INDIVIDUAL DIFFERENCES IN THE CALIBRATION OF TRUST IN AUTOMATION

A Thesis
Presented to
The Academic Faculty

by

Vlad Liviu Pop

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in Psychology

Georgia Institute of Technology
May 2013

INDIVIDUAL DIFFERENCES IN THE CALIBRATION OF TRUST IN AUTOMATION

Approved by:

Dr. Francis T. Durso, Advisor
School of Psychology
Georgia Institute of Technology

Dr. Wendy A. Rogers
School of Psychology
Georgia Institute of Technology

Dr. Rustin Meyer
School of Psychology
Georgia Institute of Technology

Date Approved: April 25, 2012

ACKNOWLEDGEMENTS

I would like to thank my advisor, Frank Durso, for all the guidance and help he provided throughout all of the steps in this process. I would like to thank my committee members, Wendy Rogers and Rustin Meyer, for their guidance and invaluable feedback. I would like to thank Stephanie Merritt for sharing the x-ray image set used in the Merritt and Ilgen (2008) study. Finally, I would like to especially thank Alex Shrewsbury for all his help with compiling the image set used as stimuli, programming the experimental task, and extracting, coding, and organizing data.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ABBREVIATIONS

df	Degrees of Freedom
F	F-Statistic
M	Mean
Mdn	Median
n	Sample Size
p	Pearson
r	Pearson Product-Moment Correlation Coefficient
R^2	Coefficient of Determination
SD	Standard Deviation
W	Shapiro-Wilk Statistic
ηp^2	Partial Eta-Squared

AICP	Automation Induced Complacency Potential
AWD	Automatic Weapons Detector

SUMMARY

A large body of research has identified that one of the major factors influencing decisions about automation use is operator's trust. Studies have shown that operator trust can be affected by individual differences in expectancy. The purpose of this study was to investigate whether operators with an expectancy that automation is trustworthy are better at calibrating their trust to changes in the capabilities of the automation. We also investigated why this individual difference affects calibration. In a baggage screening task 179 participants searched for weapons in 200 x-ray images of luggage. Participants were assisted by an automated decision aid exhibiting different levels of reliability. Measures of expectancy that automation is trustworthy were used in conjunction with subjective measures of trust and of perceived reliability to identify individual differences in trust calibration. We found that operators with high expectancy that automation is trustworthy were more sensitive to changes (both increases and decreases) in automation reliability, and that this difference was caused by attributing errors to the situation rather than the automation.

CHAPTER 1

INTRODUCTION

In an effort to reduce human error in complex systems, a variety of automated aids aimed at improving diagnostic and decision making processes are being implemented (Wiegmann, Rich, & Zhang, 2001). If these automated aids are not perfectly reliable, the decision of when to rely on automation and when to perform a task manually can greatly influence performance, safety, and profitability. For example, Wickens and Dixon (2007) found that in some cases operators relied on automated aids even when it resulted in worse diagnostic monitoring than when the operator had not used automation at all. When operators rely on automation without recognizing its limitations, misuse occurs (Parasuraman & Riley, 1997). On the other hand, if operators underutilize automation, even though using it would improve their performance, automation disuse occurs (Parasuraman & Riley, 1997). Both automation misuse and disuse are indicators of inappropriate reliance, and their consequences can be costly.

A large body of research has identified that one of the critical variables that may influence decisions about automation use is an operator's trust in the automation (see Lee & See, 2004 for a review). Operators tend to use automation that they trust while rejecting automation that they do not. For appropriate use to occur, operator trust must match the true capabilities of the automation (Lee & See, 2004). The appropriateness of trust can be evaluated in terms of its calibration. *Calibration* is "the correspondence between a person's trust in the automation and the automation's capabilities" (Lee & See, 2004; p. 55). If trust is higher than the system's capabilities, over-trust occurs, whereas if trust is lower than the system capabilities, distrust (under-trust) occurs.

The calibration of trust can be affected by characteristics of the automation, the operator, and the context (Merritt & Ilgen, 2008). Studies examining how these factors individually influence trust include how trust is affected by: automation reliability (Dzindolet, Peterson, Pomranky, Pierce, & Beck, 2003), automation display content and format (Lee & See, 2004), automation errors (Dzindolet et al., 2003; Lee & Moray 1992; Wiegmann et al., 2001), task difficulty (Madhavan, Wiegmann, & Lacson, 2006), and the consequences of misses and false alarms (Keller & Rice, 2010; Parasuraman & Riley, 1997). These studies have identified how individual factors influence trust, however we do not fully understand how the dynamic interaction of these factors influences trust. For example, Lee and Moray (1994) found that for some people changes in the capabilities of the automation led to a substantial change in trust, yet for others it led to a small change. This finding suggests that the relationship between automation characteristics and trust may be moderated by individual differences.

An individual difference that can influence initial trust level, as well as how new information and feedback are interpreted is one's expectancy that others can be trusted. In 1967, Rotter developed a measure assessing trust in parents, teachers, friends, classmates, and so on, in order to determine a person's expectancy that others can be trusted. The scale was additive such that the greater variety of people that a person trusted, the higher that person's general expectancy that people can be trusted. Interestingly, just because an individual had a high expectancy that others could be trusted did not mean that the person was more gullible (Rotter, 1967; 1971; 1980). In fact, the opposite has been shown to be true: High expectancy individuals trusted others more appropriately (Lee & See, 2004). Yamagishi, Kikuchi, and Kosugi (1999) found that when interacting with untrustworthy people, high expectancy individuals had lower estimates of trust than low expectancy individuals. They also found that high expectancy individuals predicted other's trustworthiness better than low expectancy individuals, and that high expectancy individuals were more sensitive to information revealing a lack of trustworthiness in others. The same was true in a study summarized by Rotter (1971), in which after a confederate told participants that they were being deceived, participants with a high expectancy that others could be trusted reported being suspicious whereas participants with a low expectancy did not.

The individual difference of expectancy that others can be trusted originated in the interpersonal literature, however studies showing that humans respond socially to technology have led to its application in the automation literature (Lee & See, 2004). Singh, Molloy, and Parasuraman (1993a) developed an Automation-Induced Complacency Potential (AICP) Rating Scale to assess attitudes towards everyday automation such as automated teller machines and automotive cruise control. Like Rotter's (1967) scale, this scale was also additive, such that a higher score indicated a greater variety of automation that a person found trustworthy, reliable, dependable, and safe. The scale represents the general tendency or propensity to trust automation and has been used as a measure of people's expectancy that automation can be trusted (Lee & See, 2004; Merritt & Ilgen, 2008; Singh et al., 1993a). Some studies have referred to this construct as propensity to trust, however that term is somewhat misleading as it implies that people with a high propensity to trust always have higher trust in automation, which is not the case. For example, Merritt and Ilgen (2008) found that when using automation that was highly reliable (85%), operators with a high AICP score had higher trust estimates than operators with a low AICP score, however when using automation that was not very reliable (65%) the opposite was true; operators with a high AICP score had lower trust estimates than operators with a low AICP score. This finding parallels the studies in the interpersonal literature that found when interacting with untrustworthy people, high expectancy individuals had lower estimates of trust than low expectancy individuals (Yamagishi et al., 1999). Thus we will follow the convention in the interpersonal trust literature and refer to the AICP measure as *automation expectancy* rather than propensity to trust (although it represents the same construct).

The parallel findings of Merritt and Ilgen (2008) and Yamagishi et al. (1999) suggest that individual differences in expectancy may be able to account for why changes in the capabilities of the automation in previous studies led to a substantial change in trust for some, yet for others it only led to a small change. Lee and See (2004) proposed that "one possible explanation that merits further investigation is that high [propensity to] trust individuals may be better able to adjust their trust to situations in which the automation is highly capable as well as to situations in which it is not" (p. 53). This led to the first research question explored here: Are operators with high automation expectancy better at calibrating their trust to changes in the capabilities of the automation?

The proposed question has yet to be tested empirically with automation; however, the findings of other studies investigating automation expectancy suggest that people with high automation expectancy may be better at adjusting their trust to changes in the capabilities of the automation. Singh, Molloy, and Parasuraman (1993b) found that people with high automation expectancy had higher detection rates of automation failures than people with low automation expectancy. This finding may indicate that people with high automation expectancy are more sensitive to information revealing lack of trustworthiness, however this cannot be confirmed because trust was not measured in that study. Merritt and Ilgen (2008) found that trust estimates of operators with high automation expectancy were more sensitive than operators with low expectancy, however operators were either paired with a low or high functioning automated aid whose capabilities remained constant, so we do not know if this higher sensitivity results in a better calibration of trust when the reliability of the automated aid changes.

Several models of trust development can inform questions of how operators calibrate their trust. An integrative model of trust by Mayer, Davis, and Schoorman (1995) proposed that trust develops through a feedback loop, such that after each time a trustor takes a risk with a trustee, the trustor evaluates the outcome. Mayer et al. (1995) further proposed that the outcome then indirectly influences trust through the trustee's perceptions of several trustor characteristics, including ability. In other words if the outcome is positive, prior beliefs about the trustee's abilities should be reinforced and trust should increase, whereas if the outcome is negative, prior beliefs about the trustee's abilities may be reevaluated and trust may decrease . See Figure 1 for an illustration of the feedback loop in Mayer et al.'s (1995) model.

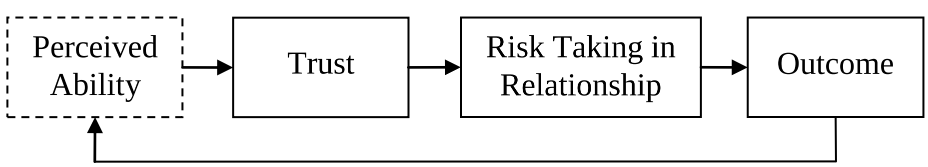

Figure 1. An example of how trust can be affected by the feedback loop proposed in Mayer et al.'s (1995) integrative trust model.

Applying Mayer et al.'s (1995) model to the current context would suggest that after each time an operator utilizes an automated aid, the operator assesses whether trusting the automation led to a correct detection decision. If trusting the automation led to a correct decision, operator trust should increase, whereas if it led to an incorrect decision, the operator may reevaluate the automation's abilities and subsequently reduce trust. Mayer et al.'s (1995) feedback loop accounts for the evolution of trust over time, however it cannot be used to explain the individual differences in trust evolution found in previous studies (e.g., Lee & Moray, 1994; Merritt & Ilgen, 2008). The explanation for individual differences in trust evolution may be related to differences in operator perceptions of automation errors.

A model proposed by Tomlinson and Mayer (2009) combines Mayer et al.'s (1995) feedback loop with Weiner's (1986) causal attribution theory to explain why negative outcomes sometimes lead to a decrease in trust and other times do not. According to Weiner's (1986) attribution theory, negative outcomes (such as errors) cause individuals to identify the outcome's cause (e.g., automation ability, luck, task difficulty). The locus of causality of this attribution can vary. An individual can attribute the cause of an error as either being internal to the automation (i.e., automation ability) or external to the automation (i.e., situation). If operators attribute the cause of errors to automation ability, Tomlinson and Mayer's (2009) model would suggest that these operators should withdraw trust, however if operators attribute the cause of errors to external factors, the same model would suggest that operator trust may not necessarily decline. See Figure 2 for an illustration of how Tomlinson and Mayer's (2009) model could be applied to explain individual differences in trust in automation.

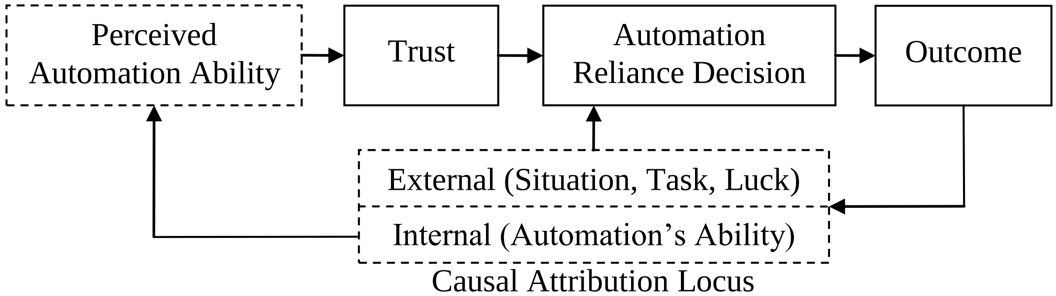

Figure 2. Individual differences in trust calibration as explained by a feedback loop varying based on the locus of control of the causal attribution of errors.

Tomlinson and Mayer (2009) developed their model to explain when trustworthiness in interpersonal relationships is in need of repair, however their model could also be applied to explain why operators with high automation expectancy may be better at calibrating their trust. As illustrated in Figure 2, trust is influenced by operator perceptions of automation abilities. This is consistent with research on trust in automation suggesting that trust calibration is based on the operator's subjective perceptions of automation capabilities and not on the actual capabilities of the automation (Lee & See, 2004; Merritt & Ilgen, 2008; Wiegmann et al., 2001). Merritt and Ilgen (2008) found that operator perceptions were affected by individual differences in automation expectancy. Specifically, they found that operators with high automation expectancy who were paired with lower functioning automation suffered the largest negative effects on perceptions of machine characteristics and on trust. The parallel of these findings and Tomlinson and Mayer's (2009) model certainly suggests that operators with high and low automation expectancy may differ in their ability to calibrate trust because of differences in the causal attribution of automation errors.

The purpose of the current study was to investigate whether operators with high automation expectancy better calibrate their trust to changes in the capabilities of the automation, and if so, why? These questions were tested using a baggage screening task in which participants searched for weapons in x-ray images of luggage with assistance from an automated decision aid. Both individual differences in automation expectancy and subjective measures (trust and perceived reliability) were examined. The first question was addressed by investigating individual differences in operators' level of trust in an automated decision aid during trials exhibiting different levels of automation reliability. Initially participants screened luggage while using an automated aid that was either 60%, 80%, or 100% reliable. The automated aid with an initial reliability of 60% increased to 80% reliability whereas the aid with an initial reliability of 100% decreased to 80% reliability. The automated aid with an initial reliability of 80% acted as a control, and its reliability did not change. Thus, in all conditions automation reliability was 80% in the second half of the experiment. The three reliability conditions are identical to the ones used by Wiegmann et al. (2001) to investigate the effects of aid reliability on user's trust. Wiegmann et al. found that operators were sensitive to automation reliability levels differing by a magnitude of 20%, thus the same conditions were used to investigate individual differences.

As previously stated, the findings of Merritt and Ilgen (2008) combined with Tomilson and Mayer's (2009) model suggest that operators with high and low automation expectancy may differ in their ability to calibrate trust because of differences in the causal attribution of automation errors. In order to test why operators with high automation expectancy may be better at calibrating their trust, the causal attribution of automation errors was manipulated for half of the participants. Previous research indicates that an individual's causal attributions depend upon the nature of available information (Kassin, 1979; Kelley, 1967; Teas & McElroy, 1986). Specifically, three types of information are used in attributing causal explanations – *consistency* (does the individual behave similarly across time and situations), *consensus* (do other individuals in the same situation behave similarly), and *distinctiveness* (does the individual behave this way only in certain situations; Kelley, 1967).

According to Kelley's (1967) theory about the causal attribution process, causal attributions are based upon particular combinations of the three aforementioned factors (Kassin, 1979; Teas & McElroy, 1986). When people are presented with the combination of information indicating that a behavior is high in consistency, low in consensus, and low in distinctiveness, they will likely attribute the cause to the individual (Kassin, 1979; Kelley, 1967; Teas & McElroy, 1986). For example if an educator is trying to determine why one of their students failed a test in their class and reviewing the student's other grades in their class indicates that the student has regularly scored poorly (high consistency), reviewing the transcript indicates that the student has failed other classes (low distinctiveness), and other students in the class performed well on the exam in question (low consensus), the educator is likely to blame the student (e.g., poor study habits, lack of attention paid in classes, etc). On the other hand, given a combination of information indicating that a behavior is high in consistency, high in consensus, and high in distinctiveness, people are likely to attribute the cause to the situation (Kassin, 1979; Kelley, 1967; Teas & McElroy, 1986). Using the same scenario as before, if reviewing the student's other grades in the class indicates that the student regularly scored well (high consistency), reviewing the transcript indicates that the student has done well in other classes (high distinctiveness), and other students in the class also performed poorly on exam in question (high consensus), the educator is likely to blame the context (i.e. difficult exam).

The current study used two information conditions to test the causal attribution hypothesis. One condition was a control in which participants did not receive any information regarding the consistency, distinctiveness, or consensus of automation errors. Providing no such information is fairly typical in studies investigating trust in automated decision aids (e.g., Lee & Moray, 1992; 1994; Madhaven et al., 2006; Wiegmann et al., 2001). The other condition was a manipulation in which participants received a combination of information that should result in automation errors being causally attributed to the automation's ability (i.e. high consistency, low distinctiveness, low consensus), rather than, for example, the situation or task.

This manipulation had not been previously used to investigate trust in automation, however, the trust literature had certainly identified the contribution of information in the development and calibration of trust. As Lee and See (2004) stated: "the availability of information at … different levels of detail and attributional abstraction may lead to high calibration" (p. 61).

CHAPTER 2

METHOD

Participants

Participants were 179 undergraduate students from the Georgia Institute of Technology, 89 of which were female and 90 male. The sample was 68.7% White, 3.9% Black, 20.9% Asian, and 6.5% other races. The mean age was 19.9 with a standard deviation of 1.86. All participants were over the age of 18, had normal or corrected to normal vision, and reported being knowledgeable about computers and comfortable using them. Participants were compensated for their participation by receiving 1.0 course credit through Experimetrix.

Apparatus

The study took place in the Cognitive Ergonomics Lab at the Georgia Institute of Technology. The task was a medium fidelity X-Ray Screening Task similar to the task used in Merritt and Ilgen's (2008) study[1]. For this task, x-ray images of luggage were presented on a monitor. Some of the bags contained weapons, and some did not. Participants had to inspect the x-rays and determine if they believed a weapon was present and that they would search the bag, or if they believed that the bag did not contain a weapon and that they would clear the bag. The slides depicted a weapon 30% of the time. This was the same rate used in Merritt and Ilgen (2008) and similar to the rate used in other studies investigating trust in automated decision aids (e.g., Dzindolet, Pierce, Beck, Dawe, & Anderson, 2001). The difficulty of the slides varied and was largely based on the number of items present in each bag. Looking ahead, accuracy on individual items ranged from 16% to 100%. Overall accuracy performance ranged from 69.5% to 97.5%, with a mean of 86.2% and a standard deviation of 5.7%.

Stimuli

Image Set

The x-ray images of luggage used as stimuli were compiled with items from the image set used by Merritt and Ilgen (2008). From this image set we chose 50 types of large bags. We did not choose small bags, such as purses and backpacks, because they were too small to hold many of the items. Items in the bag were chosen from 10 types of clothing, 18 types of knives, 18 types of guns, and 242 various types of other items, including electronics (e.g., cell phones, laptops, MP3 players, radios, PDAs, pagers, and tape players), organic materials (e.g., papers, books, shoes, wallets, and diapers), and metal objects (e.g., tools, toys, keys, and coins). Each bag was standardized at a resolution of 768 x 576, displayed on a 30" monitor with a 1024 x 768 resolution.

Bag Filling Procedure

To create each stimulus, a bag number was randomly chosen representing one of the 50 large bag types in the image set. The bag was then randomly filled with 2-5 randomly chosen items of clothing and another 8-20 randomly chosen items. For each item, one of the five views was randomly chosen. The orientation of each item was also randomly chosen to be either 0, 45°, 90°, 135°, 180°, 225°, 270°, or 315°. Finally, each bag was split up into a 4x4 grid, and each item was randomly placed in one of the 16 resulting quadrants.

The bag filling procedure was quite time consuming as it took 20-30 minutes to compile each bag depending on the number of items in it. In the interest of time, only 50 bags were compiled. These 50 bags were repeated four times for the 200 trials in this experiment. In order to minimize the chances of participants noticing that the same 50 bags were repeated four times, each bag was rotated 180°, flipped horizontally, and flipped vertically. Additionally, one of the four views of each bag (original, rotated 180°, flipped horizontally, flipped vertically) was randomly chosen to be rotated 90° clockwise. See Figure 3 for an example of the four views for each bag. Each bag view was randomly assigned to one of four trial blocks. The order of bags within each trial block was random, which left only a 0.04% chance that a participant saw the same bag twice in a row.

In each 100 trial block, 30 bags were randomly chosen for weapon placement. This was done so that even if a participant noticed that bags repeated, just because a bag did not contain a weapon before did not mean it could not contain one now, and vice versa. For example, in Figure 3, one of the four bag variants contains a weapon while three do not. For the actual placement of weapons, a randomized block design was used to place either a gun or knife in 60 of the 200 images. For both guns and knives, one out of the 18 weapon types was randomly chosen. Similarly to the placement of items, one of the five views, one of the eight orientations, and one of the 16 placement locations was randomly chosen for each weapon.

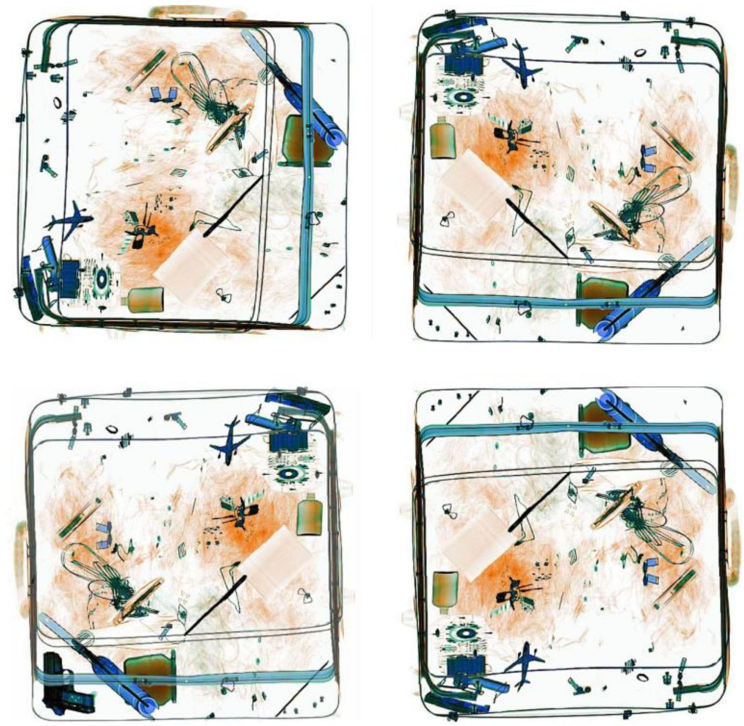

Figure 3. An example of the four different views of one bag.

Participants were told that an automated aid, the Automatic Weapons Detector (AWD), was available to assist them, however they were not told how reliable it was. The reliability of the automated aid was either 60%, 80%, or 100% depending on reliability condition. For conditions in which the automated aid was not perfectly reliable, automation errors were equally divided between misses and false alarms, and randomly assigned to bags.

Measures

Automation Expectancy

Automation Expectancy was assessed using Singh et al.'s (1993a) Automation-Induced Complacency Potential Rating Scale. This scale represents the individual difference of expectancy that automation can be trusted (Lee & See, 2004; Merritt &

Ilgen, 2008; Singh et al., 1993a) and is the same measure used by Merritt and Ilgen (2008; although they referred to it as propensity to trust). In previous research this measure has had both high internal consistency ($r > .98$) and test-retest reliability ($r = .90$; Singh et al., 1993a).

Items in this scale presented the participants with a statement, for example "I would rather purchase an item using a computer than have to deal with a sales representative on the phone because my order is more likely to be correct using the computer." Participants then rated their response on a five point Likert scale ranging from strongly disagree to strongly agree. Scores on this measure can range from 16 to 80.

Trust and Perceived Reliability

Trust was assessed using a 12 item questionnaire developed by Jian, Bisantz, and Drury (2000) to assess trust between people and automation. Example items included "I can trust the system" and "The system is dependable." Participants responded to each item using a seven point Likert scale ranging from not at all (1) to extremely (7). Additionally, perceived reliability was assessed by having each participant rate the reliability of the automated aid on a 100 point scale ranging from 0% - completely unreliable, to 100 % - completely reliable (adopted from Wiegmann et al., 2001).

Causal Attribution

Causal attribution of automation errors was assessed using an adaptation of Russel's (1982) Causal Dimension Scale. The Causal Dimension scale assess the locus of causality of attribution identified by Weiner (1979) and has high internal reliability ($\alpha = .87$; Russel, 1982). The format of the scale has been modified to reflect the investigation of the operator's causal explanation of automation errors.

The scale contained nine items. Participants indicated their opinion on each item by circling a number between 1 and 9. For example, "Is the cause something that: Reflects an aspect of the automation (9) or Reflects an aspect of the situation (1)."

Design

The experiment was a 2 (information: control, internal attribution) x 3 (reliability-group: 100%-80%, 60%-80%, 80%-80%) factorial design. Information was a between subjects manipulation of the type of information the operators received. Some operators received information indicating the consistency, consensus, and distinctiveness of the automated aid's performance while other operators received no such information. Participants were randomly assigned to an information condition.

Reliability-group was a factor of the three experimental conditions used in this study. For the first group, the reliability of the automated aid was 100% for the first half of the experiment, and unbeknownst to the operator, was reduced to 80% for the second half of the experiment. In the second group, the reliability of the automated aid was 60% for the first half of the experiment, and unbeknownst to the operator was increased to 80% for the second half. The third group was a control in which the automated aid remained at 80% reliability for the duration of the experiment. No change occurred in the presentation that could signal the change in reliability with the exception of the passage of time. Reliability-group was a between subjects manipulation and participants were randomly assigned to a reliability condition.

Procedure

Each experimental session lasted less than one hour. Participants gave informed consent, then completed questionnaires to determine automation expectancy and demographic information (e.g., age, sex, and ethnicity). Participants were then given instructions that described the task, the weapons, and the procedure to clear or search a bag. Participants were instructed to screen the x-ray images of carry-on luggage for weapons as fast as they can while still adequately inspecting each bag. After reviewing the instructions, participants completed three practice trials during which they screened bags without the assistance of the AWD. Participants then received instructions describing the AWD. The instructions for the internal attribution information group included the following information about the consistency, consensus, and distinctiveness of the automated aid's performance: "The Automated Weapons Detector does not perform the same as other automated detectors" (low consensus), " however its detection performance is the same regardless of weapon type" (low distinctiveness), "and it consistently provides the same recommendation when scanning the same bag multiple times" (high consistency).

After participants reviewed the instructions and description, they watched three trials demonstrating the AWD's capabilities to allow them to form a first impression of its characteristics. They then rated their initial trust levels in the automation and their perceived reliability. Participants screened 200 bags. Trust and perceived reliability were assessed twice, after every 100 trials. Causal attribution was assessed after participants screened all 200 bags.

The total number of trials in this experiment (200) was the same as other studies that had investigated trust in automated decision aids (e.g., Dzindolet et al., 2003; Madhaven et al., 2006; Wiegmann et al., 2001). The decision to assess trust and perceived reliability every 100 bags was based on Wiegmann et al. (2001), who used the same reliability conditions as the current study and also assessed perceived reliability every 100 trials. Wiegmann et al. (2001) found that 100 trials per reliability condition resulted in a significant difference in participants' perceived reliability when the two reliability conditions differed by a magnitude of 20%.

CHAPTER 3

RESULTS AND DISCUSSION

Automation Expectancy

Automation expectancy scores were computed by reversing the scores for negatively phrased items and then adding the scores (1-5) for the 16 non-filler statements. Scores could range from 16 to 80, with higher scores indicating higher expectancy of automation to be trustworthy, reliable, dependable, and safe. In our sample, automation expectancy scores ranged from 44 to 80 and were normally distributed, as confirmed by a Shapiro-Wilks test, $W(176) = .994$, $p = .663$. See Figure 4 for a histogram of the frequency distribution. The mean automation expectancy was 61.40 ($SD = 6.20$), and the median was 61. The average of our sample was slightly higher than those reported in previous studies ($M = 57.69$, Singh et al., 1993a; $Mdn = 56$, Singh et al., 1993b; $Mdn = 58$, Prinzel, Freeman, & Prinzel, 2001; $M = 56$, Merritt & Ilgen, 2008), which may be related to the technologically savvy Georgia Institute of Technology undergraduate student population.

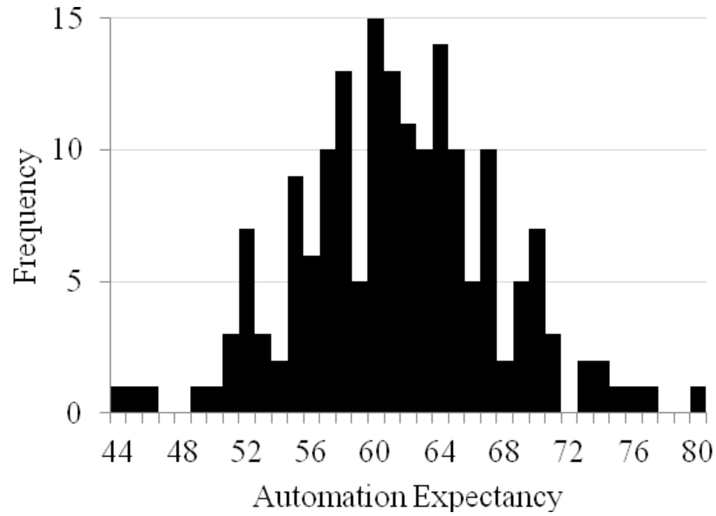

Figure 4. Frequency distribution of automation expectancy scores.

We found a significant positive correlation between automation expectancy and initial perceived reliability, r (174) = .26, p = .001, and no significant correlations between automation expectancy and perceived reliability reported after the first 100 trials, r (174) = .13, p = .075, or second 100 trials, r (174) = -.03, p = .705. The same was true for trust, with a significant positive correlation between automation expectancy and initial trust, r (174) = .29, p < .001, but no correlations between automation expectancy and trust reported after the first 100 trials, r (174) = .09, p = .264 or the second 100 trials, r (174) = -.02, p = .819. These results were consistent with Merritt and Ilgen's (2008) findings of a significant correlation between automation expectancy and initial trust but not post-task trust. Taken together, these findings indicated that the automation expectancy scale was a reliable predictor of initial trust in automation. More so, the lack of correlations between automation expectancy and post task trust indicated that having high expectancy for automation to be trustworthy did not necessarily result in the operator being over trusting or gullible.

Perceived Reliability

Data for three participants who were more than three standard deviations from the mean were excluded. The perceived reliability data for the remaining 176 participants were analyzed using linear regression with repeated measures. Reliability condition (100-80, 80-80, 60-80), information condition (Control, Causal Attribution), and automation expectancy were entered as between subject factors. Trial block was entered as a within subjects factor , with 1 representing the perceived reliability estimates reported after completing the first 100 trials, and 2 representing the perceived reliability estimates reported after completing the second 100 trials. We used partial eta-squared (ηp^2) as a measure of effect size, with .010 indicating a small effect, .059 a medium effect, and .138 a large effect (Cohen, 1988).

Overall, the pattern of data matched Wiegmann et al.'s (2001). As illustrated in Figure 5, participants were generally sensitive to different aid reliability levels as well as changes in reliability. Our findings also support Wiegmann et al.'s (2001) conclusion that trust is lost faster than it is regained, evidenced by the relatively large drop in perceived reliability when the aid went from 100% reliable to 80% compared to a much smaller increase in perceived reliability when the aid went from 60% to 80% reliable.

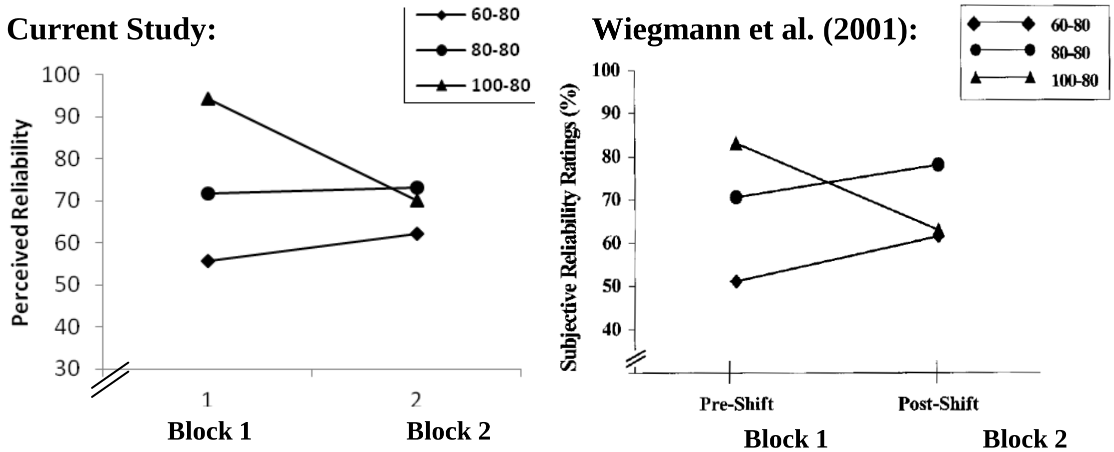

Figure 5. Perceived reliability in each trial block for the three reliability conditions in the current experiment (left) and Wiegmann et al.'s 2001 experiment (right).

There was a significant three-way interaction between block, reliability condition, and information condition, $R^2 = .04$, $F(2,164) = 3.903$, $p = .022$, $\eta p^2 = .045$, as well as a significant four-way interaction between automation expectancy, block, reliability condition, and information condition, $R^2 = .08$, $F(2,164) = 4.149$, $p = .017$, $\eta p^2 = .048$. No other effects were significant (see Table 3 in Appendix A for full results). The significant four-way interaction included the significant three-way interaction, thus only the four-way is discussed. In order to better understand the significant four-way interaction, separate three-way linear regression with repeated measures analyses were performed for each information condition. We chose to split the data by information condition so that we could first determine the relationship between the variables in the control condition and then see how the relationship was affected by the causal attribution manipulation.

In the control information condition there was a significant three-way interaction among automation expectancy, reliability condition, and block, $R^2 = .02$, $F(2,84) = 4.223$, $p = .018$, $\eta p^2 = .091$, indicating that participants were generally sensitive to different aid reliability levels as well as changes in reliability. No other effects were significant (see Table 1). The pattern of this interaction is discussed by reliability condition.

Table 1
Results of the linear regression with repeated measures used to analyze the perceived reliability data for the control information condition.

Effect	Df	F	p	ηp^2
Automation Expectancy	1,84	.398	.530	.005
Block	1,84	.020	.888	.000
Reliability Condition	2,84	1.663	.196	.038
Exp. * Block	1,84	.189	.665	.002
Exp. * Rel. Cond.	2,84	.835	.437	.020
Block * Rel. Cond.	2,84	2.314	.102	.053
Exp. * Block * Rel. Cond.	2,84	4.223	.018*	.091

* indicates that effect was significant at α =.05.

In the 100-80 reliability condition, there did not appear to be a large effect of automation expectancy in the first block. As illustrated in Figure 6, when the automation was perfectly reliable, participants had generally well calibrated yet slightly lower estimates of reliability, regardless of automation expectancy. When the automation dropped to 80% reliable in the second block, participants with a high automation expectancy had much lower estimates of reliability than participants with low automation expectancy. An additional linear regression with repeated measures analysis revealed that the block by automation expectancy interaction was not quite significant, $R^2 = .06$, $F(1,27) = 3.904$, $p = .058$, $\eta p^2 = .126$, however given the small sample size (*n=29*) and relatively large effect size we believed that the effect was truly present.

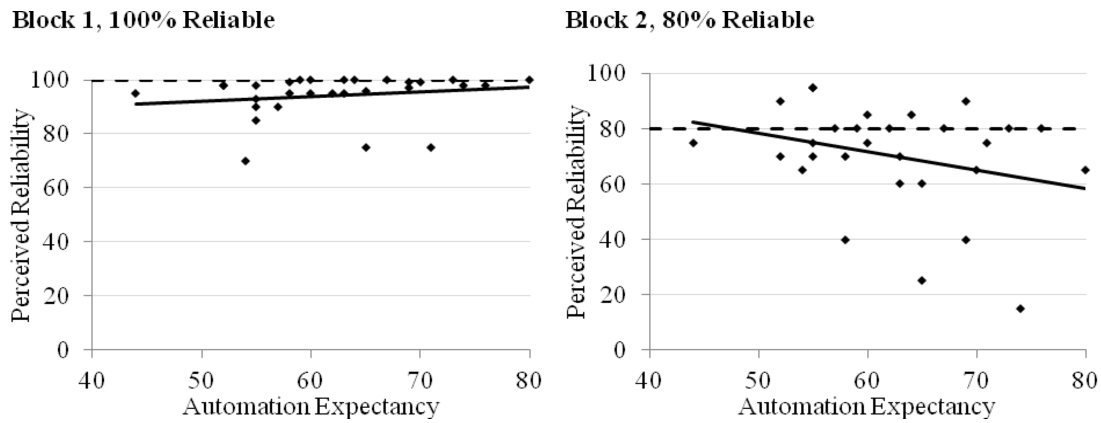

Figure 6. Actual (dotted) and perceived (solid) reliability for each trial block in the 100-80 reliability condition.

In the 80-80 reliability condition, automation expectancy did not affect perceived reliability. A separate two-way linear regression with repeated measures analysis revealed no significant results, indicating that perceived reliability did not vary as a function of automation expectancy (See Figure 7).

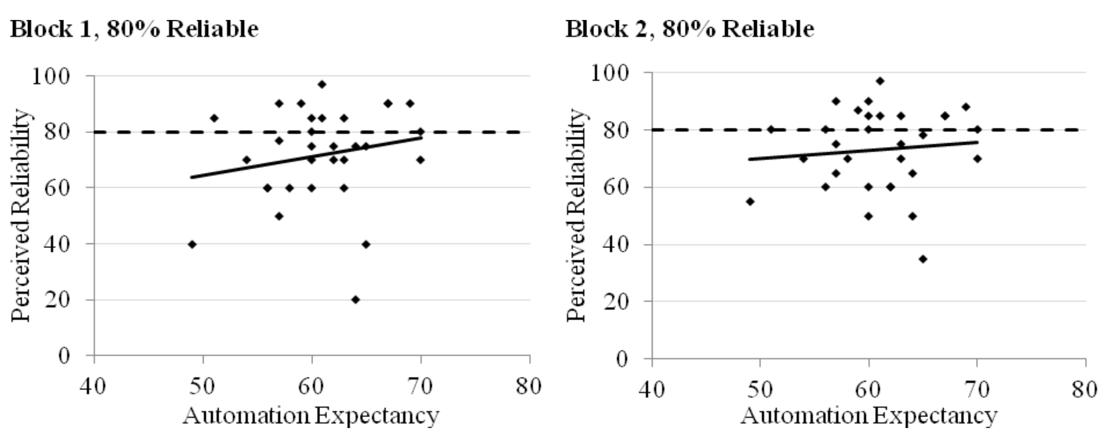

Figure 7. Actual (dotted) and perceived (solid) reliability for each trial block in the 80-80 reliability condition.

In the 60-80 reliability condition, automation expectancy did not generally affect perceived reliability during the first block. As illustrated in Figure 8, when the automation was 60% reliable, participants had generally well calibrated estimates of perceived reliability. When the automation increased to 80% reliability in the second block, participants with higher automation expectancy had higher and better calibrated estimates of reliability than participants with lower automation expectancy. An additional two-way linear regression with repeated measures analysis revealed that the block by automation expectancy interaction was not quite significant, $R^2 = .02$, $F(1,28) = 3.133$, $p = .088$, $\eta p^2 = .101$, however once again, given the small sample size (n=30) and relatively large effect size, we believed that the effect was truly present.

Taken together, the pattern of data in the control information condition indicated that higher automation expectancy resulted in better calibrated adjustment when the capability of the automation increased, but poorer calibrated adjustment when the capability decreased. As illustrated in Figure 9, in the second block, when the actual reliability of the automation was 80%, participants with high automation expectancy were more poorly calibrated than participants with low automation expectancy if the reliability in the previous block was higher (100%), but better calibrated if the reliability in the previous block was lower (60%).

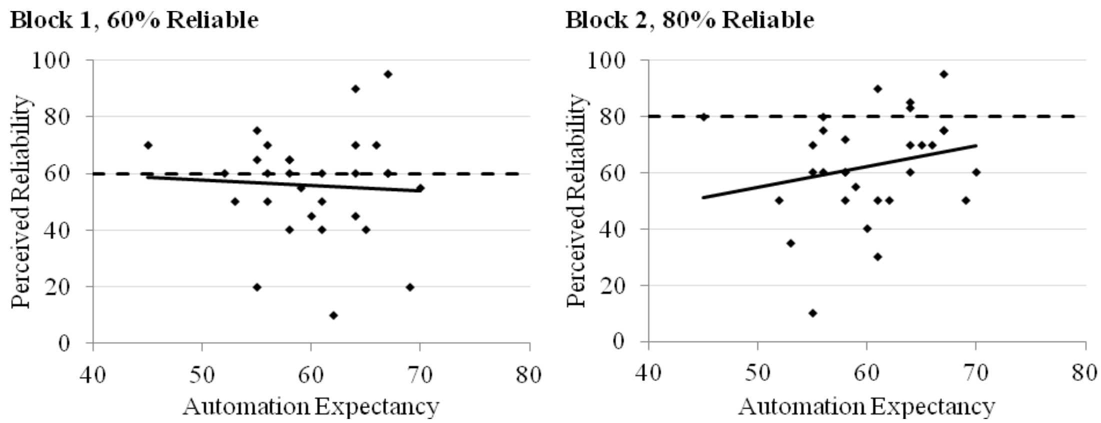

Figure 8. Actual (dotted) and perceived (solid) reliability for each trial block in the 60-80 reliability condition.

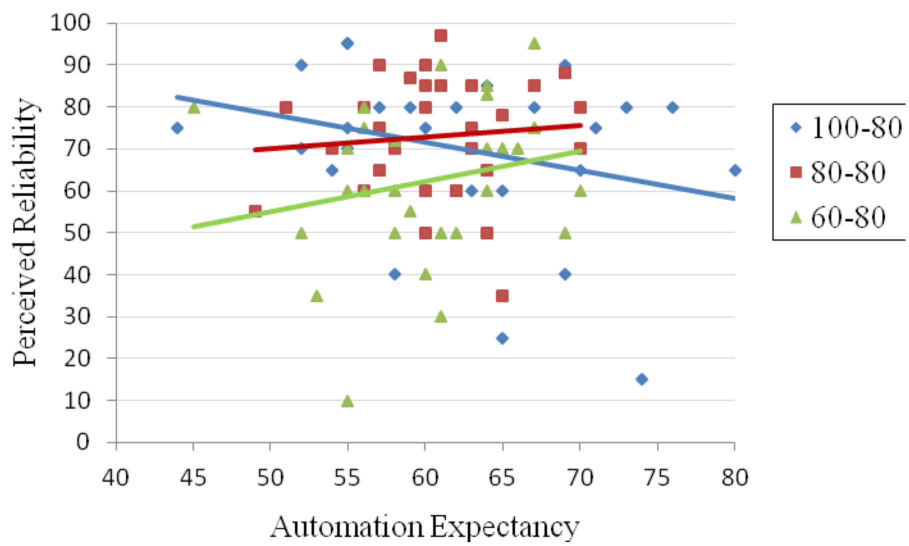

Figure 9. Perceived reliability by automation expectancy in the second block for each reliability condition in the control information condition.

The separate three-way analysis for the causal information condition revealed that the global four-way interaction was likely caused by the individual differences found in the control information condition being eliminated by the causal attribution information manipulation. The three-way interaction that was found in the control information condition was not significant when attribution biasing information was given, $R^2 = .01$,

$F(2,80) = 2.154$, $p = .123$, $\eta p^2 = .051$. In the control information condition, when the automation dropped from 100% to 80% reliable in the second block, participants with a high automation expectancy had much lower estimates of reliability (see Figure 10). In the causal attribution information condition, this was not the case. As illustrated in Figure 11, participants were relatively equally calibrated in the second block of the 100-80 reliability condition, regardless of automation expectancy. When the automation went from 60% to 80% reliable in the second block, participants in the control information condition with higher automation expectancy had higher and better calibrated estimates of reliability than participants with lower automation expectancy (see Figure 10). Once again, this was not the case in the causal attribution information condition. In fact, block 2 of the 60-80 reliability condition (illustrated in Figure 11) indicates that the opposite may be true, although we once again state that the three-way interaction was not significant.

The three-way analysis revealed a single significant effect; a two-way interaction between automation expectancy and block, $R^2 = .01$, $F(1,80) = 4.641$, $p = .034$, $\eta p^2 = .055$ (see Table 2 for full results). In order to understand this interaction (illustrated in Figure 12), separate analyses were used to investigate the effect of automation expectancy in each block. Automation expectancy was not a significant factor in either the first block, $R^2 = .0008$, $F(1,84) = 1.152$, $p = .286$, $\eta p^2 = .014$, or the second block, $R^2 = .00008$, $F(1,84) = .189$, $p = .655$, $\eta p^2 = .002$, further indicating that perceived reliability did not vary based on automation expectancy, and that the causal attribution information eliminated the individual differences observed in the control information condition.

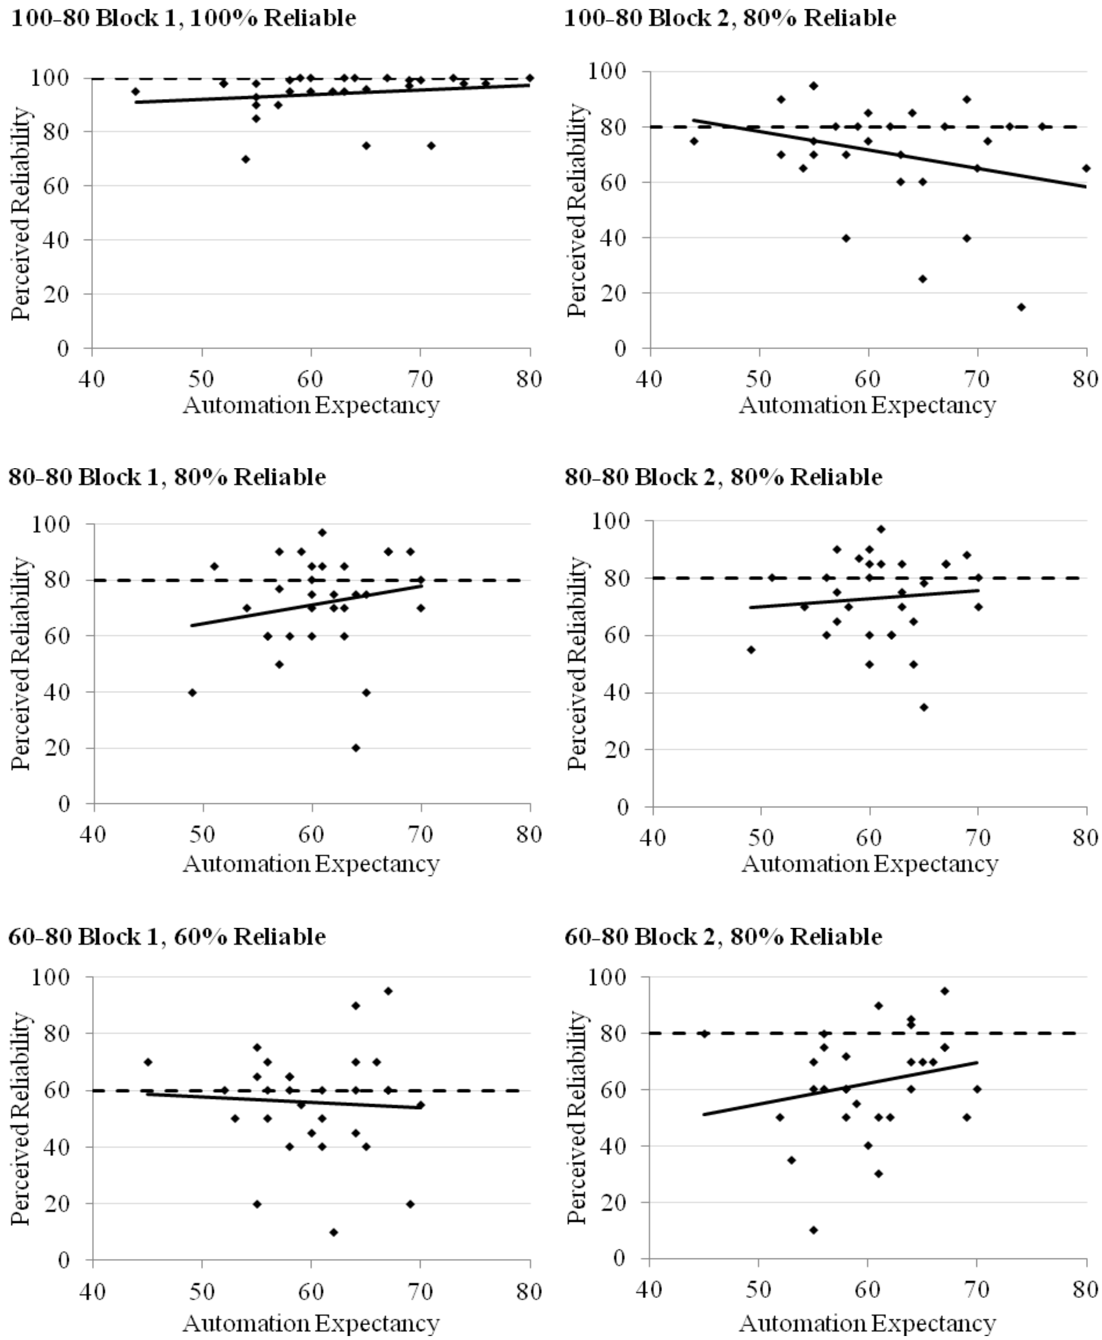

Figure 10. Actual (dotted) and perceived (solid) reliability for each trial block in the control information condition.

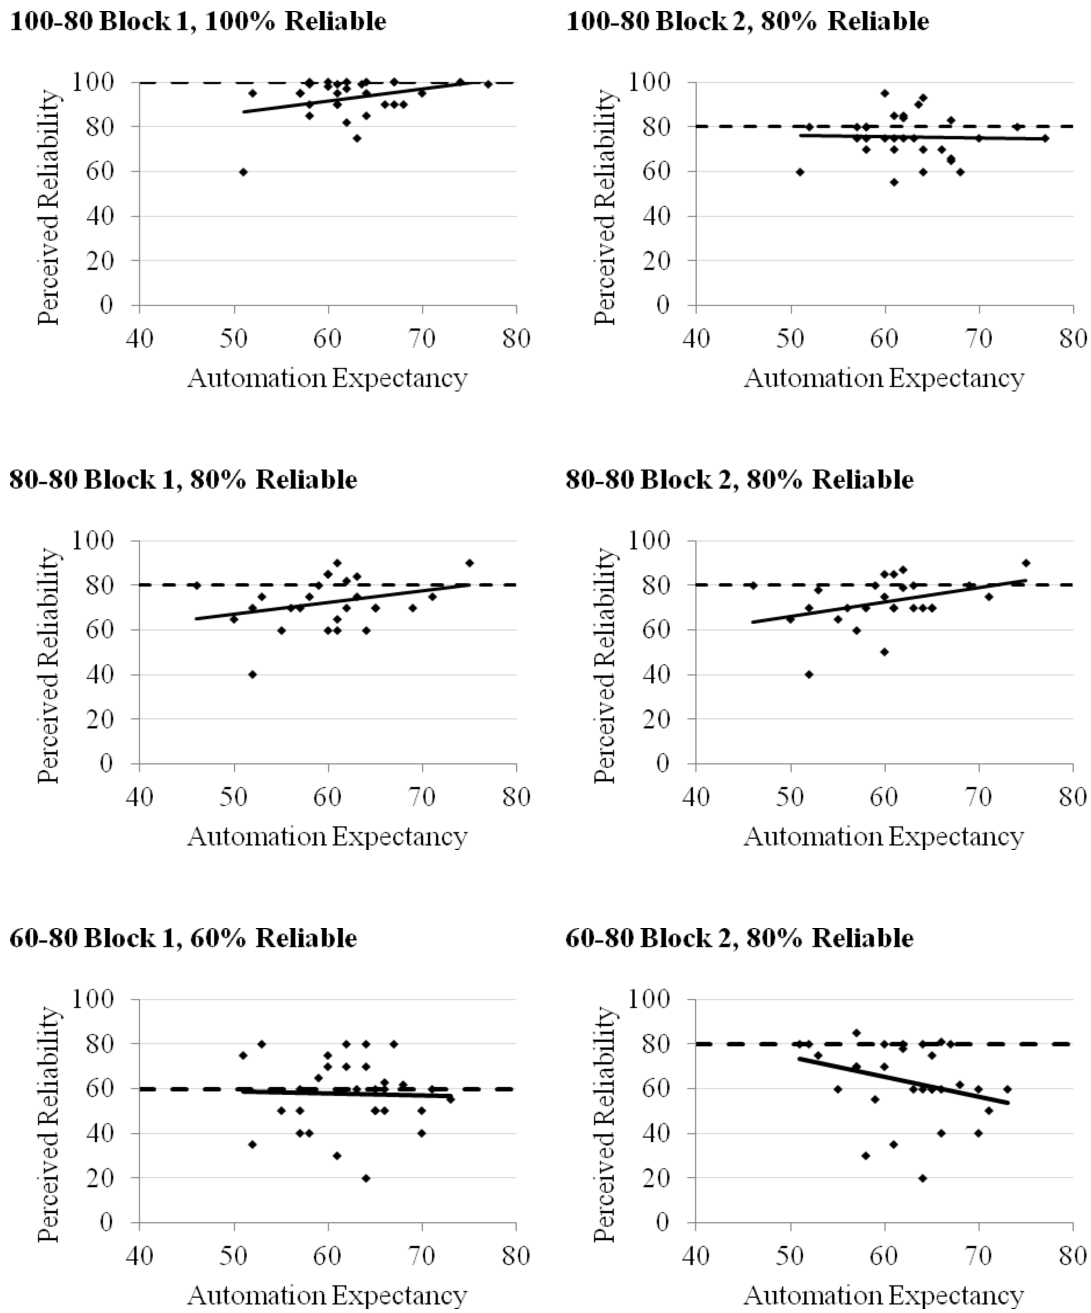

Figure 11. Actual (dotted) and perceived (solid) reliability for each trial block in the causal attribution information condition.

Table 2
Results of the linear regression with repeated measures used to analyze the perceived reliability data for the causal information condition.

Effect	Df	F	p	ηp^2
Automation Expectancy	1,80	.271	.604	.003
Block	1,80	3.446	.067	.041
Reliability Condition	2,80	1.549	.219	.037
Exp. * Block	1,80	4.641	.034*	.055
Exp. * Rel. Cond.	2,80	2.361	.101	.056
Block * Rel. Cond.	2,80	2.237	.104	.055
Exp. * Block * Rel. Cond.	2,80	2.154	.123	.051

* indicates that effect was significant at α =.05.

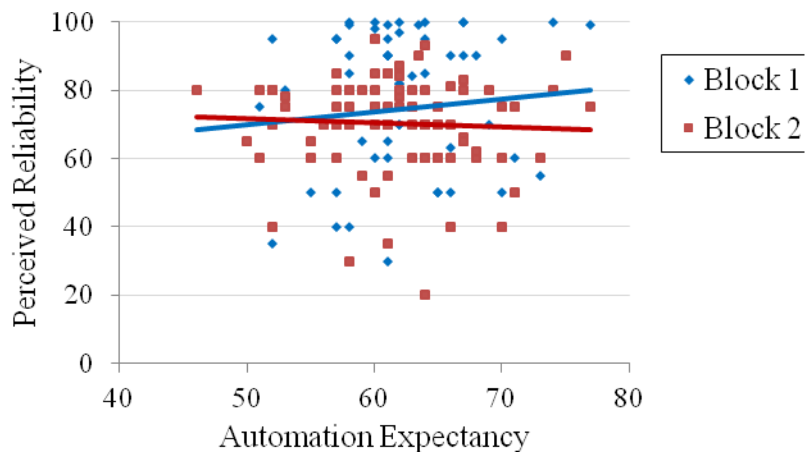

Figure 12. The significant interaction between automation expectancy and block for the causal information condition. The data are averaged across reliability conditions.

Trust

The trust scale data were analyzed using a linear regression model with repeated measures. Reliability condition, information condition, and automation expectancy were entered as between subject factors and trial block as a within subjects factor. Once again, we used partial eta-squared (ηp^2) as a measure of effect size, with .010 indicating a small effect, .059 a medium effect, and .138 a large effect (Cohen, 1988).

There was a significant interaction between reliability condition and automation expectancy, $R^2 = .02$, $F(2,164) = 3.294$, $p = .040$, $\eta p^2 = .039$. No other effects were significant (see Table 4 in Appendix A for a full results table). The significant interaction, illustrated in Figure 13, revealed that trust varied differently as a function of automation expectancy depending on reliability condition.

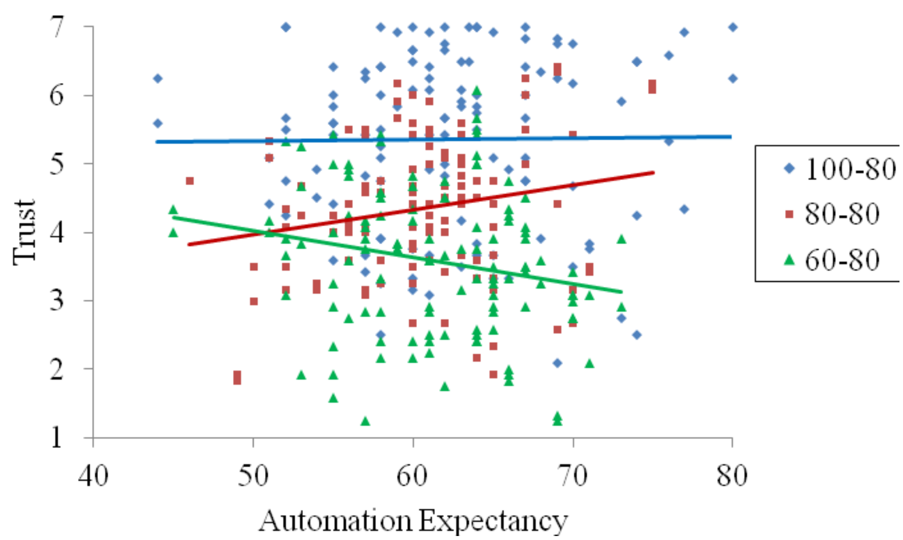

Figure 13. Trust by automation expectancy averaged over block for each reliability condition.

In the 100-80 reliability condition, trust in the automated aid was relatively high regardless of the operator's automation expectancy. This finding confirms Merritt and Ilgen’s (2008) speculation that when automation reliability is close to 100%, the optimal utilization strategy is obvious and not driven by individual differences. In the 80-80 reliability condition, operators with higher automation expectancy trusted the automation more than operators with lower automation expectancy. In the 60-80 reliability condition the opposite was true; operators with higher automation expectancy trusted the automation less than operators with lower automation expectancy. This pattern matched Merritt and Ilgen's (2008) findings that operators with higher automation expectancy suffered the largest negative effects on trust when paired with automation that had low reliability (see Figure 14). The pattern in the present interaction also matched findings of studies in the interpersonal literature that found when interacting with untrustworthy people, high expectancy individuals had lower estimates of trust than low expectancy individuals (Yamagishi et al., 1999). Taken together, these findings indicate that high expectancy individuals are more sensitive that low expectancy individuals in adjusting their trust to situations differing in trustworthiness.

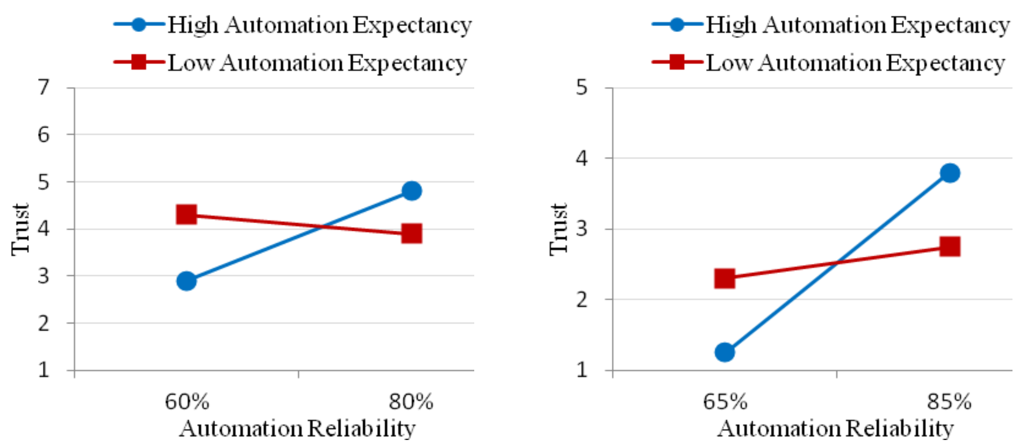

Figure 14. Interaction between automation expectancy and reliability level for the current study (left) and Merritt and Ilgen (2008; right). Note: In the current study trust was assessed on a 1 to 7 scale while in Merritt and Ilgen's (2008) study trust was assessed on a 1 to 5 scale.

Trust was highly correlated with perceived reliability, r (350) = .77, $p < .001$, so we were surprised that the significant four-way interaction found in the perceived reliability data was not significant in the trust data, $R^2 = .08$, $F(2,164) = 2.023$, $p = .136$, $\eta p^2 = .024$. We looked at the trust data and observed the same patterns as in the perceived reliability data. As illustrated in Figure 15, in the control information condition, operators with high automation expectancy were more sensitive than operators with low automation expectancy to both increases (60-80, Block 2) and decreases (100-80, Block 2) in automation reliability. Once again this individual difference diminished in the causal attribution information condition (see corresponding blocks in Figure 16). Thus it appears as if the same effects are present, however we did not have the power to detect them in the 7 point trust scale as we did in the 100 point perceived reliability scale.

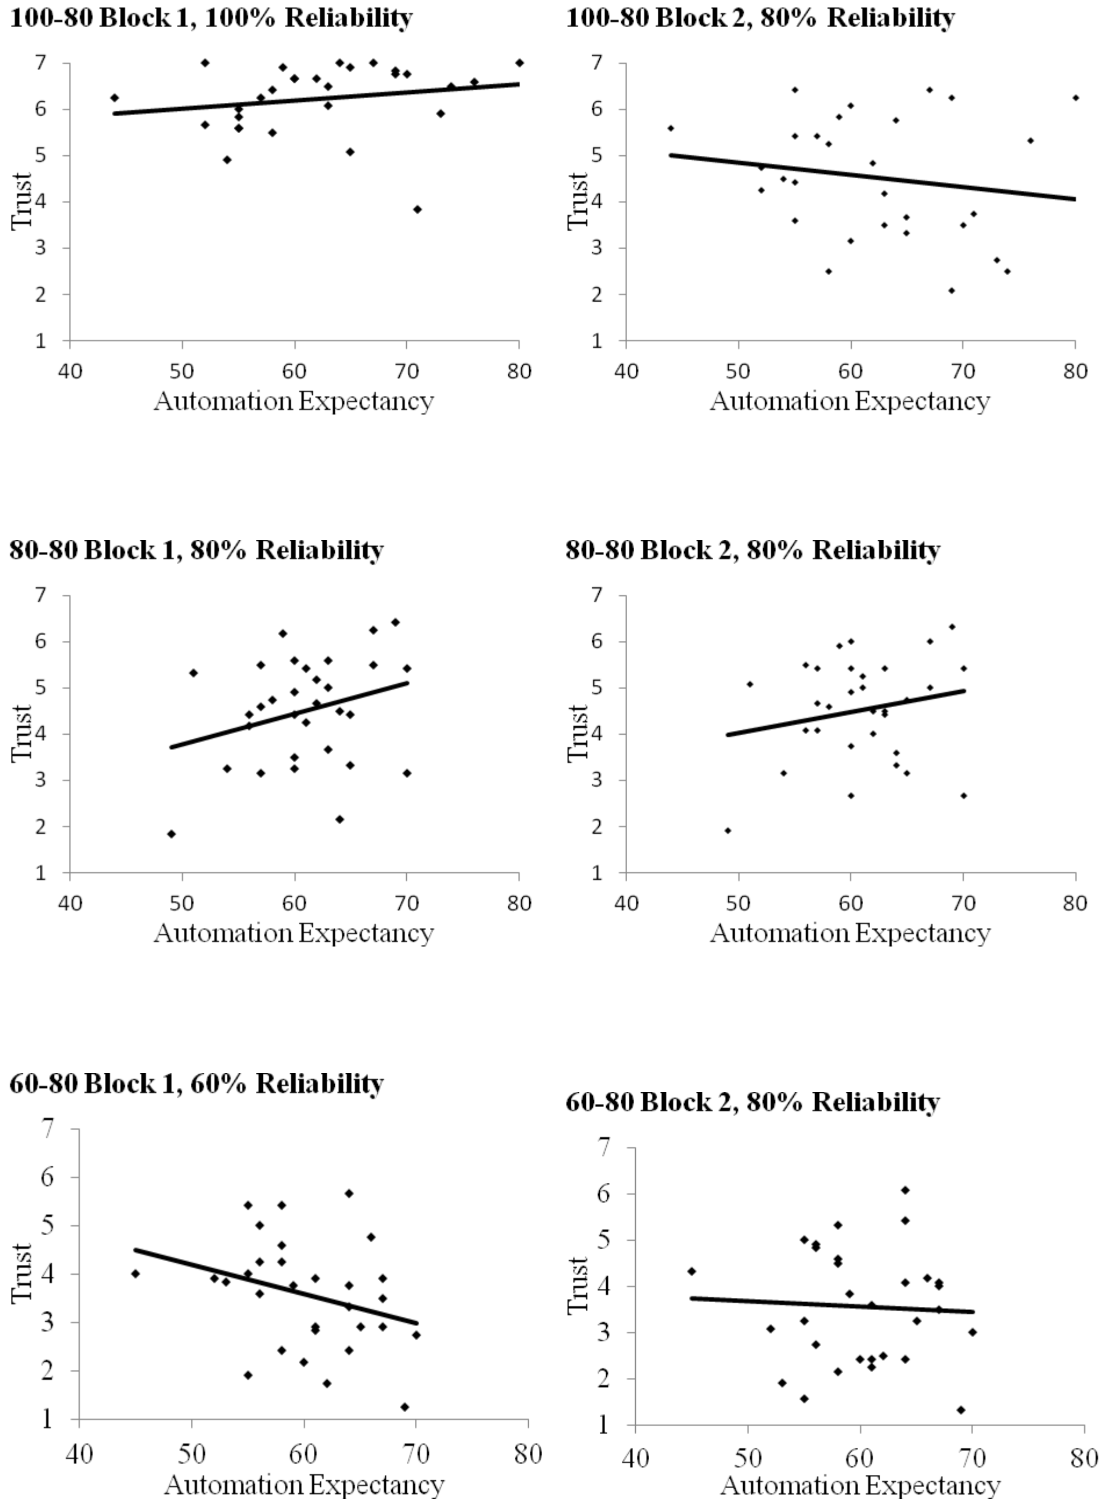

Figure 15. Trust scores for each trial block in the control information condition.

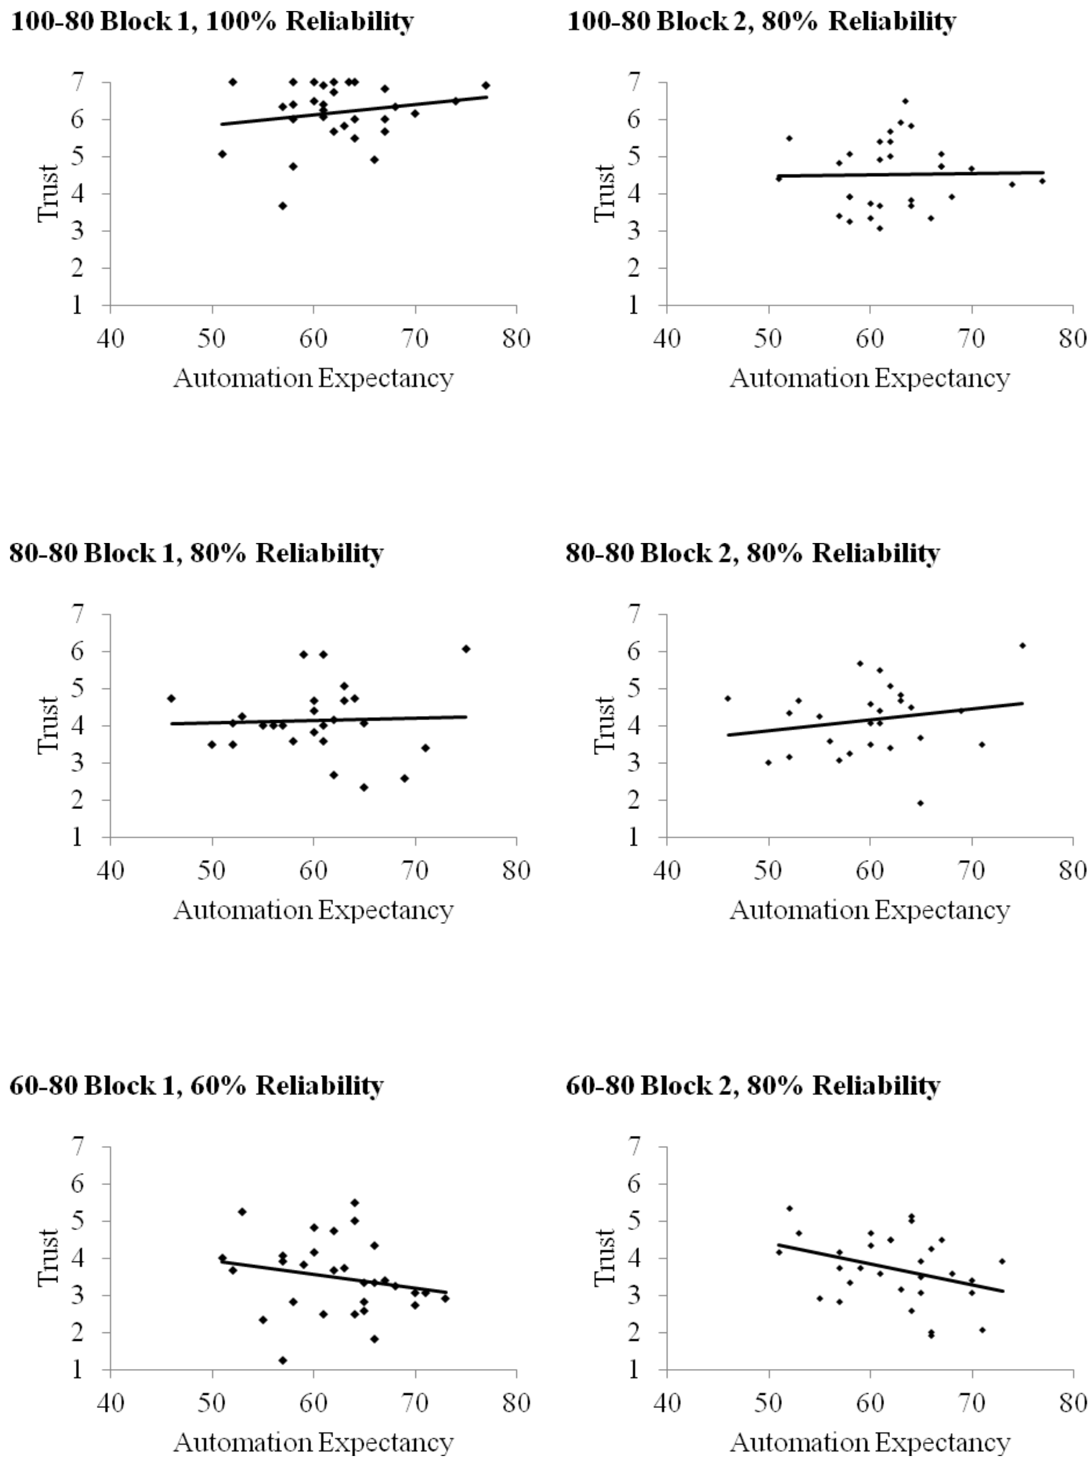

Figure 16. Trust scores for each trial block in the causal attribution information condition.

Automation Expectancy, Causal Attribution, and Trust

In order to investigate the role of causal attribution in the relationship between automation expectancy and initial trust, Baron and Kenny's (1986) mediation analysis method was used. This was the same method Dzindolet et al. (2003) used to test if the relationship between providing the operator information indicating why an automated aid might err and automation reliance was mediated by trust. Figure 17 illustrates the necessary conditions that had to be met for causal attribution to be considered a mediator; 1. A significant relation between automation expectancy and trust (*a*). 2. A significant relation between causal attribution and trust (*b*). 3. A significant relation between automation expectancy and causal attribution (*c*). 4. The relationship between automation expectancy and trust (*a*) was no longer significant after controlling for causal attribution.

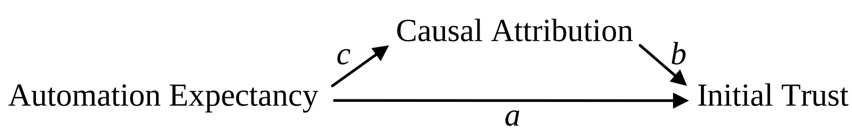

Figure 17. Conditions necessary for causal attribution to be considered a mediator.

The relationships between the automation expectancy, causal attribution, and initial trust were investigated using Pearson Correlations on data from the 90 participants in the control information condition. Data from the causal information condition were excluded to avoid confounding causal attribution with the causal information manipulation. There was a significant relationship between automation expectancy and trust (path *a*), r (88) = .26, p = .012, indicating that the more a participant expected automation to be trustworthy, the higher their initial trust. The relationship between causal attribution and trust (path *b*) was marginally significant, $r\ (88) = .20$, $p = .059$, indicating that the more participants attributed errors to the automation, the higher their initial trust. Interestingly, we did not find a significant relationship between automation expectancy and causal attribution (path *c*), $r\ (88) = .07$, $p = .493$, suggesting that participants attributed errors to the automation, regardless of their automation expectancies. The lack of a significant relationship could be related to the factors underlying automation expectancy. Automation can be reliable, dependable, and safe, regardless of whether errors are attributed as internal or external to the automation. Nevertheless, the lack of a significant relationship between automation expectancy and causal attribution violated a necessary condition of the mediation analysis, suggesting that causal attributions did not mediate the relationship between automation expectancy and initial trust.

CHAPTER 4

GENERAL DISCUSSION

The first purpose of the present study was to investigate whether operators with high automation expectancy are better at calibrating their trust to changes in the capabilities of the automation. The significant three-way interaction for perceived reliability in the control information condition indicated that higher automation expectancy resulted in better calibrated adjustment when the capability of the automation increased, but poorer calibrated adjustment when the capability decreased.

The significant interaction between reliability condition and automation expectancy for trust scores also indicated that higher automation expectancy resulted in higher sensitivity in adjusting trust to situations differing in trustworthiness. These findings may explain why previous experiments have sometimes found unstable and non-linear changes in trust. For example, Lee and Moray (1994) found that introducing errors in an automatic controller led to a substantial change in trust for some participants, yet little or no change for others. Our results indicate that this was likely due to differences in automation expectancy, but why?

The second purpose of our study was to investigate why calibration differs as a function of automation expectancy. We predicted that this effect was caused by differences in the causal attribution of automation errors. Specifically, we hypothesized that operators with high automation expectancy attributed errors as internal to the automation, leading them to re-evaluate the automation's ability and their trust in the automation, while operators with low automation expectancy attributed errors as external to the automation, thus not re-evaluating the automation's ability and their trust in the automation. Our results indicated that calibration did indeed differ as a function of automation expectancy due to differences in the causal attribution of automation errors.

In fact, the three-way interaction for perceived reliability scores in the causal information condition indicated that we were able to eliminate the effects of automation expectancy by manipulating causal attribution. However, the relationship between automation expectancy and causal attribution turned out to be the opposite of what we hypothesized. The global four-way interaction for perceived reliability indicated that manipulating the attribution of errors as internal to the automation no longer resulted in operators with high automation expectancy having better calibrated adjustment when the capability of the automation increased or poorer calibrated adjustment when the capability decreased. This finding revealed that in the control condition, operators with high automation expectancy attributed errors as external to the automation while operators with low automation expectancy attributed errors as internal to the automation. Taken together, our results indicate that operators with high automation expectancy are more sensitive in adjusting trust to situations differing in trustworthiness because they attribute errors as external to the automation (e.g., the situation or task).

The relationship between automation expectancy, causal attribution locus, and trust is not as simple as the model we proposed in the introduction (Figure 2). Trust is a multidimensional construct that is based on a variety of factors (Lee & See, 2004). Interestingly, the task and situation related reliability of automation has been defined as one of the general bases of trust in automation (Lee & Moray, 1992; Lee & See, 2004). Studies have shown that appropriate trust depends on understanding the capabilities of automation in different situations (Lee & See, 2004). Operators who attribute automation errors to task or situation characteristics foster this understanding and develop more appropriate trust, however operators that continuously attribute errors to automation characteristics do not. These findings can be used to explain the effects found in our study. Operators with high automation expectancy were more sensitive in adjusting their trust to changes in automation reliability because their external attribution of errors increased their understanding of the automation's capabilities in different situations. When causal attribution was manipulated to be internal, operators failed to get an understanding of the automation's capabilities in different situations because errors were attributed to the automation, thus we observed a decrease in the sensitivity of adjusting trust to different situations.

The results of our study are also applicable outside of the research community. Operators differing in automation expectancy may need different interventions (training) or information to successfully calibrate their trust levels. For example, operators with low automation expectancy may be given information that leads them to attribute the cause of automation errors as external to the automation (i.e. providing information indicating that automation behavior is high in consistency, high in consensus, and high in distinctiveness). Training programs can also be customized in order for operators with low automation expectancy to better understand the capabilities of automation in different situations.

As Lee and See (2004) stated, "little research has addressed the challenges of promoting appropriate trust in the face of a dynamic context that influences its capability" (p. 75). Future studies should investigate whether providing information that results in automation errors being causally attributed to the situation (i.e. high consistency, high distinctiveness, high consensus) results in operators with low automation expectancy having a higher sensitivity in adjusting trust to situations differing in reliability. Future studies may also investigate how calibrated adjustment varies as a factor of automation expectancy in a wider variety of reliability changes (e.g., automation that goes from 80% to 100% reliable, or 80% to 60% reliable).

The measurement of automation expectancy should also be examined in future studies. Researchers should consider updating items on the nearly 20 year old AICP rating scale that are somewhat outdated (e.g., using the automatic programming facility on a VCR, manually sorting through card catalogs to find items in a library). Future studies could also investigate how updates to the AICP scale impact the reliability of the underlying constructs, the internal consistency, and test-retest reliability. The relationship between the factors underlying automation expectancy (i.e., confidence, reliance, trust, safety) and causal attribution may also be explored. Ultimately, the understanding of these factors may lead to the development of a scale that can reliably measure the causal attribution of automation errors. This scale could then be used to further investigate how individual differences mediate the dynamic interaction of automation and contextual factors influencing operator trust. Understanding the dynamics of trust is critical in understanding the highly nonlinear patterns of automation reliance, and the first step in both explaining and predicting human-automation interaction.

APPENDIX A

DETAILED DATA TABLES

Table 3

Results of the linear regression with repeated measures used to analyze the perceived reliability data.

Effect	Df	F	p	ηp^2
Automation Expectancy	1,164	.658	.418	.004
Block	1,164	1.432	.233	.009
Reliability Condition	2,164	1.461	.235	.018
Information Condition	1,164	.049	.826	.000
Exp. * Block	1,164	2.558	.112	.015
Exp. * Rel. Cond.	2,164	1.266	.285	.015
Exp. * Info. Cond.	1,164	.027	.869	.000
Block * Rel. Cond.	2,164	.355	.702	.004
Block * Info. Cond.	1,164	.946	.332	.006
Rel. Cond. * Info. Cond.	2,164	1.264	.285	.015
Exp. * Block * Rel. Cond.	2,164	2.06	.131	.025
Exp. * Block * Info. Cond.	1,164	.816	.368	.005
Block * Rel. Cond * Info. Cond.	2,164	3.903	.022*	.045
Exp. * Rel. Cond. * Info Cond.	2,164	1.28	.281	.015
Exp. * Block * Rel. Cond. * Info. Cond.	2,164	4.149	.017*	.048

* indicates that effect was significant at $\alpha = .05$.

Table 4

Results of the linear regression with repeated measures used to analyze the trust data.

Effect	Df	F	p	ηp^2
Automation Expectancy	1,164	.000	.983	.000
Block	1,164	.035	.852	.000
Reliability Condition	2,164	2.438	.091	.029
Information Condition	1,164	.100	.752	.001
Exp. * Block	1,164	.270	.604	.002
Exp. * Rel. Cond.	2,164	3.294	.040*	.039
Exp. * Info. Cond.	1,164	.146	.702	.001
Block * Rel. Cond.	2,164	.270	.764	.003
Block * Info. Cond.	1,164	.015	.904	.000
Rel. Cond. * Info. Cond.	2,164	.456	.634	.006
Exp. * Block * Rel. Cond.	2,164	1.818	.166	.022
Exp. * Block * Info. Cond.	1,164	.004	.952	.000
Block * Rel. Cond * Info. Cond.	2,164	2.143	.121	.025
Exp. * Rel. Cond. * Info Cond.	2,164	.533	.588	.006
Exp. * Block * Rel. Cond. * Info. Cond.	2,164	2.023	.136	.024

* indicates that effect was significant at $\alpha = .05$.

REFERENCES

Baron, R. M., & Kenny, D. A. (1986). The mediator-moderator variable distinction in social psychological research: conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology, 51*, 1173–1182.

Cohen, J. (1988). *Statistical power analysis for the behavior sciences* (2nd ed.). New Jersey: Lawrence Erlbaum.

Dzindolet, M. T., Peterson, S. A., Pomranky, R. A., Pierce, L. G., & Beck, H. P. (2003). The role of trust in automation reliance. *International Journal of Human-Computer Studies, 58*(6), 697-718.

Dzindolet, M.T., Pierce, L.G., Beck, H.P., Dawe, L.A. & Anderson, B.W. (2001). Predicting misuse and disuse of combat identification systems. *Military Psychology*, *13*, 147 -164.

Jian, J. Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics, 4,* 53-71.

Kassin. S. M. (1979). Consensus information, prediction, and causal attribution: A review of the literature and issues. *Journal of Personality and Social Psychology, 37*(2), 1966-1981.

Keller, M. D., & Rice, S. (2010) System-wide versus component-specific trust using multiple aids. *The Journal of General Psychology, 137*, 114-128.

Kelley, H. H. (1967). Attribution theory in social psychology. *Nebraska Symposium on Motivation.* D. Levine, ed., Lincoln: University of Nebraska Press, 192-238.

Lee, J. D., & Moray, N. (1992). Trust, control strategies and allocation of function in human-machine systems. *Ergonomics, 35,* 1243–1270.

Lee, J. D., & Moray, N. (1994). Trust, self-confidence, and operators' adaptation to automation. *International Journal of Human-Computer Studies, 40,* 153–184.

Lee, J. D., & See, J. (2004). Trust in automation and technology: Designing for appropriate reliance. *Human Factors, 46,* 50–80.

Madhavan, P., Wiegmann, D. A., & Lacson, F. C. (2006). Automation failures on tasks easily performed by operators undermine trust in automated aids. *Human Factors*, *48*, 241–256.

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review, 20*, 709–734.

Merritt, S. M., & Ilgen, D. R. (2008) Not all trust is created equal: Dispositional and history-based trust in human-automation interactions. *Human Factors***,** *50*, 194-210.

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

Prinzel, L. J., Freeman, F. G., & Prinzel, H. D. (1995). Individual differences in complacency and monitoring for automation failures. *Individual Differences Research, 3*(1), 27-49.

Rotter, J. B. (1967). A new scale for the measurement of interpersonal trust. *Journal of Personality, 35*, 651– 665.

Rotter, J. B. (1971). Generalized expectancies for interpersonal trust. *American Psychologist, 26*, 443– 452.

Rotter, J. B. (1980). Interpersonal trust, trustworthiness, and gullibility. *American Psychologist, 35*, 1–7.

Russell, D. (1982). The causal dimension scale: A measure of how individuals perceive causes. *Journal of Personality and Social Psychology*, *42*(6), 1137-1145.

Singh, I. L., Molloy, R., & Parasuraman, R. (1993a). Automation-induced "complacency": Development of a complacency-potential scale. *International Journal of Aviation Psychology, 3,* 111-122.

Singh, I. L., Molloy, R., & Parasuraman, R. (1993b). Individual differences in monitoring failures of automation. *Journal of General Psychology, 120*, 357–373.

Teas, R. K., & McElroy, J. C. (1986). Causal attributions and expectancy estimates: A framework for understanding the dynamics of salesforce motivation. *Journal of Marketing, 50*, 75-86.

Tomlinson, E., & Mayer, R. (2009). The role of causal attribution dimensions in trust repair. *Academy of Management Review*, *34*(1), 85-104.

Weiner, B. (1979). A theory of motivation for some classroom experiences. *Journal of Educational Psychology*, *71*, 3-25.

Weiner, B. (1986). *An attributional model of motivation and emotion.* New York: Springer-Verlag.

Wickens, C. D., & Dixon, S. R. (2007). The benefits of imperfect diagnostic automation – a synthesis of the literature. *Theoretical Issues in Ergonomics Science*, *8*(3), 201-212.

Wiegmann, D. A., Rich, A., & Zhang, H. (2001). Automated diagnostic aids: The effects of aid reliability on users' trust and reliance. *Theoretical Issues in Ergonomics Science, 2*(4), 352-367.

Yamagishi, T., Kikuchi, M., & Kosugi, M. (1999). Trust, gullibility, and social intelligence. *Asian Journal of Social Psychology, 2*(1), 145–161.