



SCHOOL OF AVIATION
Lund University

The hindsight bias is not a bias and not about history

Sidney W. A. Dekker

Technical Report 2004-01

Lund University School of Aviation
Address: 260 70 Ljungbyhed, Sweden
Telephone: +46-435-445400
Fax: +46-435-445464
Email: research@tfhs.lu.se

Abstract. The hindsight bias is seen as a large obstacle to learning from incidents. Almost all explanations of the hindsight bias focus on how it distorts historical explanation. But perhaps the hindsight bias is not about history and not a bias. It may rather be about controlling the future. The almost inevitable urge to highlight and oversimplify past choice moments (where people went the wrong way), the drive to identify “errors”, is forward looking, not backward looking. The hindsight bias may represent an oversimplification of history that primes us for complex futures and allows us to project simple models of past lessons onto those futures, lest history repeats itself. This means that for making progress on safety and learning from incidents, people’s retrospective reconstruction, and the hindsight bias, should not be seen and combatted as the primary phenomenon.

Introduction

The hindsight bias is one of the most consistent “biases” in psychology and it has a profound influence on how we understand and learn from incidents. Fischhoff (1975) reported on a series of experiments that demonstrated how knowledge of outcome increases the postdicted likelihood of reported events, and changes the perceived relevance of event-descriptive data. In other words, hindsight makes that people overestimate. People are unaware of the effects of the hindsight bias, and consistently overestimate what or others could have known without outcome knowledge. As Weick (1995, p. 28) puts it, “people who know the outcome of a complex prior history of tangled, indeterminate events, remember that history as being much more determinant, leading ‘inevitably’ to the outcome they already knew”. Reason (1990) summarizes how hindsight changes past indeterminacy and complexity into order, structure, and oversimplified causality. Effects of the hindsight bias, especially on historical explanation, have been called “creeping determinism” (Fischhoff, 1975). Outcomes that are deemed improbable *ex ante* are judged overdetermined *ex post* (Tetlock, 1998).

Virtually all reports of the hindsight bias, whether experimental (e.g. Fischhoff, 1975; Carroll, 1978; Hawkins & Hastie, 1990; Koehler, 1991) or historical (e.g. Roese & Olson, 1996), focus on the *ex post facto* distortions brought on by knowledge of outcome, relative to uncontaminated *ex ante* judgments about the same outcome or events. The delta between these two is what makes the hindsight bias a *bias*: a departure from a rational norm. Holding this rational norm, judgments of outcome likelihood and event importance are veridical or at least more “realistic”, that is, they map more closely onto the “real” probability that an outcome will occur, and assign fairer weights to the various factors that contribute to that outcome. In research on the hindsight bias, the prevailing question is how knowledge of historical events (some outcome) influences the judged likelihood and importance of antecedent historical events. As a result of this focus, effects of hindsight are cast predominantly in a language of memory, where it is shown how the content of memory deviates systematically and predictably from the unmodified ideal; i.e. from some abstract, rational norm. For example, people “remember” a history as much more determinate (Weick, 1995).

While consistently demonstrated, the contentually subjectivist focus on hindsight as disfiguring memories of historical events precludes other, more adaptive understandings of the effect and purposes of hindsight. Embracing alternative understandings is critical to progress on safety in complex, dynamic work domains, as the hindsight bias has been deemed the greatest obstacle to drawing

constructive lessons from past failures (Fischhoff, 1975; Woods et al., 1994; Dekker, 2002). As long as the hindsight bias keeps intervening in historical understanding, we will oversimplify causality, overrate the contribution of rule- or procedure “violations” (McDonnald *et al.*, 2002), misjudge the prominence or relevance of data available to people at the time, and overestimate the likelihood of the outcome (and people’s knowledge of it). Thus we are doomed to drawing counterproductive conclusions for progress on safety (e.g. removing “bad apples”; dictating even more rules or procedures; restraining human variability through automation). Removing or constraining the effects of the hindsight bias is crucial for understanding how and why people did what they did, and how it could happen again.

For purposes of advancement of safety, concentrating mainly on the historical distortions of the hindsight bias may be addressing symptoms more than causes. Indeed, there is ample evidence that despite exhortations (for example to accident investigators) to understand the world as it looked through the eyes of the decision maker at another place and time, hindsight keeps overruling most attempts at explaining history from the inside-out (e.g. Perrow, 1984; Dekker, 2002). Fischhoff comments that “it is both unfair and self-defeating to castigate decision makers who have erred in fallible systems, without admitting that fallibility and doing something to improve the system” (1975, p. 298). But castigating decision makers appears an indelible part of failure analysis. Any random sample of accident reports will put the hindsight bias on full display. If we, besides monitoring and documenting the effect of hindsight, propose no psychological theory of *why* people disfigure history, we will come up short on solutions for keeping the hindsight bias in check, even if we all agree that it needs to be kept in check. To be sure, political and sociological theories on distorting history are in no short supply (e.g. Weber, 1949; Roese & Olson, 1996), and effects predicted by them may well change historical renderings through motivational, cultural and emotional influences even in smaller scale settings (precisely the things Fischhoff endeavoured to remove or at least experimentally control). Accident investigations are prone to such effects, as Perrow (1984) has noted:

“Formal accident investigations usually start with an assumption that the operator must have failed, and if this attribution can be made, that is the end of serious inquiry. Finding that faulty designs were responsible would entail enormous shutdown and retrofitting costs; finding that management was responsible would threaten those in charge, but finding that operators were responsible preserves the system, with some soporific injunctions about better training.” (p. 146)

But while some psychological specification of reasons for distorting effects have been put forward, for example heuristics of availability or representativeness (Tversky & Kahneman, 1974), a comprehensive functionalist psychological account has not yet been developed. Yet one key to making progress on safety may lie precisely in such a functionalist account. Why, psychologically (not politically), is it that people distort history the way they do? Answers to this question can help direct efforts to improve historical renderings of safety-critical events, thereby potentially enhancing the learning leverage extracted from them.

In keeping with recent anti-normativist traditions in psychology, one can see the hindsight bias not as a bias in the sense of departing from realism or from some rational norm. Rather, the “distortions” brought on by hindsight are inherently rational because they carry ecological utility for organisms continually adapting inside dynamic, complex and risky environments. Hindsight in this sense is not

directed at historical explanation (which makes its distortions ill-adapted and dysfunctional) but about attaining perceptions of being able to control the future (which makes its distortions adaptive and highly functional). This idea is strengthened by finding how the hindsight “bias” can override another bias in psychology related to historical and causal explanation: the fundamental attribution error. Rather than making the perspectives of actor and observer diverge (as the fundamental attribution error predicts), hindsight brings them back together. First, though, one basic ingredient of the arguments needs to be presented: errors as *ex post facto* constructs rather than stable facts of some objective reality. This ingredient is not only borne out by research on the hindsight bias, but also necessary for making further progress on keeping the hindsight bias in check.

“Errors” as *ex post facto* constructs

The new view and avoiding hindsight

A key commitment of the new view to human error (Woods et al., 1994; Dekker, 2002) is to understand why it made sense for people to do what they did. A premise is that system goals and individual goals overlap; that people do not come to work to do a bad job. Behavior is rational within situational contexts. As historian Barbara Tuchman puts it:

”Every scripture is entitled to be read in the light of the circumstances that brought it forth. To understand the choices open to people of another time, one must limit oneself to what they knew; see the past in its own clothes, as it were, not in ours” (1981, p. 75).

This position turns the exigent social and operational context into the only legitimate interpretive device. This context becomes the constraint on what meaning we, who were not there when it happened, can now give to past controversial assessments and actions. Historians are not the only ones to encourage this switch, this inversion of perspectives, this persuasion to put ourselves in the shoes of other people. In hermeneutics it is known as the difference between exegesis (reading out of the text) and eisegesis (reading into the text). The point is to read out of the text what it has to offer about *its* time and place, not to read into the text what we want it to say or reveal now. Jens Rasmussen points out that if we cannot find a satisfactory answer to questions such as “how could they not have known?”, then this is not because these people were behaving bizarrely (see Vicente, 1999). It is because *we* have chosen the wrong frame of reference for understanding their behavior. The frame of reference for understanding people’s behavior is their own normal, individual work context, the context they are embedded in and from which point of view the decisions and assessments made are mostly normal, daily, unremarkable, perhaps even unnoticeable. A challenge is to understand how assessments and actions that from the outside look like “errors” become neutralized or normalized so that from the inside they appear non-remarkable, routine, normal.

If we want to understand why people did what they did, then the adequacy of the insider’s representation of the situation cannot be called into question. The reason

is that there are no objective features in the domain on which we can base such a judgment. In fact, as soon as we make such a judgment, we have imported criteria from the outside—from another time and place, from another rationality. Ethnographers have always championed the point of view of the person on the inside. Emerson, as did Rasmussen, advised that instead of using criteria from outside the setting to examine mistake and error, we should investigate and apply local notions of competent performance that are honored and used in particular social settings (Vaughan, 1999). This excludes generic rules and motherhoods (e.g. “pilots should be immune to commercial pressures”). Such “criteria” ignore the subtle dynamics of localized skills and priority setting, they run roughshod over what would be considered “good” or “competent” or “normal” from inside actual situations. Indeed, such criteria impose a rationality from the outside, impressing a frame of context-insensitive, idealized concepts of practice upon a setting where locally tailored and subtly adjusted criteria rule instead.

Despite all of this, efforts to pull away from normativism or rationalism in human factors have lived a tortured and only very partially successful history. For example, Snook (2002) suggests removing “decision making” from the vocabulary of investigations altogether. It would be an additional way to avoid counterfactual reasoning and judgmentalism, as decisions that eventually led up to a bad outcome all too quickly become “bad” decisions (p. 206):

“Framing such tragedies as decisions immediately focuses our attention on an individual making choices..., such a framing puts us squarely on a path that leads straight back to the individual decision maker, away from the potentially powerful contextual features and right back into the jaws of the fundamental attribution error. ‘Why did they decide ...?’ quickly becomes ‘Why did they make the wrong decision?’. Hence, the attribution falls squarely onto the shoulders of the decision maker and away from potent situational factors that influence action. Framing the...puzzle as a question of meaning rather than deciding shifts the emphasis away from individual decision makers toward a point somewhere ‘out there’ where context and individual action overlap.

Yet sensemaking is not immune to retrospective, disfiguring pressure either. If what made sense to the person inside the situation still makes no sense given the outcome, then outside observers hasten to point that out (“they lost situation awareness”). Even in sensemaking, the hindsight bias is an ever-present risk. Perhaps the pull in the direction of the position of retrospective outsider is irresistible, inescapable, whether we make lexical adjustments in our investigative repertoire or not. Even with the potentially judgmental notion of “decision making” removed from the forensic psychological toolbox, it remains incredibly difficult to “see the past in its own clothes, not in ours”.

“Errors” as constructs

When looked at from the position of retrospective outsider, “human errors” can look so very real, so compelling. They failed to notice, they did not know, they should have done this or that. But from the point of view of people inside the situation, as well as potential other observers, this same “error” is often nothing more than normal work. To paraphrase Giddens, “errors” represent an active, corrective intervention in (immediate) history. It is impossible for us to give a

mere chronicle of our experiences or observations: our assumptions, past experiences and future aspirations make that we impress a certain organization on that which we just went through or saw. “Errors” are a powerful way to impose structure onto past events. “Errors” are a particular way in which we as observers or actors reconstruct the reality we just experienced. Such reconstruction, however, inserts a severe discontinuity between past and present. The present was once an uncertain, perhaps vanishingly improbable future (Tetlock, 1998). Now we see it as the only plausible outcome of a deterministic past.

But errors get “real” only when we step outside, or set ourselves outside, the stream of experience in which they occurred. While carrying out tasks in situated contexts, people do not make “errors”—their errors are not discovered (either by themselves or by outside observers) until after the fact (even if immediately after the fact, where intention and outcome have deviated). It is this outsider perspective (even if we were insiders only seconds ago) that endows history, even immediate history, with a determinism it lacked when it was still unfolding. “Errors”, then, are *ex post facto* constructs. The research base on the hindsight bias contains some of the strongest evidence for this. “Errors” are not empirical facts. As an observed “fact”, the error only exists by virtue of the observer and his position outside the stream of experience. The error does not exist because of some objective empirical reality in which it putatively takes place, since there is no such thing and if there was, we could not know it. Just as any act of observation changes the observed, our very observations of the past inherently intervene in reality, converting complex histories into more linear, more certain, fair, and disambiguated chronicles. Errors are the result of outside observers squeezing now-known events into the most plausible, or convenient deterministic scheme. In the research base on hindsight, it is not difficult to see how such retrospective restructuring embraces a liberal take on the history it aims to recount. The distance between reality as portrayed by a retrospective observer and as experienced by those who were there (even if these were once the same people) grows substantially with the rhetoric and discourse employed and the investigative practices used. Why do “errors” fulfill such an important function in our reconstructions of history, of even our own histories? Seeing errors in history may actually have little to do with historical explanation. Retrospective recounting tells us more about the observer and his or her goals and aspirations, than it does about historical events.

The hindsight bias as forward looking

In 1971, Jones & Nisbett documented the tendency of actors to attribute their actions to external situational causes, whereas observers (by contrast) attribute the same actions to causes internal to the actor(s). They called this tendency the “actor-observer divergence”. A few years later, the “fundamental attribution error” described how observers underestimate the impact of situational forces and overestimate the importance of internal dispositional factors (Ross, 1977). Such misestimations appear as a constant in analyses of past breakdowns where outside observers (e.g. accident investigators) are apt to assign reasons for outcome failures to personal shortcomings (e.g. crewmembers “lose situation awareness”, or display “ineffective crew resource management”). In aviation, specific paragraphs in accident reports are even reserved for tracing the potentially broken human components. Investigators have to explore the antecedent 24- and 72-hour histories of the humans who would later be involved in a mishap. Was there alcohol? Was there stress? Was there fatigue? Were there other predispositions in

the form of insufficient proficiency or experience? Were there previous problems in the training or operational record of these people? Were there other distractions or problems? Human error is reduced to some notion of personal “fitness for duty”. This investigative requirement institutionalizes the fundamental attribution error. Such deep-seated reflexes prompt Snook (2000) to comment that the fundamental attribution error is alive and well.

But is it?

The hindsight bias and perceived control over the future

The linearization and simplification that happens in the hindsight bias may be a form of abstraction that allows us to export and project our and others’ experiences onto future situations. Future situations can never be predicted at the same level of contextual detail as the new view encourages us to explain past situations. Predictions are possible only because we have created some kind of “model” for the situation we wish to gain control over, not because we can exhaustively foresee every contextual factor, influence, data point. This model—any model—is an abstraction away from context, an inherent simplification. The model we create—naturally, effortlessly, automatically—after past events with a bad outcome inevitably becomes a model of binary choices, bifurcations and unambiguous decision moments. That is the only useful kind of model we can take with us into the future if we want to guard against the same type of pitfalls and forks in the road.

The hindsight bias, then, is about attaining control, or a perception of control, not about explaining. It is forward-looking, not backward looking. Social psychology consistently points out how perceived uncontrollability of adverse events is among the greatest possible nonspecific human stressors, leading to a host of undesirable effects, including depression and anxiety (Johnson & Sarason, 1978) and even illness (Stern *et al.*, 1982). Such and other social psychological research posits perceived control (note: not real control, but perceived) as a highly adaptive coping mechanism, as an effective mediating variable between threats and their outcome (see e.g. Baum *et al.* (1983) who investigated control perceptions and their effects in the wake of the Three Mile Island nuclear accident). By letting individual decision makers slide into disrepute (they and their actions are uniquely bad), people can relieve the tension between broken beliefs (the system is not safe after all) and fervent hopes that it still is. That the hindsight bias may not be primary, but rather ancillary expressions of more adaptive, locally rational and useful identity-preserving strategies for the ones committing them, is consonant with observations of a range of reasoning “errors”. People keep committing them not because they are logical (i.e. globally rational) or because they only produce desired effects, but because they serve a purpose more important than all of that:

“This dynamic, this ‘striving to preserve identity’, however strange the means or effects of such striving, was recognised in psychiatry long ago. [This phenomenon] is seen not as primary, but as attempts (however misguided) at restitution, at reconstructing a world reduced by complete chaos.” (Sacks, 1998, p. 7).

However “strange the means or effects of such striving”, the hindsight bias allows us an illusion of predicting and avoid future roads to perdition. A constant

response that helps generate perceptions of control, then, has adaptive value. This applies to ourselves and our own failures as much as it applies to our observations of failures of other people. When confronted by failures that occurred to other people, we may imperatively be tripped into vicarious learning, spurred by our own urge for survival: what do I do to avoid that from happening to me? When confronted by our own performance, we have no privileged insight into our own failures, even if we would like to think we do. The past is the past, whether it is our own or somebody else's. Our observations of the past inevitably intervene and change the observed, whosever past it is. This is something that the fundamental attribution error cannot account for. Even here people appear susceptible to reframing past complexity as simple binary decisions; wrong decisions due to personal shortcomings: things *they themselves* missed, things *they themselves* should have done or should not have done.

Hindsight and actor-observer convergence

Hindsight thus blurs the distinction between actor and observer. We are all “observers” of our own performance—this is the only way we can spot errors in the first place: by setting ourselves *outside* the stream of experience. Where actor-observer distinctions blur, actor-observer divergences no longer occur, or at least become difficult to distinguish. Partly as a result of this, the fundamental attribution error is not as consistent in social psychology as the hindsight bias. The psychological perspective taken by an observer can be manipulated, even if people know the outcome, for example by instructions (Gould & Sigall, 1977), which is a hope reverberating in appeals to accident investigators to this day (e.g. Dekker, 2002).

Snook (2000) investigates how, in the fog of post-Gulf war Iraq, two helicopters carrying UN peacekeepers were shot down by American fighter jets. The situation in which the shootdown occurred was full of risk, role ambiguity, operational complexity, resource pressure, slippage between plans and practice. Yet immediately after the incident, all of this gets converted into binary simplicity (a choice to err or not to err) by DUKE—the very command onboard the airborne control center whose job it was not to have such things happen. Allowing the fighters to shoot down the helicopters was *their* “error”, yet they do not blame context at all, as the fundamental attribution error predicts they should. It was said of the DUKE that immediately after the incident:

“he hoped we had not shot down our own helicopters and that he couldn't believe anybody could make that dumb of a mistake” (Snook, 2000, p. 205).

It is DUKE himself who blames his own dumb mistake. As with the “errors” in the previous chapter, the dumb mistake is something that jumps into view only with knowledge of outcome, its “mistakenness” a function of the outcome, its “dumbness” a function of the severity of the consequences. While doing the work, helping guide the fighters, identifying the targets, all DUKE was doing was his job. Normal work. He was not sitting there making “dumb mistakes”. They are a product of hindsight, his own hindsight, directed at his own “mistakes”. The fundamental attribution error does not apply. It is overridden.

The fighter pilots too, engage in self-blame, literally converting the ambiguity, risk, uncertainty and pressure of their encounter with potentially hostile helicopters

into a linear series of decision errors, where they repeatedly and consistently took wrong turns on their road to perdition (*we* (1) misidentified, (2) engaged and (3) destroyed):

“Human error did occur. We misidentified the helicopters; we engaged them; and we destroyed them. It was a tragic and fatal mistake.” Tiger 02 quoted in Snook (2000, p. 205).

Again, the fundamental attribution error makes the wrong prediction. If it were true, then these fighter pilots would tend to blame context for their own “errors”. Indeed, it was a rich enough context—fuzzy, unfamiliar, multi-player, time-pressurized, risky—with plenty of “blameworthy” factors to go around. If that is where you would look. Yet these fighter pilots do not. *We* misidentified, *we* engaged, *we* destroyed. The pilots had the choice not to; in fact they had a series of three choices not to instigate a tragedy. But they did. Human error did occur. Of course, elements of self-identity and control are wrapped up in such an attribution, a self-identity for which fighter pilots may well be poster children.

It is interesting to note that the tendency to convert past complexity into binary simplicity—into twofold choices to identify correctly or incorrectly, to engage or not, to destroy or not—overrides the fundamental attribution error. This confirms the role of the hindsight bias as a catalyst for learning. Learning (or having learned) expresses itself most clearly by doing something differently in the future, by deciding or acting differently, by removing one’s link in the accident chain, as fighter pilot Tiger 02 puts it:

“Remove any one link in the chain and the outcome would be entirely different. I wish to God I could go back and correct my link in this chain—my actions which contributed to this disaster.” (Tiger 02, quoted in Snook, 2000, p. 205).

We cannot undo the past. We can only undo the future. But undoing the future becomes possible only when we have abstracted away past failures, when we have decontextualized them, stripped them, cleaned them from the fog and confusion of past contexts, highlighted them, blown them up into obvious choice moments that we, and others, had better get right next time around. *Prima facie*, the hindsight bias is about misassessing the contributions of past failings to bad outcomes. But if the phenomenon is really as robust as it is documented to be and if it actually manages to override the fundamental attribution error, it is probably the expression of more primary mechanisms running right beneath its surface.

The hindsight bias is a meaningful adaptation. It is not about past failures. It is about preventing future ones. In preparing for future confrontations with situations where we or others might “err” again, and do not want to, we are in some sense taking refuge from the banality of accidents thesis. The thought that accidents emerge from murky, ambiguous, everyday decision making renders us powerless to do anything meaningful about it. This is where the hindsight bias is so fundamentally adaptive. It highlights for us where we could fix things (or where we think we could fix things), so that the bad thing does not happen again. The hindsight bias is not a bias at all, in the sense of a departure from some rational norm. The hindsight bias *is* rational. It in itself represents and sustains rationality. We *have* to see the past as a binary choice, or a linear series of binary choices, because that is the only way we can have any hope of controlling the future. There is no other basis for learning, for adapting. Even if those

adaptations may consist of rather coarse adjustments; of undamped and overcontrolling regulations. Even if these adaptations occur at the cost of making oversimplified predictions. But making oversimplified predictions of how to control the future is apparently better than having no predictions at all.

Conclusion

The hindsight bias is not about history and not a bias. Rather it is about controlling the future. Or, more specifically, about giving oneself the perception of being able to control the future. Retrospective reconstruction, and the hindsight bias, should not be seen (and combatted) as the primary phenomenon. Instead, it represents and serves a larger purpose, responding to a highly pragmatic concern. The almost inevitable urge to highlight and oversimplify past choice moments (where people went the wrong way), the drive to identify “errors”, is forward looking, not backward looking. The hindsight bias may represent an oversimplification of history that primes us for complex futures and allows us to project simple models of past lessons onto those futures, lest history repeats itself.

Fischhoff concludes how, because of the hindsight bias, “the very outcome knowledge which gives us the feeling that we understand what the past was all about may prevent us from learning anything from it.” (1975, p. 299). But is the hindsight bias destructive to learning, or is it part of trying to learn? Fischhoff’s original paper carries the title “hindsight does not equal foresight”. Yet, while indeed not equal, perhaps hindsight is about foresight after all. Perhaps the “biasing” part of hindsight (the oversimplifications, the creeping determinism, the blowing out of proportion of particular data or events leading up to a bad outcome) is about creating a particular type of abstracted foresight that affords us better prediction of and control over future outcomes.

References

- Baum, A., Fleming, R., & Singer, J. E. (1983). Coping with victimization by technological disaster. *Journal of Social Studies, 39*, 117-138.
- Carroll, J. S. (1978). The effect of imagining an event on expectations for the event: An interpretation in terms of the availability heuristic. *Journal of experimental social psychology, 14*, 88-96.
- Dekker, S. W. A. (2002). *The field guide to human error investigations*. Bedford, UK: Cranfield University Press/Aldershot, UK: Ashgate.
- Fischhoff, B. (1975). Hindsight is not foresight: The effect of outcome knowledge on judgement under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance, 1*(3), 288-299.
- Gould, R. & Sigall, H. (1977). The effects of empathy and outcome on attribution: An examination of the divergent perspectives hypothesis. *Journal of Experimental Social Psychology, 13*, 480-491.
- Hawkins, S. A., & Hastie, R. (1990). Hindsight: Biased judgments of past events

after the outcomes are known. *Psychological Bulletin*, 107, 311-327.

Johnson, J. S., & Sarason, I. G. (1978). Life stress, depression and anxiety: Internal-external control as a moderator variable. *Journal of Psychosomatic Research*, 22, 205-208.

Jones, E. E., & Nisbett, R. E. (1971). The actor and the observer: Divergent perceptions of the causes of behavior. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, and B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior*, pp. 79-94. Morristown, NJ: General Learning Press.

Koehler, D. J. (1991). Explanation, imagination, and confidence in judgment. *Psychological Bulletin*, 110, 499-519.

McDonald, N., Corrigan, S., & Ward, M., 2002. Well-intentioned people in dysfunctional systems. *Keynote presented at 5th workshop on human error, safety and systems development*, Newcastle, Australia.

Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. New York, NY: Basic books.

Reason, J. T. (1990). *Human error*. Cambridge, UK: Cambridge University Press.

Roese, N., & Olson, J. (Eds.) (1996). *What might have been: The social psychology of counterfactual thinking*. Mahwah, NJ: Lawrence Erlbaum Associates.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology*, pp. 173-220. New York: Academic Press.

Sacks, O. (1998). *The man who mistook his wife for a bat*. New York, NY: Touchstone.

Snook, S. A. (2000). *Friendly fire: The accidental shootdown of US Black Hawks over Northern Iraq*. Princeton, NJ: Princeton University Press.

Stern, G. S., McCants, T. R., & Pettine, P. W. (1982). Stress and illness: Controllable and uncontrollable life events' relative contributions. *Personality and Social Psychology Bulletin*, 8, 140-145.

Tetlock, P. E. (1998). Social psychology and world politics. In D. Gilbert, S. Fiske & G. Lindzey (Eds.), *Handbook of social psychology*. New York: McGraw Hill.

Tuchman, B. W. (1981). *Practicing history: Selected essays*. New York, NY: Norton.

Tversky A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.

Vaughan, D. (1999). The dark side of organizations: Mistake, misconduct, and disaster. *Annu. Rev. Sociol.*, 25, 271-305.

Vicente, K. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum Associates.

Weber, M. (1949). Objective possibility and adequate causation in historical explanation. In *The methodology of the social sciences*, pp. 164-188. Glencoe, Ill: Free Press of Glencoe.

Weick, K. E. (1995). *Sensemaking in organizations*. London: Sage.

Woods, D. D., Johannesen, L. J., Cook, R. I. & Sarter, N. B. (1994). *Behind human error: Cognitive systems, computers and hindsight*. Dayton, OH: CSERIAC.