Running heads: SOCIAL RISK AND ATTRIBUTION

QUAYLE AND NAIDOO

Social Risk and Attribution: How Considering the Social Risk of Attributions Can Improve the Performance of Kelley's ANOVA Model in Applied Research

MICHAEL QUAYLE[1]

School of Psychology

University of KwaZulu-Natal

Scottsville, South Africa

EVASEN NAIDOO

International Training and Education Center for Health (I-TECH) South Africa

Brooklyn Square, South Africa

Classic models of attribution are increasingly used, despite serious problems with their empirical validation. This study revisits Kelley's (1967) ANOVA model of attribution and argues that it will most usefully predict attributions when attributional processes are socially "safe" and have few social consequences. The results demonstrate that attributions are most likely to be inconsistent with Kelley's predictions when attributional information and the attributions themselves are socially consequential or risky, but are more likely to be made as predicted when they are socially safe. Applications of Kelley's model, therefore, should pay attention to the extent to which attributions and attributional information are socially consequential or risky, particularly when analyzing the use of consensus information.

Attribution—which is how people answer the question "Why?"—has received substantial attention in social psychology in the past half century. Indeed, Hewstone (1983) estimated that by the late 1970s, research into

[1] Correspondence concerning this article should be addressed to Michael Quayle, School of Psychology, University of KwaZulu-Natal, Private Bag X01, Scottsville, South Africa, 3209. E-mail: quaylem@ukzn.ac.za

attribution accounted for about 11% of published work in social psychology, and most of these papers were concerned with developing models that universally and reliably predict attributions (cf. Harvey & Weary, 1984). According to searches of the PsycInfo database,[2] output in the field of attribution continued to increase until the early 1990s (see Figure 1). Although interest dipped for the next 10 years, publications per year between 2005 and 2008 were higher than ever before, suggesting a resurgence of interest in the topic. This paper revisits Kelley's (1967) classic model of attribution and demonstrates that features of the local social environment are critical to how attributions are made, suggesting that Kelley's model will be most appropriate when the attributions to be made are socially neutral and have few negative local social consequences.

-INSERT FIGURE 1 ABOUT HERE[A]-

Early interest in attribution as a cornerstone of social interaction and social psychology was stimulated by the publication of Heider's (1958) work, *The Psychology of Interpersonal Relations.* One of Heider's key arguments is that answering the question "Why?" is a frequent and fundamental component of social interaction because understanding the reason for an observed action is the basis for the perceiver's response. For example, if we infer that a shove was intentional, we may respond angrily; but if we judge it to be accidental, we may be placatory. Some have gone so far as to argue that the social process of asking and answering the question "Why?" is one of the few psychological universals (Moscovici & Hewstone, 1983).

Heider (1958) argued that people are lay psychologists who are driven to infer regular causes from observations, and thereby increase their ability to predict and control their social and physical environments. He suggested that one psychological process underlying everyday attribution may be similar to the analysis of *covariation*, which is the extent to which potential causes vary in relation to each other and to the event or observation to be explained. This idea was taken up by Kelley (1967), who argued that people—more than being naïve psychologists—are naïve *scientists*, who systematically analyze specific types of information about people and events to arrive at accurate attributions through a cognitive process akin to the statistical technique of ANOVA. Both Heider's and Kelley's ideas were taken up with enthusiasm by researchers, and these models laid much of the

---

[2]For publications per year, with keywords or words in the title starting with "attribution."

groundwork for empirical work in the field that came to be known as *attribution theory* until at least the 1980s (for a review, see Harvey & Weary, 1984).

Surprisingly, in spite of the dip in publication in the field of attribution theory in psychological literature between 1995 and 2004 (according to PsycInfo records), a broader citation search using Google Scholar reveals that Heider (1958) and Kelley (1967) have both been cited with unbroken exponential frequency since their publication (see Figure 2). While it is true that this pattern mirrors the "knowledge explosion" that has resulted in exponential publication rates in even the smallest of fields, it also demonstrates that these models—far from being forgotten relics of our theoretical past—are both cited more now than ever before. Interestingly, much like Maslow's (1943) famous hierarchy of needs, the work of Heider and Kelley is increasingly being applied outside of psychology. In fact, the ISI Citation Index reports that 49% of citations of Heider and 56% of citations of Kelley between 2005 and 2008 were in fields other than psychology, such as business (e.g., Janakiraman, Meyer, & Morales, 2006), human–computer interaction (e.g., Thatcher, Zimmer, Gundlach, & McKnight, 2008), management (e.g., Keaveney, 2008), women's studies (e.g., Smith, Tabak, Showail, Parks, & Kleist, 2005), ethics (e.g., Sanchez, Gomez, & Wated, 2008), sports science (e.g., Shapcott, Carron, Greenlees, & El Hakim, 2008), forestry (e.g., Arvai, Gregory, Ohlson, Blackwell, & Gray, 2006), and so forth.[3]

-INSERT FIGURE 2 ABOUT HERE[B]-

Despite the ever-increasing popularity of these models, there are some serious problems with their empirical validation (e.g., McArthur, 1972) that have never been solved. The present paper returns to some of the anomalies that emerged from empirical tests of Kelley's (1967) model of causal attribution and offers a simple suggestion that may improve the applicability of the model to real-life attributional situations. Specifically, we argue that attributions—and the information required to produce them—are often socially embedded and socially consequential (cf. Edwards & Potter, 1993). We suggest and empirically test the idea that when attributions are socially consequential or risky, they are most likely to deviate from Kelley's ANOVA model; but when attributional

[3] The examples all cite Kelley (1967).

processes and their outcomes are socially disengaged and unrelated to important social consequences (e.g., accountability), it is more likely that people will act as naïve scientists, in accordance with Kelley's ANOVA model.

Kelley's Model: The Lay Scientist's ANOVA

Kelley's (1967) model centers around the simple idea that an effect may be attributed to "that condition which is present when the effect is present and absent when the effect is absent" (p. 194). Kelley argued that making a confident attribution of causality requires information about entities or things (distinctiveness), time and modality (consistency), and social agreement (consensus). Kelley summarized as follows:

> The attribution to the external thing rather than to the self requires that I respond *differentially* to the thing, that I respond *consistently*, over time and over modality, and that I respond *in agreement* with a consensus of other persons' responses to it. ([C]p. 195)

Applying the model to the trivial problem of deciding whether one likes a movie because of some property of oneself or because of some property of the movie, Kelley's model has argued that

> The movie is judged to be enjoyable if I enjoy only it (or at least not all other movies), if I enjoy it even the second time, if I enjoy it on TV as well as at the drive-in theatre, and if others also enjoy it. (p. 195)

One of the features that makes this model so intuitively valid and easily applicable is that Kelley (1967) was able to represent attribution-relevant information along these dimensions to form an attribution matrix, which is commonly referred to as *Kelley's ANOVA cube* (see Figure 3). In this diagram, distinctiveness information (i.e., about different entities) is represented on the Y axis (labeled N, O, P, and Q), consistency information (i.e., about different times and modalities) is represented on the X axis (labeled T1M1, T2M2, T1M2, and T2M2), and consensus information (i.e., about different people) is represented on the Z axis (labeled self, O_1, O_2, and O_3).

-INSERT FIGURE 3 ABOUT HERE[D]-

In Kelley's (1967) illustration, Entity N (the shaded top row in the diagram) is arrived at as the most likely cause of Effect Y by the following logic: First, Self experienced Effect Y when experiencing Entity N for all combinations of time and modality (i.e., consistency); second, although Person O_1 and O_2 did not experience Entity N at all times and modalities, they did experience Effect Y each time they experienced Entity N (i.e., consensus); and third, Effect Y does not result from experiencing any other entity (i.e., distinctiveness). Therefore, because of this configuration of consistency, consensus, and distinctiveness information, Kelley argued that the cause of Effect Y will most likely be attributed to Entity N.

There are two problems with this visual presentation of the ANOVA cube. The first is that two dimensions (i.e., time and modality) are represented on one axis. However, there is general agreement that these are similar enough to collapse into the single category of *consistency information*. The second is that each axis can theoretically include unlimited numbers of entities, times, and modalities or people and so, in real-life situations, this representation soon becomes unmanageably large. Therefore Kelley (1967)—and almost all later commentators—collapsed each axis to a binary *high* or *low*, resulting in eight possible combinations where:

- The level of distinctiveness is determined by the constancy of effect in response to different entities, where multiple effects indicate *high* distinctiveness, while consistent effects across entities indicate *low* distinctiveness.
- Consistency is determined by the extent to which the response occurs in similar situations (across time and modality), where a higher frequency indicates high consistency, and a lower frequency indicates low consistency.
- The level of consensus is determined by the number of people who respond similarly in similar situations, where a larger number of people indicates high consensus, and a smaller number indicates low consensus.

The combinations of information most likely to result in attributions to the person, entity, or circumstances (Kelley, 1967, 1971; Orvis, Cunningham, & Kelley, 1975) are reported in Table 1.

-INSERT TABLE 1 ABOUT HERE[E]-

Although Kelley (1967) admitted that attribution processes that deviate from the naïve ANOVA model result in “biases, errors, and illusions” (p. 197), he maintained that “for the most part . . . [an observer] generally acts like a good scientist, examining the covariation between a given effect and various possible causes” (Kelley, 1971, p. 2). Kelley touched on arguably the most important source of “error” in his 1967 paper when he discussed the difficulties of gathering consensus information. He noted that it may be difficult to tell whether someone really agrees or whether they are merely appearing to agree and asked “What does it mean that another person’s reaction is similar to my own?” (p. 204). However, he sidestepped the issue, saying “This question is too complex to take up here, concerning as it does basic matters of interpersonal communication and comparison, ranging from emotional expression and perception to semantics and verbal labels" (p. 204).

This voluntary blind spot was not addressed as the field of attribution theory matured and, increasingly, treated attribution as a perceptual and largely individual process. The complex questions raised by the social embeddedness of attribution noted by Kelley (1967) and Heider (1958) were not taken up until the 1990s and beyond (e.g., Anderson, Beattie, & Spencer, 2001; Edwards & Potter, 1993; Hilton, Mathes, & Trabasso, 1992). However, by this time, many unresolved anomalies had accumulated in empirical tests of the model.

Empirical Anomalies With Consensus Information

The first comprehensive test of the model (McArthur, 1972) reported results that were largely consistent with Kelley's (1967) formulation, and these were consistently replicated (see Frieze & Weiner, 1971; Ruble & Feldman, 1976; Zuckerman, 1979). However, McArthur also found that consensus information was given almost no weight by participants (accounting for less than 1% of the variance), compared to distinctiveness and consistency information, which were both used largely as expected. Similarly, Major (1980) showed that when given a choice, participants were less likely to request consensus information than consistency or distinctiveness information in order to arrive at attributions. Nisbett and Borgida (1975) showed that consensus has no effect on attribution at all and suggested that consensus information may be perceived as less stable—and, therefore, less useful—than consistency or distinctiveness information because it is derived from opinions that are themselves inherently unstable (also see Cooper, Jones, & Tuller, 1972, cited in Nisbett & Borgida, 1975; Miller, Gillen, Schenker, & Radlove, 1973, cited in Nisbett & Borgida, 1975). Others found that consensus information is used, but is very sensitive to order effects (Ruble & Feldman, 1976; Zuckerman, 1979).

Ruble and Feldman (1976) noticed that McArthur (1972, 1976) and Orvis et al. (1975) presented consensus information first, followed by distinctiveness and consistency information, and hypothesized that the reduced impact of consensus information was a result of a recency effect. After counterbalancing the order of presentation, they found that all three types of information accounted for similar amounts of variance, and they concluded that consensus information only has a reduced effect on attributions when presented first (Ruble & Feldman, 1976). However, Zuckerman (1978) further explored these findings by distinguishing between attributions for non-agentic *occurrences* and agentic *action* and concluded that the order of presentation of the three types of information "affects the use of consensus information in the attribution of occurrences, but not of action" (p. 649) and that the effect of consensus is strongest for attributions for agentic action. Similarly, Pilkonis (1977) found that dispositional attributions were affected by consensus information, but situational attributions were not.

Wells and Harvey (1977) took a different approach, arguing that many studies in the field found no effect for consensus information as a result of insufficient "operationalization of either the dependent or the independent variable" (p. 280). They replicated Nisbett and Borgida's (1975) study, changing the operationalization of consensus "to reflect Kelley's conceptual criterion of perceived covariation and to represent more levels of consensus" (p. 279). They found that "consensus information significantly affected attributions in a manner consistent with attribution theory predictions" (p. 291).

Nevertheless, in spite of concerted research efforts, researchers in the field never came to agreement about why consensus information was not used by participants, as predicted by Kelley's (1967) model. In their review of the field, Harvey and Weary (1984) tried to make sense of these anomalies by arguing that the influence of consensus information depends on situational constraints (e.g., Kassin, 1979; Solomon et al., 1981 cited in Harvey & Weary, 1984). However, they[F] concluded that—contrary to Kelley's ANOVA model—"there remains the possibility that the three kinds of information are treated differently" (Kelley & Michela, 1980, p. 463).

The Social Nature of Consensus Information

In his original formulation of the model, Kelley (1967) had already foreseen—but brushed aside—the idea that consensus information may be different because of its social nature, pointing out that obtaining consensus information is a fundamentally social activity (discussed previously). He also suggested that providing consensus information is an activity vulnerable to lack of social support or lack of self-confidence. He noted that in an attempt to satisfy the consensus criterion, individuals may influence those who are vulnerable in order to produce a consensual attribution, even at the expense of veridical accuracy.

These were early presentiments (a) that people (even scientists) make attributions actively and with purpose, rather than as passive perceivers of incoming information (e.g., Hilton et al., 1992); and (b) in many settings, the information required to make attributions is not socially neutral because it is crucial to features of social relations, such as blame and accountability (Edwards & Potter, 1993). Indeed, the reason that attribution theory resulted in such concerted effort from social psychologists is that it was immediately evident that the outcome of attribution is pivotal to group processes, as it determines who can claim glory for positive outcomes and who must accept responsibility for failures. However, the focus on individual perceptual processes that quickly developed in the field neglected the socially embedded nature of attributional production. Although it is possible to imagine situations in which distinctiveness and consistency information are socially consequential, and situations in which consensus information is not socially consequential (as in self-consensus information), it is also evident that consensus is, by definition, the most socially embedded of the three types contributing to Kelley's (1967) ANOVA model.

The notion that consensus information is treated differently, depending on specific features of the social context in which attributions are produced, was supported in a content analysis of the use of Kelley's (1967) information types in the context of computer repair (Quayle, 2005). It emerged that consensus information was significantly less likely to be sought or exchanged by computer technicians and users than was consistency or distinctiveness information. Further analysis suggested that consensus information was being substantially underutilized in social settings in which seeking or reporting that kind of social information carried a risk of causing offense (Quayle, 2005). Furthermore, a conversation analysis revealed that participants were treating consensus information very carefully, with a great deal of respect for the repercussions that such information (e.g., the fact that a particular user may be at fault) may have on the immediate interaction and the longer term reputations and moral careers (cf. Goffman, 1963) of users and technicians (Quayle & Durrheim, 2006, 2008). This demonstrated, first, that consensus information in that context is social and is socially risky; and second, that even in a highly technical process (e.g., computer repair), people may be as concerned with the social features of the attributional context as with the veridical accuracy of the attribution.

These findings alerted us to the idea that specific social features of a setting may have a powerful impact on how attributional information might be sought and used. Specifically, when attributions relate to accountability and potential blame, then the gathering and giving of information—particularly consensus information—and the production of attributions is socially risky. The findings also hinted that the fundamentally social nature of consensus information may result in it being applied anomalously in many settings, compared to consistency and distinctiveness information. This is not to say that the notion of social risk may never apply to distinctiveness or consistency information, but since consensus information by definition evaluates and compares actors' actions and reactions, it seems likely to be most sensitive to social risk.

Hypotheses

Since attributions (a) take place in a social context; (b) contain social information; and (c) have social implications, it is hypothesized that people form and report attributions that are sensitive to the nature of the social context. More specifically, it is hypothesized that consensus information will be used in congruence with Kelley's (1967) model in tightly constrained social situations in which there are few social risks in doing so.[4] Conversely, when gathering or using information required to produce attributions carries social risk, it is hypothesized that the social functions of attribution will override the need to "generally act like a good scientist" (Kelley, 1967, p. 2), and socially risky information will be strategically ignored or underutilized.[5G]

Method

Participants were presented with several video vignettes showing people failing to withdraw money from an automated teller machine (ATM) and then were asked to determine the cause for the failure. Video vignettes were chosen to avoid the many problems with paper-and-pencil methods and also the problems of encoding attributional problems in language without predetermining the probable conclusion (cf. Brown & Fish, 1983). Participants' responses were given verbally and were recorded by the researcher to ensure that participants experienced the giving of attributions as a real social context.

The video vignettes were designed to reflect the combinations of consensus, consistency, and distinctiveness information most likely to result in attributions to the person or entity predicted by Kelley's (1967, 1971) ANOVA model (see Table 1). The vignettes were scripted from observations of ATM users in local malls and were all based on a movement script that ensured that they were realistic, and similar in pace and length. Each clip depicted an actor walking up to an ATM and attempting to withdraw money. There were three variables that were manipulated in the video clips: the identity of the user (i.e., consensus); whether the left or right ATM machine was used (i.e., distinctiveness); and the failure or success of the withdrawal (i.e., outcome). Although we were successful in operationalizing consensus and distinctiveness visually and spatially, our attempt to encode consistency information was less successful (as we will discuss later), and it was not included in data collection on the second site or in the final analysis.

---

[4] Of course, the context of "science," in which strict veridical accuracy is a highly valued social norm, is just such a context.

[5] Note that it is not necessary to distinguish between conscious and unconscious processes in this regard (cf. Edwards, 1993).

Operationalizing Consensus Information

Consensus information refers to how the person's behavior compares with that of others in a similar situation and was depicted by using two visibly different actors who experienced similar or different outcomes in the same situation. The researcher acted as one of the ATM users to allow the *social safety* manipulation (which will be discussed later), and an actor played the role of the second user. They dressed in a similar manner (blue jeans, long sweater, similar shoes) to minimize participants making attributions based on dress and appearance, but one of them (the researcher) wore a baseball cap as a means of differentiation. The name "Bob" or "Jon" was superimposed on the screen while the actor made the transaction. These names (after McArthur, 1972) are both "Western," masculine names that are three letters long, and they were distinctive enough to allow attribution differentiation, but similar enough to avoid unintended confounds; for example, related to any cultural connotations that might be related to each name.

Operationalizing Distinctiveness Information

Distinctiveness refers to how the person's behavior varies across situations. It was depicted by filming the vignettes so that the visual scene always contained an ATM machine on the left and another one on the right. Distinctiveness information was then operationalized by the users experiencing constant success at one ATM and constant failure at the other.

Difficulties in Operationalizing Consistency Information

Consistency refers to the variation of outcome across time or modality. We attempted to operationalize consistency by showing the users experiencing constant success during one time period and constant failure during another. We attempted to encode time differences by showing a 24-hr clock on the wall between the ATMs and superimposing subtitles over the scene that read "during office hours" or "after office hours."

However, the features used to show time were much more abstract and difficult for participants to process, compared to simply noticing left or right ATMs or recognizably different actors. Discussions with participants at the first site revealed that this manipulation was markedly less effective than were the other two spatially represented manipulations, and required doubling the length of the vignette shown to each participant. Therefore, although consistency is an important part of Kelley's (1967) model, we considered our efforts to operationalize it unsuccessful and dropped it from the study. If this design is replicated, it may be possible to differentiate between times of day by showing action in an outdoor location and manipulating the brightness of the scene to indicate day or night.

Operationalizing Outcome

Outcome was operationalized by graphically showing success or failure to withdraw money. Success was indicated by clearly showing the user putting the withdrawn money into his wallet before walking away. Failure was represented by showing the user turning back to look at the ATM, showing disappointed body language (e.g., shrugging) and then putting only the ATM card into his wallet before walking away.

These manipulations were combined to form eight vignettes, as displayed in Table 2. Each vignette started with an introduction that oriented the participants to the types of information used in the clips by displaying a still image of each type of information briefly on the screen. The vignette then displayed two instances of successful transactions, and two instances of failed transactions chosen so that each vignette provided enough information to allow a confident attribution to either one of the actors (consensus) or one of the ATM machines (distinctiveness) according to Kelley's (1967, 1971) model. The order of clips in each vignette was randomized.

-INSERT TABLE 2 ABOUT HERE[H]-

Social Risk

The final variable—the riskiness of the social context—was operationalized by manipulating the apparent identity of the interviewer. This manipulation is similar to those in experiments exploring race of interviewer effects where different interviewers are used to produce different social conditions (e.g., Davis & Silver, 2003). However, in this procedure, the researcher was constant, but made minor modifications to his appearance to present himself either as one of the users in the vignettes (Jon) or as neutral researcher.

When the researcher took on the socially risky user role, he dressed in exactly the same clothes worn by the character Jon in the vignettes (i.e., black cap, black sweater, blue jeans, black shoes), he wore a name tag with the name "Jon," he introduced himself as "Jon," and he casually mentioned that he was one of the users in the vignettes. This increased the social risk for participants, since attributing blame to "Jon" would involve casting blame on a person with whom they were interacting directly.

When appearing as a neutral researcher, he removed the cap, replaced the black sweater with a blue T-shirt, and replaced the nametag “Jon” with one bearing his own name. This reduced the social risk of attribution, since in this condition, “Jon” and “Bob” were both absent third parties. With the exception of his appearance and initial introduction, the researcher’s behavior was scripted and was as consistent as possible across conditions.

Assessing Attributional Congruence With Kelley’s Model

After watching the vignettes, participants were asked the open-ended question “So, what went wrong?” Participants provided verbal responses that were audio recorded and simultaneously noted by the researcher on a response sheet. These were later coded as congruent or incongruent with Kelley’s (1967) ANOVA model. We hypothesized that participants’ answers would be most likely to deviate from the predictions of Kelley’s (1967, 1971) model when doing so would involve the socially risky act of assigning accountability to someone physically present, and they would be most likely to conform to Kelley’s predictions when doing so was socially safe.

Sample

The first sampling site was the university [I]stand at an annual weeklong agricultural fair that draws over 200,000 visitors every year from all walks of life, at which 138 people participated. A further 96 participants were then sampled from a local university by approaching passersby in a public space on campus.

Results

The hypothesis was that participants would be likely to underutilize or ignore consensus information in situations where the attribution carried a social risk, which was operationalized as the combination of the researcher in the ATM-user role and the provision of attributional information that implicated the researcher in the failure. In other words, it was hypothesized that the social riskiness of the setting would have an effect on the congruence of the attribution made with the predictions of Kelley’s (1967) model for that combination of information. Therefore, the hypothesis would be supported by an interaction effect between social risk and the pattern of attributional information on the congruence of the answer. More specifically, this interaction should only be evident when social risk is high (i.e., when the researcher presents himself as the ATM user) and when consensus and distinctiveness are both low, suggesting that the researcher-as-user was at fault. Additionally, it was hypothesized that the accuracy of the answer would not substantially deviate from Kelley’s predictions for the other conditions. It was also expected that the majority of the answers would be correct because of the simplicity of the task and the adequate provision of the necessary information to participants.

Given that the design includes a categorical dependent variable and categorical independent variables, binary logistic regression was the most appropriate method of analysis (Hair, Anderson, Tatham, & Black, 1998). This method has the advantage that it can be used to assess interaction effects in categorical data while differentiating between dependent and independent variables (in comparison to techniques such as loglinear analysis that explore all relationships between variables without distinguishing between predictors and outcomes). Forward stepwise selection (using the likelihood-ratio method) was used to select the variables contributing the most to the prediction of incongruent or congruent responses.

As displayed in Figure 4, a fairly small proportion of participants made incongruent attributions, but there appear to be more incongruent responses in the researcher-as-user group (bottom left) than in the researcher-as-researcher group (top left). Additionally, there appear to be more incongruent responses in the attribution to person group than in the attribution to entity group.

-INSERT FIGURE 4 ABOUT HERE[J]-

As hypothesized, a backward stepwise selection procedure removed the term for information configuration, confirming that knowing the predicted target for the pattern of information received by participants did not independently predict whether they correctly attributed the cause of the failure. The optimal model consisted of social risk and the interaction between social risk and information configuration. The omnibus tests for both the saturated and optimal models were significant, although the optimal model was significant at the .01 level, while the saturated model was only significant at the .05 level (see [K]Table 4). Although the estimates of effect size for the optimal model were low (Cox & Snell's pseudo $r^2 = .033$; Nagelkerke's pseudo $r^2 = .053$), this is unsurprising since the expected effect lies with the small proportion of participants returning incongruent responses (20.2%), rather than the much larger group who report congruent attributions. The odds of receiving an incorrect answer from a participant in the experimental condition increases by 2.724 (see Table 4). However, the 95% confidence interval for the odds ratio indicates that the odds could be as high as 5.494 or as low as 1.055.

-INSERT TABLE 4 ABOUT HERE[L]-

To check that it was appropriate to combine the data from the two sites, the analysis was repeated, including site in the stepwise selection procedure. The dummy-coded variable representing the data-collection site did not enter the model, $\chi^2(1) = 0.57$, $p = .45$, and the overall model remained unchanged, retaining only the interaction effect with identical slope and log-odds coefficients.

The binary logistic regression was successful in identifying the significant interaction, but did not allow us to isolate the source of variability in incorrect answers between the four possible combinations of person and entity information possible in the interaction effect. However, the multinomial logistic regression procedure can produce valid results even when the predictor and response variables are categorical (SPSS, 2007). This justified running a multinomial logistic regression model to produce parameter estimates for the interaction effect, which are not available for SPSS binary logistic regression procedures. This was significant, $\chi^2(3) = 7.86$, $p = .049$, in line with the results of the binary logistic regression model reported previously.

The parameter estimates reveal that only the combination of researcher-as-user (the socially risky condition) and an information configuration showing the user to be at fault significantly contributed to the prediction of incongruent attributions, Wald = 4.31, $p = .038$, Exp(b) = 2.63, 1.055 ≤ odds ratio (OR) ≤ 6.550. The OR of 2.629 indicates that the odds of receiving an answer incongruent with Kelley's (1967, 1971) model from a participant who has received information showing the user to be at fault, but who also believes that the researcher was the user, increased by 2.629. These results correspond well to those of the binary logistic model, but isolate the source of variability of incorrect answers to the socially risky situation where the researcher is perceived to be the user, and the attributional information supplied shows the user to be to blame.

After attributing the cause of failure, participants were asked how certain they were of their answer on a scale from 1 to 10. Participants who reported incorrect answers reported being substantially less certain ($M = 5.99$) than did participants who answered correctly ($M = 9.26$), although this difference only approached significance, $t(221) = -1.82$, $p = .07$. However, there was no significant difference in certainty between those in the socially risky and socially safe conditions, and no significant interaction effect on certainty by social risk and correctness.

Discussion

The present research argues that Kelley's (1967) ANOVA model of attribution may not accurately describe real-life attributions in situations in which making attributions—or seeking or giving information required to make attributions—is socially consequential or risky. Although a full test of Kelley's model was not possible because of the difficulties of visually operationalizing consistency information. the analysis suggests that when comparing consensus and distinctiveness information, answers are likely to diverge from the predictions of Kelley's model when attributions point toward attributing blame to people, which makes such an attribution socially consequential or risky. Indeed, the odds of making an attribution incongruent with Kelley's predictions in the socially risky case when the information configuration pointed to the researcher was increased by approximately 2.629 times, compared to other conditions, and no other conditions deviated substantially from congruence with Kelley's predictions.

Since consensus information is most likely to have social value and currency because (with the exception of self-consensus) it is, by definition, about social agreement and its collection requires social interaction (cf. Quayle & Durrheim, 2008), its application is more likely to have social repercussions than other types of information. Therefore, empirical anomalies that emerged around the treatment of consensus information may be related to the extent to which the collection and use of the consensus information carries social risk in the context in which attributions are made. However, it is easy to imagine settings in which distinctiveness or consistency information carries similar social weight. Therefore, the notion that attributional information may be treated differently when it is socially consequential does not necessarily apply to consensus information alone.

This study cannot comment on whether attributional processes were influenced unconsciously by social features of the setting, or whether participants became aware that it would be socially tricky to assign blame to the experimenter and "held back" to be polite. We do know that participants who answered incorrectly reported less certainty in their answers, but there were no differences in reported certainty between those in the socially safe condition and those in the socially risky condition. In any case, the rating of certainty would not allow us to comment on whether reported attributions were influenced consciously or unconsciously because it is not possible to know whether participants consciously adjusting their answer would deflate their reported certainty to reduce the extent to which they appear to be "bad attributors," or inflate their reported certainty to avoid "damning with faint praise." However, it is relatively unimportant to know what people truly believe, since it is social action that is most important in understanding attributional processes in real-life settings (Edwards & Potter, 1993). For example, in a study of voting behavior, the important thing is to predict who people will vote for, not what they really believe about the candidate.

By arguing that people are scientists in miniature who value external validity above all else, Kelley's (1967) model fails to fully describe the great many contexts in which other values are important to people and play a powerful role in attributional processes, such as choosing "why" questions to ask and answer, choosing information by which attributions can be produced, and individually or collaboratively processing that information to arrive at attributions (Quayle & Durrheim, 2006, 2008). To take an extreme example, Kelley's ANOVA model of attribution would be of little use in describing the attributional processes employed in a witch hunt, in which veridical accuracy would be low on the list of social values by which attributions would be made and verified. Indeed, the context of science is one of the few settings in which adherence to veridical accuracy in spite of social consequences is not only accepted, but valued. This study adds weight to the argument that social features of attributional contexts (e.g., the consequences of seeking or using attributional information or arriving at attributions invoking accountability, blame, or even praise) are likely to be important features of how attributions are produced in real-life social settings (Edwards & Potter, 1993; Quayle & Durrheim, 2006, 2008).

A general implication of this argument is that it is likely that Kelley's (1967) model of attribution will most accurately describe real-life attributional processes in neutral settings in which it is either socially safe or even socially valued to make attributions without concern for social features of the attributional setting or the social value of the attribution itself. More specifically, the results of this study demonstrate that the application of Kelley's (1967) model of attribution may be improved by considering the social consequences of seeking or using attributional information, or arriving at particular attributions in the applied setting. Consensus information deserves special attention in this regard because it is more likely to be social in nature and to require social engagement to produce.

References

Anderson, I., Beattie, G., & Spencer, C. (2001). Can blaming victims of rape be logical? Attribution theory and discourse analytic perspectives. *Human Relations*, *54*, 445-467.

Arvai, J., Gregory, R., Ohlson, D., Blackwell, B., & Gray, R. (2006). Letdowns, wake-up calls, and constructed preferences: People's responses to fuel and wildfire risks, *Journal of Forestry*, *104*, 173-181.

Brown, R., & Fish, D. (1983). The psychological causality implicit in language. *Cognition*, *14*, 237-273.

Davis, D. W., & Silver, B. D. (2003). Stereotype threat and race of interviewer effects in a survey on political knowledge. *American Journal of Political Science*, *47*, 33-45.

Edwards, D., & Potter, J. (1993). Language and causation: A discursive action model of description and attribution. *Psychological Review*, *100*, 23-41.

Frieze, I. H., & Weiner, B. (1971). Cue utilization and attributional judgments for success and failure. *Journal of Personality*, *39*, 591-605.

Goffman, E. (1963). *Stigma*. Englewood Cliffs, NJ: Prentice Hall.

Hair, J. F., Anderson, R. E., Tatham, R. L., & Black, W. C. (1998). *Multivariate data analysis* (5th ed.). Englewood Cliffs, NJ: Prentice Hall.

Harvey, J. H., & Weary, G. (1984). Current issues in attribution theory and research. *Annual Review of Social Psychology*, *35*, 427-459.

Heider, F. (1958). *The psychology of interpersonal relations*. New York: John Wiley & Sons.

Hewstone, M. (1983). Attribution theory and common-sense explanations: An introductory overview. In M. Hewstone (Ed.), *Attribution theory: Social and functional extensions* (pp. 1-27). Oxford, UK: Basil Blackwell.

Hilton, D. J., Mathes, R. H., & Trabasso, T. R. (1992). The study of causal explanation in natural language: Analyzing reports of the Challenger disaster in the *New York Times*. In M. L. McLaughlin, M. J. Cody, & S. J. Read (Eds.), *Explaining one's self to others: Reason-giving in a social context* (pp. 41-60). Hillsdale, NJ: Lawrence Erlbaum.

Janakiraman, N., Meyer, R. J., & Morales, A. C. (2006). Spillover effects: How consumers respond to unexpected changes in price and quality. *Journal of Consumer Research*, *33*, 361-369.

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

Keaveney, S. M. (2008). The blame game: An attribution theory approach to marketer–engineer conflict in high-technology companies. *Industrial Marketing Management*, *37*, 653-663.

Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska Symposium on Motivation* ([M]pp. 192-240). Lincoln, NE: University of Nebraska Press.

Kelley, H. H. (1971). *Attribution in social interaction*. New York: General Learning.

Kelley, H. H., & Michela, J. L. (1980). Attribution theory and research. *Annual Review of Social Psychology*, *31*, 457-501.

Major, B. (1980). Information acquisition and attribution process. *Journal of Personality and Social Psychology*, *39*, 1010-1023.

Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, *50*, 370-396.

McArthur, L. Z. (1972). The how and what of why: Some determinants and consequences of causal attributions. *Journal of Personality and Social Psychology*, *22*, 171-193.

McArthur, L. Z. (1976). The lesser influence of consensus than distinctiveness information on causal attribution: A test of the person–thing hypothesis. *Journal of Personality and Social Psychology, 33*, 733-742.

Moscovici, S., & Hewstone, M. (1983). Social representations and social explanations: From the "naïve" to the "amateur" scientist. In M. Hewstone (Ed.), *Attribution theory: Social and functional extensions* (pp. 98-125). Oxford, UK: Basil Blackwell.

Nisbett, R. E., & Borgida, E. (1975). Attribution and the psychology of prediction. *Journal of Personality and Social Psychology*, *32*, 932-943.

Orvis, B. R., Cunningham, J. D., & Kelley, H. H. (1975). A closer examination of causal inference: The roles of consensus, distinctiveness, and consistency information. *Journal of Personality and Social Psychology*, *32*, 605-616.

Pilkonis, P. A. (1977). The influence of consensus information and prior expectations on dispositional and situational attributions. *Journal of Social Psychology, 101*, 267-279.

Quayle, M. (2005). *When the chips are down: Attribution in the context of computer failure and repair*. Unpublished master's thesis, University of KwaZulu-Natal, South Africa.

Quayle, M., & Durrheim, K. (2006). When the chips are down: Social and technical aspects of computer failure and repair. *Interacting with Computers*, *18*, 1260-1277.

Quayle, M., & Durrheim, K. (2008). Producing expertise and achieving attribution in the context of computer support. *British Journal of Social Psychology*, *47*, 727-762.

Ruble, T. L., & Feldman, N. S. (1976). Order of consensus, distinctiveness, and consistency information and causal attributions. *Journal of Personality and Social Psychology, 90*, 41-44.

Sanchez, J. I., Gomez, C., & Wated, G. (2008). A value-based framework for understanding managerial tolerance of bribery in Latin America. *Journal of Business Ethics*, *83*, 341-352.

Shapcott, K. M., Carron, A. V., Greenlees, I., & El Hakim, Y. (2008). Do member attributions for team outcomes represent a collective belief? *Psychology of Sport and Exercise*, *9*, 487-492.

Smith, F., Tabak, F., Showail, S., Parks, J., & Kleist, J. (2005). The name game: Employability evaluations of prototypical applicants with stereotypical feminine and masculine first names. *Sex Roles*, *52*, 63-82.

SPSS, Inc. (2007). *SPSS base 15.0 for Windows help files*. Chicago: Author.

Thatcher, J. B., Zimmer, J. C., Gundlach, M. J., & McKnight, D. H. (2008). Internal and external dimensions of computer self-efficacy: An empirical examination. *IEEE Transactions on Engineering Management*, *55*, 628-644.

Wells, G. L., & Harvey, J. H. (1977). Do people use consensus information in making causal attributions? *Journal of Personality and Social Psychology*, *35*, 279-293.

Zuckerman, M. (1978). Actions and occurrences in Kelley's Cube. *Journal of Personality and Social Psychology*, *36*, 647-656.

Zuckerman, M. (1979). Attribution of success and failure revisited, or: The motivational bias is alive and well in attribution theory. *Journal of Personality*, *47*, 245-287.

Table 1

Expected Attributions From Types of Information in Isolation: Covariation Patterns That Should Lead to Attributions to Either the Person, the Entity, or the Circumstance

	Information		
Attribution	Consensus	Distinctiveness	Consistency
Person	Low	Low	High
Entity	High	High	High
Circumstances	Low	High	Low

Table 2

Sequence of Video Clips Used to Create Each Covariation Matrix

Predicted attribution for failure	Consensus	Distinctiveness	Clip sequence			
Something to do with Jon	Low	Low	Success (Bob; Right)	Success (Bob; Left)	**Failure (Jon; Left)**	**Failure (Jon; Right)**
Something to do with Bob	Low	Low	Success (Jon; Right)	**Failure (Bob; Right)**	Success (Jon; Left)	**Failure (Bob; Left)**
Something to do with right ATM	High	High	Success (Bob; Left)	Success (Jon; Left)	**Failure (Jon; Right)**	**Failure (Bob; Right)**
Something to do with left ATM	High	High	Success (Bob; Right)	**Failure (Jon; Left)**	Success (Jon; Right)	**Failure (Bob; Left)**

Note. Failures appear in boldface.

Table 3[N]

Comparison of Model Fit Tests for the Best-Fitting Model Compared to the Saturated Model

	Omnibus tests			Model summary coefficients		
	χ^2	*df*	*p*	-2 log likelihood	Cox & Snell's r^2	Nagelkerke's r^2
Saturated model	7.86	3	.049	216.425	.035	.055
Optimal model[a]	7.55	1	.006	216.732	.033	.053

[a]The optimal model consists of social risk and the interaction between social risk and information type.

Table 4

Optimal Logistic Regression Model

Variable	*B*	*SE*	Wald	*df*	*p*	Odds ratio
Information Type × Role	0.966	.466	0.466	1	.038	2.629
Role	-1.019	.387	6.928	1	.008	0.361

Note. The reference category for the dependent variable is "1" and the model therefore predicts correct answers.

The variable "Information Type" was not included in the optimal model.

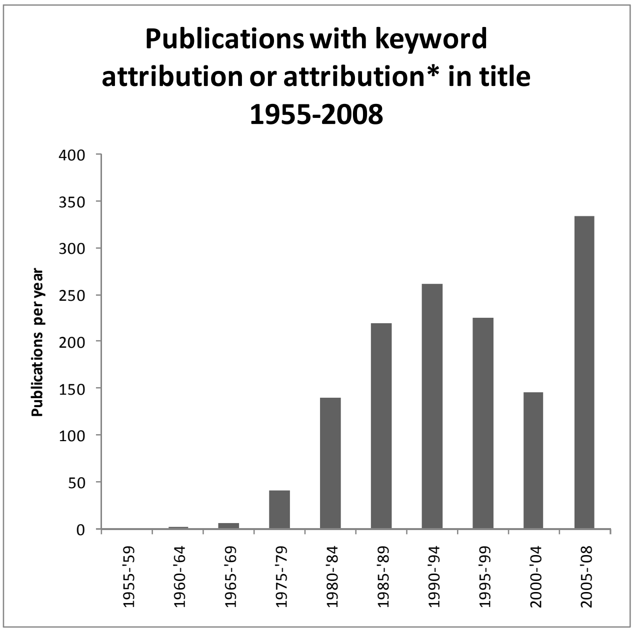

Figure 1. Publications per year with any title word starting with "attribution" in the PsycInfo database from 1955 through December 2008.

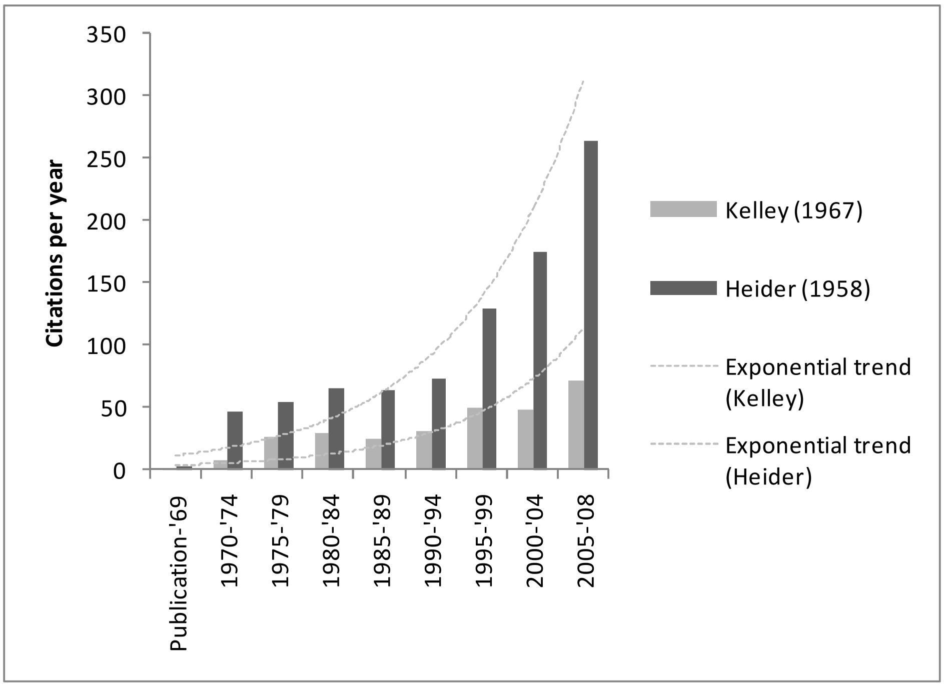

Figure 2. Citations per year of Heider (1958) and Kelley (1967) from publication up until 2008.

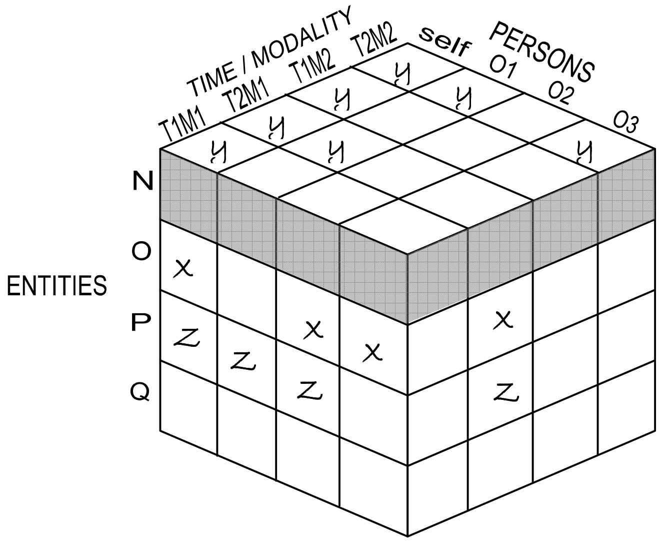

Figure 3. Kelley's ANOVA cube (adapted from Kelley 1967, p. 195).

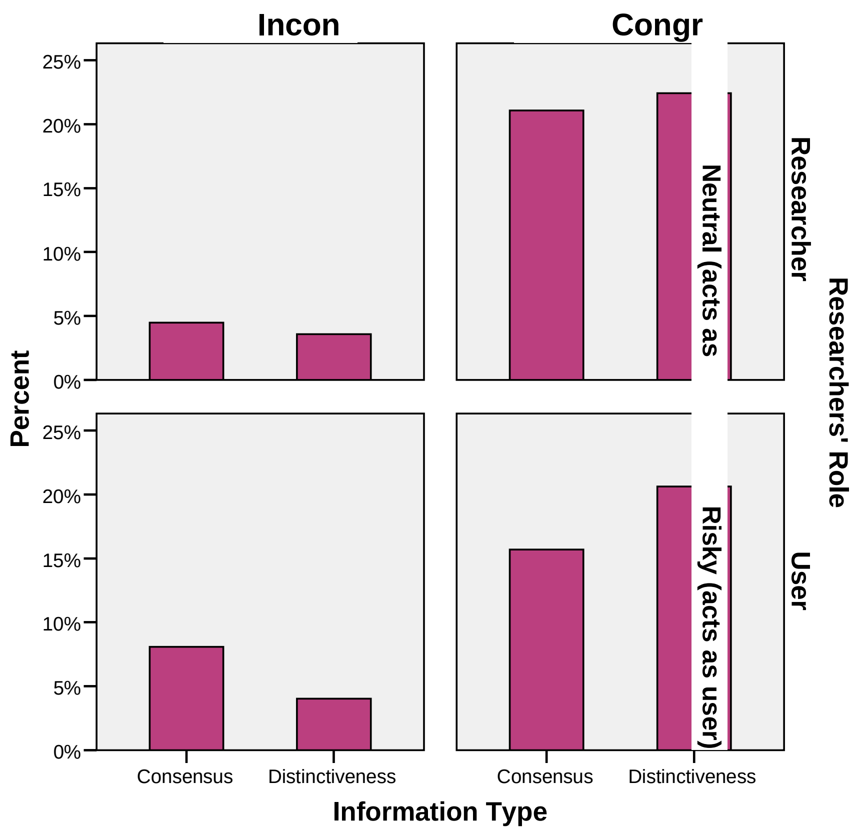

Figure 4. [O]Cross-tabulation of role, information type, and congruence of response.

[A] INSERT FIGURE 1 ABOUT HERE
[B] INSERT FIGURE 2 ABOUT HERE
[C] Page 4—Kelley (1967). Were the emphases (italics) in the original? Please indicate, one way or the other.
[D] INSERT FIGURE 3 ABOUT HERE
[E] INSERT TABLE 1 ABOUT HERE
[F] Page 7—Who do you mean by “they” here? Why are you quoting from Kelley & Michela (1980) when you are discussing a review by Harvey & Weary? Please clarify your meaning.
[G] Page 9—Footnote 5. Edwards (1993) is not in the References. Should this be Edwards & Potter (1993)? Please correct or provide the missing reference.
[H] INSERT TABLE 2 ABOUT HERE
[I] Page 12—What do you mean by “university stand”? Please clarify your meaning.
[J] INSERT FIGURE 4 ABOUT HERE
[K] Page 13—You cannot introduce Table 4 until you have presented Table 3. Please correct.
[L] INSERT TABLE 4 ABOUT HERE
[M] Page 17—Kelley (1967). Provide the volume number here.
[N] Page 21—Table 3 is not mentioned in your text at all (and you have provided no indication as to where it should be located). Please correct.
[O] Page 23—Figure 4. It appears that some work is needed on this figure. The “Neutral” and “Risky” part of the figure is not legible. Please provide a legible figure.