

3. The Viewer's Activity

The theories of filmic narration discussed in the last two chapters have little to say about the spectator, except that he or she is relatively passive. Perspectival accounts tend to treat the viewer pointillistically, as the sum total of ideal vantage points shifting from shot to shot. But this observer is merely the "all-embracing eye" of the Wodehouse passage cited at the start of Part 1: mimetic theories assign few *mental* properties to the spectator. Of mimetic theorists, only Eisenstein allows the spectator an interesting mental life, including features such as expectation and some powers of inference. Diegetic theories, for all their apparent concern with narrational effects, also tend to downplay the viewer's role. In keeping with the revision of Benveniste whereby *énonciation* gets reduced to marks of the speaker, enunciation theorists have notably ignored the spectator. When the perceiver is discussed, it is usually as the victim or dupe of narrational illusion-making. MacCabe's "metalanguage" and Metz's *discours-disguised-as-histoire* fool the viewer into taking the narration for an unmediated and "natural" representation. The passivity of the spectator in diegetic theories generally is suggested not only by the extensive borrowing of mimetic concepts of narration but also by the use of terms like the "position" or the "place" of the subject. Such metaphors lead us to conceive of the perceiver as backed into a corner by conventions of perspective, editing, narrative point of view, and psychic unity. A film, I shall suggest, does not "position" anybody. A film cues the spectator to execute a definable variety of operations.

I have no trust that starting a theory of narration with an account of the spectator's work will explain all narrational phenomena. But since the spectator's comprehension of the story is the principal aim of narration, we can usefully begin there. And by not giving the film an immediate pride of place, we can seek to avoid passive notions of viewing. But first a few disclaimers.

I shall try to explain the formal conditions under which we comprehend a film. This means that here the "spectator" is not a particular person, not even me. Nor is the spectator an "ideal reader," which in recent reader-response criticism tends to be the most fully equipped perceiver the text could imagine, the one most adequate to all the aspects of meaning presented. I adopt the term "viewer" or "spectator" to name a hypothetical entity executing the operations relevant to constructing a story out of the film's representation. My spectator, then, acts according to the protocols of story comprehension which this and following chapters will spell out. Insofar as an empirical viewer makes sense of the story, his or her activities coincide with the process I will be describing. For the comprehension of any one narrative film, of course, the "hollow" forms I will be describing must be supplemented by many sorts of particular knowledge. Moreover, my spectator is "real" in at least the sense that she or he possesses certain psychological limitations that real spectators also possess. My spectator, for instance, undergoes the *phi* phenomenon (see p. 32) and thus necessarily perceives apparent motion in films. Finally, my spectator is active; his or her experience is cued by the text, according to intersubjective protocols that may vary (as we shall see in Part 3).

The theory I advance attends to the perceptual and cognitive aspects of film viewing. While I do not deny the usefulness of psychoanalytic approaches to the spectator, I see no reason to claim for the unconscious any activities which can be explained on other grounds. In general, current film theory has underestimated the importance of the spectator's conscious and preconscious work. Study of narrative cognition may in fact be a prelude to psychoanalytic inquiry for the same reason that Freud was at pains to show that psychoanalytic theory finds its best application when cognitive explanations fall short.

As a perceptual-cognitive account, this theory does not address affective features of film viewing. This is not because I think that emotion is irrelevant to our experience of cinematic storytelling—far from it—but because I am concerned with the aspects of viewing that lead to constructing the story and its world. I am assuming that a spectator's comprehension of the films' narrative is theoretically separable from his or her emotional responses. (I suspect that psychoanalytic models may be well suited for explaining emotional aspects of film viewing.) Near the end of this chapter I will suggest how some affective features might be explicable in relation to cognitive processes.

It will come as no surprise that I do not treat the spectator's operations as necessarily modeled upon linguistic activities. I shall not speak of the spectator's "enunciating" the story as the film runs along, nor shall I assume that narrative sense is made according to the principles of metaphor and metonymy. It is by no means clearly established that human perception and cognition are fundamentally determined by the processes of natural language; indeed, much psychological evidence runs the other way, toward the view that language is an instrument of and guide for mental activity. For such reasons, I do not call the spectator's comprehension "reading" a film. It is, moreover, needlessly equivocal to speak of the spectator's activity as a "reading" when the same word is applied to the abstract propositional arguments characteristic of critical analysis and interpretation. Viewing is synoptic, tied to the time of the text's presentation, and literal; it does not require translation into verbal terms. Interpreting (reading) is disjunctive, free of the text's temporality, and symbolic; it relies upon propositional language. This chapter and this book try to explain viewing.

A Sketch for a Psychology of Filmic Perception and Cognition

Any theory of the spectator's activity must rest upon a general theory of perception and cognition. I assume here what is called a Constructivist theory of psychological activity; descended from Helmholtz, it has been the dominant view in

perceptual and cognitive psychology since the 1960s. According to Constructivist theory, perceiving and thinking are active, goal-oriented processes. (Karl Popper calls this the "searchlight" theory of mind.)² Sensory stimuli alone cannot determine a percept, since they are incomplete and ambiguous. The organism *constructs* a perceptual judgment on the basis of nonconscious *inferences*.

Inference making is a central notion in Constructivist psychology. In some cases, the inference proceeds principally "from the bottom up," in which conclusions are drawn on the basis of the perceptual input. Color perception is a good example. Other processes, such as the recognition of a familiar face, operate "from the top down." Here the organization of sensory data is primarily determined by expectation, background knowledge, problem-solving processes, and other cognitive operations. Both bottom-up and top-down processing are inferential in that perceptual "conclusions" about the stimulus are drawn, often inductively, on the basis of "premises" furnished by the data, by *internalized rules*, or by prior knowledge.³

A Constructivist theory permits no easy separation between perception and cognition. Speaking roughly, the typical act of perception is the identification of a three-dimensional world on the basis of cues. Perception becomes a process of active hypothesis-testing. The organism is tuned to pick up data from the environment. Perception tends to be anticipatory, framing more or less likely expectations about what is out there. As E. H. Gombrich puts it, "Groping comes before grasping, or seeking before seeing."⁴ The organism interrogates the environment for information which is then checked against the perceptual hypothesis. The hypothesis is thus either confirmed or disconfirmed; in the latter case, a fresh hypothesis tends to appear.⁵ Bottom-up perceptual processes, such as seeing a moving object, operate in a fast, involuntary way, but they remain similar to other inferential processes.⁶ Top-down processes are more overtly based on assumptions, expectations, and hypotheses. When we scan a crowd to look for a friend, the likelihoods furnished by context and past experience count for a good deal. Cognitive processes help frame and fix perceptual hypotheses by reckoning in probabilities weighted to the situation and to prior

knowledge. Typical cognitive activities, like sorting or remembering things, depend on inferential processes.

In all these activities, whether we call them perceptual or cognitive, organized clusters of knowledge guide our hypothesis making. These are called *schemata*. The mental image of a bird is a schema for visual recognition, and the concept of a well-formed sentence functions as a schema in speech perception.⁷ Schemata may be of various kinds—prototypes (the bird image, for instance), or templates (like filing systems), or procedural patterns (a skilled behavior, such as knowing how to ride a bicycle).⁸ We shall see shortly that schemata play an important part in story comprehension.

The dynamic nature of the Constructivist account makes it highly attractive. The perceiver in effect bets on what he or she takes to be the most likely perceptual hypothesis. Like all inferences, perceptual experience tends to be a little risky, capable of being challenged by fresh environmental situations and new schemata. After some interval, a perceptual hypothesis is confirmed or disconfirmed; if necessary, the organism shifts hypotheses or schemata. This cycle of perceptual-cognitive activity explains the ongoing, revisionist nature of perception. The theory also explains why perception is often a skilled, learned activity; as one constructs a wider repertoire of schemata, tests them against varying situations, and has them challenged by incoming data, one's perceptual and conceptual abilities become more supple and nuanced.

Visual perception has furnished the classic illustrations of Constructivist psychological theory. Taken as a purely sensory experience, seeing is a bewildering flutter of impressions. The eye fixates many times per minute, using short and fast movements (called *saccades*); the eye rotates to compensate for head and body movement; the eye trembles involuntarily; and most of the visual information we receive is peripheral anyhow. Yet we do not experience a flicker or smear of percepts. We see a stable world, smooth movements, constant patterns of light and dark. To the extent that seeing is a bottom-up process, the visual system is organized to make its inferences in an involuntary, virtually instantaneous manner. You "immediately" see a visual array as

consisting of objects distributed in three-dimensional space, and you cannot help seeing this. This automatic construction is also affected by schemata-driven processes that check hypotheses against incoming visual data. Julian Hochberg suggests that since only the fovea of the eye sees detail, the saccades purposefully explore the environment, guided by schemata that propose the most fruitful places to look. We assemble our visual world from successive glances which we constantly check against our reigning "cognitive maps." These maps tell us to ignore the eye's physiological tremor and to bring the most significant areas into foveal vision. The schemata also generate hypotheses about what we will see next.⁹ Seeing is thus not a passive absorption of stimuli. It is a constructive activity, involving very fast computations, stored concepts, and various purposes, expectations, and hypotheses.¹⁰ A comparable account can be provided for auditory perception.

No one has yet delineated a Constructivist theory of aesthetic activity, but its outlines look clear enough. The artwork is necessarily incomplete, needing to be unified and fleshed out by the active participation of the perceiver. To some extent, artworks exploit the automatic nature of bottom-up processing; in such cases, the work can create illusions. But art is also a domain of top-down procedures. The spectator brings to the artwork expectations and hypotheses born of schemata, those in turn being derived from everyday experience, other artworks, and so forth. The artwork sets limits on what the spectator does. Salient perceptual features and the overall form of the artwork function as both triggers and constraints. The artwork is made so as to encourage the application of certain schemata, even if those must eventually be discarded in the course of the perceiver's activity.

What, then, distinguishes aesthetic perception and cognition from the nonaesthetic variety? In our culture, aesthetic activity deploys such skills for nonpractical ends. In experiencing art, instead of focusing on the pragmatic results of perception, we turn our attention to the very process itself. What is nonconscious in everyday mental life becomes consciously attended to. Our schemata get shaped,

stretched, and transgressed; a delay in hypothesis confirmation can be prolonged for its own sake. And like all psychological activities, aesthetic activity has long-range effects. Art may reinforce, or modify, or even assault our normal perceptual-cognitive repertoire.

A Constructivist account would thus consider film viewing as a dynamic psychological process, manipulating a variety of factors:

1. *Perceptual capacities.* Cinema is a medium that depends upon two physiological deficiencies in our visual system. First, the retina is unable to follow rapidly changing light intensities. At critical "fusion" frequency, more than fifty flashes per second will create the impression of steady light. Second, the phenomenon known as apparent motion occurs when the eye sees a string of displays as a single moving one.¹¹ This effect depends on the fact that the eye will infer movement from an intermittent input if the jumps are not too large. Flicker fusion and apparent motion illustrate how automatic and mandatory bottom-up processing is: although we know that a film is only a stroboscopic display of fixed frames, we *cannot fail* to construct continuous light and movement. As a medium of illusion, cinema counts on our making "wrong" inferences.¹² Moreover, the theater situation helps control our miscalibration of the stimuli. Darkness reduces distracting visual information and isolates the film for our concentration. And when the perceiver is acclimated to lower light levels, fusion and apparent movement effects operate more strongly.¹³

Other bottom-up processes shape our perceptual experience, such as when we perceive color on the screen. Some tasks, such as constructing the fictional space on the basis of depth cues, involve both bottom-up and top-down processing. (See Chapter 7.)

2. *Prior knowledge and experience.* In watching a representational film, we draw on schemata derived from our transactions with the everyday world, with other artworks, and with other films. On the basis of these schemata, we make assumptions, erect expectations, and confirm or disconfirm hypotheses. Everything from recognizing objects and understanding dialogue to comprehending the film's

overall story utilizes previous knowledge. The rest of this book will consider various schemata specific to narrative comprehension.

3. *The material and structure of the film itself.* In narrative cinema, as we shall see in the next chapter, the film offers structures of information—a narrative system and a stylistic system. The narrative film is so made as to encourage the spectator to execute story-constructing activities. The film presents cues, patterns, and gaps that shape the viewer's application of schemata and the testing of hypotheses.

I have isolated these factors for convenience, but plainly they interact in any single case. Consider the role of time in film viewing. While watching a narrative film, the spectator takes as one goal the arranging of events in temporal sequence. Our prior commerce with narrative and the everyday world allows us to expect that events will occur in some determinate order, and in most films specific cues encourage us to treat each distinct action as following previously presented ones. If the narrative presents events out of chronological order, we must fall back on our ability to rearrange them according to schemata. But such films run the risk of confusing us. Moreover, cinema's viewing conditions add a constraint: under normal conditions, it is not possible to review stretches of a film as one can reread passages of prose. The relentless forward march of stimuli in a film puts an extra strain on the spectator's memory and inferential processes. A filmmaker who presents story events out of chronological order thus risks forcing the spectator to choose between reconstructing story order and losing track of current action. This is probably why most films avoid temporal reshufflings. But we have seen in recent decades that films with complex time patterns can supply audiences with new schemata or encourage them to see the film more than once. The history of film form can thus alter the perceiver's prior experiences.¹⁴ (This history is investigated in Part 3 of this book.)

In opposition to all passive notions of spectatorship, then, we should consider film viewing a complicated, even skilled, activity. Watching a movie may seem as effortless as riding a

bicycle, but both draw on a range of practiced acts. Here, perhaps, is the most significant relation between the spectator and the reader. We are accustomed to think of reading printed matter as automatic, but even after the language has been learned, reading is an immensely intricate achievement, requiring the selection of salient cues, the processing of large units, decisions about how to sample the text, anticipations, and the projection of an ongoing semantic whole.¹⁵ Comprehending a painting seems no less formidable. E. H. Gombrich has shown that the beholder needs a knowledge of the medium's constraints and conventions, a sense of the painting's purpose, the ability to fill in what is missing, and a proclivity to compare the painting with pertinent experiences of the world.¹⁶ It would be surprising if a film, with its mixtures of visual, auditory, and verbal stimuli, did not demand active and complex construction.

Narrative Comprehension

The point of the previous section is that the spectator *thinks*. To make sense of a narrative film, however, the viewer must do more than perceive movement, construe images and sounds as presenting a three-dimensional world, and understand oral or written language. The viewer must take as a central cognitive goal the construction of a more or less intelligible story. But what makes something a story? And what makes a story intelligible?

Since the early 1970s, several psychologists and linguists have sought to understand how people comprehend and recall stories.¹⁷ The research is still limited by its reductive assumptions, since the stories are simple, short, written in prose, and shorn of most aesthetic interest. Yet what data the researchers have discovered offer some pointers for theorizing. First, these studies have revealed that even five-year-old children in our culture recognize certain activities as characteristic of storytelling and story-following. Second, the patterns of comprehending and recalling a story are remarkably uniform for all age groups. People tacitly assume that a story is composed of discriminable events performed by certain

agents and linked by particular principles. People also share a sense of what is secondary to the story's point and what is essential to it. Third, and most significant from a Constructivist standpoint, people perform operations on a story. When information is missing, perceivers infer it or make guesses about it. When events are arranged out of temporal order, perceivers try to put those events in sequence. And people seek causal connections among events, both in anticipation and in retrospect.

Story research has posed but not yet solved two problems: cross-cultural comprehension, and the relation of learning to innate abilities. It seems likely that, in non-Western cultures, following a story does not take the exact forms it does in ours.¹⁸ And it seems very likely that skill in story comprehension, however much it may operate with innate mental capacities (e.g., perception of time or causality), is acquired. One study found older children more proficient than younger ones in understanding stories with events put out of temporal order. This would support culture-based theories of perception and cognition, as well as late structuralist conceptions of narrative codes' dependence upon "already read" texts. Nonetheless, neither problem is of central importance for us here. For schooled perceivers in contemporary Western culture, narrative comprehension and recall are centrally guided by the goal of creating a meaningful story out of the material presented.

Generally, the spectator comes to the film already tuned, prepared to focus energies toward story construction and to apply sets of schemata derived from context and prior experience. This effort toward meaning involves an effort toward unity. Comprehending a narrative requires assigning it some coherence. At a local level, the viewer must grasp character relations, lines of dialogue, relations between shots, and so on. More broadly, the viewer must test the narrative information for consistency: does it hang together in a way we can identify? For instance, does a series of gestures, words, and manipulations of objects add up to the action sequence we know as "buying a loaf of bread"? The viewer also finds unity by looking for relevance, testing each event for its pertinence to the action which the film (or

scene, or character action) seems to be basically setting forth. Such general criteria direct perceptual activity through anticipations and hypotheses, and they are in turn modified by the data supplied by the film.

We can specify these schemata more exactly. In comprehending a narrative film, the spectator seeks to grasp the filmic continuum as a set of events occurring in defined settings and unified by principles of temporality and causation. To understand a film's story is to grasp what happens and where, when, and why it happens. Thus any schemata for events, locations, time, and cause/effect may become pertinent to making sense of a narrative film. More rigorously, we can follow Reid Hastie in distinguishing among various types of schemata; each has a role to play in narrative comprehension.¹⁹

"Central-tendency" or *prototype* schemata, Hastie suggests, involve identifying individual members of a class according to some posited norm. In narrative comprehension, prototype schemata seem most relevant for identifying individual agents, actions, goals, and locales. Understanding *Bonnie and Clyde* involves applying prototypes of "lovers," "bank robbery," "small Southern town," and "Depression era." We cannot inventory all the possible prototype schemata that might be pertinent to narrative comprehension; each film will call on a particular configuration of them.

More useful for our purposes is the tendency for such prototypes to operate in a larger structure. Hastie calls such structures *template* schemata, or filing systems. Template schemata can add information when it is absent and test for proper classification of data. The early results of story-comprehension research suggest that in our culture perceivers do tend to presuppose a particular master schema, an abstraction of narrative structure which embodies typical expectations about how to classify events and relate parts to the whole. Perceivers tend to use this master schema as a framework for understanding, recalling, and summarizing a particular narrative. The perceiver expects each event to be discriminable and to occur in an identifiable locale. The string of events should reveal chronological order and linear causality. (For perceivers of all ages, texts with reordered

story events or ambiguous causal connections tend to reduce understanding.) Causal connections are especially important; in remembering stories, people tend to invert the order of events more frequently when the link is only sequential ("and then . . .") and not also consequential ("as a result . . .").

Several experiments yield evidence for the schematic function of a "template" of narrative structure in contemporary Western cultures. The perceiver tends to recall a deviant story as being more normal than it was. If the text as presented omits causal connections, perceivers tend to supply them when retelling the tale. This is also strong evidence for the active qualities of narrative understanding: spectators are filling in material, extrapolating and adjusting what they remember. Perceivers also agree about what can be deleted in summarizing a story.²⁰ And adults have developed strategies to deal with deviations from the master schema, the chief one of which seems to be a toleration of ambiguity. In such ways, basic structural principles continue to serve as reference points for the identification of "less intelligible" narratives. The narrative schema is like those circles, squares, and triangles which artists revise and adorn to permit the portrayal of any object; the perceiver constantly refines the basic schema to fit the narrative at hand.

Nearly all story-comprehension researchers agree that the most common template structure can be articulated as a "canonical" story format, something like this: introduction of setting and characters—explanation of a state of affairs—complicating action—ensuing events—outcome—ending. Distortions in comprehension and recall tend to occur at points when the narrative violates or ambiguates this ideal scenario. There is further evidence that goal orientation is a salient aspect of the schema of causality. One researcher found that comprehension and memory are best when the story conformed to the drive-to-a-goal pattern. When the goal was stated at the end of the tale, comprehension and recall were significantly poorer, but still not so poor as when the goal of the action was never stated. In other words, early statement of the protagonist's goal permitted the perceiver to

fill in causal and temporal connections more exactly. And in recalling and summarizing stories which stated the goal tardily, perceivers inserted goal statements earlier in the sequence. Thus the format can be recast as: setting plus characters—goal—attempts—outcome—resolution.²¹

Now, the canonic story formats have a familiar ring; they resemble the formulas for plot construction dear to Freytag, Brunetière, and a hundred lesser theorists of narrative construction. We are thus inclined to suspect a strong cultural bias at work in the experiments' descriptions of story formats. Would an African perceiver necessarily grasp a tale in terms of exposition/complication/outcome? I do not wish to dismiss the possibility of a cross-cultural canonic story, since perhaps at some level of description these formats hold good. But since the matter needs much more study, I suggest only that the formats have heuristic value for analyzing narratives produced and consumed in our culture. The contemporary Western perceiver does typically expect expository material at the outset, a state of affairs disturbed by a complication, and some character ready to function as a goal-oriented protagonist.

As a template for organizing causality and time, the canonic story accords well with a Constructivist theory of narrative cognition. The perceiver gauges how well the narrative at hand can be slotted into the schema. In most cases the perceiver does not patiently isolate each datum (in a film, each movement or shot or sound) and slowly assemble a narrative event such as "buying a loaf of bread." Instead, the spectator selects salient cues, then draws on prototype schemata (e.g., knowledge of bakeries and hunger) and template schemata specific to narrative structure (e.g., what is likely to be causally prominent). Guided by something like the canonic story, the perceiver "chunks" the film into more or less structurally significant episodes. Only some such process can explain how the perceiver understands that very different forms of surface information convey similar meanings. In a film, buying a loaf of bread might consume an instant, a scene, or several scenes. As Roland Barthes remarks, "To read a narrative continuum is in fact to arrange it—at the quick pace set by the reading material—in

a variety of structures, to strive for concepts or labels which more or less sum up the profuse sequence of observations."²² Furthermore, the top-down grasping of an event can run ahead of the data. If a hungry man enters a bakery, "buying a loaf of bread" becomes a likely hypothesis before he places his order. The categorical whole can precede the perceived part.

Prototype schemata and template schemata are employed by what Hastie calls *procedural* schemata, those operational protocols which dynamically acquire and organize information. This is perhaps most clearly seen when a template schema is inadequate for the task at hand. If the film does not correspond to the canonic story, the spectator must adjust his or her expectations and posit, however tentatively, new explanations for what is presented. Even in a more predictable narrative, the search for information and the framing of inferences follow characteristic rules. Some of these procedures were identified long ago by the Russian Formalists under the name of "motivation."²³ By what procedures does the spectator justify a given textual element? How does the element get assigned to a prototype or classified within a pertinent grid?

The spectator may justify material in terms of its relevance to story necessity. We can call this *compositional* motivation. But the spectator may apply other warrants, such as a notion of plausibility derived from some conception of the way things work in the world. "He is the sort of man who would do that": such a rationale exemplifies *realistic* motivation. Or the spectator may justify an expectation or inference on *transtextual* grounds. The clearest case is that of genre: in a Western, we expect to see gunfights, barroom brawls, and thundering hooves even if they are neither realistically introduced nor causally necessary. Finally, and most rarely, the perceiver can decide that something is present simply for its own sake—as an appealing or shocking or neutral element. The Formalists called this *artistic* motivation and thought very highly of it, since it directly focused attention on the forms and materials of the artwork.

In practice, the first three of these procedural rationales often cooperate with one another. If Marlene Dietrich sings a cabaret song, we could justify it compositionally (it's here

that the hero meets her), realistically (she plays a cabaret singer), and transtextually (Marlene sings such songs in many of her films; it's one aspect of her star persona). Sometimes too there may be disparity among these rationales, as when we consider a stray "realistic" detail as having no bearing on the unfolding action, or when in the musical genre the causal chain halts to make way for a song and dance. Most films ask the spectator to employ compositional and transtextual motivation. Realistic motivation is usually a supplementary factor, reinforcing expectations already arrived at on other grounds. Artistic motivation is a residual category and remains distinct from the others; the spectator has recourse to it only when the other sorts do not apply. Thus, the concept of motivation brings to light several procedural schemata which the spectator must actively employ.

To what extent, we might now ask, does the viewer possess *stylistic* schemata? Most narratives use their medium (language, film, graphic art, whatever) as a vehicle for narrative information. Perceivers' schemata thus tend to favor narrative patterning and to find purely stylistic patterns difficult to notice or recall. Teun van Dijk writes of the literary text: "Our memory and processing resources are able only in a very restricted way to store and retrieve these kinds of surface structural information, even if the communicative conventions require specific attention to such structures."²⁴ The same often holds true for visual phenomena: after one checks a clock, one can recall the time but not the shape of the numerals, even though those must have been registered at some stage of perception.²⁵ This suggests that when spectators are confronted with a film that emphasizes its stylistic features, they will still seek cues for constructing a story.

Film style may usually go unnoticed, but that does not entail that the spectator has no stylistic schemata. On a Constructivist account, the perceiver need be no more aware of applying an aesthetic convention than of any other cognitive operation. We have already seen that many perceptual processes are nonconscious. Perhaps owing to the stylistic uniformity of mainstream cinema, applying stylistic schemata is a top-down process that has become so practiced as to operate automatically. It is clear that on the basis of prior

experience, the spectator assumes that certain stylistic schemata will be adhered to, as when we identify a long shot or a handheld commentary on the basis of prototypes. We also employ stylistic templates. In mainstream narrative cinema, a long shot is likely to be followed by a closer view, and a musical bridge is more apt to fade out than to be cut off. Some stylistic alternatives are unlikely and some are completely ruled out. We also know that spectators accustomed to one stylistic tradition can use procedural schemata to comprehend other stylistic options (e.g., "motivate this cut by story necessity"). A Constructivist theory would emphasize that to a great degree spectators can learn to notice and recall stylistic features of any film. At later points in this book we will have occasion to consider how film style can operate as a vehicle for narration and a system in its own right.

What does the viewer do with the schemata? Plainly, many cognitive activities are performed in making sense out of narrative. The viewer posits a more or less stable set of assumptions. The spectator assumes, for example, that objects and human beings persist in space even when they are not on screen; that a character possesses the same individual identity on successive appearances; that a film in English will not suddenly lapse into Urdu. We notice such basic assumptions only when a film violates them, as when the same character is played by two quite different performers (as in Buñuel's *Obscure Object of Desire*). The viewer also makes many inferences. If our hero bursts into tears, we conclude that he is sad. Like other inductive inferences, such conclusions are open-ended, probabilistic, and subject to correction. Maybe our hero is delighted. *Memory* of course plays a role as well. Again, memory must be seen not as a simple reproduction of a prior perception, but as an act of construction, guided by schemata (as was proposed by Frederic Bartlett over fifty years ago).²⁶ There is also the cognitive task of *hypothesizing*: the spectator frames and tests expectations about upcoming story information. Since hypotheses exemplify the anticipatory quality of schema-driven perception, I shall try to clarify their operation at a little more length.

Alan Sternberg is one of the few theorists of narrative to give due weight to the process of framing and testing

hypotheses. In Sternberg's theory, the pattern of story information withheld in the work prompts the perceiver to make hypotheses of various sorts. A hypothesis may pertain to past action that the text refrains from specifying; Sternberg calls this a *curiosity* hypothesis. By contrast, a *suspense* hypothesis is one that sets up anticipations about forthcoming events. Hypotheses may also be more or less *probable*, ranging from the highly likely to the flatly improbable, and more or less *exclusive*, ranging from either/or choices to mixed sets. And since hypotheses arise in the course of time, they may be held simultaneously or successively, as when one hypothesis simply replaces another.²⁷ All of these categories will prove of use when we seek to analyze how particular films both cue and constrain the spectator's activity.

There are levels of hypothesis testing. Typically, assumptions and inferences take care of the "microscopic," moment-by-moment processing of the action, but at critical junctures we are tuned to expect particular events. Across scenes, hypotheses emerge with some clarity: will the character do *x* or *y*? A more indefinite but highly significant arc of "macroexpectation" may extend across a whole film. The narrative itself can inflect these levels, such as by playing down small-scale portions (transitions, or the secondary actions Barthes calls "catalyses") in order to stress longer-range hypotheses about significant action sequences (Barthes's "kernels" or "hinges").²⁸ Other narratives may deny us large-scale expectations; in such a case we call the text episodic. Hypotheses also vary in precision according to their placement in the text: they tend to be more "open" at the start of a text; some remain tacitly in force throughout because they are never countermanded.

The primary focus of hypothesis forming remains what Sternberg calls *suspense*—anticipating and weighing the probabilities of future narrative events. Consider the alternatives. When an action backs up an already confirmed hypothesis, it is redundant and cannot trigger the full anticipatory range of hypothesis casting. "Our whole sensory apparatus," writes Gombrich, "is basically tuned to the monitoring of unexpected change. Continuity fails to register after a time, and this is true both on the physiological and the psychological level."²⁹ But when the action presented

does not confirm a hypothesis, the perceiver is unlikely to turn back to a previous page or halt the teller's tale or stop the film to sort it all out. Instead, the perceiver presses on to see if the unfolding action will explain or modify the challenge to the hypothesis. We can call this the "wait-and-see" strategy.³⁰ Even in the detective story, which of all genres places most emphasis upon curiosity about prior events, the perceiver seldom pauses if there are snags in answering questions about why the action occurs or what the action shows about the agent. Instead, the perceiver rushes eagerly on. The detective tale offers the story of an investigation, and what the perceiver wants to know is not only "who did it" but how the detective's future actions will bring the solution to light.

The perceiver's tendency to make salient those hypotheses pertaining to upcoming information accords well with the Constructivist assumption that schemata coax us to anticipate and extrapolate. Disconfirming instances make us readjust our expectations. At the same time, the schemata need a firm foothold somewhere. The sequential nature of narrative makes the initial portions of a text crucial for the establishment of hypotheses. Sternberg borrows a term from cognitive psychology, the "primacy effect," to describe how initial information establishes "a frame of reference to which subsequent information [is] subordinated as far as possible."³¹ A character initially described as virtuous will tend to be considered so even in the face of some contrary evidence; the initial hypothesis will be qualified but not demolished unless very strong evidence is brought forward.

Between the framing of the hypothesis and the confirming or disconfirming instance in the text there intervenes some time. At a local level, the hypothesis may be validated or invalidated very soon, which is the normal case for our moment-by-moment following of the narrative. Our heroine crosses a room and opens the door; cut to a shot from outside as she walks out into the hall. But when more idiosyncratic and macrostructurally significant narrative action is at stake, the information is typically withheld for some interval. If our heroine is to encounter someone in the hall, perhaps that person's motive or purpose will be withheld from us and

revealed only later. Since the pioneering explorations of the Russian Formalists, we have recognized *retardation* as essential to narrative structure. Because the narrative unwinds in time, the fulfillment of our expectations may be considerably delayed. The narrative begins, but its forward progress must be interrupted by exposition that supplies pertinent background information. The narrative will end, but its conclusion is held back by complications, subplots, or digressions. Viktor Shklovsky called retardation "stairstep construction": a narrative text is less like an elevator than a spiral staircase which, littered with toys, dog leashes, and open umbrellas, impedes our progress.³² Sternberg goes so far as to describe the narrative text as "a dynamic system of competing and mutually blocking retardatory patterns."³³ Of course, delaying story information will not keep the spectator interested forever; our blocked expectations must be balanced by more immediate ones—usually, suspense hypotheses. In any event, delay in satisfying hypotheses can be exploited to trigger new expectations or to play off different retardatory structures against one another. If, after our heroine goes out her door, we cut to a shot of the hero across town, menaced by the killer whom the heroine seeks, we start to elaborate a hypothesis that she will come to rescue him. And if the landlord holds her palaver at her door, the retardation of each line of action by the other will intensify our interest in finding out whether our hypothesis holds good.

So ongoing and insistent is the perceiver's drive to anticipate narrative information that a confirmed hypothesis easily becomes a tacit assumption, the ground for further hypotheses. But it is still convenient to speak of the perceiver's hypotheses as being clearly validated, invalidated, or left dangling. Most are validated: no narrative surprises us at every turn. Less often, the presentation of story information leads us to disqualify some hypotheses. The character does something we thought unlikely; the chief murder suspect turns out to be the next victim; the narrative swerves off the track. Narratives are composed in order to reward, modify, frustrate, or defeat the perceiver's search for coherence.

To sum up: In our culture, the perceiver of a narrative film

comes armed and active to the task. She or he takes as a central goal the carving out of an intelligible story. To do this, the perceiver applies narrative schemata which define narrative events and unify them by principles of causality, time, and space. Prototypical story components and the structural schema of the "canonical story" assist in this effort to organize the material presented. In the course of constructing the story the perceiver uses schemata and incoming cues to make assumptions, draw inferences about current story events, and frame and test hypotheses about prior and upcoming events. Often some inferences must be revised and some hypotheses will have to be suspended while the narrative delays payoff. While hypotheses undergo constant modification, we can isolate critical moments when some are clearly confirmed, disconfirmed, or left open. In any empirical case, this whole process takes place within the terms set by the narrative itself, the spectator's perceptual-cognitive equipment, the circumstances of reception, and prior experience.

Narratives in the aesthetic realm submit all these processes to particular pressures. In ordinary perception, Ulric Neisser points out, perceptual hypotheses tend to be vague and open-ended, and they are seldom vigorously disconfirmed.³⁴ In art, however, alternative hypotheses tend to be much more explicitly defined, their set tends to be closed, and they get challenged fairly often. Aesthetic perception and cognition encounter situations which are constructed so as to upset the most common assumptions, the most valid inferences, the most probable hypotheses, and the most appropriate schemata. Narratives may arouse perceivers' anticipations and lock in mental sets before presenting information that undermines those very activities. More often than we are usually aware, narratives invoke expectations only to defeat them, plan and time our encounters with information that will upset our assumptions, encourage us to extrapolate and then chide us for going too far, parade a host of positive instances before trotting out the single and crucial exception, hold back basic data while "prattling" (Barthes's term) about irrelevancies—all the while forcing us to keep to a predetermined temporal se-

quence and (in film and other "time arts") yoking us to a fixed rate of comprehension that makes us err simply by pressure of the clock. Narrative art ruthlessly exploits the tentative, probabilistic nature of mental activity.

Although I am delineating a hypothetical spectator endowed with aspects of human perceptual and cognitive capacities, I might mention the extent to which empirical spectators can err in comprehending a film. These errors can be explained within the theory I have outlined. As I mentioned earlier, a narrative film both triggers and constrains the formation of hypotheses and inferences; it does not uniquely specify or determine them. Physiological factors such as fatigue may thwart or delay perceptual processing, making attention or memory flag. Other misunderstandings of the film may result from applying schemata which later prove to be inappropriate—such as mistaking the hero's goal or not expecting temporal reordering in a flashback. Errors of schemata selection and hypothesis forming may also spring from inadequate knowledge of the narrational norms to which the film appeals; a spectator who lacks the schemata for the 1960s "art film" will miss cues and lose patterning in a film like *8½*. Of course, a spectator may also deliberately misconstrue the story, as some French Surrealists did in their "irrational enlargements" of Hollywood B-pictures. Finally, as I have suggested, a film may contain cues and structures that encourage the viewer to make errors of comprehension; in such cases, the film "wants" a short- or long-term "misunderstanding."

To the extent that this theory showcases perception and cognition, it does not have much to say about affect. The theories of affect most compatible with the Constructivist account are those in which emotion is bound up with expectation and its interrupted or delayed fulfillment.³⁵ It should be evident that emotion is not at all alien to the process of filmic comprehension. When we bet on a hypothesis, especially under the pressure of time, confirmation can carry an emotional kick; the organism enjoys creating unity. When the narrative delays satisfying an expectation, the withholding of knowledge can arouse keener interest. When a hypothesis is disconfirmed, the setback can spur the

viewer to new bursts of activity. The mixture of anticipation, fulfillment, and blocked or retarded or twisted consequences can exercise great emotional power. The formal processes of perception and cognition—as Eisenstein well knew—can trigger affect.

Believing and Seeing

Rear Window (1954) has long been used as a small-scale model of the spectator's activity. The chairbound photographer watching without being seen; the windows across the way like movie screens; the apparent freedom from consequences yielded by the viewer's distant vantage point—these have made the film an irresistible analogy for the viewing experience. Such atomistic comparisons are valid; no film I know fits more snugly into a perspectival theory of narration. Yet *Rear Window* can also be used to exhibit the full complexity of the viewing activity. Not just the piquant situation but the very process of the unfolding action lays bare the way we typically construct the story in a fiction film. *Rear Window* is at once typical in the job it hands the spectator and extraordinary in the explicitness with which tasks are spelled out.

The story is well known. Laid up with a broken leg, L. B. Jeffries takes to watching his neighbors across the courtyard. He idly looks in on a married couple and their dog, a composer, a pair of newlyweds, a lonely middle-aged woman, a young female dancer, a sculptress, and the Thorwalds, another married couple. At the moment, Jeff's love affair with the model Lisa Fremont is deteriorating because he sees her as too unlike him: she leads a sophisticated Manhattan life while he tramps around the world pursuing adventure. In his days of rear-window spying, Jeff begins to suspect that Lars Thorwald has murdered his wife. Jeff wins Lisa to his way of thinking, but his friend Detective Lieutenant Doyle remains skeptical. After Lisa breaks into Thorwald's apartment for evidence, Thorwald realizes that Jeff has been watching him and comes to assault him. Jeff is saved in time (although he breaks his other leg), and the film ends with Jeff and Lisa apparently reconciled.

There are thus two interdependent lines of action: the mystery (What has happened to Mrs. Thorwald? Will Jeff prove that she was killed? Can Lisa find evidence?) and the romance (Will Jeff and Lisa break up?). What makes the film useful to us here is its blending of explicit and implicit appeals to spectatorial activity. The mystery plot line tends to state assumptions, hypotheses, evidence, and inferences very baldly. Characters talk about how they build up expectations, select information, and draw conclusions. The spectator shares in the characters' battle for a solution to the mystery. At certain moments in the mystery action, however, the viewer is tossed information which the characters do not possess. Thus we must sometimes qualify or challenge the characters' problem-solving processes. And in the romance plot line, the spectator must construct the story without benefit of such overt cues. For example, the spectator must discover that the absolute opposition which Jeff creates—Lisa and urban boredom versus himself and high adventure—is a false one; one can find adventure in the most mundane urban milieu.

The overall structure of the film confirms expectations grounded in the canonic story format. There is exposition which establishes time, setting, and character relations and which identifies the two primary states of affairs: life in the courtyard and the impasse in Jeff and Lisa's romance. There follows Jeff's decision to achieve a goal—to prove Mrs. Thorwald was murdered—which makes progress but which also encounters several obstacles (counterevidence, absence of tangible proof). Finally Jeff achieves the goal (resolution), and he and Lisa are seen at relative peace at the end.

The film's time scheme also encourages us to unify the action. The story begins on Wednesday and reaches its climax on Saturday, with each day carefully demarcated and a clear rhythm of alternation established between day and night. The temporal unity is strengthened by a set of dead lines: Thorwald may leave before Jeff can prove the murder took place, Thorwald may harm Lisa before the police arrive, Thorwald may kill Jeff before the police rescue him. The viewer thus constructs the story using information not only about causality but about the duration and deadlines that govern the action. Furthermore, the prologue and epilogue

are symmetrical, not only spatially (similar camera movements around the courtyard and Jeff's apartment) but in their invitation to compare present with past (e.g., the newlyweds are bickering at the end, Miss Lonelyhearts now knows the composer). The viewer can also quickly grasp the spatial premise: we will be confined to this room (with some exceptions), we will see what Jeff sees (with significant exceptions). Thus we will also have to call on stylistic schemata, such as our knowledge of eyeline-match and point-of-view editing.

From the outset, *Rear Window* emphasizes the probabilistic quality of seeing and knowing. The tour of the courtyard, with the camera panning slowly left to right from within Jeff's window, provides information that is principally visual (supplemented by music and street noises). We see the neighbors composing music, practicing dance, and so on, and we infer that these are habitual and typical activities. When the camera moves into and around Jeff's apartment, our inferences must come thick and fast. Track left to reveal his leg in a cast. (How did the leg get injured?) Then to a shattered press camera. (He's a photographer. Did he break his leg on the job?) Then to photographs of an auto race, a fire, a battered woman, an atomic bomb. (He probably did; he takes hazardous assignments.) Then down a line of cameras and flashbulbs. (Almost certainly a professional photographer.) Finally to a framed negative of a model and to a stack of magazines with the photo on the cover. (He can also do fashion photography.) Yet these inferences remain only probable until the next scene. Now most of the inferences we had made get corroborated by Jeff's phone conversation with his editor. The pattern is set: this film will encourage us to construct a story on the basis of visual information (objects, behavior) and then confirm or disconfirm that construction through verbal comment. After the solitary woman across the way toasts her phantom dinner companion and then breaks down weeping, Jeff's remark, "Miss Lonelyhearts," caps our comprehension of her situation. After Thorwald dials three digits and pauses, Jeff will mutter: "Long distance." As we watch Miss Torso's party, Jeff and Lisa spin yarns about her choice of male companions. Encouraging us to draw conclusions about

story events and then making us wait for verbal corroboration highlights the more or less likely nature of our inferential procedures. Moreover, by having the characters sum up the actions we see by a clichéd label ("nagging wife," "Park Avenue model," "the eat-drink-and-be-merry girl"), the film makes sure that we construct the proper schematic prototypes.

Rear Window asks us to generate several distinct sorts of hypotheses. We frame curiosity hypotheses about the past (Thorwald did murder his wife), and suspense hypotheses about the future (the investigation will reveal more clues). In the first four scenes, the film confines itself chiefly to the romance plot and generates vague suspense hypotheses (Jeff and Lisa will likely remain at loggerheads). After Jeff hears the scream and crash at the close of scene 4, curiosity plays a stronger part; we and Jeff watch Thorwald's mysterious comings and goings until we are prepared to label the action as "domestic murder" and ask more specific questions. And the hypotheses alter with respect to probability. As of scene 4, Jeff and Lisa's romance looks unlikely, but as the murder draws them together, their eventual alliance becomes more probable. That Thorwald has committed murder seems unlikely in scene 4, more likely after Jeff pleads his case, less likely after Doyle's deflating evidence, more likely after the neighbors' dog is murdered, and virtually certain after Jeff discovers that something was buried in the garden.

The process whereby Jeff pursues the mystery displays quite beautifully the activity of the film spectator. He begins, like you and me, in a state of tuned anticipation: "If you don't pull me out of this swamp of boredom, I'm gonna do something drastic. . . . Right now I'd welcome trouble." Jeff is prepared to perceive, but he is so eager for activity that he risks extrapolating too much. In scene 4, after Lisa leaves, he turns to the window and scans the darkened apartment block across the way. Two sensory cues pop out—a scream and a crash. He looks quickly to and fro, but no visual cues are forthcoming. Now he can apply a very loose procedural schema: he will test for lapses in normal activity. As the night goes on, he watches Thorwald leave and return to his apartment. The next morning, he possesses only the rough-

est hypothesis: "I just can't figure it. . . . I think he was taking something out of the apartment." His nurse Stella quickly suggests the hypothesis that Thorwald is running out on his wife. But then Jeff picks up new information. Thorwald looks across the way in a peculiar fashion. "That's the kind of look a man gives when he's afraid someone might be watching him." In response, Jeff uses his telephoto lens to magnify his view. Thorwald opens a suitcase and wraps a knife and saw in newspaper. By the time Lisa arrives that night, Jeff is convinced that Mrs. Thorwald has been killed.

Note that Jeff's inferential process does not completely mirror our own. In an earlier scene, at six in the morning, we have seen Thorwald leave his apartment, accompanied by a woman dressed in black. But Jeff is asleep. We thus gain information that he does not have, and it creates a hypothesis that competes with the murder hypothesis. Everything that Jeff sees and will later see could also be explained by the hypothesis that the wife took a trip. But this is not to say that we reject Jeff's hypothesis completely. The murder hypothesis, however unlikely in real life, is highly likely in a Hitchcock film (transtextual motivation), and the fact that the film would equivocate about the woman's identity makes the viewer adopt a wait-and-see attitude. Perhaps this new information can be fitted into a more global explanation. (Lisa will show us how later.) For Hitchcock the division of knowledge between characters and audience is fundamental to suspense, which he describes as "providing the audience with information that the characters do not have."³⁶ This is only a half-truth, though. In most of his films, such data are never pure; they always come filtered through a narration that withholds crucial facts. Though we are more knowledgeable than the protagonists, we remain less knowledgeable than the narration, and the film rubs our noses in the partial and ambiguous quality of our information.

At this point Jeff sums up (Thorwald left late at night, wrapped up his knife and saw) and adds new data: he did not go to work, and he avoided his wife's bedroom. It falls to Lisa to resist Jeff's inference. She suggests that Mrs. Thorwald died; but there is no sign of a doctor or undertaker. Perhaps she's sleeping or sedated; but Jeff points out that she needs constant care. Yet would Thorwald murder his wife with the

shades up? Jeff counters that he's being nonchalant to give the impression of normality. And Lisa returns to Stella's point—Thorwald is leaving his wife—adding that indeed very few husbands commit murder. Lisa's rebuke is important in reemphasizing the tentative nature of Jeff's construction of the chain of events; her question "What is it you're looking for?" implies that Jeff has run far beyond the data. But Lisa's competing schemata are defeated when Thorwald finally raises the blinds to reveal the bed stripped, the wife gone, and a steamer trunk being packed. Lisa stands staring. "Let's start from the beginning again, Jeff. Tell me everything you saw—and what you think it means." Her remark concisely reiterates the film's strategy of supplying sensory information ("everything you saw") and then forcing Jeff (and us) to interpret it ("and what you think it means").

The next morning, the nurse Stella is already convinced of Jeff's account and freely pours forth specific hypotheses ("Now just where do you suppose he cut her up? Of course—the bathtub.") But Jeff's friend Lieutenant Doyle is skeptical, reminding him that there are many alternatives, of which murder is the least plausible. Hypothesizing that the wife took a trip, Doyle sets out to gather evidence. He reports to Jeff that witnesses saw the Thorwalds leave for the railroad station at 6 A.M. the previous morning; this corroborates our view of the couple leaving. Doyle also claims that he found a postcard from Mrs. Thorwald in her husband's mail. Jeff nonetheless resists this evidence and demands proof that she got on the train and that the trunk contains only clothes. Doyle grudgingly agrees to check further. Jeff's tenacity reveals the force of the primacy effect. The first full-blown inference, although less probable than Doyle's, takes pride of place, and subsequent hypotheses are flung against this one.

That evening there is a pitched battle of mutually exclusive hypotheses. Lisa rehabilitates Jeff's original opinion. First, she modifies his initial hypothesis in order to accommodate new evidence. Jeff has seen Thorwald take his wife's handbag from a drawer and go through her jewelry. Lisa (using the highly ideological "realistic" rationale, "What women typically do") asserts that no woman would leave home without her favorite handbag and jewelry. More im-

stant, she offers an alternative account of what we saw: "We'll agree they saw a woman, but she was not Mrs. Thorwald. At least, not yet." This removes the last obstacle to our sharing Jeff's inference. (It implies, however, that the film's narration cunningly concealed how and when the second woman got into Thorwald's apartment.) Yet Doyle destroys their revised hypothesis. He dismisses Lisa's schema as "woman's intuition" and trots out hard evidence that the Thorwalds have separated. He has the clichéd prototype: "That's what's called a family problem." The trunk contained clothes and was picked up by Mrs. Thorwald in Meritaville. Doyle turns Jeff's suspicions back on him (Has *he* owned a knife and saw? "Do you tell your landlord everything?") Most compellingly, Doyle points out that no hypothesis is airtight:

JEFF: You mean to say that you can explain everything that's gone on over there and is still going on?

DOYLE: No, and neither can you. That's a secret, private world you're looking into out there. People do a lot of things in private that they couldn't possibly explain in public.

After Doyle leaves, Jeff and Lisa watch Miss Lonelyhearts in an ugly and humiliating quarrel with a man she invites in. Jeff winces and grants that Doyle might be right about private behavior; Lisa agrees that they've become ghouls. Miss Lonelyhearts's episode confirms Doyle's hypothesis, and they accept defeat.

A defeat, however, which is only momentary. In the next scene, Jeff and Lisa are brought frantically to the window by (again) a scream. The neighbors' dog has been killed, and the wife shrieks out to the courtyard at large, denouncing the callousness of her neighbors. Crowds gather at windows and balconies to look down, but Thorwald does not emerge, remaining inside his darkened apartment, quietly smoking. Jeff switches back to his earlier explanation, assuming now that Thorwald killed the dog. But, Lisa asks, why? "Because it knew too much?" The next morning, Jeff uses slides he took of the courtyard to prove that something was buried in the garden. This fits with earlier data—the dog sniffing about the flower bed and Thorwald shooing it away. At this

point, the murder inference is virtually inescapable: it explains more than any other candidate. Later scenes will reconfirm it—Thorwald will respond to a blackmail threat, Lisa will find Mrs. Thorwald's wedding ring—but now the curiosity portion of the viewer's activity is largely finished. (The only question remaining pertains to the disposal of the wife's body.) The new problem is: how to get enough evidence to convince the authorities before Thorwald flees? By the time Jeff makes the blackmail call, he can read his prey's mind as he watches: "Go on and pick it up, Thorwald. You're curious. You're wondering if it's your girlfriend calling you up—the one you killed for." This is the culmination of that story construction we saw in a more rudimentary form earlier: visual cues anchored by verbal interpretation.

If the mystery line of action follows fairly closely the characters' explicit formulations (with the crucial qualifications provided by our privileged view), the romance plot solicits the same sort of viewing activities more tacitly. In early scenes, Jeff reiterates the impossibility of any alliance between himself and Lisa: he mocks marriage, resists Stella's advice, and insists that neither he nor Lisa can change. He employs detailed contrasts to show that with respect to food, clothes, locale, taste, and safety his and Lisa's worlds are mutually exclusive. His rude remarks travesty Lisa's catered dinner and bring her to the edge of breaking off their romance. But he holds back: "Couldn't we just keep things status quo?" "Without any future?" she asks. But in the end, she agrees to see him again tomorrow night. This portion of the film corresponds to the "state of affairs" phase of our canonical story, and what it portrays is a stalemate. It is at this point that Jeff hears the scream that initiates the murder plot.

Jeff tends to utilize schemata that are strictly homogeneous: Thorwald killed/did not kill his wife; Lisa and Jeff are exactly alike/diametrically opposite. One of Doyle's telling points about the putative murder is that not everything can be neatly filed: "People do lots of things in private that they couldn't possibly explain in public." He proposes a more heterogeneous schema. In the romance plot, it is the viewer's duty to be alert for such heterogeneous categories. Jeff's sharp distinctions get blurred when Lisa develops a

taste for danger. She eagerly reports the name and address of the mysterious couple; she mulls over the case while at work; she delivers the threatening note to Thorwald and eventually breaks into his apartment. At each step, Jeff warms to her. After she supplies an identity for Thorwald's woman companion, he kisses her. When she returns breathless from her foray across the court, he grins proudly at her enthusiasm. Calling Doyle on the phone, Jeff boasts of her resourcefulness in saving herself from Thorwald. Lisa demonstrates that she can handle the adventurous life that Jeff had believed was foreign to her. From a structural standpoint, the murder plot resolves the romance plot by breaking down Jeff's rigid dichotomy and revealing Lisa's wit and courage.

Whatever the cues in this film, our expectations are funded by knowledge of other films in the tradition. We motivate transtextually. Just as our experience of thrillers leads us to expect that Jeff is right about the murder, our knowledge of narrative conventions lets us expect that the romantic problems will be resolved before the end. They are, but somewhat equivocally. The camera surveys the courtyard in a movement parallel to that of the first scene, revealing the composer and Miss Lonelyhearts together, painters cleaning Thorwald's apartment, the older couple with a new dog, Miss Torso's boyfriend/husband returning from the army, and the newlyweds bickering. The camera then moves inside to reveal Jeff asleep in his chair, now with two broken legs. Pan right to Lisa on his daybed, curled up in a blouse and blue jeans reading *Beyond the High Himalayas*. The joke depends partly on the echo with the glamorous woman on the magazine cover in the first scene; Lisa the model has become Lisa the mountaineer. But then she picks up *Harper's Bazaar* and begins thumbing through it. The spectator can draw either of two conclusions: Lisa has bridged the gap between the two worlds, or she has reinstated the old opposition by revealing that her concession to the rugged life is superficial and provisional. It is not a question here of interpreting the film one way or the other. The throwaway nature of this last shot reiterates the game of visual information and alternative hypotheses that has in-

formed the whole film. After all the demands on the spectator to draw conclusions about murder and romance, the film ends with information that calls up exclusive and equally probable hypotheses about the stability of the couple without firmly backing either one.

Arrival at this ending depends upon careful retardation, or what Hitchcock preferred to call frustration. Jeff's ultimata and the stalemate in the fourth scene leave the romance plot hanging; its resolution is delayed by the murder plot. And the murder plot in turn must be seen as having its own retardatory structure. Jeff's murder hypothesis needs proving, so he searches for clues. The first is the possibility that Mrs. Thorwald's body is in the trunk which later turns out to contain nothing but clothes. The second clue is the patch of zinnias from which Thorwald shoos away the dog; after the dog's death, Jeff infers that there's something buried there, but when Stella and Lisa dig up the flower bed, they find nothing. The crucial evidence is a third clue, the wedding ring which Jeff sees Thorwald handling and which is the target of Lisa's break-and-entry. The wedding ring is the strongest evidence—it convinces Doyle to take a second look—but from a constructional standpoint the ring could have been the first clue Jeff spotted. The trunk and the zinnia patch are red herrings, pure retardatory devices. The same goes, in a more general way, for all the neighbors living their lives adjacent to Jeff and Thorwald. Apart from their considerable value as thematic parallels to the two primary couples, the neighbors often delay the unfolding of the mystery plot by distracting Jeff's (and our) attention from the main issue. Even at a very localized level, the film uses retardation to prolong our expectations. When Jeff writes Thorwald the threatening letter, the camera adopts a distant, high-angle framing so that we cannot see what he is writing; then the camera slowly cranes in, stretching out the gradual revelation of the simple question: "What have you done with her?" *Rear Window* perfectly illustrates Sternberg's point that a narrative can be considered "a dynamic system of competing and mutually blocking retardatory patterns."

This camera movement, like others, exemplifies how the viewer may also make inferences about the stylistic patterns

that emerge. The outstanding example in *Rear Window* is the handling of the death of the neighbors' dog. Jeff and Lisa have accepted Doyle's explanation for Mrs. Thorwald's absence. A scream from outside fetches them to the window. Up to this point, all shots have been confined to Jeff's apartment or to what one could see from it. Now, as Jeff and Lisa look, we get a shot from outside their window, looking in at them. People from other apartments peer out into the courtyard, and the freedom of spatial vantage points increases. The camera frames various neighbors in medium shots and from low angles, and it is clear that neither Jeff nor Lisa sees these people. After the neighbor cries out, "Which of you did it?" and denounces life on the courtyard, she and her husband haul the dead dog up to their window. Then we return to a framing of Jeff and Lisa, the camera tracks in on them, and they notice that only Thorwald, smoking in his darkened apartment, did not go to the window. They conclude that he killed the dog.

The scene violates one tacit stylistic assumption we have formed since the start—that we will be confined to what can be seen from Jeff's window. This violation is quite unpredictable (we might speak of stylistic as well as narrative "surprise" here), but it fulfills the narrative function of giving us a slight edge on Jeff and Lisa. The shocked reactions of the neighbors would not be visible from the window, and by showing them to us the narration tells us immediately that they did not kill the dog. Thorwald is thus guilty by elimination. It is left for Jeff and Lisa to draw the positive inference from his indifference to the tumult outside. In a sense, this scene is comparable to that in which the narration let us see the woman leave Thorwald's apartment while Jeff slept. Here, too, we are freed from our restriction to what Jeff knows. In that scene, though, space helped the narration mislead us: confined to Jeff's room, we got an equivocal view of the action. Here, the narration's closer framings do not try to bamboozle us. After this scene, the camera will not leave Jeff's apartment until the final struggle with Thorwald, and again the violation of the film's spatial norm will be justified by giving us information that Jeff doesn't have (the police are

arriving). In such scenes, film style not only guides our inductive processes, but becomes, in however general a way, one object of them.

Every film trains its spectator. By the last scenes of *Rear Window* we have honed our story-construction schemata to expect a climax. We are prepared to slot actions into a narrow "outcome" format. We have been explicitly encouraged to construct certain hypotheses, accepting some and resisting others. We have been trained to take purely visual cues as signs of a narrative situation. The telltale glimpse of Thorwald escorting the woman in black, and the romance plot as a whole, have accustomed us to diverging from characters' knowledge, playing their schemata and hypotheses against more heterogeneous and nuanced ones. We are prepared to justify events and motifs compositionally, realistically, and especially transtextually. We know that retardation will be employed to delay satisfaction of our expectations. Now, in a scene of continuous tension, the film makes maximal use of all the lessons it has taught.

As Stella and Jeff watch, Lisa delivers the threatening note to Thorwald's apartment; she is nearly seen by him. Stella then turns her attention to Miss Lonelyhearts in the apartment below Thorwald's; she notices that the woman has poured out a huge quantity of sleeping pills. Our schema pops into place and dictates a (clichéd) hypothesis: Miss Lonelyhearts will try to kill herself. The first phase of this long scene will thus shift between two lines of action, Thorwald and Miss Lonelyhearts, and hence two hypotheses: he will go out in obedience to the note, she will attempt suicide. The narration cleverly lets us attend more carefully to Miss Lonelyhearts than the characters do. Lisa rushes in, and the two women decide to dig among the zinnias; Jeff will call Thorwald out of the apartment for awhile. As they pore over the phone book, cut away to Miss Lonelyhearts lowering the blind; they do not see her. And the Miss Lonelyhearts situation is left dangling while Jeff calls Thorwald. The women leave, and while they dig, Jeff calls Doyle's home and leaves a message. Jeff's attention drifts to Miss Lonelyhearts, who is at a desk writing; he mutters: "Stella was wrong about Miss

Lonelyhearts." But for the viewer, who has seen the fistful of sleeping pills and the despairing lowering of the blinds, her new action "realistically" confirms the suicide hypothesis; we label the gesture "writing a suicide note." By this crucial point we have become so well trained that we can contradict the inference of the character who has been our principal tutor.

Stella and Lisa find nothing in the garden and, setting out on her own, Lisa shinnies up the fire escape to Thorwald's apartment. Jeff watches her through his telephoto lens as she discovers Mrs. Thorwald's handbag, but not the wedding ring. When Stella returns, she notices Miss Lonelyhearts about to swallow her pills. Now the characters catch up with us and recognize that two lines of action have been running simultaneously. (What is surprise for them has been suspense for us.) Jeff can call the police to tell them of Miss Lonelyhearts, or he can call Thorwald's apartment to warn Lisa. A long shot of the apartment block brilliantly stages the problem, as Thorwald returns down the corridor, Lisa hesitates in Thorwald's living room, and Miss Lonelyhearts languishes in the apartment below. As it turns out, Miss Lonelyhearts was only (again) a delaying device, for the sound of the composer's song halts her suicide attempt; but by then Thorwald has discovered Lisa. Jeff calls the police—anticipating the action we are about to see: "A man is assaulting a woman . . ."—but he can then only watch helplessly from his wheelchair as Thorwald grabs her. Like Jeff and Stella, the viewer must wait for the outcome. After the police arrive and arrest Lisa, she signals Jeff that she has the ring. Thorwald looks across the way straight into Jeff's lens.

Thorwald's returning of Jeff's gaze makes Jeff the prey. As Jeff sends Stella out and phones Doyle, we alone see Thorwald leave his apartment. The phone rings and Jeff, wrong expectation in place, answers it unthinkingly: "Tom, I think Thorwald's left. I don't see . . ." Silence and a click. Now Jeff, not Thorwald, has given himself away. Jeff prepares for Thorwald's arrival by switching out the light and waiting with a clutch of flashbulbs (motivated realistically as the weapon to which a photographer would "naturally" turn). In

a neat reversal, we now infer Thorwald's approach by sonic cues instead of visual ones—the slam of a door and heavy footsteps in the hall. The viewer is in that situation of classic suspense astutely described by Noël Carroll: it seems more probable that Thorwald will kill Jeff, but we hope for the less likely alternative.³⁷ In realistic terms Jeff is likely to lose, but in compositional and transtextual terms we bet that he will not. He does not, although he winds up tumbling from his rear window to the courtyard into which he had stared so avidly.

Hitchcock has claimed that the tempo of the film's cutting steadily increases across the film.³⁸ The description is not accurate in detail, but certainly the early scenes insist upon longer takes (around thirty-six seconds) and leisurely camera movements, while the climactic sequences (starting with the discovery of the dead dog) are cut much faster, with the average shot length hovering between four and ten seconds. This is not idle bookkeeping, for such rapid cutting can control the pace of our viewing activity. By the end of a film, the tempo of narrative events and cutting can move faster. From the vague, multiple, heterogeneous hypotheses of the opening we have moved to a narrow set of alternatives: Lisa will or will not get caught, Miss Lonelyhearts will or will not kill herself, Jeff will or will not avoid Thorwald's onslaught. We can follow rapid editing because at this point each shot need convey only one of three sorts of information: "this reinforces alternative X," "this reinforces alternative not-X," or "still waiting." When we "lose ourselves" in a film's climactic moments, it is to a large extent because of the power of a perceptual-cognitive activity that, informed by the logic of the film up to that point, sorts information on the basis of a stringently reduced series of possible outcomes and then bets, in rapid interaction with the flow of story cues, upon certain results.

"Tell me everything you saw," Lisa urges Jeff after she finds Mrs. Thorwald gone, "and what you think it means." *Rear Window* shows how powerful can be the effect of a film that overtly grounds its formal method in basic processes of perception and cognition. Every fiction film does what *Rear Window* does: it asks us to tune our sensory capacities to

certain informational wavelengths and then translate given data into a story. Lisa's request is made, more implicitly, by every film that solicits narrative comprehension. Of course, not every film reinforces such conventional ideological categories (nagging wife, society model, adventurous photographer, lusty newlyweds, old maid) and places such trust in the connection between seeing and understanding. Some films—and we shall consider several of them later—under-

mine our conviction in our acquired schemata, open us up to improbable hypotheses, and cheat us of satisfying inferences. But no matter how much a film arouses our expectations only to frustrate them, or creates implausible alternatives that turn out to be valid, it still assumes that the spectator will initially act upon those assumptions which we use to construct a coherent everyday world.