't see the phone number."
- **Diagnosis** — their theory about why. "The phone number is too small."
- **Prescription** — their proposed fix. "Make the phone number red."

Non-designers are **excellent at observation** and unreliable at diagnosis and prescription. They are not being presumptuous by supplying all three; prescription is the only vocabulary they have. "Make it bigger" is what "I did not notice it" sounds like when you do not have the words *hierarchy* or *contrast*.

The rule follows directly: **never argue with an observation, and never simply execute a prescription.**

If they did not see the phone number, they did not see it. That is data and it is not negotiable. Whether the fix is size, weight, position, colour or removing the three things competing with it, is yours to determine — and the actual answer is frequently to delete something else rather than to enlarge the number.

The one question that does the work

"What were you doing, or looking for, when you noticed that?"

This reliably converts a prescription back into the observation underneath it. "Make the logo bigger" becomes one of several quite different problems:

What they said, and what is underneath it

The prescription

- Make the logo bigger
- Make the phone number red
- It needs more energy
- Can we try a different font

The observation

- I could not tell whose this was
- I did not see the phone number
- Nothing pulled my eye first
- Something looks wrong and I cannot say what

Prescription is the only vocabulary most people have. Never argue with an observation, and never simply execute a prescription, because the fix is often to delete something else.

- "I could not tell at a glance whose this was" — a brand-presence problem, possibly solved by colour or position rather than size.
- "My board will ask why the logo is small" — a political problem, and honesty about that changes the conversation entirely.
- "It feels unbalanced there" — a composition problem that has nothing to do with the logo.

Each needs a different fix. Enlarging the logo solves the second one and damages the design for the other two.

Feedback needs a shape, or it eats the project

A trickle of comments across a week is how a design dies by accumulation: each change is reasonable, no one is watching the whole, and the result is a committee artefact nobody chose.

What works:

1. **Batch it.** One round, all comments together, from everybody, by a stated date.
2. **Name who decides.** Not who comments — who decides, when two people want opposite things. Establish this at the brief, because discovering it

during round three is how projects stall.

3. **Set a round count.** "Two rounds of revisions included" is a normal professional term and it protects both sides.
4. **Ask for problems, not solutions.** Say it explicitly when you send the work: "tell me what does not work or what you did not notice, and leave the fixes to me." Most clients are relieved; supplying fixes is work they did not want.

Show the brief with the work

Attach one sentence when you present: *the one thing this has to do is X, for Y.* This changes the question the reviewer is answering from "do I like this" to "does this do X", which is a question they can actually answer well, and it makes irrelevant feedback visibly irrelevant without anyone needing to say so.

Feedback from other designers is a different instrument

Designers give you diagnosis, which is valuable, and they diagnose against craft standards rather than against your brief. A designer will tell you the spacing is inconsistent and will not tell you that nobody outside the trade understood the piece.

Use both and do not confuse them. **Strangers tell you whether it works. Designers tell you why it does not.** A critique from a peer who has not read the brief is a critique of the craft alone, and craft is not what the piece was for.

The limitation, which cuts both ways

This framework can become a way of never being wrong. If every disagreement is reclassified as the client's inability to diagnose, you have built a machine for dismissing them, and you will be wrong sometimes in exactly the cases where it matters.

Two honest checks. First: clients often know things you do not — about their customers, their market, their history, the piece that failed three years ago for reasons nobody wrote down. When a client insists on something that seems

wrong, the question to ask is what they know that you do not, and sometimes the answer is substantial.

Second: credibility is a finite budget. Spend it defending the things the message depends on, and concede the rest quickly and without sulking. A designer who fights over a typeface has nothing left when the client wants to add a fourth message to a poster, which is the fight that actually matters.

BEFORE YOU MOVE ON

A client says "make the phone number bigger" and the designer instead deletes two competing elements, leaving the number the same size. Why can this be the correct response?

- A. Clients rarely mean their suggestions literally, so a designer should generally treat stated changes as approximate.
- B. Smaller elements are cheaper to remove than to enlarge, so the fix that touches least of the layout is preferred.
- C. The reliable part of the feedback was that the number went unnoticed, and noticing depends on what surrounds it as much as on its size.
- D. Enlarging any one element disrupts the type scale, and maintaining a consistent scale across the whole system takes priority over accommodating an individual client request.

Argue from the document, not from taste

THE ONE THING TO KEEP

Defend decisions from the agreed brief rather than from taste, in four sentences — goal, mechanism, evidence, cost of the alternative — then concede everything the message does not depend on, and when overruled state the risk once in writing and execute it properly.

Why "I think it looks better" always loses

When a designer defends a decision on taste, they have moved the discussion onto ground where the client's opinion is exactly as valid as theirs — and the client is paying. That argument cannot be won and should not be.

The brief exists so that you never have to have it. A decision defended from an agreed document is a decision defended from something the client already said yes to.

"We agreed the one thing this has to do is get people to the open day. The date is the only element that survives at arm's length. Enlarging the logo means the date stops being the only one, and then nothing is."

That is not a matter of opinion. It is an argument against a shared standard, and the client can engage with it — by disagreeing about the goal, which is a conversation worth having, or by accepting it.

The shape of a good defence

Four sentences, in this order:

1. **The goal, restated from the brief.** "The one thing is the open day."
2. **The mechanism.** "Ranking these three levels is what makes one thing survive a glance."

3. **The evidence, if you have any.** "Five people I showed it to all said 'open day, Saturday'."
4. **The cost of the alternative, stated plainly and without drama.** "We can enlarge the logo. The likely effect is that people remember the brand and not the date."

Then stop talking. The commonest failure at this point is continuing to argue after you have made the argument, which reads as anxiety and invites the client to reassure you by finding a compromise neither of you wants.

Show the rejected option once

Showing the alternative you considered and set aside is genuinely useful: it demonstrates the decision was made rather than defaulted to, and it lets the client see the trade-off rather than take your word for it.

Do not show three options where two are deliberately weak in order to steer the choice. Clients recognise it more often than designers think, and once they have recognised it every subsequent presentation is read as a manipulation. The same judgement applies here as to your audience: a technique that works only while hidden is one to be suspicious of.

Concede fast on everything that is not load-bearing

Go through the feedback and mark each item: **does the message depend on this?**

Most items will not. The typeface, the exact blue, the corner radius, the order of two items of equal weight — concede these immediately and pleasantly. It costs the work almost nothing and buys you the standing to hold the one or two things that matter.

A designer who defends every decision is not principled; they are unable to distinguish, which is a competence problem rather than an integrity one, and clients read it correctly.

When they overrule you

They are allowed to. It is their business, their money and their risk, and they may know things you do not.

What to do:

1. **State the risk once, clearly, in writing.** "Noted — adding the fourth message means the open-day date is unlikely to be what people remember. Happy to proceed."
2. **Then do it properly.** A grudging execution is worse than either option and it is visible. If you are doing it, do it well.
3. **Keep both versions.** Sometimes the result comes back and the conversation reopens, and having the other version ready is worth more than having been right.

What not to do: sabotage, sulk, execute it badly to prove a point, or bring it up afterwards. All of these are more damaging to you than the design decision ever was.

The limitation

Arguing from the brief only works if there is a brief, it was written down, and it was agreed. Without that, every discussion is taste against taste, and you will lose most of them for the structural reason above.

There is also a category this cannot reach. Sometimes the real objection is not about the work at all: somebody needs to be seen to contribute, somebody's superior has an opinion, somebody was not consulted earlier and this is where it surfaces. No argument about hierarchy addresses any of that.

The move is to recognise it rather than to keep making design arguments at a problem that is not a design problem. Asking directly — "is there someone else who needs to be comfortable with this?" — is usually met with relief, and it converts an unwinnable debate into a scheduling issue.

BEFORE YOU MOVE ON

Why is presenting three options, two of them deliberately weak, a worse tactic than presenting one decision with its rejected alternative?

- A. It only works while it is unrecognised, and once seen it makes every later presentation read as a manipulation.
 - B. Producing three options costs time that could be spent improving the chosen one, so the approach is inefficient.
 - C. Clients asked to choose feel ownership of the outcome, which makes them harder to overrule later in the project.
 - D. Weak options can be chosen by accident, leaving the designer committed to work they deliberately made bad.
-

60 • 10 min

If explaining it would stop it working, it is manipulation

THE ONE THING TO KEEP

Persuasion supplies reasons a person could evaluate while manipulation exploits a process they cannot inspect, so ask whether the technique would survive being explained — and remember that a fake scarcity claim works once per person and the damage never appears in the conversion figure.

A line you can actually apply

Everything in this course is persuasion technique, and persuasion is not a dirty word: ranking a design so one message survives, opening with a promise, showing a real cost, choosing a photograph that carries information. A client hiring you to communicate is hiring you to persuade.

So the question is where it becomes something else, and "you know it when you see it" is not a usable standard when a client is asking you to add a count-

down timer on Thursday.

A test that holds up:

***Persuasion gives the person reasons they could evaluate.
Manipulation exploits a process they cannot inspect, and it works better when they do not know it is happening.***

The operational version: **would this still work if you explained it to them?**

- "We made the date the biggest thing so you would remember it." Still works. That is design.
- "We put a fake countdown there so you would panic." Stops working, and they are angry. That is manipulation.

It is not a perfect boundary — some honest techniques weaken slightly when explained, and some manipulations survive being named. It is good enough to decide almost everything you will meet.

Would it still work if you explained it

	The claim can be checked	The claim cannot be checked
Persuasion	Two Europeans. Kind, date is biggest so you remember it.	
Manipulation	Drip pricing. Real time manipulation. A timer that resets on reload.	

The boundary is not perfect, and it decides almost everything you will meet. What these patterns share is that they work once per person, and the damage never appears in a conversion figure.

The common patterns, and why each works

1. **False urgency and scarcity** — "only 2 left", a timer that resets on reload. It works on loss aversion: a potential loss weighs more than an equivalent gain, so a closing window produces a decision faster than a reason would. What makes it manipulation is not the timer, it is that the claim is unverifiable and usually untrue.

2. **Confirmshaming** — the decline option reading "No thanks, I don't want to save money." It attaches a small social cost to a private click, borrowing a mechanism built for actual social situations where nobody is watching.
3. **Asymmetric effort** — one click to subscribe, a phone call during office hours to cancel. Each step is individually defensible and the asymmetry is the whole design.
4. **Drip pricing** — a low price first, fees revealed at step four. The first number anchors, later numbers are judged as deviations from it, and the person has already invested effort by the time the real figure arrives.
5. **Preselected consent** — boxes ticked by default. Defaults are enormously powerful because they require no decision, and a default chosen for the business rather than for the person converts inattention into agreement.
6. **Engineered outrage** — a framing chosen because anger spreads, not because it is the most accurate framing available. The visual version is a thumbnail or crop selecting the least representative instant.

The visual craft in each case is often excellent, which is the uncomfortable part. These are well-designed.

The regulation is moving and differs by country

Regulators have taken an interest. The European Union's Digital Services Act contains provisions aimed at manipulative interface design; consumer-protection authorities in several countries have brought enforcement actions over subscription traps and hidden fees; India's consumer authority has published guidelines naming a list of specific dark patterns.

The landscape differs by jurisdiction, is actively changing, and none of this is legal advice. What is durable regardless of the law: these patterns work once per person. The countdown that turned out to be fake is a customer who now discounts everything on the page, and you cannot detect that in a conversion figure because it shows up as people who never came back.

The grey area is real and you will live in it

Some cases are genuinely unclear, and it is more honest to say so than to pretend the test resolves everything.

- **Real scarcity, stated.** Twelve seats genuinely exist. Saying so is information.
- **A real deadline.** The open day is on Saturday whether anyone books or not.
- **A default that helps the person.** A safe setting preselected is still exploiting the default effect, in their favour.
- **A thumbnail choosing the most striking true moment.** Selection is unavoidable; there is no unselected frame.

Where these sit depends on whether the claim is true and whether the interest served is the audience's. That judgement does not go away, and a technique is not laundered by being on the acceptable side of a rule.

What to do when a client asks for one

They usually are not being sinister. They saw it work somewhere and it is now on the list.

What works better than a lecture: name the risk in their terms, and offer the honest version of the same thing. "A fake timer converts a bit better this week and burns the people who notice. If the deadline is real, let us say the real one — that is the same mechanism and it survives being checked."

If they insist, you have the ordinary professional decision about what you will put your name to. It is worth knowing in advance where your line is, because it is much harder to locate at four in the afternoon on the day of a deadline.

BEFORE YOU MOVE ON

A countdown timer on a page shows a real deadline that exists regardless of who visits. Why does this sit on a different side of the line from a timer that resets on reload?

- A. A real deadline can be verified against the business's other materials, so customers are less likely to complain about it.
 - B. The real one states a fact the person could act on, and explaining it to them would not stop it working.
 - C. Resetting timers are prohibited in several jurisdictions, while genuine deadlines are not, so the legal exposure differs.
 - D. A genuine deadline creates urgency through information rather than emotion, so it does not rely on loss aversion at all.
-

61 • 9 min

Somebody is in this piece, and it is not you

THE ONE THING TO KEEP

Work about people you are not fails through the single story, decoration, backdrop and composite — ask whether the subject would recognise it as fair, consult early and pay for it, and remember that consultation surfaces what you did not know rather than transferring the responsibility.

The commonest failures, named

Most work about other people goes wrong in a small number of recognisable ways, and naming them makes them avoidable.

- **The single story.** One narrative stands in for a whole group, so every piece about a place is about poverty, or struggle, or resilience. Each individual piece may be true. The aggregate is a caricature, and the people in it are the ones who have to live next to it.

- **The poverty aesthetic.** Desaturating a photograph, shooting from above, choosing the most degraded surroundings — a set of visual choices that signal *hardship* and were made by somebody who was leaving afterwards.
- **Decoration.** A person included for their visible difference and given nothing to do. In casting this is a face in the third row; in writing it is a name with no line.
- **The backdrop.** Somebody's home, work or ceremony used as texture behind a story about someone else.
- **The composite.** A character assembled from several real people and presented as one, which resembles nobody and cannot be checked.

The question that catches most of it

Would the person in this piece recognise themselves, and would they recognise it as fair?

Not *would they be flattered*. Fair. It is the same test as for a photograph's crop, applied to a whole piece, and it catches the majority of ordinary failures before they reach anyone.

Consultation, done properly

The principle from disability rights — *nothing about us without us* — generalises cleanly. If a piece is about a group you are not in, people from that group should be involved before it is finished, not shown it afterwards as a courtesy.

What that means in practice:

1. **Early, not at sign-off.** Consultation after the edit is locked is a request for approval, not consultation, and everybody can tell.
2. **Pay people.** Expertise about a community is expertise. Asking for unpaid labour from people with less power than you is part of the problem you are trying not to reproduce.

3. **Accept a no.** If the answer comes back that the piece should not be made this way, the consultation has worked. Treating only agreement as a valid outcome means you were seeking cover.
4. **Show it before publication**, in full, with time to respond.

Language, names and the small things that are not small

- Spell names correctly, including diacritics, and ask rather than guess. A misspelt name in a caption is a statement about how much attention the subject was worth.
- Use the terms people use about themselves, which frequently differ from the terms institutions use about them, and which change.
- Ask about honorifics and forms of address rather than defaulting to the convention you grew up with.
- Avoid the passive constructions that remove agency — *were given*, *were helped*, *benefited from* — unless that is genuinely what happened.

Casting and stock imagery

When you are choosing images rather than making them, the same failures appear in a different form. A stock library will happily supply you with a diverse group of people arranged around a laptop, none of whom has anything to do with your subject.

The better path is the one already argued: photograph the actual people, with consent, doing the actual thing. If that is impossible and stock is genuinely the answer, choose images where the person is doing something specific rather than modelling an attribute, and check the licence conditions about portraying identifiable people in a negative context.

The limitation, which is important

Consultation is not a permission slip, and this is where well-intentioned processes go wrong.

One person does not speak for a group. Two people from the same community will frequently disagree, and there is no arrangement of consultants that converts a contested judgement into a safe one. A consultation used as protection — so that responsibility can be transferred if criticism arrives — is worse than none, because it spends somebody's time and name on your defence.

What consultation actually does is narrower and still worth a great deal: it surfaces things you did not know, it catches errors of fact and framing, and it gives people a say in how they are shown. The decision stays yours, the responsibility stays yours, and if the piece is wrong you cannot hand it to the people you asked.

The other limitation is worth saying plainly: sometimes the right call is that this is not your story to tell, and the better contribution is to help somebody inside it tell their own. That is a real answer, not a failure of nerve, and it is available more often than people acting in good faith tend to consider.

BEFORE YOU MOVE ON

Why is consulting a community representative after the edit is locked worse than not consulting at all?

- A. Late changes are expensive, so the consultation cannot lead to meaningful revision and everyone knows it.
- B. A single representative cannot speak for a community, so the exercise produces an unrepresentative view.
- C. Showing unfinished work builds trust, and presenting a finished piece signals that the maker never intended to collaborate in the first place.
- D. It spends the person's time and name on protecting the work rather than improving it.

62 · 10 min

What the tool can be trusted with, and what it cannot

THE ONE THING TO KEEP

Sort AI tasks by what a mistake costs and whether you would notice, because these tools fail fluently and remove the uncertainty signal you normally rely on — disclosure only survives if it is visible in the frame, the copyright position differs by country and is unresolved, and the name on the piece owns every frame.

Sort the jobs by what a mistake costs

Generative and machine-learning tools are now in every part of this pipeline. The useful way to think about them is not whether they are good, but **what happens when they are wrong, and whether you would notice.**

Low cost, error visible — use freely:

- Transcription of an interview, which you are going to read anyway.
- Background removal, masking and rotoscoping.
- Upscaling for a low-resolution source.
- Rough cuts and shot logging.
- Twenty headline variants to react against.
- A first draft of alt text, which you then correct.

High cost, error invisible — do not delegate:

- Anything evidential. A generated image is evidence of nothing, and putting one where a reader expects a photograph is a false claim regardless of intent.
- Facts, figures, names and dates. Fluent text is produced with the same confidence whether or not it is correct, so errors do not look like errors.

- Translation of anything that matters, where the characteristic failure is a confident reversal of meaning.
- Captions on sensitive material, where automatic recognition invents a plausible sentence from overlapping speech.

The pattern: these tools fail **fluently**. A human assistant who did not know something would sound uncertain. This does not, so the usual signal you rely on to know where to check is absent, and you have to check by category instead.

Disclosure, and what actually survives

Labelling requirements for synthetic media exist in several jurisdictions, with different scope and different timetables — India's rules for synthetic media, China's labelling requirements, the European Union's transparency provisions — and they are being amended as they are implemented. This is not legal advice and the answer for your case depends on where you are and where you publish.

The durable design fact, independent of any of it: **metadata does not survive a screenshot**. Provenance standards such as C2PA attach a signed record of how an image was made, and that record is stripped the moment somebody photographs their screen, re-uploads, or saves through a tool that does not carry it. If the disclosure matters, it has to be visible in the frame.

The copyright position is genuinely unresolved

Two separate disagreements are live, and neither has a settled answer.

Whether training on copyrighted work without permission is lawful. This is being litigated in several countries with different statutory starting points. The United States argues it through fair use, a flexible test applied case by case. The United Kingdom, India and other fair-dealing systems have enumerated purposes rather than a general test. Japan and Singapore have specific statutory provisions for text and data mining. The European Union has an exception with an opt-out mechanism. These are not different readings of

one rule; they are different rules, and cases decided under one say little about another.

Whether a generated image can be owned. Several jurisdictions have taken the position that a work requires human authorship, which leaves purely generated output outside copyright and therefore unprotectable — a commercial problem if a client expects exclusivity. Others have reached different conclusions, and the treatment of human-directed and human-edited work is unclear nearly everywhere.

Anyone who tells you this is settled is describing one country at one moment, or selling something. The practical stance for a working designer: **tell the client what is unresolved rather than making a promise you cannot keep**, and if exclusivity matters to the deliverable, say plainly that a generated asset may not be able to carry it.

The labour question, stated without resolving it

These tools remove work that used to be entry-level — stock illustration, basic retouching, simple animation, first-draft copy — and entry-level work is how people learned. Whether what replaces it is better, worse or merely different is genuinely contested, and confident predictions in either direction are not supported by anything.

What is not contested: it is happening, and the people it affects first are the people with the least security. That is worth saying in a course whose audience includes many of them, rather than being smoothed over.

The rule that does not change

The person whose name is on the piece is responsible for every frame of it, whatever made the frame.

Not the tool, not the model, not the supplier. If the caption is wrong, you published a wrong caption. If the hands have six fingers, you shipped six fingers. If the translated line says the opposite of what the client meant, that is your line now.

This single rule resolves most day-to-day questions without needing to know the law, because it puts the checking back where it belongs — on you, before it goes out, at full size, in the language it will be read in.

BEFORE YOU MOVE ON

Why is C2PA provenance metadata an incomplete solution for disclosing that an image was generated?

- A. It is attached to the file, so it is lost the moment someone screenshots or re-uploads the image.
 - B. Most platforms do not yet read the standard, so the signature is present but never displayed to viewers.
 - C. The signature can be forged by anyone with the generating tool, so a viewer cannot rely on its absence meaning anything.
 - D. Provenance records describe the editing history rather than the origin, so they cannot establish whether an image was generated or captured.
-

63 • 9 min

Take apart the thing that worked on you

THE ONE THING TO KEEP

Study the pieces that actually worked on you by writing out the promise, the beats, the joining words and above all what was left out — and construct the feedback that practice needs, because creative work returns an aggregated number too late to connect to any decision you made.

The habit that separates people who improve

Most people consume a great deal of visual work and learn almost nothing from it, because watching something is not studying it. The habit that changes

this is small: when something works on you, **stop and find out why while the effect is still fresh.**

The crucial part is *on you*. Not a piece that won awards, not a piece somebody told you was good — a piece that actually made you stop scrolling, or stay to the end, or believe something. That reaction is a measurement you have access to and nobody else does, and it is about to fade.

The deconstruction, in seven questions

Take one piece and answer these in writing. It takes fifteen minutes.

1. **What did the first second promise?** Write the promise as a sentence.
2. **Was it paid, and where?**
3. **What are the beats?** List them and write the joining word between each pair — *but*, *therefore*, or *and then*. Count the *and then*s.
4. **What is level one?** Squint or blur it. What survives?
5. **Where did my eye go first, and what put it there?** Size, contrast, position, isolation, direction.
6. **What was left out?** This is the hardest and the most valuable — the shot they did not use, the information they withheld until later, the thing a weaker version would have explained.
7. **What would I have done that they did not?** Then ask why they did not, and take seriously the possibility that they had a reason.

Question six teaches more than the other six together, because addition is visible and subtraction is not. Most of the craft in a strong piece is in what is absent.

A file of examples with reasons

Collecting work you admire is common and mostly useless. A folder of screen-shots is a mood board; it teaches taste by exposure and nothing else.

What works is the same folder with **one sentence per item naming the mechanism**: not "nice poster" but "one message survives the blur because everything except the date is at level three". Now the file is searchable by

problem. When you are stuck on hierarchy you can look up six examples of hierarchy with notes, rather than scrolling a wall of images hoping to be inspired.

Practice needs feedback, and creative work does not supply it

The conditions under which practice reliably produces improvement are well described: a specific sub-skill rather than the whole activity, repetition just beyond current ability, and **feedback close enough in time to connect to what you did.**

Creative work fails the third condition badly. You make a piece, it goes out, and what comes back is a number affected by timing, luck and distribution, weeks later, aggregated across every decision you made. There is no way to learn from that which choice was the good one. This is precisely why people can do this work for years and get no better — not from lack of effort, but because the loop never closed.

So you have to construct the feedback:

- **Redo with a known answer.** Take a piece that worked, recreate your own version from the same raw ingredients, and compare. The comparison is the feedback.
- **Predict before you look.** Write down what you expect the retention curve to do and where. Then look. Being wrong in a specific place is worth more than any amount of general viewing.
- **Ask one person for one thing.** "Where did you get bored?" answered honestly beats an hour of general praise.
- **Separate the sub-skill.** Spend a week on titles alone. Retitle twenty of your own old pieces. Nothing else.

Four drills that cost nothing

1. **Retitle ten pieces you have already published**, applying the front-loading rule, and note which original titles were doing nothing.
2. **Re-cut a 90-second film in your head from its own material**, writing the beat list and your joining words, then compare to the real

order.

3. **Redraw a bad diagram** — a leaflet, a form, a sign — by first writing the one thing it must convey.
4. **Blur-test five designs a day** for a week, from anywhere: a shop window, a packet, a poster at a station. Say what survives before you unblur.

The limitation

Deconstruction is reconstruction. You are inventing an account of somebody's intentions from the outside, and a good deal of what you attribute to a decision was an accident, a constraint, a client demand, or the only shot that came out usable. Your explanation will often be more coherent than what happened.

That is fine for learning and dangerous for arguing. The exercise builds your vocabulary for deliberate choices, which is its purpose. It does not entitle you to say what anyone did or why, and a portfolio of confident reverse-engineered explanations is a recognisable kind of wrong.

Hold your accounts loosely, keep collecting them, and judge them by whether they help you make the next thing rather than by whether they are true about the last one.

BEFORE YOU MOVE ON

Why do people often make visual work for years without improving, even while working hard?

- A. Repetition entrenches early habits, so long practice makes existing weaknesses more automatic rather than correcting them.
- B. Creative judgement is largely innate, so effort produces smaller gains in visual work than in technical fields.
- C. The response to a piece arrives late, aggregated and confounded, so it cannot be connected to any individual decision.
- D. Client work forces compromises, so practitioners rarely get to execute their own best judgement and therefore never test it.

64 • 10 min

Fifteen pieces with a point of view beat fifty with a filter

THE ONE THING TO KEEP

A portfolio is a prediction someone is making about you, so a small body of work with a visible repeated question is worth more than a large one with none.

What someone is doing when they look at your work

They are not admiring it. They are reducing risk. Somebody has to hand over money and a deadline to a person they do not know, and every piece they look at is being used to answer one question: *can I predict what I will get?*

That reframing explains a result that otherwise looks unfair. A portfolio of fifty unrelated pieces — a wedding video, a crypto logo, a restaurant menu, a game UI, three filters — reads as "I can do anything", which the viewer translates instantly to "I have no judgement, and I will do whatever you say". A portfolio of fifteen pieces that clearly share a question reads as a person who has decided something. The second is easier to hire because it is easier to imagine.

This is why the volume advice is only half right. Making a lot is how you get good. Showing a lot, undifferentiated, is how you stay unhired.

What a point of view actually is

It is not a niche, though it can be. It is a repeated question your work keeps answering.

- "I photograph small businesses the way they photograph their own families."

- "I make explainer motion for things people are embarrassed not to understand."
- "I design for shops that print in one colour on cheap paper."

Notice that each of these excludes work, and that is the function. A statement that excludes nothing communicates nothing — the same rule as the third lesson, applied to a body of work rather than a page.

You mostly do not invent this. You find it. Lay everything you have made in a grid — a Google Slides deck, a wall of printouts, a single big canvas in Photopea — and look for what recurs that you did not plan: a subject you keep returning to, a treatment you keep reaching for, the kind of problem you enjoy. That accidental pattern is usually the honest thesis, and it is better than one you decide on in the abstract because you have already proven you will keep doing it.

Show the deciding, not only the result

The pretty picture is evidence. The thing being bought is your judgement. So each case study, in four short parts:

1. **The problem**, in one sentence, in the client's terms.
2. **The constraint** — the budget, the two-day deadline, the fact that it had to work photocopied in black and white.
3. **What you decided, and why**. This is the part almost nobody writes and the part that gets work.
4. **What changed**. Numbers if you have them; an honest observation if you do not.

Three or four process images — the ugly boxes from lesson six, the rejected direction, the version the client chose over your favourite — do more than another hero shot. They demonstrate that a method exists.

Two honesty points. Do not present client work as yours alone if it was a team's; name what you did. And do not publish work under NDA, or a person's likeness, without permission — a portfolio is publication, and the release

you got for a campaign may not cover reuse on your own site. When in doubt, ask; the answer is usually yes and costs one message.

Doing this on nothing

- **Free hosting that works:** a Notion page, GitHub Pages, a Carrd or similar free tier, a Google Sites page. Behance and Dribbble are free and searched by clients.
- **A PDF is a portfolio.** A well-ordered twelve-page PDF, under 10MB, emailed on request, is completely respectable and loads on a bad connection.
- **On a phone only:** a deliberately sequenced Instagram grid plus a link. It is not a great portfolio, but it is a real one, and it beats waiting until you can afford a website.

What matters far more than the platform is order. The first three pieces determine whether anyone sees the fourth. Put your best work first, not your newest, and do not include anything you would not want to be hired to repeat — because that is precisely what including it invites.

The limitations, honestly

A point of view found too early is a costume. If you have made twelve things, you do not have a thesis yet, you have a preference, and announcing it will lock you out of the accidents that would have shown you something better. Make a lot of work first, badly, in public if you can bear it. Independent creative deals with the business side of getting paid while you do that.

And the thesis is allowed to change. It should change on the scale of years, not weeks. Changing it every month is not evolution; it is the fifty-pieces problem with extra steps.

The last honest note is about the work itself. What has changed in this trade is not that judgement stopped mattering — it is that execution got cheaper and faster for everyone, which raises the value of deciding what is worth executing. That decision is the whole of this course, and it is the thing a body of work is evidence of.

Do this today

Choose the fifteen pieces you would show, and only fifteen. Write one sentence that is true of at least ten of them. If you cannot, you have found your next six months of work — and if you can, that sentence goes at the top of the page, above everything else.

BEFORE YOU MOVE ON

A freelancer with fifty varied pieces online gets plenty of portfolio views but very few enquiries, and the enquiries that arrive are for low-budget work outside their interests. What does this pattern most likely indicate?

- A. Fifty pieces is too many to browse, so viewers give up before reaching the strongest work and the fix is trimming the count.
- B. The pricing is not stated, so viewers assume the freelancer is cheap and enquire accordingly.
- C. The work shows no repeated question, so viewers cannot predict what they would get and default to enquiring about small, low-risk jobs.
- D. The strongest pieces are buried further down, so reordering the page by quality will convert the same traffic into better enquiries.

CASE STUDY · MODULE 8

The countdown that was nearly real

A freelance designer in Indore, three days before a coaching centre's admission page goes live.

The request arrived on a Thursday afternoon, which is when these requests arrive. The centre's marketing lead wanted a countdown timer reading seats closing, and a line saying only 3 seats left.

The designer had the brief document in front of her. The real cut-off was the board's admission date of 30 June, six weeks away, and the timer was to be set to seventy-two hours and reset on reload. The figure of three had not been counted from anything; it had been chosen because it looked urgent.

The test that decides it

Everything on that page was persuasion, including the things she had already done: ranking the fee and the date above everything else, putting the specific noun first in the headline, and choosing a photograph of the actual classroom instead of a stock one. Persuasion is what a client hires a designer for.

The usable line is whether the technique survives being explained. We made the date the biggest thing so you would remember it still works when you say it out loud. We put a fake countdown there so you would panic stops working and makes the person angry, and the difference is not how exciting each one is. It is whether the reader is being given a reason they can evaluate or a process they cannot inspect.

A fake timer also fails a second, plainer test: the claim cannot be checked and is not true.

The decision, and the cost on both sides

Refusing was not free. The centre was about a fifth of her year's income, and the page would go live with a timer on it whether she built it or not, because somebody else would.

Agreeing was not free either, and the argument that actually moved the room was about the centre rather than about her. Its whole position was we tell you your child's real rank. The parents most likely to reload a page and notice a re-setting timer are the parents who compare four centres before choosing, which is the audience the positioning was built for.

She had one piece of evidence and used it. The previous season's page, which had carried a similar scarcity line, converted at 3.1 per cent, and a Google review naming the countdown was still the second review on the listing.

Then she offered the honest version of the same mechanism, which is the move that usually works better than a lecture. The 30 June deadline is real. The batch size of 60 is real. The current enrolment of 41 is countable and sits in a register that already exists. Printing the real count uses the same loss aversion and survives being checked.

That had a cost of its own, and it belonged to the director, not to her: a live count is a number that will look bad in a weak year, in public, and once it is up, taking it down is itself a signal.

How the argument was made

She used four sentences, in order, and then stopped talking.

The goal, restated from the brief: the one thing this page has to do is get comparing parents to book a visit. The mechanism: the position we agreed is that we tell people the real number, and this page is where that is either true or not. The evidence: last season's page and the review that is still second on the listing. The cost of the alternative, said plainly and without drama: we can put the timer up, and the likely effect is that the parents who compare four centres, who are the ones we are for, find out that we do not.

Continuing to argue after the argument has been made reads as anxiety and invites a compromise nobody wants, so she stopped.

She also conceded everything else in the same meeting and conceded it quickly. The marketing lead wanted a different blue, the corner radius changed and the testimonials moved above the fee. None of those touched what the page claimed, and holding any of them would have spent the standing she needed

for the one thing that mattered. A designer who defends every decision is not principled, and clients read it correctly.

STOP HERE

Before turning the page: what would you have done, and what would it have cost if you were wrong? Write it down. The point of the next section is to compare.

WHAT HAPPENED

The timer went. The page carried the real cut-off date and a line reading 41 of 60 places taken, updated nightly from the register.

Conversion came in at 2.7 per cent against the previous season's 3.1. It went down, and the designer had written in advance that it might, which is the difference between a stated limitation and an excuse.

What she also had was the centre's own call log. Of 63 callers, nine mentioned the seat count without being asked, and two said they had checked the page twice over a week to see whether it had moved. Checking twice is the behaviour a resetting counter destroys the first time somebody notices, and it is not visible in a conversion figure.

The director kept the live count the following season and the decision was written into the brief so it would not be re-argued every June.

One page, one season, no control group: the argument that held was never the number. It was that the mechanism still worked when it was explained.

Ranking the fee, front-loading the headline and using a real photograph are all persuasion. What separates them from the countdown?

Whether the technique survives being explained. Saying we made the date the biggest thing so you would remember it leaves it working, because the viewer has been given a reason they can evaluate. Saying we put a fake timer there so you would decide faster stops it working and makes them angry, because it relied on a process they could not inspect and on a claim that was not true.

The honest page converted lower. Why was that not an argument for the timer?

The conversion figure measures one visit by one person at one moment and cannot see what a false claim costs, which is people who noticed and never came back. That cost is real and unmeasured, so a comparison between the two numbers is not a comparison between the two options. The call log recorded the behaviour that matters here: people returning to check a count, which only happens while the count is believed.

She offered a replacement instead of refusing outright. Why is that usually more effective than explaining why the request is wrong?

Clients asking for these patterns are usually not being sinister; they saw one work somewhere and added it to a list. Naming the risk in their own terms and supplying the honest version of the same mechanism gives them a way to keep what they wanted without the exposure, which is a decision they can take. A lecture asks them to concede, and people defend positions they have been asked to concede.

MODULE 8 · TEST

Check what stuck

6 questions, one right answer each. This one is for you, not for a record: answers and the reason behind each are at the back. If two questions from the same module go wrong, re-read that module before the exam rather than after.

- 1 You are about to run a five-second test on five strangers. Which pair of habits would waste the session?
 - A. Explaining the design first, and asking whether they like it.
 - B. Testing five people rather than twenty, and not recording what was said.
 - C. Showing it on a phone, and not giving them long enough to look.
 - D. Using family members, and testing two versions in the same sitting.

- 2 A client says make the logo bigger. What is the rule for handling it?
 - A. Explain the hierarchy so they understand why the change cannot be made.
 - B. Execute it, since it is their brand and their money.
 - C. Never argue with an observation, and never simply execute a prescription.
 - D. Ask them to rank it against their other comments before you act on any of them.

- 3 Why does defending a decision with I think it looks better reliably lose the argument?
 - A. Clients have been trained by long experience to distrust a designer's aesthetic claims about a piece.
 - B. It is an informal register and reads as unprofessional in a client meeting.
 - C. It cannot be minuted, so it is forgotten by the time of the next round.
 - D. It moves the discussion onto ground where their opinion is exactly as valid as yours.

- 4 A client asks for a countdown timer that resets on reload. Which test decides whether this is persuasion or manipulation?
- A. Whether it is lawful in the countries the page will be read in.
 - B. Whether it raises conversion by more than it costs to build and to maintain afterwards.
 - C. Whether most people would say it was acceptable if they were surveyed.
 - D. Whether it would still work if you explained it to the person it is aimed at.
- 5 You consult people from a community your piece is about, before the edit is locked. What has the consultation actually done?
- A. Surfaced what you did not know and given people a say, while the decision stays yours.
 - B. Transferred responsibility for the piece to the people who were consulted.
 - C. Produced a single agreed position that can be cited later if any criticism of the piece arrives.
 - D. Replaced the need for individual consent from the people who appear in it.
- 6 Why is fail fluently the reason to sort AI tasks by what an error costs rather than by how good the tool is?
- A. Because these tools run more slowly on the tasks where accuracy matters most.
 - B. Because generated output is reliably wrong on names, numbers and dates.
 - C. Because the uncertainty signal you normally use to know where to check is absent.
 - D. Because licence terms generally forbid using generated material in evidential contexts.

EXAMINATION

Making It Say Something — final examination

This paper covers the whole course: extracting a real brief and naming the viewer's before state, ranking a design so one message survives a blur, joining beats with but and therefore, gathering material with consent and a licence you can point at, writing words that carry the picture, choosing type, colour and image for what they say, delivering into the surface and the system that carries it, and taking a piece from private judgement into public responsibility. You get 25 questions drawn from a larger pool, so a retake is a different paper. The pass mark is 70 per cent. There is no timer: read the question twice and take as long as you need. If you do not pass, you may retake after thirty minutes with fresh questions. A teacher can open the next course for you at any time regardless of this paper, and that is recorded as an override rather than as a pass.

On the website a sitting is 25 questions drawn from this pool and the pass mark is 70%. There is no time limit, there and here. Passing opens the next course in the track.

1 1. Before you open the software

You cover the most beautiful element of your own poster with your hand and the message survives intact. What have you learned?

- A. That element is decoration rather than load-bearing, which is allowed but should be known.
- B. That element is redundant and has to be deleted before the piece can be finished.
- C. The hierarchy is broken, because covering any element at all should damage the message somewhere.
- D. The message is too simple, since a single element should not be able to carry it.

2 1. Before you open the software

After a discovery call you type the answers back in your own words, in a shared document, and ask the client to confirm. What does that step actually buy?

- A. It satisfies the professional requirement to keep written records of client instructions.
- B. It slows the project down enough for the client to reconsider a weak request.
- C. It turns a vague conversation into an object you can point at three weeks later.
- D. It transfers responsibility for the outcome from you to the person who confirmed it.

3 1. Before you open the software

Writing one eighty-word portrait sharpens your decisions and narrows your sight at the same time. Which pair of defences addresses that cost?

- A. Write the portrait in the client's words, and confirm it with them before designing.
- B. Write two portraits when the audience splits, and name in one line who it is not for.
- C. Widen the portrait until it covers the whole audience, and check it against the demographic.
- D. Test the portrait on three customers, and update it after every round of feedback.

4 1. Before you open the software

A campaign is not moving anybody. You can spend the week on layout and language, or on evidence. Which stage of failure does better layout never fix?

- A. Attention, because a viewer who never took the piece in cannot be reached by any argument.
- B. Comprehension, because language and hierarchy are the only levers that act on it.
- C. Distribution, because no design decision changes how many people are shown the piece.
- D. Belief, because they saw it, understood it, and did not accept the claim.

- 5 1. Before you open the software

A client sincerely tells you their customers buy on craftsmanship. Reviews of their competitors say people are buying on price. Why does the discovery meeting not catch this on its own?

- A. Clients understate price sensitivity because it reflects badly on their own product.
- B. The questions were asked in the wrong order, so the answers were rehearsed.
- C. Every answer came from the same source, so there is nothing to disagree with.
- D. Discovery questions ask about the future, and buying behaviour is a fact about the past.

- 6 1. Before you open the software

A mood board of forty beautiful images tends to drift into copying even when nobody intends it. Which habit prevents the drift?

- A. Keeping the board to work from outside your own field and country.
- B. Writing under each item the one thing you are taking from it.
- C. Restricting the board to material that carries an open licence.
- D. Removing anything a client might recognise before you present it.

- 7 1. Before you open the software

Which of these has to be decided before production rather than after, if the campaign is to be attributable at all?

- A. Whether the outcome measure will be enquiries, walk-ins or bookings for the client.
- B. Whether the client considers the result good enough to commission a second piece.
- C. Which vanity measures the platform will report once the piece has been published.
- D. A separate phone number or code word, so a call can be traced to the poster.

8 2. Making one message survive

A restaurant menu, a train timetable and a spec sheet all carry many facts at once. What replaces the one-message rule for them?

- A. One structure: consistent groups, a clear order, and hierarchy applied inside each group.
- B. One message per page, with the remaining facts moved to a separate document.
- C. No hierarchy at all, since every item on a list of this kind is by definition equally important.
- D. A single hero element, chosen as the item the business most wants people to buy.

9 2. Making one message survive

A layout is mirrored for an Urdu edition and the result reads backwards to the people it is for. Which of the five emphasis levers failed to carry across?

- A. Isolation, because the amount of space around an element is read differently.
- B. Contrast, because reading direction changes which tones are perceived as strong.
- C. Weight, because bold and regular are not distinguished in all scripts.
- D. Position, because it assumes a reading direction and the strong entry moves.

10 2. Making one message survive

Your text passes the 4.5 to 1 contrast floor on a calibrated monitor. Why aim nearer 7 to 1 for work that will be outdoors or on cheap paper?

- A. Higher ratios are required by accessibility law in most countries for public material.
- B. Every real viewing condition reduces effective contrast, and none increases it.
- C. Ink on uncoated paper darkens, which reverses the ratio between text and ground.
- D. The 4.5 figure applies only to screens and has no equivalent for print at all.

11 2. Making one message survive

An auto-playing video sits beside a block of body text and the text is not being read. What explains the strength of that effect?

- A. Video is higher in contrast than type, so it wins on the strongest emphasis lever.
- B. Readers assume the moving picture contains the same information in an easier and faster form.
- C. Peripheral vision is disproportionately sensitive to motion and keeps re-capturing the eye.
- D. The video's audio competes with subvocalisation, which is how most people read prose.

12 2. Making one message survive

An address, opening hours and a phone number all have to appear and none can be deleted. How does grouping reduce the cost they impose?

- A. Three items placed close with consistent alignment are perceived as one object.
- B. A boxed group can be set at a smaller size than loose items of the same rank.
- C. Grouping moves them to level three, where relationships no longer have to be managed.
- D. Items in a group share a single alignment, which halves the number of decisions.

13 2. Making one message survive

You mirror a layout horizontally and edges you believed were aligned are visibly not. Why does the flip reveal what staring did not?

- A. Mirroring exposes rounding errors introduced by the software's transform tools.
- B. You can no longer read the text, so you see the composition as shape.
- C. The eye is more sensitive to alignment on the right than on the left of a page.
- D. Flipping changes the entry point, so the composition is scanned in a new order.

14 2. Making one message survive

The hand-drawn illustration you spent an afternoon on is competing with the headline. What is usually the correct move?

- A. Delete it, because an element that competes with the message has already failed.
- B. Keep it and enlarge the headline until the hierarchy is restored between them.
- C. Demote it to level three, where it adds texture without taking the message's place.
- D. Move it behind the headline at low opacity, so it reads as background rather than figure.

15 3. The shape of a story

A safety notice, an ingredients panel and a set of assembly instructions are all dramatised into three-beat stories. What goes wrong?

- A. The added structure takes up length, and the information no longer fits into the space that is available.
- B. Audiences arrive already wanting the information, so structure slows the scan and invites suspicion.
- C. Regulators require factual material to be presented without narrative framing.
- D. Three beats cannot carry more than one instruction, so most of the content is lost.

16 3. The shape of a story

A thumbnail promises a bamboo house and the video is eight minutes about the channel. Beyond the ethics, what is the arithmetic of that trade?

- A. The piece is demonetised, because platforms detect the mismatch between the title and the content.
- B. Nothing measurable changes, which is why the practice persists across every platform.
- C. It spikes the first metric and depresses the second: high clicks, short watch time, unfollows.
- D. The audience grows faster, but the new viewers are worth less to advertisers.

17 3. The shape of a story

You can swap beats two and three of a rough cut and nothing breaks. What does that tell you?

- A. They are not causally linked, so one of them is probably deletable.
- B. The edit is well balanced, since either order communicates the same content.
- C. The transitions are too soft, which is hiding the change between the beats.
- D. The beats are the same length, so the viewer cannot perceive the difference.

18 3. The shape of a story

You are compressing a ninety-second structure into fifteen seconds. What happens to the ask?

- A. It is dropped, because fifteen seconds cannot earn the right to ask for anything.
- B. It replaces the new state, since the result can be described in the caption instead.
- C. It moves to the front, so that viewers who leave early have still seen it.
- D. It goes inside the new state, with the number on screen while the result is shown.

19 3. The shape of a story

Why is a precise claim that turns out to be false worse than a vague one that turns out to be false?

- A. Precise claims are more likely to breach advertising rules in most jurisdictions.
- B. The precision that makes a true detail persuasive makes a false one checkable and quotable.
- C. Vague claims are forgotten quickly, so nobody ever notices when one of them turns out to be false.
- D. Precise claims attract a narrower audience, so the damage is concentrated.

20 3. The shape of a story

A brand film says there were difficult times and then says it now produces four hundred units a week. What gives the invented hardship away?

- A. The claim about units is unverifiable, so the whole passage reads as marketing.
- B. The hardship is placed before the success, which is the order that every fabricated arc uses.
- C. Struggle is described in general terms while the resolution is described in concrete ones.
- D. The piece uses the passive voice for the difficulty and the active voice for the result.

21 3. The shape of a story

An editor keeps trimming the final card to under a second because on the tenth viewing it feels interminable. What is the error?

- A. The card is too long, and the text should be shortened rather than the duration extended.
- B. Platforms trim trailing stills automatically, so the extra seconds are lost anyway.
- C. The loop makes duration irrelevant, since the piece restarts before the card is read.
- D. They already know what it says, and the viewer is reading it for the first time.

22 4. Getting the material

You are filming a market, a wedding or a real conversation. Why does a locked twelve-frame storyboard hurt that work?

- A. Observational shoots have no beats, so a board describes a structure that does not exist.
- B. You chase your drawings instead of watching, and the best moment happens during a setup.
- C. Boards are drawn from memory of similar events, which is unreliable for unfamiliar settings.
- D. The board has to be re-approved by the client whenever the real situation diverges from it.

23 4. Getting the material

A subject gives you a date, a quantity and a sequence with complete confidence, and two of the three turn out to be wrong. What should you conclude?

- A. They are covering something, and the rest of the interview needs to be treated sceptically.
- B. The question was leading, which produced an answer shaped to fit your expectation.
- C. Memory reconstructs rather than replays, so anything checkable has to be checked.
- D. The interview was too long, and accuracy falls once a subject has been speaking for an hour.

24 4. Getting the material

You cover the picture and the meaning is unchanged; you cover the words and the meaning is unchanged. What does the image cost you?

- A. Nothing, since a redundant image reinforces the words and improves recall of them.
- B. Only load time, which matters on slow connections and not otherwise.
- C. It is doing continuity work by default, which is a legitimate third job.
- D. One channel is running twice while the viewer's attention pays for two.

25 4. Getting the material

You are filming staff for an employer's video. Why is a verbal yes in front of the manager a weak form of consent?

- A. Verbal consent is not admissible, so it has to be collected on paper in every country.
- B. Employees rarely understand what a corporate video will be used for afterwards.
- C. Consent given under an employment relationship is under pressure by construction.
- D. An employer already holds a licence to staff likenesses, so consent is redundant.

26 4. Getting the material

You publish a photograph with the location withheld, uploaded to your own website. What is the specific failure to check for?

- A. The EXIF block, which usually carries exact coordinates, the device and a timestamp.
- B. The colour profile, which can identify the camera model from its rendering.
- C. The file name, which most cameras derive from the shooting location.
- D. The compression signature, which links the file to other images from the same device.

27 4. Getting the material

A colleague tells you that short excerpts of any work are fine because of fair use. What is the accurate position for someone without a legal budget?

- A. Fair use is an international standard, so the same four-factor test applies everywhere.
- B. Excerpts under thirty seconds are exempt in most systems, which is where the rule comes from.
- C. Platform takedown systems apply the law directly, so passing their checks means that you are safe.
- D. Countries have different, non-equivalent exceptions, decided after the fact rather than before.

28 4. Getting the material

Your audit shows beat six has no material because the event it claims never happened. How does that differ from beat three, which needs a shot you can still film?

- A. Beat six is an asset gap and beat three is a structure gap, so beat three is the one rewritten.
- B. Beat six is a structure gap: no amount of footage repairs it, so the beat is rewritten.
- C. Both are asset gaps, and the difference is only in how expensive they are to fill.
- D. Beat six can be filled with generated material, since no real footage of it can exist.

29 5. Words that carry the picture

Two headlines carry the same information. One is How we cut our reject rate in half using a ninety-pound tool. Why is The ninety-pound tool that halved our reject rate the stronger of the two?

- A. It is shorter, and shorter headlines are read more often across every surface.
- B. It opens with a number, which search and social systems weight more heavily.
- C. Its most specific element sits in the first four words, so it survives truncation.
- D. It uses the active voice, which is more direct and is therefore more persuasive to a reader.

30 5. Words that carry the picture

Every mail client shows a snippet after the subject line. Most senders spend it on View this email in your browser. What is the opportunity?

- A. A second headline of roughly equal length, which can pair with the subject.
- B. A place to repeat the subject line, which raises recognition in a crowded inbox.
- C. A field that is indexed by search, unlike the subject line itself.
- D. A space for the unsubscribe link, which keeps it out of the message body.

31 5. Words that carry the picture

Your narration contains the figure thirty-one point four per cent, and the precision matters. What is the correct handling?

- A. Keep the exact figure spoken, and slow the delivery so the listener can absorb it.
- B. Round it in the narration and put the exact figure on screen.
- C. Repeat the exact figure twice, since a listener cannot rewind to hear it again.
- D. Move it into the caption and remove the claim from the narration entirely.

32 5. Words that carry the picture

Automatic captions are usually faster to correct than to retype, and they fail in a predictable pattern. Which failure is the dangerous one?

- A. Proper nouns, because the model has never encountered the client's surname.
- B. Currencies and measurements, because the units are dropped or misheard.
- C. Overlapping speech, because it produces a fluent sentence nobody actually said.
- D. Accents outside the training distribution, because whole passages come back empty.

33 5. Words that carry the picture

An image already has a caption that describes it fully. What should its alt text do?

- A. Repeat the caption, so that the description is available in both places.
- B. Begin with Photo of, so a screen-reader listener knows what kind of element it is.
- C. Editorialise, since a listener cannot see the image and needs help interpreting it.
- D. Be short or empty, so the listener does not hear the same thing twice.

34 5. Words that carry the picture

You have cut a third from a draft and it has become harder to read rather than easier. What was removed that should have stayed?

- A. The redundancy that lets a distracted reader who missed a clause catch up.
- B. The qualifiers, whose removal made every remaining claim sound overstated.
- C. The transitions, which are the only way a reader can tell paragraphs apart.
- D. The examples, which are the part a reader remembers after finishing.

35 5. Words that carry the picture

A draft scores as easy on a readability formula and nobody can follow it. How should such scores be used?

- A. As a target, since a lower score reliably means a wider audience can read the text.
- B. As a smoke alarm for drafts that have drifted into sixty-word sentences.
- C. As a check on vocabulary, since they measure how common each word is in use.
- D. Not at all, since the underlying research does not apply outside American English.

36 6. Type, colour and image as meaning

A brand system written before any work has been made is broken within a month and then ignored. What is the better order?

- A. Write the system from published guidelines, then adapt it once work exists to test it.
- B. Write a minimal system first and expand it every time a new situation is met.
- C. Make five real pieces, then write down the parts that actually worked.
- D. Copy a system from a business you admire and change only the colour values.

37 6. Type, colour and image as meaning

You commit to an elegant Latin typeface for a campaign that will later carry Devanagari and Tamil. What should you have checked first?

- A. Whether the licence for that face permits embedding in a PDF and in an exported video file.
- B. Whether the face was designed by a foundry with a South Asian catalogue.
- C. Whether the face has a variable version, which covers more weights per script.
- D. Whether it has glyphs for those scripts, or readers get tofu and a mismatched fallback.

38 6. Type, colour and image as meaning

A comfortable text setting is pasted into a wider column and becomes hard to read, with no other change. What has been broken?

- A. The measure and the line height are coupled, and a longer sweep needs a bigger target.
- B. The x-height is now too small in relation to the column width, so the text reads as grey.
- C. The justification has too few word spaces to distribute, which opens rivers.
- D. The type size should scale with column width, and it has stayed fixed.

39 6. Type, colour and image as meaning

A client cites the study where a red button outperformed a green one and wants every call to action in red. What is the accurate reply?

- A. Red is culturally loaded in this particular market, so the effect would run in the opposite direction.
- B. Button colour matters less than button size, so the budget is better spent elsewhere.
- C. That result is one test on one page, and what predicts noticing is contrast with the surroundings.
- D. Colour psychology holds for consumer products and not for services, which is what they sell.

40 6. Type, colour and image as meaning

A mid-yellow looks brighter to you than a particular blue, yet the yellow fails a contrast check on white and the blue passes. Why can you not settle this by eye?

- A. Monitors render saturated yellow inaccurately, so the perception is a calibration artefact.
- B. The ratio is computed from relative luminance, which does not track perceived lightness.
- C. Contrast checkers measure against a mid-grey reference rather than against white.
- D. The blue is more saturated, and saturation is weighted separately in the formula.

41 6. Type, colour and image as meaning

A wide lens used close to a face reads as aggressive and a long lens from a distance reads as observational. What has changed?

- A. The lens coating, which alters contrast and micro-detail in the subject's features.
- B. The depth of field, which removes context and makes the subject appear isolated.
- C. The exposure, since a close subject is lit differently from a distant one.
- D. The distance ratio between near and far parts of the subject, not the face.

42 6. Type, colour and image as meaning

You draw a process as five boxes and arrows. What is the specific way a diagram can mislead?

- A. It asserts that the world has five distinct components with clean interfaces.
- B. It removes appearance entirely, so readers cannot recognise the parts in real life.
- C. Arrows imply a fixed order, which is rarely how a real process runs.
- D. Boxes of equal size imply equal importance, which flattens the hierarchy.

43 7. The surface and the system carrying it

A vivid on-screen blue comes back from the printer duller and slightly shifted, and the client blames the printer. What happened?

- A. The file was supplied at 150 dpi rather than 300, so the ink could not hold that much saturation.
- B. The bleed was missing, so the press trimmed into the most saturated area.
- C. Screens mix light and presses mix ink, and some screen colours have no CMYK equivalent.
- D. The paper was uncoated, which absorbs ink and lowers every colour equally.

- 44 7. The surface and the system carrying it

A creator improves click-through by twelve per cent and retention drops by nine. How should that be read?

- A. As an improvement, since more people now reach the piece than before.
- B. As a neutral trade, since the two numbers measure different audiences.
- C. As a sound problem, since new viewers arrive with the sound off.
- D. As a failure moved later in the funnel, where it is harder to see.

- 45 7. The surface and the system carrying it

A platform changes its interface overnight and a year of correctly composed vertical work now has a button on its subject. What is the defence?

- A. Keep a master at the widest sensible framing, so a new crop is an export.
- B. Compose everything to the published safe-area template of the largest platform.
- C. Add a permanent border, so the interface never overlaps the picture area.
- D. Export several aspect ratios at the time of delivery and archive all of them.

- 46 7. The surface and the system carrying it

Left alone, a platform picks the cover frame for you. Why is the automatic choice usually poor?

- A. It samples from the middle of the file, which is rarely the point at which the subject appears.
- B. It prefers frames with faces, which crowds out product and process shots.
- C. Frame zero is often the darkest, and automatic selection prefers sharp over meaningful.
- D. It re-encodes the frame at a lower quality than the video it was taken from.

47 7. The surface and the system carrying it

A muted, desaturated style stood out in a feed last year and is invisible this year, unchanged. What does that demonstrate?

- A. Audiences habituate to any style within about a year regardless of its quality.
- B. Conspicuousness is a relationship with the neighbours, not a property of the work.
- C. Platform compression has changed, and desaturated work now loses more in re-encoding.
- D. The feed now shows work at a smaller size, and muted work suffers more at small sizes.

48 7. The surface and the system carrying it

Completion rate correlates with satisfying pieces, and hard optimisation against it produces pieces that withhold the answer until the final second. What is that an instance of?

- A. A ranking change, since platforms periodically reweight the signals they use.
- B. Survivorship, since only the creators who succeeded explain what they did.
- C. Selection pressure on creators, which is a separate effect from any measurement problem.
- D. A measure used as a target ceasing to be a good measure, as Goodhart described.

49 7. The surface and the system carrying it

A retention curve rises at 0:41. What are the two possible readings, and how do you tell them apart?

- A. A loop or a rewatch; check whether something at 0:41 was unreadable the first time.
- B. A recommendation spike or a share; check the traffic sources for that timestamp.
- C. A caption change or a music change; check which of the two lands at 0:41.
- D. A sampling artefact or a real rise; check whether the view count is above a thousand.

50 8. Judgement, and the work that accumulates

You will read that five testers find eighty-five per cent of problems. What is the honest version of that claim?

- A. It holds for interfaces and not for static graphic work, which has fewer failure modes.
- B. It requires the five testers to be drawn from the intended audience rather than from friends.
- C. It comes from a model assuming problems are independent and equally likely to be noticed.
- D. It was measured on paper prototypes, and screen work needs roughly twice as many testers.

51 8. Judgement, and the work that accumulates

A peer reviews your work and reports inconsistent spacing. A stranger reports not knowing what the piece was for. How should the two be used?

- A. Strangers tell you whether it works; designers tell you why it does not.
- B. Designers are more reliable overall, because they can diagnose as well as observe.
- C. Strangers are more reliable overall, because they represent the real audience.
- D. Both are observations, so they should be batched and weighted by how often they recur.

52 8. Judgement, and the work that accumulates

You have argued from the brief and the client has overruled you. What is the professional sequence from here?

- A. Execute the change and raise the outcome afterwards, once the results are known.
- B. Ask for the decision in writing from them, so responsibility is clearly recorded.
- C. Offer a compromise version, so that neither position is fully adopted or rejected.
- D. State the risk once in writing, do it properly, and keep both versions.

53 8. Judgement, and the work that accumulates

A checkout shows a low price first and reveals fees at step four. Which mechanism is being used?

- A. Loss aversion, since the buyer now stands to lose the price they were promised.
- B. Anchoring plus sunk effort: later numbers are judged against the first, after investment.
- C. Social proof, since other buyers are implied to have accepted the same fees.
- D. Default bias, since continuing requires no decision at all while stopping requires a deliberate one.

54 8. Judgement, and the work that accumulates

Each individual piece a studio makes about one district is accurate, and the body of work is still a caricature. Which failure is that?

- A. The composite, where a character assembled from several people resembles nobody.
- B. The backdrop, where somebody's home is texture behind a story about someone else.
- C. The single story, where one narrative stands in for a whole group across many pieces.
- D. Decoration, where a person is included for visible difference and given nothing to do.

55 8. Judgement, and the work that accumulates

You embed provenance metadata recording that an image was generated. Why is a visible label in the frame still required?

- A. Metadata does not survive a screenshot, a re-upload or a tool that does not carry it.
- B. Provenance standards are not yet readable by the software most viewers use.
- C. Metadata can be edited by anyone, so it carries no evidential weight at all.
- D. Labelling rules in every country specify the exact wording that must appear on the image.

56 8. Judgement, and the work that accumulates

Of the seven deconstruction questions, why does what was left out teach more than the other six together?

- A. It is the only one that cannot be answered without access to the original material.
- B. Omissions are what platforms reward, so they explain the piece's distribution.
- C. It identifies the maker's constraints, which are the real cause of most of the decisions taken.
- D. Addition is visible and subtraction is not, and most of the craft is in what is absent.

Answers

Each entry gives the correct option, then what each wrong one reveals about how somebody was thinking. The second part is worth more than the first: a wrong answer here is a named misreading, not a mistake.

Lesson checks

1. Nice work. What is it for? — C

A — Treats an imbalance in emphasis as a problem with the strongest element rather than with what was left unemphasised.

B — Assumes the reader will give the piece more attention than a real passer-by does.

D — Uses audience mismatch to dismiss a result about prominence, which would show up with any viewer.

2. Nobody is going to hand you a real brief — A

B — Mistakes a client's lack of vocabulary for a delegation of the decision, leaving nothing to point at when they reject it.

C — Assumes the obstacle is time pressure rather than the client having no articulated position to retrieve.

D — Offers a choice between variations rather than between premises, which produces a preference but no brief.

3. "Young urban professionals" cannot be persuaded — B

A — Mistakes the benefit for the reader when the benefit is for the maker: the portrait's job is to settle the designer's decisions, not to appear in the piece.

C — Confuses the imagined viewer with the targeting: the portrait guides the making and does not change who is served the piece.

D — A portrait is still partly invented; the gain is decidability, not evidential strength, and treating it as data is how teams stop checking it.

4. You cannot move someone without knowing where they are standing — C

A — Reaches for the most enjoyable fix; the premise already states the offer is understood, so rewriting for clarity repairs something that is not broken.

B — Treats page views as attention but ignores that the stated symptom is people understanding and still not acting, which is a later stage of the funnel.

D — Inverts the remedy: raising the ask increases the belief required, which is precisely the resource already in short supply.

5. The constraints were true all along; you just found out late — B

A — Assumes quantity unlocks budget; the per-unit economics change but the total budget is a separate line that nobody has revisited.

C — Confuses volume with audience breadth; the same 5,000 flyers can be handed to exactly the same kind of person, and broadening the message weakens it.

D — A real scheduling consequence, but it is the smaller one — the list's value here is exposing that the technique, not just the timeline, is now wrong.

6. Six questions, and what to do when they cannot answer — B

A — Attributes the effect to rapport; the technique works on the first meeting with a guarded client, so rapport is not the mechanism.

C — Treats a working document as a commercial instrument; the brief's value is informational, and using it as leverage discourages the corrections you want.

D — Optimises for blame rather than accuracy, and a brief written to be defensible tends to be vague in exactly the places that matter.

7. Reference, and the line you actually cross — B

A — Explains the outcome by the person rather than the method, which leaves nothing you can actually do differently tomorrow.

C — Inverts the usual dynamic — more sources tends to dilute any single one; the failure here is depth of borrowing, not count.

D — Treats surface substitution as transformation, which is exactly the move that produces a copy with new paint.

8. Agree the measure before you make the thing — C

A — Fixes the ending, but almost nobody reached the ending — a 4% completion rate means the failure happened long before the ask.

B — Possible but unsupported, and reaching for a measurement fault before reading the diagnostic figure is how teams avoid the finding.

D — Blames distribution for a retention curve that a mismatched opening produces just as readily, and it is the explanation that requires no change to the work.

9. A design that says four things says none — D

A — Treats a ranking problem as a palette problem, which is why the brighter colours were rejected.

B — Assumes reading distance removes the need for an entry point, when scan order matters at any distance.

C — Converts a diagnosable structural fault into a matter of taste, so the same fault survives whichever option wins.

10. There are only five ways to make something look important — A

B — Blames a specific colour for an effect that would occur identically with any single high-contrast hue used alongside four other emphasis devices.

C — Treats one lever as the culprit when the headline could shrink and still dominate through the remaining four.

D — Overstates position: centred layouts can carry three clear levels perfectly well when the other levers are spent sparingly.

11. The eye reports differences, never values — A

B — Raises a real localisation concern but not the failure here, which affects viewers inside the same culture as the designers.

C — Invokes fatigue rather than perception; the states are indistinguishable at first glance, not after prolonged use.

D — Addresses convention rather than perception: a legend does not help a viewer who cannot tell the two marks apart in the first place.

12. You see far less at once than you think — B

A — A true and important caveat about gaze data in general, but it does not explain why this particular layout reproduced the pattern.

C — Questions the statistics when the problem is that the finding was designed into the page, and a larger sample would reproduce it just as faithfully.

D — Correctly notes a cultural boundary, but a right-to-left audience would produce a mirrored version of the same uninformative result.

13. Every element you add taxes the ones already there — D

A — True as far as it goes, but it explains why one fix is bad rather than why grouping works, which is the mechanism in question.

B — Describes a hierarchy benefit that follows from the grouping, not the perceptual reason the chaos subsides.

C — Borrows contrast language for a spacing effect; the block's legibility changes little, while the count of objects competing for the first glance changes a great deal.

14. The empty part is doing the most work — D

A — Assumes a size problem; the errors are about which field a label belongs to, which equal spacing causes at any type size.

B — Offers a different pattern with well-documented problems of its own, and side-steps rather than explains the perceptual cause.

C — Proposes adding elements when the existing gaps could carry the structure unaided, which adds cost without addressing why the current gaps fail.

15. Four cheap tests that fail a design before an audience does — A

B — Questions the instrument when three independent human reports already contradict the conclusion the instrument would support.

C — Dismisses the only evidence about meaning on a distinction that rarely explains a total failure to state what is on offer.

D — Jumps to a remedy before the diagnosis, and adding an element to a piece whose central claim is unclear usually produces two unclear claims.

16. The bit you are proudest of is usually the bit in the way — C

A — Reasons from sunk cost, which is the exact bias the lesson identifies as the reason the element survived too long in the first place.

B — Substitutes a commercial motive for a design one, and builds the hierarchy around the client's perception rather than the viewer's glance.

D — Turns a judgement into a quota; some pieces need no ornament at all, and a mandatory flourish is how house styles become tics.

17. "And then" is the weakest word you own — B

A — Treats felt pacing as a substitute for structural necessity, which fails as soon as the sound is off.

C — Optimises the order of interchangeable items without making any of them depend on another.

D — Assumes drop-off is caused by duration rather than by every beat being individually optional.

18. The first second is a promise, and you have to pay it — A

B — Explains short watch time by duration alone, which would not also stall follower growth on high-view videos.

C — Locates the mismatch in surface quality rather than in the content of the promise being made.

D — Attributes an internally consistent pattern of metrics to an external system rather than reading what the metrics say.

19. If nothing changes, it is not a story — C

A — Treats a structural absence as a production problem; pace makes an unchanging sequence shorter, not memorable.

B — Confuses what a viewer should do with what a viewer will remember, and an ask cannot repair ninety seconds that contained no event.

D — Blames delivery for a script that contains no change, and a better voice would narrate the same unchanging sequence.

20. Where the ninety seconds actually go — B

A — Makes a stylistic claim that varies entirely by genre and does not address what the running time is being spent on.

C — Overstates a real finding: early drop-off is steep but substantial audiences do continue, and the eight-second turn fails them specifically.

D — Assumes the shortfall is coverage when the turn may be well covered and simply crowded out by establishing material.

21. One checkable detail outperforms ten adjectives — B

A — Generalises about numbers when an unverifiable statistic persuades no better than an adjective; the mechanism is exposure to contradiction, not numeracy.

C — Names a genuine secondary benefit of concrete detail but not the reason the claim is believed at the moment it is read.

D — Assumes a comparison the reader rarely performs, and the sentence persuades even when no competitor figure is available to compare against.

22. Nothing is at stake, so nobody is watching — C

A — Explains by trend rather than mechanism, which predicts the effect will fade and gives no guidance on when an admission backfires.

B — Assumes the reader computes the percentage; the persuasive work is done before any arithmetic, by the willingness to expose a checkable figure.

D — Describes a plausible defensive benefit while missing why the reliability claim itself becomes more believable at the moment it is read.

23. The person who changes is the one it is about — B

A — Reads a passed test as a failure; adding brand presence to a story with no change produces a branded film with no change.

C — Reaches for casting when the structural problem would persist with any casting, however closely matched to the audience.

D — Proposes covering a structural gap with commentary, which produces a film that explains why we should care about an event that did not occur.

24. Endings, and the four asks that cancel each other out — B

A — Singles out one element for a suppression effect that the mechanism attributes to the number of competing options, not to any one of them.

C — Attributes the drop to a credibility judgement viewers are unlikely to make in the final two seconds of a short video.

D — Explains by channel quality while ignoring that the phone number is present in both versions and performed worse in the crowded one.

25. Draw ugly boxes now or pay for it later — C

A — Blames tool defaults for a constraint that comes from the cost of abandoning invested work.

B — Assumes exploration scales with time spent rather than with how cheap it is to discard an option.

D — Confuses a viewing-size problem with the reason no alternative layouts were generated at all.

26. Most interview footage is unusable, and the question is why — C

A — Treats it as a beginner's crutch rather than a permanent constraint of the format; experienced interviewers do this too, for the same mechanical reason.

B — Confuses volume with usability, and more material actually slows an edit; the point is self-containment, not length.

D — Invents an audience-perception effect; the real cost of a fragmentary answer is that it is structurally unusable, not that it sounds less credible.

27. Half your footage is wallpaper — A

B — Treats it as an attention problem; the damage happens even when the viewer takes in the whole sentence, because the claim itself is overwritten.

C — Jumps to an external consequence and ignores the immediate effect on the ordinary viewer, who never screenshots anything.

D — Reframes a single-shot credibility failure as a long-run brand-consistency issue, which misses why this one cut fails on its own.

28. Consent is a scope, not a signature — A

B — Asserts a uniform legal outcome across countries where the law actually differs, and the common failure is a relationship breaking down rather than a case being fought.

C — Imports a commercial norm from paid talent work; the dispute usually comes from an unpaid participant surprised by a use, not from a fee argument.

D — Invents a platform requirement; no mainstream host inspects releases, and the problem surfaces between you and the participant.

29. Anonymity fails by combination, not by name — C

A — Identifies a real secondary benefit but not the reason the crop is stronger; a crop protects even when the background is empty.

B — Names a genuine social effect while missing the technical one; the file would still be attackable if nobody were curious.

D — Misdescribes how compression works — codecs do not restore information that was destroyed before encoding — and so points at the wrong mechanism entirely.

30. Free to download is not free to use — D

A — Confuses BY with NC and adds a payment condition that no Creative Commons licence contains.

B — Misreads ShareAlike as NoDerivatives; BY-SA explicitly allows modification, it only constrains what licence the result carries.

C — Invents a restriction; CC licences cover copyright only and say nothing about the separate question of a subject's image rights.

31. A number is a claim with the argument removed — B

A — Explains by volatility rather than by encoding; a stable index would still be fine truncated and a volatile sales bar chart still would not.

C — Assumes expertise cancels a visual distortion; the effect of a truncated axis is robust and does not depend on the reader's sophistication.

D — Appeals to convention instead of the graphical reason, which leaves you unable to judge any chart type the conventions do not cover.

32. Find the missing shot now, not in the edit — B

A — Offers a plausible production effect that would make the extra takes merely useless, not disproportionately costly to sift.

C — Identifies a cost that genuinely scales with volume but only linearly, so it cannot explain the disproportionate slowdown.

D — Describes a real downstream symptom of too much choice while skipping the arithmetic that causes the work to expand in the first place.

33. A headline is not a summary — B

A — Invents a ranking effect; truncation is a display limit applied after ranking and is indifferent to the opening word's grammar.

C — Treats it as a length problem, but the same reordering helps even when both versions are far over the limit.

D — Correctly notes that concrete nouns attract attention but assumes the reader sees the whole string, which is exactly what truncation prevents.

34. Everything you write has a fold in it — A

B — Attributes the effect to distribution rather than to the reader's choice, which would predict no difference among people who do see both.

C — Invents a timing mechanism; the fold decision happens immediately, not after a delay long enough for curiosity to decay.

D — Shifts the explanation to search discovery, which does not account for the different behaviour of readers already looking at the same post.

35. A listener cannot go back — A

B — Singles out the date when the same sentence fails with the year removed; the problem is the suspended structure, not any one element.

C — Assumes the clause boundary is the difficulty, but intonation marks it clearly; the cost is memory load, not segmentation.

D — Explains by duration and patience, which would predict that any long sentence fails equally, including the unpacked two-sentence version of the same length.

36. Most of your audience has the sound off — C

A — Points at capitalisation, but the same slowdown occurs with sentence-case captions where both lines begin identically.

B — Treats it as a length violation, yet the break damages reading even when both arrangements sit well inside the character limit.

D — Generalises a real cost of line changes into a rule against two-line captions, which broadcast practice contradicts and which does not distinguish good breaks from bad ones.

37. The caption decides what the photograph means — B

A — Assumes automatic detection of decorativeness; without an explicit empty value many screen readers fall back to announcing the file name.

C — Optimises for a checker rather than a listener, and automated tools do not in fact penalise length.

D — Judges the decision by search value when alt text exists for a listener, and this reasoning would also wrongly justify stuffing keywords into meaningful images.

38. Your first draft contains the route, not the destination — A

B — Attributes the difficulty to layout, but the same dense text is hard to follow when set as one continuous block.

C — Substitutes an engagement explanation for a comprehension one; the reader here is trying to follow, not deciding whether to care.

D — Names a genuine effect of heavy compression but one that applies to word choice, not to the removal of examples and transitions the question describes.

39. You cannot un-know the thing you are explaining — C

A — Attributes the benefit to reading speed, which would mean a slow careful re-read by the author could substitute for it.

B — Identifies a real editing bias, but attachment concerns what you are willing to remove, not what you are able to perceive as unclear.

D — Frames it as a sampling argument, which one reader cannot support, and misses that the author's judgement is not merely less reliable but systematically blind here.

40. Most people reading your English did not grow up in it — D

A — Assumes the synonyms are rarer words, but the cost appears even when all three are common; the problem is reference, not frequency.

B — Names a genuine downstream problem for translation while missing what happens to a reader reading the English directly.

C — Explains by recall rather than by comprehension, which predicts a delayed difficulty instead of the immediate confusion that actually occurs.

41. Your brand is not the logo, it is what you repeat — B

A — Assumes inconsistency comes from insufficient palette range rather than from undefined usage.

C — Substitutes a rendering and accessibility issue for a decision-making one.

D — Expects a tool to supply the usage rules that were never written, when the tool can only store what has been decided.

42. A typeface means where you have seen it before — D

A — Assumes the typeface will render the Indic text at all; the problem occurs before any shaping, because there are no glyphs to distort.

B — Makes an aesthetic judgement about Devanagari typography rather than naming the technical failure that happens first.

C — Invents a licensing structure; licences are generally about use and distribution, not about which of a font's existing glyphs you use.

43. Point size is the wrong number to argue about — C

A — Frames it as cognitive load per line, which would be fixed by reading more slowly; the failure is a targeting error the reader cannot slow their way out of.

B — Notes a real navigational aid but one unaffected by column width, since the same paragraphs remain however wide the column is.

D — Invents a relationship; x-height is a property of the typeface at a given size and does not change when the column changes.

44. Red is not danger everywhere — D

A — Raises a valid concern about evidence quality but would be answered by a larger sample; the deeper problem survives any amount of data from this one page.

B — Correctly notes that associations are local, but the result would still fail to generalise within a single market once the page around the button changed.

C — Asserts a different dominant factor, but the test held the label constant, so the observed difference did genuinely come from the colour change.

45. One in twelve men cannot see the difference you just used — D

A — Invents a medium mismatch; the formula is defined for screen luminance and the yellow would fail on paper too.

B — Describes a restriction the formula does not have; it accepts any two colours and returns a ratio for them.

C — Imagines a correction factor for colour vision that the ratio does not contain, and misidentifies which confusions deuteranomaly actually causes.

46. Every photograph is mostly the part you cut off — A

B — Describes an evidential difference that actually supports the rule, rather than explaining why the line sits away from the harm.

C — Gets the reason for the rule's placement right but stops at detection, missing that cropping is the more powerful distortion.

D — Assumes readers correct for framing, when the excluded context is by definition invisible to them and cannot be discounted.

47. The photograph that fits anywhere belongs nowhere — A

B — Identifies a technical cost that applies equally to a good photograph, so it cannot explain why the generic one specifically harms the page.

C — Offers an aesthetic consistency argument, which does not account for the measurable attention effect that occurs even on otherwise strong pages.

D — Overstates the restrictions on the main free libraries and answers a licensing question rather than the attention question being asked.

48. Three pictures, three different promises — D

A — Names a real cost of surplus detail, but a cleanly lit studio photograph with no clutter would still fail to show the stitch.

B — Invents a trust ordering; the choice is driven by what each form can depict, not by what it signals about the author.

C — Raises a reproduction property that is often true and entirely beside the point, since the photograph fails at any size.

49. The same design fails in five different ways — C

A — Attributes both a colour shift and a trim error to resolution, which affects sharpness rather than either.

B — Assigns a predictable gamut limitation to a variable material choice, and explains nothing about the white edge.

D — Turns a specific weakness in CMYK handling into a blanket claim, which hides the two separate faults actually present.

50. The thumbnail and the title are one thing — A

B — Discards a real signal instead of reading the two together, which is what exposes the mismatch.

C — Treats two metrics that moved in opposite directions at the same moment as unrelated coincidence.

D — Invokes a rule about platform mechanics to avoid judging what the thumbnail is actually promising.

51. Something is always sitting on top of your picture — A

B — Understates what the tools do — most track across the shot — and blames a limitation that better tracking would solve, which it does not.

C — Confuses a resolution artefact with a compositional decision; the crop would land in the same wrong place at any quality.

D — Describes a real further crop that happens downstream, but it does not explain why the tool's own choice of region was wrong to begin with.

52. The still is what recruits the play — A

B — Close in effect but wrong in mechanism; the cache is held per URL by the platform, not copied into each individual post.

C — Blames browser caching, which would resolve for new viewers who never downloaded the old file — yet they see it too.

D — Invents a dependency between search indexes and social cards, which are separately fetched and stored by each platform.

53. It has to work silent and reward being unmuted — D

A — Normalisation changes level uniformly, so it cannot selectively remove the low end that gives the voice its body.

B — Reverses how perceptual codecs allocate bits; they generally preserve low frequencies, which are cheap to encode and perceptually prominent.

C — Identifies real speech-range tuning but treats the bass as merely de-emphasised, when the driver cannot produce those frequencies at all.

54. Nothing is ever seen on a white wall — C

A — Treats saturation as intrinsically more attention-getting, which the original success of the muted style directly contradicts.

B — Assumes viewers attribute styles to originators, which almost never happens in a feed seen at a glance by strangers.

D — Invents a technical artefact to explain an effect that is about the surrounding context rather than about image quality.

55. The system is optimising something, and it is not your work — C

A — Describes an audience-targeting effect that would narrow reach without making individual pieces worse, which is not the drift being asked about.

B — Assumes a weighting nobody can verify, and the drift occurs even where completion is only one signal among several.

D — Names one real tactic and mistakes it for the general mechanism, which also produces looping, padding and withheld payoffs.

56. Which number answers which question — D

A — Claims a difference in measurement accuracy; both come from the same underlying data, and the dip's value lies elsewhere.

B — Argues the aggregate is pessimistic rather than uninformative, which does not explain why a located dip can be acted on and a total cannot.

C — Assumes one cause for the dip; the common causes are editorial, and treating it as a technical fault sends you looking in the wrong place.

57. Three tests, no budget, twenty minutes — D

A — Treats it as a timing error that a stricter stopwatch would fix, when the damage persists however briefly you speak.

B — Names a real social pressure, but the test would still fail with a tester who felt entirely free to disagree.

C — Invokes channel interference to describe what is simply prior knowledge, and would wrongly imply a written explanation is safe.

58. "Make the logo bigger" is a symptom, not an instruction — C

A — Generalises into a licence to reinterpret anything, which is how designers stop hearing clients at all.

B — Invents an economy of effort; the choice is about what makes the number noticeable, not about which edit is quicker.

D — Elevates a system rule above the problem being reported, which would justify ignoring the observation rather than solving it.

59. Argue from the document, not from taste — A

B — Objects on grounds of effort, which would be answered by producing the weak options quickly; the problem survives that fix.

C — Treats client ownership as a drawback when it is usually an advantage, and misses what the staged choice costs you.

D — Names a real risk of the tactic, but one that a careful designer avoids, so it does not explain why the approach is wrong in principle.

60. If explaining it would stop it working, it is manipulation — B

A — Judges by the chance of being caught, which would make an unverifiable but true claim as bad as a false one.

C — Substitutes the current regulatory position for the underlying judgement, leaving you unable to assess a pattern no regulator has addressed yet.

D — Claims a different psychological mechanism when both timers work through loss aversion; the difference is the truth of the claim, not the route it takes.

61. Somebody is in this piece, and it is not you — D

A — Identifies the practical constraint but stops there; the deeper problem is what the exercise is then being used for.

B — Raises a real limitation of any consultation, early or late, so it does not explain why timing makes it worse.

C — Frames it as a relationship signal, which suggests the fix is better communication rather than actually moving the consultation earlier.

62. What the tool can be trusted with, and what it cannot — A

B — Describes a temporary adoption gap that wider support would close, whereas the underlying weakness persists no matter how many platforms adopt it.

C — Misstates how signed provenance works; the cryptographic record resists forgery, and the practical failure lies elsewhere.

D — Misdescribes the standard's contents; the record is designed to cover origin as well as subsequent edits.

63. Take apart the thing that worked on you — C

A — Describes a consequence of practising without feedback while skipping the cause, which is that the feedback never arrives in a usable form.

B — Attributes the plateau to fixed ability, which the improvement people do make once feedback is constructed directly contradicts.

D — Blames external constraint; the same plateau appears in unconstrained personal projects where every decision was the maker's own.

64. Fifteen pieces with a point of view beat fifty with a filter — C

A — Identifies volume as the problem rather than the absence of any pattern, so a smaller but equally scattered set would perform the same.

B — Explains budget mismatch by a missing price rather than by what the work itself signals about judgement.

D — Assumes the issue is which pieces are seen first rather than what the collection as a whole communicates.

Module test — 1. Before you open the software**1. C**

The viewer already holds a belief with their own experience behind it. Your claim arrives as a competing one from a party with an obvious interest, so it loses. What moves a belief is evidence the viewer can check without trusting you first: a photograph of the dispatch board, a named customer with a phone number, a missed-delivery figure published against your own interest.

2. D

People who cannot author a brief are excellent at correcting one. Proposing converts an unanswerable question into an editing task, which is a thing busy people can do in two minutes, and the correction you get back is usually better than any answer the questioning would have produced.

3. B

Almost every decision you make while working has the same shape: would this person stop, and would they believe it. A category cannot stop or believe, so it has nothing to say. What decides the question is a written portrait naming the device, the posture, the state they are in, and the past disappointment your work has to survive.

4. A

Each of the seven constraints is a filter that rejects whole classes of idea in seconds. The production route carries its own list of impossibilities: the laser printer cannot do full bleed, the plotter cannot do fine gradients, one ink cannot do a reversed-out photograph. An idea that fails a constraint is not a good idea you cannot afford; it is an idea for a different job.

5. D

Structure, mechanism and constraint are yours to take: where the emphasis sits, why the piece works, that it was made in two colours. Surface is not. The recognition test is the practical version of that line, because taking structure produces work that looks like nothing in your reference pile, while taking surface produces a copy with new words.

6. C

Small campaigns produce small numbers, and small numbers are noisy. A difference of four sits inside the ordinary week-to-week spread, so it is not evidence about the flyer. You need either a larger volume or a longer period, and the honest move is to say in advance which of the two you are relying on.

Module test — 2. Making one message survive

1. B

The eye needs a single entry point and finds it by looking for the strongest contrast. Three equal candidates give it no reason to prefer any, so the scan order becomes arbitrary, and an arbitrary scan is uncomfortable enough that the reader resolves it by leaving. The fix is ranking rather than removal: two of the three drop a level.

2. A

Emphasis is comparative and a surface holds a fixed amount of perceivable difference. Spend it all on level one and levels two and three have nothing left to separate them, so the hierarchy collapses into the shouting thing and everything else. Two levers per element is the working limit, and the fix is almost never to enlarge level one.

3. D

Red and green of identical lightness are strikingly different in colour and nearly identical in tone. Anything built on hue alone vanishes in greyscale, in a one-colour photocopy, and for the substantial minority with red-green colour vision deficiency. The fix is never to avoid colour; it is to make tone, shape or position the first signal and colour the second.

4. C

The F came from eye-tracking of long, largely unformatted prose with little visual hierarchy. Under those conditions it is real and repeatable. It describes behaviour in the absence of design: give people a strong entry point, a photograph or a pull quote and the pattern changes. Applied as a template it produces exactly the unformatted page it describes.

5. A

Removal answers a question rearrangement cannot: what is this element actually doing. Going through every element in turn usually finds two or three the design is better without, and at least one you assumed was load-bearing and is not. It is also the version of the conversation a client can judge, because two files side by side beat an opinion about hierarchy.

6. D

Proximity is read pre-consciously as grouping, so every gap is a statement about what belongs with what. Equal gaps state that everything belongs with everything, which is the same as stating nothing. Move the caption to 8px from its image and 40px from the heading and it snaps into place, with nothing else changed.

Module test — 3. The shape of a story

1. C

And then means the second beat does not depend on the first, so it can be removed with nothing broken. A viewer feels that even though they could not name it: leaving early costs them nothing, so they leave. Rewriting until the joint is but or therefore is what makes the middle of a piece impossible to cut.

2. B

A viewer answers those two questions almost without thinking, and if either answer is missing or ambiguous they leave, because on a feed leaving costs nothing. Everything called a hook is a way of delivering that answer faster, which is why the test of an opening is not how startling it is but whether the rest of the piece pays what it implied.

3. D

The minimum unit of a story is a state, a thing that happens, and a different state. When the start and end states match you have a description. Descriptions are legitimate objects, and they behave differently in attention: a menu, a price list or a location film earns its viewing by being useful, beautiful or short, and no amount of music supplies the missing change.

4. A

Clear narration runs at roughly 140 to 160 words a minute, so ninety seconds is about 210 to 240 words. At 400 you are not holding a long ninety-second script, you are holding a different film, and the useful moment to decide that is while the change is a document edit rather than a re-record.

5. C

Claims from an interested party carry a discount because the obvious explanation is that they were made because they suited the speaker. A statement that damages the speaker removes that explanation, so the credit transfers to the surrounding claims. The limits are real: the fault has to be true, survivable, paired with what was done about it, and not the thing the customer is buying.

6. B

If what survives is somebody who had a problem, tried something and ended up different, you have a story that a brand appears in. If nothing survives, you had an advertisement decorated with anecdotes. It is not an argument for hiding the brand; it is a diagnostic for whether anything other than the brand is being carried.

Module test — 4. Getting the material

1. D

A polished draft gets lovely and is then rejected at the end, because it reads as a decision already taken. Unfinished work reads as an invitation, so the objection arrives while it still costs thirty seconds to act on. The same mechanism is why twelve throwaway thumbnails produce options you would never try after forty minutes in an editor.

2. C

Eleven years of what. The orphaned answer now needs a caption or a voiceover to repair it, which costs time and costs the piece a moment of somebody's real voice. The fix is a sentence said before you start: fold my question into your answer, so it comes out as I have been repairing sewing machines for eleven years.

3. A

When picture and voice disagree, most viewers believe the picture and rewrite the voice to fit. They will not be able to name what bothered them afterwards; they will simply trust the piece less. This is why grabbing a cutaway for its prettiness is dangerous: it carries whatever it happens to contain, including the thing in the corner you did not notice.

4. B

The complaints that arrive late are almost never from people denying they agreed. They are from people who agreed to one thing and found another. So the scope has to be written in words the person actually reads and kept narrow enough to be true, and a withdrawal route has to be agreed before anybody needs it.

5. D

All four statements are broadly true about their clauses, and only ShareAlike is a copyleft term that propagates. A ShareAlike photograph used in a commercial brochure can oblige you to license the brochure the same way, which is very unlikely to be what your client agreed to. NonCommercial forbids the use; ShareAlike permits it and changes your work.

6. A

A line encodes value by position, so starting the axis somewhere other than zero shifts every point equally and can be legitimate where zero is meaningless, as with a temperature or an index. A bar encodes value by length, so cutting its base makes the drawn length stop matching the quantity, and the comparison the reader performs is then wrong by construction.

Module test — 5. Words that carry the picture

1. A

A headline that describes the piece perfectly and is not read has failed completely. One that gets the first line read has succeeded even if it left things out, provided the piece then pays what it implied. Clarity is a tactic in service of that job rather than a value in itself, which is why it usually wins: confusion is a reason to move on.

2. D

Both ask for a tap, and they ask for different things. Trust is what a stranger does not have, and interest is what they may already have. Giving the premise away looks backwards until you notice what each version is requesting, which is also why the fold matters in proportion to how much of your audience is strangers.

3. B

A listener cannot backtrack without an act of will, so a held-open structure has to be carried in memory until it closes. They mostly do reconstruct it, and the reconstruction costs them the first half of your next sentence. The rule that generates the repair is to finish each grammatical thought before starting the next one.

4. C

The beginner's mistake is fitting the words on screen. The real limit is whether they can be read in the time they are there, by somebody who is also looking at the picture. When speech outruns the reading limit you condense the caption, and that is a skill rather than a cheat: a verbatim caption that cannot be read has served nobody.

5. B

Captions have two jobs: fixing which of an image's possible meanings applies, which is necessary and cheap, and adding what no arrangement of pixels could carry. The same lathe, four months after the workshop flooded, earns its place. The fourth option describes alt text, which is a different job for a different reader and is not a caption.

6. D

Knowing something changes how the signal seems from the inside, which is why tappers in the Stanford study predicted half their listeners would name the tune and about one in forty did. Everything you check will look clear to you, and it will look clear precisely because of what you know. Only a person who does not know is a measurement.

Module test — 6. Type, colour and image as meaning

1. C

Recognition is built by repetition, and the reason people fail at it is memory rather than laziness. In six weeks you will not recall whether the orange was one hex value or another, or whether headings were 32 or 34, so you re-decide slightly differently every time. The drift is invisible per post and obvious across a set, which is why the system goes on one page with actual values.

2. A

There is nothing in the shape of a letter that is inherently trustworthy or modern. Meaning is the accumulated weight of where a reader has met the face: in textbooks, on legal documents, on a decade's advertising. That makes it learned, local and changeable, so the useful question is never what does this font mean but where has my reader seen it.

3. D

Point size is a leftover from metal type and describes the body rather than the letters. What the eye responds to is x-height, the height of a lowercase x and therefore of most of the text, and two faces at the same point size can differ in it by a third. This is why make it 14pt is not a specification and setting a paragraph at real size is.

4. B

Association is learned by repeated cultural pairing, so red means luck in one place and loss in another, and white is a wedding in one and a funeral in another. What does not vary is physiology: saturation attracts, pure blue is poor for small text, red on blue vibrates, and around one in twelve men has red-green colour vision deficiency. Reserve your confidence for those.

5. C

A photograph is defined as much by its edges as by its contents, and everything outside the rectangle was excluded by somebody. A crowd photographed tightly around a confrontation is a riot; the same crowd wide is forty people and an argument. Neither is manipulated. The agency line is drawn where it is because adding and removing pixels is checkable, not because it tracks the harm.

6. A

To be sellable into many contexts an image must be specific to none, so it occupies the shape of a photograph and delivers the content of a decoration. Readers learn quickly which regions of a page reliably contain nothing and stop allocating attention there. The cost is that a slot which could have carried evidence carried nothing.

Module test — 7. The surface and the system carrying it

1. D

Everything else is that rule applied to different situations: 28 point on a slide read from twenty metres, three words on a thumbnail shown at 150 pixels, heavier type on a feed image seen at 400 pixels after re-encoding. A design that works on your monitor is a design for that monitor at sixty centimetres, and resizing tools move elements rather than rethinking which elements exist.

2. B

Title and thumbnail are one unit with two channels. The title carries the specific thing a sentence can hold, such as a number or a place; the image carries what a sentence cannot, such as scale or a state or a real expression. When they complement each other the pair opens a question the viewer wants closed. When they duplicate, they make a statement, and a statement is easy to walk past.

3. A

The visible frame is smaller than the exported one, and the amount is not a specification you can hold anyone to. Platforms redraw their interfaces without announcement, iOS and Android differ, and a tall phone with a notch crops differently from a short one. Two minutes of posting privately and screenshotting catches the failure that no amount of reading catches.

4. C

The first time a URL is shared, the platform reads the page's tags and stores the result, sometimes for a long time. The practical consequence is about sequence rather than markup: get the preview right before the link goes anywhere, and use the platform's card validator to force a re-fetch. Doing that before publication costs a minute; doing it afterwards costs an apology.

5. D

Autoplay is muted by policy, so sound-on is an earned state and the silent version has to carry the whole meaning rather than a hint of it. Bolting subtitles onto a narration film makes it readable rather than watchable, because the argument has moved into text and the eye cannot study the picture and the words at once. That is a structural decision, not an export step.

6. B

A cliff in the first seconds says the opening did not answer what this is and whether it is for me. Paired with a high click-through it says the cover and title attracted people the piece was not for, which is the signature of an inflated promise. It also rules out a structural problem later on, because almost nobody reached the later part.

Module test — 8. Judgement, and the work that accumulates

1. A

Saying what the piece is about installs the answer and ends the test, which is the commonest way a tester is wasted. Liking is answered socially: people say yes to be kind, including about a design that told them nothing. Ask what it said, what they would do, and what they noticed first, and hand it over in silence.

2. C

Non-designers are excellent at observation and unreliable at diagnosis and prescription, and prescription is the only vocabulary most of them have. Make it bigger is what I did not notice it sounds like without the words hierarchy or contrast. The observation is data and is not negotiable; which lever fixes it is yours, and the answer is often to delete something else.

3. D

On taste, the client's view is as good as yours and they are paying, so the argument cannot be won and should not be. The brief exists so that you never have to have it: a decision defended from an agreed document is defended from something they already said yes to, and they can engage with it by disagreeing about the goal, which is a conversation worth having.

4. D

We made the date the biggest thing so you would remember it survives being said out loud. We put a fake countdown there so you would panic stops working and makes the person angry. Persuasion supplies reasons a person can evaluate; manipulation exploits a process they cannot inspect. The boundary is imperfect and it decides almost everything you will meet.

5. A

One person does not speak for a group, and two people from the same community will often disagree, so no arrangement of consultants converts a contested judgement into a safe one. A consultation used as protection is worse than none, because it spends somebody's time and name on your defence. What it genuinely does is catch errors of fact and framing you could not have caught.

6. C

A human assistant who did not know something would sound uncertain, and you would check there. These tools produce the same confident register whether or not the output is correct, so the cue that tells you where to look is gone and you have to check by category instead: anything evidential, any fact, figure, name or date, any translation that matters.

Making It Say Something — final examination

1. A 1. Before you open the software

The removal test sorts a piece into furniture and structure. Decoration is permitted and is often what makes work feel like it came from somewhere; the failure is not knowing which is which. The second half of the test is the sharper one: cover the message instead, and if the piece still feels finished, you never had one.

2. C 1. Before you open the software

Two people leaving a meeting with a shared understanding is the exception; two private understandings is the norm. Writing it back converts the conversation into a thing both parties agreed to, so when round four asks for something that contradicts answer three, you have a calm and non-accusatory object rather than a competing memory.

3. B 1. Before you open the software

Every decision becomes tuned to one set of constraints, so a viewer without those constraints gets a piece shaped by problems they do not have. The buyer and the person who approves the purchase are often different people with opposite fears, which is what the second portrait catches. Naming who it is not for is a decision, and it stops the client later asking why the piece spends so long on credibility.

4. D 1. Before you open the software

The three stages fail differently and take different repairs. Attention is fixed by the first beat, the size, the placement. Comprehension is fixed by language and ranking. Belief is fixed by evidence, a guarantee, a third party or an admission, and it usually means going back to the client to ask for something they may not want to give, which is why teams avoid it.

5. C 1. Before you open the software

The process finds what the client knows and can say, which is not the same as what is true. The only correction is a second source, and it need not be expensive: ten minutes of competitor reviews, the person who answers the phone, three sentences from a real customer. When the second source contradicts them, show them the three sentences rather than asserting your opinion.

6. B 1. Before you open the software

Without the note, the board becomes an aesthetic mood and you unconsciously take surface. The note forces the useful distinction: the three-tone limit rather than the woodcut style, the on-screen-text rhythm rather than the music. The other options are all worth doing and none of them stops the specific drift, because the drift happens while you are looking, not while you are sourcing.

7. D 1. *Before you open the software*

Attribution is a mechanism built into the artwork, not an analysis performed afterwards. A different phone number, a code word to mention, a separate short link per surface: each costs nothing at design time and is unavailable once the file has gone to print. If none of them is possible, say so in advance, because a limitation named before the fact is a position and the same sentence afterwards is an excuse.

8. A 2. *Making one message survive*

Somebody scanning a menu for the vegetarian section is running a search, not receiving a message, and the only thing the piece is judged on is how short the search is. Forcing a hero onto it makes it slower to use. The same applies to a form, a ballot paper and a slide of exam dates.

9. D 2. *Making one message survive*

Size, weight, contrast and isolation carry across scripts. Position does not: in a left-to-right script the strong entry is top-left, and in an Urdu or Arabic layout it is top-right. Mirroring the geometry without rethinking the entry point produces a layout whose claims about reading order are now wrong.

10. B 2. *Making one message survive*

A cheap phone at thirty per cent brightness in daylight, ink spreading into uncoated stock so thin type closes up, a projector with a dying lamp: all of these take contrast away. The floor is a floor, so design with headroom. The cost of extra contrast is aesthetic restraint, and the cost of too little is that the piece does not work for part of your audience, who never tell you.

11. C 2. *Making one message survive*

The periphery reports coarse information only, and movement most of all, which is a legacy of needing to notice things approaching from the side. One moving element in a still layout captures the eye regardless of size or position, and keeps doing it. Use motion once, for the thing you want looked at, and honour a reduced-motion preference, because for some viewers continuous movement causes genuine symptoms.

12. A 2. *Making one message survive*

The visual system groups by proximity and similarity before you are aware of it, so five loose items become two groups, and two groups are one relationship rather than ten. The discipline it costs is spacing: the gap between groups has to be clearly larger than the gap within one, or the grouping does not read and you have five items with inconsistent gaps.

13. B 2. *Making one message survive*

After two hours you read your own layout from memory, going straight to each element because you already know the route. Mirroring removes the reading and leaves the masses, which is why painters have used a mirror for centuries. Do not then ship the mirrored version: reading direction matters and the flip reverses the entry point.

14. C 2. *Making one message survive*

The distinction is positional rather than moral. A distinctive element at level three gives the work a voice; the same element at level one takes the message's place. Sanding every piece down to a centred sans-serif headline on white is the failure mode of designers who have learned restraint and nothing else, and being forgotten produces no complaints and therefore no feedback.

15. B 3. *The shape of a story*

Structure is for the audience who has not yet decided to care. When they have already decided, the job is to get out of the way and organise it. Dramatising makes the piece slower to scan and slightly untrustworthy, because the reader senses the persuasion and starts to wonder what is being hidden.

16. C 3. *The shape of a story*

Bait is a withdrawal from an account you cannot refill, and it shows up as a shape rather than as a verdict. Platforms have moved steadily towards weighting watch time and satisfaction over clicks, and audiences remember the last three seconds long after they have forgotten the first. The line is not influencing versus not influencing; it is whether the piece pays what the opening promised.

17. A 3. *The shape of a story*

A chain of state changes cannot be reordered without damage, because each new state creates the next problem. A list of events can. Writing the state each beat leaves the viewer in, one line per beat, finds the pairs that leave the same state, and that routinely removes a quarter of a rough edit on evidence rather than on feel.

18. D 3. *The shape of a story*

At fifteen seconds the rough division is two seconds of promise, five of problem, five of turn and three for the new state and the ask together, so there is no room for a separate ask. The structure does not change; the sections compress, and two of them merge. At six minutes the same shape stretches by subdividing the turn, not by growing the setup.

19. B 3. *The shape of a story*

A detail is believed because it could be contradicted: two millimetres, the second word, eleven at night. That is the whole mechanism, and it runs in both directions. A vague claim that fails disappoints; a precise claim that fails is a lie with a number attached, and it is the kind of thing customers photograph and post.

20. C 3. *The shape of a story*

That asymmetry is the signature of a story whose first half was written afterwards, and audiences detect it reliably without being able to name it. Real stakes are specific: what was risked at what scale, what was given up, what went wrong. If there genuinely were none, say so plainly, which is unusual enough to be interesting in itself.

21. D 3. *The shape of a story*

It is a reliable illusion produced by familiarity. Two to three seconds of stillness is roughly the time to read and act on a short line, and it is one drag of the clip's edge in any free editor. The ending has three jobs, and holding the card is what lets the third one land after the first two have been done.

22. B 4. *Getting the material*

Planning what has to exist is different from planning what it looks like. For observational work you board the structure and a short list of shots you must not leave without, and keep the frames open. Over-boarding converts a shoot into an exercise in matching pictures you drew before you arrived.

23. C 4. *Getting the material*

An interview is testimony, not record. Each retelling is assembled from fragments plus expectation, which is neither dishonesty nor unusual. Checking is not doubting them; it is protecting them, because the number that goes out under their name with your edit around it is the one a stranger will argue with.

24. D 4. *Getting the material*

Pairing a picture with words helps recall when the picture adds information the words do not have, and helps very little when it restates them. Identical is not reinforcement. An image should illustrate, prove or bridge, and anything else is wallpaper, which feels productive to shoot because it is easy and looks nice.

25. C 4. *Getting the material*

The person who says no is saying no to their manager, in front of colleagues. Give them a route to decline that does not require them to decline to your face: a form collected later, an opt-out address, a deadline after which silence means no. The same principle governs children, where a guardian agrees and the young person should still be asked.

26. A 4. *Getting the material*

Most social platforms strip EXIF on upload; your own site, a PDF, an email attachment and a file-sharing link generally do not, and the platforms that strip it are protecting their users rather than yours. One command with exiftool removes it. Documents have the same problem elsewhere: a PDF from a word processor usually carries the author name and a file path.

27. D 4. *Getting the material*

The United States has fair use, a flexible test applied case by case. The United Kingdom, India and many Commonwealth systems have fair dealing, which covers enumerated purposes such as criticism, review and reporting, and a promotional video is not among them. Platform takedown runs on its own logic regardless. Treat an uncleared excerpt as a risk you are taking knowingly rather than a right you hold.

28. B 4. *Getting the material*

Before the audit both feel like needing more material. After it, one is a shopping list and the other is a rewrite, and confusing them is how people shoot for two more days to prop up a beat that should be deleted. The audit's second half matters too: mark what you will not use, because selection cost scales with the pairs you could compare, not with the number of takes.

29. C 5. *Words that carry the picture*

Text is cut from the right, in more places than people expect: search results, subject lines, link previews, a poster read at an angle. The first version truncates to How we cut our reject rate in, which carries nothing distinguishing. The rule is to reorder rather than to shorten, and the test is whether every noun in the line could be photographed.

30. A 5. *Words that carry the picture*

The preheader is free space and a second fold. Set deliberately, the subject names the thing and the preview adds the detail that makes it matter. The related mechanical fix for the body is to delete your first paragraph and check whether anything is missing: usually the piece starts better, because paragraph two is where you stopped introducing and started saying something.

31. B 5. *Words that carry the picture*

Five pieces of information arriving in two seconds are not retained. Nearly a third is one piece, retained by everybody, and in a spoken piece it is more accurate in the only sense that matters, which is what the listener ends up believing. Two channels, each doing what it is good at: the ear gets the shape and the eye gets the decimal.

32. C 5. *Words that carry the picture*

The first three are real and they look like errors, so a correction pass catches them. Overlapping speech produces something grammatical and plausible, which does not look like an error and therefore survives the pass. It is the same property that makes fluent machine output risky everywhere: the signal you normally use to know where to check is missing.

33. D 5. Words that carry the picture

Alt text is not a caption: a caption is for everybody and alt text conveys what a sighted reader gets from looking. Do not open with photo of, because the screen reader has already said it is an image. A decorative image takes empty alt text so it is skipped; an informative one takes as much text as the information needs, which for a chart can be a paragraph.

34. A 5. Words that carry the picture

Compression removes redundancy, and redundancy is partly what makes text comprehensible, especially for a reader who is tired, on a phone, or reading in a second language. The distinction to hold is between repetition that helps and repetition that is you hearing yourself: an example that makes an idea concrete earns its place, a second statement of an idea already stated does not.

35. B 5. Words that carry the picture

These formulas count syllables per word and words per sentence, and nothing else. They therefore reward chopping a clear sentence into fragments and punish the correct technical term for having four syllables. Optimising a proxy is how you end up with text that is measurably simple and functionally useless, and the real costs sit in idiom, phrasal verbs and elegant variation, which the formula cannot see.

36. C 6. Type, colour and image as meaning

A system invented in advance contains rules for situations you have never met, you break them all, and you learn to ignore the document. A system derived from your own output is a description before it is a rule, which is why you obey it. The part people leave out and most need is the NEVER list, where the mistakes you have already made once are recorded.

37. D 6. Type, colour and image as meaning

A renderer with no shape for a character substitutes an empty rectangle, which is the tofu the Noto project is named for removing. The automatic fallback is often worse than the rectangle, because it does not match anything else on the page. Type a sentence of each script you need into the specimen before committing: thirty seconds now against a rebuild later.

38. A 6. *Type, colour and image as meaning*

At the end of each line the eye jumps back and down, approximately, and corrects. A longer line means a longer jump, so it lands on the wrong line more often, and more line height gives it a larger target to aim at. Forty-five to seventy-five characters with 1.4 to 1.6 line height is the working range, and the two settings move together.

39. C 6. *Type, colour and image as meaning*

The general law does not follow from a single site with one context and one set of surrounding colours. A red button on a red-heavy page is invisible and a green one is conspicuous, and on a green page the reverse holds. The effect attributed to hue is mostly the effect of the difference, which is unexciting and does hold up.

40. B 6. *Type, colour and image as meaning*

Relative luminance is a weighted measure of emitted or reflected light, weighted because the eye is far more sensitive to green than to blue. A colour can look vivid and bright and still sit close to white in luminance. The measurement is free, and the 4.5 to 1 figure is an accepted floor rather than a guarantee: meet it, then look at the thing on a real screen in real light.

41. D 6. *Type, colour and image as meaning*

Close up, a nose is proportionally much nearer the camera than an ear, so the near part is drawn larger; far away, that ratio approaches one and the features flatten. The face has not changed; the relationship between the camera and the face has. Height, distance, the chosen instant and what is in focus are all claims made before any editing.

42. A 6. *Type, colour and image as meaning*

Very often the real thing is a continuum, or the boxes overlap, or one of them is three things nobody has separated. The diagram becomes clearer than the truth and the reader remembers the boxes. This is not an argument against diagrams, since a simplification you can name beats a fog; it is an argument for one line saying what you simplified.

43. C 7. *The surface and the system carrying it*

Vivid blues, bright greens and neon oranges exist on a screen and have no press equivalent, so they are reproduced as the nearest available ink mixture. Convert and inspect before sending. Resolution and bleed are the other two reliable print failures: about 300 dpi for the size being printed, and 3mm past the trim with important content 5mm inside it.

44. D 7. *The surface and the system carrying it*

Read click-through and retention together, always. Optimising purely for clicks selects for the audience that clicks rather than the audience that stays, and recommendation systems have shifted towards watch time and satisfaction, so a high-click, low-retention pattern is actively penalised rather than merely unpleasant.

45. A 7. *The surface and the system carrying it*

There is no specification you can hold a platform to, so the defence is structural rather than technical. Archive the project file and a wide master, not only the delivered video. Exporting fixed crops in advance does not help, because the change you are defending against is a crop you did not know you would need.

46. C 7. *The surface and the system carrying it*

Editors fade in, so frame zero is frequently the darkest frame in the piece, and an automatic heuristic optimises for exposure and sharpness, which are not the same as meaning. A crisp empty workbench beats a slightly soft human face on every criterion a machine has and none that matter. Better still, shoot a deliberate still on location rather than hunting for a frame.

47. B 7. *The surface and the system carrying it*

The look stood out because it differed from its surroundings, and adopting it moves you into the surroundings, which is why chasing the current style is self-defeating. The method that follows is to mock the design into a screenshot of a real feed, at the real size, in both the light and the dark theme, and look at it for one second.

48. D 7. *The surface and the system carrying it*

A proxy correlates with the real thing across the range of behaviour that existed when it was chosen. Optimising hard pushes you outside that range, into the region where the correlation breaks. Nobody has to be dishonest for it to happen, which is why the durable rule is to optimise the form and never the promise.

49. A 7. *The surface and the system carrying it*

A rise means people went back. A deliberate one, where a viewer returns for a number, a detail or a joke, is a strong signal. An accidental one usually means something was illegible or too fast the first time, which is a defect rather than a compliment. The curve is the only number that points at a place in the file, so go and watch 0:41.

50. C 8. *Judgement, and the work that accumulates*

Real problems are not equally likely, and distinct groups of users hit different ones. So five is enough to find the obvious failures and not enough to conclude there are none. That is still an excellent trade for twenty minutes, provided you do not treat a clean run as evidence that the piece works.

51. A 8. *Judgement, and the work that accumulates*

They are different instruments and confusing them wastes both. A designer diagnoses against craft standards and will not tell you that nobody outside the trade understood the piece; a critique from a peer who has not read the brief is a critique of craft alone, and craft is not what the piece was for.

52. D 8. *Judgement, and the work that accumulates*

It is their business, their money and their risk, and they may know things you do not. A grudging execution is worse than either option and is visible. Keeping the other version is worth more than having been right, because sometimes the result comes back and the conversation reopens. Sabotage, sulking and bringing it up afterwards all damage you more than the decision ever did.

53. B 8. *Judgement, and the work that accumulates*

Drip pricing works because the first figure sets the reference point and every later one is judged as a deviation from it, by somebody who has already spent effort getting to step four. It is worth being able to name each pattern's mechanism, because the argument that persuades a client is about how the mechanism behaves after they are noticed, and each of them works once per person.

54. C 8. *Judgement, and the work that accumulates*

The failure is at the level of the aggregate rather than the piece, which is why each maker can defend their own work and the people in it still have to live next to the result. The question that catches most of these is whether the subject would recognise themselves and recognise it as fair, which is not the same as whether they would be flattered.

55. A 8. *Judgement, and the work that accumulates*

A signed record of how an image was made is stripped the moment somebody photographs their screen or saves it through something that does not preserve it.

Labelling duties differ by country and keep moving, and this is not legal advice; the design fact is independent of all of it. If the disclosure matters, it has to be in the picture.

56. D 8. *Judgement, and the work that accumulates*

You can see every element that is present and none of the ones that were considered and dropped: the shot not used, the information withheld until later, the thing a weaker version would have explained. Hold your account of it loosely, because deconstruction is reconstruction and a good deal of what you attribute to a decision was an accident or a constraint.