

THE POST-EVIDENCE WORLD

Photography Was Never Just a Witness

Cedric

*For my father,
who was the cleverest person I've known.*

*For my grandfather,
who was the cleverest person with his hands I've known.*

LIST OF ILLUSTRATIONS

All photographs are the author's own work unless otherwise noted.

Chapter Three

The Schwebbahn above the river Wupper, Wuppertal, Germany, as seen from the driver's cab.

Photograph by the author, March 2026.

Chapter Three

The author photographing in Rute, Spain, using a twin-lens reflex camera.

Photograph by the author's wife, taken without the subject's knowledge.

Chapter Six

The author's grandfather in his workshop.

Photograph by the author. Reproduced with the family's permission.

A U T H O R ' S N O T E

This book grew from a newsletter. Since 2024 I have been publishing essays on photography at imaybewrongphoto.substack.com, under the title 'I May Be Wrong.' The newsletter was, and remains, an attempt to think in public about what photography is and what it does: not instruction, not technique, but the harder and more interesting questions underneath.

Some of the thinking developed there found its way into this book, substantially rewritten for a general reader and for the specific argument the book makes. The taxonomy in Chapter Three, the argument about photographic presence in Chapter One, several of the observations about photographers and their relationship to their practice in Chapters Three and Six: these have roots in posts that appeared in the newsletter over the past two years. Readers who know the archive will recognise the terrain, though not the form.

Chapters Four and Five, and the historical and legal material throughout, are largely original to this book, written specifically for the argument rather than developed from existing work.

The photographs in this book are my own work, with the exception of the image of me in Chapter Three, which was taken by my wife. She had no idea she was illustrating a chapter about the photographer as interpreter. Some coincidences are too good to stage.

Cedric

The Last Certain Image

In the spring of 2004, a series of photographs taken inside Abu Ghraib prison in Iraq began circulating through news organizations around the world. They showed American soldiers humiliating, degrading and in some cases torturing Iraqi detainees. The images were specific: specific people, specific acts, a specific place. They were also, in the language that matters here, indexed to reality. Light had entered a camera lens, had struck a sensor, and had recorded what was there. The photographs existed because those events had taken place.

What followed was consequential in a way that is worth reconstructing carefully, not least because some of the details matter for what comes later. The photographs had not been taken by journalists. They had been taken by soldiers, on personal cameras, as a form of record-keeping or amusement or both. One of those soldiers, Army reservist Joe Darby, spent three weeks deciding what to do with a CD of images a colleague had handed him before passing them to the Army's Criminal Investigation Command in January 2004. The photographs that changed the course of the war existed because soldiers had photographed their own acts, not because anyone had travelled to Abu Ghraib to bear witness. They became evidence despite themselves. Seventeen soldiers and officers were removed from duty. Eleven were court-martialled and convicted on charges including assault, maltreatment and dereliction of duty. Donald Rumsfeld, then Secretary of Defence, offered his resignation twice; George W. Bush declined it both times, but the photographs had made the offer unavoidable. Brigadier General Janis Karpinski, who had commanded the prison, was demoted. Multiple independent investigations were launched, their findings pointing progressively higher up the chain of command without, in the end, reaching it. The International Committee of the Red Cross, which had been raising concerns about conditions at Abu Ghraib in private for months, found that its concerns were now front-page news in every country in the Western world. Allies

expressed outrage. Public support for the war, already declining, fell further and did not recover.

None of that happened because someone wrote a compelling description of what was occurring inside the prison. It happened because of photographs. The distinction matters. Accounts of prisoner abuse in Iraq had already appeared in human rights reports and journalistic investigations before the images were published. They had produced concern in some quarters and denial in others, but they had not produced Senate hearings or court martials. Photographs produced those things, because photographs occupied a different category of evidence. They were not claims about what had happened. They were, in a sense that felt legally and morally binding, proof.

This is what I want to examine before anything else: the mechanism that made those photographs do what they did. The politics, the ethics of their publication, the subsequent argument about whether the soldiers depicted were scapegoats for a broader institutional failure: all of that matters, and none of it is the point here. The point is why photographs had that power, and what it rested on.

* * *

Photography has always worked differently from other ways of representing the world. The difference is not aesthetic and it is not a matter of quality or intent. It is physical.

A painting of a scene is the product of a mind. The painter observed or imagined something, translated it through their hand, and produced an object that reflects their interpretation. The painting might be extraordinarily accurate, might be based on careful study and direct observation, but its connection to what it depicts is mediated entirely by human intelligence and human skill. A novelist describing a burning village is doing the same thing at a further remove: processing observations or research or imagination into language, which a reader then processes back into mental imagery. The chain from event to representation runs entirely through human minds.

Photography changed this relationship in a way that seemed, for a century and a half, to be philosophically significant. A camera does not interpret. It records. Light travels from objects in the world, enters a lens and leaves a trace on whatever light-sensitive surface is waiting for it. The trace is caused by the thing it depicts. A photograph of a burning village could not exist without the burning village. A photograph of a man's face in pain could not exist without the man, without the pain, without the moment. The causal chain from event to image is not mediated by human imagination. It runs through physics.

Philosophers call this the indexical nature of photographs: an index is a sign that points to its cause, the way a footprint points to a foot or smoke points to fire. Photographs are indices of the events that produced them. This is what separates them from paintings or drawings or written descriptions, which are icons or symbols, representations that resemble or refer to things without being caused by them. You can paint something that never existed. You cannot, in any traditional sense, photograph it.

For most of photography's history, this property was understood intuitively even by people who had never heard the word indexical and had no interest in the philosophy of representation. You did not need a theory of photographic ontology to understand that a photograph was different from a painting. You felt it when you looked at old family photographs: the specific quality of presence that comes from knowing that the light which made this image once bounced off the face of someone who is now dead. You felt it in the difference between seeing a painting of a war and seeing a photograph of one. The painting might be technically superior, compositionally more powerful, more emotionally affecting. The photograph carried something the painting could not: the knowledge that this happened.

That knowledge is what gave the Abu Ghraib photographs their force. Their composition was artless, their aesthetic quality nil, the person who made them nobody in particular. Their force came entirely from the fact that they were caused by the events they depicted. Someone had been there. The camera had been there. The light had been there. And the images existed because all three had been present at the same moment.

* * *

Photography requires something other arts do not: you have to be there. This might seem obvious, but it is actually a significant constraint that sets everything about what photographs can mean and what they can do.

Consider what painters and writers are free to do that photographers are not. A painter can work entirely from imagination or memory, can construct scenes that never took place, can paint a landscape they have never visited based on description or invention. A novelist can describe people they have never met, events that never happened, places that exist only in their own mind. Their presence during the act of creation is required, but presence at the subject is not.

Photography removes that freedom. The camera must be positioned somewhere real, aimed at something that exists, recording light from actual objects. Even the most staged and constructed photograph requires real people, real objects, real light. You cannot photograph a mountain from memory. You cannot photograph a person you have never encountered. The camera must be there, and someone must have decided to bring it.

This is not merely a technical constraint. It is the source of photography's moral and evidential weight. When Nick Ut photographed nine-year-old Kim Phuc running down Route 1 in Trang Bang in June 1972, her clothes burned away by napalm, he was present. Not in an abstract sense, but in the specific, physical, irreversible sense that he was standing on that road at that moment, and the camera he was holding recorded light that had bounced off an actual child in actual pain. The image exists because he was there and because she was there and because that moment happened. It could not have existed otherwise.

That photograph appeared on the front pages of newspapers across the world. It is credited with shifting American public opinion about the Vietnam War, though such attributions are always more complicated than they appear. What is documentable is that it contributed to a public reckoning about what was being done in the name of the United States, and that its contribution rested entirely on what it was: not an illustration of the suffering of Vietnamese civilians but a record of it. The difference mattered. A painting of the same

scene, however harrowing, would have been a statement about the war. The photograph was evidence of it.

The photographer's presence was not incidental to this. It was the mechanism. Ut had to be on that road. He had to be there when Kim Phuc ran past him. He had to make the decision to raise the camera and press the shutter at that moment rather than any other. The image carries the trace of all those decisions, and it carries the trace of all those decisions because a human being was present to make them. Photography is the art form that cannot be faked from a distance. The photographer is always, in some irreducible sense, a witness.

* * *

There is a contract implicit in all of this, and it is worth naming precisely because it is one we entered into without quite noticing, and are now leaving the same way.

The contract was never stated explicitly. No government passed a law declaring photographs to be evidence. No philosophical consensus was formally ratified. It was more like a collective agreement that developed over a century and a half of photographic culture: photographs occupy a special category. They are not the same as other representations. When a photograph shows something, it has a claim on our belief that a painting or a written account does not automatically have. The camera was there. The light was recorded. Something real happened.

Legal systems formalized this agreement in their own way. Courts developed rules of evidence that treated photographs differently from other kinds of documentation, treating them as a form of direct proof rather than hearsay or opinion. Journalism built verification practices on top of the same assumption, practices that asked not just whether an image was compositionally convincing but whether the photographer could be placed at the scene and whether the chain of custody for the image was clean. Science used photography as a tool for establishing facts, not just illustrating them. The indexical nature of photographs was, in all these domains, treated as a given.

This was never philosophically naive, exactly. Anyone who thought carefully about photography understood that a photograph is also an interpretation: the photographer chose where to stand, when to press the shutter, what to include and exclude within the frame. They understood that photographs could be cropped, staged, selectively published, used to tell partial or misleading stories while containing nothing technically false. The history of photography includes propaganda, manipulation and deliberate deception going back almost to the medium's invention. Josef Stalin was erasing political enemies from official photographs in the 1930s. The darkroom was always, in the hands of a skilled operator, a tool for alteration.

But these limitations did not, in practice, undermine the contract. They were understood as abuses of a system whose basic integrity remained intact. A manipulated photograph was still understood as a photograph that had been tampered with, which means the concept of an untampered photograph remained coherent and meaningful. The baseline was still: the camera records what is there. Deceptions were deceptions against that baseline, which is precisely why they were understood as deceptions at all.

The contract depended on that baseline remaining stable. And for a century and a half, it did.

* * *

We are now producing around four trillion photographs a year. The number is almost impossible to hold in the mind, which is partly the point. To put it in some kind of proportion: in 2000, somewhere in the region of eighty billion photographs were taken globally. By 2024, that figure had increased by a factor of fifty. Most of those images are never looked at by anyone other than the person who made them, if they are looked at at all. They exist as data on servers, as thumbnails in camera rolls, as the visual residue of a civilization that has learned to photograph everything as a reflex rather than a decision.

Something has happened to the way we receive photographs inside this deluge, and I think it is at least as important as anything that has happened to the way photographs are made. Photographs no longer arrive as events. They

arrive as weather, as ambient noise, as the default texture of being alive in the early twenty-first century. And somewhere inside that volume, the act of sustained looking has changed. Not because people are stupid or careless, but because the infrastructure we use to distribute and consume images was specifically designed to prevent the kind of attention that photographs, at their most powerful, require.

Looking at a photograph properly is slow work. It requires allowing what the image depicts to land, to make a claim on you, to mean something. The Abu Ghraib photographs required their viewers to sit with the reality of what they showed, to resist the impulse to scroll past, to permit the knowledge that this happened to settle into comprehension and then into moral response. That settling takes time. It requires what I can only call presence on the part of the viewer: a willingness to be arrested by the image, to let the indexical chain from event to camera to screen connect with something in the person looking at it.

Infinite scroll was engineered to prevent exactly that. The logic of a feed is that there is always something else to see, that nothing justifies sustained attention at the cost of the next thing, that the measure of an image's success is the speed with which it can produce an emotional response before being replaced. I notice this in my own behaviour, which is what makes it worth taking seriously rather than simply lamenting. I am a photographer. I have spent a significant portion of my adult life thinking about what photographs are and what they do. I still catch myself scrolling past images that, encountered in any other context, would have stopped me.

The volume does not just numb people who have no particular investment in photography. It numbs people who should know better.

What has been lost is not the photograph itself. The indexed image, the trace of light from a real event, still exists and still carries its ontological weight whether or not anyone is paying attention. What has been lost is the cultural practice of receiving it properly. For a century and a half, there was a rough collective understanding that photographs made a particular kind of claim on us, that they were different from illustrations or paintings or descriptions precisely because they were caused by the world rather than imagined from it. That understanding has not been argued away. It has been made functionally

irrelevant by abundance and by the architecture of the platforms through which most photographs now travel.

* * *

Here is the question I want to hold alongside everything above: what would happen to the Abu Ghraib photographs if they were released today?

The images themselves would not have changed. They would still show American soldiers and Iraqi detainees in the same rooms at the same moments. The indexical chain from those events to those images would be as intact as it was in 2004. Whatever happened in that prison happened, and the photographs that recorded it would still be the photographs that recorded it.

What would have changed is everything around the images. In 2004, the question a viewer was asked to consider was whether the photographs were genuine records of real events. The answer was yes, and the consequences followed from that answer. Today, a significant portion of the population would be asked to consider a different question first: whether the photographs were real at all. Were they generated? Were they fabricated by people opposed to the war, by political opponents of the administration, by foreign intelligence services with an interest in destabilizing American public opinion? Are these actual photographs or are they synthetic imagery designed to look like photographs?

Those questions would not be asked in good faith by everyone who raised them. Some would be asked cynically, as a strategy for neutralizing evidence that powerful people would prefer to see neutralized. But some would be asked in genuine confusion, by people who have absorbed, correctly, that images can no longer be trusted to be what they appear to be, and who have not yet developed reliable tools for knowing when to trust them anyway.

The mechanism that gave those photographs their force in 2004 depended on a shared assumption: that photographs are caused by the things they depict, and that this causal relationship gives them a different kind of claim on our belief than other representations. That assumption is no longer shared. It has not been argued away or philosophically refuted. It has been industrially

undermined by the development of technologies that can produce images with photographic appearance without photographic cause, and it has been psychologically undermined by the sheer volume of imagery through which everyone now moves every day.

The contract has not been formally cancelled. Nobody sent a notice.

There was no announcement, no moment at which everyone agreed that the old understanding no longer applied. Instead it has been quietly eroded from multiple directions at once, in ways that are individually incremental and collectively transformative. The camera is still recording light. The images are still being made. The photographs still look like photographs. But something essential about what they mean, what they can do, what they can prove, has changed without the vocabulary or the institutions that depend on photographs catching up to the change.

That change is what the chapters that follow try to account for. Not photography as a specialist concern, not what artificial intelligence means for working photographers, not the aesthetics of synthetic imagery. What it means for all of us that the last medium we trusted as automatic evidence has stopped working in the way we assumed it did. The Abu Ghraib photographs did something in 2004 that photographs can no longer reliably do. Understanding why they could do it, and why they no longer can, is the question everything that follows tries to answer.

The Convenient Lie We All Agreed To

There is a photograph taken in 1920 that shows Vladimir Lenin addressing a crowd in Sverdlov Square in Moscow. Leon Trotsky stands on the steps of the podium to Lenin's right, clearly visible, his presence as unremarkable as the cobblestones beneath him. By the early 1930s, Stalin's technicians had removed Trotsky from the image entirely. In the altered version, the steps where he stood are empty. The crowd remains. Lenin remains. Trotsky has been made not to have been there.

What is striking about this, viewed from the distance of a century, is how legible the deception has become. We can place the original and the altered photograph side by side and see exactly what was done and why. We know who was removed, when and on whose orders. The manipulation is not just visible in retrospect; it is comprehensible as an act, motivated by a specific political terror and directed at a specific political memory. Trotsky's presence at that podium is not in serious historical dispute. The photograph that records it has survived. The alteration has survived. Both exist, and the existence of both is what allows us to call one true and the other a lie.

This recoverability is worth pausing on, because it reveals something important about what Stalin's technicians were actually doing. They were not creating a synthetic image. They were fighting against one. The original photograph recorded something real: Trotsky was there, the camera was there, light travelled from his body to the lens and left a trace. The alteration was an attempt to unmake that trace, to impose a different reality over the one the camera had recorded. The effort required was substantial. The traces were detectable to a careful eye, and remain detectable today. The manipulation worked against the nature of photography rather than stepping around it.

Stalin's image manipulators were, in a perverse sense, paying tribute to photographic truth. You only need to falsify a record if the record carries weight.

The erasures were acts of respect for the indexical nature of photographs, conducted by people who understood precisely what they were fighting against and who found the fight worth the considerable effort it required. The fact that we can reconstruct what was removed, that the original evidence has survived to contradict the lie, is a consequence of that fight being conducted on photography's own terms. The lie was imposed on top of a truth, and the truth was always there underneath it, waiting to be recovered.

This is the first thing to understand about the history of photographic deception: it was, for most of photography's existence, a history of manipulation rather than fabrication. Someone was always present. Something always happened. The camera always recorded something real. What was subsequently done to the image, the cropping, the retouching, the erasure, the selective publication, operated on an original that existed and could in principle be produced to contradict the alteration. The photograph's indexical nature was the thing being fought against, which means it was never fully defeated.

That changed, within living memory, more completely than most people have registered. But to understand how it changed, and what was lost in the changing, it is first necessary to understand what the photograph's authority actually rested on. The philosophical account will come in due course. The more urgent question is the social one: how did the convenient lie get established? Why did the world agree to treat photographs as automatic evidence, even as the people most involved with photographs understood that they never quite were?

* * *

The authority of photographs was never a philosophical conclusion. No one reasoned their way to it from first principles. It accumulated, over the better part of two centuries, as a practical response to what photographs could demonstrably do that other forms of representation could not.

Courts began accepting photographs as evidence in the latter half of the nineteenth century, with considerable caution and considerable argument about their reliability. Judges who worried about staging, retouching and misrepresentation were raising genuine concerns. Over time, practical necessity, combined with the growing technical sophistication of authentication methods, pushed legal systems toward treating photographs as a special category of evidence: fallible, certainly, but carrying a presumption of accuracy that written testimony or illustrated documentation could only aspire to. By the twentieth century, the photograph as legal evidence was so thoroughly established that its assumptions had become invisible, built into rules of evidence and courtroom procedure in ways that most practitioners had stopped examining.

Journalism built its own version of the same agreement. The photograph as proof became so fundamental to how news was communicated, how atrocities were documented, how the powerful were held accountable, that the infrastructure of photographic verification, the chain of custody, the corroborating eyewitness, the caption that specified who took what where and when, was less a safeguard against deception than a ritual acknowledgement that the photograph was, in principle, the kind of thing that could be verified. The underlying assumption was still that the camera had been somewhere real, had recorded something that existed. Verification checked the details. It did not question the foundation.

Science, medicine, law enforcement, military intelligence, historical scholarship: every discipline that needed to establish facts about the physical world found in photography a tool that seemed to offer something no previous representational medium had managed. A drawing of a fossil could reflect the draughtsman's interpretive choices. A photograph of one was the fossil itself, translated into light. A painting of a crime scene reflected the painter's reconstruction. A photograph was the scene, so far as light and optics could carry it. The operating presumption was that photographic imperfections were knowable and correctable, whereas the imperfections of other representations were structural and irreducible. That was enough.

This presumption was a useful fiction. Useful because it allowed institutions to function, evidence to be weighed, history to be documented,

accountability to be enforced. A fiction because it elided the considerable gap between what a photograph records and what any given reader takes it to mean. The physics were sound. The useful fiction lived in the gaps: between the photograph and the encounter that produced it, between the image and the intent behind it, between what the camera captured and what the caption claimed. In all of those gaps, as a handful of clear-eyed critics were about to demonstrate, something other than simple truth had always been living.

* * *

Roland Barthes published *Camera Lucida* in 1980, the year before he died. It is a short book, personal in register, built around grief for his recently dead mother and organized around a single proposition about what photographs fundamentally are. He called this proposition the noeme of photography, the thing that photography thinks, and he expressed it in three words of French that have resisted clean translation into English: *ça a été*. This has been. The photograph, Barthes argued, is a certificate of presence. Whatever it shows was there. Light travelled from it to the camera. The referent and the image are bound together by a causal chain that no other representational medium can replicate. Looking at a photograph, any photograph, we know at some level that we are looking at something that existed. The photograph carries presence itself, preserved in light, rather than merely representing it.

Barthes was right. This is important to say plainly before complicating it, because the temptation when reading him in light of what has happened to photography since is to position him as naively mistaken, as a thinker who missed what was coming. He didn't miss it. He described the mechanism with extraordinary precision. The noeme he identified is a genuine property of photographs, grounded in physics rather than sentiment. When the Abu Ghraib photographs did what they did in 2004, they did it because Barthes was correct: the indexical chain from those events to those images gave them a claim on belief that no written account could replicate. His philosophy of photography stands as a precise description of something real.

What Barthes did not account for, could not account for writing in 1980, is that the mechanism he described could be socially undermined while remaining physically intact. The causal chain is still there. The physics have not changed. What has changed is the viewer's ability to know, when looking at any given image, whether the chain exists. Barthes built his theory on an assumption so obvious in 1980 that it barely registered as an assumption: that something which looks like a photograph is a photograph. That if the *ça a été* feeling arises when you look at an image, the feeling is reliable. The referent was there. The camera was present. The indexical bond is genuine.

That assumption is no longer safe. The feeling Barthes described, the sense of presence that photographs produce, can now be generated in the viewer by images that have no indexical nature whatsoever, produced by statistical models trained on millions of other images rather than by the passage of light from real objects. The feeling and the fact have been separated. Barthes's theory remains sound; the conditions under which it operated have changed. He mapped the territory with extraordinary care, just before the territory changed. His precision is now the measure of the loss.

* * *

Susan Sontag's *On Photography* appeared three years before *Camera Lucida*, in 1977, and approached the same territory from a different angle. Where Barthes was phenomenological, interested in what photographs feel like and what that feeling reveals about their nature, Sontag was political and moral, interested in what photographs do to the people who look at them and the people they depict. She was suspicious of photography in ways that Barthes was not, arguing that the proliferation of images numbs rather than sensitizes, that the camera is an instrument of acquisition and aggression as much as documentation, that photographs flatten the complexity of events into surfaces that can be consumed without being understood.

Much of this argument holds up well. Her account of how photographic abundance produces compassion fatigue was prescient in 1977 and reads as prophecy now, when four trillion images are made each year and the

infrastructure of their distribution has been specifically engineered to prevent the sustained attention that any single image might deserve. Her scepticism about the photograph as transparent window onto reality was philosophically sound and practically important. The useful fiction needed questioning, and she questioned it seriously.

Sontag made a specific error in *On Photography* worth examining rather than dismissing. Her larger argument survives it. The error itself demonstrates, with uncomfortable precision, the exact failure mode that photographs invite. Her treatment of Diane Arbus is the case in point.

Arbus photographed people who existed at the margins of mid-twentieth-century American social life: people with dwarfism, nudists, identical twins, transvestites, carnival performers, the institutionalized. Her images are formally rigorous and emotionally unsettling. Sontag found them morally troubling in a specific way: she read them as the product of a photographer who had treated her subjects as specimens, who had exploited their vulnerability for aesthetic purposes, whose relationship to the people she photographed was essentially predatory. The photographs, in Sontag's reading, were evidence of that relationship. Looking at them, she felt she understood what had happened in the encounter between Arbus and her subjects, and what she understood disturbed her.

The feeling was real. The conclusion drawn from it went beyond what the photographs could support. Arbus spent extended time with the people she photographed, often months. Many of them sought her out. The encounters were frequently collaborative rather than exploitative, the relationships ongoing rather than transactional. None of this is visible in the photographs. Nor could it be: a photograph records a moment, not the history that preceded it or the relationship that made it possible. Sontag was reading the finished image as evidence of the encounter that produced it, inferring the photographer's intent and moral character from the visual qualities of the work. She was treating the photographs as certificates of presence not just for what they showed, but for what she felt about the showing.

This is the failure mode that photographs uniquely invite, and it predates synthetic imagery by decades. The indexical nature of photographs, the genuine

fact that they are caused by the things they depict, produces in the viewer a feeling of direct access to reality that extends beyond what the photograph can actually deliver. We feel, looking at a photograph, that we are in contact not just with the image but with the encounter, the intent, the meaning, the moral weight of what was recorded. Barthes identified this feeling accurately as the noeme of photography. What he did not emphasize, and what Sontag's Arbus reading demonstrates, is that the feeling can operate on territory the photograph has no authority over. The sense of evidential access spreads beyond the indexical fact into interpretation, inference and moral judgement, and carries the same feeling of certainty all the way.

Sontag was one of the most rigorous critical minds of the twentieth century, writing seriously and at length about the nature and limits of photographic truth. That she fell into this particular trap says something about the power of the trap. If photographs could mislead Sontag about what they were evidence of, while she was actively thinking about the ways photographs mislead, the rest of us are not in a strong position.

* * *

Barthes and Sontag were writing at a specific historical moment, and the limitations of their accounts are partly limitations of that moment. In 1977 and 1980, the most sophisticated photographic deception available was still working within the old model: manipulation of real photographs, erasure of real presences, staging of real scenes. Stalin's technicians. The darkroom's capacity for alteration. The careful selection of which real images to publish and which to suppress. All of it operated on the assumption that there was an original, a real event, a genuine indexical chain, that was being worked against or managed or concealed.

What neither Barthes nor Sontag needed to account for was the possibility that the indexical chain might be bypassed altogether. That something could be produced which carried the visual grammar of photography, the grain and focus and perspectival logic that signals photographic origin, without any causal connection to a real event. They were theorizing the nature

and limits of a medium whose fundamental mechanism, the recording of light from real objects, they could take as given. The philosophical and political questions they were asking presupposed that the photograph was a photograph.

That presupposition has now been removed, unevenly and incompletely, but removed. Real photographs still exist in vast quantities and still carry their indexical weight. The *ça a été* is still a genuine property of genuine photographs. The viewer can no longer know, however, looking at any given image, whether the feeling of photographic presence they experience corresponds to an actual indexical chain or to a simulation of one. The feeling and the fact have been decoupled. What Barthes correctly identified as the noeme of photography has become, in the hands of generative technology, an aesthetic property that can be reproduced without the underlying reality that gave it meaning.

The useful fiction, the social agreement to treat photographs as automatic evidence, was always contestable. Stalin's erasures proved it. Sontag's misreading of Arbus proved it in a different way. The fiction survived those proofs because the contestation still took place on photography's own terms: you could manipulate a photograph, but the manipulation was imposed on an original that existed, detectable in principle if not always in practice, recoverable if the original had survived. The agreement to trust photographs was an agreement to treat them as a special category of evidence, not infallible, but different in kind from other representations. That difference in kind is what is now being eroded.

What happens to the four distinct activities that we still call by the single word photography when the mechanism that distinguished them from other kinds of image-making can no longer be assumed? The categories themselves begin to blur, in ways that are invisible by design and consequential beyond anything that Stalin's technicians, working painstakingly to remove a single figure from a single image, could have imagined.

CHAPTER THREE

Four Kinds of Photography and Why the Distinction Is Collapsing



Last week I was in Wuppertal, in the Ruhr valley, following the Schwebebahn, the suspended monorail that has run above the river Wupper since 1901. It is one of those pieces of infrastructure that makes you wonder why the rest of the world didn't follow the same logic: the train hangs below its track rather than sitting on top of it, which means it cannot derail, and it follows the river's curves through the city without needing any of the land on either bank. I had been photographing it for most of the afternoon from the outside, from bridges and riverbanks, trying to find angles that conveyed its particular combination of industrial weight and unlikely grace.

At some point I boarded it and made my way to the front carriage, where you can stand behind the glass behind the driver and watch the track approach. The evening light was fading and the darkened glass deepened it further. Another train came toward us from the opposite direction, its headlights visible

along the elevated track ahead. I raised my phone and waited. As the trains drew level, I pressed the shutter.

The image that came back was not what I had seen. My phone, operating on default settings I had not fully audited, had decided that the low light and the movement required its computational assistance. It took several frames in rapid succession and merged them algorithmically, a process designed to reduce noise and improve sharpness by combining information from multiple exposures. In the case of a stationary scene, this works tolerably well. In the case of a train moving at speed above a river in failing light, it produced something else entirely: the oncoming train appeared in the image three or four times, slightly offset at each repetition, a smeared ghost of itself moving across the frame.

I knew immediately that the image was wrong. The knowledge was not analytical, not the result of examining the photograph and identifying the technical artefact. I knew it because I had been there. I had seen one train. The image showed a different thing, a thing that had never existed above the Wupper that evening. The record that remained on my phone did not correspond to what I had witnessed, and I was the only person who could know that, because I was the only person present. Any viewer looking at that image without my testimony would have no reliable way of knowing whether they were looking at a computational failure, a deliberate multiple exposure, a synthetic creation, or simply something that had actually looked that way. The photograph had lost the ability to answer that question on its own.

I still have the image. The record of the machine's intervention is preserved. There is a certain irony in that: the photograph that demonstrates how photographs can lie about what the photographer witnessed is itself a piece of photographic evidence, indexed to a moment that really occurred, a moment in which a machine quietly replaced my seeing with its own.

* * *

Photography is no longer one thing. It has not been for a while, but we are still using the same word for activities that are increasingly distinct in their nature, their purpose and their relationship to reality. The confusion this creates

is not merely semantic. When people argue about what photography is, what it can do, what it means for artificial intelligence to threaten it, they are frequently arguing past each other because they are talking about different activities that happen to share a name and a visual grammar.

My afternoon in Wuppertal moved between at least three of those activities without my fully registering the transitions. When I was standing on the riverbank composing a frame, making decisions about where to stand and when to press the shutter, responding to the specific quality of the light and the particular relationship between the hanging train and the water below it, I was doing something that has a reasonably clear definition: a human being was present, was seeing, was choosing. When I stood behind the driver's glass and waited for the oncoming train, I was doing the same thing. When my phone took over the processing of what I had seen and produced something I had not seen, I was doing something else, though the word I would use to describe it in either case is the same. I took a photograph.

The taxonomy I want to propose has four categories, defined by two questions. Who made the image: a human or a machine? And who is the intended audience: a human or a machine? The intersections of those two questions produce four distinct kinds of photography, and understanding what happens when they bleed into each other is, I think, more important than most of the arguments currently being conducted about artificial intelligence and the future of the medium.

* * *

The first category is the one most people mean when they say photography: images made by humans for humans. Someone sees something, decides it is worth capturing, makes choices about how to capture it, and shares the result with other people who look at it and respond. The entire chain involves human intention, human seeing, human interpretation, human taste. The photographer is present and makes choices. The audience engages with those choices. This is what I was doing on the riverbank in Wuppertal when the light was right and I was paying attention.

This category is shrinking proportionally but not disappearing. Something more interesting is happening to it: it is becoming more self-conscious, more intentional, more explicitly defined against the other categories that now surround it. Photographers who work in this way are increasingly aware that they are making a choice to do so, that the default has shifted, that human-for-humans photography now requires a kind of deliberate resistance to the infrastructure that surrounds image-making. The value of the image is migrating from its aesthetic qualities toward the fact of human presence: someone was there, saw this, decided this was worth recording. That decision, which once went without saying, has become the thing worth saying.

The second category has existed longer than most people realize, and has been expanding rapidly for decades. These are images made by machines for machines: captured automatically, processed algorithmically, used to trigger other automated systems or feed databases that machines query. The cameras in a hospital scanning a patient's tissue for early signs of cancer. The satellite imaging systems processing land surface data for patterns that no human will ever examine directly. The traffic cameras feeding number plate recognition systems. The drone surveys of construction sites generating point-cloud data for structural analysis software. None of these images are made for human viewing. They exist as data, as input to computational processes, and human visual aesthetics are entirely irrelevant to their function. What matters is information density, consistency and machine-readable format.

This category is easy to overlook in conversations about photography's future because it operates largely outside the cultural spaces where photography is discussed. Nobody is arguing about the artistic merit of breast cancer screening imagery or the authenticity of satellite data. Yet it matters here precisely because it represents photography stripped entirely of the human witness. The causal chain from event to image is intact. Light from real objects is still being recorded. The indexical nature that Barthes identified is still present. But the purpose of the recording has nothing to do with human testimony, human memory, or human evidence. These images are witnesses to nothing in any human sense of the word.

The third category is the newest and the most obviously disruptive. These are images made by machines for humans: photographs in appearance without photographs in origin. AI-generated imagery that looks photographic but depicts nothing that existed. The computational photography in modern smartphones that composites multiple exposures, applies algorithmic improvements that never existed in the original scene, and adjusts colours based on what the device thinks the scene should look like rather than what it actually looked like. Deepfakes and synthetic media. The ghost trains in my image from Wuppertal.

This category occupies so much of the current argument about photography and artificial intelligence that it has become almost synonymous with the crisis. And the concern is legitimate: if images can look photographic without being caused by real events, the indexical chain that gave photographs their evidential weight is broken. But I want to suggest that the most disruptive aspect of this category is not the images that are entirely synthetic. Those, at least, have a clear status: they are fabrications, and the fabrication is total. The more corrosive problem is the images that are partially synthetic, that sit in the territory between human seeing and machine processing, that are neither fully one thing nor the other. My Wuppertal photograph is not a fabrication in the conventional sense. I was there. The train was there. Something real happened. The image that records it is a collaboration between my act of witnessing and a computational process that altered what was witnessed without my instruction or consent.

The fourth category receives less attention than it deserves. These are images made by humans for machines: photographs taken by people, with all the intention and presence that implies, but whose ultimate destination is computational rather than human. Images uploaded to platforms that use them as training data, photographs sold to AI companies for building generative models, images a photographer shared with a human audience and that were subsequently scraped without permission and fed into systems that will learn from them to generate competing images.

This category inverts the traditional purpose of photography in a way that is genuinely strange to contemplate. The photographer is present. They see

something. They make decisions. They produce an image intended for human viewing. And then, without their knowledge or consent, that image becomes raw material for machines to process into the capability to generate other images, images that require no presence, no seeing, no decision. The photographer trained themselves through years of work and accumulated seeing. Their work then trains the machines. The testimony becomes a template.

* * *

These four categories would be manageable if they stayed distinct. A photographer could know which kind of photography they were doing. A viewer could know which kind of image they were looking at. The institutions that depend on photographs, the legal systems, the journalism, the historical record, could develop appropriate responses to each category. The crisis is not the existence of the categories. The crisis is the bleed between them.



The bleed I find most difficult to accept is the one I experienced in Wuppertal. Human-for-humans photography becoming machine-for-humans photography without acknowledgment. I intended to make a record of what I saw. I brought my presence, my attention, my compositional instincts, my knowledge of the Schwebbahn's history and its particular visual character. I pressed the shutter at the moment that seemed right. All of that is human. And then the machine took over and produced something I had not seen, without flagging that it had done so, presenting the result as if it were a continuous act of human witnessing.

This is not the same as a darkroom manipulation of the kind that Stalin's technicians performed. Those alterations were imposed on an existing indexical image by a human being who knew what they were doing and why. The image beneath the alteration remained, in principle, recoverable. What my phone did was different: it intervened in the image-making process itself, before any human-made image existed to be altered, and produced something that has no single original, no undoctored version, no recoverable truth beneath it. The computational process and the human act of witnessing are fused in the result in ways that cannot be separated. I cannot point to a line in the image and say: this is what I saw, this is what the machine added.

That inability is the condition that most people now photograph in, whether they know it or not. Every smartphone camera in use today operates this way to some degree. The extent varies by device, by setting, by scene, but the principle is universal: the camera is not recording what the photographer sees. It is collaborating with the photographer to produce an optimized version of an approximation of what the photographer saw, shaped by algorithms whose priorities are not the same as the photographer's and whose interventions are largely invisible. The photographer and the machine are co-authors of every image, but only one of them knows the other is there.

The second bleed is slower and less visible. Human-for-humans photography becoming human-for-machines photography without permission. This is the category inversion I described earlier: the image made for a human audience that becomes, through scraping and training, raw material for systems that generate images that require no human presence at all. The photographer

who spent twenty years developing their eye, their specific way of seeing a particular kind of light or a particular kind of human face, finds that those twenty years of accumulated seeing have been harvested and processed into capability that machines can now deploy without the seeing, without the presence, without the years.

What is being stolen in this process is not just the images. It is the witness.

The specific, irreplaceable human fact of having been somewhere and seen something and chosen to record it. That fact is what gave the photographs their value, what made them evidence in the Barthesian sense, what connected them to the world. The machine that learns from them inherits the visual patterns without inheriting the presence. The resulting images look like photographs. They carry the aesthetic grammar of photographic witnessing. They are, precisely, certificates of absence masquerading as certificates of presence.

The third bleed is approaching and its implications are the most consequential of the three. Machine-for-humans photography will soon become indistinguishable from human-for-humans photography for most viewers in most contexts. The point at which a viewer cannot reliably determine, looking at an image, whether a human being was present when it was made is already visible on the horizon. When it arrives, the contract that Chapter 1 described, the social agreement to treat photographs as automatic evidence, will have nothing left to stand on. Real photographs will lose credibility not because they have changed but because they are surrounded by images that cannot be distinguished from them, and uncertainty contaminates everything it touches.

In Wuppertal, standing behind the driver's glass, I knew my image was wrong because I was there. That knowledge, personal and non-transferable, is becoming the only reliable form of photographic truth available. The photograph cannot carry it. The caption cannot verify it. The chain of custody, the documented provenance, the metadata: all of these can be fabricated or lost.

What remains is the testimony of the person who was present, which is to say what remains is interpretation rather than evidence. The photographer as witness is giving way to the photographer as interpreter. The next two chapters

examine what that replacement costs, and the final chapter considers whether the interpreter might, in the end, be a more honest account of what photographers always were.

The Last Wall

Photography did not fall first.

Understanding the order in which media of record have lost their automatic authority matters, because the story is not a sudden rupture but a slow accumulation, a process that has been running for most of a century, and photography's collapse is its completion rather than its beginning. Every medium we have ever relied upon to establish what happened has gone through its own version of this crisis, each one at the point where the cost of deception fell low enough to be practical. We are living through the last stage of something that started long before digital cameras or generative algorithms existed.

The sequence runs roughly as follows. Written text never acquired the kind of automatic evidential authority that later media would briefly enjoy, because the capacity to fabricate text is as old as writing itself. Propaganda, forgery, the rewriting of historical records, the strategic manipulation of official documents: these are as ancient as the need to establish facts about the world. Courts and institutions developed elaborate procedures for authenticating written documents precisely because nobody assumed that words on a page were automatically trustworthy. The scepticism was built in from the beginning.

Recorded sound arrived in the second half of the nineteenth century as something genuinely new: a medium that could capture the human voice and preserve it. By the early twentieth century the phonograph had become a mass medium, and audio recordings were being used as evidence in ways that their creators had not anticipated. A voice recorded was a voice preserved, which seemed to give recordings a kind of indexical weight that written transcripts lacked. You could hear it. You could hear the specific timbre and cadence of a particular person saying particular words at a particular moment. The recording carried something that a written account of the same words could not.

This trust lasted roughly a century. The specific moment at which it publicly cracked in the United States came in November 1973, when

investigators in the Watergate proceedings discovered an eighteen-and-a-half-minute gap in a White House tape recording of a conversation between President Nixon and his chief of staff, three days after the Watergate break-in. The tape did not contain silence. It contained an active buzz, later determined by forensic audio experts to be the product of at least five separate erasure operations performed by hand.

Speech had been there and had been removed. The gap was audible evidence of its own making.

What happened next is instructive. The forensic panel established that the erasures had been deliberately made. They could not recover what had been erased. The conversation that might have been the most direct evidence of presidential involvement in a cover-up was gone, and gone in a way that was provably intentional. The nation was left with the audible remains of an act of destruction and no way of knowing what had been destroyed. The headphone packaging used by the investigators at the time carried the slogan "Hearing is Believing." That slogan had just been made to look naive.

The Watergate tape gap is a useful marker because it represents audio manipulation working within the old model: destruction of an existing recording rather than fabrication of a new one. The underlying trust in audio evidence survived it, because the manipulation was ultimately detectable as manipulation, even if the content could not be recovered. What audio recordings could do as evidence was damaged but not ended. Courts continued to admit them. Investigators continued to rely on them. The assumption that a recording was what it claimed to be remained operationally intact, because the technology required to fabricate a convincing voice recording from scratch did not yet exist in a form accessible to anyone without significant resources.

That assumption is now being removed with the same efficiency and at the same scale as photographic synthesis. Audio deepfakes, voice cloning systems trained on recordings of a target's speech, can produce new sentences spoken in a voice indistinguishable from the original. The technology required is available to anyone with a consumer device and a few minutes of source material. An audio historian quoted by CBC Radio put the situation with

admirable directness: we may end up looking back at the last hundred and fifty years as a blip in time, back when we had figured out how to record sound but not yet how to fake it. A century and a half of audio evidence, from the first phonograph recordings to the Nixon tapes, now sits in a category whose evidential status has been quietly reclassified from reliable to contingent.

* * *

Video came later and fell faster. The moving image arrived at the end of the nineteenth century as an extension of photography's indexical promise: not just a trace of a moment in light but a trace of duration, of movement, of the world unfolding in time. Filmed footage seemed to offer something that even photographs could not: the specific, uninterruptable sequence of events as they actually occurred. A photograph could capture a moment. Film could capture what happened between moments.

The trust this generated was real, and the infrastructure built on top of it was substantial. News footage became the primary vehicle through which people understood the major events of the twentieth century. The Zapruder film of Kennedy's assassination. The footage of the moon landing. The filmed testimony of Nuremberg defendants. The news broadcasts of the Vietnam War that Sontag wrote about, those images arriving nightly into American living rooms and shifting public opinion about the war. Video was the medium of record for the world as it happened.

The film and television industries understood from the beginning that what you see is not necessarily what was recorded: editing, lighting and staging had been tools of the craft since the silent era. But these were understood as tools of storytelling and presentation, operating within a framework where the underlying footage itself was authentic. You could cut, arrange, and contextualize authentic footage in misleading ways. You could not fabricate footage of events that never happened, at least not without expensive production processes that left obvious traces.

Computer-generated imagery began eroding this distinction in cinema during the 1990s, in ways that were visible and understood as fiction. When a

film replaced a location with a digital environment or added a creature that had never existed, audiences knew they were looking at something constructed. The CGI revolution in Hollywood was not a threat to evidential trust because it was transparently applied to acknowledged fiction. Nobody argued that the dinosaurs in Jurassic Park were evidence of anything. The technology was available, but its application was contained within frames that audiences understood to be representational rather than documentary.

The containment failed in 2017. A Reddit user identified only as "deepfakes" published algorithms, built on the Generative Adversarial Network architecture first developed in 2014, that allowed face-swapping in video at a quality that crossed the threshold of casual detectability. The technology had academic precursors going back to 1997, when researchers developed a program capable of reanimating a filmed face to match a different audio track, and had been advancing steadily in research settings throughout the 2010s. The 2017 moment was not a scientific breakthrough so much as a democratization: the capability moved from specialist laboratories to consumer hardware, and the cost of fabricating convincing video footage collapsed from the budget of a film studio to the resources of an individual with a laptop.

The initial applications were largely pornographic, celebrities' faces imposed on other bodies without consent, a pattern that has since become one of the primary vehicles for the technology's harm. But the implications for evidential video were immediately apparent to anyone paying attention. A technology capable of putting one person's face on another person's body in convincing video was a technology capable of putting words in any person's mouth, actions in any person's hands, presence at any event in the world. The five-year period from 2017 to 2022 saw the rapid industrialization of this capability, until fabricated video of public figures saying things they had not said became a routine feature of political disinformation campaigns. The wall that filmed footage had represented, the wall of moving visual evidence as automatic record of what actually happened, had come down.

* * *

Against this background, photography's special status looks less like an inherent property of the medium and more like a temporary advantage that circumstances granted it. It lasted as long as it did not because photographs are more indexical than audio recordings or film footage, but because the specific cost of fabricating convincing photographic imagery remained high enough to prevent industrialized deception until recently. The techniques required, skilled retouching, sophisticated composite construction, convincing lighting and perspective synthesis, demanded either significant human labour or computational resources that were unavailable outside specialist institutions.

The cost fell. It fell in stages: the mass availability of image editing software in the 1990s, the improvement of generative models throughout the 2010s, and then the arrival of diffusion-model image generation in the early 2020s that put high-quality photographic synthesis within reach of anyone who could type a description of what they wanted to see. The threshold that had separated human-made photographs from fabricated images collapsed, and it collapsed in the same way that the thresholds for text fabrication, audio fabrication and video fabrication had collapsed before it: through a combination of technological improvement and cost reduction that moved the capability from specialists to the general population.

What made photography's fall feel different from the falls of the other media was not its speed or its completeness. It was the specific thing that was lost. Written text never carried automatic evidential authority, so its manipulation, though consequential, did not feel like a categorical loss. Audio recording carried some authority, and losing it damaged specific domains of evidence, legal proceedings in particular, but audio evidence had always been supplementary to other forms of proof rather than primary. Filmed footage carried more authority, and its corruption was more consequential, but the film and television industries' long familiarity with the tools of construction meant that audiences had some prior preparation for scepticism about moving images.

Photography had been treated as categorically different from all of these. The indexical argument, the fact that photographs are caused by the things they depict rather than constructed in the imagination, had given them a status in legal systems, in journalism, in historical record and in ordinary human

cognition that no other medium had achieved. People did not merely trust photographs more than they trusted other representations. They trusted photographs in a way that felt qualitatively distinct from trust in other representations. The *ça a été* feeling that Barthes identified was not a stronger version of the feeling produced by a convincing written account or a well-made audio recording. It was a different kind of feeling, grounded in a different kind of relationship between the image and the world.

When that feeling can be produced by images that have no causal connection to any event in the world, the loss is not just of a useful tool for establishing facts. It is the loss of the only medium that ever claimed, with genuine philosophical justification, to be something other than representation. Every other medium that society has used to record what happened has been representational: words, drawings, even audio recordings and film footage are, in the final analysis, humans or machines translating experience into a communicable form. Photography was different because the translation was performed by physics rather than intelligence. Light from real objects made real marks. The marks were caused by the world.

That difference is now gone as a reliable property of images in circulation. Real photographs still exist. The causal chain still runs through physics when a camera records light from actual objects. But a viewer looking at an image in the world has no reliable way of knowing, from the image alone, whether the *ça a été* applies. The feeling of photographic presence has been detached from the fact of photographic presence, and the detachment is invisible by design.

* * *

The sequence matters because it explains why the current moment feels more disorienting than previous collapses of evidential trust. What is happening to photography is not an isolated crisis of a single medium. It is the completion of a process that has been running through every medium of record in succession. Text, audio, video, photography: each one gained a degree of automatic evidential authority as its technology matured, and each one lost that

authority when the cost of deception fell low enough to make industrialized fabrication practical. We are reaching the end of the sequence because photography was the last medium on the list.

The sequence also explains a psychological phenomenon that the technology commentary tends to underestimate: the accumulation of scepticism across media. Each previous collapse of evidential trust did not merely undermine the specific medium involved. It trained a generation of readers and viewers in a generalized wariness about all media of record. Someone who came of age during the deepfake video crisis already approached moving images with a degree of suspicion that would have been unusual in someone who grew up watching news footage in the 1960s. That trained scepticism did not remain contained to the medium that produced it. It spread.

The result is an epistemic environment in which the question of whether any given record of any event is authentic has become a standing question rather than a question raised only by exceptional circumstances. Most people most of the time cannot answer it, because the tools for verification are not available to ordinary readers and viewers, and even where they exist, the effort required to apply them to every piece of content encountered in the course of a day is prohibitive. The practical response, which is becoming the default response, is to apply a prior scepticism to all media of record and reserve belief for sources whose trustworthiness has been established through relationships rather than through the properties of the medium itself.

This is what the collapse of photography's special status actually means in the lives of people who have never thought carefully about indexicality or the philosophy of representation. It means that the last category of evidence they had been trained to treat as automatically authoritative has joined the others. It means that the question they have been avoiding, the question of how to know what is real when no medium can be trusted on its own terms, can no longer be deferred.

Photography was the wall. The wall is down. What comes next is the subject of the two chapters that follow.

What Collapse Actually Costs

On 16 March 2022, three weeks into Russia's full-scale invasion of Ukraine, a video appeared on the hacked website of the Ukrainian television channel Ukraine 24 and spread rapidly across social media platforms. It showed President Volodymyr Zelensky apparently addressing the nation, telling his soldiers to lay down their arms and surrender. The video was about a minute long. It was a deepfake.

The fabrication was not particularly sophisticated. The voice quality was wrong, the head movements unconvincing, the audio out of register with the image. It was identified as fake almost immediately by Ukrainian users and international researchers, removed by Facebook and YouTube within hours, and personally rebutted by Zelensky from his Telegram channel before the day was out. One version accumulated around 120,000 views on Twitter before being taken down. The Ukrainian government had, in fact, warned citizens on 2 March that Russian forces might deploy deepfake surrender videos, and had prepared a public response in advance. By most measures, the operation failed.

What was noted at the time by Sam Gregory of the human rights organization Witness, and what deserves more weight than it received, was the secondary effect. The deepfake did not need to convince anyone to cause harm. Its existence was sufficient to ensure that some people, encountering genuine Zelensky footage in the future, would wonder whether what they were watching was real. The video was a weapon aimed at the general condition of knowing rather than at any specific belief. Its target was the reliability of the medium through which all messages would subsequently travel.

The Zelensky case was a fabrication deployed against a target who could fight back, in a media environment with sufficient independent scrutiny to debunk it quickly, in a country whose government had thought to prepare citizens in advance. The deepfake failed under those conditions. What the case illustrates is not how effective such fabrications can be at their most powerful,

but how comprehensively the intention behind them has shifted. The goal was not to deceive everyone. It was to introduce enough uncertainty into the information environment that certainty became, for some people, permanently unavailable.

What those functions are, and what it costs to lose them, runs through specific institutions, specific legal proceedings, specific acts of documentation, and the accumulation of ordinary human record. The landing is hard.

* * *

Courts have depended on photographs and video as evidence for more than a century, and the rules governing that dependence were built on an assumption that is now visibly fraying. The assumption is simple: that an image or recording either is what it claims to be or has been altered, and that the difference between these two states is, in principle, determinable by sufficiently careful examination.

That assumption rests on a distinction already established: the distinction between manipulation, which operates on an existing authentic record, and fabrication, which produces something that has no authentic original beneath it. A manipulated photograph can be compared against an original. A fabricated one has no original to compare against. The legal system's evidentiary procedures were built for the first problem. They are not equipped for the second.

The cases now appearing in courts around the world illustrate two distinct failures, which are worth distinguishing because they require different responses. The first is the submission of fabricated evidence as genuine. In September 2024, a California judge, Victoria Kolakowski of the Alameda County Superior Court, dismissed a housing dispute case after identifying a video submitted by the plaintiffs as AI-generated. The supposed witness in the video, whose voice was disjointed and whose facial expressions repeated at regular intervals, was a deepfake of a real person who had appeared in other, authentic footage in the same case. The case, *Mendonces v. Cushman & Wakefield*, appears to be among the first documented instances of a suspected deepfake being

submitted as authentic evidence and detected. Judge Kolakowski dismissed the case on the grounds of the plaintiffs' use of fabricated material. The plaintiffs sought reconsideration, arguing she had suspected but not proven the evidence was AI-generated. The reconsideration was denied.

The second failure is subtler and in some ways more corrosive. It involves not the submission of fake evidence but the challenge of real evidence on grounds that it might be fake. Legal scholars Bobby Chesney and Danielle Citron gave this phenomenon a name in 2019: the liar's dividend. The concept is this: once it is widely understood that realistic fabrication is possible, any piece of genuine evidence can be challenged simply by raising the possibility that it might be synthetic. The challenger does not need to prove fabrication. The mere possibility, plausibly raised, is enough to introduce reasonable doubt.

The liar's dividend has already appeared in American courtrooms in ways that should disturb anyone who depends on evidence to function. During the trial of Kyle Rittenhouse in 2021, the prosecution sought to zoom in on a video that had already been admitted into evidence. The defence objected, arguing that the Apple pinch-to-zoom function uses artificial intelligence to process images, and that using it might alter the underlying footage. The court required the prosecution to produce expert testimony showing the function would not alter the video. Having no expert prepared, the prosecution was not permitted to zoom in on admitted evidence. The challenge was not substantiated. There was no actual evidence that pinch-to-zoom had manipulated anything. The possibility was enough.

During the trial of a defendant charged in connection with the January 6th Capitol riot, a prosecutor presented video footage of the defendant at the scene. On cross-examination, defence counsel questioned an FBI agent about whether the video might have been AI-manipulated. The court allowed the questioning even though, when pressed, counsel could not provide any basis for the theory. The video was genuine. The question was permitted anyway. Judges are increasingly in the position of deciding, without adequate training or tools, whether specific evidence is what it claims to be, while lawyers are discovering that raising the spectre of AI manipulation is a strategy available regardless of whether the evidence has actually been manipulated.

The US federal courts are now actively considering amendments to the Federal Rules of Evidence to address these problems. At a meeting of the Advisory Committee on Evidence Rules in November 2024, a proposed new rule was considered that would shift the burden of proof when a party successfully argues that evidence might have been AI-altered: the proponent of the challenged evidence would then be required to demonstrate that its probative value outweighs the risk of prejudice. The committee took no action. Some members argued the existing rules were adequate. Others suggested more judicial experience was needed before codifying an approach. The process of adaptation is underway, and it is slower than the technology it is trying to address.

* * *

Outside courtrooms, the stakes are different but not smaller. Journalism has always depended on evidence, and photographic evidence in particular, to establish that what reporters describe actually occurred. The photograph was the thing that moved an account from assertion to record: not just that a journalist said something happened, but that a camera was present and recorded it. The weakening of that function is not hypothetical. It is happening in the places where it matters most.

In April 2024, a deepfake audio recording circulated in Baltimore depicting Eric Eiswert, the principal of Pikesville High School, apparently making racist and antisemitic remarks about students and colleagues. The recording spread rapidly. Eiswert received death threats. He was put on administrative leave while the recording was investigated. It was eventually identified as a deepfake, apparently created by a school athletic director who had a grievance against Eiswert. The athletic director was subsequently charged. Eiswert's career survived, though not without damage. He had done nothing. The recording had been entirely fabricated. The harm it caused, and the months it took to establish that the recording was fake, are both measures of how the infrastructure of accountability can be inverted: a tool designed to document

wrongdoing was used to fabricate it, and the fabrication was sufficient to trigger the responses that genuine wrongdoing would have triggered.

The Eiswert case involved an individual and a local institution. The same mechanism, operating at national and international scale, has consequences that are harder to contain and harder to recover from. Following Myanmar's military coup in February 2021, the ruling junta posted a video purportedly showing a former government minister confessing to having bribed Aung San Suu Kyi. The video attracted immediate scrutiny: audiences and analysts noted visual distortions and odd vocal quality. Questions about whether it was a deepfake spread across social media. In the context of a military coup, where the junta had strong reasons to fabricate evidence against the detained civilian leader it had overthrown, the scepticism was understandable. Meanwhile, the junta and its supporters were simultaneously casting doubt on genuine footage documenting human rights abuses committed by the military, claiming those recordings were fabricated.

The situation in Myanmar illustrates a specific feature of the liar's dividend that is worth holding separately from the courtroom cases. In a domestic legal context, the dividend operates within a framework of rules, procedures and institutional authority that, however imperfect, provides some check on its abuse. A judge can exclude a baseless challenge. An appellate court can review a ruling. The rules of evidence create a structure within which the exploitation of synthetic uncertainty can be managed, even if not yet adequately.

In a conflict zone, or in any context where the institutions capable of adjudicating competing claims about evidence are themselves parties to the conflict, no such framework exists. A military regime that has committed atrocities and simultaneously developed the capacity to raise doubts about documentation of those atrocities has acquired something of considerable value: the ability to make accountability structurally difficult regardless of the quantity or quality of the evidence against it. The doubts do not need to be convincing. They need only be sufficient to slow the mechanisms of international response, to give potential allies political cover for inaction, to ensure that the question of what actually happened remains formally open long enough for consequences to

become practically impossible. The liar's dividend, at this scale, is not an evidentiary problem. It is a political one.

The Gabon case of 2019 is the purest illustration of this dynamic, and it is useful precisely because the question of what the truth was remains open. President Ali Bongo Ondimba, absent from public life after suffering a stroke, appeared in a video address to mark the new year. His opponents claimed it was a deepfake, that the government was concealing his incapacitation behind a fabrication. Independent analysts found no technical evidence of manipulation. The video was, in all probability, genuine. But the belief that it might not be was sufficient: five days after it appeared, a group of military officers announced a coup attempt, citing the video as evidence of the government's deception. The coup failed. Bongo subsequently made other public appearances. What the episode demonstrated, regardless of the truth, is that the mere existence of deepfake technology had made a real video of a real president unable to function as evidence of his own reality. The doubt was available, and once available, it was used.

* * *

There is a version of the argument being made in some quarters that the current crisis of photographic and documentary evidence is simply a new phase of an old problem: information has always been contested, powerful actors have always tried to manage evidence in their favour, and democratic societies have always had to develop the tools to maintain accountability in the face of those efforts. On this view, the current situation is serious but not qualitatively different from what came before. Better technology for verification, better media literacy, better institutional procedures: these are the responses required, and they are the responses being developed.

This view underestimates one thing. Every previous crisis of evidential trust involved a contest over specific pieces of evidence: this photograph was altered, that recording was fabricated, this document was forged. The contest was specific because fabrication was expensive and effortful, which meant it was deployed selectively and its products were individually identifiable as targets for

investigation. The adversarial relationship between fabrication and verification was, in practice, a relationship between specific acts of deception and specific acts of detection.

What the industrialization of synthetic media has done is change the terms of that contest in a fundamental way. When fabrication is cheap and scalable, the relationship between fabrication and verification is no longer adversarial in the same sense. It is asymmetric. Fabrication can now be produced faster than verification can be applied, distributed faster than corrections can follow, and deployed at a volume that overwhelms the institutional capacity to respond. The question shifts from whether specific evidence is genuine to whether any evidence can be treated as presumptively genuine. The answer, increasingly, is no. That shift in the default assumption is what makes the current situation qualitatively different from previous crises of evidential trust, and what makes the responses adequate to those previous crises insufficient for this one.

The cost is measured not only in specific cases where fake evidence is submitted or genuine evidence is dismissed. It is measured in the cumulative effect on the epistemic environment in which all evidence operates. Every case in which a deepfake is submitted and not detected contributes to the baseline uncertainty. Every case in which genuine evidence is successfully challenged as potentially synthetic, whether the challenge succeeds or not, trains lawyers and judges and juries to treat documentary evidence with a scepticism that was, until recently, reserved for contested or suspicious material. The uncertainty spreads not only through the specific cases in which it is raised but through the general atmosphere in which evidence functions. That general atmosphere is changing, and the institutions that depend on evidence are changing more slowly than the atmosphere they operate in.

* * *

There is one further dimension of this that the legal and political cases do not fully capture: the loss to historical record. The cases discussed above involve living people and current events, where at least the possibility of verification

exists and where the institutions and individuals harmed have some recourse. But evidence also serves a function that extends beyond the resolution of present disputes. It is the raw material from which future understanding is constructed.

I have spent time in my own photography thinking about this in personal terms. I have a family archive that goes back several generations: photographs of my great-grandparents in their twenties, faces of people whose names are now uncertain even to those of us who carry their genetic material, lives that left almost no record beyond two dates and a faded image. These photographs carry their indexical weight silently. The people depicted were there. The light from their faces made these marks. Whatever else has been lost, that physical fact remains. The images are evidence of presence in a sense that no written account could replicate.

The archive that future generations will inherit from ours will be different in kind. The photographs taken with smartphones today, half of them processed by algorithms whose interventions are invisible, some fraction of them generated synthetically and distributed without markers, accumulated in volumes that make individual verification impossible: this archive will not carry the same kind of evidential weight. Not because the genuine photographs within it are less genuine, but because the mixture will be unresolvable. The provenance of specific images will be uncertain. The question of whether a given face in a given place at a given moment was actually there will be, for most images, unanswerable.

For the history of public events, where documentation is extensive and professional, this is already driving efforts to develop technical standards for provenance tracking, metadata authentication and watermarking that can establish chains of custody for important images. These efforts are serious and deserve support. For the history of private lives, where no such infrastructure exists and where the archive is a phone camera roll rather than a news organization's server, the problem is effectively unaddressed. The ordinary human record, the visual evidence of how people actually lived in the early twenty-first century, is already compromised in ways that will only become

visible to future historians who try to understand what they are looking at and find the question unanswerable.

This is what the collapse of photographic trust costs at the scale of ordinary life, rather than at the scale of legal proceedings or political accountability. It costs the reliability of the record. The evidence of having been somewhere, seen something, lived through a moment: that evidence now requires more than a photograph to establish. It requires the kind of contextual verification and personal testimony that photographs were supposed to make unnecessary. The photographer as witness, the person who was there and whose presence the camera recorded, has to step forward and say: I was there, and this is what I saw. The photograph alone no longer says it for them.

The photograph alone no longer says it for them.

CHAPTER SIX

The Interpreter and the Chain of Custody



A few years before my grandfather died, I took a photograph of him in his workshop. The room was exactly what a lifetime of working with your hands produces: a pegboard wall covered with tools arranged with the unconscious order of someone who has reached for the same things in the same places for decades, a workbench worn smooth at the edges, a large blue vice bolted to the bench with the permanence of something that has never needed to move. He

was in his nineties by then. The hands that had built things were still his hands, but the body they belonged to had changed.

I set up the lighting and I asked him to stand at the bench. I didn't tell him what to do. I asked him to do something he used to do, and left the choice to him. He picked up a coil of metal wire and rested his other hand on the vice and looked down at the work in front of him. What I had wanted was a record of who he had been, not a document of who he was that afternoon. What he gave me was his own answer to that question: the specific things his hands reached for, the specific posture his body took, drawn from somewhere deeper than instruction. He was too old by then to be doing what the pose suggested, and we both knew it. The expression he wore was partly performance and partly something else, something that came from the fact that his hands had been in that position thousands of times before and the muscle memory of it was still there even when the capacity for the work was not.

The photograph that resulted is indexed to reality in every technical sense. He was there. The light was there. The wire and the vice and the tools on the wall were there. The camera recorded light from all of it. Barthes's *ça a été* operates without qualification: this happened, the image exists because it happened, the physical fact of his presence in that workshop on that afternoon is preserved in the image with the same ontological force as any photograph ever made. And yet the image is also interpreted memory, a collaboration between two kinds of knowing: my understanding of who he had been, expressed as a question, and his understanding of who he had been, expressed through his choice of what to hold and how to stand. The meaning the photograph carries was built from both. It shows something that was no longer literally true at the moment the shutter opened, and it is more true for that.

I would not have anything showing who he really was without it. That is the photograph's strength, and it comes directly from the interpretation rather than from any failure of it. A candid snapshot of an old man in a tracksuit leaning on a bench would have been more honest in one narrow sense and far less true in every sense that matters. The interpretation was not mine alone. He participated in it, answered the question in his own terms, brought his own memory to the act. The photograph records that collaboration.

* * *

Photography was never only witnessing. It was always also this: the photographer's knowledge, memory and understanding brought to bear on whatever the camera was pointed at. The witness account of photography, the idea that the photographer simply records what is there and lets the camera carry the meaning, was always a partial description at best and a convenient fiction at worst. Every decision a photographer makes, where to stand, when to press the shutter, what to include and exclude, how to light the subject, what to ask the subject to do, is an interpretive act. The camera records the result of those acts. It does not make them.

What has changed is not the reality of interpretation but the visibility of the fiction that surrounded it. The automatic authority that photographs carried, the social agreement that they occupied a special evidential category, provided cover for that fiction. If photographs were simply witnesses, then the photographer's interpretive choices were invisible inside the image's claim to truth. The photograph said: this is what was there. The photographer's role in constructing what was there, in choosing the moment and the frame and the light and the subject's position, was absorbed into that claim rather than acknowledged as constitutive of it.

The collapse of automatic photographic authority has removed that cover. It has done so destructively, through the industrialization of fabrication and the consequent contamination of the evidential category that photographs occupied. But the removal also reveals something that serious photographers have always known and that the witness fiction obscured: that the value of a photograph is inseparable from the intelligence, knowledge and judgement of the person who made it. The image of my grandfather at his bench is worth something because I knew him, because I understood what needed to be preserved, because I made deliberate choices about how to make that understanding visible. The camera was the instrument of those choices. It was not their source.

Calling this interpretation rather than witnessing is not a demotion. It is a more accurate description of what photography has always done at its best.

The witnesses were always also interpreters. The collapse has simply made the acknowledgement unavoidable.

* * *

Barthes described the *ça a été* as the thing that photography thinks, the irreducible property that sets photographs apart from all other representations. He was right about the property. What he did not need to consider, writing in 1980, was whether the viewer could know that the property applied to any given image they encountered. The *ça a été* is still real when a photograph is genuine. My grandfather's hands on that vice carried light to the camera and left a trace. That physical fact is unchanged by anything that has happened to photographic culture since.

What has changed is the viewer's relationship to that fact. The automatic trust that allowed viewers to receive photographic presence passively, to feel the *ça a été* and take it as given that the feeling was reliable, is no longer available as a default. The viewer now has to choose to extend trust, to decide, on the basis of something other than the image itself, whether the indexical chain they are being invited to feel is genuine. That choice, which used to be made unconsciously in almost every photographic encounter, has become explicit. It requires information that the photograph cannot supply about itself.

This is where the chain of custody becomes the operative concept. Not the technical chain of custody that journalism and legal systems have developed for specific high-value images, though those matter and deserve the investment being made in them. The chain of custody that matters for ordinary photographic encounters is relational: it runs through the identity and traceable record of the person who made the image. Trust in a photograph, in the post-automatic world, is trust in a photographer. That trust is earned through consistency, through transparency about method and perspective, through the accumulated record of a practice that can be examined and evaluated. It is the same kind of trust that has always governed belief in human testimony, and it is the kind of trust that photography's automatic authority had, for a century and a half, allowed both photographers and viewers to bypass.

The bypassing was convenient but it was also impoverishing. Automatic authority required nothing of the viewer beyond recognition that the image looked like a photograph. It required nothing of the photographer beyond the technical capacity to make one. The relational trust that replaces it is more demanding in both directions. The photographer who wants to be believed has to be knowable, has to provide the context that allows viewers to evaluate the interpretation rather than simply receive it. The viewer who wants to understand what a photograph means has to be willing to do the work of situating the image within the knowledge and perspective of the person who made it. Neither of these demands is new to serious photography or serious looking. What is new is that they apply to the entire category rather than just to the work that announces itself as interpretive.

* * *

There is a version of this argument that leads somewhere bleak: if all trust is relational and personal, then we have simply returned to a pre-photographic world in which everyone believes what they choose to believe and evidence is whatever powerful people say it is. The chain of custody dissolves into the chain of authority, and those with the most authority get to decide what the evidence shows. On this reading, the collapse of automatic photographic trust is not a clarification but a regression, the removal of a check on power that, imperfect as it was, served an important function.

The reading is not entirely wrong. The liar's dividend, the ability of powerful actors to dismiss genuine documentation by raising the spectre of fabrication, is real and consequential, as the Myanmar and Gabon cases demonstrate. The loss of automatic photographic authority is being exploited, and will continue to be exploited, by people whose interests are served by the ability to make accountability structurally difficult. That cost is serious.

But the pessimistic reading misses something. The automatic authority of photographs was not a neutral check on power. It was a specific social agreement that benefited some parties more than others, that was more easily mobilized by institutions than by individuals, and that could itself be

weaponized through the selective deployment of genuine images in misleading contexts. The witness fiction did not only constrain the powerful. It also constrained the photographed, who had no recourse against the authority of images made without their consent or against the interpretations placed on them. The Abu Ghraib photographs did what they did partly because of automatic photographic authority. So did every photograph ever used to criminalize, stigmatize, or misrepresent the people it depicted.

A world in which photographic trust is relational and earned rather than automatic is a world in which those questions have to be asked explicitly. Who made this? What did they know? What did they want? What is their record? Those questions are not a retreat from evidence. They are the questions that should always have accompanied evidence and that the convenience of automatic authority allowed us to skip. The interpreter, who acknowledges their perspective and their choices and provides the context that allows their work to be evaluated, is a more honest figure than the witness who claimed to simply show what was there.

* * *

What does this mean in practice, for people who make photographs and for people who look at them?

For photographers, it means that the work of interpretation has to be made visible rather than hidden inside the claim to witnessing. The choices, what was included and excluded, why this moment rather than another, what was known about the subject that shaped the framing, become part of what the image offers rather than noise that contaminates its objectivity. The photograph of my grandfather at his bench is more useful, not less, when accompanied by the knowledge that I staged it, that the scene was constructed from memory rather than observed in the moment. That knowledge does not undermine the image. It completes it. Without it, the viewer sees an old man at a workbench. With it, they see interpreted memory, a grandson's understanding of a grandfather made visible through a deliberate photographic act. The second thing is richer than the first.

This does not require every photograph to be accompanied by an essay. It requires photographers to think of themselves as interpreters whose perspective is part of what they are offering, and to provide enough context, in captions, in the framing of series, in the visible record of their practice, for viewers to situate the work. The specific form that context takes varies with the work. A photojournalist documenting an atrocity provides chain of custody through professional affiliation, verifiable location data, corroborating sources and the accumulated credibility of an organization whose practices can be examined. A personal photographer like me provides it through the accumulated transparency of a practice, through writing about what I do and why, through being knowable enough that viewers can evaluate what they are seeing.

For viewers, it means developing the habit of asking about the photographer rather than simply receiving the photograph. Who made this? What is their relationship to what they are showing? What context are they providing and what are they withholding? These are not questions of suspicion. They are questions of engagement. A viewer who asks them is treating the photograph as what it is: an interpretive act by a specific person with a specific perspective, offering something that requires their own intelligence and judgement to complete. The passive receipt of automatic evidence has always been the lesser form of photographic engagement. The more demanding form was always available. It is now the only form available.

* * *

When I look at the photograph of my grandfather, I see something a viewer who did not know him cannot see: the specific quality of his concentration, the way his hands carry the authority of decades of work even in a pose that is no longer authentic, the particular relationship between the man and the tools that surrounded him in that space. I see interpreted memory. The photograph is the record of my interpretation, preserved in light.

That is not less than witnessing. It is photography doing what photography has always done at its most valuable: not recording the surface of the world but making visible what someone understood about it. The

understanding requires a person. The person requires a history and a perspective and a relationship to what they are photographing. All of that is now, necessarily and honestly, part of what the viewer engages with.

The post-evidence world is not a world without evidence. It is a world in which evidence requires the same effort and the same relational intelligence that all other forms of human testimony have always required. Photography's century and a half of automatic authority was the exception, not the rule. The exception is over. What remains is the work, which was always the more important thing, and the responsibility that serious photography has always carried and that the convenience of automatic trust had, for too long, allowed to be forgotten.

The girl on the road in 1972 is still running. The photograph that recorded her is still indexed to the reality of that road, that moment, that child's pain. The causal chain from that event to that image has not been broken by anything that has happened since. What has changed is the world that receives it. That world can no longer extend its trust to the image automatically. It has to extend it knowingly, through an understanding of who was there, why they were there, what they were trying to show, and whether the record of their practice justifies the belief that what they are offering is worth the trust it asks for.

Photographers who can answer those questions honestly have nothing to fear from the end of automatic photographic authority. The witness was always an interpreter. The collapse has simply made the acknowledgement unavoidable. Learning to make it clearly, and to receive it thoughtfully, is what photography requires now.