Decision Support Systems 29 (2000) 195–206

Decision Support Systems

www.elsevier.com/locate/dsw

# Countering the anchoring and adjustment bias with decision support systems

Joey F. George [a,*], Kevin Duffy [b], Manju Ahuja [a,1]

[a] College of Business, Florida State University, Tallahassee, FL 32306, USA
[b] College of Business and Economics, West Virginia University, Morgantown, WV 26506, USA

Accepted 5 April 2000

**Abstract**

Psychologists have identified several limitations to, and biases in, human decision-making processes. One such bias is the anchoring and adjustment effect, which has been demonstrated to be robust both inside and outside the experimental laboratory. Some decision support systems (DSS) have been designed to lessen the effects of decision-making limitations with promising results. This study tested a DSS designed to mitigate the effects of the anchoring and adjustment bias. The results show that anchoring and adjustment remains robust within the context of automated decision support. Implications that follow these results are offered. 

*Keywords:* Bias; Heuristics; Anchoring and adjustment; Decision support systems

## 1. Introduction

Decision support systems (DSS) have long been developed to supplement limited human information processing capabilities [2,12]. For example, several systems have been developed that support specific decision-making strategies, such as multi-criteria decision-making. Human decision-making, however, has been found to suffer from limitations other than limited information processing capabilities. Psychologists, such as Tversky and Kahneman [14], have worked during the past several decades to uncover systematic biases in human decision-making processes. Many of these biases, such as framing, representativeness, and availability, have become well-known in the literature (e.g., Ref. [6]). Yet few DSS seem to have been developed with the specific consideration to counter these well-known decision-making biases. In cases where such DSS have been built or considered [9,13], decision-making behavior has been successfully altered through interaction with the system.

One particularly interesting decision-making bias is anchoring and adjustment, described as follows [14, p. 1128]:

> In many situations, people make estimates by starting from an initial value that is adjusted to yield the final answer. The initial value, or start-

* Corresponding author. Tel.: +1-850-644-7449.

*E-mail addresses:* jgeorge@garnet.acns.fsu.edu (J.F. George), duffy@be.wvu.edu (K. Duffy), mahuja@garnet.acns.fsu.edu (M. Ahuja).

1 Tel.: +1-850-644-0916.

0167-9236/00/$ - see front matter 
PII: S0167-9236(00)00074-9

ing point, may be suggested by the formulation of the problem, or it may be the result of a particular computation. In either case, adjustments are typically insufficient. That is, different starting points yield different estimates, which are biased toward the initial values.

To demonstrate the anchoring and adjustment effect, Tversky and Kahneman asked subjects to estimate the percentage of United Nations member states that were African nations. The researchers spun a wheel of fortune, with numbers ranging from 0 to 100, in the presence of the subjects. Subjects were first asked to determine if the random value from the wheel was too high or too low of an estimate, and then to actually estimate the percentage of African membership at the U.N. The random, arbitrary numbers had an effect on subjects' responses. For example, those who received 10 as the value guessed 25% African membership, while those receiving 65 guessed 45%.

In a DSS, individuals typically have control over the specific amount and speed of information they desire to view. Further, DSS allow individuals to instantly obtain available information as needed to evaluate and reevaluate their decisions. Given the absence of time constraints, and given that a DSS user has the ability to carefully examine data and model relationships, it seems intuitive that the potential effects of an anchor should be diminished by the scrutiny that can be applied to data while using a DSS. Thus, we suggest that the effect of anchoring and adjustment could be mitigated through the use of a DSS. The purpose of the study described in this paper was to test this possibility.

We built a real estate appraisal DSS and asked subjects to use the available information to estimate the value of a parcel of real estate. Some subjects were presented with anchors, in the form of the owner's asking price, and some were not. In an attempt to lessen the effects of the anchoring and adjustment bias, some subjects received warnings when their estimated values were within a certain range of the anchor, while others did not.

In the next section, we review the literature on mitigating the effects of decision-making biases and limitations, and in particular, on the use of DSS to counter these effects. We also review some of the literature that illustrates the robustness of the anchoring and adjustment bias in real world problem-solving. We present two hypotheses derived from the literature and describe the laboratory study designed to test these hypotheses. Next, we present the details of the research design followed in this experimental study. The results, described and discussed in the final sections of the paper, demonstrate the robustness of the anchoring and adjustment effect.

## 2. Literature review

Psychologists and others have been interested in determining how to mitigate or eliminate the effects of decision-making biases and heuristics for almost as long as such biases have been reported [1,5,10]. Fischhoff called the efforts to diminish the effects of biases ''debiasing'' [5]. Debiasing efforts can include warnings, feedback, and training. Where the judge, and not the task, is considered to be faulty, Fischhoff describes four levels of debiasing activities: (1) warnings about the possibility of bias; (2) descriptions of the direction of bias; (3) feedback about the subject's behavior, which personalizes the implications of the warning; and (4) an extended program of feedback, coaching, and whatever else it takes to allow the subject to achieve mastery of the task [5, p. 426].

Fischhoff indicates that training and feedback can be a fruitful activity, saying that a ''variety of training efforts have been undertaken with an admirable success rate'' [5, p. 437]. For Alpert and Raiffa [1], feedback improved the performance of subjects making decisions under uncertainty, although it did not completely eliminate the subjects' overconfidence. Sharp et al. [10] also reported an improvement in subject performance after subjects were given feedback. In a more recent study, Connolly and Dean [4] found that, while training efforts appeared to have little impact on overtightness of expected outcome distributions, explicit attention to establishing good upper and lower bounds as part of the problem-solving process did reduce the overtightness effect. Overtightness is defined as ''too small a range between the times judged 'improbably long' and 'improbably short' — that is, an overtight estimated distribution'' [4, p. 1031].

Just as decision-making models and database access are easily built into DSS, various debiasing strategies described by Fischhoff, intended to counter biases and limitations, can easily be made part of a DSS. Although we are not aware of any research investigating the use of DSS to address the anchoring and adjustment bias, DSS have been designed specifically to address other decision-making biases and heuristics. Todd and Benbasat [13] designed a DSS that decreased the cognitive effort necessary for decision-makers to use a particular decision-making strategy, elimination-by-aspects (EBA), when solving a problem. The conjunctive and EBA strategies are two of four prototypical strategies used for dealing with multi-attribute, multi-alternative preferential choice problems. The conjunctive strategy involves evaluating each attribute against a minimum threshold level. Information is evaluated by alternative. If an attribute of an alternative does not meet the minimum threshold level, the alternative is dropped from further consideration. The EBA strategy is similar, but whereas the conjunctive strategy involves examining each attribute for each alternative in sequence, EBA requires examining the values for a given attribute across all alternatives at once. As with the conjunctive strategy, alternatives that do not meet the threshold level for that attribute are eliminated from further consideration (see Ref. [13] for more details).

From a cognitive effort perspective, decision-makers given a choice will typically choose to use a conjunctive strategy because it takes less cognitive effort than the rival EBA strategy. Subjects in Ref. [13] who used the DSS chose the EBA strategy, as opposed to the conjunctive strategy, because the DSS made the EBA strategy less effortful than the alternative. As would be expected, subjects who did not have the DSS available chose the conjunctive strategy.

Another decision-making limitation identified by Tversky and Kahneman is called neglect of base rates, or the base-rate fallacy. Decision-makers succumbing to the base-rate fallacy ignore prior probabilities associated with the belief in a hypothesis, even though those probabilities remain relevant in the face of new information. This fallacy was investigated from a decision support perspective by Roy and Lerch [9]. They found that the base-rate fallacy could be reduced if the written problem to be solved was also accompanied by a graphic probability mapping of the same problem. The base-rate fallacy bias was even further reduced if the written version of the problem that accompanied the graph did not contain all of the information pertinent to the problem. Although the authors did not build a DSS that combined probability problems and graphs, their work demonstrates that a DSS that did could be useful for such problems or for related problems where causality or temporal order is important.

Block and Harper [3], in an attempt to debias the effects of anchoring and adjustment, found that warnings about its influence acted to reduce the effect but could not eliminate it completely. They found that the anchoring and adjustment effect remained strong, even when subjects were warned by the researchers before they began work that they could be too overconfident in the accuracy of their estimates. Warnings somewhat decreased the effect but did not completely eliminate it.

Although they did not investigate what factors might mitigate the effects of anchoring and adjustment, Northcraft and Neale [7] did demonstrate the robustness of the effect outside the laboratory. The study involved the estimation of the value of a house for sale. In the study, subjects were taken to a house and were permitted to examine the property. They were all given identical information about it, except for the listing price, which was varied across subjects. The listing price was 4% above or below the house's actual appraised value for some subjects, and 12% above or below for others. The information they all received included the following:

- standard Multiple Listing Service (MLS) listing sheet for the property;
- MLS summary data for the city and the immediate neighborhood for the preceding 6 months;
- comparative information about other property in the same neighborhood;
- standard MLS information for other properties for sale in the same neighborhood.

Subjects were asked to supply four values: (1) an appraised value for the house; (2) a fair advertising price for the house; (3) a listing price; and (4) an amount reflecting the lowest offer they would accept as seller of the house. Northcraft and Neale [7] found that the listing price significantly biased the estimates given by the subjects, whether the price was 4% or 12% above or below the house's appraised value, and whether or not the subjects were students or experienced real estate agents.

The robustness of the anchoring and adjustment bias has been further demonstrated by several researchers in recent years. Ritov [8] examined anchoring effects in negotiation and found no indication of a decline in the bias as negotiators gained experience. Whyte and Sebenius [15] examined the anchoring and adjustment effect in situations where multiple anchors of varying degrees of relevance to the estimate required were available. They demonstrated that anchoring had a powerful effect, even though the anchor was unrelated to the estimation task to be performed, and other relevant and appropriate anchors were provided. They also found that groups were susceptible to the effects of an arbitrary anchor and were as incapable of disregarding such information as individuals.

## 3. Research hypotheses

Given the demonstrated robustness of the anchoring and adjustment effect, even in the face of warnings that the bias may be operating, we would expect anchoring and adjustment to be a factor in generating estimated values in most relevant situations. The question is whether this bias would continue to operate in the context of a computer-based information system, where subjects would have the ability to call for information on demand and examine it as much or as little as they desired. Given what we know about the robustness of anchoring and adjustment, we would still expect it to operate in the context of a DSS. Based on the success of various debiasing strategies as outlined above, we would expect warnings about the influence of an anchor to reduce the effect of the bias, although we would not expect the bias to be completely eliminated [3]. This leads to the following two hypotheses.

**Hypothesis 1.** Estimated values provided by subjects who are exposed to anchors will differ significantly from those of subjects who are not exposed to anchors.

**Hypothesis 2.** The introduction of warnings, based on the proximity of the estimated value to the anchor, will significantly reduce the effect of the anchor in determining subsequent values while using a DSS, but it will not completely eliminate the anchoring and adjustment effect.

## 4. Research method

### *4.1. The appraisal system*

To test the effects of DSS use on the anchoring and adjustment bias, we designed and built a DSS for house appraisals. The system was developed using Visual Basic 4.0. Information about a house for sale was obtained through the assistance of a local realtor and entered into the system. This information is similar to that used in Ref. [7]. Instead of taking subjects to the house in question, however, several digital photographs of the house were taken and provided as part of the information available to subjects via the DSS.

The final version of the program consisted of 23 screens, plus the possibility of a ''warning'' screen at four different points in the process of completing the task. The screens were assembled in a logical sequential flow, starting with an introductory screen, followed by a ''pricing hints'' screen (Fig. 1). Next, the program requested minimal personal and background information from each subject (e.g., age, gender, length of time living in the area, and whether or not the subject had purchased a house locally or elsewhere). An additional question asked whether the subject was currently, or had been, a realtor. Both subjects who answered yes to this question were directed to a screen to capture additional information about their experience as realtors.

Following the gathering of background information, the subjects were shown a ''main information'' screen for the house. This screen consisted of a photo of the exterior of the house and control buttons. The buttons could be clicked to provide screens showing additional photographs of the house (a total

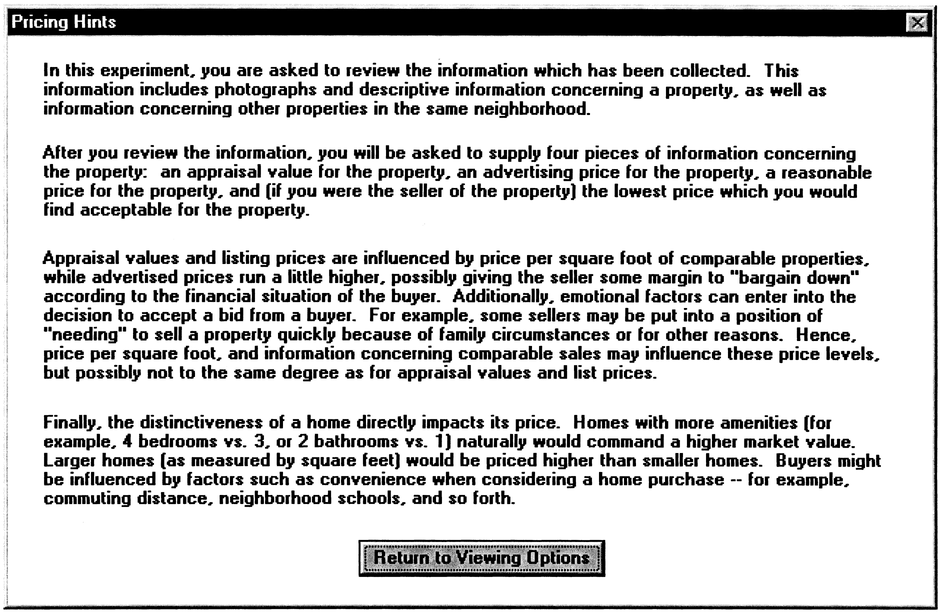

Pricing Hints

In this experiment, you are asked to review the information which has been collected. This information includes photographs and descriptive information concerning a property, as well as information concerning other properties in the same neighborhood.

After you review the information, you will be asked to supply four pieces of information concerning the property: an appraisal value for the property, an advertising price for the property, a reasonable price for the property, and (if you were the seller of the property) the lowest price which you would find acceptable for the property.

Appraisal values and listing prices are influenced by price per square foot of comparable properties, while advertised prices run a little higher, possibly giving the seller some margin to "bargain down" according to the financial situation of the buyer. Additionally, emotional factors can enter into the decision to accept a bid from a buyer. For example, some sellers may be put into a position of "needing" to sell a property quickly because of family circumstances or for other reasons. Hence, price per square foot, and information concerning comparable sales may influence these price levels, but possibly not to the same degree as for appraisal values and list prices.

Finally, the distinctiveness of a home directly impacts its price. Homes with more amenities (for example, 4 bedrooms vs. 3, or 2 bathrooms vs. 1) naturally would command a higher market value. Larger homes (as measured by square feet) would be priced higher than smaller homes. Buyers might be influenced by factors such as convenience when considering a home purchase -- for example, commuting distance, neighborhood schools, and so forth.

Return to Viewing Options

Fig. 1. Pricing hints screen.

of eight interior and exterior shots of the house were provided), as well as text information concerning the number of bedrooms, square footage of the house and properties, and amenities within the house.

Other options available from the main information screen included: (1) a pricing hints screen; (2) information (in the form of two tables) on other houses currently for sale; (3) information concerning houses

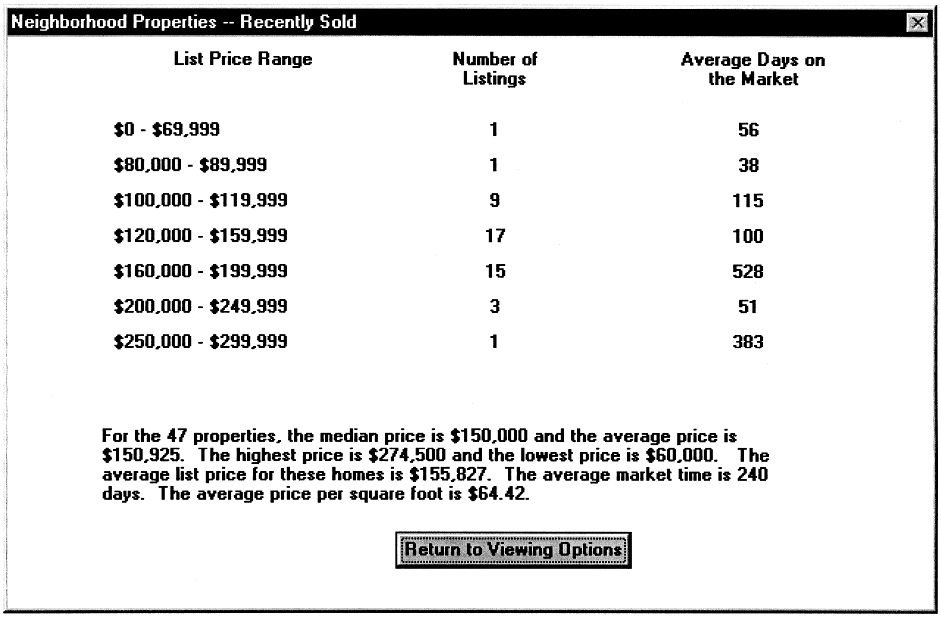

Neighborhood Properties -- Recently Sold

| List Price Range | Number of Listings | Average Days on the Market |
|---|---|---|
| $0 - $69,999 | 1 | 56 |
| $80,000 - $89,999 | 1 | 38 |
| $100,000 - $119,999 | 9 | 115 |
| $120,000 - $159,999 | 17 | 100 |
| $160,000 - $199,999 | 15 | 528 |
| $200,000 - $249,999 | 3 | 51 |
| $250,000 - $299,999 | 1 | 383 |

For the 47 properties, the median price is $150,000 and the average price is $150,925. The highest price is $274,500 and the lowest price is $60,000. The average list price for these homes is $155,827. The average market time is 240 days. The average price per square foot is $64.42.

Return to Viewing Options

Fig. 2. Table for recently sold properties.

recently sold or for which a sale was in progress; and (4) information concerning houses which had been listed but remained unsold. Each table had three columns: a list price range, the number of listings within this range, and the average days on the market for houses in this particular range (Fig. 2).

When subjects finished viewing information about the house, the program moved to an input screen. This screen asked each subject to input the same four values asked for in Ref. [7]: (1) an appraised value for the house; (2) a fair advertising price for the house; (3) a listing price; and (4) an amount reflecting the lowest offer they would accept as seller of the house. Subjects could also click on a button to launch the Microsoft Windows calculator program.

Subjects could return to review previously shown information. The system had the ability to issue warnings if subjects in certain treatment groups came too close to the initial value provided as an anchor. Although the program was structured to provide a logical flow to the sequencing of information, subjects were free to return to prior screens, or to leave the input screen and return to the house information options. The program permitted a subject to view any screen without any time restriction; a subject could return to a particular screen an unlimited number of times. However, forward movement was restricted such that the subjects could not proceed beyond the input screen (and risk the possibility of terminating the experiment) without supplying a value for each of the four variables.

Once subjects had supplied information for each of the variables, they were shown a screen which asked them to indicate, via check-boxes, the information which they had used in deriving their pricing estimates. This screen contained check-boxes pertaining to the information provided for this house, information concerning neighborhood properties, and other information. Additionally, this screen provided subjects with a free-form entry text box, into which they could type comments concerning the experiment or other information used in their pricing decisions. The final screen in the program thanked subjects for participating in the experiment and asked them to click on an ''exit'' button in order to terminate the program.

Because of the experimental nature of the project, the researchers were required to provide a means to terminate the program in the event that a subject decided not to participate further. This was accomplished by coding the Escape key to cancel the program. Alternately, all data for subjects who completed the program were written to an ASCII data file when the subject reached the final screen. The information recorded for each subject consisted of all information that a subject was asked to input, as well as an indication of the treatment group to which the subject had been assigned, the number of screens viewed, and the order in which screens were viewed. Additionally, the first and last values for each of the four variables were recorded, as was the number of values which the subject input for each variable. (The data would show, for example, how many times a particular subject changed his or her mind concerning his or her estimate of the appraisal value for the house.)

### 4.2. Research design

The research design was a $2 \times 2$ full-factorial model with a control. The independent factors were the magnitude of an anchor (either high or low) and the presence of a warning (either there or not). The design included a control group where there was no anchor provided and hence no warnings. High anchors were 12% above the house's actual appraised value. Low anchors were 12% below the house's appraised value. The actual appraised value was US$214,900. Warnings were activated at two different levels: when subjects' estimated values were within a range that covered the anchor provided $\pm 10\%$, or when their estimated values were within a range that covered the anchor provided $\pm 20\%$.

Because warnings were issued only when the subjects came too close to the anchor value, we could not assign the number of subjects under the treatment groups of ''warnings'' vs. ''no warnings''. In one of the pilot tests that preceded the experiment, warnings were initiated only when subjects' estimated values were in a range of the anchor $\pm 10\%$. Fewer subjects than anticipated in the warnings treatment did not receive warnings in the pilot, as their estimates were outside of this range. (Later analysis showed their estimates were affected by the anchors they received, but in many cases, the estimates were outside a range that covered the anchor $\pm 10\%$ of the anchor's value.) In order to ensure that enough subjects in the warnings treatment in the actual experiment received warnings, to ensure a large enough sample for analysis, the range was expanded so that subjects received warnings if their estimates were as much as $\pm 20\%$ of the anchor. Even with this expanded range, eight of the 55 subjects in the warning treatment did not receive warnings. Whether warnings were triggered in the $\pm 10\%$ or the $\pm 20\%$ range had no significant effect on the final estimated values for subjects in the warnings treatment. Therefore, no further distinctions will be made based on the level at which warnings were triggered.

Subjects were randomly assigned to a treatment via a random number generator incorporated into the software. Subjects who were shown an anchor value would receive either a high anchor or a low anchor; the high or low value anchors were likewise assigned on the basis of a random number generator.

### *4.3. Pilot tests*

The experiment was pilot tested in a Decision Support Facility with 21 senior MIS students recruited from a required capstone course. Students were offered extra credit for participating in the experiment. The pilot was conducted in an attempt to decide between a very strongly worded and boldly presented warning (the warning included bright red lettering of 26 point type) and a milder warning. Both warnings cautioned subjects against anchoring their estimated values around the anchor values they had been provided. The results of the first pilot showed no difference between the two warnings.

An additional change to the software at this point involved removing the restriction that the ''lowest acceptable price'' variable should be greater than or equal to the ''appraisal value'' variable. Subjects participating in the pilot believed emotional factors (such as divorce) or professional circumstances (such as job transfer) could influence a seller to settle for a loss on a property.

A second pilot was run with 33 senior MIS students recruited from a required database course. These students, too, were offered extra credit for participating in the experiment. The second pilot supported the results of the first pilot in showing no difference between the two warnings. The researchers decided at this point to use the milder warning in running the experiment because it was less intrusive and had the same apparent impact as the bold warning. No other modifications were made to the program following the second pilot.

### *4.4. The experiment*

The experiment was run at a large southeastern university with 131 students recruited from sections of junior and senior level required IS courses, during the summer and fall terms of 1997. These students were offered extra credit for participating in the experiment.

One of the researchers visited each section of the selected classes. A brief description of the project was provided and the students were invited to participate in the experiment. Specifically, the researcher explained that the researchers were investigating the possible effectiveness of the World Wide Web as a means of conducting a real estate transaction. Further, the students were informed that participation in the project would involve a time commitment of approximately 30 min (this estimate had been derived from the pilot studies). A sign-up sheet, offering 10 slots during three daily time periods across multiple weeks, was circulated to the students. Students were asked to sign-up to participate on a day and time period compatible with their class schedules.

The experiment was conducted in the College's Decision Support Facility. This room has been set up to permit individuals to work without influencing, or being influenced by, their neighbors. Students were asked to sit anywhere in the room. Prior to beginning the experiment, an introductory script and an informed consent form were read aloud to the group of participants. The researcher offered a cash prize incentive to the three participants who came closest to the actual value of the home.

Different versions of the software were randomly installed on the personal computers in the decision support facility. Once the versions were installed, the record as to which version was on which machine(s) was discarded, in order to prevent the researcher from steering subjects toward any particular seat.

Before beginning the experiment, students were given an opportunity to ask questions. Following the question period, and once signed consent forms had been collected, students were asked to begin the project. When a student completed the experiment, she or he was free to go.

## 5. Results

A total of 131 subjects participated in the experiment. Each subject was randomly assigned to one of five treatment groups, as shown in Table 1. However, as was pointed out previously, eight of the 55 subjects in the warnings treatment never saw any warnings, as the estimates they provided were outside of the range of values that triggered warnings. The responses of these eight subjects cannot be considered together with those of subjects who did see warnings, even though they were assigned to the same treatment. This altered the effective distribution of subjects to cells (Table 1). Table 1 also provides the mean and standard deviations for each treatment and control group for all four values that subjects provided.

To test whether the anchoring and adjustment bias occurred in subjects using the appraisal DSS, we compared the four estimations of the house's value provided by subjects using a MANOVA. The factor was the anchor provided, either none (in the control group), high, or low. The statistical significance of the $F$ scores that resulted ranged from 0.016 for the Roy's largest root value ($df = 4$, 126) to 0.108 for the Pillai's trace ($df = 8$, 252). Tests of between subject effects were all significant at the 0.05 level: appraised value ($F(2, 128) = 4.027$, $p < 0.05$); list price ($F(2, 128) = 5.920$, $p < 0.01$); market value ($F(2, 128) = 5.495$, $p < 0.01$); and lowest acceptable price ($F(2, 128) = 6.102$, $p < 0.01$). Hypothesis 1 is supported. Means for the control group and the high and low anchor groups are provided in Table 2.

To test whether the warnings to subjects had any effect on the presence or magnitude of the anchoring and adjustment bias, we conducted a MANOVA to

Table 1
Means, standard deviations, and subject distribution for dependent variables

| Treatment | Appraisal value | List price | Purchase price | Lowest acceptable offer | Cell size (research design) | Cell size (effective distribution) |
|---|---|---|---|---|---|---|
| *Control* | | | | | | |
| Mean | 192,550 | 208,731 | 194,954 | 188,650 | $N = 26$ | $N = 26$ |
| Standard deviation | 24,110 | 25,719 | 26,006 | 27,533 | | |
| *Low anchor, no warning* | | | | | | |
| Mean | 187,842 | 202,394 | 187,076 | 180,242 | $N = 23$ | $N = 26$ |
| Standard deviation | 31,873 | 44,811 | 31,333 | 32,149 | | |
| *High anchor, no warning* | | | | | | |
| Mean | 202,029 | 220,686 | 206,707 | 197,765 | $N = 27$ | $N = 32$ |
| Standard deviation | 30,862 | 25,719 | 20,814 | 23,921 | | |
| *Low anchor, warning* | | | | | | |
| Mean | 182,265 | 198,220 | 185,650 | 175,859 | $N = 33$ | $N = 30$ |
| Standard deviation | 34,177 | 30,593 | 36,587 | 31,421 | | |
| *High anchor, warning* | | | | | | |
| Mean | 200,253 | 219,658 | 200,173 | 194,785 | $N = 22$ | $N = 17$ |
| Standard deviation | 25,091 | 18,989 | 19,073 | 20,261 | | |
| *Total* | | | | | | |
| Mean | 192,575 | 209,405 | 194,808 | 187,075 | | |
| Standard deviation | 30,513 | 31,102 | 28,905 | 28,749 | | |

Table 2
Descriptive statistics for control, high anchor, and low anchor groups

| | No anchor | High anchor | Low anchor |
|---|---|---|---|
| *Appraised value* | | | |
| Mean | 192,550 | 201,413 | 184,854 |
| Standard deviation | 24,110 | 28,736 | 32,947 |
| *List price* | | | |
| Mean | 208,731 | 220,329 | 200,158 |
| Standard deviation | 25,719 | 20,909 | 37,559 |
| *Market value* | | | |
| Mean | 194,954 | 204,440 | 186,312 |
| Standard deviation | 26,006 | 20,271 | 33,950 |
| *Lowest acceptable price* | | | |
| Mean | 188,650 | 196,731 | 177,894 |
| Standard deviation | 27,533 | 22,549 | 31,547 |

compare the four estimates of the house's value provided by subjects. The factor was whether or not warnings were provided. None of the $F$ scores were statistically significant at the 0.05 level. Warnings, then, did not eliminate the anchoring and adjustment effect.

However, it is possible to test whether warnings served to lessen the effects of the anchoring and adjustment bias. If warnings work to make subjects aware of the bias, then subjects who received warnings should have modified their estimates away from the anchor. While it is possible that all subjects may have adjusted their estimates during the process of moving through the system, those receiving warnings should have moved their estimates further away than those who did not receive warnings.

We can examine this question by looking first at the distance between the anchor and the initial estimates given by subjects for all four values of the house, and then by looking at the distance between the anchor and the final estimates for all four values. One would expect there to be no significant differences, in terms of the distance between the estimate and the anchor, between those with warnings and those without for the initial estimates made, since initial estimates were made before any warnings were issued. That is indeed the case. A one-way ANOVA with warnings as the factor resulted in no statistically significant $F$ scores.

If the warnings had some effect, however, one would expect larger distances between the anchor and the final estimates for those who had been warned than for those who had not. In reviewing the data in Table 1, it is clear that subjects who received warnings entered relatively lower final values, compared to their peers who did not receive warnings. Subjects who received warnings did in fact provide final values further away from the anchors they were provided than those who did not receive warnings. This is consistent for high and low anchors across all four values estimated for the house. Although the differences are in the right direction, the question is whether the differences are statistically significant. A one-way ANOVA with warnings as the factor resulted in no statistically significant $F$ scores, so Hypothesis 2 is not supported.

To determine if the warnings had any effect on the subjects and their deliberations, we conducted a post hoc analysis, in which we compared the number of times subjects entered estimates for each of the four values they were asked to generate, comparing those who received warnings to those who did not. Subjects in the control group were excluded from the analysis. Descriptive statistics for the number of

Table 3
Descriptive statistics for warning and no warning groups, for number of estimates entered for each value

| | No warnings | Warnings |
|---|---|---|
| *Appraised value* | | |
| Mean | 1.8 | 2.2 |
| Standard deviation | 0.96 | 1.7 |
| Range | 1 to 4 | 1 to 8 |
| *List price* | | |
| Mean | 1.6 | 2.2 |
| Standard deviation | 0.9 | 2.2 |
| Range | 1 to 4 | 1 to 11 |
| *Market value* | | |
| Mean | 1.6 | 1.7 |
| Standard deviation | 0.88 | 1.5 |
| Range | 1 to 4 | 1 to 8 |
| *Lowest acceptable price* | | |
| Mean | 1.5 | 1.5 |
| Standard deviation | 0.8 | 1.1 |
| Range | 1 to 4 | 1 to 7 |

times subjects entered estimates are contained in Table 3.

As Table 3 shows, subjects who received warnings averaged entering more estimates for each value than those who did not receive warnings. The only value for which the differences were statistically significant, at the $\alpha < 0.1$ level, was list price ($F(1, 103) = 2.988$, $p = 0.087$). The presence of warnings, then, did seem to have some impact on the attention paid to estimating the value of the house, at least for the house's list price.

## 6. Discussion

Northcraft and Neale [7] demonstrated that the anchoring and adjustment bias was not a simple parlor trick but instead was a limitation of decision-making that occurred in the context of real business problems with both novice and expert decision-makers. This study demonstrated that moving the context of a real business problem within a computer-based DSS did not lessen the strength of the anchoring and adjustment bias. As the study demonstrated, the presence and direction of the anchor had a profound effect on subjects' estimated values of a house. Those with anchors 12% higher than the appraised value of the house generated values for the house close to the anchor they were given. Those with an anchor 12% lower generated values close to the anchor they were given. The overall effect was just as robust as that demonstrated in Refs. [7,14].

One attractive feature of a DSS is the ability of the designer to take into account the biases and limitations inherent in the decision-making process and design accordingly. In this study, some subjects received warnings when their estimates for the value of the house were too close to the anchor they had received. The attempt was to motivate subjects to change their final estimates, once they had been made aware of how the stated asking price was affecting their evaluation decisions. For at least the list price value, subjects who received warnings did change their estimated values more times than those who did not receive warnings. However, there was no main effect for warnings on any of the final values subjects entered, nor were there any statistically significant differences in terms of the distance between the anchor and the final values for warned and unwarned subjects. Unlike the findings in Ref. [3], where subjects received general verbal warnings before beginning their work, the effect of anchoring and adjustment in this study was not significantly lessened. Since we compared different warning messages in our pilot studies, and since there were no differences between loud and more reserved warning messages, we are confident the type of warning message itself was not responsible for the lack of effect. Rather, anchoring and adjustment seems to be robust and resistant to the influence of warning messages alone.

It is difficult to draw clear implications from a single study, but taken in conjunction with the findings from related studies, there are at least two lessons here for designers of DSS. The first is that the rationalization of a decision-making process by formalizing it within a computer-based information system does not seem to make the process itself more rational. The bias that operates without the information system continues to operate within it. Second, there may well be ways to mitigate the effects of decision-making biases, but the techniques employed to do so may themselves be limited and require substantial interventions. The first step in making progress in this area is to recognize that the decision-making literature is of direct consequence to designing and building DSS. The next step is to focus research on how the biases discussed in this literature and present in DSS can be mitigated or even eliminated.

Although anchoring and adjustment appears to be resistant to warnings alone, it could be susceptible to other debiasing strategies. There may be many different remedies DSS designers could employ to deal with limitations in decision-making. One set of possible remedies would involve providing warnings and explanations as to why a particular solution may not be as objective as the decision-maker might believe. Such remedies would fall under the decision support design principle of guidance, as suggested in Ref. [11]. However, warnings are only the first step in Fischhoff's schedule of debiasing strategies [5]. Other remedies built into DSS could involve the other three steps in Fischhoff's schedule of strategies: issuing descriptions of the direction of bias; feedback which personalizes the implications of the warnings; and extended training programs. Further research is warranted to explore merits of these techniques.

A second set of remedies might prevent the decision-maker from going any further with the process until he or she made changes to what the system's heuristics considered to be biased estimates or processes. Such remedies would be fall under Silver's restrictiveness design principle [11].

As is the case with any laboratory study, this research has its limitations. One limitation relates to the use of student subjects. However, the task was salient for the business school students used as subjects, and they treated the task seriously. Also, financial incentives were used to reduce errors caused by insufficient attention. Further, Ref. [15] compared susceptibility of students and experienced managers to anchoring. They found the same kind of anchoring effects with experienced managers that were observed among students. Also, as pointed out earlier, Ref. [7] found similar anchoring effects with both students and experienced real estate agents.

In summary, this research applied the theory of anchoring and adjustment to the context of DSS and tested the ability of DSS to reduce this bias. We found that the anchoring and adjustment effects were robust within the context of a DSS. Anchoring and adjustment could pose a significant risk to the quality of certain types of decision-making, and future research should explore ways of mitigating this bias in computer-based DSS.

## Acknowledgements

The authors would like to thank Joe Valacich and Mun Yi for their helpful comments on earlier versions of this paper. We would also like to thank Richard Snow for his invaluable assistance.

## References

[1] M. Alpert, H. Raiffa, A progress report on the training of probability assessors, in: D. Kahneman, P. Slovic, A. Tversky (Eds.), Judgment under Uncertainty: Heuristics and Biases, Cambridge Univ. Press, Cambridge, England, 1982, pp. 294–305.

[2] S.L. Alter, Decision Support Systems, Addison-Wesley, Reading, MA, 1980.

[3] R.A. Block, D.R. Harper, Overconfidence in estimation: testing the anchoring-and-adjustment hypothesis, Organ. Behav. Hum. Decis. Processes 49 (1991) 188–207.

[4] T. Connolly, D. Dean, Decomposed versus holistic estimates of effort required for software writing tasks, Manage. Sci. 43 (7) (1997) 1029–1045.

[5] B. Fischhoff, Debiasing, in: D. Kahneman, P. Slovic, A. Tversky (Eds.), Judgment under Uncertainty: Heuristics and Biases, Cambridge Univ. Press, Cambridge, England, 1982, pp. 422–444.

[6] R.M. Hogarth, Judgement and Choice, 2/e, Wiley, Chichester, 1987.

[7] G.B. Northcraft, M.A. Neale, Experts, amateurs, and real estate: an anchoring-and-adjustment perspective on property pricing decisions, Organ. Behav. Hum. Decis. Processes 39 (1987) 84–97.

[8] I. Ritov, Anchoring in simulated competitive market negotiation, Organ. Behav. Hum. Decis. Processes 67 (1) (1996) 16–25.

[9] M.C. Roy, F.J. Lerch, Overcoming ineffective mental representations in base-rate problems, Inf. Syst. Res. 7 (2) (1996) 233–247.

[10] G.L. Sharp, B.L. Cutler, S.D. Penrod, Performance feedback improves the resolution of confidence judgments, Organ. Behav. Hum. Decis. Processes 42 (1988) 271–283.

[11] M. Silver, Decision support systems: directed and nondirected change, Inf. Syst. Res. 1 (1) (1990) 47–70.

[12] R.H. Sprague, A framework for the development of decision support systems, MIS Q. 4 (4) (1980).

[13] P. Todd, I. Benbasat, An experimental investigation of the impact of computer-based decision aids on decision-making strategies, Inf. Syst. Res. 2 (2) (1991) 87–115.

[14] A. Tversky, D. Kahneman, Judgment under uncertainty: heuristics and biases, Science 185 (1974) 1124–1131.

[15] G. Whyte, J. Sebenius, The effect of multiple anchors on anchoring in individual and group judgement, Organ. Behav. Hum. Decis. Processes 69 (1) (1997) 75–85.

Joey F. George (PhD, University of California at Irvine, 1986; AB, Stanford University, 1979) is Professor of Information Systems and the Thomas L. Williams Jr. Eminent Scholar in IS at Florida State University. His research interests focus on the use of information systems in the workplace, including computer-based monitoring, group support systems, and deception in computer-mediated communication.

Kevin P. Duffy is Assistant Professor of MIS at the Division of Business Administration, College of Business and Economics, West Virginia University. He did his doctoral work in MIS at the Florida State University. His research interests include information systems planning, change management, and organizational learning and unlearning.

Manju Ahuja is an Assistant Professor of MIS at Florida State University. Her publications have appeared in journals such as *Organization Science, Communications of the ACM* and the *Journal of Computer-Mediated Communications*. She obtained her PhD in MIS from the University of Pittsburgh. She taught at Pennsylvania State University and University of Pittsburgh before joining Florida State University in 1996. Prior to starting a PhD program, Manju Ahuja worked as a Systems Analyst designing on-line database systems for several years. She is actively involved in research on issues related to virtual organizations, knowledge management, social networks, and use of information technologies for collaborative work.