CHAPTER

17

INTENTIONAL FORGETTING IN PERSPECTIVE: COMMENTS, CONJECTURES, AND SOME DIRECTED REMEMBERING

Robert A. Bjork
University of California, Los Angeles

This is a remarkable book. Over the years, I, like other authors in this volume, have felt the need to start articles and talks on intentional forgetting with the argument that forgetting is not simply a failure of humans as information-processing devices, but, rather, that forgetting is an essential component of *any* information-processing system, living or artificial; that there must be some means to forget, or erase, or inhibit, or segregate out-of-date information. To illustrate that argument, it has also seemed necessary to point to examples of everyday situations where forgetting is consistent with our goals, especially the need to update our memories. In this volume as well, Colin MacLeod's impressive overview of directed-forgetting research begins in similar fashion, as does the E. L. Bjork, R. A. Bjork, and Anderson chapter on "Varieties of Goal-Directed Forgetting."

Why do such preambles seem necessary? One reason, I think, is that we are aware that the typical person in our audiences—even many of our professional colleagues—assumes, at least tacitly, that remembering is good and forgetting is bad. A second reason, related to the first, is that we realize that the notion of *intentional* forgetting is a strange one, at least outside of Freudian/psychodynamic contexts. Forgetting does not seem like something we would *want* to do in the first place, and it does not seem like something we *could* do, even if we wanted to. The observation, familiar to all of us, that it seems impossible *not* to think of something, such as an elephant, when instructed to do so (cf. the work of Wegner and his colleagues; e.g., Wegner, 1994), gets interpreted as a kind of evidence that forgetting is not something over which we have any control.

An incident from early in my own involvement in research on intentional forgetting illustrates the point. In 1967, and again in 1968, I gave talks on intentional forgetting at Indiana University, the second of which I titled "Intentional Forgetting, Part 2: Forget Part 1." I began that talk with a disclaimer, saying that I didn't really mean what the title said, but that I couldn't resist using it. At that point, someone in the audience said, "Well, it's too late now." The reason that comment was funny then, and is funny now, is because it presumes the speaker was able to do a preposterous thing, that is, to forget on demand.

One reason I find this book remarkable is that its chapters, collectively, demonstrate that our justifying preambles have, in a sense, been too weak. The picture of intentional forgetting that emerges from this volume—in human, animal, social, legal, and clinical contexts—is that intentional forgetting is as variegated, as complicated, and nearly as prevalent as intentional remembering—and also inextricably intertwined with intentional remembering.

This book is also remarkable in its timeliness. As Golding and Long point out in their broad overview chapter, the various approaches to the study of intentional forgetting have been largely isolated from each other. With some noteworthy exceptions, the separate research traditions represented in this volume have not so much passed like ships in the night as run side by side like ships in the night. In certain instances, the lack of communication is understandable, possibly even justifiable, because the underlying issues diverge substantially across those separate traditions, but in other instances, the lack of cross-fertilization seems a missed opportunity to broaden our separate perspectives and to share empirical findings and procedural innovations in a timely way.

In the first of the major sections that follow I give my perspective on the history of research on directed forgetting in human memory. I then comment on ignoring, disregarding, and discounting information in social and legal contexts; on motivated forgetting in clinical contexts; and on other approaches to the study of intentional forgetting.

DIRECTED FORGETTING IN HUMAN MEMORY: A PERSONAL PERSPECTIVE

Directed Forgetting as a Puzzle

My involvement in directed-forgetting research began as an effort on my part to control for memory load. As a graduate student at Stanford University, I was intrigued by some results reported by Bennet Murdock in 1963. Murdock presented short lists of paired associates to subjects and, at the end of each list, tested for the recall of one of the pairs in each list by

presenting its stimulus member as a cue for recall of its response member. Holding constant the number of pairs presented prior to a tested pair, Murdock manipulated (across lists) the number of subsequent pairs in the list; and, conversely, holding constant the number of subsequent pairs, he also manipulated the number of prior pairs in a list. As the number of pairs that either preceded or succeeded a tested pair increased, performance on that pair decreased, which Murdock attributed to increasing within-list proactive and retroactive interference, respectively.

What seemed unclear to me was whether the decrease in recall as a function of the number of prior pairs should actually be attributed to proactive interference of the classic interference-theory type. Given that subjects had to learn the proactive pairs as they were presented, because any one of them might be tested at the end of the list, there were reasons to expect increased proactive interference as the number of such pairs increased. However, the number of prior pairs also increased a subject's memory load, that is, the total number of pairs that were a potential target of the end-of-list test. It occurred to me that one might be able to control for memory load with a signal to subjects that they could forget the presignal pairs in a given list; that is, a signal that the pair to be tested at the end of the list would *not* be a presignal pair. As long as some lists were to contain such a signal and others were not, in unpredictable fashion, the subjects would still need to try to learn the to-be-forgotten pairs *as* they were presented, which would be sufficient reason to expect proactive interference of the classic type. With respect to a subject's memory load of pairs that were candidates to be tested at the end of a list, however, such a signal would let the subject eliminate the presignal pairs.

As tends to be true of our autobiographical memories for events during our graduate careers, particularly those happening relatively early in our graduate careers, I have a clear image of the reaction of my graduate advisor and mentor, William Kaye Estes, to the notion of signaling subjects to forget (by means of a background color change). His reaction was, as I recall it, best described as bemused. He nonetheless encouraged me to go ahead with the experiment, and quite strongly. Looking back, however, I think he mostly thought that running such an experiment would do me no harm and might even advance my general education as a researcher.

When I later showed Professor Estes my actual results, I am not sure either of us knew quite what to think. I had expected, as a young mathematical modeler, to be able to tease out the relative contributions of memory load and "genuine" proactive interference, but what I found was no proactive interference effects at all of the to-be-forgotten pairs. In that sense, subjects seemed able to forget when instructed to do so. Initially, if briefly (see MacLeod's review), we even had to entertain the possibility that some kind of erasure mechanism might be operating. Such a possibility seemed

unlikely, but was suggested not only by the elimination of proactive interference, but also by subjects' apparent inability to recall to-be-forgotten pairs. When, near the very end of the experiment (and contrary to the assurances I had given the subjects), I presented the stimulus member of a to-be-forgotten pair as a test for its associated response (without informing subjects that I was doing so), recall of the correct response was negligible.

A second experiment, carried out in collaboration with Gordon Bower, also produced results that seemed surprising and puzzling at the time. We again presented short lists of paired associates, some of which (in unpredictable fashion) contained cues to forget the pairs presented prior to the cue, but we occasionally repeated the stimulus member of a to-be-forgotten (TBF) pair as the stimulus member of a to-be-remembered (TBR) pair, but with a new response (in essence, an A–B, A–D manipulation at the level of individual TBF and TBR pairs). When such TBR pairs were tested, we found that recall was actually better than corresponding pairs where the stimulus term was not a repetition of a to-be-forgotten stimulus. That result, if still a bit surprising, is more interpretable now (see MacLeod's review of the work by Timmins, 1973, and by Geiselman & Bagheri, 1985, on repetition effects in directed forgetting) than it was then.

Those two experiments, referred to as "early pilot studies" in the Bjork (1970) article, comprised my earliest work on directed forgetting. My interests in directed forgetting then simmered for a couple of years, while I attended to such matters as fitting Markov learning models and completing a doctoral dissertation, but were rekindled shortly after I arrived at Michigan in 1966 as a beginning assistant professor. David LaBerge, however, who had been my advisor during my one year of graduate work at the University of Minnesota, learned about my initial results when he visited Stanford, and he decided to pursue the forget-cue procedure. With the assistance of Ross Legrand, and with little or no input from me, he carried out a follow-up experiment using a procedure designed to eliminate—or at least greatly reduce—opportunities for rehearsal. After concluding that the results were interesting and generally consistent with what I had found earlier, he sent me an initial write-up of the experiment, which, to my surprise, had me listed as first author. My contributions to the version that we eventually published (Bjork, LaBerge, & Legrand, 1968) were modest, but David LaBerge, over my objections, continued to insist that I should be first author.

Directed Forgetting as a Tool

As is captured so well in MacLeod's review, the earliest research focused on directed-forgetting phenomena as a kind of puzzle that was interesting for its own sake. At Michigan, a number of faculty members (particularly Elizabeth Bjork, David Krantz, Arthur Melton, and Walter Reitman), postdoc-

toral visitors (particularly David Elmes and Addison Woodward), and graduate students (particularly Arthur Glenberg, Robert Jongeward, and Alexander Pollatsek) shared my interest in and curiosity about directed forgetting. We generated new procedures and results at a rapid rate, and new results and procedural variations from other laboratories also began to appear in abundance, as MacLeod summarized so well in his review of the "golden age" period.

Looking back at that period, I realize now that the way I came to view directed-forgetting research was heavily influenced by my collaborators and colleagues at Michigan, and by the information-processing zeitgeist that then prevailed in the Human Performance Center where I had my office. The influence of two colleagues in particular, Walter Reitman and Arthur Melton, warrants mention.

Walter Reitman was very interested in directed forgetting from the very first moment that I mentioned some of my results to him, but *why* he was so interested seemed strange to me. He found directed forgetting interesting from a control-process standpoint, and he thought it important to do certain experiments in which subjects were *asked* to control their forgetting/remembering strategies in various ways. He was also interested in such unsavory things as getting verbal reports from subjects as to their own metacognitive operations. Given my own training, such approaches seemed to smack of mentalism or introspectionism. At one point, we cotaught a graduate seminar—a seminar that was unusually well attended; one attraction, apparently, being the opportunity to watch Walter and me argue with each other, which we did frequently and with considerable vigor. The term *directed forgetting* was coined in one of those seminar sessions and soon became our preferred term; it had, as we viewed it then, an attractive double meaning—that the forgetting in question was "directed" in the sense of being cued or instructed and was also "directed" in the sense of being aimed at the to-be-forgotten information.

Walter Reitman's influence on me, if by no means immediate, was eventually substantial. In fact, as I mention at the end of these comments, I think that issues of control, intent, and resource allocation are among the important remaining issues in the study of directed forgetting in humans (see also the comments on intentionality in the E. L. Bjork, R. A. Bjork, & Anderson chapter in this volume).

Arthur Melton's influence was also substantial, but of a quite different character. He influenced me to look at directed forgetting in the broader context of forgetting and interference processes, and in my interactions with him, I came to realize that certain of the questions suggested by directed-forgetting phenomena were very similar to those faced by the interference theorists of another era. (See the E. L. Bjork, R. A. Bjork, & Anderson chapter in this volume for an example of some of those similarities.)

The information-processing approach of that era, based as it was on a kind of flow-chart analogy between how humans process information and the processing architecture of the typical digital computer, influenced my thinking and the thinking of most memory researchers. As I stressed elsewhere (Bjork, 1989), for example, the computer metaphor led one to think in terms of explanatory mechanisms such as scanning, grouping, sorting, tagging, and so forth, and not in terms of mechanisms like inhibition. During that period, I took pains in colloquium talks to disabuse my audience of the idea that directed-forgetting results had anything to do with actual inhibition or erasure, or with clinical phenomena, such as repression.

Partly as a consequence of the prevailing information-processing zeitgeist, my orientation gradually shifted across my years at Michigan from trying to understand the dynamics of directed forgetting, *per se*, to viewing cues to forget and remember as tools to study processes such as rehearsal and the differentiation of items in human memory. In my 1972 review of the directed-forgetting literature, for example, I noted that the widespread interest in the directed-forgetting paradigm seemed motivated, at least in part, by the impression that it revealed "some curious and previously unappreciated abilities," but I argued that the paradigm was important "not primarily because it raises new questions or illustrates surprising capacities, but rather because it has the potential of contributing new leverage on some old and important problems in the study of human memory" (p. 218). Looking back at my own directed-forgetting publications during the 1970s—that is, publications of experimental studies where a directed-forgetting manipulation was involved—the label *directed forgetting*, or some equivalent expression, often does not even appear in the titles of those articles (see Bjork, 1975, 1978; Bjork & Geiselman, 1978; Bjork & Landauer, 1979; Woodward, Bjork, & Jongeward, 1973).

Basically, I argued in 1972 that directed-forgetting procedures constituted a new tool for memory researchers. It is interesting that now, a quarter century later, MacLeod concludes his review in this volume with the assertion that directed forgetting is "an established technique in the set available to memory researchers"—a technique that "as we develop new tests of memory and new ideas about its operation . . . will no doubt be called on to help in answering questions along the way" (p. 52).

The Search for a Missing Mechanism

With respect to the goal of understanding directed forgetting *per se*, virtually all of the basic phenomena that had been reported as of the mid 1970s, the point at which I moved from Michigan to the University of California, Los Angeles, seemed consistent with the two-process theory I had advocated earlier (Bjork, 1970). According to that theory, subjects, in response to a

forget cue, "devote all further rehearsal, mnemonic, and integrative activities exclusively to the R-items; and they differentially group, organize, or code R-items in a way that functionally segregates them from F-items in memory" (Bjork, 1972, p. 229). On certain empirical grounds, however, and in terms of a kind of plausibility argument that applied to certain experimental situations, I began to feel ever more strongly that those two processes were not the sole mechanisms involved in directed forgetting. My misgivings were shared by Edward Geiselman, who joined me at UCLA, first as a postdoctoral fellow, then as a faculty colleague.

We were suspicious about the theory in two general respects. First, as we argued in the Geiselman, Bjork, and Fishman (1983) article, the theory emphasized mechanisms of selective remembering, not selective forgetting, and an accumulating body of diverse findings seemed to imply that some kind of active forgetting process was triggered by cues to forget. Certain effects of repeating or re-exposing TBF items, for example, seemed most readily interpreted as evidence that those items, in response to the initial cue to forget, had been inhibited (see Bjork, Abramowitz, & Krantz, 1970, and E. L. Bjork, R. A. Bjork, & Glenberg, 1973, as cited in Bjork, 1989; Reed, 1970; Timmins, 1973). The early work of Weiner and his colleagues (e.g., Weiner & Reed, 1969), with linkages to motivated forgetting, such as repression, also suggested that active inhibition might be involved (see MacLeod's summary), as did the results obtained by Johnson (1971), who examined the pupillary response to forget cues. And there seemed to be parallels between certain posthypnotic-amnesia phenomena, where active forgetting was more clearly involved, and corresponding directed-forgetting phenomena.

Second, in certain cases where the two-process theory was logically consistent with the results, it seemed implausible that subjects could actually execute the presumed selective-rehearsal and selective-grouping operations in the time available (often only 1 sec or so). In an experiment by Jongeward, Woodward, and Bjork (1975), for example, sets of four words were presented, one word at a time at a 2.3-sec rate; there was then a 3-sec rehearsal period; and, finally, a 1-sec cue to remember the first two, the last two, all, or none of the words in that set. Thus, for a total of 12.2 secs, subjects had to process the words in a given set without knowing which of those words, if any, they would be asked to remember or forget, and they then had 1 sec to interpret quite a complicated cue before the next set of four words began. They were also instructed to restrict any rehearsal activities to the current set of words, which they said, during the debriefing, they were only too happy to do. At the end of lists consisting of eight such four-word sets, the subjects recalled about 35% of the TBR words and intruded less than 5% of the TBF words.

In an initial effort to find the "missing mechanism," Bjork and Geiselman (1978), using a variation on the item-by-item cuing procedure, obtained

evidence that the momentary retrieval of TBR items in response to a remember cue, together with the absence of such a retrieval of TBF items in response to a forget cue, played a critical role in the differentiation of TBR and TBF items in memory. That mechanism, however, seemed more descriptive than explanatory, and also seemed special purpose in nature; that is, it was not clear how that mechanism might explain some of the other results that seemed to suggest that a forget cue triggered an inhibitory process.

Our next approach was to examine the effects of a forget cue on the incidental memory of to-be-judged items that had been interleaved with the to-be-learned items (Geiselman et al., 1983). The fact that a cue to forget or remember the to-be-learned items had the same effects on the to-be-judged items as on the to-be-learned items themselves, even though subjects were not trying to learn the to-be-judged items (and, hence, had no reason to rehearse those items either before or after a cue of either type), provided strong evidence for an active forgetting process of some kind. The details of our results implicated retrieval inhibition as that process (for excellent summaries of our results and arguments, see the chapters by MacLeod and E. L. Bjork, R. A. Bjork, & Anderson in this volume).

Our initial evidence suggesting that retrieval inhibition was the "missing mechanism" in directed forgetting led Ed Geiselman and his students to test the retrieval-inhibition idea in various ways, as summarized in MacLeod's chapter, and led me to reinterpret some earlier unpublished results (see Bjork, 1989). At the point that the existing data seemed convincing that retrieval inhibition was, indeed, a mechanism in directed forgetting, Harold Gelfand, who was then visiting UCLA, raised an interesting question: What are the necessary conditions for retrieval inhibition to happen? Is it something that simply happens, in magic-wand fashion, when subjects are presented an explicit and unambiguous cue that information presented prior to the cue is wrong, or was presented in error, and should be forgotten? Or is it necessary that some type of new to-be-learned material is presented—something that will serve to replace the to-be-forgotten material?

To address those questions, we (Gelfand & Bjork, 1985, as cited in Bjork, 1989) asked subjects to recall a first list of words they had earlier been instructed to forget (or to remember). The critical experimental manipulation involved what happened *after* the instruction to forget or remember that list and *before* the recall of that list was then tested. During that interval, which was fixed in duration, the subjects either did nothing (while the experimenter fumbled around with some folders), or carried out an incidental-learning task involving a list of adjectives, or learned a second list. It was only in the condition where a second list was learned to replace the to-be-forgotten list that we found evidence of retrieval inhibition. Thus, from those results, it appears that retrieval inhibition is a by-product of new learning, and not simply a product of an intent or instructional set to forget.

Gelfand and Bjork's results continue to influence my own thinking about the dynamics of directed forgetting, particularly having to do with the updating of memory, and with the linkage of directed-forgetting phenomena to other phenomena, such as "unlearning" in retroactive interference. Those results are also suggestive of mechanisms that might underlie the repression of real-world memories (see the speculations of E. L. Bjork, R. A. Bjork, & Anderson in this volume). Were I to construct my own table of "Principal Findings," analogous to MacLeod's Table 1.2, the fact that retrieval inhibition is apparently a by-product of new learning would be high on the list.

By the end of the 1980s, I had not only come to believe—in marked contrast with my views of an earlier era—that inhibitory processes were involved in directed forgetting, but that inhibition, particularly retrieval inhibition, is a prevalent and adaptive mechanism in human memory more generally (see Bjork, 1989). About that same time, other results from other paradigms in research on attention, language, and memory soon began to implicate inhibitory mechanisms as well. From that point to the present, in collaboration with a number of individuals at UCLA, especially Michael Anderson, Elizabeth Bjork, Laura Da Costa, John Shaw, and Bobbie Spellman, we have taken a broader look at the role of inhibition in human memory (see, e.g., Anderson, Bjork, & Bjork, 1994; Anderson & Spellman, 1995; E. L. Bjork, R. A. Bjork, & Anderson in this volume; Shaw, Bjork, & Handal, 1995).

Procedures, Precedents, and Processes in Directed Forgetting

In his review of directed-forgetting research in this volume, MacLeod covers findings that derive from two basic procedures: the *item procedure*, in which individual items are cued one at a time; and the *list procedure*, in which there is a single cue to forget one of two sets of items, where those sets are defined temporally and/or by item type. He rightly credits Muther's (1965) study—of which I was unaware during much of my own early work—with being the first to employ the item-by-item procedure in directed forgetting. If I can lay any claim at all to having been the first to use a directed-forgetting procedure, that claim is limited to the standard version of the list procedure, where an explicit cue to subjects to forget the items studied prior to the cue is presented before the items they are then asked to learn, if any.

If I was unaware of Muther's work, I was well aware, at least by some early point in my work at Michigan, of John Brown's (1954) study,[1] which

[1]In a footnote, MacLeod mentions that the same John Brown, in his 1958 article, used a short-term forgetting procedure virtually identical to that used by Peterson and Peterson a year later in their well-known 1959 article. As an historical footnote of sorts, I should perhaps reveal that I deserve substantial credit for the term "Brown–Peterson paradigm," which was

MacLeod cites as an early precursor of work on directed forgetting. On each of a series of trials, Brown instructed subjects, either before (preinput cuing) or after (postinput cuing) the presentation of two interleaved or successive sets of items, whether to recall one or both of those sets, and, if both, in what order. Although it now seems fair to cite Brown's study as a precursor of the work of Epstein and his colleagues on postinput cuing, I did not, at the time, think of Brown's experiment as a directed-forgetting experiment. Rather, I thought of it, as Brown did, as an experiment on effects of set-to-learn and output interference.

The various ways one might interpret Brown's early experiment raises the issue, I think, of what should and should not be called directed forgetting—and, hence, what should and should not be called an experiment on directed forgetting. We have reached a level of procedural and process sophistication, in my opinion, where there is some need to sharpen our terminology. In particular, I think precuing procedures warrant another name of some kind ("directed ignoring"?), and I think only a subset of postcuing procedures should qualify as directed-forgetting procedures. If every experiment where subjects are told, implicitly or explicitly, that they need not try to learn or remember some of what is to be presented, or are told they need not recall some of what they have tried to learn, is to be thought of as an experiment on directed forgetting, the term becomes so broad as to become meaningless from a process standpoint. Most of our laboratory paradigms in the study of human memory involve either preinput instructions to subjects that some to-be-processed items need not be learned, as in the case of digit shadowing and other rehearsal-preventing distractor activities, or postinput instructions to recall only a sample of what was learned, as in the case of cuing the recall of one of a set of paired associates, or both.

At a minimum, I think the term *directed forgetting* should be reserved for situations in which (a) there has been a prior attempt, however brief or extended, to learn the material that is now to be forgotten and (b) there is an explicit (or totally unambiguous implicit) cue to forget that material.

Beyond that minimal definition, it may prove increasingly necessary to distinguish among the procedures that *do* qualify as directed-forgetting procedures. As MacLeod concludes in his review, it has become apparent, especially given the work of Basden, Basden, and Gargano (1993; see Basden & Basden, this volume), that the item and list procedures invoke quite different processes, or at least a different weighting of the three processing mechanisms—selective rehearsal, selective grouping, and retrieval inhibi-

---

introduced in the late 1960s and is now standard. Not that I, as a new assistant professor, had any real influence on the field during the late 1960s; rather, *I* was the one who convinced Art Melton that John Brown, as well as the Petersons, deserved credit for that procedural innovation, and Art Melton *did* have an influence on the field.

tion—that have been implicated in research on directed forgetting. Within the list-procedure category, I also think the mechanisms involved differ quite dramatically depending on whether the cue to forget a set of items studied earlier is presented *before* or *after* the set of items that are to be remembered (and recalled). In particular, I think the fact that a forget cue of the former type, that is, one presented prior to the to-be-remembered items, is the much more effective of the two (see, e.g., Bjork, 1970) is attributable, in part, to the critical role the *subsequent* learning of the to-be-remembered items plays in the inhibition of the to-be-forgotten items (see my earlier comments on Gelfand & Bjork's, 1985, results).

Stated in a different way, I think directed-forgetting procedures should no longer be thought of as *a* tool in the study of human memory dynamics, but, rather, as a set of tools, where different procedures are the tools of choice to gain leverage on somewhat different problems. Thus, if it is one's goal to understand the "inhibitory regulation of working memory" (Hasher & Zacks, 1988; Zacks & Hasher, 1994), for example, and how that process may change with aging, the item procedure seems a particularly useful tool (see, however, Zacks, Radvansky, & Hasher, 1996, where some inferential leverage is gained by using both the item and list methods). If, on the other had, one's primary goal is to understand interference and inhibition processes in the updating of human memory, the list method seems the tool of choice. In short, as MacLeod documents so well in his review, directed forgetting as a process is multifaceted; and, I would add, as a research instrument directed-forgetting procedures are multifaceted as well.

INTENTIONAL FORGETTING IN SOCIAL AND LEGAL CONTEXTS

The six chapters in this volume on intentional forgetting in social and legal contexts, together with Golding and Long's overview of those topics, provide a thorough review of an interesting and impressive body of research. The cognitive approach and social/legal approach to the study of intentional forgetting are, by a considerable margin, the two approaches to the study of intentional forgetting that are most heavily represented in this volume. Unfortunately, those two lines of research, as Golding and Long point out, have been largely isolated from each other. They assert, in particular, that research on intentional forgetting in the social/legal tradition has been mostly uninformed by developments in the research on directed forgetting in humans, but the converse is at least as true, in my opinion.

To some extent, the isolation of the cognitive and social/legal traditions is understandable, even defensible. The issues that provide the historical motivations to use instructions to forget, ignore, disregard, suppress, and

discount information in the social and legal research traditions are somewhat different from those motivating research in the cognitive tradition, though there are clear points of intersection. In general, however, I think research on intentional forgetting in social and legal contexts would profit from a full understanding of the theoretical and procedural developments in research on directed forgetting in humans, and I think there are innovations and results in the social/legal research tradition that have clear implications for those of us concerned with the basic mechanisms of directed forgetting in humans.

A point that recurs with some frequency in the chapters on intentional forgetting in social and legal contexts is that instructions to disregard, ignore, and so forth are frequently unsuccessful in achieving their stated purpose. That is, subjects' judgments or impressions or attitudes are frequently unaffected by such instructions, or even show "boomerang" effects. Where the authors *do* make allusions to the directed-forgetting literature, they often say that such findings seem at odds with the successful directed forgetting typically obtained with standard directed-forgetting procedures using simple verbal materials. What strikes me, however, is how consistent the results from the social/legal domain are with the results in the directed-forgetting domain—provided that the experimental conditions are comparable.

As far as whether the conditions of intentional forgetting are the same from one situation to another, a number of considerations seem important, some of which I discuss below. In that context, I point to some noteworthy parallels in the findings obtained in the cognitive and social/legal domains. My comments are necessarily brief. For a perceptive, systematic, and thorough effort to link and organize the cognitive and social-judgment literatures on intentional forgetting, the interested reader should see Johnson's review (1994). My view of the relations between the paradigms and phenomena in the cognitive and social/legal domains has much in common with her analysis.

Forgetting Versus Ignoring, Disregarding, and Discounting

If, as I mentioned earlier, there is an issue as to what should and should not be considered directed forgetting in human memory research, there is a far broader and more complicated issue as to what should and should not be considered intentional forgetting in social and legal contexts. One aspect of the issue is whether the subjects in a given experiment are, in fact, actually asked to forget something that was said or that happened earlier, or, instead, are asked to be suspicious of that event, or to pretend it never happened, or to discount its relevance or importance.

Terminology in the social/legal research literature—if not always the discussions, analyses, and interpretations in that literature—reflects a sen-

sitivity to the issue of what constitutes intentional forgetting. Terms such as *ignore, disregard, discredit, discount,* and *suppress* have shades of meaning that seem generally well understood. The basic point, however, is that the typical instruction administered in social/legal experiments on intentional forgetting does not actually ask the subject to forget something, but, rather, to give that "something" little or no weight in some decision or judgment.

Instructed Forgetting Versus Intentional Forgetting

A second consideration is whether *whatever* the subject is being asked to do is consistent with the subject's own goals. In the typical directed-forgetting experiment in the human memory domain, the subject's own goals tend to be consistent with following the experimenter's instructions. Under normal circumstances, subjects prefer to have fewer things to remember, to be able to focus rehearsal and mnemonic activities on to-be-recalled information, and to do whatever will help them perform well, all of which are goals fostered by instructions that some items can be forgotten.

In certain social/legal paradigms, however, an instruction to disregard may not be something the subject wants to hear, and may not be consistent with what the subject views as his or her most important task. For example, a juror who is told that some testimony or evidence that seems highly pertinent to the case is inadmissible and should be disregarded, may, consciously or unconsciously, weigh the goal of arriving at the most accurate verdict as more important than honoring the judge's instruction to disregard (forget). In a variety of other real-world and experimental situations as well, the forgetting asked for may be consistent with the asker's goals, but may not fit well with the goals of the recipient of that request. For example, when the stated reason for an instruction to forget or disregard is that the information in question is "confidential," there may well be such a conflicting agenda (see, e.g., Golding, Fowler, Long, & Latta, 1990).

Of course, such considerations are hardly news to social psychologists. In the context of research on intentional forgetting, however, such considerations seem important for two reasons. First of all, in the literature on instructions to ignore, disregard, and discount in social and legal experiments, as summarized in this volume, there is a striking, if unsurprising pattern: The more the information that is the target of such instructions is discredited, proven to be false, and so forth, the more evidence there is that subjects actually follow those instructions in their judgments or recall performance.

That pattern makes me wonder whether in research on hindsight effects, where subjects, in making a judgment or prediction of some kind, seem unable to set aside some outcome information they now know, even when asked to do so (see Golding and Long's review), the results might differ if

there were a way to change the task so that *ignoring* what one now knows is consistent with one's own self-esteem or the opinion of others. As Campbell and Tesser (1983) conjectured, hindsight biases may derive in part from motivational factors—in particular, the fact that having good foresight is an attribute that impresses others. Given that there *are* situations in the world, however, where we know more about something than we might like other people to know we know (about television soap operas, for example; or even, in some subcultures, mathematical reasoning), it might be possible to construct a judgment situation where subjects are more motivated to ignore just-presented outcome information.

A second consideration, and one that may be more important from a theoretical standpoint, is that the *intent* to forget may have an active, intrinsic, and direct effect on to-be-forgotten items in memory. That is, in addition to the indirect effects of intent to forget, such as reducing or eliminating a subject's inclination to continue to think about, rehearse, or integrate to-be-forgotten information with to-be-remembered information, the intent to forget may play an active role of some kind in the inhibition of to-be-forgotten information. It is an open research question, but there are some human-memory results (see E. L. Bjork, R. A. Bjork, & Anderson, this volume) suggesting that inhibition may require the allocation of resources, which, in turn, would require the intent to allocate those resources; the work of Macrae and his colleagues on the role of effort and resource allocation in the suppression of stereotype information also suggests such a possibility (e.g., Macrae, Bodenhausen, Milne, & Wheeler, 1996; for a summary, see Bodenhausen, Macrae, & Milne, this volume).

Direct Versus Indirect Measures of "Forgetting"

Typically, in the social/legal research tradition, whether an instruction to disregard, ignore, or discount some information has an effect on subjects' memory for that information is assessed not by a test of the subjects' later ability to recall that information, but by some other measure, such as whether that instruction alters a legal or social judgment, evaluation, or impression of some kind. The absence of any such effects is taken as evidence that the instruction was ineffective. At one level of analysis, such a conclusion seems unassailable. Often, however, that conclusion is accompanied by a claim, overt or implied, that the instruction also had no effect on subjects' episodic recall of the event or information in question. That claim might also be true, but the lack of an effect of the instruction on judgments and impressions does not necessarily constitute evidence for that claim.

The problem is that judgments, evaluations, and impressions are frequently *indirect* tests of memory, which are often sensitive to prior events

when direct measures, such as recall, are not (see, e.g., Richardson-Klavehn & Bjork, 1988). That is, such tasks typically do not refer the subject back to the episodic events of interest; rather, they require an overall decision or judgment of some kind (for a more complete discussion, see Johnson, 1994). Of particular relevance are the recent findings in the human-memory literature that an instruction to forget can impede recall of the to-be-forgotten information without lessening the effect of that information as measured by indirect tests, such as word-fragment completion (see MacLeod's summary of those findings in this volume). It is entirely possible, therefore, that in certain of the social and legal experiments where an instruction to disregard has been declared completely ineffective, as measured by some type of judgment-task performance, a test of recall of the to-be-disregarded information would have shown impaired recall.

Whether, and under what circumstances, such a dissociation might appear is a matter of considerable theoretical and practical interest. If an instruction to forget or to disregard can impair recollection of the event or information that is the target of that instruction, without lessening its impact as measured by other means, there are circumstances where the impact of to-be-forgotten information on certain judgments may be larger than the impact of comparable to-be-remembered information. Using a variation on the "false fame" paradigm of Jacoby and his colleagues (e.g., Jacoby, Kelley, Brown, & Jasechko, 1989), we recently found support for such a conjecture (E. L. Bjork, R. A. Bjork, Stallings, & Kimball, 1996). When a nonfamous name was made familiar by virtue of its having been presented earlier in the experiment on a list of such names, subjects were more likely to attribute that familiarity, falsely, to real-world fame when that earlier list was to be forgotten than when it was to be remembered.

When Intentional Forgetting Succeeds and Fails: Other Considerations

When considered together, the cognitive and social/legal research literatures suggest that there are a number of procedural determinants of whether instructions to forget or to disregard succeed or fail. In her review, Johnson (1994) attempted to characterize in detail the conditions that yield successful and unsuccessful intentional forgetting. At a somewhat more global level, I am struck by the importance of the following factors.

The Relatedness of What Is to Be Forgotten and to Be Remembered. In general, instructions to disregard or to ignore in social/legal contexts seem ineffective when the to-be-forgotten information has already been integrated with the information to be remembered. In the directed-forgetting literature, an instruction to forget also tends to be ineffective when to-be-

forgotten and to-be-remembered items are closely related and/or integrated with each other (see, e.g., Golding, Long, & MacLeod, 1994). Certain findings that derive from item-by-item cuing procedures seem like exceptions to that rule, but with the item-by-item procedure, to-be-forgotten items may never be encoded at a level that would integrate them with to-be-remembered items, even if the items themselves have a natural relationship.

The Temporal Positioning of Instructions to Forget or Disregard. Both literatures suggest that, in a sense, we can reach a point where it is too late to instruct someone to forget or to disregard. Instructions to forget or to disregard become less effective as time passes because presentation of the information to be forgotten or disregarded passes, especially when that delay is filled with the presentation and processing of information that is to be remembered.

The Reexposure of To-Be-Forgotten Information. Another point of consistency is that reexposure to information that was earlier to be forgotten or disregarded can have dire consequences with respect to the goal of forgetting or disregarding successfully. Such reexposures can reinstate that information and its interference with subsequent recall of information to be remembered and/or its influence on subsequent judgments where that information is to be disregarded.

The Replacement of To-Be-Forgotten Information With To-Be-Remembered Information. Finally, although the conclusion seems more tentative, it appears that giving subjects a set to replace to-be-forgotten information with new to-be-remembered information may play an important role in successful intentional forgetting. Gelfand and Bjork's (1985) results, as mentioned earlier, suggest that new (replacement) learning plays an important role in the inhibition of old to-be-forgotten information, and findings in the social/legal context also suggest that to-be-disregarded information is more effectively nullified in subsequent judgments if alternative information is provided as a replacement for that information.

MOTIVATED FORGETTING IN CLINICAL CONTEXTS

From an historical standpoint, one might expect that the oldest tradition of research on intentional forgetting would be in the context of research on clinical disorders. Conceptually, the notion that intentional or motivated forgetting plays a role in psychodynamic disorders, as in the repression mechanisms hypothesized by Freud and others, traces back a century or

so. From an empirical research standpoint, however, the use of explicit instructions to forget or to remember appears to be a very recent innovation (the exception being in the case of hypnosis, where the use of explicit instructions to hypnotized subjects to forget what happened while they were hypnotized has a very long history; see, e.g., Nagge, 1935).

An early body of controlled research on "motivated forgetting" examined whether materials with a negative valence were forgotten more rapidly than were neutral or positive materials (see, e.g., Ratliff, 1938). The forgetting examined in those experiments was not motivated in the sense of being instructed, but, rather, was motivated in the sense of satisfying a presumed (theory-based) need to forget or repress negative or unpleasant memories. It is a fair summary of that early research, I think, to say that the evidence of more rapid forgetting of negatively valenced materials in those controlled studies was, overall, unconvincing. In an interesting and potentially important development, the motivated-forgetting approach has been resurrected by Cloitre and other researchers, but now augmented by the use of explicit directed-forgetting procedures.

Directed-Forgetting Procedures as a Tool in Clinical Research

As summarized by Cloitre (this volume), directed-forgetting procedures are emerging as a potentially fruitful tool in clinical research. At several different locations, groups of researchers recently carried out a number of experiments in which subjects with certain clinically diagnosed disorders, or subjects who are high or low on certain measured personality traits, or subjects who have a history of trauma of some specific kind, are given explicit instructions to forget or to remember words or phrases that are related or unrelated to those disorders, traits, or traumas (those groups are Cloitre and her collaborators, and McNally and his collaborators, as cited by Cloitre; and Brewin, Myers, Power, and their collaborators in England—see, e.g., Brewin, in press; Myers, Brewin, & Power, 1996). The logic behind these experiments (see Cloitre's discussion) is that subjects who have a certain disorder, or who possess a certain trait, or who have been the victim of a particular trauma, such as rape, are likely to show a pattern of recall performance that is diagnostic and, from a theoretical standpoint, informative. As Cloitre clarifies, there are theory-based reasons, depending on the disorder, trait, or trauma of interest, to expect that relevant materials might be better or more poorly recalled than neutral materials, and that such differences might interact with whether those materials were to be forgotten or to be remembered.

Two aspects of this new line of research make it seem especially promising to me. First, as Cloitre points out—and something that had not occurred

to me prior to reading her discussion—directed-forgetting procedures engage attentional, rehearsal, retrieval, and inhibitory processes in a way that makes enhanced or impaired recall of to-be-remembered materials potentially as informative as enhanced or impaired recall of to-be-forgotten materials. That is, one might think the sole observation of interest would be whether, compared to the recall of neutral materials or the performance of control subjects, the recall of to-be-forgotten materials that have personal valence for a given subject is impaired or facilitated. As Cloitre points out, however, there are also theory-based reasons to expect that the levels of recall of to-be-remembered neutral and valenced materials might differ as well, and that such differences, if any, especially when viewed against the levels of performance on corresponding to-be-forgotten materials, have the potential to be diagnostic and informative.

A second reason the approach seems promising at this early stage is that the initial results do reveal some substantial differences in the recall of materials to be forgotten and to be remembered as a function of subject type, and whether those materials are neutral or subject relevant in some way. Those differences, thus far, have not always corresponded with the experimenters' expectations and predictions, but that fact may simply illustrate that the approach has the potential to inform and refine theorizing. And, as Cloitre points out, researchers may have only scratched the surface thus far with this new methodology; any systematic differences that might appear on comparisons between the list and item-by-item directed-forgetting procedures have the potential to be diagnostic, as do possible differences on direct and indirect measures of retention.

Inhibition and Recovery of Traumatic Memories

E. L. Bjork, R. A. Bjork, and Anderson conclude their chapter in this volume with some speculations on the possible relevance of inhibitory processes identified in laboratory paradigms to the forgetting and recovery of real-world traumatic memories. They carefully label their comments as "speculations," because whether the phenomena and processes identified in highly constrained and emotionally quite neutral laboratory settings are generalizable to the memory of real-world traumas is open to serious question at this point. Certain parallels are tempting and suggestive, however, as these authors outline.

It is not appropriate here to reiterate those possible connections, but a couple comments seem in order. First, and quite the opposite of what I thought years ago, the processes triggered by directed-forgetting procedures in the laboratory may bear some meaningful and instructive relations to the real-world processes of inhibition and recovery that are of interest to clinical researchers. Second, and related to the first point, there is much

to be gained by a dialogue between cognitive scientists interested in forgetting mechanisms and clinical researchers interested in the memory dynamics that characterize certain disorders. It seems clear now that the research paradigms and theoretical characterizations of memory processes developed by cognitive psychologists are of potential importance to clinical researchers. It also seems clear that the work of clinical researchers has the potential to inform the theoretical work of cognitive scientists, in much the same way, perhaps, that research on patients with organic amnesias of one kind or another has isolated and illustrated certain memory processes more clearly than is observable in standard experimentation with normal subjects.

OTHER APPROACHES TO INTENTIONAL FORGETTING

In this section I comment on the other approaches to the study of intentional forgetting that are represented in this volume. Where my comments are brief, it is not because there is little to say, but because my perspective, in terms of how the issues and phenomena that characterize a given approach relate to the issues and phenomena in research on directed forgetting in adult waking humans, is already well represented by the authors of one or more of the prior chapters in this volume.

Directed Forgetting in Animals

The body of research on directed forgetting in animals, which is summarized and analyzed by Grant and by Zentall, Roper, Kaiser, and Sherburne in this volume (and by Roper & Zentall, 1993, in an earlier review), is impressive in two respects. First, it is a body of research marked by careful analysis and clever experimentation. Second, it is impressive in terms of the sheer amount of research that has been completed over the last 20 years or so; I am particularly struck by Zentall, Roper, Kaiser, and Sherburne's statement that "directed forgetting is perhaps the most thoroughly examined animal memory phenomenon that has been studied for its cognitive implications" (pp. 3–4).

The early research on directed forgetting in animals was stimulated by the early research on directed forgetting in humans, but the specific experimental procedures, appropriately enough, soon came to differ in some important ways from those used with humans. The most common procedure used by animal researchers, the matching-to-sample procedure, where cues following a given stimulus sample serve to signal whether comparison stimuli will or will not follow that sample (remember and forget cues, respectively), does not correspond in its details to any of the procedures used

with humans. The authors mention that the procedure is closest to the item-by-item procedure used with humans, but the differences are substantial in how and when TBR and TBF items are tested in the two cases.

At a more general level, however, in terms of process-model considerations, there does seem to be a meaningful correspondence between the matching-to-sample procedures used with animals and the item-by-item procedures used with humans. In both instances, the procedures seem particularly well suited to examining the role of rehearsal/maintenance processes in distinguishing between TBR and TBF items, which has been a major focus of the research on directed forgetting in animals. What I cannot help wondering, however, given some of my own interests, is whether animals could take advantage of cues to forget in the sense of updating memory, that is, in the sense of replacing TBF stimuli with TBR stimuli.

To address that question, it seems plausible that the directed-forgetting version of Sternberg's memory-scanning task, introduced long ago by Bjork, Abramowitz, and Krantz (1970; as cited in Bjork, 1989), and resurrected more recently by other researchers (see Neumann, Cherau, Hood, & Steinnagel, 1993; Neumann & DeSchepper, 1992; Zacks et al., 1996), could be adapted to animals. The standard memory-scanning paradigm consists of a series of trials, on each of which a small set of items is presented one at a time. After each such set of to-be-remembered items, a probe item is presented, to which the subject is to respond "yes" or "no" as quickly as possible depending on whether the probe item does or does not match one of the items in that set. In the directed-forgetting version of the task, some lists, in unpredictable fashion, contain a signal to subjects to forget the items presented prior to the signal, that the judgment to be made is whether the probe matches one of the items that will follow that signal. In effect, the forget cue is a signal to the subject that the list is starting over. Human subjects can indeed start over, or "forget," in the sense that their "yes" response times are unaffected by the presence (or number) of presignal to-be-forgotten items in a given list. Their "no" responses are slower, however, when the probe is a to-be-forgotten item than when it is an item not presented at all on that trial.

What makes adapting the paradigm to animals seem feasible are some results obtained by Wright and his collaborators (e.g., Wright, Santiago, Sands, Kendrick, & Cook, 1985). Using pigeons, monkeys, and humans as subjects, Wright and his colleagues found strikingly similar effects of probe delay and serial input position on the memory-scanning performance of each species (percent correct judgments, not reaction time, was the measure of interest). The search sets consisted of four visual stimuli in each case, but the materials were scaled to each species (simple pictures for pigeons, more complicated pictures for monkeys, and kaleidoscope patterns for humans), as was the range of probe delays (0 to 10 sec for pigeons; 0 to 30 sec for monkeys; and 0 to 100 sec for humans). Their actual results, though very

interesting, are not what is important for present purposes; it is the fact that those results were similar across species that suggests that the directed-forgetting version of the memory-scanning task may be adaptable to animals.

Posthypnotic Amnesia

Golding and Long provide a concise and informative summary of research on posthypnotic amnesia—that is, on the effects of instructions to hypnotized individuals that, after awakening, they will be unable to recall the events that occurred during hypnosis. In the modern era of research on posthypnotic amnesia, there has been a productive interaction between researchers interested in posthypnotic amnesia and researchers interested in directed forgetting in waking humans, in part because those researchers have sometimes been the same people. One issue that has motivated recent research on posthypnotic amnesia, in fact, is whether the mechanisms are the same as those identified in nonhypnotic directed forgetting. As Golding and Long mention, there is currently a difference of opinion on that matter among researchers.

Whether the mechanisms are basically the same or differ in some important way, there are some striking similarities between the empirical phenomena that characterize posthypnotic amnesia and analogous effects of directed forgetting in waking subjects. In both cases, for example, there is disrupted and disorganized retrieval of to-be-forgotten items when subjects are asked to try to recall those items; and subjects exhibit a kind of source amnesia for where and when the TBF items that are recalled were presented. Also, in both cases, the influence of to-be-forgotten events and information on indirect measures of memory is nonetheless preserved, and TBF items are as readily relearned as TBR items.

Thought Suppression

In the influential work by Wegner and his colleagues on thought suppression (e.g., Wegner, 1994; Wegner, Schneider, Carter, & White, 1987), the basic procedure involves asking subjects *not* to think about some target item, such as a white bear or the Statue of Liberty. As a means of trying to comply with that instruction, subjects typically initiate a process of self-distraction; that is, they try to keep their minds focused on some *other* salient thought or personal experience. That strategy intermittently fails, however, apparently because subjects, owing to some kind of nonvolitional monitoring process, intermittently check whether they are complying with the instruction. There is also evidence that attempting to suppress a thought can make it hyperaccessible as measured by other means and that, after a period of suppression, there can be "rebound" effects where thoughts that were suppressed become more accessible than they would have been without a period of attempted suppression.

It is not feasible or appropriate here to summarize the myriad variations on that basic procedure that have now been explored, nor to describe fully the details of the phenomena that have been obtained. Golding and Long provide a brief overview in their chapter, and Bodenhausen, Macrae, and Milne summarize a series of interesting experiments on a related topic, stereotype suppression, in their chapter. There are some comparisons of the basic methodologies of thought-suppression and directed-forgetting experiments that merit comment, however, and some speculations on the relationship of the processes that underlie thought suppression and directed forgetting also seem warranted.

At one level, an instruction not to think about something, such as a white bear, seems quite similar to an instruction to forget something that was just presented for study. There are some substantial differences, however. In the directed-forgetting case, for example, the instruction refers to a well-defined episodic event, not a concept or experience or stereotype that already exists in memory. Also, it is the experimenter-provided task of learning the items that *are* to be remembered that diverts attention from the to-be-forgotten items, not a subject-generated effort to think of something other than the to-be-suppressed thought. Another difference is that subjects in directed-forgetting experiments are not asked to monitor how well they are avoiding the retrieval of to-be-forgotten items (although, it occurs to me, they *could* be asked to do so), whereas subjects in thought-suppression experiments are typically asked to indicate, by ringing a bell or some other means, when the to-be-suppressed thought has intruded into consciousness.

Possibly because of such differences, the phenomena and presumed processes in research on thought suppression and directed forgetting have not seemed to be all that closely related, although some similarities have been noted. In a chapter on thought suppression as a mental-control technique, for example, Wegner, Eich, and Bjork (1994) stressed that the success of both thought suppression and directed forgetting depends, in part, on avoiding environmental cues that point to or re-present the "suppressed" or "forgotten" thought or information. In general, however, the processes implicated in research on thought suppression have not seemed to inform research on directed forgetting, and vice versa.

Recently, however, Martin Conway (personal communication, 1996) pointed out that thought suppression and directed forgetting may be more closely related than researchers in either area have tended to think. As background for Conway's conjecture, I need first to summarize Wegner's (1994) "ironic processes" theory of thought suppression (see also Golding & Long's summary in this volume). Wegner argues that two processes are involved: one intentionally initiated, effortful, and resource demanding; the other involuntary and less demanding of cognitive resources. In thought-suppression experiments, the first process takes the form of distracting

one's self by thinking about other matters, and is initiated as a means of complying with the instruction to suppress the thought in question. The other process, a monitoring process, is one that runs in the background, but intermittently and involuntarily captures consciousness, so to speak. Its function is to check whether the goal of the first process is being achieved; in the thought-suppression context, however, that involves accessing the very thought that is to be suppressed. An important assumption of Wegner's theory is that the intentional process is more demanding of central cognitive resources; thus, under cognitive-load or secondary-task conditions, it suffers more than does the relatively automatic monitoring process, leading to less effective thought suppression.

Conway's conjecture is that directed forgetting itself may involve a similar interplay of processes and that, in particular, directed forgetting may itself require central resources. That conjecture is prompted by several considerations. First, the results obtained by Gelfand and Bjork (1985), discussed earlier, suggest that *both* the intent to forget (i.e., the allocation of central resources to that goal) *and* new learning of to-be-remembered information are necessary to suppress or inhibit to-be-forgotten information. Second, as E. L. Bjork, R. A. Bjork, and Anderson speculate in this volume, the mechanism that underlies the repression of traumatic memories may be retrieval-induced (or learning-induced) retrieval inhibition; thus, for example, victims of abuse may be motivated to retrieve positive or neutral memories of an abusive family member, which, eventually, would inhibit access to the recall of abuse incidents.

In Conway's view, "retrieval inhibition is exactly the mechanism needed to create Freud's 'screen memories.' " The intention to suppress thoughts or to forget prior events would have much in common and, in fact, would interact in certain ways. Drawing on Wegner's ideas, Conway speculated that "when control-process resources are depleted, repressed materials may be brought to mind by a monitoring process, especially in the presence of relevant environmental stimuli, and at that point directed forgetting could be used to again repress the offensive materials." As an example, he described the behavior of a serial killer interviewed by Christianson and Engelberg (1997), who, "when reminded of his murders, would immediately pick any reading materials to hand (e.g., the label of contents on some tinned food), and intensely and repetitively read that material until the 'memories passed.' " (See Conway, 1997, for more examples.)

Directed Forgetting in Development and Aging

A final approach that merits comment is the use of directed-forgetting procedures to examine issues in the development of information-processing abilities, or in how those abilities change with aging. In their overview

chapter, Golding and Long summarize the literature on directed forgetting in retarded and normal children and adolescents, much of it contributed by Bray and his colleagues; they also summarize the findings in research on directed forgetting in elderly subjects.

To document the usefulness of directed-forgetting procedures as a tool in research, there may be no better examples than the research on intentional-forgetting processes in development and aging. In terms of identifying what memory-control processes, such as selective rehearsal, are and are not available at a given point in normal or retarded development, directed-forgetting procedures have proved fruitful and instructive. And in terms of identifying changes in the management of attentional and working-memory resources that accompany aging, especially having to do with the inhibitory control of those resources (see Hasher & Zacks, 1988; Zacks & Hasher, 1994; Zacks et al., 1996), directed-forgetting procedures have also proved an especially useful tool.

With respect to using directed forgetting as a tool, the recent work of Zacks et al. (1996) illustrates an important point: Theoretical leverage can be gained by using more than a single directed-forgetting procedure. In a series of experiments, Zacks et al. tested younger and older subjects using both the item-by-item and the list methods of directed forgetting. In addition, with the list method, where some lists contained a cue to subjects to forget the items in the list presented prior to the cue, they examined both the recall of postcue items (Experiment 2) and the speed at which subjects could identify whether an end-of-list probe item had or had not occurred among the postcue items (Experiment 3); the latter case is the directed-forgetting version of the memory-scanning paradigm, which I described more fully in my comments on directed forgetting in animals.

Among their other results, Zacks et al. found that older subjects, with item-by-item cuing, were less able than younger subjects to exclude to-be-forgotten items from recall. With the list method, however, they *were* able to exclude to-be-forgotten items; that is, with "start-over" cuing in a list context, the older subjects, like the younger subjects, did not show proactive interference owing to the presence or number of to-be-forgotten items in their recall of, or reaction time to, to-be-remembered items.

For present purposes, it is not important exactly what that finding, and the other findings obtained by Zacks et al., may imply as to the information-processing capabilities that are and are not preserved in aging (for a full discussion, see their article). The important point is that the pattern of effects across different directed-forgetting procedures is *likely* to differ in certain ways in a given situation, and those differences are likely to be instructive. As MacLeod documents in his review chapter, and as I argued earlier in this chapter, the different ways of administering instructions to remember and to forget can put somewhat different demands on subjects,

which, in turn, can trigger somewhat different processing strategies and mechanisms.

Finally, the research by Zacks and her colleagues illustrates yet another important point: The use of directed-forgetting procedures as a tool in research on special populations such as the elderly has the potential to inform theory development more generally. In their use of the directed-forgetting version of the memory-scanning paradigm, they replicated—for both younger and older subjects—a pattern of results that Bjork et al. (1970; as cited in Bjork, 1989) found puzzling: "Yes" response times to to-be-remembered probes were an increasing function of the size of the to-be-remembered set, but independent of the number of to-be-forgotten items in the list; and "no" response times to to-be-forgotten probes were slower than "no" response times to items that had not appeared at all in the list, but were also independent of the number of to-be-forgotten items in the list.

That pattern is puzzling because the "yes" response times suggest that subjects can restrict their memory scanning to the to-be-remembered set, which then predicts that the two types of "no" responses should not differ. Furthermore, given that the two types of "no" responses *do* differ, it would seem that the "no" responses to to-be-forgotten items should vary with the number of to-be-forgotten items. As a possible interpretation of that puzzling pattern, I speculated in my 1989 paper that "no" responses to to-be-forgotten probes might be slowed not because the scanning of the to-be-remembered set was any slower once such probes were encoded, but because to-be-forgotten items were inhibited, resulting in a slower encoding of to-be-forgotten probes. I put that hypothesis forward in a very tentative and qualified way, but Neumann and his colleagues (Neumann & DeSchepper, 1992; Neumann et al., 1993) replicated the same pattern of results and interpreted other aspects of their results as consistent with an inhibited-encoding interpretation.

For reasons that Zacks et al. document, however, such a view would seem to predict that the difference in response time to the two types of "no" probes, that is, to-be-forgotten probes and nonpresented probes, should be larger for the younger subjects, the opposite of what they found. Whether their arguments against the inhibited-encoding view, and in support of an alternative interpretation they favor, are totally convincing or not, their results provide a convincing demonstration that the results of directed-forgetting experiments in the context of development, aging, clinical disorders, and so forth, have the potential to inform theory development more broadly.

CONCLUDING COMMENTS

I have grown accustomed to thinking of directed-forgetting procedures as a new and somewhat novel approach in research on human memory. In the context of this volume, however, as Golding and Long point out, research

on directed forgetting in human memory is the old and established field. It is not true in every instance, but the historical roots of much of the work reviewed in this volume go back to the early work on directed forgetting. For those of us who, 30 years or so ago, began tinkering with instructions to people to forget things, and trying to defend why doing that made any sense at all, it would have been hard to imagine that any book like this one might exist some day, let alone a book with the range and wealth of findings and ideas that are reported in this volume.

So what does the future hold? On the one hand, I am convinced by this volume that directed-forgetting procedures, as a research tool, will prove increasingly useful in a variety of contexts. On the other hand, I think there is much that remains to be understood about the basic mechanisms that underlie the intentional ignoring, forgetting, suppressing, or discounting of information, prior events, or current thoughts. Toward that end, we need to refine our characterization of such processes, which, though closely related, may involve somewhat different mechanisms and may serve somewhat different functions in our interactions with the world.

With respect to directed forgetting in humans, I think there is much that remains to be understood about the inhibitory mechanisms that now appear implicated by recent findings. Do central control processes play a direct role in such inhibition? Or is retrieval inhibition solely a by-product of new learning or the retrieval of competing information? And in what ways are the inhibitory mechanisms in directed forgetting similar to or different from the inhibitory mechanisms that are intrinsic to human attention? Some of the answers to those questions may come, in part, from neuroscience approaches to human memory. The results of research on frontal-lobe functions and the interplay of the hippocampus and neocortex suggest, to me at least, that directed-forgetting tasks, in combination with imaging techniques, might prove revealing.

Finally, whatever twists and turns future research on intentional forgetting might take, one thing is clear: Intentional forgetting, as I said at the outset of this chapter, is inextricably intertwined with intentional remembering. If our goal is to understand fully how information is selected, encoded, learned, and remembered, we also need to understand how information is ignored, inhibited, discounted, and forgotten.

REFERENCES

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1063–1087.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review, 102,* 68–100.

Basden, B. H., Basden, D. R., & Gargano, G. J. (1993). Directed forgetting in implicit and explicit memory tests: A comparison of methods. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 3*, 603–616.

Bjork, E. L., Bjork, R. A., Stallings, L., & Kimball, D. R. (1996, November). *Enhanced false fame owing to instructions to forget.* Paper presented at the meetings of the Psychonomic Society, Chicago.

Bjork, R. A. (1970). Positive forgetting: The noninterference of items intentionally forgotten. *Journal of Verbal Learning and Verbal Behavior, 9*, 255–268.

Bjork, R. A. (1972). Theoretical implications of directed forgetting. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 217–235). Washington, DC: Winston.

Bjork, R. A. (1975). Short-term storage: The ordered output of a central processor. In F. Restle, R. M. Shiffrin, N. J. Castellan, H. R. Lindeman, & D. B. Pisoni (Eds.), *Cognitive theory* (Vol. 1, pp. 151–171). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bjork, R. A. (1978). The updating of human memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 12, pp. 235–259). New York: Academic Press.

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honor of Endel Tulving* (pp. 309–330). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bjork, R. A., & Geiselman, R. E. (1978). Constituent processes in the differentiation of items in memory. *Journal of Experimental Psychology: Human Learning and Memory, 4*, 347–361.

Bjork, R. A., LaBerge, D., & Legrand, R. (1968). The modification of short-term memory through instructions to forget. *Psychonomic Science, 10*, 55–56.

Bjork, R. A., & Landauer, T. K. (1978). On keeping track of the present status of people and things. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (pp. 52–60). London: Academic Press.

Brewin, C. R. (in press). Clinical and experimental approaches to understanding repression. In D. Read & D. S. Lindsay (Eds.), *Recollections of trauma: Scientific research and clinical practice.* New York: Plenum Press.

Brown, J. (1954). The nature of set-to-learn and intra-material interference in immediate memory. *Quarterly Journal of Experimental Psychology, 6*, 141–148.

Brown, J. (1958). Some tests of the decay theory of immediate memory. *Quarterly Journal of Experimental Psychology, 10*, 12–21.

Campbell, J. D., & Tesser, A. (1983). Motivational interpretations of hindsight bias: An individual difference analysis. *Journal of Personality, 51*, 605–620.

Christianson, S.-A., & Engelberg, E. (1997). Remembering and forgetting traumatic experiences: A matter of survival. In M. A. Conway (Ed.), *Recovered memories and false memories* (pp. 230–250). Oxford, England: Oxford University Press.

Conway, M. A. (1997). Past and present: Recovered memories and false memories. In M. A. Conway (Ed.), *Recovered memories and false memories* (pp. 150–191). Oxford, England: Oxford University Press.

Geiselman, R. E. (1977). Effects of sentence ordering on thematic decisions to remember and forget prose. *Memory & Cognition, 13*, 323–330.

Geiselman, R. E., & Bagheri, B. (1985). Repetition effects in directed forgetting: Evidence for retrieval inhibition. *Memory & Cognition, 13*, 51–62.

Geiselman, R. E., Bjork, R. A., & Fishman, D. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic amnesia. *Journal of Experimental Psychology: General, 112*, 58–72.

Gelfand, H., & Bjork, R. A. (1985, November). *On the locus of retrieval inhibition in directed forgetting.* Paper presented at the meetings of the Psychonomic Society, Boston, MA.

Golding, J. M., Fowler, S. B., Long, D. L., & Latta, H. (1990). Instructions to disregard potentially useful information: The effects of pragmatics on evaluative judgments and recall. *Journal of Memory and Language, 29*, 212–227.

Golding, J. M., Long, D. L., & MacLeod, C. M. (1994). You can't always forget want you want: Directed forgetting of related words. *Journal of Memory and Language, 33*, 493–510.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 193–225). San Diego, CA: Academic Press.

Jacoby, L. L., Kelley, C., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology, 56*, 326–338.

Johnson, D. A. (1971). Pupillary responses during a short-term memory task: Cognitive processing, arousal, or both? *Journal of Experimental Psychology, 90*, 311–318.

Johnson, H. M. (1994). Processes of successful intentional forgetting. *Psychological Bulletin, 116*, 274–292.

Jongeward, R. H., Woodward, A. E., & Bjork, R. A. (1975). The relative roles of input and output mechanisms in directed forgetting. *Memory & Cognition, 3*, 51–57.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Wheeler, V. (1996). On resisting the temptation for simplification: Counterintentional effects of stereotype suppression on social memory. *Social Cognition, 14*, 1–20.

Murdock, B. B. (1963). Short-term retention of single paired associates. *Journal of Experimental Psychology, 65*, 433–443.

Muther, W. S. (1965). Erasure or partitioning in short-term memory. *Psychonomic Science, 3*, 429–430.

Myers, L. B., Brewin, C. R., & Power, M. J. (1996). *Repressive coping and the directed forgetting of emotional material.* Manuscript submitted for publication.

Nagge, J. W. (1935). An experimental test of the theory of associative interference. *Journal of Experimental Psychology, 18*, 663–682.

Neumann, E., Cherau, J. F., Hood, K. L., & Steinnagel, S. I. (1993). Does inhibition spread in a manner analogous to spreading activation? *Memory, 1*, 81–105.

Neumann, E., & DeSchepper, B. G. (1992). An activation based fan effect: Evidence for an active suppression mechanism in selective attention. *Canadian Journal of Psychology, 46*, 1–40.

Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology, 58*, 193–198.

Ratliff, M. M. (1938). The varying functions of affectively toned olfactory, visual and auditory cues in recall. *American Journal of Psychology, 51*, 695–701.

Reed, H. (1970). Studies of the interference processes in short-term memory. *Journal of Experimental Psychology, 84*, 452–457.

Richardson-Klavehn, A., & Bjork, R. A. (1988). Measures of memory. *Annual Review of Psychology, 39*, 475–543.

Roper, K. L., & Zentall, T. R. (1993). Directed forgetting in animals. *Psychological Bulletin, 113*, 513–532.

Shaw, J. S., Bjork, R. A., & Handal, A. (1995). Retrieval-induced forgetting in an eyewitness paradigm. *Psychonomic Bulletin and Review, 2*, 249–253.

Timmins, W. K. (1973). Repetition of intentionally forgotten items. *Journal of Verbal Learning and Verbal Behavior, 12*, 168–173.

Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review, 101*, 34–52.

Wegner, D. M., Eich, E., & Bjork, R. A. (1994). Thought suppression. In D. Druckman & R. A. Bjork (Eds.), *Learning, remembering, believing: Enhancing human performance* (pp. 277–293). Washington, DC: National Academy Press.

Wegner, D. M., Schneider, D. J., Carter, S. R., III, & White, T. L. (1987). Paradoxical effects of thought suppression. *Journal of Personality and Social Psychology, 53*, 5–13.

Weiner, B., & Reed, H. (1969). Effects of the instructional sets to remember and to forget on short-term retention: Studies of rehearsal control and retrieval inhibition (repression). *Journal of Experimental Psychology, 79*, 226–232.

Woodward, A. E., Bjork, R. A., & Jongeward, R. H. (1973). Recall and recognition as a function of primary rehearsal. *Journal of Verbal Learning and Verbal Behavior, 12*, 608–617.

Wright, A. A., Santiago, H. C., Sands, S. F., Kendrick, D. F., & Cook, R. G. (1985). Memory processing of serial lists by pigeons, monkeys, and people. *Science, 229*, 287–289.

Zacks, R. T., & Hasher, L. (1994). Directed ignoring: Inhibitory regulation of working memory. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 241–264). Orlando, FL: Academic Press.

Zacks, R. T., Radvansky, G. A., & Hasher, L. (1996). Studies of directed forgetting in older adults. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 143–156.