THE INTERNATIONAL JOURNAL OF AVIATION PSYCHOLOGY, *8*(1), 47–63

Automation Bias: Decision Making and Performance in High-Tech Cockpits

Kathleen L. Mosier
San Jose State University and NASA Ames Research Center
Moffett Field, CA

Linda J. Skitka
University of Illinois at Chicago

Susan Heers
Western Aerospace Labs and NASA Ames Research Center
Moffett Field, CA

Mark Burdick
University of Illinois at Chicago

Automated aids and decision support tools are rapidly becoming indispensable tools in high-technology cockpits and are assuming increasing control of "cognitive" flight tasks, such as calculating fuel-efficient routes, navigating, or detecting and diagnosing system malfunctions and abnormalities. This study was designed to investigate *automation bias,* a recently documented factor in the use of automated aids and decision support systems. The term refers to omission and commission errors resulting from the use of automated cues as a heuristic replacement for vigilant information seeking and processing. Glass-cockpit pilots flew flight scenarios involving automation events or opportunities for automation-related omission and commission errors. Although experimentally manipulated accountability demands did not significantly impact performance, post hoc analyses revealed that those pilots who reported an internalized perception of "accountability" for their performance and strategies of interaction with the automation were significantly more likely to double-check automated functioning against other cues and less likely to commit errors than those

---

Requests for reprints should be sent to Kathleen Mosier, Department of Psychology, San Francisco State University, 1600 Holloway, San Francisco, CA 94132.

> who did not share this perception. Pilots were also likely to erroneously "remember" the presence of expected cues when describing their decision-making processes.

In the "glass cockpit," as well as in domains as diverse as medical diagnosis, nuclear power plants, manufacturing and process control, and Air Traffic Control (ATC), human decision-makers are being exposed to and are required to utilize automated aids to perform their jobs and make decisions. As these and other environments become more complex and data-intensive, the use of automated aids and decision support tools is likely to become even more commonplace and critical to performance. Aircraft flight management systems, for example, are designed not only to keep the aircraft on course, but also to assume increasing control of cognitive flight tasks, such as calculating fuel-efficient routes, navigating, or detecting and diagnosing system malfunctions and abnormalities. An inescapable facet of these automated aids is that they change the way pilots perform tasks and make decisions. Most of the changes have proven to be beneficial. The advantages these systems offer in terms of increased efficiency and data storage and manipulation, for example, is self-evident—computers can assimilate more information and process it faster than humans.

Other repercussions of automated aids, however, are less positive. Several researchers have documented problems in the use of advanced automated systems in aviation, including mode misunderstandings and mode errors, failures to understand automation behavior, confusion or lack of awareness concerning what automated systems are doing and why, and difficulty tracing the functioning or reasoning processes of automated agents (e.g., Billings, 1996; Sarter & Woods, 1993). Additionally, automated systems introduce opportunities for new decision-making heuristics and associated biases. *Automation bias,* a recently documented factor in the use of automated aids and decision support systems, refers to errors resulting from the use of automated cues as a heuristic replacement for vigilant information seeking and processing. This article describes research examining the nature and implications of automation bias in the cockpit, its correlates, and possible countermeasures.

THE EVOLUTION OF AUTOMATION BIAS

Automation—Just Another Cue?

Pilots regularly rely upon probabilistic information to evaluate what they cannot access directly (Wickens & Flach, 1988). In traditional aircraft, crewmembers are trained and develop skills in situation assessment through the use of system and environmental cues. Fuel gauges, engine indicators, weather forecasts, altimeters, and other instruments provide cues with which the human operator can assess and make decisions concerning the state of the aircraft and its surrounding environment. In most situations, processing is facilitated by intercorrelations among cues (Wickens & Flach, 1988). Expert pilots exploit this intercorrelation in their cue search and may perceive sets of correlated cues as a single perceptual chunk (Ebbeson & Konecni, 1980). They frequently use feature matching and pattern matching as diagnostic strategies. Pilots know and look for *patterns* or *combinations* of cues that are most ecologically valid, reliable, or relevant for diagnosing particular situations, and they are able to incorporate contextual information to formulate a workable action plan based on their assessment of these cues (e.g., Kaempf & Klein, 1994).

When an automated aid is introduced, however, it disrupts the pattern of cue utilization. Automated aids present new, powerful cues. These cues are readily available, are widely believed to be accurate, and are a highly salient source of information. The computational, system observation, and diagnostic capabilities of automated aids are advertised as being superior to that of their human operators. Additionally, the implementation of automated aids may be accompanied by diminished access to other traditionally utilized cues. Research on the use of probabilistic cues, such as those found in the aircraft cockpit, has demonstrated that decision makers may focus on salient cues and ignore other critical but less obvious information, especially under time pressure (Wickens & Flach, 1988), and also that the activation of one "most powerful" cue may be sufficient for people to confidently generate a decision, particularly if time is short (Gigerenzer, Hoffrage, & Kleinbölting, 1991). In combination, these factors contribute to the perception that automated aids do not provide "just another cue," but rather are more important, more diagnostic, and more reliable than previously utilized, conventional information sources.

Automation as Heuristic

The availability of automation and automated decision aids feeds into the general human tendency to travel the road of least cognitive effort. Typically, people try to engage in the least amount of cognitive work they can get away with (Fiske & Taylor, 1994), will prefer strategies that are easy to justify (and do not involve analyzing relative weights or numerical computations) and will often utilize heuristics (or cognitive shortcuts) to reduce effort and information load. To be successful, heuristics must provide quick and simple ways of dealing with a great deal of information and must be reasonably accurate most of the time (for reviews, see Kahneman, Slovic, & Tversky, 1982, and Fiske & Taylor, 1994).

Automated aids in high-tech environments, such as the aircraft cockpit, provide decision makers with a new heuristic, or shortcut, for decision making and task performance. These systems are, in fact, designed specifically to decrease human workload by performing many cognitive tasks—including information synthesis, system monitoring, diagnosis, planning, and prediction—in addition to controlling the physical placement of the aircraft. However, the presence of automated cues also diminishes the likelihood that decision makers will either make the cognitive effort to seek other diagnostic information or process all available information in cognitively complex ways. Parasuraman and Riley (1997) describe this tendency toward overreliance as "automation misuse." In addition, automated cues increase the probability that decision makers will cut off situation assessment prematurely when prompted to take a course of action by a computer or automated aid. For example, Layton, Smith, and McCoy (1994) examined pilot use of a graphical flight planning tool and found that computer generation of a suggestion or recommendation early in the course of problem evaluation significantly impacted decision processes and biased pilots towards the computer's suggestion, even when the computer's brittleness (e.g., in terms of an inadequate model of the "world") resulted in a poor recommendation with potential adverse consequences.

To the extent that other available information besides automation *is* processed, decision makers may show either assimilation or discounting biases. Other indicators may be interpreted as being more consistent with the automated information than they really are (assimilation), especially if other cues in the environment are ambiguous (e.g., Darley & Gross, 1983; Glick, Zion, & Nelson, 1988). Conversely, cues that are completely inconsistent with automated information may be discounted. Additionally, confirmational biases lead information processors to overattend to consistent information and ignore other data or to process new information in a manner that confirms one's preexisting belief and avoids recognition of the need for its revision (Hamilton, 1981; Wickens & Flach, 1988). To the extent that decision makers are prompted toward a particular problem or diagnosis when automated aids bring it to their attention, other information is less likely to be used to create a new impression or to modify the impression of the problem created by the automation.

Because automated aids and decision support tools, when used correctly, are extremely consistent and provide cues that are highly correlated with other information, using them heuristically will generally be effective. Diagnoses are made correctly, power plants function efficiently, and airplanes fly safely using automated aids. However, indiscriminate use of these tools will result in errors, as would indiscriminate use of any other heuristic; in some cases, indiscriminate use may have the paradoxical effect of increasing errors rather than eliminating them. In particular, inappropriate usage of automated systems in decision making and task performance may result in *automation bias,* a term describing errors made when human operators use automated cues as a heuristic replacement for vigilant information seeking and processing (Mosier & Skitka, 1996). Two types of automation-related errors, *omission errors* and *commission errors,* are negative consequences of this bias.

Automation omission errors result when decision makers do not take appropriate action because they are not informed of an imminent problem or situation by automated aids. For example, a China Airlines B747–SP, flying at 41,000 ft., lost power in its #4 engine. The autopilot, which was set for pitch guidance and altitude hold, attempted to correct for the loss by holding the left wing down, masking the approaching loss of control of the airplane. The crew did not realize that there was a problem with the engine and took no action to deal with it. When the captain disengaged the autopilot, the airplane rolled to the right, yawed, then entered a steep descent in clouds. Extensive damage occurred during descent and recovery (NTSB Report AAR–86–03, in Billings, 1996). Evidence of the tendency to make automation-related omission errors has also been found in Aviation Safety Reporting System reports, as discussed by Mosier, Skitka, and Korte (1994). In a nonrandom sample of 166 events, they found that the most likely flight phase for omission errors to occur was the cruise phase. In many of the incidents, the crews set up a system (not always correctly) to perform a task and then trusted the system to do it. Once the crews delegated a task to the automation, they did not check other cues to catch inconsistencies or mistakes in task performance. The 1983 incident in which a Korean Airlines B–747 was shot down by Soviet fighters included a similar causal chain. In this case, the crew selected a magnetic heading and followed it throughout the flight rather than coupling the navigational system's inertial reference system to the autopilot. The flight did not follow its originally planned flight path but rather maintained the selected heading until it was shot down. The crew relied on automation that had been inappropriately set up, and they never checked their progress manually, allowing the flight to stray well into Soviet airspace ("Analysis of flight data," 1993).

Automation commission errors are errors made when decision makers inappropriately follow automated information or directives (e.g., when other information in the environment contradicts or is inconsistent with the automated cue) that have begun surfacing recently as by-products of automated systems. Some airlines, for example, are dealing with automation-induced commission errors in pilots' responses to an overly-sensitive warning system that often indicated a cargo compartment fire when in fact there was none. Even though the presence or absence of the fire was verifiable on other cockpit indicators, many pilots consistently followed the warning directive and flooded their cargo bays with fire retardant.

Experimental evidence of automation-induced commission errors was provided by a full-mission simulation in the NASA Ames Advanced Concepts Flight Simulator (ACFS; Mosier, Palmer, & Degani, 1992). During takeoff, crews received contradictory fire indications. An auto-sensing electronic checklist suggested that the crew shut down the #1 engine, which was supposedly on fire.

Traditional engine parameters indicated that the #1 engine was recovering and that the #2 engine was actually more severely damaged. Seventy-five percent of the crews in the auto-sensing condition incorrectly shut down the #1 engine, whereas only 25% with the traditional paper checklist did likewise. Analysis of the crews' audiotapes indicated that crews in the automated condition tended to discuss less information before deciding whether to shut down the engine, suggesting that automated cues can curtail information search.

The implementation of automated aids in the aircraft cockpit, then, carries the potential for possible adverse influence, in that the situation assessment process may get short-circuited by automated cues; that is, operators may stop short at the automated display or not double-check the operation of the automated system via other available cues. If we are to anticipate and reduce automation bias in the cockpit, we must discover what patterns of automation bias will be displayed by flight crews in a flight situation and how these errors can be ameliorated. Questions related to this are: What cognitive processes are activated when crews utilize information from automated decision aids? Are there internal or individual variables that affect susceptibility to or resistance against automation bias? What kinds of interventions will be successful in ameliorating automation bias and eliminating resultant errors?

Accountability

A wide body of social psychological research has found that many cognitive biases and resultant errors can be ameliorated by imposing predecisional accountability, which sensitizes decision makers to the need to construct compelling justifications for their choices and *how they make them.* Increasing accountability can successfully mitigate decision-making biases such as primacy effects (Tetlock, 1983), the fundamental attribution error (Tetlock, 1985), over-confidence effects (Tetlock & Kim, 1987), and the "sunk cost" effect (Simonson & Nye, 1992). Accountability demands cause decision makers to employ more multidimensional, self-critical, and vigilant information seeking, as well as more complex data processing; it has also been shown to reduce cognitive "freezing" or premature closure on judgmental problems (Kruglanski & Freund, 1983) and to lead decision makers to employ more consistent patterns of cue utilization (Hagafors & Brehmer, 1983). In sum, accountability increases vigilance in decision making and increases the tendency to use all available information for situation assessment.

The extent to which these effects generalize to human performance situations has only begun to be empirically explored. In a recent study intended to parallel and serve as an analog to the present investigation, Skitka, Mosier, and Burdick (1996) reported overall omission–commission error rates of approximately 55% in a sample of students performing a complex, low-fidelity flight task. In this study, the imposition of accountability demands significantly improved performance. Specifically, making participants accountable—either for the accuracy of responses or for overall performance—reduced omission and commission errors. Because improved performance entailed the verification of automated information against other cues, differences in performance could be directly traced to increased vigilance and attention to all available information. It was apparent, therefore, that imposed accountability had the effect of making operators more likely to verify the accuracy of automated information and less likely to rely heuristically on automated cues.

The use of automated cues as a shortcut in decision making may result in omission or commission errors (automation bias). The imposition of predecisional accountability demands has been found to successfully ameliorate biased decision making in several social psychological contexts, and it has recently been found to increase vigilance and reduce automation-related errors for students in a low-fidelity flight task. It is not clear, however, that these results would generalize to a population of experienced pilots using automation in the cockpit. In this study, investigations were extended to glass-cockpit pilots in a part-task flight environment. We hypothesized that (a) the rate of occurrence of automation-related errors among pilots would be comparable to that found with students, and (b) accountability demands in the aircraft cockpit environment would promote more careful, data-based decision-making strategies and increased vigilance in decision making, making pilots less susceptible to automation bias and more likely to check all relevant information before making decisions or taking actions. We examined patterns of errors associated with automation bias and the relation between errors and accountability, and we gathered preliminary information on cognitive processes associated with this phenomenon.

METHOD

Participants

Participants in this study were 25 commercial glass-cockpit pilots (i.e., pilots of automated aircraft, including Boeing 737–300, 757, 767, 747–400, MD–11). The average age of the pilots was 47, mean total flight experience was 12,370 hr, and the average career flying time was 23 years. All pilots were currently qualified in their aircraft.

Experimental Task and Equipment

The part-task flight simulation facility used in this experiment was modeled after the ACFS at NASA Ames Research Center. It employs two Silicon Graphics color monitors to present glass displays of primary and secondary flight displays, navigation and communication information, and electronic checklists, as well as Engine Indicating and Crew Alerting System (EICAS) and Flight Management System (FMS) instruments (see Figure 1). The displays are similar to what would be found in a B747–400. Pilots interacted with the controls and displays of the aircraft through a touchscreen overlaying the instruments needed to accomplish the flight. They received clearances and information from an air traffic controller, who was stationed in the area outside the simulator booth.

Two facets of the simulator were new to all of the pilots: communications mode and electronic checklists. Communications with ATC were conducted via datalink or digitized communication. Incoming messages were signaled by a chime and were accessed from the flight management computer and displayed on its monitor. Clearances that involved changes in flight parameters or communications frequencies could be automatically loaded into the appropriate control unit by pressing the "LOAD" button. Autoloaded clearances appeared on the relevant display for confirmation. Altitude clearances appeared in the altitude window of the Mode Control Panel (MCP), and a flight level change was initiated if the pilot engaged an altitude change mode (e.g., flight level change or vertical speed). Heading clearances loaded into the heading window of the MCP and initiated the com-

FIGURE 1 Mini-ACFS displays.

manded turn via Heading Select mode. Similarly, speed clearances loaded into the speed window and accomplished the commanded speed via Speed Intervention mode. Commanded speed and altitude changes were also displayed in magenta above the speed and altitude tapes on the Primary Flight Display (PFD). Heading changes were reflected on the Navigation display. Frequency changes became active upon loading and appeared on the communications panel display. Pilots could downlink requests via datalink and were also equipped with headsets so that communications could be conducted orally if desired.

Electronic checklists (see Mosier et al., 1992) for normal and abnormal procedures were displayed on the monitor in front of the pilot, and were accessed through a menu display. Additionally, whenever an EICAS message indicated an abnormal situation (e.g., #1 Engine Fire), touching the message itself would bring up the corresponding checklist. The checklists were "manual-sensing," which meant that the checklist could sense the state of some items after the pilot had touched the display. If the checklist sensed that the item was not completed, the text and triangle indicators turned yellow; if it was sensed as completed, they turned green. The pilot had the option to override skipped or unaccomplished items.

In addition to the aircraft displays, a secondary tracking task (on the on the bottom right-hand corner of the left monitor; see Figure 1) was presented on one of the monitors to provide a means of increasing the workload of the pilots and was incorporated into the automation events described below. This task involved using a joystick to keep displayed cross-hairs inside the boundaries of a blue circle. When the cross-hairs traversed the boundaries of the circle, the circle turned red. Feedback on how much time the subject was able to stay within the target circle was accumulated and displayed to the subject.

Design and Procedures

Pilots were assigned to either the Accountable or NonAccountable condition. Accountable participants were told that: (a) the purpose of the study was to analyze the *performance* of the pilot in heavily automated cockpits; (b) we would be monitoring and evaluating their performance with respect to the use of automated flight systems; (c) we would be collecting performance data; and (d) they would be asked to *justify their performance and strategies in the use of the automated systems* in an interview following the experiment. All of the computer monitors in the lab were left on during the briefing, and a video camera was turned on just prior to the experimental legs. Participants in the NonAccountable condition were told only that the purpose of the study was to analyze the *role* of the pilot in heavily automated cockpits. To reinforce nonaccountability, these pilots were also advised that, due to a breakdown of the data collection computer, we would be unable to collect any performance data that day. They were asked to fly the experimental legs and to contribute important subjective data on written questionnaires after the simulation. No mention of a formal interview or justification of performance was made. Monitors at the experimenter observation station were turned off while participants were in the room, and the video camera was pointed away from the subjects and was not turned on.

Participants were trained individually on each of the components of the experimental task and were given time to practice. Following training, they flew two legs (order was counterbalanced): Los Angeles (LAX) to San Francisco (SFO), and SFO to Sacramento (SMF). The flight route was preloaded into the FMS prior to beginning the trial. Subjects were instructed to communicate with ATC through datalink, although voice communications would be available to them if needed. Most clearances from ATC (e.g., a change in altitude, speed, heading, or frequency) could be autoloaded into the appropriate flight system, and correct loading could be verified by checking the MCP, PFD, or navigation display as appropriate. During several training legs, all participants practiced autoloading each type of clearance, and verifying proper loading and execution on the appropriate display. Pilots manually performed the secondary tracking task from SFO–SMF whenever they were above 5,000 ft. The secondary task was "automated" below 5,000 ft on the SFO–SMF leg, and at all times during the LAX–SFO leg. While this task was automated, the cross-hairs remained inside the boundaries of the circle without pilot input.

Automation events. Four automation failures during these legs offered the possibility for pilots to make omission errors if they did not verify proper automation functioning:

1. An altitude clearance misloaded into the flight control system, and was reflected by incorrect numbers on the MCP and PFD.
2. A commanded heading change loaded correctly but was incorrectly executed by the flight system, and the improper execution was reflected on the navigational display.
3. A frequency change misloaded into the communications system and was reflected by incorrect numbers on the communications display.
4. The tracking task automation failed at 7,000 ft during the LAX–SFO flight, and was signaled by the boundary circle turning red as the cross-hairs traversed it.

Verification information for each flight-related event was available on the appropriate display, as it would be in the aircraft.

One event offered the possibility for a commission error and concerned a false automated warning of an engine fire. The pilots were presented with an automated message suggesting that an engine was on fire, which was contradicted by normal engine parameters and the absence of any other indicators. The pilots had to determine whether there really was a fire and decide whether or not to shut down the supposedly affected engine. This event was intended to be a conceptual replication of the false engine fire event in the Mosier et al. (1992) study. On approach to SFO, at an altitude of approximately 250 ft., pilots were commanded (verbally) to execute a missed approach and go-around due to traffic on the runway. On climb-out, an electronically-generated "Engine Fire" message appeared on the EICAS. Touching the message brought up the electronic checklist for #1 engine fire and shutdown.

During training, subjects were reminded that an engine fire would be indicated by a number of cues, including an *aural warning,* a red *master warning light* or yellow *caution light* on the right monitor (top right corner), an illuminated red *fire handle* on the right monitor (top right corner), possible *deterioration of engine indicators* shown on the left monitor, and an *electronically generated EICAS message* on the right monitor. All of these cues were demonstrated and their locations pointed out as part of the training process. During the experimental missed approach, subjects received *only* the automated EICAS message, which was linked to the electronic checklist, and which was not verified by any of the other expected cues.

Postexperimental questionnaires requested data on flight experience and, in Likert-scale format, probed perceptions of accountability (e.g., belief that their performance and strategies would be evaluated, and that they would have to justify their performance) and attitudes toward the experimental task, computers, and automation (e.g., task difficulty, comfort and confidence during the experiment, reliability of computers and automated aids). Pilots were also asked several questions concerning the "engine fire" event, the cues that are typically necessary to diagnose the presence of an engine fire, the cues present during the event, and any other factors affecting their handling of it. Finally, participants were debriefed on the nature of our manipulations and the purpose of the study. No "justification" interviews were conducted.

RESULTS

Omission Error Events

Descriptive analyses revealed overall omission rates for flight-related events of approximately 55%,[1] as predicted, replicating the patterns exhibited by students in the Skitka et al. (1996) analog study. The altitude load failure and the heading

---

[1] Due to several factors—including individual pilot strategies and inputs that precluded event occurrence (e.g., hand dialing the heading clearance rather than autoloading it)—only 21 pilots experienced every omission error event.

capture failure, the two events arguably most critical to aircraft operation safety, remained undetected by 44% and 48% of the participants respectively. The frequency misload was undetected by 71% of pilot participants. Only three pilots detected all three flight-related events; five pilots failed to detect any of the three flight-related events. The tracking task automation failure, which was completely irrelevant to flight functioning, was detected by all of the participants, with reaction times ranging up to 4 min.

Contrary to predictions, the number of omission errors did not vary significantly as a function of experimentally-manipulated accountability. Omission errors were correlated with total flight hours, $r(20) = .49$, $p < .05$, and with years of flight experience, $r(20) = .46, p < .05$, suggesting that increased experience decreased the likelihood of catching the automation failures. To ascertain other underlying factors discriminating participants who were more likely to verify automated tasks (and thus catch errors) from those less likely to do so, pilots were classified according to the number of omission errors they committed. Those who missed two or three of three flight-related events were categorized as "high-bias" participants ($n = 11$), and those who missed none or only one event were placed into the "low-bias" group ($n = 10$). Analyses of variance were conducted using bias group as the independent variable and responses on the debriefing questionnaire (1–7 scale) as dependent variables.

Bias groups were statistically equivalent on items such as comfort with the experiment, confidence in their strategies, and confidence in computers. Low-bias participants, however, reported more nervousness, $F(1, 19) = 7.08$, $p < .015$, a higher sense of being evaluated on their performance, $F(1, 19) = 2.21, p < .001$, a higher sense of being evaluated on strategies in use of the automation, $F(1, 19) = 9.63, p < .006$, and a stronger need to justify their interaction with the automation, $F(1, 19) = 6.24$, $p < .02$. Results indicated that those subjects who felt more accountable (whether they were in the accountable condition or not) were less likely to make omission errors than those who did not feel as accountable for their performance.

Commission Errors

All of the pilots ($N = 21$) who experienced the false engine fire message did ultimately shut down the engine. This was contrary to responses on the debriefing questionnaire indicating that an EICAS message without other cues would not be sufficient to diagnose "definitely a fire," and that it would be safer, in the presence of only an EICAS message, to retard the throttle of the indicated engine and complete the go-around procedure with the engine running rather than to shut down the suspect engine. Further analyses of the debriefing questionnaires, however, revealed that unanticipated factors may have entered into the shutdown decision. As part of the debriefing process, pilots were asked which cues were present during the go-around at SFO. All of the pilots correctly remembered that an engine fire EICAS message was present. Unexpectedly, however, 67% of them also "remembered" at least one additional cue that was not actually present, and most of these "remembered" more than one phantom indicator.[2] Additionally, the number of phantom cues that pilots recalled was related to their strategy in handling the engine shutdown. Specifically, the faster pilots initiated and completed the shutdown process, the greater the number of indicators they recalled. Number of indicators recalled accounted for approximately 20% of the variance in the speed of the shutdown process, including bringing the engine fire checklist to the view screen, hitting the engine stop switch, and pulling the engine fire handle, $r = -.45$, $p < .05$.

DISCUSSION

Results of this study suggest that automation bias is a significant factor in pilot interaction with automated aids, and that pilots are not utilizing all available information when performing tasks and making decisions in conjunction with automation. Pilots exhibited the same overall rate of automation-related errors as the student population in the Skitka et al. (1996) study, demonstrating that expertise does not insulate individuals from automation bias. In fact, experience and expertise, which might be predicted to make pilots more vigilant and less susceptible to automation bias, were related to a greater tendency to use only automated cues. One possible explanation for this may be that, because automated systems tend to be highly reliable, more experience with them reinforces the notion that other cues are merely redundant and unnecessary. Additionally, the higher-experienced pilots in this study tended to be those more senior within their airline, and were currently flying as captains. They may be used to delegating the cross-checking role to their first officers rather than doing it themselves.

Although this study does not provide evidence that externally imposed accountability affects the decision-making behavior of professional pilots, it does demonstrate that the internalization of "accountability" for performance and strategies in the use of automated systems impacts automation bias. Perceived accountability was positively correlated with increased verification of automated functioning and fewer omission errors. In other words, pilots who *reported a higher internalized sense of accountability* for their interactions with automation verified correct automation functioning more often and committed fewer errors than other pilots. These results suggest that the sense that one is accountable for one's interaction with automation encourages vigilance, proactive strategies, and the use of all information in interactions with automated systems. The fact that the perception of

---

[2] Three of the 21 participants were excluded from the memory analyses due to the occurrence of other indicators, such as a master warning for exceeding maximum recommended speed with flaps extended, during the engine fire event.

accountability was not correspondent with our external manipulation indicates the need to establish the degree to which accountability is a variable that can be significantly influenced in pilots and other professional decision makers, who are already functioning at a high level of personal responsibility for their conduct. Alternatively, the lack of experimentally determined accountability effects could be the artifact of a small sample, or the experimental manipulation could have been confounded by a sense of accountability induced by the experience of participating in a NASA study. The perception of accountability might also be part of some innate cognitive style or personality construct, a hypothesis that will be addressed in a future study.

Several aspects of the experimental tasks seemed to affect the tendency of participants to verify the functioning of the automation; these included task importance, predictability, and feedback. Descriptive data suggested that pilots were more likely to catch automation events that involved altitude and heading than one that involved a frequency error, corresponding with the typical ordering of flight priorities as (a) aviate, (b) navigate, (c) communicate. The fact that the tracking task automation failure so readily captured pilot attention can be explained in part by the fact that the task display offered trend information (e.g., pilots could see when the cursor was starting to drift out of the target area), as well as immediate salient feedback on errors. In many real-world automation-aided tasks, pilots do not receive either predictive information or immediate feedback on accuracy and errors. Incorporating these features in the design of new automated systems may be a way to facilitate vigilant use and aid in the detection of automation failures.

The finding that all pilots responded to the false engine fire event by shutting down the indicated engine was unanticipated. Even more surprising was the fact that their actions were contrary to responses on the debriefing questionnaires indicating that some *combination* of cues would be necessary to diagnose "definitely a fire," and that it would be safer, in the presence of only an EICAS message, to retard the throttle of the indicated engine rather than shutting it down. Apparently, these pilots were acting contrary to their self-described strategies. Part of the explanation for their actions may rest in the "false memory" data.

Pilots were definitely aware of the cues typically associated with an engine fire and, in this time-pressured and cognitively demanding situation, were evidently remembering a pattern of cues that was consistent with what *should have been present* during the event. Real correlations among the cues—during an engine fire, all of them should have been present—contributed to the illusion of their presence. In similar situations during their careers (i.e., engine fire situations calling for the shutdown of an engine), the pilots in this study would have been wrong if they had *not* remembered the presence of several cues.[3] The more quickly they acted, the

---

[3] Thanks to Brunswik Society member Michael Birnbaum of Cal State University, Fullerton, for this observation.

more "sure" they had to be that their actions were correct, and the more phantom cues they remembered as having prompted their actions. Their judgments and their memories were biased in the direction of confirming their expectations, and, as Hamilton (1981) aptly remarked concerning correlated cues and expectations, "[they] wouldn't have seen it if [they] hadn't believed it" (p. 137).

It is possible that this false memory phenomenon may be an additional mechanism by which automation bias is bolstered. Pilots may not ever be aware that they are using automated information as a short cut, either because the automated cues *are* in fact consistent with other information (and no error results), or because errors are not traced back to a failure to cross-check automated cues. In the false engine fire event, for example, most pilots believed they had acted based on several cues, and thus would not be prompted to change their strategies in the use of automated information. Consistent with this analysis, pilots tended to falsely remember cues that fit into their expectations but were not physically handled during the engine shutdown. For example, they were less likely to remember the presence of an illuminated fire handle ($n = 2$), which they had to manually press during engine shutdown (and would have been forced to observe directly) than a master warning light ($n = 8$), which was not directly manipulated as part of the shutdown procedure. Explicit manipulation of a control made it more salient, and less likely to be incorrectly remembered.

This study represents an initial step in discovering patterns of automation bias and strategies to ameliorate it. Follow-up studies will investigate cognitive and personality constructs that may be associated with a tendency toward automation bias, examine whether the relationship between accountability and automation bias remains constant in a crew (i.e., two-pilot) situation, and explore the extent to which the false memory phenomenon persists in a crew context. These studies will also determine whether other strategies, such as training pilots to verify automated functioning and information specifically as a defense against automation bias, or providing display prompts that encourage verification, can reduce errors associated with automation bias.

ACKNOWLEDGMENTS

This work was supported by NASA Ames Research Center, Moffett Field, CA, NCC2–798, NCC2–837, and NAS2–832. Earlier versions of this manuscript were presented at the Second Automation Technology and Human Performance Conference in Cocoa Beach, FL, March, 1996, and at the Annual Conference of the Society for Judgment and Decision Making, Chicago, IL, November, 1996. Kathleen Mosier joined the faculty of San Francisco State University in Fall 1997. Susan Heers is currently with Monterey Technologies, Inc.

REFERENCES

Analysis of flight data and cabin voice tapes for Korean Airlines Flight KE–007. (1993). *Aviation Week and Space Technology, 138*(25), 17.

Billings, C. E. (1996). *Human-centered aviation automation: Principles and guidelines* (Tech. Mem. No. 110381). Moffett Field, CA: NASA Ames Research Center.

Darley, J. M. & Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology, 44,* 20–33.

Ebbeson, E. D., & Konecni, V. (1980). On external validity in decision-making research. In T. Wallsten (Ed.), *Cognitive processes in choice and decision making.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Fiske, S. T., & Taylor, S. E. (1994). *Social cognition* (2nd Ed.). New York: McGraw-Hill.

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98,* 506–528.

Glick, P., Zion, C., & Nelson, C. (1988). What mediates sex discrimination in hiring decisions? *Journal of Personality and Social Psychology, 55,* 178–186.

Hagafors, R. & Brehmer, B. (1983). Does having to justify one's decisions change the nature of the decision process? *Organizational Behavior and Human Performance, 31,* 223–232.

Hamilton, D. L. (1981). *Cognitive processes in stereotyping and intergroup behavior.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kaempf, G. L. & Klein, G. (1994). Aeronautical decision making. In N. Johnston, N. McDonald, & R. Fuller (Eds.), *Aviation psychology in practice* (pp. 223–254). Aldershot, UK: Avebury Technical.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases.* New York: Cambridge University Press.

Kruglanski, A. W., & Freund, T. (1983). The freezing and unfreezing of lay inferences: Effects on impressional primacy, ethnic stereotyping, and numerical anchoring. *Journal of Experimental Social Psychology, 14,* 448–468.

Layton, C., Smith, P. J., & McCoy, C. E. (1994). Design of a cooperative problem-solving system for en-route flight planning: An empirical evaluation. *Human Factors, 36*(1), 94–119.

Mosier, K. L., Palmer, E. A., & Degani, A. (1992). Electronic checklists: Implications for decision making. In *Proceedings of the Human Factors Society 36th Annual Meeting* (pp. 7–11). Santa Monica, CA: Human Factors Society.

Mosier, K. L., & Skitka, L. J. (1996). Human decision makers and automated decision aids: Made for each other? In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and Applications* (pp. 201–220). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Mosier, K. L., Skitka, L. J., & Korte, K. J. (1994). Cognitive and social psychological issues in flight crew/automation interaction. In M. Mouloua & R. Parasuraman (Eds.), *Human performance in automated systems: Current research and trends* (pp. 191–197). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors, 39*(2), 230–253.

Sarter, N. R., & Woods, D. D. (1993). *Cognitive engineering in aerospace application: Pilot interaction with cockpit automation* (NASA Contractor Report 177617). Moffett Field, CA: NASA Ames Research Center.

Simonson, I. & Nye, P. 1992). The effect of accountability on susceptibility to decision errors. *Organizational Behavior and Human Decision Processes, 51,* 416–446.

Skitka, L. J., Mosier, K. L., & Burdick, M. D. (1996). *Accountability and automation bias.* Manuscript submitted for publication.

Tetlock, P. E. (1983). Accountability and the perseverance of first impressions. *Social Psychology Quarterly, 48,* 227–236.

Tetlock, P. E. (1985). Accountability: A social check on the fundamental attribution error. *Social Psychology Quarterly, 48,* 74–83.

Tetlock, P. E., & Kim, J. I. (1987). Accountability and judgment processes in a personality prediction task. *Journal of Personality and Social Psychology, 52,* 700–709.

Wickens, C. D., & Flach, J. M. (1988). Information processing. In E. L. Wiener & D. C. Nagel (Eds.), *Human factors in aviation* (pp. 111–156). San Diego, CA: Academic.

Manuscript first received February 1997