User Interface Design and the Halo Effect: Some Preliminary Evidence

Completed Research Paper

Daniel S. Soper
Department of Information Systems and Decision Sciences
Mihaylo College of Business and Economics
California State University, Fullerton
dsoper@fullerton.edu

Abstract

This paper draws on work from cognitive psychology and human-computer interaction to inquire into the extent to which assessments of web interface design are contaminated by the halo effect. The halo effect is a human cognitive bias in which a person's overall affect toward an entity produces overtly positive or negative distortions in evaluations of the entity's individual characteristics. These delusions and distortions in judgment can directly erode one's capacity to make rational decisions or perform objective evaluations, and therefore have critical implications for managers and organizations alike. Inquiry into the halo effect was carried out by means of a controlled, randomized experiment which incorporated 3 polarizing issues, 7 web interface design variations, and more than 300 research subjects. After controlling for other cognitive biases and several additional factors that are known to influence judgment and decision-making, the results revealed substantial halo-based contamination in assessments of web interface design.

Keywords

Halo effect, cognitive bias, user interface design, web design.

Introduction

Over the course of the past several decades, the information systems (IS) field has become increasingly concerned with phenomena that lie at the intersection of technology, human cognition, and behavior (e.g., Venkatesh and Morris 2000; Wells et al. 2011). Paralleling this increased interest in psychological and behavioral phenomena has been a rise in IS research oriented toward human interactions with technology, particularly in the areas of web design and usability (e.g., Agarwal and Venkatesh 2002; Deng and Poole 2010; Soper and Mitra 2013). This paper contributes to the ongoing integration of these streams of research by inquiring into the extent to which a cognitive bias known as the *halo effect* produces systematic errors in assessments of user interface design.

Most people are experientially familiar with the halo effect. Put simply, the halo effect (or *halo error*) is the tendency of human beings to judge an entity's individual traits according to their overall feelings or affect toward the entity (Bingham 1939; Thorndike 1920). If, for example, you hold a particular friend in high regard, then you would be likely to assign high ratings to her individual characteristics (such as intelligence, efficiency, cleanliness, reliability, etc.), despite the fact that she may not truly be deserving of such ratings. The halo effect has been shown not only to influence our perceptions of other people, but also the way in which we perceive organizations, government entities, and even certain types of consumer goods and technologies. What is not known, however, is the extent to which the halo effect colors our judgments of user interface design. Given the critical role of software and websites in modern societies, this is not an insignificant or trivial oversight. In an effort to reconcile this issue, the current study employs a controlled, randomized experiment involving 7 web interface design variations and more than 300 research subjects in order to assess the extent to which web interface design evaluations are distorted by halo-based contamination.

The Halo Effect

The halo effect is a human cognitive bias in which one's overall affect toward an entity produces overtly positive or negative distortions in evaluations of the entity's individual characteristics. As a cognitive bias, the halo effect is one of several known phenomena which lead to systematic errors in human judgment. Together, such errors cause human beings to behave very differently from what might otherwise be expected under a normative model of rational decision-making. Halo effect theory can be readily represented using a structural model, as shown in Figure 1 below.

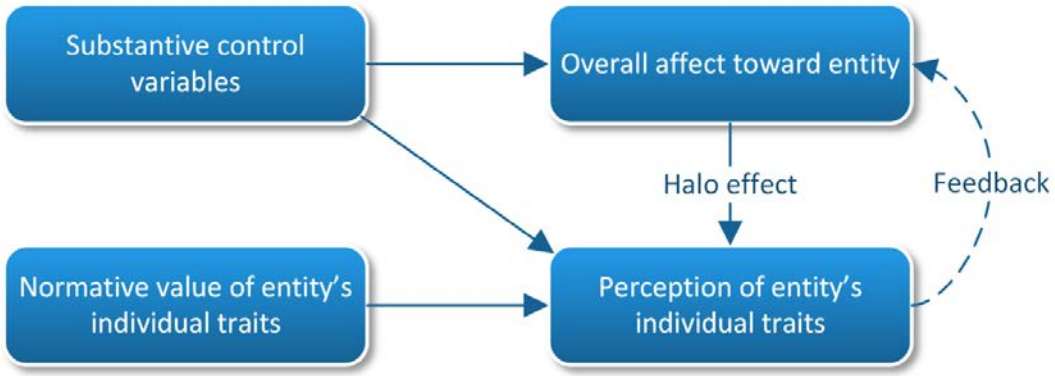

Figure 1. A representation of halo effect theory.

As shown in the figure, the theory posits that the way in which a human mind perceives an entity's individual traits depends not only upon the objective or normative value of those traits, but also upon the person's overall affect toward the entity. The extent to which a person's perceptions of an entity's individual traits are distorted is hence directly related to the nature and direction of the person's overall affect toward the entity being judged. A positive (negative) broad impression can hence be expected to produce positively (negatively) biased estimates of individual traits. Inter-judge random variation notwithstanding, the magnitude of the halo effect can thus be quantified as the extent to which a judge's rating of an individual trait is influenced by his or her overall affect toward the entity, after controlling for the normative value of that trait.

The Halo Effect in Human Beings

The halo effect was first described in the context of one person judging the traits and characteristics of another. Some of the most visible and widely disseminated findings in this stream of research have emerged from studies focusing on physical characteristics. In their highly cited paper, Dion et al. (1972) demonstrated that physically attractive men and women were presumed to have more socially desirable personalities, higher levels of occupational status, more marital competence, more social and professional happiness, and higher levels of total happiness than persons of average or below average attractiveness. Similarly, Landy and Sigall (1974) found that assessments of writing quality were distorted by the attractiveness of the author. When asked to evaluate the quality of an essay along several dimensions, judges who were shown a photo of an attractive woman and told that she was the author consistently rated the essay more favorably than judges who were shown a photo of an unattractive author. The ratings of judges who were not shown a photo fell between these two extremes. A later replication and extension of this work found that the halo effect described by Landy and Sigall was largely isolated to the condition in which male judges were evaluating an essay written by an author whom they believed to be female (Kaplan 1978). Kaplan's study thus revealed the gender of the rater and ratee to be a potentially important factor in research involving the halo effect. The halo effect has also been observed in other domains involving the evaluation of one person's traits by another. Early work by Remmers (1934), for example, found strong evidence of the halo effect in both high school and university students' evaluations of their instructors, with the effect being especially pronounced among university students. It has further been shown that the scores assigned by examiners for a person's performance on the most widely used IQ tests are strongly contaminated by the halo effect when the examiner is provided *a priori* with an overall

assessment of the person's level of intelligence (Sattler et al. 1970). With respect to its ability to distort our judgments of other people, the halo effect is clearly both a pervasive and pernicious phenomenon.

The Halo Effect and Technology

Beyond causing distortions in evaluations of individuals and organizations, an array of recent studies indicates that assessments of technology are also commonly contaminated by halo error. Kim (2009), for example, concluded that positive past experiences with a technology are often projected onto similar new technologies by way of a halo effect. This position was echoed in more specific terms by Massey and Montoya-Weiss (2006), who found that perceptions regarding the utility of knowledge conversion technologies can influence future perceptions of other knowledge conversion technologies. Perceptions also form the basis of Fred Davis' (1989) Technology Acceptance Model (TAM), and since all of the major constructs in TAM (i.e., ease of use, usefulness, and usage intentions) are self-reported, Szajna (1996) found that those constructs are subject to contamination by the halo effect. In an explicit example of self-assessments of usage intentions, Lee and Chen (2011) expressed similar concerns with respect to halo error in participants' self-assessments of past usage of and future intentions to use virtual worlds.

In addition to influencing assessments of specific technologies and systems, halo effects have also been suspected to produce distortions in evaluations of web design. Hartmann et al. (2007), for example, speculated that halo error may influence judgments of website usability and aesthetics, while Deng and Poole (2010) speculated that the initial emotional response evoked by the visual characteristics of a website may be carried by way of a halo effect to evaluations of other aspects of the website. While the studies above speculate that halo effect contamination may be present in assessments of web design, it is important to note that in each case such speculations were only secondary conjecture – heretofore the role of the halo effect in assessments of web design has never been explicitly or directly tested.

Hypotheses

One of the foundational concepts of halo effect theory is that our perceptions of an entity's individual characteristics or traits are derived largely -- but not entirely -- from objective reality. Put another way, the normative values of an entity's individual traits serve as the perceptual foundation upon which evaluations or judgments of those traits are based. The theory simply predicts that our final perceptions of individual traits will deviate from reality because of the heuristic nature of human cognitive processing. We therefore expect perceptions of web interface design characteristics to be linked to objective reality. Namely:

> **H1**: Perceptions of web interface design traits are positively related to the normative values of those traits.

The relationship between the normative values of an entity's individual traits and our perceptions of those traits notwithstanding, the primary prediction of halo effect theory holds that our perceptions of an entity's individual traits are influenced by our general feelings toward the entity, and we are hence unable to independently assess those traits without our ratings being colored by our general feelings or affect toward the entity. Importantly, this overall affect emerges from an aggregation of whatever information we have available about the entity, be it years of personal experience, or something as simple as a photograph. In light of halo effect theory, we thus expect perceptions of web interface design characteristics to be distorted by a person's overall affect toward the website being judged. Namely:

> **H2**: Perceptions of web interface design traits are positively related to the judge's overall affect toward the website whose traits are being judged.

The priming effect has been shown to exert a strong influence on affective states (Kolb and Whishaw 2008). Like the halo effect, the priming effect is a cognitive bias that can cause people to behave differently than what would otherwise be expected. In human cognition, the priming effect is related to implicit memory, and is characterized by a person being exposed to a stimulus which unconsciously influences her reaction to or feelings toward future stimuli (Kolb and Whishaw 2008; Tulving et al. 1982). Inquiring into a person's affective states loads those feelings into the person's working memory, thus drawing attention to her affective feelings. This act unconsciously alters the person's cognitive processes, and hence influences her future responses. Isolating and controlling for the influence of the priming effect

is therefore critical to obtaining an accurate measurement of any halo effects. Since the priming effect predicts that a person's self-reported affective feelings toward a web interface will vary according to whether she is asked about those feelings before or after judging the interface's individual traits, we hypothesize the presence of psychological priming so as to control for its anticipated effects in our measurement model. Thus:

> **H3**: A subject's self-reported overall affect toward a website varies according to whether she is asked about her affective feelings before or after judging the website's design traits.

Beyond psychological priming, we also sought to isolate and control for substantive inter-judge differences that could significantly influence a judge's affect or decision-making when evaluating a web interface. In this regard, age, gender, and level of experience have all been identified in the psychology, managerial, and IS literatures as factors which commonly influence affect, judgment, and decision-making (Charles et al. 2001; MacPherson et al. 2002; Riedl et al. 2010; Venkatesh et al. 2000). For this reason, age, gender, and a subject's level of web experience were included as control variables in the research model.

Methodology

Inquiry into the hypotheses developed above was carried out by means of a controlled, randomized experiment that was broadly structured around two large groups of subjects; namely, a baseline (control) group and a halo effect (treatment) group. In total, the experiment incorporated 3 polarizing issues, 7 web interface variations, and 330 research subjects. The following subsections provide specific details regarding the design and execution of the experiment.

Interface Design Characteristics

Following a detailed review of the literature, Aladwania and Palvia (2002) concluded that the appearance of a web interface is a unique construct with respect to assessments of website quality, and should be treated separately from other constructs such as content or technical adequacy. In deference to this conceptualization, a web interface template was created for the experiment which allowed the content of the web page (e.g., the title, text, etc.) to be modified while preserving the layout and appearance of the web interface itself. In this way, subjects could be asked to evaluate the design characteristics of a web interface while being subjected to experimentally manipulated variations in content. Given that subjects had never before seen the interface or its content, their overall affect toward the website would emerge solely from these informational cues. Holding the design of the interface constant thus ensured that differences in subjects' affect toward the website would be linked specifically to the experimentally manipulated variations in content. An example of the interface template being subjected to the content substitution process is illustrated in Figure 2 below.

Figure 2. Interface content substitution process.

Inasmuch as the halo effect is concerned with distorted evaluations of an entity's individual characteristics, it was necessary to identify a set of specific interface design characteristics that subjects could be asked to evaluate in order for the experiment to be properly aligned with halo effect theory. A pre-validated, five-item subscale designed to measure the attractiveness of a web interface design was

therefore adopted (Aladwania & Palvia, 2002). In accordance with the original instrument, subjects in the experiment were asked to respond to the evaluative statements in the subscale using a seven-point Likert-type scale anchored at *1 = strongly disagree* and *7 = strongly agree*. Minor modifications were made to the wording of the items in the subscale in order to adapt those items to the context of the current experiment (see Table 1).

Original Statement (Aladwania & Palvia, 2002)	Modified Statement Used in Current Experiment
____'s website looks attractive.	This website looks attractive.
____'s website looks organized.	This website looks organized.
____'s website uses fonts properly.	This website uses fonts properly.
____'s website uses colors properly.	This website uses colors properly.
____'s website uses multimedia features properly.	This website uses multimedia features properly.

Table 1. Original and modified subscale items.

Baseline Group

Guided by the foundational literature on halo effect theory, a baseline (control) group was incorporated into the experiment in order to establish normative ratings for the design characteristics of the web interface that was used in the study. Following Bingham (1939) and Johnson and Vidulich (1956), normative values were obtained by asking members of the baseline group to evaluate the web interface along only one dimension. Specifically, normative values were computed as the mean of baseline subjects' ratings for each dimension. Further, since any intelligible content displayed on the interface might serve as an informational cue which could trigger an affective association toward that content (and by extension, distortions in judgment) it was necessary for each web interface evaluated by subjects in the baseline group to be populated with neutral, unintelligible content. For purposes of the current study, the original, untranslated text from the ancient epic poem *Beowulf* (Anonymous, 8th-11th century) was thus used as the content basis for the web interface evaluated by subjects in the baseline group. Since this epic poem was written in Old English (Ænglisc), which has now been a dead language for more than 800 years, its original text is essentially unintelligible to modern readers, highly specialized linguists notwithstanding. An example of such unintelligible content can be seen in Figure 2 above. As a final contrivance aimed at obtaining normative ratings, subjects in the baseline group were specifically admonished to disregard the content displayed on the web interface, and instead focus their attention solely on its design characteristics.

Halo Effect Group

In addition to the baseline group, the experiment also incorporated a halo effect (treatment) group, the purpose of which was to generate responses that could be compared to those of the baseline group for evidence of halo effect contamination. Members of the halo effect group were subjected to conditions designed to allow any potential halo effects to emerge and be measured. In contrast to the baseline group, subjects in the halo effect group were asked to evaluate the interface along all five evaluative dimensions. Following past research (D. M. Johnson & Vidulich, 1956), this created a situation in which subjects were required to rate multiple individual characteristics of a single entity – in this case, a web interface. Whereas subjects in the baseline group were exposed to unintelligible content, the content appearing on the web interface evaluated by the halo effect group was intentionally polarizing in nature. Per halo effect theory, this experimental manipulation was intended to provide a strong psychoemotional anchor – unrelated to the task at hand – that could induce preexisting affective associations and subsequent distortions in judgment. To this end, three issues acknowledged by scholars to be socially polarizing in the United States were included in the current study; namely, abortion rights (Mouw & Sobel, 2001), illegal immigration (Newton, 2008), and gun control laws (Winkler, 2011).

For each of the three polarizing issues included in the experiment, two diametrically opposed variations in content were developed for use with the web interface evaluated by subjects in the halo effect group. The specific polarizing issue to which a halo effect subject was exposed was assigned iteratively, as was the

position – either supportive or oppositional – of the content seen by that subject. In order to determine whether and to what extent any halo effects were related to subjects' preexisting affective impressions, it was necessary to inquire into each subject's position with respect to the polarizing issue with which he or she was presented. This was accomplished by asking each subject to specify the extent to which they agreed with a statement about the polarizing issue to which he or she had been assigned. Subject responses to these statements were recorded using a seven-point Likert-type scale anchored at *1 = strongly disagree* and *7 = strongly agree*. With a view toward determining the extent to which priming effects may influence halo effect contamination (Kolb & Whishaw, 2008), half of the subjects in the halo effect group were asked to specify their feelings toward the polarizing issue prior to rating the web interface's design characteristics, while the remaining half were asked to specify their feelings after already having rated the interface.

Subjects

The halo effect group involved 3 polarizing issues and 2 possible positions for each issue (either supportive or not supportive), thus yielding a total of 6 possible interface configurations. For purposes of statistical validity a minimum of thirty subjects was required for each possible configuration, thus bringing the minimum sample size for the halo effect group to 180 subjects. The baseline group involved a single interface and five possible evaluative dimensions. For purposes of statistical validity each possible evaluative dimension required a minimum of thirty responses, thus bringing the total minimum number of subjects to 150, and the minimum sample size for the overall experiment to 330 subjects.

Considering that the target population was adult web users in the United States, the leading global online advertising firm was engaged to craft a targeted campaign for the purpose of soliciting volunteers for the research study. The firm's advanced technology allowed subject recruitment to be explicitly constrained to web users in the United States who were at least 18 years old. IP address restrictions were also enforced to help prevent the same subject from participating in the study more than once.

Experimental Execution

All of the abovementioned elements were aggregated into an interactive, web-based software system which was designed to carry out the experiment and store all of the data generated therefrom. After collecting demographic data (age, gender, and level of web experience), the system guided subjects through a training process which pointed out the salient features of the software program and instructed subjects as to how they could respond to the questions they would encounter during the experiment. Subjects were then iteratively assigned to either the baseline or halo effect group. In light of the structural model of halo effect theory shown in Figure 1, and in light of the hypotheses developed in the previous section, the research model depicted in Figure 3 below formed the principal foundation of inquiry for the current study.

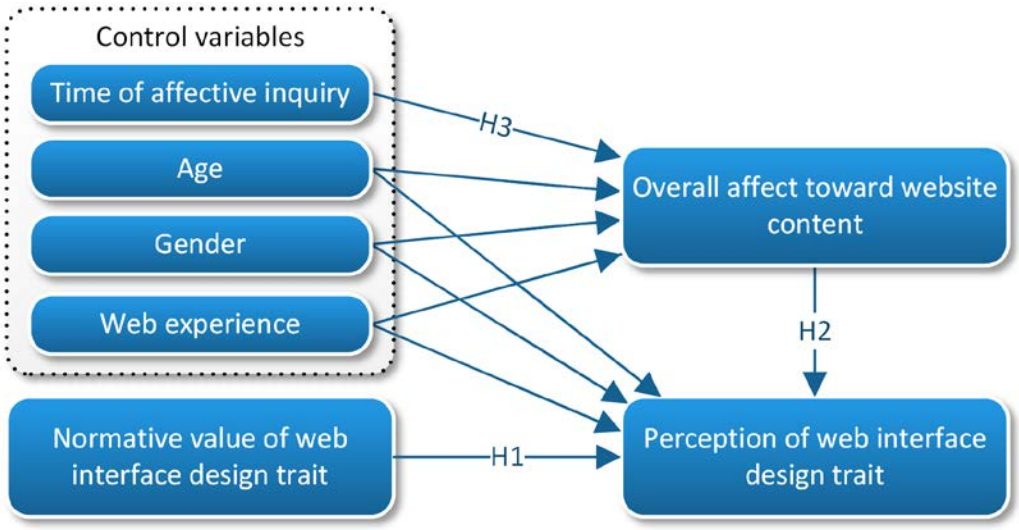

Figure 3. Research model.

Results

Initial estimation of the structural model revealed statistically significant relationships among all of the hypothesized inter-construct pads, excepting for those associated with web experience. Since a subject's degree of web experience was not observed to be related to her overall affect toward a website's content or her perceptions of the design of a website's interface, those paths were removed from the structural model, after which the model was re-estimated. The final parameter estimates for the structural model are shown in Figure 4 below.

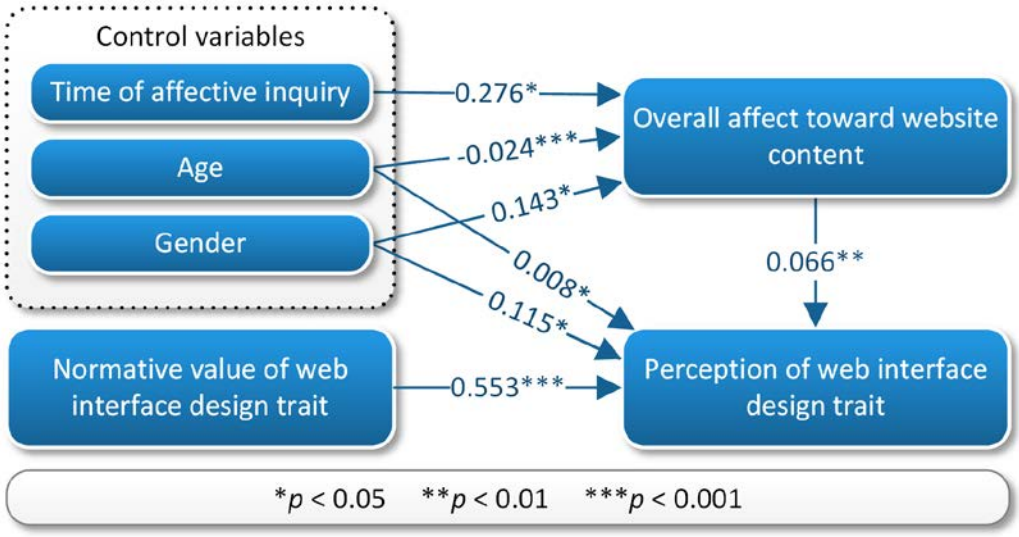

Figure 4. Parameter estimates and significances.

In light of the structural equation modeling methodology that was used to obtain the parameter estimates shown in Figure 4, it was also necessary to evaluate the extent to which the research model fit the observed data. To this end, Table 2 below provides the values of several of the commonest measures of fit as they apply to the research model described above.

Measure of Model Fit	Value
Chi-square (*p*):	1.759 (0.881)
Comparative fit index (CFI):	0.999
Normed fit index (NFI):	0.972
Root mean square error of approximation (RMSEA):	0.006
Probability of close fit (PCLOSE):	0.999

Table 2. Fit statistics for the research model.

Each of the measures of fit shown in the table above implies a very good to excellent fit between the research model and the observed data, and we therefore conclude that the research model was properly specified.

Tests of Hypotheses

Our first hypothesis tested the proposition that perceptions of web interface design traits are positively related to the normative values of those traits. These normative values were thus hypothesized to serve as the perceptual foundation upon which evaluations or judgments of such traits would be based. Results obtained from the estimation of the structural model revealed a strong positive relationship between the normative values of web interface design traits and subjects' perceptions of those traits ($p < 0.001$). Hypothesis 1 was therefore supported.

Our second hypothesis addressed the primary prediction of halo effect theory, which holds that human beings find it cognitively difficult to separate our overall affect toward an entity from our evaluations of the entity's individual traits, with our affect being derived from an aggregation of whatever information we have available about the entity. Perceptions of web interface design traits were therefore hypothesized to be distorted by a subject's overall affect toward the content of the website whose traits were being judged. The results obtained from estimating the structural model revealed a significant positive relationship between a subject's overall affect toward the content of a website and her perceptions of the website's interface design traits ($p < 0.01$). Hypothesis 2 was therefore supported.

The third hypothesis sought to assess and control for the role of psychological priming, which has been shown to exert a strong influence on a person's affective state. A person's self-reported affective feelings toward the content of a website were therefore hypothesized to vary according to whether she was asked about those feelings before or after judging the individual characteristics of the website's interface. Results obtained from the estimation of the structural model revealed a strong positive relationship between the time of affective inquiry and a subject's self-reported overall affect toward the content of a website ($p < 0.05$). Hypothesis 3 was therefore supported.

Discussion, Limitations, and Concluding Remarks

The study described herein inquired into the extent to which evaluations of web interface design are distorted by halo effect-based contamination. For this purpose a controlled, randomized experiment was used which incorporated 3 polarizing issues, 7 web interface design variations, and more than 300 research subjects. After controlling for other cognitive biases and several additional factors that are known to influence judgment and decision-making, the quantitative results revealed substantial halo effect contamination in assessments of web interface design.

The presence of halo effect contamination in evaluations of interface design – and potentially many additional facets of human computer interaction – has manifold implications for both researchers and managers alike. With respect to the former, researchers must not expect evaluations of interface designs to be accurate if those evaluations are made using traditional, multi-item measurement scales. More broadly, and much more profoundly, a century of research clearly warns that the results of *all* studies which employ multi-item measurement scales but do not control for halo error should be looked upon with a healthy dose of skepticism. Although several disciplines have begun to heed this clarion call, due to the lack of familiarity with the halo effect and its attendant implications, the IS field has not yet responded

to this problem. The halo effect is particularly troubling for IS research since so many of our most hallowed and well-established theoretical constructs (such as *perceived usefulness* and *perceived ease-of-use*) are almost always measured using multi-item scales. Adopting methods similar to those used to establish the normative ratings for the current study should therefore be considered in the context of multi-item measurement.

The results of the current study also have important implications for managers. In the modern era, a website now commonly serves as the public face of an organization. To this end, the design of an organization's website has critical implications for the way in which customers and other interested parties will perceive the organization, with such perceptions influencing the organization's prospects for success. Since the design, layout, and appearance of a website can directly influence the way in which people think about the underlying organization, managers should be highly interested in ensuring that the design of their organization's website is well-aligned with the expectations of its users. In many cases, however, proposed designs are evaluated by parties that have a vested interest in the website or in the underlying organization itself. The results reported herein suggest that the overall affect of these parties may unconsciously contaminate their assessments of the website's design, thereby leading to erroneous conclusions. In order to avoid such problems, managers should strongly consider adopting mechanisms that have been shown to minimize halo error when soliciting assessments of proposed website designs.

There are several limitations to this work which merit acknowledgement. First, the findings were derived only from the responses of adult web users. Second, the research subjects were all English-speaking web users. Although halo error has been identified in many different countries and cultures, variations in web design are evident according to whether a website's primary language is written from left-to-right, right-to-left, top-to-bottom, etc. (Soper & Mitra, 2013). Third, we relied upon the findings of past research to create experimental conditions that were designed to maximize the halo effect. The extent to which the magnitude of the observed distortions in judgment would vary under conditions that were not intentionally designed to maximize the halo effect remains unknown. Finally, the web-based interfaces used in the experiment represent only a subset of the immense and ever-changing variety of technology interfaces with which human beings interact on a daily basis. Whether and to what extent the findings reported here are generalizable to application software interfaces, mobile interfaces, game interfaces, etc. remains unknown. Together, these limitations represent rich and fertile ground for future research.

The limitations noted above notwithstanding, the results of the study reveal that the halo effect plays a highly significant role in assessments of interface design. This finding extends the boundaries of halo effect theory into the realm of human-computer interaction, and establishes halo effect contamination as a legitimate factor for HCI research concerned with human judgment and decision-making. In light of the substantial halo-based errors in judgment detected here, it is suggested that the results of past HCI-related research studies which employed multi-item measurement scales may need to be reconsidered.

References

Ackoff, R. L. 1961. "Management Misinformation Systems," *Management Science* (14:4), pp. 147-156.

Benbasat, I., and Zmud, R. W. 2003. "The Identity Crisis within the IS Discipline: Defining and Communicating the Discipline's Core Properties," *MIS Quarterly* (27:2), pp. 183-194.

Bonini, C. P. 1963. *Simulation of Information and Decision Systems in the Firm*, Englewood Cliffs, NJ: Prentice-Hall.

Broadbent, M., Weill, P., O'Brien, T., and Neo, B. S. 1996. "Firm Context and Patterns of IT Infrastructure Capability," in *Proceedings of the 14th International Conference on Information Systems*, J. I. DeGross, S. Jarvenpaa, and A. Srinivasan (eds.), Cleveland, OH, pp. 174-194.

Carroll, J. 2005. "The Blacksburgh Electronic Village: A Study in Community Computing," in *Digitial Cities III: Information Technologies for Social Capital*, P. van den Besselaar and S. Kiozumi (eds.), New York: Springer-Verlag, pp. 43-65.

Agarwal, R., and Venkatesh, V. 2002. "Assessing a Firm's Web Presence: A Heuristic Evaluation Procedure for the Measurement of Usability," *Information Systems Research* (13:2), pp. 168-186.

Bingham, W.V. 1939. "Halo, Invalid and Valid," *Journal of Applied Psychology* (23:2), pp. 221-228.

Charles, S.T., Reynolds, C.A., and Gatz, M. 2001. "Age-Related Differences and Change in Positive and Negative Affect over 23 Years," *Journal of Personality and Social Psychology* (80:1), pp. 136-151.

Davis, F. 1989. "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly* (13:3), pp. 319-340.
Decety, J., Yang, C.-Y., and Cheng, Y. 2010. "Physicians Down-Regulate Their Pain Empathy Response: An Event-Related Brain Potential Study," *Neuroimage* (50:4), pp. 1676-1682.
Deng, L., and Poole, M.S. 2010. "Affect in Web Interfaces: A Study of the Impacts of Web Page Visual Complexity and Order," *MIS Quarterly* (34:4), pp. 711-730.
Dion, K., Berscheid, E., and Walster, E. 1972. "What Is Beautiful Is Good," *Journal of Personality and Social Psychology* (24:3), pp. 285-290.
Fujita, F., Diener, E., and Sandvik, E. 1991. "Gender Differences in Negative Affect and Well-Being: The Case for Emotional Intensity. ," *Journal of Personality and Social Psychology* (61:3), pp. 427-434.
Hartmann, J., Sutcliffe, A., and Angeli, A.D. 2007. "Investigating Attractiveness in Web User Interfaces," *ACM SIGCHI Conference on Human Factors in Computing Systems*, San Jose, CA: ACM, pp. 387-396.
Hooks, G., and Mosher, C. 2005. "Outrages against Personal Dignity: Rationalizing Abuse and Torture in the War on Terror," *Social Forces* (83:4), pp. 1627-1645.
Kaplan, R.M. 1978. "Is Beauty Talent? Sex Interaction in the Attractiveness Halo Effect," *Sex Roles* (4:2), pp. 195-204.
Keen, S. 1991. *Faces of the Enemy: Reflections of the Hostile Imagination*. Harper & Row: San Francisco, CA.
Kim, S.S. 2009. "The Integrative Framework of Technology Use: An Extension and Test," *MIS Quarterly* (33:3), pp. 513-538.
Kolb, B., and Whishaw, I.Q. 2008. *Fundamentals of Human Neuropsychology*. New York, NY: Worth Publishers.
Landy, D., and Sigall, H. 1974. "Beauty Is Talent: Task Evaluation as a Function of the Performer's Physical Attractiveness," *Journal of Personality and Social Psychology* (29:3), pp. 299-304.
Lee, Y., and Chen, A.N.K. 2011. "Usability Design and Psychological Ownership of a Virtual World," *Journal of Management Information Systems* (28:3), pp. 269-307.
Levenson, R.W., Carstensen, L.L., and Gottman, J.M. 1994. "Influence of Age and Gender on Affect, Physiology, and Their Interrelations," *Journal of Personality and Social Psychology* (67:1), pp. 56-68.
Lucker, G.W., Beane, W.E., and Helmreich, R.L. 1981. "The Strength of the Halo Effect in Physical Attractiveness Research," *The Journal of Psychology: Interdisciplinary and Applied* (107:1), pp. 69-75.
MacPherson, S.E., Phillips, L.H., and Della Sala, S. 2002. "Age, Executive Function and Social Decision Making: A Dorsolateral Prefrontal Theory of Cognitive Aging," *Psychology and Aging* (17:4), pp. 598-609.
Massey, A.P., and Montoya-Weiss, M.M. 2006. "Unraveling the Temporal Fabric of Knowledge Conversion: A Model of Media Selection and Use," *MIS Quarterly* (30:1), pp. 99-114.
Morse, D.S., Edwardsen, E.A., and Gordon, H.S. 2008. "Missed Opportunities for Interval Empathy in Lung Cancer Communication," *Journal of the American Medical Association: Internal Medicine* (168:17), pp. 1853-1858.
Powell, M., and Ansic, D. 1997. "Gender Differences in Risk Behaviour in Financial Decision-Making," *Journal of Economic Psychology* (18:6), pp. 605-628.
Remmers, H.H. 1934. "Reliability and Halo Effect of High School and College Students' Judgments of Their Teachers," *Journal of Applied Psychology* (18:5), pp. 619-630.
Riedl, R., Hubert, M., and Kenning, P. 2010. "Are There Neural Gender Differences in Online Trust? An Fmri Study on the Perceived Trustworthiness of Ebay Offers," *MIS Quarterly* (34:2), pp. 397-428.
Salas, E., Rosen, M.A., and DiazGranados, D. 2010. "Expertise-Based Intuition and Decision Making in Organizations," *Journal of Management* (36:4), pp. 941-973.
Sattler, J.M., Hillix, W.A., and Neher, L.A. 1970. "Halo Effect in Examiner Scoring of Intelligence Test Responses," *Journal of Consulting and Clinical Psychology* (34:2), pp. 172-176.
Soper, D.S., and Mitra, S. 2013. "An Inquiry into Mental Models of Web Interface Design," *19th Americas Conference on Information Systems (AMCIS)*, Chicago, Illinois.
Szajna, B. 1996. "Empirical Evaluation of the Revised Technology Acceptance Model," *Management Science* (42:1), pp. 85-92.
Taylor, R.N. 1975. "Age and Experience as Determinants of Managerial Information Processing and Decision Making Performance," *Academy of Management Journal* (18:1), pp. 74-81.
Thorndike, E.L. 1920. "A Constant Error in Psychological Ratings," *Journal of Applied Psychology* (4:1), pp. 25-29.
Tulving, E., Schacter, D.L., and Stark, H.A. 1982. "Priming Effects in Word Fragment Completion Are Independent of Recognition Memory," *Journal of Experimental Psychology* (8:4), pp. 336-342.
Venkatesh, V., and Morris, M.G. 2000. "Why Don't Men Ever Stop to Ask for Directions? Gender, Social Influence, and Their Role in Technology Acceptance and Usage Behavior," *MIS Quarterly* (24:1), pp. 115-139.

Venkatesh, V., Morris, M.G., and Ackerman, P.L. 2000. "A Longitudinal Field Investigation of Gender Differences in Individual Technology Adoption Decision-Making Processes," *Organizational Behavior and Human Decision Processes* (83:1), pp. 33-60.

Wark, G.R., and Krebs, D.L. 1996. "Gender and Dilemma Differences in Real-Life Moral Judgment," *Developmental Psychology* (32:2), pp. 220-230.

Wells, J.D., Valacich, J.S., and Hess, T.J. 2011. "What Signals Are You Sending? How Website Quality Influences Perceptions of Product Quality and Purchase Intentions," *MIS Quarterly* (35:2), pp. 373-396.

Wu, B.T.W., and Petroshius, S.M. 1987. "The Halo Effect in Store Image Measurement," *Journal of the Academy of Marketing Science* (15:3), pp. 44-51.