Forwards or Backwards

https://forwardorbackwards.com



Learning Photography Course – 2. The Exposure Triangle

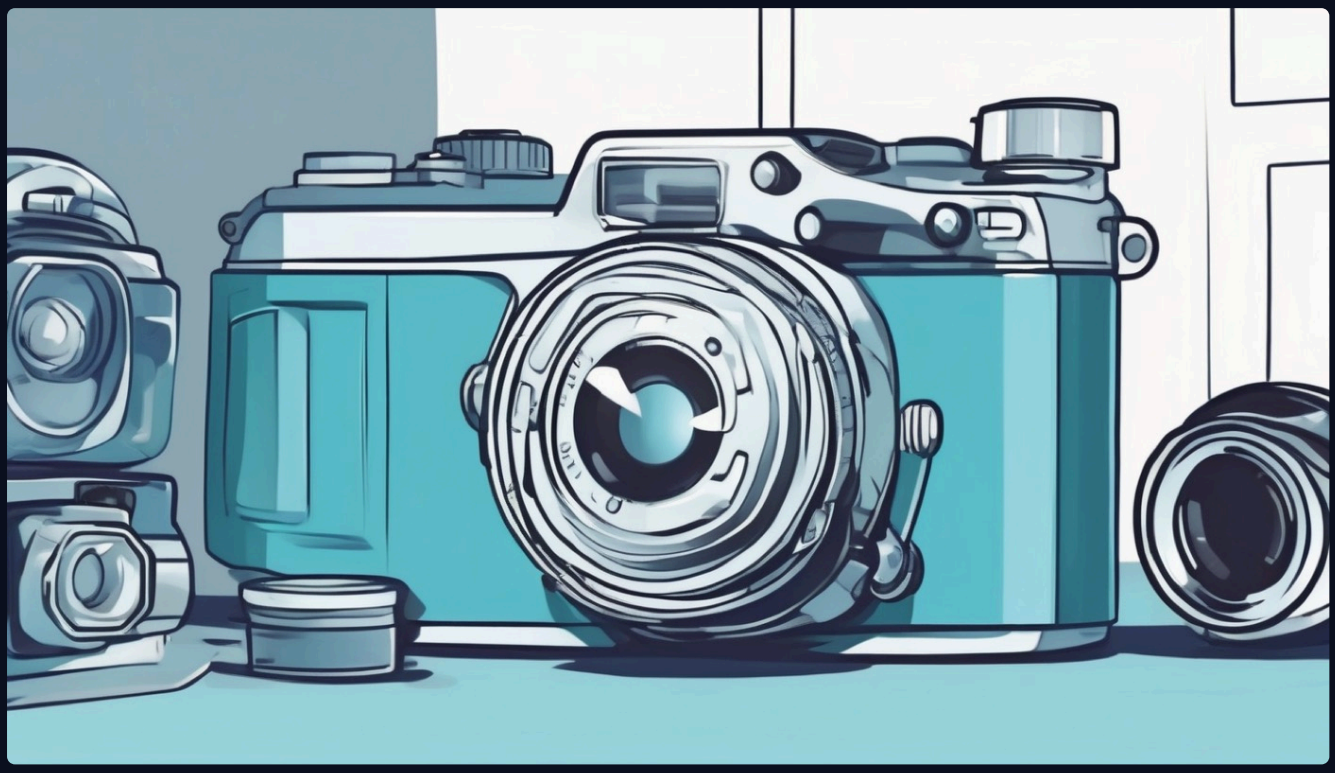
Watch on YouTube: https://www.youtube.com/watch?v=u4ItiX0P0VQ

Chapters

1. 1 The stop
2. 2 Aperture
3. 3 Shutter speed
4. 4 ISO
5. 5 The triangle as one system
6. 6 The shooting modes
7. 7 Metering
8. 8 Exposure compensation

9. 9 The histogram
10. 10 Reciprocity, and where it fails

1. 1 The stop



There is exactly one piece of arithmetic in photography, and once you have it, the rest of this module is bookkeeping.

Here it is. A stop is a doubling, or a halving, of the quantity of light. That is the whole thing. One stop more light is twice as much. One stop less is half as much. There is no second meaning hiding behind the word, and every control you are about to meet is measured in it.

What makes that useful is not the doubling itself. It is that the doubling lets you compare things which otherwise have nothing in common.

Consider what you are actually holding. The aperture is a hole, described as a fraction of the lens's focal length. The shutter is a duration, measured in seconds. The ISO setting is an index applied to a measurement, and it has no physical unit at all. A hole, a length of time, and an index. There is no sensible way to add those together. You cannot subtract a sixtieth of a second from f/8. The question does not mean anything.

But you can say that opening the aperture by one stop, slowing the shutter by one stop, and raising the ISO by one stop each do the same thing to the brightness of the resulting picture. Each of them doubles it. Once all three are expressed in stops they become interchangeable, and the exposure triangle stops being a diagram in a beginner's guide and becomes a currency.

That is what the stop is for. It is an exchange rate between a distance, a time and an index.

It is worth pausing on how unusual that is. In most of engineering, if you want to compare two unlike quantities you have to do real work — convert units, carry dimensions, check that the algebra survives. Photography gets to skip all of it, because somebody decided a very long time ago to measure every control in the same relative unit rather than in its own absolute one. The lens is not marked in millimetres of opening. The shutter dial is not marked in milliseconds, even though milliseconds are exactly what it is doing. Both are marked in a series where one step is one factor of two, and that decision is the reason a photographer can think about exposure at all while a subject is walking away.

Now, cameras do not only move in whole stops. Most cameras made in the last thirty years move in thirds, and many offer halves as an option. A third of a stop is a factor of two raised to the power of one third, which is about one point two six. A half stop is the square root of two, about one point four one four. Keep both of those numbers in mind, because they are about to appear somewhere that will look like a coincidence and is not one.

Third stops matter more than they sound. The difference between a picture that is correctly exposed and one that is a third of a stop bright is invisible to most people in most circumstances. The difference between correct and a full stop bright is not invisible at all — a full stop is a doubling, and a doubling is obvious. Thirds give you somewhere to stand between "fine" and "clearly wrong", and most real exposure adjustment happens in exactly that space.

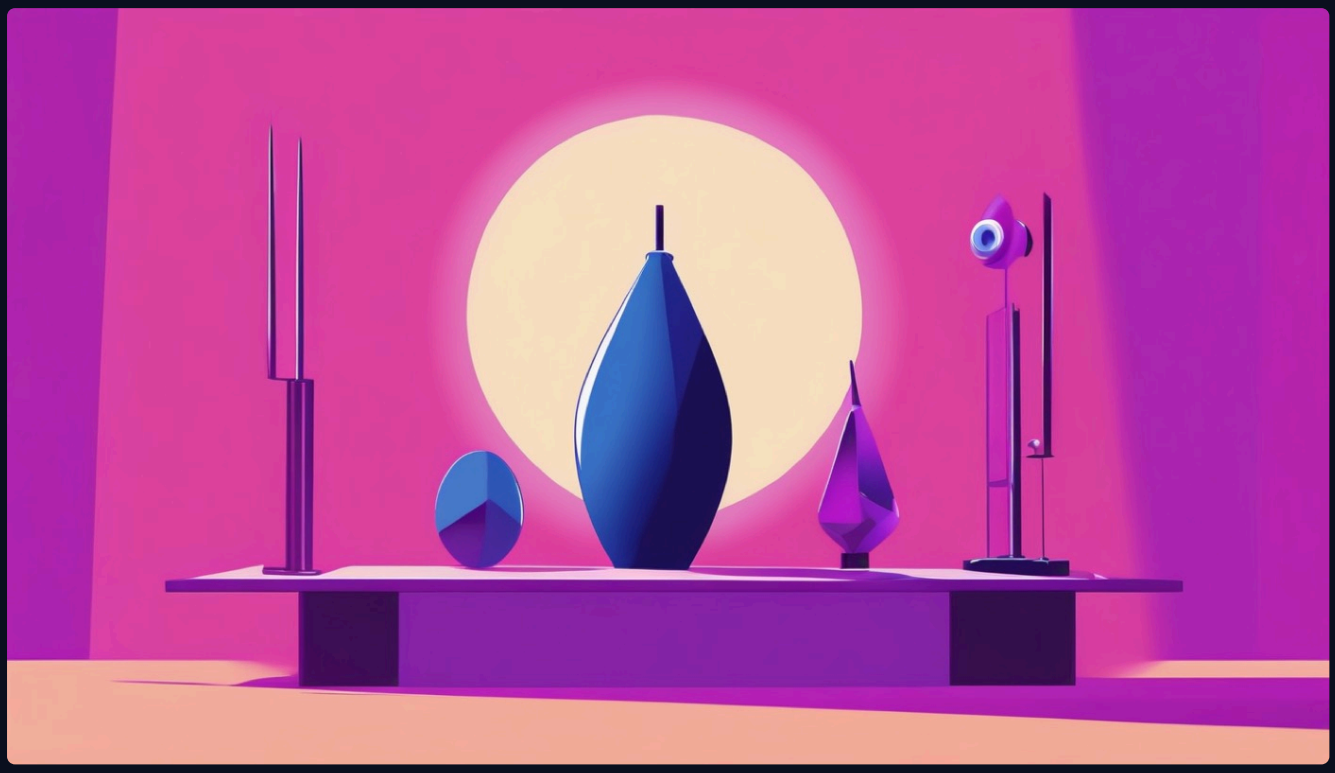
There is a practical consequence to all this that will save you frustration with any camera you ever pick up, including one made before you were born.

Every camera made since about nineteen fifty tells you where it is in stops, and almost none of them use the word. On a dial with detents, one click is usually a third of a stop. In a viewfinder, that little horizontal scale with a zero in the middle and marks running either side is a stop scale, and each mark is typically a third of a stop. When the indicator sits two marks to the right of zero, the camera is telling you it intends to give two thirds of a stop more light than its meter believes is correct.

That is the whole language. Learn to read that one scale and you can read any camera: a nineteen-sixties rangefinder, a modern mirrorless body, a borrowed camera with menus in a language you do not speak. The numbers on the dials differ. The scale does not.

And it is the reason this module comes second rather than fifth. Nothing else in photography makes sense until the unit does.

2. 2 Aperture



The aperture is a hole, and the hole is adjustable.

Inside the lens there is an iris — a ring of overlapping blades that slide across one another to open and close a roughly circular opening. You can usually see it if you look into the front of a lens while changing the setting: a many-sided hole growing and shrinking. It is a genuinely old piece of engineering, essentially unchanged in principle for a century and a half, and it does exactly one job. It decides how much of the light arriving at the front of the lens is allowed to continue to the back.

The number that describes it is the f-number, and it is the most misread number in photography. Here is the definition. Everything confusing about it dissolves once the definition is in place.

The f-number is the focal length of the lens divided by the diameter of the entrance pupil — the effective opening as seen from the front.

So on a fifty millimetre lens, f/2 means an opening of fifty divided by two, which is twenty-five millimetres across. On that same lens, f/16 means fifty divided by sixteen, which is three point one two five millimetres. That is the arithmetic of this chapter, and it is now finished.

Look at what the definition forces to be true. The f-number is a fraction with the focal length on top and the opening on the bottom. A bigger number on the bottom makes a smaller opening, in exactly the way that a sixteenth of a cake is smaller than a half of one. Nothing about the scale is backwards. It only looks backwards if you read "f/16" as a size rather than as a division — and the slash in the notation is trying quite hard to tell you it is a division.

There is a second thing the definition explains, which almost never gets said out loud. It explains why f/4 means the same brightness on every lens ever made.

A fifty millimetre lens at f/4 has a twelve and a half millimetre opening. A four hundred millimetre lens at f/4 has a hundred millimetre opening — eight times the diameter, sixty-four times the area, an enormous and expensive piece of glass. Yet both deliver the same illumination to the film or the sensor, because the longer lens is also spreading its image over a correspondingly larger area. The f-number is not a measure of the hole. It is a measure of the hole relative to the job it is doing. That is why it transfers between lenses, and why a light meter can give you a usable answer without knowing what lens you own.

Now, the sequence. Full stops of aperture run one, one point four, two, two point eight, four, five point six, eight, eleven, sixteen, twenty-two. Written out like that it looks like somebody made it up on a bad afternoon. It is in fact the tidiest thing in the entire system.

The light passing through a hole depends on the hole's area. Area varies with the square of the diameter. So to halve the area you do not halve the diameter — you divide it by the square root of two, which is one point four one four. Multiply along the sequence by one point four one four and you get one, one point four one four, two, two point eight two eight, four, five point six five seven, eight, eleven point three one, sixteen, twenty-two point six three.

Which is the sequence engraved on the lens, rounded. And the rounding is where the strangeness came from. Your camera says eleven because eleven point three one would be fussy. It says twenty-two because twenty-two point six three would be worse. The numbers were never arbitrary. They were tidied, and the tidying hid the pattern.

Notice too that the sequence alternates. One, two, four, eight, sixteen — the powers of two. And one point four, two point eight, five point six, eleven, twenty-two — those same powers of two multiplied by one point four one four. Two interleaved doubling sequences, shuffled together. Once you see that, you can extend the scale in either direction in your head, which is occasionally useful and permanently satisfying.

And there is one more thing the aperture does, which is the reason it is the setting most photographers reach for first.

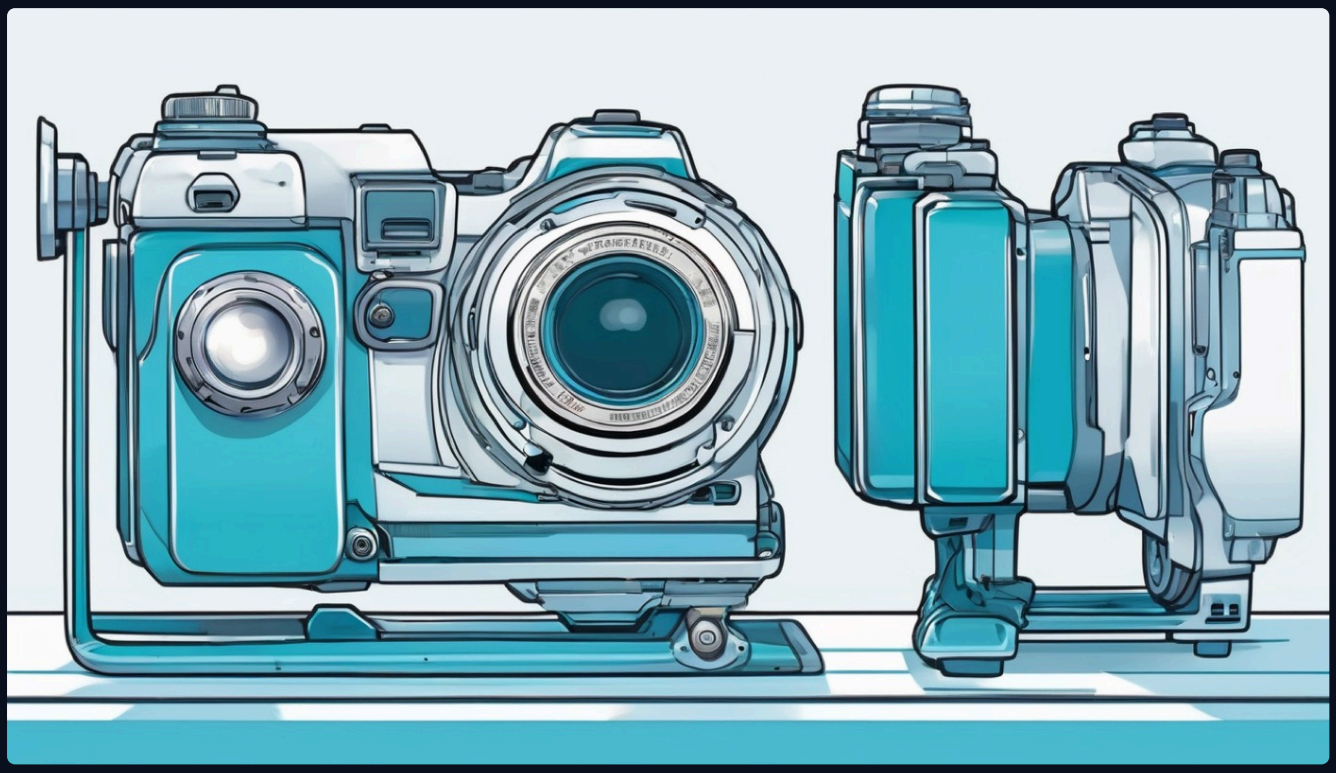
Changing the aperture changes how much of the scene is acceptably sharp — the depth of field. A wide aperture, which is a small f-number, gives a shallow slice of sharpness: the eyes sharp, the ears already softening, the background dissolved to nothing. A narrow aperture, a large f-number, gives a deep one: the fence post in front and the mountain behind it both acceptably crisp.

Why that happens is a chapter of module five, and we are not going to do it here. But the consequence belongs here, because it changes how you think about the triangle. The aperture is

rarely a neutral choice. When you open up to gather light, you are simultaneously thinning the sharp zone, whether or not that is what you wanted. You have made two decisions by turning one ring.

That is true of no other control in photography, and it is why aperture is usually the setting a photographer decides first and defends hardest.

3. 3 Shutter speed



The shutter decides how long the light is allowed to arrive.

There are two designs worth knowing, because they behave differently, and the difference turns up in the pictures.

A leaf shutter sits inside the lens, close to the aperture, and opens and closes rather like a second iris. Because it opens from the centre outward, every part of the frame is exposed during the same interval. That has a consequence for flash which is genuinely useful, and we will pick it up in the lighting module.

A focal-plane shutter sits just in front of the sensor or the film, and it is two curtains travelling one after the other. The first curtain opens; after the set interval, the second curtain follows it across and closes the gap behind it. At slower speeds the first curtain has entirely cleared the frame before the second one sets off, so for a moment the whole frame stands open. At fast speeds it does not: the second curtain starts before the first has finished, and what actually travels across the frame is a moving slit. Different parts of the picture are exposed at slightly different moments — a few milliseconds apart, but not simultaneously.

Almost every interchangeable-lens camera uses a focal-plane shutter, for a practical reason. Putting one shutter in the body is cheaper than putting a shutter in every lens, and it lets you change lenses freely. The leaf shutter survives in some medium-format and technical cameras, where what it buys is worth what it costs.

The speed sequence halves, as by now you would expect: one second, a half, a quarter, an eighth, a fifteenth, a thirtieth, a sixtieth, a hundred and twenty-fifth, a two hundred and fiftieth, a five hundredth, a thousandth, and onward. And it carries exactly the same rounding habit as the aperture scale. A fifteenth is really a sixteenth. A hundred and twenty-fifth is really a hundred and twenty-eighth. Halve repeatedly from one and you get one, two, four, eight, sixteen, thirty-two, sixty-four, one hundred and twenty-eight, two hundred and fifty-six — and the camera rounds every one of them to something that fits on a dial.

Which is the same lesson as the aperture scale: the odd-looking numbers are not odd. They are powers of two wearing a friendlier coat.

Now, blur. There are two completely different kinds, they are constantly confused, and the confusion matters because they have different cures.

Subject motion blur is the subject moving during the exposure. A running child, a passing car, the tip of a bird's wing. The only cure is a shorter exposure. Nothing else works — not a tripod, not image stabilisation, not a steadier grip, not holding your breath — because the thing that moved was not the camera. This is the single most common misunderstanding about stabilisation, and it accounts for an enormous number of disappointing indoor photographs of children.

Camera shake is you moving during the exposure. The whole frame smears together, because the whole frame moved together. The cures are a shorter exposure, or a tripod, or stabilisation, and all three genuinely work.

The old guide for camera shake is the reciprocal rule: use a shutter speed at least as fast as one over the focal length. A fifty millimetre lens wants a fiftieth of a second or faster. A two hundred millimetre lens wants a two-hundredth.

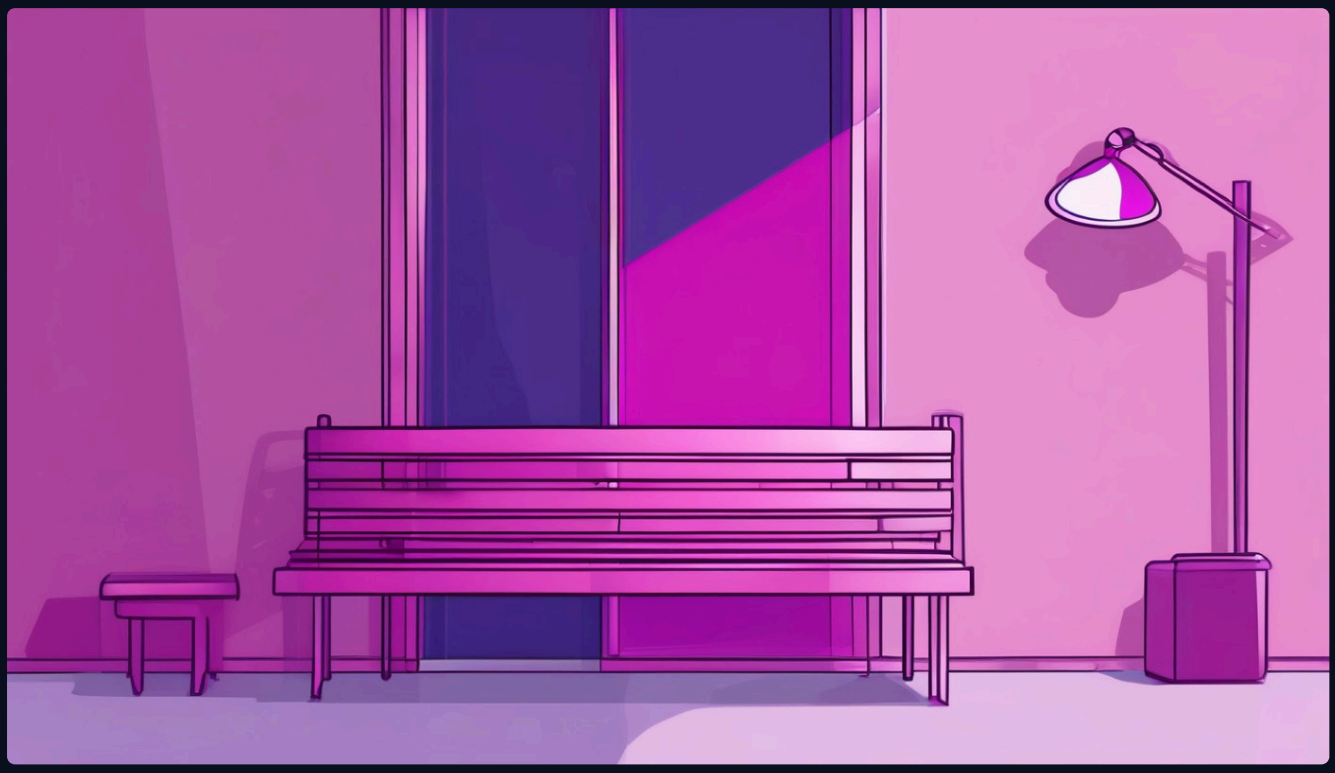
I want to be plain that this is a rule of thumb inherited from thirty-five millimetre film, not a standard anybody publishes, and that it was calibrated against what looked acceptable in a modest print. Examine the same file at full size on a screen and you will often find you need a stop or two more than the rule suggests. It is a floor, not a guarantee.

Stabilisation changed the arithmetic. Moving lens elements, or the sensor itself, to counteract small involuntary movements buys back real shutter speed, and manufacturers quote the benefit in stops.

I am going to tell you what I am not going to do, which is quote a number. The measurement behind those claims comes from a standardised industry test, and I have not put my hands on that standard, so I am not going to repeat a figure I cannot source. What is safe to say, and what actually matters when you are holding the camera, is this: stabilisation helps with camera shake, it helps substantially, and it does absolutely nothing whatsoever about a moving subject.

That second half is the half everyone forgets, and it is the half that explains the blurred photograph.

4. 4 ISO



ISO is the setting that causes the most trouble, because the usual explanation of it is wrong.

The usual explanation is that raising the ISO makes the sensor more sensitive to light. It does not. It cannot. A sensor's response to arriving photons is a property of the silicon, fixed when the wafer was fabricated, and no menu setting reaches into the physics and changes it.

Here is what ISO actually is, and the standard states it directly. ISO twelve two three two, in its two thousand and nineteen revision, defines the exposure index as K divided by H sub m — where H sub m is the average exposure at the focal plane, measured in lux-seconds, and K is a constant fixed at ten lux-seconds. Both of those come from the standard, which covers digital still cameras.

Read that as English and it says: the ISO number is an index that tells the camera what quantity of light the finished picture is to be built around. Change the index and you change what the camera considers a correct exposure, and therefore how much it amplifies the signal it collected. You have not changed the collecting. You have changed the bookkeeping applied to it afterwards.

That distinction sounds academic right up until you see what it costs, and then it becomes the most practical idea in the module.

Noise. The graininess in a photograph comes from several places at once, and they are not equally the camera's fault.

The largest contributor in most conditions is not a fault at all. Light arrives as discrete photons at random intervals, so any measurement of light carries an inherent randomness — count photons for a fixed time twice and you will get two slightly different answers. That is shot noise, and it is present in the light before the camera is involved at all. Its relative size falls as you collect more light: gather four times the photons and the noise grows only twofold, so the ratio improves. More light is a cleaner picture, always, and that single sentence explains most of what follows.

Then there is read noise, generated by the electronics as the collected charge is measured, and pattern noise, which is fixed variation between one photosite and its neighbour.

Raising the ISO adds no shot noise, because shot noise came in with the light. What raising ISO does is amplify a signal collected from less light — and a smaller signal sits proportionally closer to the noise floor. The picture looks noisier because you gathered less light. You gathered less light because you used a faster shutter or a smaller aperture. And you were able to do that because you raised the ISO.

So the honest sentence is this. High ISO does not create noise. High ISO is what you do when you cannot collect much light, and not collecting much light is what makes noise visible. The setting is the symptom, not the disease.

One consequence is worth knowing about your own camera. Every camera has a base ISO, usually the lowest full setting on the dial, at which the signal is largest relative to the noise and the dynamic range is at its widest. That is the best that camera will ever perform, and no setting above it improves on it. Everything higher is a trade you have chosen deliberately — usually to freeze something that was moving, or to make a photograph at all.

Which brings the practical point, because this is where beginners get stuck and lose pictures. A noisy photograph of the thing is worth more than a clean photograph of nothing. Cameras made in the last decade are extraordinarily good at high ISO, noise reduction has improved further still, and a slightly grainy frame of a moment that happened once beats a pristine blurred smear every time.

Raise it when you need it. That is what it is for.

5. 5 The triangle as one system



Now we can put the three together, and all that work on the stop pays for itself.

Because all three controls are measured in stops, they can be traded against one another. Open the aperture one stop and halve the shutter time, and the picture has exactly the same brightness as before. Raise the ISO one stop and close the aperture one stop, same result. These are equivalent exposures: the same quantity of light, indexed the same way, arrived at by a different route.

But — and this is the entire point of the chapter — they are not the same photograph.

Take one scene and expose it three ways.

At f/2, a thousandth of a second, ISO one hundred: a shallow zone of sharpness, and anything moving is frozen solid.

At f/8, a sixtieth of a second, ISO one hundred: identical brightness, but now deep sharpness front to back, and anything that moved has smeared.

At f/8, a thousandth of a second, ISO sixteen hundred: identical brightness again, deep sharpness, frozen motion, and more visible noise.

Three identical brightnesses. Three completely different photographs.

Brightness is the one thing the triangle holds constant, and brightness is very nearly the only thing you do not need to care about, because it is the one thing you can adjust afterwards. Everything you

do care about — what is sharp, what is frozen, how clean the file is — is what varies between those three.

Which gives you a working method, and it is far simpler than the diagram suggests. Do not attempt to balance three things simultaneously. Decide which two you actually care about, set those two, and let the third land where it must.

A portrait where you want the background soft: the aperture is the decision. Set it wide. Pick a shutter fast enough not to shake. Let the ISO go wherever it goes.

A waterfall you want to smooth: the shutter is the decision. Set it long. Put the camera on something solid. Close the aperture until the exposure balances at base ISO.

A gymnast indoors: the shutter is a decision and it must be fast. The aperture is a decision and it must be wide. The ISO is therefore not a decision at all — it is a consequence, and it will be high. That is not a failure. That is the arithmetic telling you the truth about the room.

Two worked examples, properly.

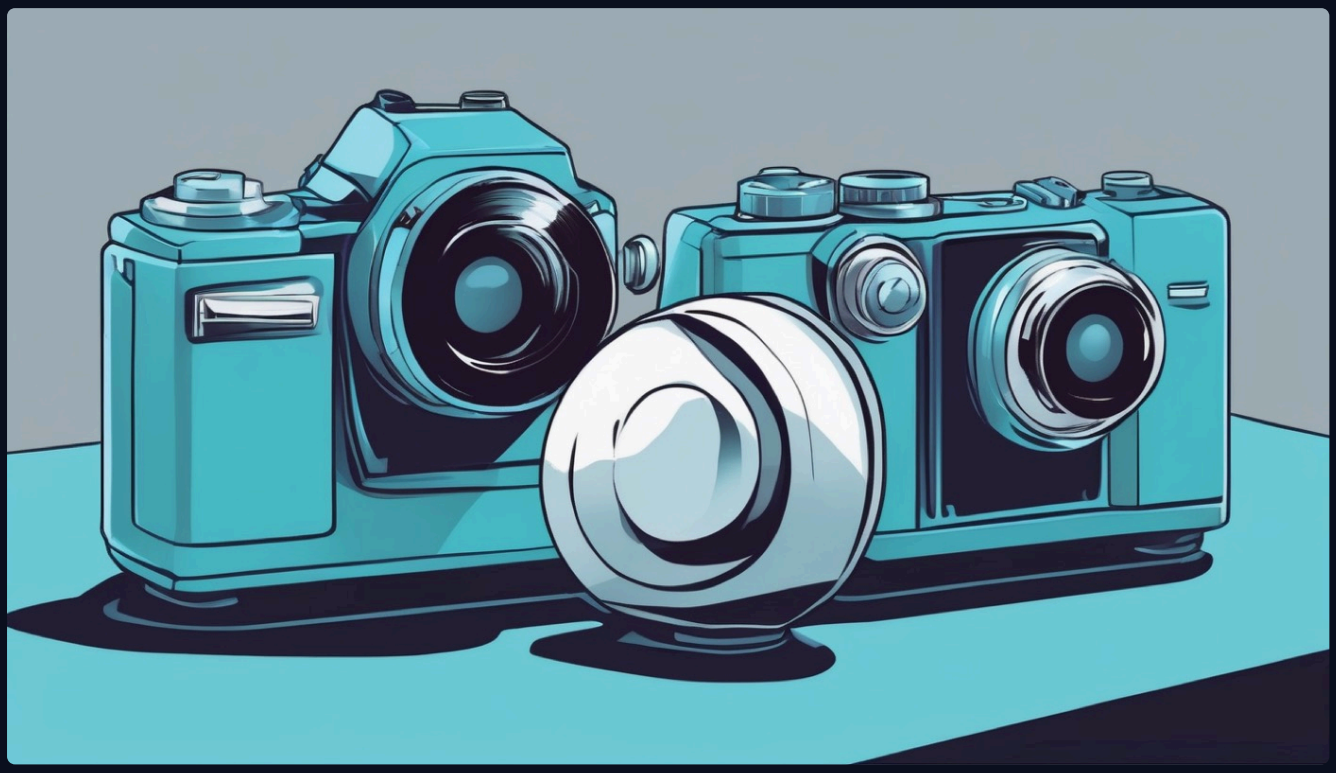
Indoors, a child moving, light from a window. To freeze a child you want somewhere around a two-hundred-and-fiftieth of a second. The room is dim, so you want the lens wide open, say f/1.8. Both of those are decisions about the picture, not about the exposure. The ISO is then whatever the meter says makes those two work — sixteen hundred, perhaps, or thirty-two hundred. You have not compromised anything. You have spent ISO, which is the cheapest of the three, in order to protect the two you cared about.

A landscape at dusk. Nothing in the frame is moving, so shutter speed costs you nothing and is free to be anything at all. You want front-to-back sharpness, so f/11. You want the cleanest file the camera can give, so base ISO. The shutter is then whatever those two demand — four seconds, say — and since the camera is on a tripod and nothing in the frame moves, four seconds is free.

Same triangle. Opposite reasoning. In both cases: two decisions, one consequence.

That is the method, and it will carry you through the remaining eight modules.

6. 6 The shooting modes



The mode dial is the camera asking which of the three you would like to hold.

Aperture priority — A on most cameras, Av on Canon — means you set the aperture and the camera chooses the shutter to match. It is the right default for the majority of photography, for the reason we established a moment ago. Aperture is usually the decision that changes the picture most, so handing it to the camera means handing away precisely the choice you most wanted to make.

Shutter priority — S, or Tv — is the reverse. You set the shutter, the camera sets the aperture. It is correct whenever motion is the point: panning a cyclist, smoothing water, freezing a bird's wing.

There is a catch with shutter priority worth understanding, because it fails in a way aperture priority mostly does not. The aperture range is bounded by the lens — perhaps f/3.5 to f/22, a range of about six stops. The shutter range is enormous by comparison: thirty seconds to a four-thousandth or beyond, some nineteen or twenty stops. So if you demand a two-thousandth of a second in poor light, the camera will open to its widest aperture and simply fail to get there. It will usually blink a warning at you, which is easy to miss when you are looking at the subject.

Shutter priority can write a cheque the lens cannot cash. Aperture priority rarely can.

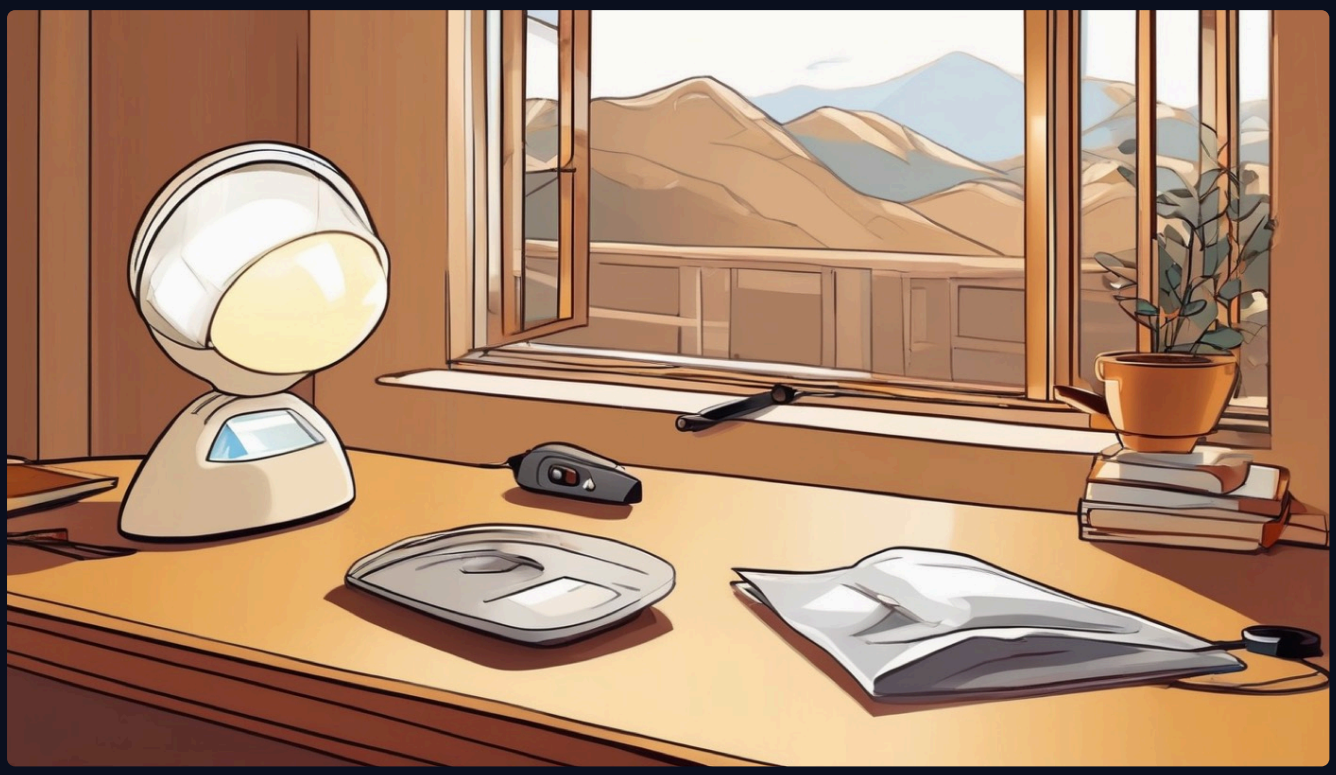
Manual means you set both, and on a modern camera it is far less austere than it sounds, because of auto-ISO. Set the aperture you want. Set the shutter you want. Let the camera move the ISO so that those two can stand.

That is not a dilution of manual. It is manual applied to the two settings that change the picture, with the consequence automated — which is exactly the method from the last chapter, written down as camera settings. For a great deal of real work it is the most sensible position on the dial.

And Program. P sets both, and lets you shift the pairing — turn a wheel and it slides along the equivalent exposures, offering a wider aperture with a faster shutter, or the reverse, at constant brightness. It is not the shameful mode. It is a perfectly reasonable place to be when the light is changing faster than you can think about it, and when you would rather spend your attention on the subject than on the dials.

Nobody is grading you on which letter the dial points at. All the modes set the same three numbers. They differ only in which of them you have reserved for yourself.

7. 7 Metering



Every automatic exposure rests on a measurement, and it is worth knowing what that measurement assumes, because the assumption is exactly where automatic exposure goes wrong.

The meter in your camera is a reflected-light meter. It reads the light coming back off the scene toward the lens. And it has no way of knowing whether it is looking at a dark subject under bright light or a bright subject under dim light, because both send back the same amount of light. The information is simply not there to be had.

So it assumes. It assumes the scene averages out to a mid-tone, and it sets the exposure so that the average comes out in the middle.

For most scenes that assumption is perfectly sound, because most scenes contain a spread of light and dark that does average to something middling. A street, a room, a face against a background — all fine.

It breaks in two places, and they are mirror images of each other.

Photograph a snowfield. The meter sees far more light returning than an average scene returns, concludes the scene is too bright, and cuts the exposure. Your snow comes out grey.

Photograph a black cat on black cloth. The meter sees very little light returning, concludes the scene is too dark, and opens up. Your black comes out grey.

In both cases the meter did precisely what it was designed to do. In both cases the photograph is wrong. The meter is not broken. It is answering a different question from the one you asked.

I should be straight with you about a number here, because you will meet it everywhere. Middle grey is quoted as eighteen per cent reflectance almost universally, and it is also quoted as twelve per cent, or twelve point seven, by people who know exactly what they are talking about. The difference comes from how the calibration constant is expressed — whether against reflectance or against luminance. I have not traced this one to the standard itself, so I am not going to assert either figure to you.

What matters for taking photographs is not the percentage. It is knowing that the meter is driving everything toward the middle, and recognising the scenes where the middle is the wrong place to be.

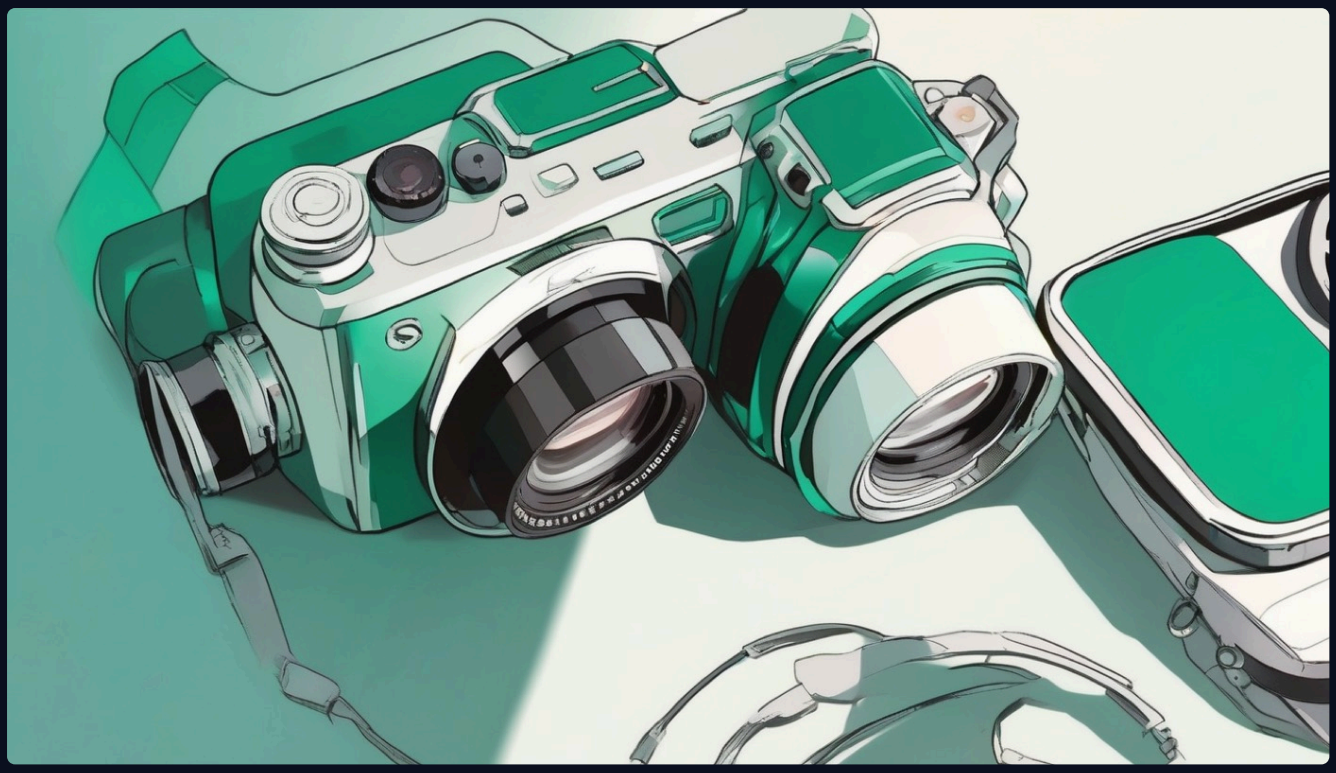
The camera offers you patterns to choose between. Evaluative or matrix metering divides the frame into zones and weights them with a good deal of manufacturer cleverness, often including a stored library of scene types; it is the best general default and it handles the snowfield rather better than its ancestors did. Centre-weighted reads the whole frame but biases the middle, and its virtue is that it is predictable — you can learn what it will do and then rely on it. Spot reads a small circle, typically a few per cent of the frame, and it is the instrument for when you know exactly which part of the scene must be correct and are willing to let the rest fall where it falls.

And there is a way to sidestep the assumption altogether.

An incident meter is a separate handheld device with a white dome, held at the subject and pointed back toward the camera. It measures the light falling on the subject rather than the light coming off it. So it does not care whether the subject is a snowfield or a black cat — it gives the same reading for both, which is the correct reading for both, because the light falling on them is the same light.

That is why studios still use them, six decades after cameras began metering for themselves. Not nostalgia. A different and better question.

8. 8 Exposure compensation



Exposure compensation is the control for telling the camera it is wrong.

It works in any automatic mode. You dial in plus or minus some number of stops, and the camera shifts its exposure by that much from whatever its meter concluded. On many bodies it is a dedicated dial on the top plate, marked from minus three to plus three; on others it is a button held down while turning a command wheel.

Which way to push is the part people get backwards, and the two hard cases are the two we have just met.

Snow. The meter will underexpose it toward grey, so you add light: plus one, plus one and two thirds, sometimes plus two. You are overriding a correction that should never have been applied in the first place.

A dark subject filling the frame. The meter will overexpose it toward grey, so you subtract: minus one, minus one and a third.

The version worth memorising is that you compensate toward the tone you actually want. Snow should be white, so push up. Black should be black, so pull down. Said that way it stops being a thing to remember and becomes a thing you can work out at the time, which is far more durable.

Compensation interacts with manual-plus-auto-ISO in a way that surprises people the first time they meet it. In manual with auto-ISO you have already fixed the aperture and the shutter, so

compensation cannot move either of them. It moves the only variable left, which is the ISO.

That is usually exactly what you want, and it means exposure compensation still functions in manual mode — which sounds like a contradiction until you remember that something in that exposure is still automatic. Compensation adjusts whatever the camera is choosing, whichever setting that happens to be.

And now the most common unforced error in all of photography, which I would like you to hear once so that it happens to you fewer times.

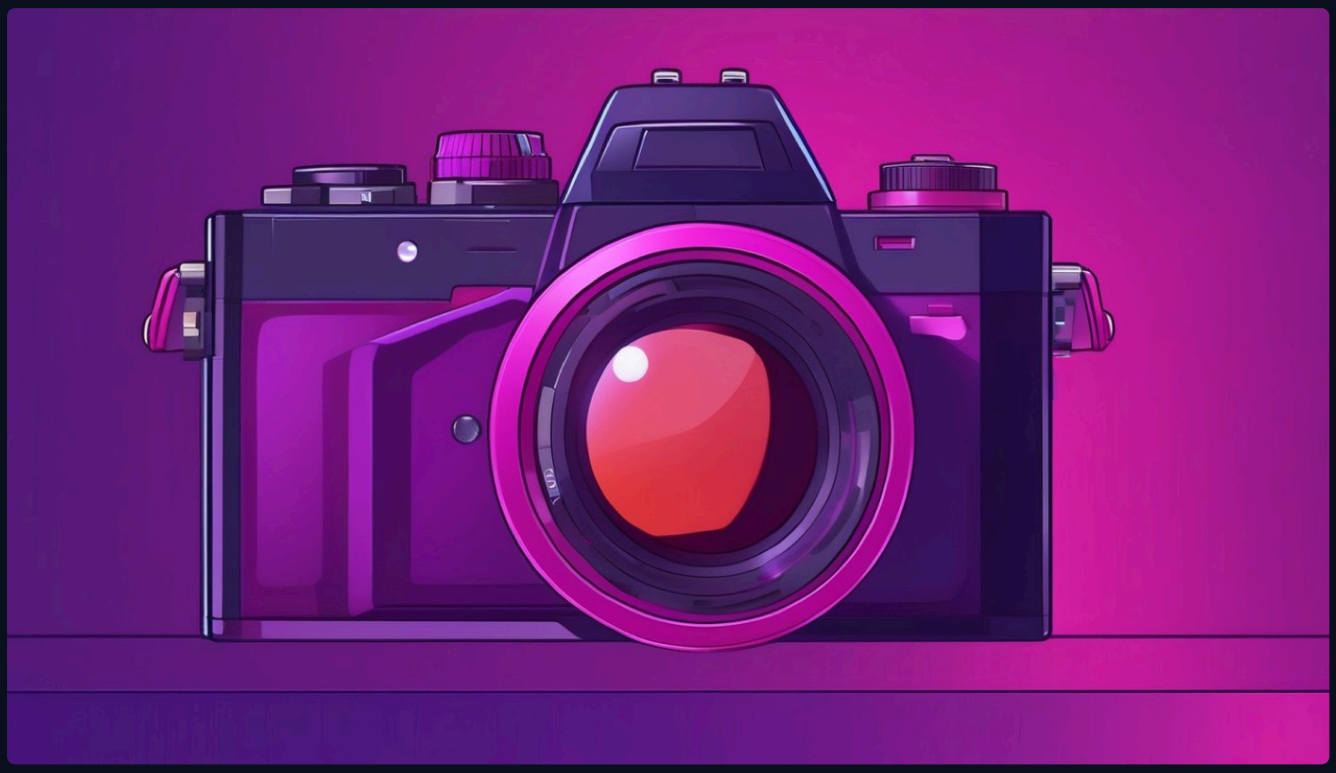
Exposure compensation stays where you left it. You dial in plus two for a snow scene. You go indoors. Every photograph for the rest of the day is two stops bright. On many cameras it survives being switched off and on again. There is no warning, because the camera thinks you meant it — and you did mean it, forty minutes ago.

Check it. Make checking it the first thing you do when you pick the camera up, alongside checking the ISO and the mode. Three seconds, and it retires an entire category of ruined afternoons.

Finally, bracketing. If the scene genuinely matters and you are unsure, take three frames — one at the meter's reading, one under, one over — and choose later at your leisure. The camera will do this for you automatically, usually from a single button.

It is not a failure of nerve. It is insurance, it costs essentially nothing on digital, and on a scene you cannot return to it is simply the sensible thing to do.

9. 9 The histogram



The histogram is the only honest feedback the camera gives you about exposure. The picture on the screen is not it.

The screen's apparent brightness depends on the screen's brightness setting and on the light you happen to be standing in. A photograph that looks perfect on the back of the camera in a dim room is routinely a stop underexposed when you get it home; a photograph that looks hopeless in bright sunshine is often fine. The histogram is indifferent to all of that, because it is computed from the data rather than from your eyes.

What it shows is the distribution of tones across the whole frame. Dark tones to the left, bright tones to the right, and the height at any point tells you how much of the frame sits at that brightness.

It is an account of the picture's tones, and it is not a measure of quality. There is no correct shape.

That is worth saying firmly, because a great deal of advice implies there is — that a good histogram is a nice hill in the middle. A photograph of a black cat in a coal cellar should be bunched hard to the left, and if it is not, something has gone wrong. A high-key portrait on a white background should be bunched to the right. The shape should match the scene, and the scene is yours to choose.

What you are actually looking for is clipping — data piled up hard against either end, where the frame has run out of range.

And the two ends are not equal. This is the practical heart of the chapter.

Highlights clipped to pure white are gone. Every pixel in that region recorded the same maximum value, and there is no longer any information about what was brighter than what. No editing recovers it, because there is nothing to recover — the data is not hiding, it was never written.

Shadows clipped to black are often partly recoverable, because there is usually some small signal down there, buried in noise. Lifting it costs you noise rather than information — a grainy shadow rather than an empty one.

So when you are forced to lose one end, lose the shadows.

That asymmetry is where "expose to the right" comes from. Push the exposure as bright as you can without clipping the highlights, gathering as much light as possible, and pull it back in editing afterwards. More light means a stronger signal against the noise — which is the principle from the ISO chapter arriving from a different direction.

It is sound, and it has limits. You are spending your highlight headroom, and if the scene's brightness range changes while you are shooting — the sun comes out, the subject turns toward the light — you will clip something that mattered. It is a technique with a cost, not a rule to follow blindly.

And there is one catch that explains a great deal of confusion among people who shoot raw.

The histogram your camera displays is computed from the embedded JPEG preview, not from the raw file. That preview has had a tone curve, a colour profile and the camera's contrast settings applied to it, and it has less headroom than the raw data underneath. So the raw file very often holds highlight detail that the camera's display insists is clipped.

Your real headroom is a little wider than the screen admits. How much wider depends on your camera, and the only way to find out is to test your own: photograph a bright scene, clip it on the display, and see what the raw file gives back. Half an hour, once, and then you know for the life of the camera.

10. 10 Reciprocity, and where it fails



We have assumed something for this entire module, and it is time to say it out loud — and then break it.

The reciprocity law says that exposure is the product of intensity and time, and that the product is all that matters. Half the light for twice as long yields the same exposure. That assumption is precisely what makes the stop work as a currency, and it is what allows you to trade aperture against shutter at all. Without it, none of the last nine chapters would be true.

For nearly everything you will ever do, it holds.

On film, past about a second, it stops holding.

At long exposures, light arrives so slowly that film's response to it falls away, and the film behaves as though it is slower than its rated speed. Meter a long exposure on film, follow the meter faithfully, and you will get a thin, underexposed negative. The longer the exposure, the worse the shortfall — and the meter gives you no hint at all, because the meter is measuring light, and the light is exactly what it says it is. The film is the part that changed.

Ilford publish the correction, and because they publish it there is no need for anyone to guess.

In their technical information sheet on film reciprocity failure compensation, dated December two thousand and twenty-three, the correction is given as a power law: the corrected time equals the metered time raised to a film-specific factor. They give that factor for twelve of their films.

For HP5 Plus, the factor is one point three one.

So: meter ten seconds. Ten raised to the power one point three one is twenty point four. You give it twenty point four seconds.

Look at the size of that. Your ten-second exposure was short by more than ten seconds — more than a full stop — and nothing in the camera would have told you.

For FP4 Plus the factor is one point two six, and the same metered ten seconds becomes about eighteen point two.

Ilford also state the threshold plainly, and it is a mercy: exposures of one second or less require no compensation at all. Below a second, reciprocity holds, and you can ignore this entire chapter.

And they add a caveat worth carrying, because it is the part most summaries drop. Very long exposures can increase contrast, because reciprocity failure is not uniform across the tonal range — the shadows fall away faster than the highlights do. So a very long exposure may also want a reduced development time. The correction is not only to the exposure. It reaches into the processing as well.

Digital does not do any of this. There is no reciprocity failure in silicon. A photosite collecting for thirty seconds counts photons in exactly the way it counts them in a thirtieth of a second, and the law holds all the way down.

What digital does instead, on long exposures, is accumulate thermal noise — the sensor warms as it works, and warm silicon generates spurious signal — along with read noise from the electronics. Different mechanism entirely, and different cures: dark-frame subtraction, which photographs the noise with the shutter closed and subtracts it, or simply a colder night.

It is worth being precise about that, because "long exposures are noisy" and "long exposures are underexposed" are two different problems, and only one of them belongs to film.

So the law holds, until it does not, and the point where it lets go has been measured and published by the people who make the film. That is the most useful kind of exception: one that somebody has already done the work on, so that you can look it up rather than discover it on a ruined negative.

Which is the whole module in a sentence. Three controls, one unit that makes them comparable, one law that ties them together — and one clearly marked place where that law lets go.

Module three is every kind of camera: what genuinely differs between them, and what only appears to.

This module is accurate as at the sixteenth of September, two thousand and twenty-six.

Sources and credits

Primary sources

- ISO 12232:2019 — exposure index = K / H sub m, with K fixed at 10 lux-seconds; digital still cameras.
- ILFORD / HARMAN technology Limited, Film Reciprocity Failure Compensation, December 2023 (sheet v2, May 2024) — Tc = Tm^P; HP5 PLUS P = 1.31; FP4 PLUS P = 1.26; no compensation required at one second or less.

Not regulated financial advice.