

AN INCREDIBLY LUCID TREK INTO THE FORMIDABLE CRAFT OF FILM EDITING, WHICH IN MY OPINION IS THE CORE OF THE CINEMATIC ART FORM. WALTER MURCH'S DEPTH OF INSIGHT INTO THIS SUBJECT IS ASTONISHING, AND THE BOOK IS A MUST FOR ANYONE WHO IS INTERESTED IN TRULY UNDERSTANDING THE FILMMAKING PROCESS.

—GEORGE LUCAS

THE FIRST TIME I SAW MURCH HE WAS ASLEEP ON HIS KEM, HAVING WORKED ALL NIGHT ON *THE CONVERSATION*. HE'S THE ONLY EDITOR I KNOW WHO ACTUALLY SLEEPS WITH HIS FILMS. IN FACT, MURCH HAS A LOVE AFFAIR WITH EVERY FILM HE'S INVOLVED WITH. *IN THE BLINK OF AN EYE* IS HIS CHRONICLE OF PASSIONS—IT'S WHERE HE FINALLY TELLS US SOME OF THE INTIMATE DETAILS OF THESE AFFAIRS. NOT FOR THOSE WHO CANNOT LOVE FILM.

—PHILIP KAUFMAN

THIS SMALL BOOK CONTAINS A WEALTH OF FIRST-HAND KNOWLEDGE ABOUT THE MYSTERIES OF GIVING BIRTH TO A FILM. I THINK IT IS OF IMMENSE VALUE TO PROFESSIONALS AND FILM BUFFS ALIKE.


—FRED ZINNEBANN

SILMAN-JAMES PRESS

IN THE BLINK OF AN EYE is celebrated editor Walter Murch's vivid, multifaceted, thought-provoking essay of film editing. Starting with what might seem to be the most basic editing question—Why do cuts work?—he treats the reader to a marvelous “ride” through the esthetics and practical concerns of cutting film. Along the way, he offers his insights on such subjects as continuity and discontinuity in editing, dreaming, and real life; the criteria of a good cut; and the blink of the eye as both an analog to and an emotional cue for the cut. New to this second edition is Murch's lengthy meditation on the current state of digital editing.

WALTER MURCH is an editor, sound designer, director, and screenwriter who has been honored repeatedly by both the American and British Motion Picture Academies for his picture editing and sound mixing. Among the many notable films he has edited and/or mixed are *The Conversation*, *American Graffiti*, *Julia*, *Apocalypse Now*, *The Godfather Part II* and *Part III*, *The Unbearable Lightness of Being*, *House of Cards*, *Ghost*, *Crumb*, *The English Patient*, and *The Talented Mr. Ripley*.

FILMMAKING/\$13.95
ISBN: 1-879505-62-2



Many of the thoughts that follow, although presented to the public in a lecture, are therefore more truly cautionary notes to myself, working methods I have developed for coping with my own particular volcanoes and glaciers. As such, they are insights into one person's search for balance, and are perhaps interesting to others more for the glimpses of the search itself than for the specific methods that search has produced.

I would like to thank Ken Sallows for providing me with the transcription of the original lecture and the opportunity to present it to a wider audience. For cosmetic reasons, I have made certain revisions and added some footnotes to what was, for the most part, an extemporaneous dialogue between myself and the audience, whom I thank for their interest and participation. I have also updated some technical points and added an afterword that considers the impact that non-linear, digital editing has had on the process of filmmaking.

Special thanks also to Hilary Furlong (then of the Australian Film Commission), who was instrumental in bringing me to Australia, where the lecture was originally given.

Walter Murch
Rome, August 1995

Cuts and Shadow Cuts

It is frequently at the edges of things that we learn most about the middle: ice and steam can reveal more about the nature of water than water alone ever could. While it is true that any film worth making is going to be unique, and the conditions under which films are made are so variable that it is misleading to speak about what is "normal," *Apocalypse Now*, by almost any criteria—schedule, budget, artistic ambition, technical innovation—qualifies as the cinematic equivalent of ice and steam. Just considering the length of time it took to complete the film (I was editing picture for one year and spent another year preparing and mixing the sound), it turned out to be the longest post-production of any picture I have worked on, but that may consequently spill some light on what "normal" is, or might be.¹

One of the reasons for that length was simply the amount of film that had been printed: 1,250,000 feet,

¹ And I had come on relatively late in the process. Richie Marks and Jerry Greenberg had already been editing for nine months when I joined them in August 1977, a few months after the end of shooting, and the three of us worked together until Jerry left in the spring of 1978. Richie and I then continued together, joined by Lisa Fruchtmann, until I began to work on the soundtrack.

which works out to be just over 230 hours. Since the finished film runs just under two hours and twenty-five minutes in length, that gives a ratio of ninety-five to one. That is to say, ninety-five "unseen" minutes for every minute that found its way into the finished product. By comparison, the average ratio for theatrical features is around twenty to one.

Traveling across that ninety-five-to-one landscape was a little like forging through a thick forest, bursting upon open grassland for a while, then plunging into a forest again because there were areas, such as the helicopter sequences, where the coverage was extremely high, and other scenes where the coverage was correspondingly low. I think the Colonel Kilgore scenes alone were over 220,000 feet—and since that represents twenty-five minutes of film in the finished product, the ratio there was around one hundred to one. But many of the connecting scenes had only a master shot: Francis had used so much film and time on the big events that he compensated with minimal coverage on some of these linking scenes.

Take one of the big scenes as an example: The helicopter attack on "Charlie's Point," where Wagner's *Ride of the Valkyries* is played, was staged as an actual event and consequently filmed as a documentary rather than a series of specially composed shots. It was choreography on a vast scale of men, machines, cameras, and landscape—like some kind of diabolical toy that you could wind up and then let go. Once Francis said, "Action," the filming resembled actual combat: Eight cameras turning simultaneously (some

on the ground and some in helicopters) each loaded with a thousand-foot (eleven-minute) roll of film.

At the end of one of these shots, unless there had been an obvious problem, the camera positions were changed and the whole thing was repeated. Then repeated again, and then again. They kept on going until, I guess, they felt that they had enough material, each take generating something like 8,000 feet (an hour and a half). No single take was the same as any other—very much like documentary coverage.

Anyway, at the end of it all, when the film was safely in the theaters, I sat down and figured out the total number of days that we (the editors) had worked, divided that number by the number of cuts that were in the finished product, and came up with the rate of cuts per editor per day—which turned out to be . . . 1.47!

Meaning that, if we had somehow known *exactly* where we were going at the beginning, we would have arrived there in the same number of months if each of us had made just under one-and-a-half splices per day. In other words, if I had sat down at my bench in the morning, made one cut, thought about the next cut, and gone home, then come in the next day, made the cut I thought about the day before, made another cut, and gone home, it would have taken me the same year it actually took to edit my sections of the film.

Since it takes under ten seconds to make one-and-a-half splices, the admittedly special case of *Apocalypse Now* serves to throw into exaggerated relief the fact that editing—even on a "normal" film²—is not so

²By comparison, an average theatrical feature might have a cuts-per-day figure of eight.

much a *putting together* as it is a *discovery of a path*, and that the overwhelming majority of an editor's time is not spent actually splicing film. The more film there is to work with, of course, the greater the number of pathways that can be considered, and the possibilities compound upon each other and consequently demand more time for evaluation. This is true for any film with a high shooting ratio, but in the particular case of *Apocalypse* the effect was magnified by a sensitive subject matter and a daring and unusual structure, technical innovations at every level, and the obligation felt by all concerned to do the very best work they were capable of. And perhaps most of all by the fact that this was, for Francis, a personal film, despite the large budget and the vast canvas of the subject. Regrettably few films combine such qualities and aspirations.

For every splice in the finished film there were probably fifteen "shadow" splices—splices made, considered, and then undone or lifted from the film. But even allowing for that, the remaining eleven hours and fifty-eight minutes of each working day were spent in activities that, in their various ways, served to clear and illuminate the path ahead of us: screenings, discussions, rewinding, re-screenings, meetings, scheduling, filing trims, note-taking, bookkeeping, and lots of plain deliberative thought. A vast amount of preparation, really, to arrive at the innocuously brief moment of decisive action: the cut—the moment of transition from one shot to the next—something that, appropriately enough, should look almost self-evidently simple and effortless, if it is even noticed at all.

Why Do Cuts Work?

Well, the fact is that *Apocalypse Now*, as well as every other theatrical film (except perhaps Hitchcock's *Rope*³), is made up of many different pieces of film joined together into a mosaic of images. The mysterious part of it, though, is that the joining of those pieces—the "cut" in American terminology⁴—actually does seem to work, even though it represents a total and instantaneous displacement of one field of vision with another, a displacement that sometimes also entails a jump forward or backward in time as well as space.

It works; but it could easily have been otherwise, since nothing in our day-to-day experience seems to prepare us for such a thing. Instead, from the moment we get up in the morning until we close our eyes at night, the visual reality we perceive is a continuous

³ A film composed of only ten shots, each ten minutes long, invisibly joined together, so that the impression is of a complete lack of editing.

⁴ I was aware, talking to an Australian audience, of the bias inherent in our respective languages. In the States, film is "cut," which puts the emphasis on *separation*. In Australia (and in Great Britain), film is "joined," with the emphasis on *bringing together*.

stream of linked images: In fact, for millions of years—tens, hundreds of millions of years—life on Earth has experienced the world this way. Then suddenly, at the beginning of the twentieth century, human beings were confronted with something else—edited film.

Under these circumstances, it wouldn't have been at all surprising to find that our brains had been "wired" by evolution and experience to reject film editing. If that had been the case, then the single-shot movies of the Lumière Brothers—or films like Hitchcock's *Rope*—would have become the standard. For a number of practical (as well as artistic) reasons, it is good that it did not.

The truth of the matter is that film is actually being "cut" twenty-four times a second. Each frame is a displacement from the previous one—it is just that in a continuous shot, the space/time displacement from frame to frame is small enough (twenty milliseconds) for the audience to see it as *motion within a context* rather than as twenty-four different contexts a second. On the other hand, when the visual displacement is great enough (as at the moment of the cut), we are forced to re-evaluate the new image as a *different context*: miraculously, most of the time we have no problem in doing this.

What we *do* seem to have difficulty accepting are the kind of displacements that are neither subtle nor total: Cutting from a full-figure master shot, for instance, to a slightly tighter shot that frames the actors from the ankles up. The new shot in this case is different enough to signal that *something* has changed, but not different enough to make us re-evaluate its

context: The displacement of the image is neither motion nor change of context, and the collision of these two ideas produces a mental jarring—a jump—that is comparatively disturbing.⁵

At any rate, the discovery early in this century that certain kinds of cutting "worked" led almost immediately to the discovery that films could be shot discontinuously, which was the cinematic equivalent of the discovery of flight: In a practical sense, films were no longer "earthbound" in time and space. If we could make films only by assembling all the elements simultaneously, as in the theater, the range of possible subjects would be comparatively narrow. Instead, Discontinuity is King: It is the central fact during the production phase of filmmaking, and almost all decisions are directly related to it in one way or another—how to overcome its difficulties and/or how to best take advantage of its strengths.⁶

The other consideration is that even if everything *were* available simultaneously, it is just very difficult

⁵ A beehive can apparently be moved two inches each night without disorienting the bees the next morning. Surprisingly, if it is moved two *miles*, the bees also have no problem: They are forced by the total displacement of their environment to re-orient their sense of direction, which they can do easily enough. But if the hive is moved two *yards*, the bees will become fatally confused. The environment does not seem different to them, so they do not re-orient themselves, and as a result, they will not recognize their own hive when they return from foraging, hovering instead in the empty space where the hive used to be, while the hive itself sits just two yards away.

⁶ When Stanley Kubrick was directing *The Shining*, he wanted to shoot the film in continuity and to have all sets and actors available all the time. He took over almost the entire studio at Elstree (London), built all the sets simultaneously, and they sat there, pre-lit, for however long it took him to shoot the film. But *The Shining* remains a special exception to the general rule of discontinuity.

to shoot long, continuous takes and have all the contributing elements work each time. European filmmakers tend to shoot more complex master shots than the Americans, but even if you are Ingmar Bergman, there's a limit to what you can handle: Right at the end, some special effect might not work or someone might forget their lines or some lamp might blow a fuse, and now the whole thing has to be done again. The longer the take, of course, the greater the chances of a mistake.

So there is a considerable logistical problem of getting everything together at the same time, and then just as serious a problem in getting it all to "work" every time. The result is that, for practical reasons alone, we don't follow the pattern of the Lumière Brothers or of *Rope*.

On the other hand, apart from matters of convenience, discontinuity also allows us to choose the best camera angle for each emotion and story point, which we can edit together for a cumulatively greater impact. If we were limited to a continuous stream of images, this would be difficult, and films would not be as sharp and to the point as they are.⁷

⁷ Visual discontinuity—although not in the temporal sense—is the most striking feature of Ancient Egyptian painting. Each part of the human body was represented by its most characteristic and revealing angle: head in profile, shoulders frontal, arms and legs in profile, torso frontal—and then all these different angles were combined in one figure. To us today, with our preference for the unifying laws of perspective, this gives an almost comic "twisted" look to the people of Ancient Egypt—but it may be that in some remote future, our films, with their combination of many different angles (each being the most "revealing" for its particular subject), will look just as comic and twisted.

And yet, beyond even these considerations, cutting is more than just the convenient means by which discontinuity is rendered continuous. It is in *and for itself*—by the very force of its paradoxical suddenness—a positive influence in the creation of a film. We would want to cut even if discontinuity were not of such great practical value.

So the central fact of all this is that cuts *do work*. But the question still remains: *Why?* It is kind of like the bumble-bee, which should not be able to fly, but does.

We will get back to this mystery in a few moments.

"Cut Out the Bad Bits"

Many years ago, my wife, Aggie, and I went back to England for our first anniversary (she is English, although we'd been married in the United States), and I met some of her childhood friends for the first time.

"Well, what is it that you do?" one of them asked, and I replied that I was studying film editing. "Oh, editing," he said, "that's where you cut out the bad bits." Of course, I became (politely) incensed: "It is much more than that. Editing is structure, color, dynamics, manipulation of time, all of these other things, etc., etc." What he had in mind was home movies: "Oop, there's a bad bit, cut it out and paste the rest back together." Actually, twenty-five years down the road, I've come to respect his unwitting wisdom.

Because, in a certain sense, editing *is* cutting out the bad bits, the tough question is, *What makes a bad bit?* When you are shooting a home movie and the camera wanders, that's obviously a bad bit, and it's clear that you want to cut it out. The goal of a home movie is usually pretty simple: an unstructured record of events in continuous time. The goal of nar-

rative films is much more complicated because of the fragmented time structure and the need to indicate internal states of being, and so it becomes proportionately more complicated to identify what is a "bad bit." And what is bad in one film may be good in another. In fact, one way of looking at the process of making a film is to think of it as the search to identify what—for the particular film you are working on—is a uniquely "bad bit." So, the editor embarks on the search to identify these "bad bits" and cut them out, provided that doing so does not disrupt the structure of the "good bits" that are left.

Which leads me to chimpanzees.

About forty years ago, after the double-helix structure of DNA was discovered, biologists hoped that they now had a kind of map of the genetic architecture of each organism. Of course, they didn't expect the structure of the DNA to look like the organism they were studying (the way a map of England *looks* like England), but rather that each point in the organism would somehow correspond to an equivalent point in the DNA.

That's not what they found, though. For instance, when they began to compare them closely, they were surprised to discover that the DNA for the human and the chimpanzee were surprisingly similar. So much so—ninety-nine percent identical—as to be inadequate to explain all of the obvious differences between us.

So where do the differences come from?

Biologists were eventually forced to realize that there must be something else—still under much dis-

cussion—that controlled the *order* in which the various pieces of information stored in the DNA would be activated and the *rates* at which that information would be activated as the organism grew.

In the early stages of fetal development, it is difficult to tell the difference between human and chimp embryos. And yet, as they grow, they reach a point where differences become apparent, and from that point on, the differences become more and more obvious. For instance, the choice of what comes first, the brain or the skull. In human beings, the priority is brain first, skull next, because the emphasis is on maximizing the size of the brain. Any time you look at a newborn human infant you can see that the skull is not yet fully closed around the top of the still-growing brain.

With chimpanzees, the priority is reversed: skull first, *then* brain—probably for reasons that have to do with the harsher environment into which the chimp is born. The command from the chimp's sequence is, "Fill up this empty space with as much brain as you can." But there's only so much brain you can get in there before you can't fill it up anymore. At any rate, it seems to be more important for a chimp to be born with a hard head than a big brain. There's a similar interplay between an endless list of things: The thumb and the fingers, skeletal posture, certain bones being fully formed before certain muscular developments, etc.

My point is that the information in the DNA can be seen as uncut film and the mysterious sequencing code as the editor. You could sit in one room with a

pile of dailies and another editor could sit in the next room with exactly the same footage and both of you would make different films out of the same material. Each is going to make different choices about how to structure it, which is to say *when* and *in what order* to release those various pieces of information.

Do we know, for instance, that the gun is loaded *before* Madame X gets into her car, or is that something we only learn *after* she is in the car? Either choice creates a different sense of the scene. And so you proceed, piling one difference on top of another. Reversing the comparison, you can look at the human and the chimp as different films edited from the same set of dailies.⁸

I'm not assigning relative values here to a chimpanzee or a human being. Let's just say that each is appropriate to the environment in which it belongs: I would be wrong swinging from a branch in the middle of the jungle, and a chimpanzee would be wrong writing this book. The point is not their intrinsic value, but rather the inadvisability of changing one's mind in the process of creating one of them. Don't start making a chimpanzee and then decide to turn it into a human being instead. That produces a stitched-together Frankenstein's monster, and we've all seen its equivalent in the theaters: Film "X" would have been a nice little movie, perfectly suited to its "environment," but in the middle of production someone got an inflated idea about its possibilities, and, as a result, it became boring and pretentious. It was

⁸ By the same token, a chimpanzee and a cockroach are made from different "dailies" to begin with.

a chimpanzee film that someone tried to turn it into a human-being film, and it came out being neither.

Or film "Y," which was an ambitious project that tried to deal with complex, subtle issues, but the studio got to it and ordered additional material to be shot, filled with action and sex, and, as a result, a great potential was reduced to something less, neither human nor chimp.

Most with the Least

You can never judge the quality of a sound mix simply by counting the number of tracks it took to produce it. Terrible mixes have been produced from a hundred tracks. By the same token, wonderful mixes have been made from only three tracks. It depends on the initial choices that were made, the quality of the sounds, and how capable the blend of those sounds was of exciting emotions hidden in the hearts of the audience. The underlying principle: Always try to do the most with the least—with the emphasis on try. You may not always succeed, but *attempt* to produce the greatest effect in the viewer's mind by the least number of things on screen. Why? Because you want to do only what is necessary to engage the imagination of the audience—suggestion is always more effective than exposition. Past a certain point, the more effort you put into wealth of detail, the more you encourage the audience to become spectators rather than participants. The same principle applies to all the various crafts of filmmaking: acting, art direction, photography, music, costume, etc.

And, of course, it applies to editing as well. You would never say that a certain film was well-edited

because it had more cuts in it. Frequently, it takes more work and discernment to decide where *not* to cut—don't feel you have to cut just because you are being paid to. You are being paid to make decisions, and as far as whether to cut or not, the editor is actually making twenty-four decisions a second: "No. No. No. No. No. No. No. No. No. No. No. Yes!"

An overactive editor, who changes shots too frequently, is like a tour guide who can't stop pointing things out: "And up there we have the Sistine Ceiling, and over here we have the Mona Lisa, and, by the way, look at these floor tiles . . ." If you are on a tour, you do want the guide to point things out for you, of course, but some of the time you just want to walk around and see what *you* see. If the guide—that is to say, the editor—doesn't have the confidence to let people themselves occasionally choose what they want to look at, or to leave things to their imagination, then he is pursuing a goal (complete control) that in the end is self-defeating. People will eventually feel constrained and then resentful from the constant pressure of his hand on the backs of their necks.

Well, if what I'm saying is to do more with less, then is there any way to say how much less? Is it possible to take this right to its absurd logical conclusion and say, "Don't cut at all?" Now we've come back to our first problem: Film is cut for practical reasons *and* film is cut because cutting—that sudden disruption of reality—can be an effective tool in itself. So, if the goal is as few cuts as possible, when you *have* to make a cut, what is it that makes it a good one?

The Rule of Six

The first thing discussed in film-school editing classes is what I'm going to call three-dimensional continuity: In shot A, a man opens a door, walks half-way across the room, and then the film cuts to the next shot, B, picking him up at that same halfway point and continuing with him the rest of the way across the room, where he sits down at his desk, or something.

For many years, particularly in the early years of sound film, that was the rule. You struggled to preserve continuity of three-dimensional space, and it was seen as a failure of rigor or skill to violate it.⁹ Jumping people around in space was just not done, except, perhaps, in extreme circumstances—fights or earthquakes—where there was a lot of violent action going on.

I actually place this three-dimensional continuity at the bottom of a list of six *criteria* for what makes a

⁹ The problem with this thinking can be seen in any multi-camera situation-comedy on television. Because the cameras are filming simultaneously, the actors are necessarily always "correct" as far as their spatial continuity and relation to each other is concerned, but that absolutely does not prevent bad cuts from being made all the time.

good cut. At the top of the list is Emotion, the thing you come to last, if at all, at film school largely because it's the hardest thing to define and deal with. *How do you want the audience to feel?* If they are feeling what you want them to feel all the way through the film, you've done about as much as you can ever do. What they finally remember is not the editing, not the camerawork, not the performances, not even the story—it's how they felt.

An ideal cut (for me) is the one that satisfies all the following six criteria at once: 1) it is true to the emotion of the moment; 2) it advances the story; 3) it occurs at a moment that is rhythmically interesting and "right"; 4) it acknowledges what you might call "eye-trace"—the concern with the location and movement of the audience's focus of interest within the frame; 5) it respects "planarity"—the grammar of three dimensions transposed by photography to two (the questions of stage-line, etc.); 6) and it respects the three-dimensional continuity of the actual space (where people are in the room and in relation to one another).

1) Emotion	51%
2) Story	23%
3) Rhythm	10%
4) Eye-trace	7%
5) Two-dimensional plane of screen	5%
6) Three-dimensional space of action	4%

Emotion, at the top of the list, is the thing that you should try to preserve at all costs. If you find you have to sacrifice certain of those six things to

make a cut, sacrifice your way up, item by item, from the bottom.

For instance, if you are considering a range of possible edits for a particular moment in the film, and you find that there is one cut that gives the right emotion *and* moves the story forward, *and* is rhythmically satisfying, *and* respects eye-trace and planarity, *but* it fails to preserve the continuity of three-dimensional space, then, by all means, that is the cut you should make. If none of the other edits has the right emotion, then sacrificing spatial continuity is well worth it.

The values I put after each item are slightly tongue-in-cheek, but not completely: Notice that the top two on the list (emotion and story) are worth far more than the bottom four (rhythm, eye-trace, planarity, spatial continuity), and when you come right down to it, under most circumstances, the top of the list—emotion—is worth more than all five of the things underneath it.

And, in fact, there is a practical side to this, which is that if the emotion is right and the story is advanced in a unique, interesting way, in the right rhythm, the audience will tend to be unaware of (or unconcerned about) editorial problems with lower-order items like eye-trace, stage-line, spatial continuity, etc. The general principle seems to be that satisfying the criteria of items higher on the list tends to obscure problems with items lower on the list, but not vice-versa: For instance, getting Number 4 (eye-trace) working properly will minimize a problem with Number 5 (stage-line), whereas if Number 5 (stage-line) is correct but

Number 4 (eye-trace) is not taken into consideration, the cut will be unsuccessful.

Now, in practice, you will find that those top three things on the list—emotion, story, rhythm—are extremely tightly connected. The forces that bind them together are like the bonds between the protons and neutrons in the nucleus of the atom. Those are, by far, the tightest bonds, and the forces connecting the lower three grow progressively weaker as you go down the list.

Most of the time you will be able to satisfy all six criteria: the three-dimensional space and the two-dimensional plane of the screen and the eye-trace, and the rhythm and story and emotion will all fall into place. And, of course, you should always aim for this, if possible—never accept less when more is available to you.

What I'm suggesting is a list of priorities. If you have to give up something, don't ever give up emotion before story. Don't give up story before rhythm, don't give up rhythm before eye-trace, don't give up eye-trace before planarity, and don't give up planarity before spatial continuity.

Misdirection

Underlying these considerations is the central preoccupation of a film editor, which should be to put himself/herself in place of the audience. What is the audience going to be thinking at any particular moment? Where are they going to be looking? What do you want them to think about? What do they need to think about? And, of course, what do you want them to feel? If you keep this in mind (and it's the preoccupation of every magician), then you are a kind of magician. Not in the supernatural sense, just an everyday, working magician.

Houdini's job was to create a sense of wonder, and to do that he didn't want you to look *here* (to the right) because that's where he was undoing his chains, so he found a way to make you look *there* (to the left). He was "misdirecting" you, as magicians say. He was doing something that would cause ninety-nine percent of you to look over here when he wanted you to. And an editor can do that and does do that—and *should* do that.

Sometimes, though, you can get caught up in the details and lose track of the overview. When that hap-

pens to me, it is usually because I have been looking at the image as the miniature it is in the editing room, rather than seeing it as the mural that it will become when projected in a theater. Something that will quickly restore the correct perspective is to imagine yourself very small, and the screen very large, and pretend that you are watching the finished film in a thousand-seat theater filled with people, and that the film is beyond the possibility of any further changes. If you still like what you see, it is probably okay. If not, you will now most likely have a better idea of how to correct the problem, whatever it is. One of the tricks I use to help me achieve this perspective is to cut out little paper dolls—a man and a woman—and put one on each side of the editing screen: The size of the dolls (a few inches high) is proportionately correct to make the screen seem as if it is thirty feet wide.

Seeing Around the Edge of the Frame

The film editor is one of the few people working on the production of a film who does not know the exact conditions under which it was shot (or has the *ability* not to know) *and* who can at the same time have a tremendous influence on the film.

If you have been on and around the set most of the time, as the actors, the producer, director, cameraman, art director, etc., have been, you can get caught up in the sometimes bloody practicalities of gestation and delivery. And then when you see the dailies, you can't help, in your mind's eye, seeing around the edge of the frame—you can imagine everything that was there, physically and emotionally, just beyond what was actually photographed.

"We worked like hell to get that shot, it has to be in the film." You (the director, in this case) are convinced that what you got was what you wanted, but there's a possibility that you may be forcing yourself to see it that way because it cost so much—in money, time, angst—to get it.

By the same token, there are occasions when you shoot something that you dislike, when everyone is in a bad mood, and you say under protest: "All right, I'll do this, we'll get this one close-up, and then it's a wrap." Later on, when you look at that take, all you can remember was the hateful moment it was shot, and so you may be blind to the potentials it might have in a different context.

The editor, on the other hand, should try to see only what's on the screen, as the audience will. Only in this way can the images be freed from the context of their creation. By focusing on the screen, the editor will, hopefully, use the moments that should be used, even if they may have been shot under duress, and reject moments that should be rejected, even though they cost a terrible amount of money and pain.

I guess I'm urging the preservation of a certain kind of virginity. Don't *unnecessarily* allow yourself to be impregnated by the conditions of shooting. Try to keep up with what's going on but try to have as little specific knowledge of it as possible because, ultimately, the audience knows nothing about any of this—and you are the ombudsman for the audience.

The director, of course, is the person most familiar with all of the things that went on during the shoot, so he is the most burdened with this surplus, beyond-the-frame information. Between the end of shooting and before the first cut is finished, the very best thing that can happen to the director (and the film) is that he say goodbye to everyone and disappear for two weeks—up to the mountains or down to the sea or out to Mars or somewhere—and try to discharge this surplus.

Wherever he goes, he should try to think, as much as possible, about things that have absolutely nothing to do with the film. It is difficult, but it is necessary to create a barrier, a cellular wall between shooting and editing. Fred Zinnemann would go climbing in the Alps after the end of shooting, just to put himself in a potentially life-threatening situation where he had to be *there*, not day-dreaming about the film's problems.

Then, after a few weeks, he would come down from the Alps, back to earth; he would sit down in a dark room, alone, the arc light would ignite, and he would watch his film. He would still be, inherently, brimming with those images from beyond the edge of the frame (a director will never be fully able to forget them), but if he had gone straight from shooting to editing, the confusion would be worse and he would have gotten the two different thought processes of shooting and editing irrevocably mixed up.

Do everything you can to help the director erect this barrier for himself so that when he first sees the film, he can say, "All right, I'm going to pretend that I had nothing to do with this film. It needs some work. What needs to be done?"

And so you try as hard as you can to separate out what you wish from what is actually there, never abandoning your ultimate dreams for the film, but trying as hard as you can to see what is actually on the screen.

Dreaming in Pairs

In many ways, the film editor performs the same role for the director as the text editor does for the writer of a book—to encourage certain courses of action, to counsel against others, to discuss whether to include specific material in the finished work or whether new material needs to be added. At the end of the day, though, it is the writer who then goes off and puts the words together.

But in film, the editor also has the responsibility for actually assembling the images (that is to say, the “words”) in a certain order and in a certain rhythm. And here it becomes the *director’s* role to offer advice and counsel much as he would to an actor interpreting a part. So it seems that the film editor/director relationship oscillates back and forth during the course of the project, the numerator becoming the denominator and vice versa.

In dream therapy there is a technique that pairs the patient—the *dreamer*, in this case—with someone who is there to *listen* to the dream. As soon as possible after waking, the dreamer gets together with his listener to review the dreams of the previous night.

Frequently there is nothing, or just a single disappointing image, but this is usually enough to begin the process. Once the image is described, the listener’s job is to propose an imaginary sequence of events based on that fragment. An airplane, for instance, is all that is remembered. The listener immediately proposes that it must have been an airliner flying over Tahiti filled with golf balls for a tournament in Indonesia. No sooner has this description been offered than the dreamer finds himself protesting: “No, it was a bi-plane, flying over the battlefields of France, and Hannibal was shooting arrows at it from his legion of elephants.” In other words, the dream itself, hidden in the memory, rises to its own defense when it hears itself being challenged by an alternate version, and so reveals itself. This revelation about bi-planes and elephants can in turn prompt the listener to elaborate another improvisation, which will coax out another aspect of the hidden dream, and so on, until as much of the dream is revealed as possible.

The relationship between director and editor is somewhat similar in that the director is generally the dreamer and the editor is the listener. But even for the most well-prepared of directors, there are limits to the imagination and memory, particularly at the level of fine detail, and so it is the editor’s job to propose alternate scenarios as bait to encourage the sleeping dream to rise to its defense and thus reveal itself more fully. And these scenarios unfold themselves at the largest level (should such-and-such a scene be removed from the film for the good of the whole?) and at the most detailed (should this shot end on this frame

or 1/24th of a second later on the *next* frame?). But sometimes it is the editor who is the dreamer and the director who is the listener, and it is he who now offers the bait to tempt the collective dream to reveal more of itself.

As any fisherman can tell you, it is the quality of the bait that determines the kind of fish you catch.

Team Work: Multiple Editors

Not only does the editor collaborate with the director, there are frequent times when two or more editors are working simultaneously, sometimes with equal authority. This seems odd to many people, who do not see the same thing happening with directors of photography or production designers. But for some reason, which has to do with the collaborative mentality of editors and with the fact that the time pressure of post-production is not quite so unforgiving in its consequences as it is during production, multiple editors are often employed. I have worked, and enjoyed, collaborating with other editors on many films: *The Conversation*, *Apocalypse Now*, *The Unbearable Lightness of Being*, and *Godfather, Part III*.

The main advantage to collaborative editing is speed; the main risk is lack of coherence. But if there are upward of 350,000 feet of workprint (sixty-five hours), you are probably going to need to take that risk and have two editors, or at least an associate editor working under supervision. But problems can

sometimes arise if there is just one editor on a film and he develops a locked viewpoint about the material. This is particularly troublesome if the director and the editor have not worked together before and have no time to develop a common language. In this case, it might not be a bad idea to consider having multiple editors.

The Godfather was the first film on which Francis worked with two editors. Originally there had been just a single editor, but the problem of locked viewpoint became acute and he was let go after several months. The decision was made to reconstitute what had been done up to that point and start again, but because they had effectively lost those months, and it looked as though the film was going to be almost three hours long with an inflexible deadline, it made sense to hire two editors. The film was still shooting and there was just a lot of work to do: Each editor had a ninety-minute film to complete in twenty-four weeks. But unlike the later *Godfather, Part II* or *Apocalypse*, the work was split strictly in half. Bill Reynolds cut the first part and Peter Zinner cut the last part. There's a specific point where Bill's section ends and Peter's begins.

On *Godfather, Part II*, although the responsibility for editing was divided up in a checkerboard pattern, scenes were initially cut and recut by the same person.¹⁰ But when Francis began to play with the structure of the film, people found themselves recutting what others had originally edited.

¹⁰ The editors of *Godfather, Part II*, were Peter Zinner, Barry Malkin, and Richard Marks.

The interest on a \$25 million film is around \$250,000 a month. If having two editors can help you release that film a month earlier, they will have repaid a good portion, if not all, of their salaries for the whole film. It is simply a matter of how much you want to achieve in the time you have available. If you end up with a cut-per-day rate of 1.47, as we did on *Apocalypse*, that means that many different avenues have been explored to get to the final product. If that's what you want to do, you probably need more than one editor.

The Decisive Moment

While Phil Kaufman was shooting *The Unbearable Lightness of Being* in France, I was editing it in Berkeley, California—6,000 miles away. The dailies would be shipped approximately every two weeks, and I would then sit and watch ten hours or so of film, taking notes, making sure it was in sync, getting ready to code it, etc.

But in addition to the usual procedures, I also would select at least one representative frame from every setup (camera position) and take a still photograph of it off the workprint. We then had these photos developed and printed at the local “one hour” place, like family snapshots, and they were put onto panels arranged according to scene. Whenever a setup had complex staging or a moving camera, it was necessary to take more than one photo (I think the most that I ever took for *Unbearable* was seven, which was for a very complicated party scene)—usually it was three, and most of the time it was just one.

We had to use a special reproducing negative to make these pictures, since an ordinary snapshot negative produces too much contrast. The speed of the


film is low—about ASA 2 or something—so the time exposure has to be quite long, but it worked out well: The photos were actually pretty close approximations of the real color balance and contrast ratio of the film.

The photographs are a great help in later discussions with the director about what was shot and how it was shot—they resolve those kinds of discussions very quickly.

They also provide a record of some details beyond the ability of even the best continuity person: The particulars of an actor’s hairline, or a little peculiarity of costume, the way the collar flips up or down, or how rosy somebody’s complexion was, or whether their hat left a mark on their forehead when they took it off—those kinds of things.


They are also a great resource for the publicity department or anyone else coming onto the film at a later date. You can instantly see and cross-reference characters in all kinds of different emotional states, as well as the photography, the costumes, and the scenery.

Also, just because of the way they were displayed, the pictures tended to collide against each other in interesting ways. On *Unbearable*, we had perhaps sixteen panels of photos, 130 photos to a panel, and each panel was laid out like a page of a book: You “read” the photos from left to right and then down a row, left to right again, etc., just like reading text, and when you got to the bottom of one panel, you went up to the top of the next and read




Two frames of Teresa (played by Juliette Binoche) from *The Unbearable Lightness of Being*. The three-digit numbers in the lower left-hand corner of each picture (620, 296) refer to the setup number of the shot from which the frame was taken, and the number (2.2) in the adjacent rectangle of #620 identifies that frame as occupying a certain place in the sequence—frame two in a series of two.






Two frames of Sabina (played by Lena Olin) from *The Unbearable Lightness of Being*. The numbering system here is the same as for Teresa's photos on the previous pages. Both these frames come from the impromptu photo session between Sabina and Teresa.





Two frames of Teresa from *The Unbearable Lightness of Being*. In the case of #542, one frame was sufficient to give a complete idea of the setup. Number 635, however, required three photos—of which this is number two—because of the complex nature of the shot. This comes from the same scene as #634 of Sabina and the two shots were cut together in the finished film.



across the first line, etc. So the juncture *between* those panels was an interesting thing to look at, because it juxtaposed frames that were never meant to go together and yet there they were, right next to each other. And sometimes you got sparks out of that, it would cause you to think about things, editorial leaps, that otherwise you might never have thought of without this system.

But the most interesting asset of the photos for me was that they provided the hieroglyphs for a language of emotions.

What word expresses the concept of ironic anger tinged with melancholy? There isn't a word for it, in English anyway, but you can see that specific emotion represented in this photograph.

Or the photo may represent a kind of nervous anticipation: The character is scared and lustful at the same time, and yet she is confused because that lust is for another woman. And that woman is sleeping with her husband. So what does that mean?

Whatever it means, it is there in her expression, in the angle of her head and her hair and her neck and the tension in the muscles and the set of her mouth and what is in her eyes. And if you can simply point to an expression on an actor's face, you have a way around some of the difficulties of language in dealing with the subtleties of nameless emotions. You, as the director, can say, "That's what I want. The sequence we are working on should have more of that, you know. I want it to embody the nameless but familiar emotion I see in that photograph."

The editor's job now is to choose the right images and make those images follow one another at the right rate to express something like what is captured in that photograph.

In choosing a representative frame, what you're looking for is an image that distills the essence of the thousands of frames that make up the shot in question, what Cartier-Bresson—referring to still photography—called the "decisive moment." So I think, more often than not, the image that I chose wound up in the film. And also, more often than not, quite close to the cut point.

When you look at dailies, there's a pitfall similar to the one that you may fall into during casting sessions. For the actor who is walking through the door, this casting session is the one and only time that he is going to present him or herself to you. This is a tremendously important moment for him, but for you, this may be the sixtieth person you've seen that day. Inevitably there is a kind of mental glaze that may form after a while that stops you thinking as sharply as you should.

Well, dailies are like that, too. Each shot is auditioning for your approval. Take five: "How about me? I can do this." Then take seven comes in the door: "How about this?" Take nine: "Or this?"

And to keep your awareness, to really be alive to the possibilities in each shot, you have to keep jabbing yourself. You try to remain fresh and see the wonderful things and make records of those as well as the things that may not be so wonderful. Which is what you have to do when you are casting a film.

But if you have to choose a representative set of stills from every setup, you will automatically start thinking differently—you have to be analytical right from the beginning, which is what you should be when you are looking at dailies. But, everyone being human and dailies sometimes going on as long as they do, we sometimes tend to just sit there and let the dailies roll over us. What this photo system does is just tip you up out of your chair a bit. It is an encouragement to do what you should be doing anyway. And it is the beginning of the editorial process. You are already beginning to edit at the point that you say, “I like this frame rather than that frame.”

Methods and Machines: Marble and Clay

The tools you choose to edit with can have a determining effect on the final product. But when I started editing in 1965, there was really just one option, at least in Hollywood: the Moviola, a “standup” editing machine—virtually unchanged since the 1930s—that looks something like a sewing machine (and sounds like one). In the early days of Zoetrope Studios, which was founded on a European model, we used imported Steenbecks or KEMs, “horizontal” editing machines from Germany that were quieter and handled the film more gently, had two large screens and two sound tracks, and were better able to handle large amounts of film. Now, of course, the landscape has been permanently altered by the computer—electronic digital editing machines, such as Avid and Lightworks, that harness together a video display terminal and a computer with a large-scale memory bank where the film’s images and sounds can be stored digitally.¹¹

¹¹ For a more complete survey of the current non-linear systems, see the Afterword: Non-Linear, Digital Editing—The Unbearable Lightness.

I feel equally comfortable working on a Moviola, a KEM Universal, or an Avid. It depends on the film, the film's budget and schedule, and on my hunch about the style of the director and how long the average take is going to be. *The Conversation* was edited on a KEM, *Julia* on a Moviola, *Apocalypse Now* on a KEM, whereas I used both on *Unbearable*—it started out on a Moviola and changed to a KEM. With the KEM, I arrange things opposite of the way they are normally done, so that the screen directly in front of me is my search screen and the soundtrack in front of me is in sync with that. The screen on the left side and the sound on the right side are in sync with each other and hold the assembled film. If I were working on a bench, that would be what's going through the synchronizer.

In fact, speaking of benches, I should add that I work standing up: My KEM is raised about fifteen inches off the floor to put it at the right height for me. One of the things I always liked about the Moviola is that you stand up to work, holding the Moviola in a kind of embrace—dancing with it, in a way—so it was frustrating for me to have to sit down at the KEM. I edited both *The Conversation* and *Apocalypse* that way, sitting down, but a voice in the back of my head kept telling me something wasn't quite right. And so when I came to edit *Unbearable*, I had the KEM raised up on two plywood boxes.

Editing is a kind of surgery—and have you ever seen a surgeon sitting to perform an operation? Editing is also like cooking—and no one sits down at the stove to cook. But most of all, editing is a kind of

dance—the finished film is a kind of crystallized dance—and when have you ever seen a dancer sitting down to dance?

Other than the “standing/sitting” question, the differences between the Moviola system and the KEM system boil down to sculptural ones: The Moviola system “emulsifies” the film into little bits (individual shots) and then the editor reassembles it out of those bits, like making something out of clay. You take a little bit of clay and you stick it here and you take another little bit of clay and you stick it *there*. At the beginning of the process there is nothing in front of you, then there is something in front of you, and then there is finally the finished thing all built up out of little clay bricks, little pellets of information.

With the KEM system, I don't ever break the film down into individual shots—I leave it in ten-minute rolls in the order in which it came from the lab. In sculptural terms, this is like a block of marble—the sculpture is already there, hidden within the stone, and you reveal it by taking away, rather than building it up piece by piece from nothing, as you do with clay. It is really the difference between “random-access” editing and its opposite, whatever that is called—“linear-access,” let's say.

Computerized digital editing and, strangely enough, good old-fashioned Moviola editing with an assistant, are both random-access, non-linear systems: You ask for something specific and that thing—that thing alone—is delivered to you as quickly as possible. You are only shown what you ask for. The Avid is faster at it than the Moviola, but the process is the same.

That's a drawback for me because your choices can then only be as good as your requests, and sometimes that is not enough. There is a higher level that comes through *recognition*: You may not be able to articulate what you want, but you can recognize it when you see it.

What do I mean by that? Well, if you learn to speak a foreign language, you will find that there is a gap between how well you can speak it and how well you can understand it when it is spoken to you. A human being's ability to understand a foreign language is always greater than his ability to speak it.

And when you make a film, you are trying to learn a foreign language—it just happens to be a unique language that is only spoken by this one film. If you have to articulate everything, as you do with a random-access system like video/computer or Moviola/assistant, you are limited by what and how much you can articulate and how good your original notes were. Whereas the advantage of the KEM's linear system is that I do not always have to be speaking to it—there are times when *it* speaks to *me*. The system is constantly presenting things for consideration, and a sort of dialogue takes place. I might say, "I want to see that close-up of Teresa, number 317, in roll 45." But I'll put that roll on the machine, and as I spool down to number 317 (which may be hundreds of feet from the start), the machine shows me everything at high speed down to that point, saying in effect: "How about this instead? Or this?" And I find, more often than not, long before I get down to shot 317, that I've had three other ideas triggered by the material that I have seen flashing by me.

"Oh, this other shot is much better than the one I thought I wanted." As soon as I saw it, I recognized it as a possibility, whereas I couldn't articulate it as a choice.

When you look at dailies the first time, you have a relatively fixed idea—based on the script—of what you are looking for. Later on, though, you review some of your original notes, and they will say, for instance: "212-4: NG." What does that mean? It meant that *at the time* you thought take four of Slate 212 was No Good, and you didn't bother to make a note of why you thought so. Well, many times in the re-editing, what you thought was originally unusable may come to be your salvation.

If it was a question of only one shot, or two dozen shots, you could probably deal with the problem of second-guessing those original notes, but, in fact, an ordinary film will have 700, 1,000, 2,000 setups with more than two printed takes per setup on average, and so there may easily be two to four thousand shots that you have to have an opinion about. That's a lot of second-guessing, so you tend to fall back on your first impressions in the original notes. Which are valuable, but limited, if that is all you have.

Whereas with the KEM system, because the film is all stored in these big rolls in an *almost* arbitrary way, you are learning something new about the material as you search for what you think you want. You are actually doing creative work, and you may find what you *really* want rather than what you *thought* you wanted. This system is useful enough in the original assembly of the scene, in helping you to familiar-

ize yourself with the material, but it becomes particularly valuable in the recutting, where your original notes—heavily influenced by the script—become less and less useful as the film finds its own voice.

There are limits, of course: A system that is too linear (which means that you have to spend too much time searching before you find what you want) would be burdensome. You would quickly become overwhelmed and/or bored with it. So there is a golden mean somewhere. If the system is completely random-access, that is a defect, in my opinion. But if it is too linear, that's a defect as well. What I've found, personally, is that given the attitude I have toward the material, given the speed that I work, and given the mechanical speed of the KEM itself, keeping the dailies in ten-minute reels in shooting order adds just about the right amount of chaos that I need in order to work the way I want to.

The digital systems, Avid and Lightworks, are promising because they offer the potential to merge the best of the Moviola's non-linear approach with the best of the KEM's ability to handle and review large amounts of material quickly. At this point, there is still some procedural awkwardness with the digital systems, but I expect those to be smoothed out over time.

In any case, there are certain things that remain the same for me no matter what system I am using. I would always review the material twice: once at the beginning, the day after the material was shot, noting down my first impressions and including any notes the director cares to give me. And then when I was

ready to cut a particular scene, I would collect all the relevant material and review it again, making notes in more detail than the first time.

When you look at rushes the second time, you have evolved and the film has evolved. You will see different things than you saw the first time, because you may have assembled scenes that hadn't been shot the first time you saw the material, and strengths or problems may be emerging with characters and events as they unfold.

In an ideal world, what I would like to do is assemble a first cut and then stop and just look at all the dailies again, fresh. Whether I would ever actually be able to do that is another question: The present schedule of films, at any rate, prohibits such a thing. This is where the hidden virtues of the linear (KEM) system come in—because of the way the material is stored, in ten-minute rolls of film, it is constantly being reviewed. If this were gardening, I would be talking about the virtues of turning over the soil and aerating it.

In the actual editing of a scene, I will keep on working until I can no longer "see myself" in the material. When I review my first assembly of a scene, more often than not I can still vividly (too vividly!) recall making the decisions that led to each of the cuts. But as the scene is reworked and refined, it reaches a point, hopefully, where the shots themselves seem to create each other: This shot "makes" the next shot, which "makes" the next shot, etc. In this way, the Walter Murch who decided things initially gradu-

ally recedes until, finally, there comes a point where he has become invisible and the characters take over, the shots, the emotion, the story seem to take over. Sometimes—the best times—this process reaches the point where I can look at the scene and say, “I didn’t have anything to do with that—it just created itself.”

As far as the color balance of the final prints goes, I’ve had a few good experiences, but many have been frustrating for one reason or another. The worst thing is that the labs show the film at twenty-four frames per second (and sometimes even at thirty-two fps) without the ability to stop or go backward. You are sitting there and saying: “That shot should be redder.” “Which one?” “The close-up of the foot.” But already eight shots have gone by. Most often, the impression you get is that they are doing this as a courtesy just to keep you quiet.

One way that works for me, where you can really see what’s going on and get specific with your notes about color, is to take the workprint and the first answer print and put them in sync over a light box of the correct color temperature. There is something about a low level of light coming through a piece of film that enables you to see tonalities that are almost invisible if you’re looking at a bright screen of projected light. There may be a residual greenness, you know, but you look at the film in a viewing theater and you’ll say, “I don’t know. Is that shot a little green? Or is it too blue?” And, of course, before you have decided, you are looking at something else. If you work with the light-box system, the green will just leap out, especially if you’ve got your original

workprint there for comparison. And you can stop, go forward, go backward, etc.

If you are fortunate enough to be working with a good color timer at a good laboratory, of course, it is like any sort of successful collaboration.

Test Screenings: Referred Pain

Toward the end of the editing process on *Julia*, Fred Zinnemann observed that he felt the director and the editor, alone with the film for months and months, could only go ninety percent of the way toward the finished film—that what was needed for the last ten percent was “the participation of the audience,” whom he saw as his final collaborators. Not in the sense that he would respond to them blindly, but that he felt their presence was helpful as a corrective, to keep certain obsessions from becoming corrosive and to point out blind spots that may have developed through over-familiarity with the material.

This has certainly been my experience: All of the films I have worked on have been tested before large audiences except for *The Conversation* and *Unbearable Lightness*. We had screenings of them, of course, but we never had wide-open, public screenings. Francis Coppola in particular has always been an enthusiastic supporter of screening his films almost at any stage, almost no matter how unfinished they were.

Rough screenings would be for small groups of about ten people whom he knew, mixed with two or three people who were strangers. The strangers would have no previous idea of what this film was about, and he would question them afterward, on a one-to-one basis, to compare their opinions to the reaction of the people who did know about the film.

Fred Zinnemann, by contrast, would show *Julia* to a preview audience only when it was technically completely finished, with a cut negative and optical sound track. He was perfectly prepared to change it after that, but he doesn't believe that general audiences can completely discount visible splices, color mismatches, and incomplete sound tracks, and I agree with him.

Even with technically finished films, public previews are tricky things. You can learn a tremendous amount from them, but you have to be cautious about direct interpretations of what people have to say to you, particularly on those cards they fill out after the screening. I'm extremely suspicious of those. The most helpful thing of all is simply learning how *you* feel when the film is being shown to 600 people who have never seen it before. Emotionally, it seems like some big hand has come and grabbed you by the hair, picked you up, and put you down ninety degrees to one side. And you think, “Oh God, look at *that*.” It's as if up to this moment you have been constructing a building but always standing in front of it to evaluate it. Now all of a sudden you are looking at the side of the building and seeing things you seem to have never seen before.

You shouldn't blindly follow what you learn from these test screenings any more than you should anything else. What can you learn from the differences *between* the previous screenings and this one? Given these two headings, where is the North Pole? Test screenings are just a way to find out where you are.

There was one procedure on *Julia* that, unfortunately, I have never seen repeated. We had a person sitting at a table in the lobby of the preview theater with a sign in front of him that said, "If you want to talk to us on the telephone after a few days, leave your number here." And then those conversations were transcribed and added into the survey. If you are going to do previews and listen to what people have to say, that's the way to do it—after they have had a day or two to let the film sink in. Don't look at what people write in the heat of the moment—you get a reaction, but it is a skewed reaction. There's a lot of what is medically called "referred pain" in that process.

When you go to a doctor and tell him that you have a pain in your elbow, it is the quack who takes out his scalpel and starts to operate on the elbow. Then you wind up with not only the original pain but probably a pain in your wrist and your shoulder as well. Whereas an experienced doctor studies you, takes an x-ray, and determines that the cause of the pain is probably a pinched nerve up in your shoulder—you just happen to feel it in your elbow. The pain in the shoulder has been "referred" to the elbow. Audience reactions are like that. When you ask the direct question, "What was your least favorite scene?" and eighty percent of the people are in agree-

ment about one scene they do not like, the impulse is to "fix" the scene or cut it out. But the chances are that that scene is fine. Instead, the problem may be that the audience simply didn't understand something that they needed to know for the scene to work.

So, instead of fixing the scene itself, you might clarify some exposition that happens five minutes earlier. Don't necessarily operate on the elbow: instead, discover if nerves are being pinched somewhere else. But the audience will never tell you that directly. They will simply tell you where the pain is, not the source of the pain.

Editing decisions become particularly acute in the last days before the film is released, since changes made now will be permanent. If you, as the editor, have a particularly strong feeling about something at this stage, you should try to make your point as forcefully and convincingly as you can—perhaps you stay late and do a test version of your idea, sketch something out—but you also need to have discretion, a sense of who you are dealing with, and present your ideas to the director or producer in the right context. And how you go about this has to do with your whole working history, how you were hired, how much you respect the director, how much the director respects you.

I remember one moment particularly, after the previews of *Julia*, when Fred Zinnemann and I were debating what final changes to make in the structure of the beginning, which seemed to have been difficult for the audience to follow. The opening reel of the film had a nested sequence of flashbacks—a

memory of a memory of a memory of a memory—and it was perhaps one too many. I suggested eliminating of one scene that occupied a unique time-frame in the film's structure (one that was never reprised), and we decided to remove this, since it meant that the scenes that were left would consequently sort themselves into a more graspable sequence. As I was undoing the splices (and they made a little screech as they came apart, almost as if they were crying out in pain), Zinnemann looked thoughtfully at what was happening and observed in an almost offhand way, "You know, when I first read this scene in the script, I knew that I could do this film."

I hesitated briefly, looked at him, and then continued undoing the splices. But my heart was in my throat because at that stage in the process you do not know; you can only have *faith* that what you are doing is the right thing. Were we mistakenly cutting out the heart of the film, or were we snipping off the umbilical cord?

In retrospect, I believe it *was* the umbilical cord and that we were right to remove it: The scene did have an essential function at one point, which was to connect Fred Zinnemann to the project, but once that connection had been made and Zinnemann's sensibility had flowed through that scene into all the other scenes in the film, it could finally be removed without any harm.

But things like that do give you pause.

Don't Worry, It's Only a Movie

Earlier I asked the question, "Why do cuts work?" We *know* that they do. And yet it is still surprising when you think about it because of the violence of what is actually taking place: At the instant of the cut, there is a total and instantaneous discontinuity of the field of vision.

I recall once coming back to the editing room after a few weeks in the mixing theater (where all movements are smooth and incremental) and being appalled at the brutality of the process of cutting. The "patient" is pinned to the slab and: Whack! Either/Or! This not That! In or Out! We chop up the poor film in a miniature guillotine and then stick the dismembered pieces together like Dr. Frankenstein's monster. The difference (the miraculous difference) is that out of this apparent butchery our creation can sometimes gain not only a life but a soul as well. It is all the more amazing because the instantaneous displacement achieved by the cut is not anything that we experience in ordinary life.

We are accustomed to such things, of course, in music (Beethoven was the innovator and master of this) as well as in our own thoughts—the way one realization will suddenly overwhelm everything else, to be, in turn, replaced by yet another. But in the dramatic arts—theater, ballet, opera—there didn't seem to be any way to achieve total instantaneous displacement: stage machinery can only move so fast, after all. *So why do cuts work?* Do they have some hidden foundation in our own experience, or are they an invention that suits the convenience of filmmakers and people have just, somehow, become used to them?

Well, although “day-to-day” reality appears to be continuous, there *is* that other world in which we spend perhaps a third of our lives: the “night-to-night” reality of dreams. And the images in dreams are much more fragmented, intersecting in much stranger and more abrupt ways than the images of waking reality—ways that approximate, at least, the interaction produced by cutting.

Perhaps the explanation is as simple as that: We accept the cut because it resembles the way images are juxtaposed in our dreams. In fact, the abruptness of the cut may be one of the key determinants in actually *producing* the similarity between films and dreams. In the darkness of the theater, we say to ourselves, in effect, “This looks like reality, but it cannot be reality because it is so visually discontinuous; therefore, it must be a dream.”

(Along those lines, it is revealing that the words a parent uses to comfort a child frightened by a nightmare—“Don't worry, darling, it's only a dream”—are

almost the same words used to comfort a child frightened by a film—“Don't worry, darling, it's only a movie.” Frightening dreams and films have a similar power to overwhelm the defenses that are otherwise effective against equally frightening books, paintings, music. For instance, it is hard to imagine this phrase: “Don't worry, darling, it's only a painting.”)

The problem with all this is that the comparison of films and dreams is interesting, probably true, but relatively barren of practical fruits: We still know so little about the nature of dreams that the observation comes to a stop once it has been made.

Something to consider, though, is the possibility that there may be a part of our waking reality where we actually do experience something like cuts, and where daylight images are somehow brought in closer, more discontinuous, juxtaposition than might otherwise seem to be the case.

I began to get a glimmer of this on my first picture-editing job—*The Conversation* (1974)—when I kept finding that Gene Hackman (Harry Caul in the film) would blink very close to the point where I had decided to cut. It was interesting, but I didn't know what to make of it.

Then, one morning after I had been working all night, I went out to get some breakfast and happened to walk past the window of a Christian Science Reading Room, where the front page of the *Monitor* featured an interview with John Huston. I stopped to read it, and one thing struck me forcefully because it related exactly to this question of the blink:

"To me, the perfect film is as though it were unwinding behind your eyes, and your eyes were projecting it themselves, so that you were seeing what you wished to see. Film is like thought. It's the closest to thought process of any art.

"Look at that lamp across the room. Now look back at me. Look back at that lamp. Now look back at me again. Do you see what you did? You *blinked*. Those are *cuts*. After the first look, you know that there's no reason to pan continuously from me to the lamp because you know what's in between. Your mind cut the scene. First you behold the lamp. *Cut*. Then you behold me."¹²

What Huston asks us to consider is a physiological mechanism—the blink—that interrupts the apparent visual continuity of our perceptions: My head may move smoothly from one side of the room to the other, but, in fact, I am cutting the flow of visual images into significant bits, the better to juxtapose and compare those bits—"lamp" and "face" in Huston's example—without irrelevant information getting in the way.

Of course there are limits to the kind of juxtapositions I can make this way—I can't jump forward or backward in time and space (that is the prerogative of dreams and films).¹³ But even so, the visual displacements available to me just by turning my head (from the Grand Canyon in front of me to the forest behind me, or even from one side of this room to the other) are sometimes quite great.

¹² *Christian Science Monitor*, August 11, 1973. John Huston interviewed by Louise Sweeney.

¹³ But see footnote #16.

After I read that article, I started observing people, watching when they blinked, and I began to discover something much different than what they tell you in high-school biology, which is that the blink is simply a means to moisten the surface of the eye. If that's all it is, then for each environment and each individual there would be a purely mechanical, predictable interval between blinks depending on the humidity, temperature, wind speed, etc. You would only blink when your eye began to get too dry, and that would be a constant number of seconds for each environment. This is clearly not the case: People will sometimes keep their eyes open for minutes at a time—at other times they will blink repeatedly—with many variations in between. The question then is, "What is causing them to blink?"

On the one hand, I'm sure you've all been confronted by someone who was so angry that he didn't blink at all: This is a person, I believe, in the grip of a single thought that he holds (and that holds him), inhibiting the urge and need to blink.¹⁴ And then there is the opposite kind of anger that causes someone to blink every second or so: This time, the person is being assailed simultaneously by many conflicting emotions and thoughts, and is desperately (but unconsciously) using those blinks to try to separate these thoughts, sort things out, and regain some kind of control.

¹⁴ There is that telling phrase from classic cowboy (and now diplomatic) stand-offs: "he blinked." The loser in this mental game of chicken could not hold fast to his single position and instead allowed some other thought to intrude at the critical moment. The blink signals the moment he relinquished his primary thought.

So it seems to me that our rate of blinking is somehow geared more to our emotional state and to the nature and frequency of our thoughts than to the atmospheric environment we happen to find ourselves in. Even if there is no head movement (as there was in Huston's example), the blink is either *something that helps an internal separation of thought to take place*, or it is *an involuntary reflex accompanying the mental separation that is taking place anyway*.¹⁵

And not only is the *rate* of blinking significant, but so is the actual *instant* of the blink itself. Start a conversation with somebody and watch when they blink. I believe you will find that your listener will blink at the precise moment he or she "gets" the idea of what you are saying, not an instant earlier or later. Why would this be? Well, speech is full of unobserved grace notes and elaborations—the conversational equivalents of "Dear Sir" and "Yours Sincerely"—and the essence of what we have to say is often sandwiched between an introduction and a conclusion. The blink will take place either when the listener realizes our "introduction" is finished and that now we are going to say something significant, or it will happen when he feels we are "winding down" and not going to say anything more significant for the moment.

And that blink will occur where a cut could have happened, had the conversation been filmed. Not a frame earlier or later.

So we entertain an idea, or a linked sequence of ideas, and we blink to separate and punctuate that idea from what follows. Similarly—in film—a shot

¹⁵ Dr. John Stern of Washington University in St. Louis has recently (1987) published experimental work in the psycho-physiology of the blink that seems to confirm this.

presents us with an idea, or a sequence of ideas, and the cut is a "blink" that separates and punctuates those ideas.¹⁶ At the moment you decide to cut, what you are saying is, in effect, "I am going to bring this idea to an end and start something new." It is important to emphasize that the cut by *itself* does not create the "blink moment"—the tail does not wag the dog. If the cut is well-placed, however, the more extreme the visual discontinuity—from dark interior to bright exterior, for instance—the more thorough the effect of punctuation will be.

At any rate, I believe "filmic" juxtapositions are taking place in the real world not only when we dream but also when we are awake. And, in fact, I would go so far as to say that these juxtapositions are not accidental mental artifacts but part of the method we use to make sense of the world: We must render visual reality discontinuous, otherwise perceived reality would resemble an almost incomprehensible string of letters without word separation or punctuation. When we sit in the dark theater, then we find edited film a (surprisingly) familiar experience. "More like thought than anything else," in Huston's words.¹⁷

¹⁶ This can occur regardless of how big or small the "idea" happens to be. For instance, the idea could be as simple as "she moves quickly to the left."

¹⁷ William Stokoe makes an intriguing comparison between the techniques of film editing and American Sign Language: "In signed language, narrative is no longer linear. Instead, the essence is to cut from a normal view to a close-up to a distant shot to a close-up again, even including flashback and flash-forward scenes, exactly as a movie editor works. Not only is signing arranged more like edited film than like written narration, but also each signer is placed very much as a camera: the field of vision and angle of view are directed but variable." William Stokoe, *Language in Four Dimensions*, New York Academy of Sciences (1979).

Dragnet

If it is true that our rates and rhythms of blinking refer directly to the rhythm and sequence of our inner emotions and thoughts, then those rates and rhythms are insights into our inner selves and, therefore, as characteristic of each of us as our signatures. So if an actor is successful at projecting himself into the emotions and thoughts of a character, his blinks will *naturally and spontaneously* occur at the point that the character's blinks would have occurred in real life.¹⁸

I believe this is what I was finding with Hackman's performance in *The Conversation*—he had assumed the character of Harry Caul, was thinking a series of Harry's thoughts the way Harry would think them, and, therefore, was blinking in rhythm with those thoughts. And since I was absorbing the rhythms he

¹⁸ One of the things about unsuccessful acting is that the actor's blinks seem to come at the "wrong" times. Although you may not notice this consciously, the rhythm of the actor's blinks don't match the rhythm of thoughts you would expect from the character he is playing. In fact, a bad actor is probably not thinking anything like what the character would be thinking. Instead: "I wonder what the director thinks of me, I wonder if I look okay," or "What's my next line?"

was giving me and trying to think similar thoughts myself, my cut points were naturally aligning themselves with his "blink points." In a sense, I had re-routed my neural circuitry so that the semi-involuntary command to blink caused me instead to hit the stop button on the editing machine.

To that same end, one of the disciplines I follow is to choose the "out point" of a shot by marking it in real time. If I can't do this—if I can't hit that same frame repeatedly at twenty-four frames per second—I know there is something wrong in my approach to the shot, and I adjust my thinking until I find a frame I *can* hit. I never permit myself to select the "out point" by inching back and forth, comparing one frame with another to get the best match. That method—for me, at any rate—is guaranteed to produce a rhythmic "tone deafness" in the film.

Anyway, another one of your tasks as an editor is this "sensitizing" of yourself to the rhythms that the (good) actor gives you, and then finding ways to extend these rhythms into territory not covered by the actor himself, so that the pacing of the film as a whole is an elaboration of those patterns of thinking and feeling. And one of the many ways you assume those rhythms is by noticing—consciously or unconsciously—where the actor blinks.

There is a way of editing that ignores all of these questions, what I would call the "Dragnet" system, from the 1950s TV series of the same name.

The policy of the show seemed to be to keep every word of dialogue on screen. When someone had

finished speaking, there was a brief pause and then a cut to the person, who was now about to talk, and when he in turn finished speaking there was a cut back to the first person who nodded his head or said something, and then when *that* person was finished, they cut back again, etc. It extended to single words. "Have you been downtown yet?" *Cut*. "No." *Cut*. "When are you going downtown?" *Cut*. "Tomorrow." *Cut*. "Have you seen your son?" *Cut*. "No, he didn't come home last night." *Cut*. "What time does he usually come home?" *Cut*. "Two o'clock." At the time, when it first came out, this technique created a sensation for its apparently hard-boiled, police-blotter realism.

The "Dragnet" system is a simple way to edit, but it is a shallow simplicity that doesn't reflect the grammar of complex exchanges that go on all the time in even the most ordinary conversations. If you're observing a dialogue between two people, you will not focus your attention solely on the person who is speaking. Instead, while *that person is still talking*, you will turn to look at the listener to find out what he thinks of what is being said. The question is, "When exactly do you turn?"

There are places in a conversation where it seems we almost physically *cannot* blink or turn our heads (since we are still receiving important information), and there are other places where we *must* blink or turn away in order to make better sense of what we have received. And I would suggest that there are similar points in every scene where the cut *cannot* or *must* occur, and for the same reasons. Every shot has po-

tential "cut points" the way a tree has branches, and once you have identified them, you will choose different points depending on what the audience has been thinking up to that moment and what you want them to think next.

For instance, by cutting away from a certain character *before* he finishes speaking, I might encourage the audience to think only about the face value of what he said. On the other hand, if I linger on the character *after* he finishes speaking, I allow the audience to see, from the expression in his eyes, that he is probably not telling the truth, and they will think differently about him and what he said. But since it takes a *certain amount of time* to make that observation, I cannot cut away from the character too early: Either I cut away while he is speaking (branch number one) or I hold until the audience realizes he is lying (branch number two), but *I cannot cut in between those two branches*—to do so would either seem too long or not long enough. The branch points are fixed organically by the rhythm of the shot itself and by what the audience has been thinking up to that moment in the film,¹⁹ but I am free to select one or the other of them (or yet another one further on) depending on what realization I want the audience to make.

In this way, you should be able to cut from the speaker to the listener and vice versa in psychologically interesting, complex, and "correct" patterns that reflect the kinds of shifts of attention and realization that go on in real life: In this way, you establish a

¹⁹ One way to shift the actual branch points themselves is to place the shot in a different context, where the audience will be thinking (and noticing) different things.

rhythm that counterpoints and underscores the ideas being expressed or considered. And one of the tools to identify exactly where these cut points, these "branches," may be is to compare them to our patterns of blinking, which have been underscoring the rhythm of our thoughts for tens of thousands, perhaps millions, of years of human history. Where you feel comfortable blinking—if you are really listening to what is being said—is where the cut will feel right.

So there are really three problems wrapped up together:

- 1) identifying a series of potential cut points (and comparisons with the blink can help you do this),
- 2) determining what effect each cut point will have on the audience, and
- 3) choosing which of those effects is the correct one for the film.

I believe the sequence of thoughts—that is to say, the rhythm and rate of cutting—should be appropriate to whatever the audience is watching at the moment. The average "real-world" rate of blinking is somewhere between the extremes of four and forty blinks per minute. If you are in an actual fight, you will be blinking dozens of times a minute because you are thinking dozens of conflicting thoughts a minute—and so when you are watching a fight in a film, there should be dozens of cuts per minute.²⁰ In

²⁰ This would make the audience participate emotionally in the fight itself. If, on the other hand, you wanted to create an objective distance—to have the audience observe the fight as a phenomenon in itself—then you would reduce the number of cuts considerably.

fact, statistically the two rates—of real-life blinking and of film cutting—are close enough for comparison: Depending on how it is staged, a convincing action sequence might have around twenty-five cuts a minute, whereas a dialogue scene would still feel "normal" (in an American film) averaging six cuts per minute or less.

You should be right with the blinks, perhaps leading them ever so slightly. I certainly don't expect the audience to blink at every cut—the cut point should be a *potential* blink point. In a sense, by cutting, by this sudden displacement of the visual field, you are blinking *for* the audience: You achieve the immediate juxtaposition of two concepts for them—what they achieve in the real world by blinking, as in Huston's example.

Your job is partly to anticipate, partly to control the thought processes of the audience. To give them what they want and/or what they need just before they have to "ask" for it—to be surprising yet self-evident at the same time. If you are too far behind or ahead of them, you create problems, but if you are right with them, leading them ever so slightly, the flow of events feels natural and exciting at the same time.

A Galaxy of Winking Dots

Along these lines, it would be fascinating to take an infrared film of an audience and find out when and in what patterns people blink when they are watching a movie. My hunch is that if an audience is really in the grip of a film, they are going to be thinking (and therefore blinking) with the rhythm of the film.

There is a wonderful effect that you can produce if you shine infrared light directly out in line with the lens of a camera. All animal eyes (including human eyes) will bounce a portion of that light directly back into the camera, and you will see bright glowing dots where the eyes are: It is a version of the "red-eye" effect in family snapshots taken with flashbulbs.

If you took a high-contrast infrared motion picture of an audience watching a film, placing the camera on stage and aligning the light source directly with the camera, you would see a galaxy of these dots against a field of black. And when someone in the audience blinked, you would see a momentary interruption in a pair of these dots.

If it were true, if there *were* times when those thousand dots winked more or less in unison, the filmmaker would have an extremely powerful tool at his disposal. Coherent blinking would be a strong indication that the audience was thinking together, and that the film was working. But when the blinking became scattered, it would indicate that he may have lost his audience, that they had begun to think about where to go for dinner, or whether their car was parked in a safe place, etc.

When people are deeply "in" a film, you'll notice that nobody coughs at certain moments, even though they may have a cold. If the coughing were purely an autonomic response to smoke or congestion, it would be randomly constant, no matter what was happening on screen. But the audience holds back at certain moments, and I'm suggesting that blinking is something like coughing in this sense. There is a famous live recording of pianist Sviatoslav Richter playing Mussorgsky's *Pictures at an Exhibition* during a flu epidemic in Bulgaria many years ago. It is just as plain as day what's going on: While he was playing certain passages, no one coughed. At those moments, he was able to suppress, with his artistry, the coughing impulse of 1,500 sick people.

I think this subconscious attention to the blink is also something that you would probably find as a hidden factor in everyday life. One thing that may make you nervous about a particular person is that you feel, without knowing it, that his blinking is wrong. "He's blinking too much" or "He's not blinking enough" or "He's blinking at the wrong time." Which means he is not really listening to you, thinking along with you.

Whereas somebody who is really focused on what you are saying will blink at the “right” places at the “right” rate, and you will feel comfortable in this person’s presence. I think we know these things intuitively, subconsciously, without having to be told, and I wouldn’t be surprised to find that it is part of our built-in strategy for dealing with each other.

When we suggest that someone is a bad actor, we are certainly not saying that he is a bad human being; we are just saying that this person is not as fully *in* the character as he wants us to believe, and he’s nervous about it. You can see this clearly in political campaigns, where there is sometimes a vivid distinction between who somebody is and who they want the voters to believe they are: Something will always be “wrong” with the rate and moment that these people blink.

That brings me back to one of the central responsibilities of the editor, which is to establish an interesting, coherent rhythm of emotion and thought—on the tiniest and the largest scales—that allows the audience to trust, to give themselves to the film. Without their knowing why, a poorly edited film will cause the audience to hold back, unconsciously saying to themselves, “There’s something scattered and nervous about the way the film is thinking, the way it presents itself. I don’t want to think that way; therefore, I’m not going to give as much of myself to the film as I might.” Whereas a good film that is well-edited seems like an exciting extension and elaboration of the audience’s own feelings and thoughts, and they will therefore give themselves to it, as it gives itself to them.

AFTER WORD

DIGITAL
FILM EDITING

Past, Present, and Imagined Future