Journal of Experimental Psychology: General
1996, Vol. 125, No. 4, 387–402

0096-3445/96/$3.00

A New Look at Anchoring Effects: Basic Anchoring and Its Antecedents

Timothy D. Wilson, Christopher E. Houston, and Kathryn M. Etling
University of Virginia

Nancy Brekke
Lake Forest College

In previous anchoring studies people were asked to consider an anchor as a possible answer to the target question or were given informative anchors. The authors predicted that *basic anchoring effects* can occur, whereby uninformative numerical anchors influence a judgment even when people are not asked to compare this number to the target value. Five studies supported these hypotheses: Basic anchoring occurs if people pay sufficient attention to the anchor value; knowledgeable people are less susceptible to basic anchoring effects; anchoring appears to operate unintentionally and nonconsciously in that it is difficult to avoid even when people are forewarned. The possible mechanisms of basic anchoring and the relation between these mechanisms and other processes of judgment and correction are discussed.

How many physicians practice medicine in your community? Suppose that before answering this question you had just written down the street address of a friend, which happened to be 3459. Would the street address influence your answer to the question about the number of physicians? Would your estimate be higher than if your friend's address happened to be 63? The purpose of this article is to explore the nature of anchoring effects such as this hypothetical one whereby an arbitrary number in memory influences an unrelated judgment. Despite years of research on anchoring and adjustment, there are unanswered questions about the conditions under which anchoring processes are initiated. Is writing down an address, for example, sufficient to cause an anchoring effect, or must people be explicitly asked to compare the address to the number of physicians?

The term *anchoring* has been used to describe a number of phenomena, including the effects of exposure to one stimulus (e.g., a heavy weight) on psychophysical judgments of another stimulus (a lighter weight; Helson, 1964); the process whereby exposure to one stimulus or situation (e.g., an extremely happy, idyllic situation) changes people's perception of the endpoints of the scale used when making a rating on that dimension (e.g., one's current level of happiness; Parducci, 1968); and the case in which people's answers to a question are influenced by thinking about an arbitrary value as a possible answer to the question (e.g. Tversky & Kahneman, 1974).

Timothy D. Wilson, Christopher E. Houston, and Kathryn M. Etling, Department of Psychology, University of Virginia; Nancy Brekke, Department of Psychology, Lake Forest College.

The preparation of this article and the research it reports were supported by National Institute of Mental Health Grant MH41841 and National Science Foundation Grant SES-8708007.

We would like to thank Daniel Kahneman, Jay Meyers, and Chris Wetzel for helpful comments on a draft of this article.

Correspondence concerning this article should be addressed to Timothy D. Wilson, Department of Psychology, Gilmer Hall, University of Virginia, Charlottesville, Virginia 22903–2477. Electronic mail may be sent via Internet to tdw@virginia.edu.

We are concerned with this last type of anchoring effect, in which people are influenced by arbitrary anchor values. In a classic study by Tversky and Kahneman (1974), for example, people estimated whether a number that resulted from the spin of a "wheel of fortune" was more or less than the percentage of African countries in the United Nations and then guessed the correct percentage. People's guesses were substantially lower if they began with a low anchor than if they began with a high anchor. Anchoring effects such as these have been found with many types of responses, such as answers to general knowledge questions (Russo & Schoemaker, 1989), utility assessment (Hershey & Schoemaker, 1985), causal attribution (Quattrone, 1982), the detection of human deception (Zuckerman, Koestner, Colella, & Alton, 1984), predictions of future performance (Switzer & Sniezek, 1991), predictions of the likelihood of future events (Plous, 1989), and task persistence (Cervone & Peake, 1986). At a more theoretical level, anchoring is a central part of explanations of diverse phenomena, such as why people adjust insufficiently from an initial causal attribution (Quattrone, 1982), make conservative judgments of probability (Edwards, 1968; Slovic & Lichtenstein, 1971), believe that past events were inevitable (Fischhoff, 1975), and believe false information (Gilbert, 1991). In this age of specialization and fragmentation of psychology, it is rare to find a single, relatively simple process that explains such diverse phenomena.

It is thus ironic that anchoring processes are themselves poorly understood (Chapman & Johnson, 1995; Jacowitz & Kahneman, 1995; Strack & Mussweiler, 1995). Although several recent explanations have been offered, we suggest that an important part of the process has been ignored: the conditions under which the anchoring processes are initiated whereby people consider an arbitrary value as a possible answer to a question. Virtually all previous studies have focused on a later stage of the process, after an anchor value has been compared to the target value, and thus do not speak to the question of when arbitrary numbers will be used as a point of comparison.

Most explanations of anchoring have focused on the way in which people integrate an anchor and target value to formulate an answer. Once the number 3,459 is considered as a possible answer to the question about physicians, for example, how do people integrate this number into their estimate of the correct answer? A number of integration and adjustment processes have been suggested, such as averaging (Lopes, 1985), insufficient adjustment (Tversky & Kahneman, 1974), and adjustment until people are just within the range of plausible values (Quattrone, Lawrence, Finkel, & Andrus, 1984). We refer to this part of the anchoring process as integration and adjustment, because it involves integrating the anchor and target values, adjusting one's response away from the anchor value (see right-hand box in Figure 1), or both. Our list of integration processes is by no means exhaustive (see Chapman & Johnson, 1995; Jacowitz & Kahneman, 1995; Strack & Mussweiler, 1995, for a more complete review). For our purposes, the important point is that most explanations focus on this stage of the process, whereby an anchor has already been considered as a possible answer to the target question. That is, integration and adjustment cannot occur until there is something to integrate and something from which to adjust.[1]

Chapman and Johnson (1995) have argued that an earlier part of the anchoring process is critical: the retrieval from memory of features of the target (see middle box in Figure 1). They propose a confirmatory search mechanism, whereby people focus more on reasons why the target is similar to the anchor than on reasons why the target is different from the anchor. Once the number 3,459 is used as a point of comparison, for example, people might retrieve from memory reasons why the number of doctors in their community is consistent with this number, which influences their estimate of the correct answer. Thus, before integration and adjustment processes come into play, the way in which people retrieve information about the target may be important. Again, however, the question of why people compare an anchor to a target value in the first place is not addressed.

In many studies the anchoring process is initiated by explicitly asking people to compare the anchor value to the target value. In Tversky and Kahneman's (1974) wheel of fortune study, for example, people were first asked to judge whether the anchor value was higher or lower than the percentage of African countries in the United Nations (see also studies by Chapman & Johnson, 1995; Cervone & Peake, 1986; Joyce & Biddle, 1981; Quattrone et al., 1984; Russo & Schoemaker, 1989). In a sense the experimenters started the process for participants; thus these studies do not address the question of whether people would have compared the anchor to the target value in the absence of instructions to do so.

In other studies people were not asked to compare the anchor value to the target question, but the anchor value was informative rather than arbitrary. That is, the anchor was not described as randomly chosen or arbitrary, and participants might have inferred that it had probative value (e.g., Davis, Hoch, & Ragsdale, 1986; Lovie, 1985; Northcraft & Neale, 1987; Zuckerman et al., 1984). For example, Northcraft and Neale (1987) gave real estate agents a packet of information about a local house and asked them to estimate its appraised value. Included in this packet was the listing price for the house, which the authors varied from low to high. The listing price acted as an anchor, such that the lower it was, the lower the real estate agents' estimates of the appraised value of the house. In this case the anchor was not completely arbitrary (i.e., listing prices are often based on appraised values, and the owners of the house presumably did not choose the listing price randomly). Thus, it is not surprising that people used the listing price when forming their judgments.

An important question thus remains unanswered: Will anchoring effects occur when people have an uninformative number in short-term memory and are not asked to compare this number to the target value? This question is of considerable importance to the generalizability of laboratory demonstrations of anchoring. In everyday life it is relatively uncommon to be provided with an arbitrary number and asked to compare it to an unrelated judgment. It is much more common to have a number on one's mind, such as the address of a friend you just wrote down or a phone number you just dialed, when forming a judgment on an unrelated task. If people can be demonstrated to use arbitrary values as anchors even when not explicitly asked to do so, anchoring effects are probably much more common in everyday reasoning than previously demonstrated. This question is also of considerable methodological importance, because it is common for research participants to have an arbitrary value in mind (e.g., a question number or answer to a previous question) when responding to another, unrelated item.

Kahneman and Knetsch (1993) suggested a model in which anchoring can occur when the anchor is arbitrary and people are not asked to consider the anchor as a possible answer. They proposed a backward priming mechanism, whereby the need to answer a question about the target causes people to consider as a possible answer any number in short-term memory, no matter what the source of this number happens to be. This process is automatic and unreflective, argued Kahneman and Knetsch (1993), such that a number stored in short-term memory due to one task (e.g., thinking about a street address) is brought to mind automatically and considered as a possible answer to an unrelated question (e.g., the number of physicians in one's town).

In contrast, anchoring may be more similar to standard priming effects, whereby attention to a stimulus increases

[1] It may be that there is another stage in the anchoring process, whereby people must decide which value to express and how to translate this value onto a response scale. Conversational norms (Grice, 1975; Jacowitz & Kahneman, 1995) and self-presentational concerns (e.g., Schlenker, 1980) could come into play here, whereby people communicate the value that they think conveys the most information or puts them in the best light. We have omitted this stage from Figure 1, however, because it is clear that anchoring is not due only to conversational norms or self-presentational concerns. People's judgments are influenced by anchor values even when it is clear that the anchor value is completely uninformative and when individual responses are anonymous (e.g., Tversky & Kahneman, 1974; Strack & Mussweiler, 1995).

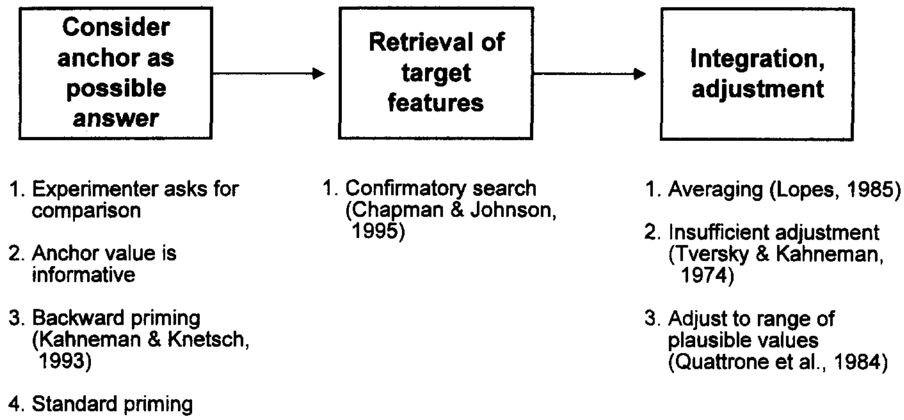

Figure 1. The anchoring process.

the activation potential of a category or value, increasing the likelihood that this value will be used when judging a subsequently encountered stimulus (e.g., Higgins, 1996). The difference between the Kahneman and Knetsch (1993) account and standard priming effects concerns the point at which activation of the anchor occurs: when it is initially encoded (the standard view) or when people are asked a question about the target value (the backward priming view). This distinction, while important, is not critical to our purposes. The main goal of our studies was to see if there is such a thing as a *basic anchoring effect*, which we define as the case whereby people's judgments of a target are influenced by an anchor that is completely uninformative and people are not asked to consider the anchor as a possible target value.[2]

If it turns out that there are basic anchoring effects, several interesting questions arise. Will any number in short-term memory influence subsequent judgments? It seems obvious that the answer is no. Surely people will be less likely to consider arbitrary values as answers if they can easily retrieve an answer from memory (such as the number of children they have) or if they can easily calculate the answer (such as adding 2 plus 2; see Kahneman & Knetsch, 1993). As obvious as this prediction is, however, it is still worth testing in an area known for its nonobvious results. In one of the present studies we predicted that people who are highly knowledgeable will be less susceptible to basic anchoring effects, because they can retrieve an answer directly from long-term memory. Evidence for this prediction has been obtained in some previous anchoring studies, in which high certainty about an answer was associated with smaller anchoring effects (e.g., Chapman & Johnson, 1994). In these studies, however, participants were instructed to compare the anchor value to their answer to the target questions. It is not yet clear whether knowledge or certainty moderates basic anchoring effects.

Even if people cannot retrieve an answer from long-term memory or compute it easily, it seems doubtful that their estimate will be anchored on any number that happens to be in short-term memory. There are many, many arbitrary numbers in our minds throughout a given day, such as the temperature that was just announced on the radio, the numbers on a computer keyboard that we just pressed, the numbers on the dial of a clock we just consulted, or the page numbers of a book or questionnaire we read. It seems unlikely that numbers considered as briefly as these would be used to make an unrelated judgment. We thus hypothesized that numbers will only intrude on unrelated judgments if people pay sufficient attention to them.

A final issue is whether anchoring processes are the result of a rational, conscious decision process in which people judge the relevance of the anchor to the judgment or a nonconscious process in which a number in short-term memory is unintentionally considered as a possible answer to an unrelated question. Previous studies have not ruled out the possibility that people use anchors consciously and deliberately, because people have been found to use anchors only when they are asked to use them or when the anchors are informative. We hypothesized that basic anchoring occurs unintentionally and nonconsciously.

[2] Schwarz and Wyer (1985) conducted a study that came close to demonstrating basic anchoring effects. They found that the way in which people rank-ordered the importance of a set of environmental issues—either from most important to least important or from least important to most important—anchored their absolute judgments of the importance of these and related issues. Participants who began their rankings with the most important issue subsequently rated all the issues as more important than did people who began their rankings with the least important issue. Most relevant to the current concerns was the finding, in Schwarz and Wyer's (1985) Study 2, that the order in which people ranked the importance of one topic (the importance of attributes in choosing a marriage partner) anchored their ratings of the importance of a completely unrelated topic (environmental issues). Thus, people's responses to one topic anchored their judgments on an unrelated topic, even when they were not asked to compare their first response to the second one. Nonetheless, the domain of people's judgments was the same in the two tasks: ratings of importance. The question remains whether completely arbitrary values can lead to anchoring effects when people are not asked to compare these values to the target judgment and when the judgment is on a different dimension than the anchor variable.

One implication of this hypothesis is that anchoring effects are difficult to prevent. Wilson and Brekke (1994) outlined the conditions that must be fulfilled in order for people to avoid making contaminated judgments, in this case preventing the arbitrary anchor from influencing their judgments: People must be aware that bias has occurred, be motivated to correct the bias, know the direction and magnitude of the bias, and have sufficient control over their responses to be able to correct for the bias. If anchoring processes occur unintentionally and outside of awareness, it should be difficult for people to know the direction and magnitude of the effect. Even if people are aware that anchoring effects occur and are motivated to prevent them, it may be difficult to tell how much their response has been affected and thus difficult to adjust correctly for the bias. For example, even if people knew that their estimate of the percentage of African countries in the United Nations was influenced by the number that came up on the wheel of fortune, they would not know how much to adjust their answer to correct for the bias. Consequently, anchoring effects may be difficult to avoid even when people are aware that they occur.

To summarize, the present studies tested four hypotheses:

> *Hypothesis 1.* Basic anchoring effects will occur, such that a number that is completely uninformative will influence people's judgments, even when people are not asked to compare this number to the target value.

> *Hypothesis 2.* People who are knowledgeable about the target question will be less influenced by arbitrary anchors.

> *Hypothesis 3.* People must pay sufficient attention to a numerical value in order for basic anchoring effects to occur. If a number is considered only briefly, anchoring will not occur.

> *Hypothesis 4.* Anchoring processes are unintentional and nonconscious; therefore, it is very difficult to avoid anchoring effects, even when motivated to do so and forewarned about them.

We began our investigations with a test of Hypotheses 1 and 2. People in the anchoring condition drew a random number. As in past studies, some were asked to compare the number to the target question (the number of countries in the United Nations). Others were asked to compare the number to an unrelated question (the number of physicians listed in the local Yellow Pages). We predicted that anchoring effects would occur in both conditions. If so, this would be an example of a basic anchoring effect (Hypothesis 1), because anchoring would occur even when a value was uninformative and people did not compare it to the target value. We tested Hypothesis 2 by seeing whether people's level of knowledge about the target question moderated the effects.

Study 1

Method

We distributed questionnaires to 116 students in undergraduate psychology classes and summer school classes in several disciplines at the University of Virginia. Participants were told that the study was like a quiz show, in which they would be asked to choose a random number from a bag containing "a large range of numbers" and to judge whether this number was less than, equal to, or greater than the answer to a general knowledge question. In the anchoring conditions, the number people chose was always 1,930. In the control condition, the slip of paper said, "You have chosen a special piece of paper, entitling you to skip a round in the game. In the space where it says 'THE NUMBER ON MY PAPER IS,' please write 'none.' Then go on to the next page."

Participants were randomly assigned to compare the anchor number to the question they would answer next (the dependent measure) or to an irrelevant question. In the relevant condition, people were asked to judge whether their number was less than, equal to, or greater than the number of countries in the United Nations, which was the main dependent measure. In the irrelevant condition, people were asked to judge whether their number was less than, equal to, or greater than an unrelated question: how many physicians and surgeons there were in the local phone book. In this condition, people thus received an irrelevant anchor and were never asked to compare it to the target question (about the United Nations).

If people were in the control (no-anchor) condition they received the same instructions as people in either the relevant or irrelevant conditions. Because they did not receive an anchor number, however, they were instructed to leave this question blank. Giving people in the control condition the same instructions as people in the relevant or irrelevant conditions allowed us to maintain a fully crossed, 2 (control vs. anchor) × 2 (relevant vs. irrelevant) design, though it should be kept in mind that the control-relevant and control-irrelevant cells were essentially the same, because no anchor value was received in either cell. All participants then responded to the dependent measure, how many countries there are in the United Nations. Finally, people rated how knowledgeable they were about the number of countries in the United Nations and how confident they were in their answer to this question, both on 9-point scales.

Results and Discussion

In this and all subsequent studies, the distribution of responses on the main dependent measure was skewed and there were large differences in variance across conditions. To increase the homogeneity of the variances and normalize the distributions, we performed logarithmic transformations of the data (Winer, 1971) and used these transformed scores in all analyses. The means we report in all figures were computed from the untransformed data. In most cases the results were very similar when the means of the logarithmic scores were converted back to raw scores; one exception is noted below.

We conducted a multiple regression that conformed to a 2 (control vs. anchor) × 2 (relevant vs. irrelevant anchor) × Level of Knowledge analysis of variance (ANOVA), with the last factor included as a continuous variable (Aiken & West, 1991). This analysis revealed a significant main effect of the anchoring manipulation, $t(108) = 2.16$, $p < .05$. People who received the anchor value guessed that there were more countries in the United Nations than did people in the control condition. However, the effect of this manipulation was moderated by people's knowledge, as reflected by a significant Anchor × Knowledge interaction, $t(108) =$

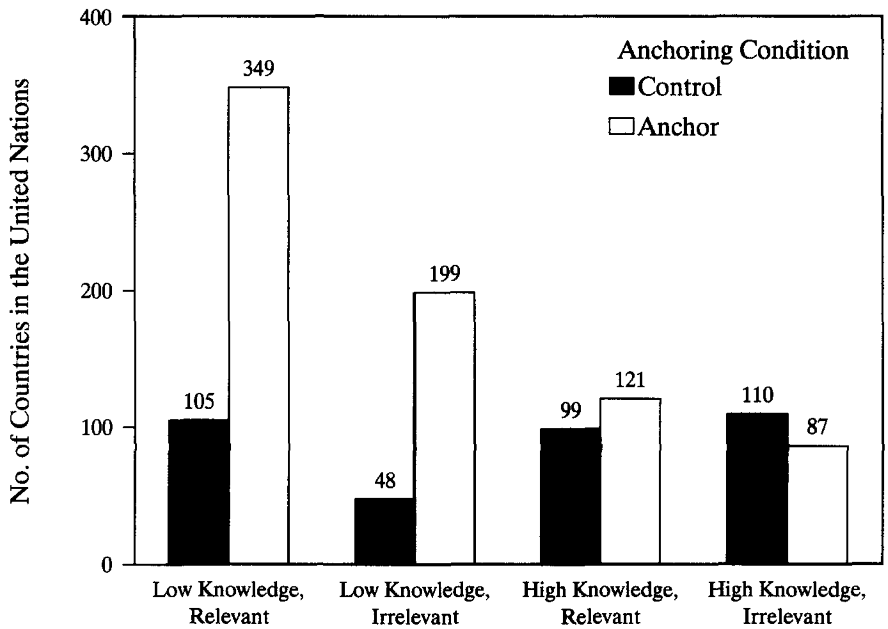

Figure 2. Study 1: Ratings of the number of countries in the United Nations as a function of the anchoring condition, relevance of the anchor to the target question, and prior knowledge.

2.05, $p < .04$. Figure 2 displays the means based on a median split on the knowledge variable and shows that anchoring effects occurred only among people low in knowledge. In addition, the main effect of the relevance manipulation was significant, $t(108) = 2.61$, $p = .01$, reflecting the fact that people who received the relevant anchor guessed a higher number than those who received the irrelevant anchor. This effect was qualified by a nearly significant Knowledge × Relevance interaction, $t(108) = 1.78$, $p = .08$, reflecting the fact that the just-described effect of relevance was found only among people low in knowledge.

Most relevant to our current concerns, anchoring effects occurred in both the relevant and irrelevant conditions among people low in knowledge. Considering only people low in knowledge, the main effect of the anchor manipulation was significant, $F(1, 59) = 6.59$, $p = .01$, $r = .32$, and this main effect was not qualified by the relevance of the anchor, interaction $F(1, 59) < 1$, *ns*.[3] These results are consistent with Hypothesis 1: It was not necessary to ask people to consider the anchor as a possible answer to the target question in order for anchoring effects to occur. Hypothesis 2 was also confirmed, in that the anchor influenced unknowledgeable but not knowledgeable people. It is interesting to note that at the time the study was conducted the correct answer to the question about the number of countries in the United Nations was 159; thus knowledgeable people gave estimates that were fairly accurate—at least more so than unknowledgeable people who received the anchor.

As shown in Figure 2, the magnitude of the anchoring effect was larger when people compared the anchor to the relevant question than when they compared it to the irrelevant question ($rs = .37$ and $.25$, respectively). The difference in the magnitude of the effect was not reliable, however, as indicated by the fact that the Anchor × Relevance interaction was nonsignificant among unknowledgeable people. Instead, there was a main effect of relevance, $F(1, 59) = 9.71$, $p < .01$, $r = .38$, reflecting the fact that those who received the initial question about the United Nations gave higher answers, in both the control and anchoring conditions. It may be that the extra thought about the question yielded more extreme responses, similar to the effects of thought on the extremity of attitudes (Tesser, 1978).

The main result of Study 1 was that a basic anchoring effect occurred among unknowledgeable people. As far as we know this is the first demonstration of an anchoring effect when the number is uninformative and people are not asked to compare the anchor to the target value. A number of questions about this effect, however, remain. First, asking people in the irrelevant condition to compare the anchor to an unrelated question might have caused them to compare it to the target question as well. It is thus not clear whether anchoring effects occur when people attend to a number but are not asked to consider it as a possible answer to any subsequent question. Second, it is not clear how much attention people must pay to a value for it to cause an anchoring effect. As discussed earlier it seems implausible that every number that enters short-term memory anchors the next judgment we make. Study 2 addressed these ques-

[3] All effect sizes are reported in correlational units. Cohen (1988) suggested that *rs* of .5, .3, and .1 be considered large, moderate, and small effects, respectively.

tions by introducing the anchor value in a different way and manipulating how much people attended to the anchor.

Study 2

Method

Participants were 549 students in undergraduate psychology classes at the University of Virginia. The procedure was identical to that of Study 1, with the following exceptions: We told participants in the anchoring conditions that we needed to assign each of them a unique identification (ID) number. An adhesive note was attached to each questionnaire with a number written on it, which, we said, had been randomly chosen. In fact, the number ranged from 1928 to 1935 for all participants (we varied the number slightly in case people glanced at their neighbors' questionnaires). Participants were asked to copy this number on their questionnaire. On the next page they were told that in order to save time, people would be randomly assigned to answer different sets of questions at the end of the questionnaire. We asked people to check their ID number on the adhesive note to see which set of questions they should answer.

People were asked to check different properties of their ID number that required different amounts of attention to the number. In the red–blue condition, people were asked to note whether the number was written in red or blue ink. At a later point, we said, people in the red condition would answer one set of questions, whereas people in the blue condition would answer a second set. In fact, the number was written in blue for all participants. In the four-digits condition, participants checked whether the number was a four-digit number. In the GT-100 condition, participants checked whether the number was greater than 100. In the GT-1920–1940 condition, participants checked whether their number was greater than either 1920 or 1940 (no differences were found between those who received 1920 vs. 1940, so we collapsed across these groups). Finally, we included a comparison condition that was similar to the relevant condition in Study 1, in which people checked whether their ID number was greater than, equal to, or less than the number of physicians listed in the local Yellow Pages (which was the dependent variable in this study). This is the only condition in which people were instructed to consider their ID numbers as an answer to a question.

We assumed that the different judgments people made about their ID number required different amounts of attention to it. The lowest amount of attention was required in the red–blue and four-digits conditions, because here people did not even have to process their ID as a number; they simply had to perceive the color of the ink or count the number of digits. In the GT-100 condition, people had to process the fact that their ID number (which ranged from 1928 to 1935) was a number greater than 100, which presumably could be done easily and quickly. In the GT-1920–1940 condition, people had to attend to their ID number more carefully, because 1920 and 1940 were closer in value to their ID numbers. People in the comparison condition had to pay the most attention to their ID number, because they had to write their ID number again and check whether it was less than or greater than the number of physicians in the phone book—a judgment that presumably took more thought.

To maintain our cover story about different people answering different questions, people were told to complete the page with demographic questions and questions about their level of knowledge only if their ID number met the specified criterion, for example, if it was written in blue. (All participants' numbers met the requisite criterion, such that everyone completed this page.) In addition to the questions asked in Study 1, we assessed people's awareness of how much the anchor values influenced their ratings. People in the anchoring conditions rated the extent to which they thought their answer to the physicians question was influenced by writing down their ID number and how much they thought other people's answers were influenced by writing down their ID numbers. Both of these questions were answered on 9-point scales with the endpoints labeled "decreased it a great deal" and "increased it a great deal." The midpoint of the scale, 5, was labeled "had no effect." Finally, an additional page of unrelated general knowledge questions was included at the end, ostensibly to be completed by people whose ID numbers met a different criterion (e.g., written in red).

Results and Discussion

The main dependent measure was the number of physicians people thought were listed in the phone book. One reason we chose this question was that people would presumably have little knowledge about the correct answer. Given that knowledge moderated the anchoring effects in Study 1, we wanted a question about which people knew little. As it turned out there were a few people in this and subsequent studies who reported that they were quite knowledgeable about the topic (budding physicians, perhaps). Given the results of Study 2, we eliminated them from the analyses. Our decision rule was to eliminate anyone in this and all subsequent studies whose reported knowledge was above the midpoint of the scale.[4] The number of such people was too small to include knowledge as a factor in the design (e.g., it was 9% of the participants in Study 2). Including the knowledgeable participants did not, for the most part, appreciably change the results, though as would be expected from the results of Study 1, it did weaken the effects of the anchoring manipulations in some cases.

As predicted, the more people had to pay attention to the anchor value, the greater the anchoring effect (see Figure 3). A one-way ANOVA revealed a significant overall effect of condition, $F(5, 494) = 12.18, p < .0001$. A linear trend was also significant, $F(1, 494) = 50.86$, $p < .0001$, and accounted for 84% of the between-groups variance. The remaining between-groups variance was nonsignificant, $F(4, 494) = 2.51, p > .05$. Thus, the more attention people had to pay to anchor value, the larger the anchoring effect. A Newman–Keuls test of the differences between means revealed that all of the anchoring conditions except the four-digits condition were significantly higher than the control condition at $p < .05$. The mean in the comparison condition differed significantly from all other means; none of the other anchoring conditions differed significantly from each other. The effect was relatively large in the comparison

[4] We measured knowledge with two questions in each study: how knowledgeable people said they were about the topic and how confident they were in their answers to the main dependent measure, both answered on 9-point scales. Because people's answers to these questions were significantly correlated in each study (mean $r = .60$), we used as our measure of knowledge people's average response to the two questions and eliminated people whose average response was above the midpoint of 5.

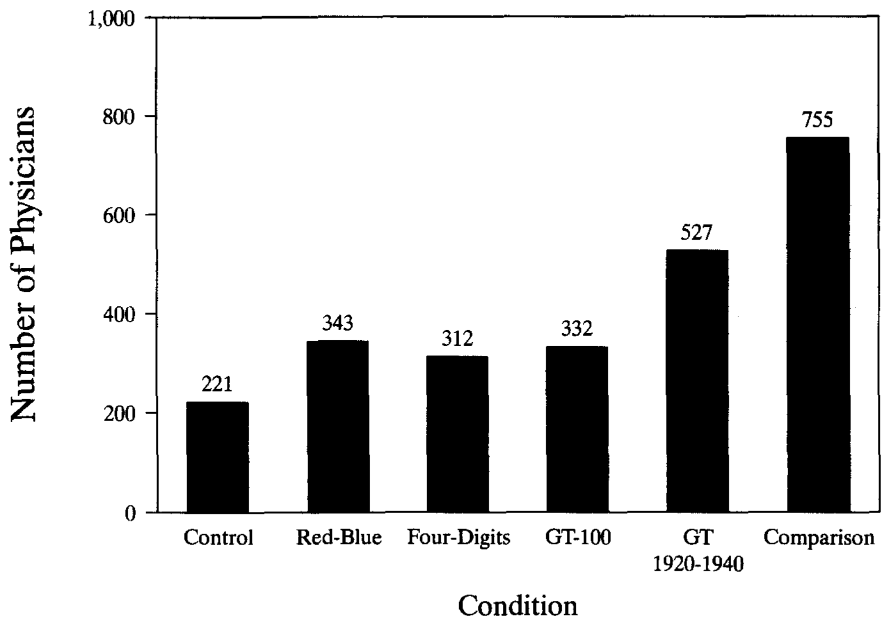

Figure 3. Study 2: Ratings of the number of physicians in the yellow pages as a function of the anchoring condition. GT = greater than.

condition (r = .48), moderate in the GT-1920–1940 condition (r = .29), and small in the other conditions (e.g., r = .16 in the red–blue condition).

We asked people how much they thought their answer and other people's answers to the doctor question were influenced by writing down their ID number. Consistent with research indicating that people think their own judgments are less biased than other people's judgments (Wilson & Brekke, 1994), people in the anchoring conditions reported that they were less influenced than others were, Ms = 5.22 versus 5.43, $t(426) = 3.89$, $p < .001$. When rating the influence of the anchor on their own responses, 79% of the participants circled a 5, which was labeled "have no effect." The correlation between people's ratings of how much they were influenced and their answer to the doctor question was significant, $r(413) = .30$, $p < .001$. Thus, the higher people's estimates of the number of doctors, the more they believed they were influenced by the anchor value. As noted by Nisbett and Ross (1980), however, such correlations do not necessarily demonstrate that people were aware of their own cognitive processes. It is possible that people simply inferred, after the fact, that if their answer to the doctor question was high they must have been influenced somewhat by the anchor.

To examine people's awareness of anchoring effects further, we tested whether there was a significant anchoring effect among the 79% of participants who reported that the anchor value had no effect. This group gave a significantly higher estimate of the number of physicians than people in the control condition did, Ms = 322 vs. 221, $t(409) = 3.01$, $p < .005$, suggesting that there was widespread unawareness of the effects of the anchor (though we cannot rule out the possibility that some participants were aware of these effects).

The results of Study 2 suggest that the key to obtaining basic anchoring effects is the amount of attention people pay to the anchor value. A quick appraisal of an arbitrary number, such as counting its number of digits, does not cause this number to anchor judgments on an unrelated question. Paying slightly more attention to the number, however, such as judging whether it is higher or lower than a closely related number, is sufficient to produce basic anchoring effects.

It is possible, however, that the key is not how much attention people pay to the number but how deeply they process it (Craik & Lockhart, 1972; Jelicic & Bonke, 1991). When anchoring effects occurred in Study 2, not only did people have to attend to the anchor value more, they had to perform computations on the number (such as judging whether it was higher or lower than a closely related number). Perhaps performing computations on a number places it more centrally in people's "workspace," making it more likely to be considered as a possible answer on the next problem encountered. We tested this possibility in Study 3 by crossing a computation manipulation with one that required greater attention to the anchor value.

Study 3

Method

The participants were 110 undergraduate psychology students at the University of Virginia whose reported knowledge about the dependent measure—the percentage of current students at their university who would get cancer in the next 40 years—was below the midpoint; this represented 84% of those people who participated. The procedure was identical to Study 2 except for these changes: We utilized a different cover story, namely that people would complete two questionnaires from two separate research projects, which were being handed out together to save time. The first questionnaire was labeled "Handwriting Analysis Survey." It explained that, although psychologists have been interested in trying to make inferences about people's personalities from their handwriting, little systematic research has been conducted on this topic. Participants were asked to write down some specified numerals or letters to provide a sample of handwriting for initial coding. There were five pages of numerals or letters, which differed by condition. In the five-pages condition, each page had seven numbers, between 4,421 and 4,579. In the one-page condition only the last page contained numbers, between 4,497 and 4,503; the first four pages in this condition contained common words such as "sofa" and "file."

Crosscutting the amount of numbers, we manipulated whether people performed computations on the numbers. In the no-computation condition, we listed every number and people were asked simply to copy it, in their own handwriting, in a space next to the number. In the computation condition only one number was listed at the top of each page. Participants were instructed to write that number down and then to write down a series of numbers that either preceded or followed that number; for example, to "write down the number that comes before the original number." The actual numbers that people wrote down in the computation and no-computation conditions were identical. The only difference was that in the no-computation condition people copied down numbers that were written on the page, whereas in the computation condition people had to perform simple computations (counting up or down from a written value) before writing down the numbers. Finally, we included a control condition in which people copied down five pages of words (i.e., they did not receive a numerical anchor).

After completing the handwriting analysis survey all participants filled out a questionnaire that was ostensibly from the school of medicine and asked general information questions about health. The first question, which was the major dependent variable, asked how many current students at the University of Virginia would be expected to contract cancer within the next 40 years. We included four filler questions as well, such as how many students would be treated for alcohol-related problems and how many students have tested positive for the HIV virus. We also asked the same follow-up and demographic questions we asked in Study 2.

Results and Discussion

As shown in Figure 4, people who received five pages of numbers gave higher estimates of the number of students who would get cancer than did students who received one page or no numbers at all. There was no effect, however, of whether people performed computations on these numbers. A 2 (number of pages of numbers) × 2 (computation) ANOVA revealed a significant main effect of the number of pages, $F(1, 92) = 4.98$, $p = .03$. Neither the main effect of computation nor the interaction was significant, $Fs(1, 92) < 1.32$. Similarly, when the no-number control condition was added to the analyses, a contrast revealed that people who completed five pages of numbers gave significantly higher estimates than people who completed either one page or no pages of numbers, $F(1, 105) = 3.99$, $p < .05$. (This contrast assigned weights of $+3$ to the two-pages cells and -2 to the remaining cells.) The effect size of this contrast was small to moderate, $r = .19$. There was no significant difference between people who completed one page versus no page of numbers, $F(1, 105) < 1$.[5]

Consistent with the results of Study 2, people in the anchoring conditions reported that the anchor value influenced their cancer estimates significantly less than the anchor value influenced the cancer estimates of others, $Ms = 5.15$ versus 5.74, $t(94) = 6.04$, $p < .001$. When rating the influence of the anchor on their own responses, 86% of the participants circled a 5, which was labeled "have no effect." There was a significant correlation between their reports of influence and their estimates of the number of students who would get cancer, $r(93) = .21$, $p < .05$, though as discussed earlier, it is unclear whether this represents an awareness of the anchoring process or a post hoc inference about how the anchor must have influenced them. To explore this further, we compared the magnitude of the anchoring effect among people who reported that the anchor did not influence them. Of these participants, those who completed five pages of numbers gave a significantly higher estimate of the number of students who would get cancer than people who completed one page of numbers gave, $Ms = 3,110$ vs. $1,402$, $t(75) = 2.41$, $p < .02$. Thus, as in Study 2, most people appear to have been unaware that the anchor influenced their responses.

The results of Study 3 illustrate once again that arbitrary numbers can cause anchoring effects, even when people are not asked to consider these numbers as a possible answer to the target question (Hypothesis 1). The results also are consistent with Hypothesis 3, that in order for such anchoring effects to occur people must pay sufficient attention to the arbitrary numbers; copying down five pages of numbers (a total of 35 numbers) anchored people's judgments,

[5] We should note two things about the results of Study 3. First, this was the only study in which there was somewhat of a discrepancy between the means based on the raw scores and those based on the logarithmic transformations of these scores. The pattern of means of the logarithmic scores was the same as shown in Figure 4 except in the five-pages computation cell, which was somewhat lower. It should be recalled, however, that all of the ANOVAs and contrasts were performed on the logarithmic scores, not the raw scores. Thus, the conclusion that anchoring effects occurred in the five-pages condition but not the one-page condition still holds. Second, it is interesting to note that there were no significant effects of the anchoring manipulation on answers to the filler questions. One reason for this finding might be that after people had answered the question about cancer rates, they had two numbers in memory: the manipulated anchor value and their answer to the question about cancer. Having two competing numbers in memory may well dilute anchoring effects. In addition, there are probably optimal anchoring values that differ from people's unanchored estimates but are not so extreme as to be implausible. Jacowitz and Kahneman (1995), for example, chose values that were at the 15th and 85th percentiles of the distribution of a group that received no anchors. Our anchor value in Study 3 was chosen in the same manner; it was at the 85th percentile of pretest participants who answered the question about cancer. It is unlikely that this same number was at as optimal a level to get anchoring effects on the filler questions as it was for the cancer question.

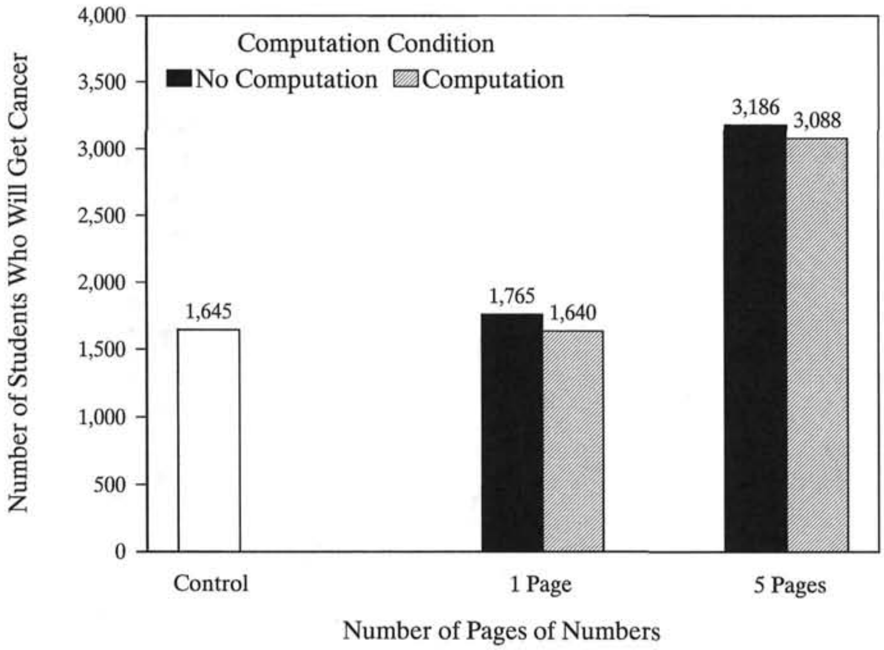

Figure 4. Study 3: Ratings of the number of students who will get cancer in the next 40 years as a function of the anchoring condition.

whereas copying down one page (a total of 7 numbers) did not. Further, we found that whether people performed simple computations had no reliable effect on anchoring.

These results do not, of course, prove that performing computations will never matter. Our computation manipulation was relatively weak, and it is possible that deeper processing of numerical values would lead to stronger anchoring effects. We can say, however, that computation is not a necessary condition for anchoring effects. In our no-computation conditions people simply copied down numbers that were typed on a page, and this was sufficient to produce basic anchoring effects.

The results of Studies 2 and 3 suggest that people are not very cognizant of anchoring effects. When asked how much their answers were influenced by the anchor, people gave relatively low estimates. They recognized that anchor values can influence other people's responses, but reported that they, themselves, were less influenced. These results are consistent with the assumption that anchoring effects are unintentional and nonconscious, which is the rationale for Hypothesis 4, that these effects are difficult to avoid. Studies 4 and 5 tested Hypothesis 4 more directly. In Study 4 we provided people with an incentive to avoid an anchoring effect, to make sure that people were motivated to do so. In Study 5 we gave people detailed (one might even say heavy-handed) warnings about anchoring effects. We hypothesized that neither attempt to prevent anchoring effects would be successful, because anchoring processes operate at a level outside of awareness, making it difficult to know how much one has been influenced by an anchor value.

Study 4

Method

The participants were the 58 members of an undergraduate social psychology class at the University of Virginia whose reported knowledge about the number of physicians in Charlottesville was below the midpoint. This represented 80% of those people who participated. The procedure was identical to that of Study 2, in which people judged the number of physicians and surgeons listed in the Yellow Pages of the phone book, except for the following changes. Participants were told that we would award a prize—dinner for two at a local restaurant—to the person whose answer to one of the questions was most accurate. In the incentive condition, people were told that the prize question was the dependent variable, namely the number of physicians. In the no-incentive condition, people were told that the prize question was a filler item that followed the one about physicians for all participants, namely how many different varieties of rice are grown in the world. Crosscutting the incentive manipulation, half of the people received an anchor on an adhesive note and judged whether it was greater than, equal to, or less than the number of physicians. The others did not receive an anchor. At the conclusion of the study we awarded $50 to the person whose answer to the question about physicians was most accurate.

Results and Discussion

As shown in Figure 5, people who received an anchor once again gave significantly higher estimates of the number of physicians in the phone book. There was no evidence that providing an incentive to be correct moderated the effect of the anchor. A 2 (anchor) × 2 (incentive) ANOVA

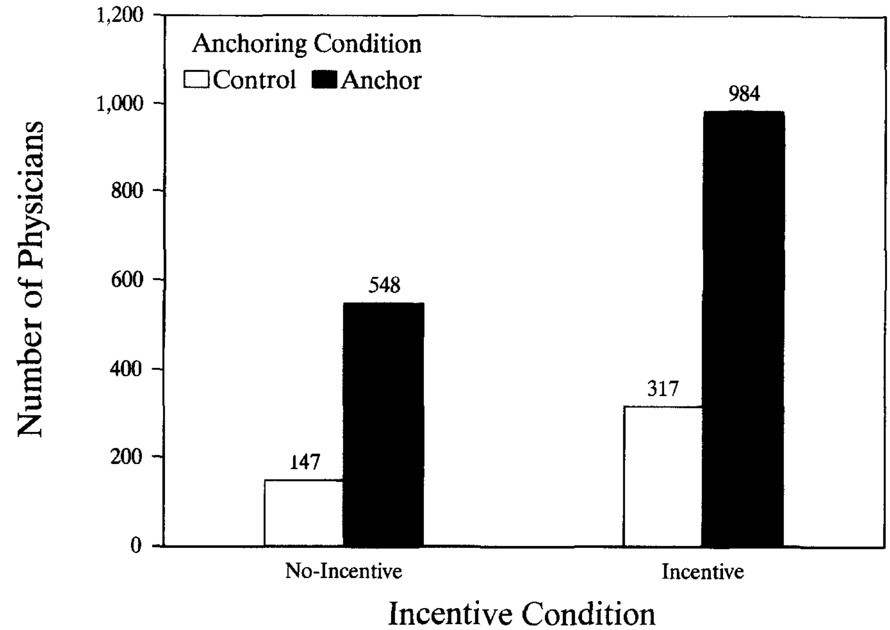

Figure 5. Study 4: Ratings of the number of physicians in the yellow pages as a function of the anchoring condition and the offer of a prize for being the most accurate.

revealed a strong effect of the anchoring manipulation, $F(1, 52) = 13.16$, $p = .001$, $r = .45$, but no Anchoring × Incentive interaction, $F(1, 52) < 1$. This result, we should note, is consistent with a study of Chapman and Johnson's (1995) that was conducted concurrently with this one: Offering people monetary incentives had no influence on the magnitude of anchoring effects. Interestingly, people who received the incentive gave higher estimates in both the control and anchoring conditions, as reflected by a significant main effect of incentive, $F(1, 52) = 4.41$, $p = .04$, $r = .28$. There were no significant effects of the incentive or the anchor manipulations on the filler question about the number of varieties of rice grown in the world, $Fs(1, 53) < 2.37$, $ps > .12$. (See Footnote 5 for possible explanations of the failure to find anchoring effects on filler questions.)

Of further interest is that when no incentive was offered, people in the anchoring conditions reported that they were more influenced by the anchor than were other people, $Ms = 5.75$ and 5.33, respectively. In the incentive conditions, however, people reported that they were less influenced by the anchor than other people, $Ms = 5.18$ and 5.63. A 2 (incentive) × 2 (rating self vs. others) between–within ANOVA revealed a significant Incentive × Rating interaction, $F(1, 21) = 5.18$, $p = .03$. In other words, when a prize was offered people reported that they were less influenced by the anchor (relative to others) than when no prize was offered, even though there was no evidence that the incentive moderated the anchoring effect.

We should note that the results in the no-incentive condition are the only case, in this series of studies, in which people believed that the anchoring manipulation would affect themselves more than others. It should be recalled that all participants were offered a prize for being accurate; the only difference was whether the prize was offered for accuracy on the main dependent measure (the question about doctors in the incentive condition) or the filler question (the question about varieties of rice in the no-incentive condition). We can only speculate that offering a prize suggested to people that answers to the questions are often inaccurate or biased, making them more likely to consider the fact that their answers might be influenced by the anchor. When the prize was offered for accuracy on the main dependent measure (in the incentive condition), however, people appeared to have exerted an extra effort to be accurate. Interestingly, this led to the perception that they were less influenced by the anchoring manipulation but not to an actual reduction in the effect of the anchor.

As in the previous studies there was a moderate correlation between people's belief that they were influenced by the anchor and their estimate of the target (the number of doctors), $r(20) = .31$. This correlation was not significant, however, possibly due to the smaller sample size in Study 3 than in Studies 1 and 2. Only 5 of the 22 people in the anchoring conditions reported that the anchor had no effect; thus we cannot estimate reliably how much this subgroup was influenced by the anchor.

Unexpectedly, there was a main effect of the prize manipulation: People who were offered a prize for accuracy on the doctor question reported that there were more doctors than did people who were offered a prize for the filler question. Perhaps the extra attention or effort to be accurate caused people to think more about the question and to decide that their initial estimates were too low. The important point for our purposes is that this extra effort did not eliminate the effect of the anchor, presumably because anchoring processes occur unintentionally and nonconsciously, even when people are trying to avoid them. As suggested earlier, it is very difficult to fix a mental process that people cannot directly observe.

What if people were informed more directly about the nature of this mental process? As argued by Wilson and Brekke (1994), the motivation to avoid unwanted influences is a necessary but not sufficient condition to avoid biased responses. People must also be aware that their responses are biased, be aware of the direction and the magnitude of the bias, and have the ability to correct for this bias. In Study 4 people were not informed about the nature of anchoring effects or how they influence responses. We mentioned nothing about avoiding the influence of the anchor value, much less the form this influence might take (i.e., raising or lowering people's answers to the dependent measure). In Study 5 we used a variety of forewarning manipulations in which people were informed about the nature of anchoring effects and asked to avoid being influenced by them. We also delivered these manipulations at a variety of points in the process (e.g., before receiving the anchor, after receiving it but before comparing it to the target question, and after comparing it to the target question) to see if being forewarned earlier in the process makes it easier to avoid anchoring effects.

Such manipulations come closer to satisfying Wilson and Brekke's (1994) stated conditions for avoiding bias. They still may not be successful, however, at completely undoing anchoring effects because they leave one condition unsatisfied: Even if people are informed about anchoring effects and the direction in which an anchor can influence their responses, they still do not know the magnitude of the effect or precisely how much to adjust their responses to undo the effects of an anchor. Thus, we predicted that the forewarning manipulations would not eliminate anchoring effects.

Study 5

Method

Participants were the 408 undergraduate psychology students at the University of Virginia whose reported knowledge about the number of physicians in Charlottesville was below the midpoint. This represented 90% of those people who participated. The procedure was identical to that of Study 4, except for the following changes: No prizes or incentives were offered for being correct. We included our standard control condition in which people did not receive an anchor and our standard comparison condition in which people received an anchor and were asked to compare it to the number of doctors in the phone book. We included several additional conditions that were identical to the anchoring condition, except that we informed people about the nature of anchoring effects and asked them not to be influenced by the anchor.

We varied both the nature of our warning and the point at which it was delivered by randomly assigning participants to one of the following conditions: In the contamination condition we provided written instructions, right after the general instructions but before giving people the anchor value, indicating that "a number in people's heads can influence their answers to subsequent questions." We gave a hypothetical example of such an effect (unrelated to the question about the number of physicians in the phone book) but did not specify its direction, that is, whether numbers increase or decrease people's answers to subsequent questions. People then read the following instructions:

> When you answer the questions on the following pages, *please be careful not to have this contamination effect happen to you.* We would like the most accurate estimates that you can come up with.

In the underestimation condition these instructions were identical, except that we informed people (wrongly) that large numbers in their heads cause them to decrease their estimates of answers to subsequent questions. In the overestimation-general condition these instructions were identical, except that we informed people (correctly) that large numbers in their heads cause them to increase their estimates of answers to subsequent questions. In the overestimation-specific condition the instructions were identical to the preceding ones, except that we gave specific instructions about how this might influence their answers. We told people that they would be asked to write down an ID number we provided and that this ID number might influence their answers to subsequent questions. In other conditions people received these same instructions, but at different points. In the after-ID condition, they did not receive the overestimation-specific instructions until after they had received and copied down their ID number. In the after-comparison condition, they did not receive these instructions until after they had received their ID number and answered whether it was greater or less than the number of doctors in the phone book (but before they gave their exact estimate). In the after-initial-estimate condition, people did not receive these instructions until after making their specific estimate of the number of doctors in the phone book. They were then given the opportunity to change their answer to this question.

Results and Discussion

As shown in Figure 6, our forewarning manipulations, regardless of their form or placement, were unsuccessful. A one-way ANOVA on people's estimates of the number of doctors revealed a highly significant overall effect, $F(8, 399) = 6.75$, $p < .0001$. A Newman–Keuls test of the difference between pairs of means revealed that every anchoring condition differed significantly from the control condition at $p < .05$. None of the anchoring conditions differed significantly from each other, regardless of whether people were forewarned about anchoring effects, how they were forewarned, or when they were forewarned.

There were indications that the forewarning manipulations influenced people's responses to some degree, although not sufficiently to undo the anchoring effects completely. One such indication comes from the after-initial-estimate condition, in which people gave two estimates of the number of doctors in the phone book: one before being informed about anchoring effects and one immediately after. The mean estimate shown in Figure 6 is the revised estimate people gave after being warned. If we compare this mean to their estimate before being warned, $M = 745$, we can see that people did lower their estimates significantly (although insufficiently), $t(44) = 4.41$, $p < .001$. Second, when people were told that the anchor value would decrease their estimates (in the underestimation condition), they gave a relatively high answer, $M = 641$. When they were told

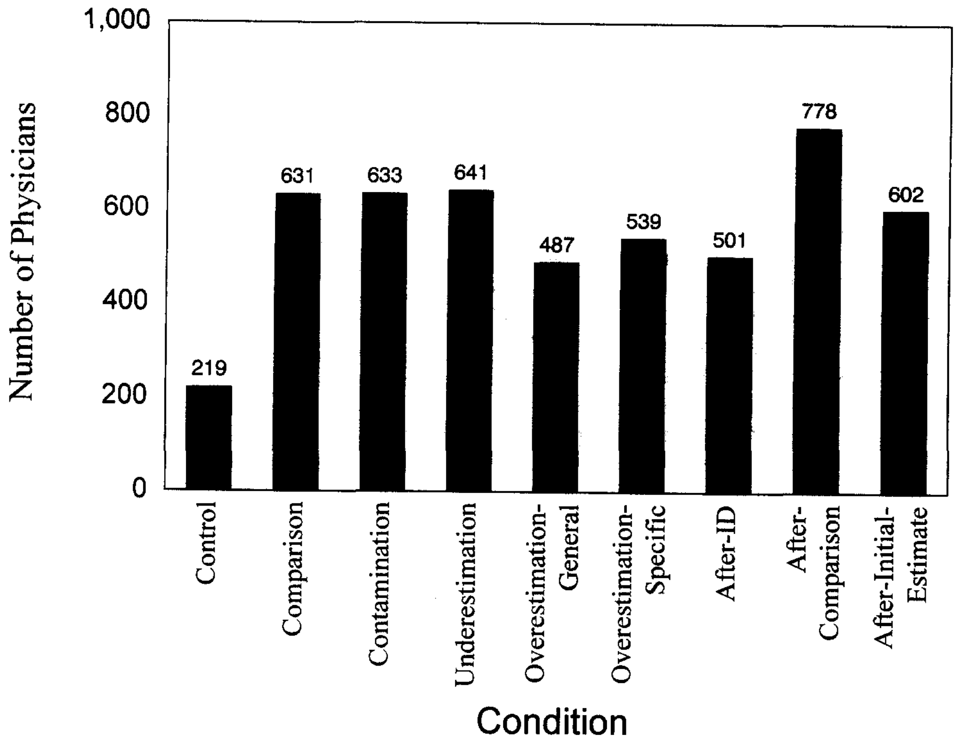

Figure 6. Study 5: Ratings of the number of physicians in the yellow pages as a function of the anchoring condition and forewarning.

that the anchor value would increase their estimate, before comparing the anchor to the doctor question (in the overestimation, overestimation-specific, and overestimation-after-ID conditions), their estimates were lower (*M*s = 487, 539, and 501, respectively). As noted earlier, a post hoc Newman–Keuls comparison indicated that there was no significant difference between any of the means in the anchoring conditions. A focused contrast, however, that assigned a weight of +3 to the underestimation condition and a −1 to the overestimation, overestimation-specific, and overestimation-after-ID conditions, was significant, $F(1, 379) = 4.73, p < .05$.

It is possible that these changes in people's responses were due to demand characteristics, whereby people felt compelled to raise their answer when told that the anchor value would decrease their estimates and to lower their answer when told that the anchor value would increase their estimates. Given that responses were anonymous, however, we suspect that people genuinely tried to give unbiased responses and altered their responses to some degree in the forewarning conditions. It is apparent from Figure 6, however, that these adjustments were inadequate. We suggest that because anchoring effects occur unintentionally and unconsciously, it was difficult for people to know the extent to which an anchor value influenced their estimates. As a result, they were at the mercy of their naive theories about how susceptible they were to anchoring effects (Petty & Wegener, 1993; Wilson & Brekke, 1994).

Consistent with the prior studies, people's theories underestimated the amount of bias, in that people reported that they were rather immune to anchoring effects. The mean amount of reported effect of the anchor value on one's own responses was 5.23, which was close to the midpoint of 5, labeled "have no effect." People reported that others were influenced significantly more by the anchor than they were, *M*s = 5.53, $t(361) = 4.78, p < .001$. As in the previous studies there was a moderate correlation between ratings of how much they were influenced by the anchor and their estimate of the number of doctors, $r(365) = .26, p < .001$; however, once again there were signs of inaccuracy among those who reported that the anchor had no effect: These participants gave a significantly higher estimate of the number of doctors than people in the control condition, *M*s = 365 versus 219, $t(193) = 2.30, p < .05$.

Interestingly, people's reports about the influence of the anchor on their own responses did not differ according to which forewarning condition they were in, $F(7, 359) < 1$, *ns*. Condition did influence, however, people's reports about the influence of the anchor on the responses of others, $F(7, 354) = 3.11, p = .003$. People's lowest estimate, not surprisingly, was in the underestimation condition, in which they were told that numbers cause people to lower their estimates of numerical values ($M = 4.92$). The highest estimates were in the conditions in which they were forewarned relatively late, namely the after-comparison and after-initial-estimate conditions, *M*s = 5.98 and 6.06, respectively. Thus, our forewarning manipulations were successful in convincing participants that people in general are susceptible in anchoring effects, but not very successful in convincing participants that they themselves were susceptible.

General Discussion

The present studies were concerned with the conditions under which anchoring effects occur. Each of our four hypotheses was supported. First, we found that completely arbitrary numbers can anchor people's judgments, even when there is no logical reason to consider these numbers as answers to the target questions. As far as we know these are the first demonstrations of basic anchoring effects. Second, we found in Study 1 that the amount of knowledge people have about the target question moderated these effects. Not surprisingly, people who were knowledgeable were uninfluenced by an arbitrary anchor, presumably because they could retrieve from memory the answer they believed to be correct. Third, we found that people must pay sufficient attention to an arbitrary number for basic anchoring effects to occur. Fourth, anchoring processes appear to occur unintentionally and nonconsciously, in that most people reported that they were uninfluenced by the anchor, and neither the offer of an incentive to be accurate (Study 4) nor forewarning people about anchoring effects (Study 5) eliminated these effects.

We should note that there is still some ambiguity about Hypothesis 2 concerning the amount of attention to an anchor that is necessary in order for basic anchoring effects to occur. How much attention is sufficient? In Study 2 anchoring effects occurred when people judged whether their ID number was greater than 100. In Study 3, however, anchoring effects did not occur in the condition in which people copied down one page of seven numbers, which would seem roughly equivalent in attention, if not greater than, judging whether one number is greater than 100.

This inconsistency may be more apparent than real. In Study 3, in which copying one page of numbers did not produce an anchoring effect, there was more of a psychological separation between the manipulation and the dependent measures. People believed that they were copying numbers as part of a study of handwriting and answering the question about cancer as part of an unrelated health survey. In Study 2, in which judging whether the anchor value was greater than 100 produced an anchoring effect, people believed that the manipulation was part of the same study. It may be that when the tasks are separate in an individual's mind that numbers used in the first task are less likely to come to mind when thinking of possible answers on the second task, unless a considerable amount of attention has been paid to the numbers on the first task. That is, when people answer numerical questions it may be that numbers considered most recently on a similar task are more likely to come to mind than numbers considered as part of what was thought to be a separate task.

This account is admittedly speculative and points to aspects of anchoring processes that remain unclear. Exactly what are the conditions under which basic anchoring effects will occur? Under what conditions will a number in short-term memory be considered as a possible answer to an unrelated question? We found that the amount of attention paid to prior numbers makes a difference, but the exact parameters of attention—how much is enough and whether the perceived relevance of the task makes a difference—are not entirely clear. Kahneman and Knetsch (1993) and Chapman and Johnson (1994) found that another moderator is the format of the number in short-term memory: People were more likely to use anchor values that were in the same format (e.g., dollar amounts) than numbers in a different format (e.g., percentages). Investigating such moderator variables further is clearly important.

Another possible limitation of the present studies is that we used high anchor values in all studies, designed to increase people's estimates to the target questions. It is possible that the processes we have investigated are limited to increases resulting from high anchor values and do not apply to decreases resulting from low anchor values. Anchoring effects resulting from low values are common in the literature, however (e.g., Jacowitz & Kahneman, 1995; Quattrone et al., 1984), and there is no reason to assume an asymmetry in basic anchoring effects. Jacowitz and Kahneman (1995) found that high anchors had larger effects than low anchors, but significant effects were found in both directions. (One reason the effects are larger for high anchors might be because of "floor effects" with low anchors; for many questions, there is less room to go down from one's initial estimate than there is to go up.)

Why do basic anchoring effects occur? As mentioned, there are two versions of a priming hypothesis that can account for the results. Kahneman and Knetsch (1993) advocate a backward priming hypothesis, whereby the need to answer a question invokes a search for possible answers. Once this search is initiated, any plausible value in short-term memory is considered as a possible answer, triggering the anchoring process (see Figure 1). Second, it is possible that a more standard priming mechanism is responsible for the results, whereby the activation potential of the anchor value is increased as soon as people attend to it (Higgins, 1996).

The present studies were not designed to distinguish between a standard and backward priming interpretation of anchoring. Nonetheless the similarities between the present results and standard priming results are worth noting, because there is a strong prima facie similarity between basic anchoring and standard priming effects. Consider the well-known case whereby people assimilate their judgments of a person to a trait category that has been arbitrarily primed (e.g., Higgins, 1996; Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1979). One way of describing anchoring effects would be to say that people's numerical estimates are assimilated to numerical values that have been primed, just as people's impressions of others are assimilated to trait categories that have been primed. Another way of describing priming effects would be to say that when people are asked to form an impression of someone, they call to mind as possible answers trait constructs that are in short-term memory, due to a priming manipulation, and their impressions are influenced by this anchor value.

It is worth noting, however, an important difference between the results of priming studies and the present studies. Priming effects (specifically, the assimilation of a judgment to a primed category) are most likely to occur when people are unaware that the accessible information has been primed arbitrarily. If people are aware of the source of the primed information, they are more likely to try to avoid being influenced by it, which often results in an overcorrection (i.e., contrast effects; see Lombardi, Higgins, & Bargh, 1987; Martin, 1986; Newman & Uleman, 1990; Strack, Schwarz, Bless, Kubler, & Wanke, 1993). Researchers in priming studies often go to great lengths to disguise the connection between the priming manipulation and the judgment, to ensure that assimilation effects occur.

We found anchoring effects even when people were blatantly provided with anchor values and explicitly told not to use these values when answering subsequent questions. People could hardly have been made more aware of the potential bias of the anchor values in Study 5, in which we explicitly forewarned people about anchoring effects and instructed them to avoid them. Assimilation effects still occurred, whereby people's responses were in the direction of the anchor values. By comparison, Strack et al. (1993) found that people who were reminded about the priming task right before the judgment task showed contrast effects, whereby they adjusted their impressions too far away from the primed trait.

One explanation for these different findings is as follows: When people are aware of a potential bias on their judgments, the way in which they try to correct their judgments depends on their naive theories about the direction and magnitude of that bias (Petty & Wegener, 1993; Wilson & Brekke, 1994). It is possible that people have different naive theories about the effects of numerical anchors versus semantic primes. When they become aware of primes, they believe they are in danger of an assimilation effect and thus they adjust their impressions (often too far) in the opposite direction. When they become aware of numerical anchors, they appear not to adjust much at all, possibly because their naive theories are that their judgments could not possibly be influenced by unrelated, arbitrary numbers. Consistent with this argument, most participants in our studies reported that the anchor value had very little or no influence on their judgments, even when substantial anchoring effects were found.

We cannot establish at this point whether anchoring and priming effects are due to the same mechanisms (or whether standard or backward priming mechanisms are involved). It seems clear, however, that anchoring should be studied in the broader context of research on judgment and inference. The processes depicted in Figure 1 are similar to several recent models of the way in which people attempt to correct their judgments (e.g., Gilbert, 1991; Martin, 1986; Petty & Wegener, 1993; Schwarz & Bless, 1992; Strack, 1992). Gilbert (1991), for example, has proposed a model of belief formation that holds that people are predisposed to believe everything they initially comprehend. When people have reason to believe that a belief is false, they attempt to "unaccept" the belief with an effortful correction process.

An extension of this model to social comparison processes by Gilbert, Giesler, and Morris (1995) is particularly relevant to basic anchoring effects. Much previous research on social comparison has focused on the target of social comparison: Under what conditions, for example, do people engage in downward social comparison, choosing to compare themselves to people who are performing at a lower level than themselves (e.g., Wills, 1981; Wood, 1989)? Gilbert et al. (1995) suggested that the comparison process is more spontaneous and unintentional than previously thought. When in the presence of others, we spontaneously compare our performance to theirs, they suggested, even if the others' performance is irrelevant to our own. Only afterwards do we think about how appropriate the comparison was and, if necessary, engage in an effortful correction process whereby we try to undo or ignore the comparison.

The similarity of anchoring processes to social comparison processes is striking: In both cases there is an initial comparison between two values (i.e., between an anchor and a target value or another person and oneself) that seems to occur spontaneously and unintentionally. In both cases this comparison process occurs even when the point of comparison (the anchor value or another person's performance) is completely arbitrary and uninformative.

Interestingly, in both models the conditions under which a stimulus is spontaneously used as a point of comparison are not entirely clear. Gilbert et al. (1995) pointed out that we do not compare ourselves to every person we meet, just as we argue that a person's judgment is not anchored by every number encountered. Gilbert et al. suggested that social comparison processes are most likely to occur when people explicitly judge the performance of another person, when the other person is in close spatial and temporal proximity, and when the other person's performance is extreme. These conditions are similar to those that appear to initiate basic anchor effects. In particular, the first one—explicitly judging another person's performance—is similar to our finding that people must pay sufficient attention to a number for basic anchoring effects to occur. Simply observing another person's performance may not be enough to trigger the comparison process, and simply observing a number may not be sufficient to trigger an anchoring process.

Despite these similarities, there are two important differences between social comparison and anchoring processes. First, people's proclivity to adjust for the two different kinds of bias appears to differ. As long as people have sufficient motivation and cognitive capacity, they can adequately adjust for invalid social comparison information. Gilbert et al. (1995), for example, found that when people realized that another person's performance was irrelevant to their own, they were able to correct their judgments adequately. As we have seen, people find it difficult to adjust sufficiently for invalid anchors, even when offered incentives to do so or forewarned about anchoring effects. The reason for this difference may be the same as in our discussion of the priming literature: People have different theories about the magnitude and direction of social comparison than they do about anchoring effects. Our naive theories tell us that we might have inadvertently compared ourselves to someone else in an inappropriate manner and that we should go back and undo this comparison. The present studies imply that there is no such theory about the effects of numeric anchors. People do not believe that they are influenced by arbitrary anchors and see no need to adjust their responses to target questions.

The second difference is, perhaps, more fundamental: The initial judgment that needs correcting is in a different direction in social comparison versus anchoring. Social comparison processes typically result in contrast effects, whereby the assessment of our own performance is in a direction away from another person's (e.g., seeing another person do well lowers our assessments of our own abilities). Anchoring almost always results in assimilation effects, whereby people's answers to a target question are in a direction toward the anchor value (e.g., being presented with a high anchor increases one's estimates of the percentage of students who will get cancer). Given this important difference in the direction of the initial effect, it would be premature to conclude that anchoring and social comparison processes are similar.

Clearly, much of this discussion about the differences between anchoring, priming, and social comparison is speculative, and it would be premature to make firm conclusions about the similarities and differences. Based on the results of the present studies we can conclude that anchoring effects are more general than previously believed. If people pay at least a minimal amount of attention to completely arbitrary numbers, these numbers can anchor numerical answers to unrelated questions, even when there is no logical reason for people to use the numbers. These effects were not limited to relatively unimportant general knowledge questions, but were also found on questions about the risk of getting life-threatening diseases. Further, motivating people to avoid the effects or forewarning them did not prevent the effects from occurring. The methodological implications of these results are also clear: Researchers should be wary of asking their participants to attend to a numerical value and then give a numerical estimate on an unrelated question.

References

Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions.* Newbury Park, CA: Sage.

Cervone, D., & Peake, P. (1986). Anchoring, efficacy, and action: The influence of judgmental heuristics on self-efficacy judgments and behavior. *Journal of Personality and Social Psychology, 50,* 492–501.

Chapman, G. B., & Johnson, E. J. (1994). The limits of anchoring. *Journal of Behavioral Decision Making, 7,* 223–242.

Chapman, G. B., & Johnson, E. J. (1995). *Anchoring, confirmatory search, and the construction of values.* Manuscript submitted for publication.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11,* 671–684.

Davis, H. L., Hoch, S. J., & Ragsdale, E. K. E. (1986). An anchoring and adjustment model of spousal predictions. *Journal of Consumer Research, 13,* 25–37.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal representation of human judgment* (pp. 17–52). New York: Wiley.

Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance, 1,* 288–299.

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist, 46,* 107–119.

Gilbert, D. T., Giesler, R. B., & Morris, K. A. (1995). When comparisons arise. *Journal of Personality and Social Psychology, 69,* 227–236.

Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds.), *Syntax and semantics: Speech acts* (Vol. 3, pp. 41–58). New York: Academic Press.

Helson, H. (1964). *Adaptation level theory: An experimental and systematic approach to behavior.* New York: Harper.

Hershey, J. C., & Schoemaker, P. J. H. (1985). Probability versus certainty equivalence methods in utility measurement: Are they equivalent? *Management Science, 31,* 1213–1231.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. Kruglanski (Eds.), *Social Psychology: Handbook of basic principles,* (pp. 133–168). New York: Guilford.

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology, 13,* 141–154.

Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin, 21,* 1161–1166.

Jelicic, M., & Bonke, B. (1991). Levels of processing affects performance on explicit and implicit memory tasks. *Perceptual and Motor Skills, 72,* 1263–1266.

Joyce, E., & Biddle, G. (1981). Anchoring and adjustment in probabilistic inference in auditing. *Journal of Accounting Research, 19,* 120–145.

Kahneman, D., & Knetsch, J. (1993). *Strong influences and shallow inferences: An analysis of some anchoring effects.* Unpublished manuscript, University of California, Berkeley.

Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization: Assimilation versus contrast as a function of awareness of the priming task. *Personality and Social Psychology Bulletin, 13,* 411–429.

Lopes, L. L. (1985). Averaging rules and adjustment processes in Bayesian inference. *Bulletin of the Psychonomic Society, 23,* 509–512.

Lovie, P. (1985). A note on an unexpected anchoring bias in intuitive statistical inference. *Cognition, 21,* 69–72.

Martin, L. L. (1986). Set/reset: Use and disuse of concepts in impression formation. *Journal of Personality and Social Psychology, 51,* 493–504.

Newman, L. S., & Uleman, J. S. (1990). Assimilation and contrast effects in spontaneous trait inference. *Personality and Social Psychology Bulletin, 16,* 224–240.

Nisbett, R. E., & Ross, L. D. (1980). *Human inference: Strategies and shortcomings of social judgment.* Englewood Cliffs, NJ: Prentice Hall.

Northcraft, G. B., & Neale, M. A. (1987). Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes, 39,* 84–97.

Parducci, A. (1968). The relativism of absolute judgment. *Scientific American, 219,* 84–90.

Petty, R. E., & Wegener, D. T. (1993). Flexible correction processes in social judgment: Correcting for context-induced contrast. *Journal of Experimental Social Psychology, 29,* 137–165.

Plous, S. (1989). Thinking the unthinkable: The effect of anchoring on likelihood estimates of nuclear war. *Journal of Applied Social Psychology, 19,* 67–91.

Quattrone, G. A. (1982). Overattribution and unit formation: When behavior engulfs the person. *Journal of Personality and Social Psychology, 42,* 593–607.

Quattrone, G. A., Lawrence, C. P., Finkel, S. E., & Andrus, D. C. (1984). Explorations in anchoring: The effects of prior range, anchor extremity, and suggestive hints. Unpublished manuscript, Stanford University, Stanford, California.

Russo, J. E., & Schoemaker, P. J. H. (1989). *Decision traps: Ten barriers to brilliant decision making and how to overcome them.* New York: Doubleday.

Schlenker, B. R. (1980). *Impression management: The self-concept, social identity, and interpersonal relations.* Monterey, CA: Brooks/Cole.

Schwarz, N., & Bless, H. (1992). Constructing reality and its alternatives: An inclusion/exclusion model of assimilation and contrast effects in social judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 217–245). Hillsdale, NJ: Erlbaum.

Schwarz, N., & Wyer, R. S. (1985). Effects of rank ordering stimuli on magnitude ratings of these and other stimuli. *Journal of Experimental Social Psychology, 21,* 19–29.

Slovic, P., & Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance, 6,* 649–744.

Srull, T. K., & Wyer, R. S., Jr. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology, 37,* 1660–1672.

Strack, F. (1992). Sensation- and information-based social judgments: Toward a general model. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 249–275). Hillsdale, NJ: Erlbaum.

Strack, F., & Mussweiler, T. (1995). *The enigmatic anchoring effect: A pervasive phenomenon in search of a viable explanation.* Unpublished manuscript. University of Trier, Germany.

Strack, F., Schwarz, N., Bless, H., Kubler, A., & Wanke, M. (1993). Awareness of the influence as a determinant of assimilation versus contrast. *European Journal of Social Psychology, 23,* 53–62.

Switzer, F., & Sniezek, J. A. (1991). Judgment processes in motivation: Anchoring and adjustment effects on judgment and behavior. *Organizational Behavior and Human Decision Processes, 49,* 208–229.

Tesser, A. (1978). Self-generated attitude change. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 11, pp. 289–338). New York: Academic Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185,* 1124–1131.

Wills, T. A. (1981). Downward social comparison principles in social psychology. *Psychological Bulletin, 90,* 245–271.

Wilson, T. D., & Brekke, N. C. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116,* 117–142.

Winer, B. J. (1971). *Statistical principles in experimental design* (2nd ed.). New York: McGraw-Hill.

Wood, J. V. (1989). Theory and research concerning social comparison of personal attributes. *Psychological Bulletin, 106,* 231–248.

Zuckerman, M., Koestner, R., Colella, M. J., & Alton, A. O. (1984). Anchoring in the detection of deception and leakage. *Journal of Personality and Social Psychology, 47,* 301–311.

Received February 3, 1995
Revision received November 14, 1995
Accepted March 4, 1996 ■