

6. Incorporating the Irrelevant: Anchors in Judgments of Belief and Value

Gretchen B. Chapman and Eric J. Johnson

Imagine walking down a supermarket aisle and passing an end-of-aisle display of canned tomato soup. A sign on the display says, "Limit 12 per customer." Would such a sign influence the number of cans you would buy? Would you buy more cans than if the sign said "No limit per customer"? Our intuitions say no, but empirical evidence indicates that purchase behaviors are influenced by such a sign (Wansink, Kent, & Hoch, 1998). Consider another example: A wheel of fortune is spun and stops at the number 65. You are then asked if the percentage of African countries in the United Nations is above or below that number. Could this exercise influence your estimate of the relevant percentage? Although it may seem unlikely, the evidence is that such anchors have an effect: Groups who received larger numbers determined by a wheel of fortune gave higher estimates than groups who received lower numbers, demonstrating that irrelevant anchors influenced these estimates (Tversky & Kahneman, 1974).

"Anchoring and adjustment" is one of three well-known heuristics described by Tversky and Kahneman (1974) in a classic paper that also describes the representativeness and availability heuristics. Like the other heuristics, anchoring and adjustment can be a useful way of making judgments. Imagine that you are trying to set a value on an antique chair that you have inherited from a distant aunt. You might recall seeing a very similar chair in slightly better condition at a local antique dealer. You might start with that price as an anchor, and incorporate the difference in quality. This seems to be a useful and effort-saving use of the anchoring and adjustment heuristic. Now, however, imagine that you had seen (on Public Television's *Antiques Road Show*) a not-so-similar chair that, unlike yours, is signed by the designer and worth, as a result, many thousands of dollars more. If you were to use this as an anchor, and if you did not properly incorporate the fact that your chair did not have a signature, you might end up with an estimate that was too high, or biased. Thus anchoring can be a useful heuristic, but it can also result in biased answers.

Research suggests that people use an anchor-and-adjust strategy to solve a variety of estimation problems. For example, Rottenstreich and Tversky (1997) proposed that when judging the likelihood of a disjunctive event

This work was supported in part by NSF grant SBR97-96042 to the first author.

(e.g., the likelihood of being a chemistry or biology major), people anchor on an estimate for one event (e.g., chemistry major) and adjust to take into account the other event as well. Similarly, Kruger (1999) suggested that when answering questions such as, “How does your driving ability compare to that of your peers?” people first anchor on their own abilities, and then adjust for the skills of their peers. Griffin and Tversky (1992) proposed that when making confidence judgments, people anchor on the extremeness of the evidence confronting them and then adjust for the credibility of the evidence. In all these cases, adjustment is often insufficient, resulting in a bias.

Tversky and Kahneman (1974, p. 1128) presented anchoring as a process in which “people make estimates by starting from an initial value that is adjusted to yield a final answer [and] . . . adjustments are typically insufficient.” Notions of anchoring were first introduced to decision-making research in early descriptions of preference reversals (Slovic, 1967; Slovic & Lichtenstein, 1968; Lichtenstein & Slovic, 1971). In judging the attractiveness of a gamble, the gamble attribute most compatible with the response mode seemed to be an anchor. For example, in pricing a gamble, subjects would anchor on the monetary outcome of the gamble and make adjustments from there.

The concept of anchoring and adjustment has had widespread impact. However, the mechanisms of anchoring have been systematically explored only recently. In this chapter, we review what is currently known about the causes and effects of anchoring. We start by offering some definitions, and then identify some stylized facts about this heuristic. We next examine two families of causes of anchoring. We close by reviewing other phenomena related to anchoring and potential applications.

ANCHORING DEFINITIONS

Because it has been used in many different areas, the term *anchoring* has been used to mean somewhat different things. We group these definitions into three types: One refers to an *anchoring procedure* in which a salient but uninformative number is presented to subjects. A second meaning is an *experimental result*, in which the uninformative number influences the judgments. Finally, anchoring and adjustment is sometimes used to refer to the *psychological process* by which the uninformative number has its effect.

Our discussion of anchoring might benefit from some standardized nomenclature that emphasizes these distinctions. We define an *anchoring procedure* as one in which a salient but uninformative number is presented to subjects before they make a numeric judgment. Most anchoring studies follow a two-step procedure introduced by Tversky and Kahneman (1974). Subjects are first asked to compare the anchor to the target value, stating whether the anchor is higher or lower than the target. For example, subjects are asked whether the percentage of African countries in the United Nations is more or less than 10%. Second, subjects are asked to give a numeric estimate of the target – for example, to state

the percentage of African countries in the United Nations. Some anchoring studies (e.g., Wilson, Houston, Brekke, & Etling, 1996) use other procedures that do not include the initial comparison judgment. Differences in procedure may be important, because similar effects obtained with different procedures may not necessarily represent the same phenomenon or underlying mechanism.

All anchoring procedures involve presentation of an anchor. We concentrate on numeric anchors that are uninformative but salient to the decision maker. Thus, a number can be identified as an anchor before looking to see whether it influences judgment. There are two reasons for focusing on uninformative anchors. First, the influence of uninformative anchors is clearly a bias. If respondents report that an anchor is irrelevant to the judgment at hand, yet it influences their judgment, it is hard to argue that this reflects the rational use of relevant information. A second reason for focusing on uninformative anchors is to rule out one potentially uninteresting cause of anchoring effects. Subjects might attend to anchors and incorporate them into their judgments because they reason that the experimenter would not have mentioned the anchor were it not informative or relevant (Grice, 1975). This reduces the anchoring effect to a sort of demand effect (e.g., Schwarz, 1994; Schwarz & Bless, 1992; Sudman, Bradburn, & Schwarz, 1996). In order to avoid this type of explanation, a number of investigators use anchors that are obviously uninformative. For example, Russo and Shoemaker (1989, p. 90) asked participants to estimate the year Attila the Hun was defeated in Europe after considering an anchor constructed from their phone numbers. Because these numbers are in no plausible way related to the year of Attila's defeat, any influence is clearly an unwanted bias. Similarly, the output of a wheel of fortune is not plausibly predictive of the membership of African countries in the United Nations.

Informative numbers might be anchors. When anchors are informative, experiments that show anchoring often rely on demonstrating order effects – specifically, that earlier items of information receive more weight. For example, Kahneman and Tversky (1974) used a multiplication problem and contrasted estimates of two orders of the same product: $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$ and $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$. Here, the product of the first few numbers is relevant information with respect to the final estimate. However, the disparity of the two estimates, which depends on the order in which the sequences are presented, indicates that the earlier numbers receive too much weight in the estimation of the final product. Demonstrations that use meaningful information as anchors often raise a number of important but auxiliary questions, such as the appropriate weighting of a sequence of information and differences between sequential and simultaneous presentation. Therefore, we concentrate on cases in which anchors are irrelevant to the task at hand. As shown later, even when judges agree that the numbers are irrelevant, they do have an impact.

We define *anchoring* as an experimental result or outcome; the influence of an anchor that renders the final judgment too close to the anchor. Thus, *anchoring* is defined as assimilation rather than contrast (Sherif, Sherif, & Nebergall, 1965;

Sherif & Hovland, 1961). The vast majority of decision-making studies on anchoring have found a positive relation between anchors and judgments. We reserve discussion of the cognitive process underlying anchoring effects for a later section. First, we review key results about the anchoring effect.

NECESSARY CONDITIONS FOR ANCHORING

Not all uninformative numbers produce anchoring effects. Instead, certain features of the anchor, target, and judgment task are required.

Attention to the Anchor

As described earlier, the vast majority of anchoring studies follow a two-step procedure in which an initial comparison task is followed by numeric estimation of the target. This procedure assures that subjects attend to the anchor and compare it to the target. In fact, the initial comparative judgment is not necessary to produce anchoring. Wilson et al. (1996) found that anchoring could be achieved without the comparative judgment, provided that subjects devoted sufficient attention to the anchor. For example, doing five pages of computations on large numbers increased a later target judgment about cancer incidence (relative to a no-anchor control condition), but doing one page of computations did not. Thus, anchoring occurred even when subjects did not explicitly compare the anchor to the target value, but the anchor had to be made very salient by extensive processing. Wilson et al. (1996) suggest that it is necessary for the irrelevant anchor to come to mind as a potential response to the target question.

Anchor–Target Compatibility

In most studies of anchoring, the anchor is an arbitrary number on the same scale as the target response. Several studies have examined whether this condition is necessary for anchoring to occur. For example, Chapman and Johnson (1994) asked subjects to judge the value of various consumer items by indicating the amount of money or additional life expectancy they would demand in exchange for giving up the item. Before specifying an amount, subjects first considered an anchor expressed in dollars or years of life expectancy. Anchoring occurred if the anchor and response were on the same scale (e.g., both life expectancy) but not if they were on different scales (e.g., monetary anchor and life expectancy response). Kahneman and Knetsch (1993) report a similar result. They asked Toronto residents whether they would pay \$25 (low anchor) or \$200 (high anchor) to clean up specified lakes to maintain fish populations. Some subjects were then asked to estimate the amount the average Toronto resident would contribute. These subjects showed an anchoring effect, giving mean estimates of \$14 and \$36 in the low and high anchor conditions, respectively. Other subjects were instead asked to estimate the percentage of Toronto residents who would pay \$100. These subjects did not show an anchoring effect

(estimating 24% and 29% in the low and high anchor conditions, respectively). Thus, anchoring occurred when anchors and responses were both expressed on a dollar scale, but not when anchors were expressed as dollars but responses as percentages.

Strack and Mussweiler (1997) found that it is not sufficient for the anchor and response to be on the same numeric scale; to achieve strong anchoring, anchor and response must also express the same underlying dimension. They asked subjects to estimate the height or width of the Brandenburg Gate after considering a numeric anchor described as the height or width of the gate. Anchoring was much stronger if both the anchor and the target judgment represented the height (or both the width); it was weaker if one was height and the other was width. All anchors and responses were expressed as meters, but this agreement alone did not determine the size of the anchoring effect. To obtain a large anchoring effect, the anchor had to represent the same dimension (width or height) as the target.

In a similar experiment (Wilson et al., 1996, Experiment 1) subjects were asked to estimate the number of countries in the United Nations (UN). Before making this judgment, some subjects were asked to judge whether a high anchor was more or less than the number of countries in the UN, whereas other subjects were asked whether the same anchor was more or less than the number of physicians and surgeons in the phone book. Anchoring occurred both when the anchor had been compared to the target question (about the UN) and when compared to the irrelevant question (about physicians). The switch in questions (UN to physicians) produced an anchoring effect that was numerically but not statistically smaller than the anchoring effect in the no-switch condition. In this study, the anchor and target were both expressed on the same numeric scale (counting numbers), but not the same dimension or quantity (number of physicians or number of countries). The change in dimensions tended to reduce the anchoring effect, but the reduction was not as large as in the Strack and Mussweiler (1997) study.

Extreme Anchors

Several studies have found that anchoring occurs even when the anchors are extreme or represent implausible responses to the target question. Strack and Mussweiler (1997) asked subjects, for example, to estimate the year that Einstein first visited the United States after considering anchors of 1215 or 1992. These implausible anchors produced anchoring effects just as large as more plausible anchors (e.g., 1905, 1939). Similarly, Chapman and Johnson (1994) asked subjects to indicate minimum selling prices for monetary lotteries after considering a monetary anchor. Some of the anchors were higher than the most one could win or less than the least one could win in the lottery. Although they were implausible responses for the selling price question, they nevertheless produced an anchoring effect. In a second experiment, anchors of up to 28 times the lottery expected value (EV) also produced anchoring, but at the far extremes

these anchors had a smaller proportional effect. Quattrone, Lawrence, Finkel, and Andrus (1981) found similar results. Thus, even extreme anchors produce anchoring effects, although it is possible that plausible and implausible anchors have their effects via different mechanisms (Strack & Mussweiler, 1997).

Awareness

Given that anchoring seems to be pervasive, it seems natural to ask if subjects are aware of its influence. Wilson et al. (1996) asked their participants whether they were influenced by the anchor. There was a moderate and significant correlation between reported awareness and the size of the anchoring effect. Nevertheless, the vast majority of subjects reported that they were not influenced by the anchor but still showed an anchoring effect. Thus, the relationship between anchoring and awareness of the anchor's effect was weak, and awareness was not necessary for anchoring. Making participants aware of the anchor's effect does not decrease anchoring. Both Wilson et al. (1996) and Quattrone et al. (1981) found that warning participants not to be influenced by the anchor was unsuccessful. Related to the issue of awareness is whether anchoring occurs even when subjects view the anchor as uninformative. Chapman and Johnson (1999, Experiments 3 and 4) found conflicting results on perceived informativeness. In one study, participants who perceived the randomly generated anchor to be informative showed a larger anchoring effect; whereas in a second study, they showed a smaller effect. In both studies, however, a significant anchoring effect was achieved even among participants who reported that the anchor was uninformative.

Incentives

By many accounts (e.g., Petty & Wegener, 1993; Wilson & Brekke, 1994), subjects cannot avoid judgment biases unless they are aware of them. Thus, given the modest relationship between anchoring and awareness, it would not be surprising if incentives did not reduce anchoring. In fact, the evidence about the influence of incentives on anchoring is mostly negative. Chapman and Johnson (unpublished data) used a procedure in which half the subjects were paid according to the preferences implied by their judgments of simple lotteries. There was no reduction in anchoring when subjects' payoffs depended on their responses. Wilson et al. (1996) found that offering an accuracy prize decreased subjects' ratings of the anchor's influence without actually decreasing the anchoring effect. Tversky and Kahneman (1974) also report that payment for accuracy did not reduce anchoring. Wright and Anderson (1989), however, did find a marginal ($p < .09$) reduction in anchoring due to incentives. Their incentive manipulation included both monetary payment for accuracy and public posting of accuracy scores, which may explain their greater success. Three negative findings and one marginal finding, however, lead to an overall conclusion that incentives reduce anchoring very little if at all.

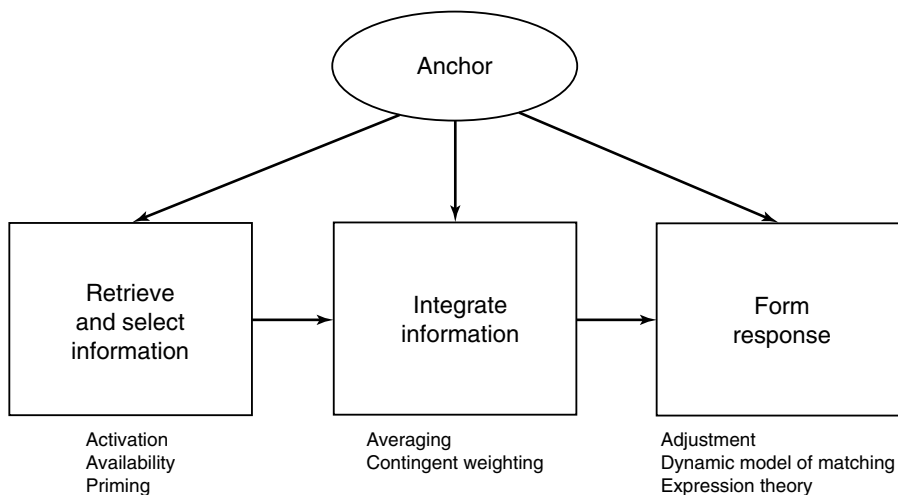
Summary

In summary, anchoring effects are common when the anchor has received sufficient attention. This effect occurs even for extreme anchors and even when respondents are unaware of the effect, have been warned to avoid the effect, or are motivated to be accurate. Anchors are most influential if they are relevant to the target judgment; that is, if they are expressed on the response scale and represent the same underlying dimension, thus comprising a potential answer to the target question.

CAUSES OF ANCHORING

What are the psychological mechanisms that cause anchoring? Figure 6.1 illustrates three stages at which an anchoring process might occur. First, information regarding the target is retrieved through search of memory or the environment. The presence of the anchor could influence what information is retrieved; for example, information similar to the anchor might be primed selectively. Second, the information must be integrated to form an overall target judgment. The anchor may affect that integration process, for example, by giving greater weight to information compatible with the anchor. Or perhaps the anchor itself is included as one of the pieces of information to be integrated. Finally, the judgment must be expressed on an external scale (e.g., dollars, meters). The anchor might influence how the internal judgment is expressed on the external scale. For example, the anchor might facilitate use of the portion of the scale closest to the anchor. Of course, anchoring may have multiple causes, and the relevant mechanisms may occur at more than one stage. An understanding of the locus

Figure 6.1. Three stages at which an anchoring mechanism could occur. The bottom of the figure shows classification of several potential anchoring processes.



of anchoring effects is important in understanding how to debias this effect. For example, if anchoring occurs at the retrieval stage, then debiasing efforts aimed at later stages are likely to be unsuccessful.

ANCHORING AS ADJUSTMENT

Anchoring effects have most often been explained in conjunction with the idea of insufficient adjustment away from the anchor. The name *anchoring and adjustment* implies a particular cognitive process whereby decision makers first focus on the anchor and then make a series of dynamic adjustments toward their final estimate. Because these adjustments are insufficient, the final answer is biased toward the anchor. Tversky and Kahneman (1974) described this type of adjustment process as occurring when subjects answered the United Nations question. This type of account raises a central question: Why are adjustments insufficient (Lopes, 1982)? Most accounts focus either on uncertainty for the true value or a lack of cognitive effort.

The first class of explanations uses the idea that judges are uncertain about the value they want to report. For example, Quattrone et al. (1981) proposed that subjects adjust the anchor until shortly after it enters a range of plausible values for the target item. Thus, when adjusting from a high anchor, decision makers stop at the high end of plausible values, but stop at the low end when adjusting from low anchors. Their study showed a larger anchoring effect for questions judged to have a broad range of plausible values.

A more formal model with a similar underlying idea was developed by Busemeyer and Goldstein (1992; see also Busemeyer & Townsend, 1993) as an anchoring and adjustment account of preference reversals. This dynamic model of matching is not a proposed mechanism of anchoring per se, but rather an anchoring account of preference reversals. (The relation between anchoring and preference reversals is discussed in a later section.) This model is similar to that put forward by Quattrone et al. (1981) in that it posits a process in which decision makers test for a match between the anchor and the target and make adjustments if a match is not achieved. Premature matches lead to insufficient adjustment.

A second class of explanations for insufficient adjustment follow from the notion that adjustment is effortful. Consequently, lack of effort or lack of cognitive resources cause adjustment to be terminated too soon, resulting in a final response that is too close to the anchor. One line of evidence supporting such an adjustment process comes from studies that use cognitive load or busyness manipulations. These paradigms examine situations in which judges place too much weight on one type of information, much like an anchor in our description (Gilbert, Miller, & Ross, 1998). These include the overreliance on information about behavior, as opposed to situational constraints (i.e. the correspondence bias; Gilbert, Pelham, & Krull, 1988), and overreliance on one's own knowledge, as opposed to information available to others (i.e. the "curse of knowledge";

Keysar & Barr, Chapter 8, this volume). These phenomena can be understood as instances of anchoring on one piece of information and consequently underweighting subsequent knowledge by a process of insufficient adjustment.

For example, Kruger (1999) proposed an anchoring account of the “above average” effect – the result that, when asked to compare themselves with their peers, people tend to judge themselves as above average. Kruger argues that people anchor on their own abilities and then make adjustments to account for their peers’ abilities. Whereas focusing on an anchor is an automatic process, adjustment is more effortful, and consequently tends to be insufficient. Kruger found that increasing cognitive load (by asking subjects to rehearse a string of digits) increased the above-average bias, suggesting that limiting cognitive resources makes adjustment more insufficient.

Other studies (e.g., Gilbert et al., 1988) provide similar evidence that increases in cognitive load affect the second stage (adjustment) of the process, and not the first stage (anchoring), suggesting that the second stage is more effortful and less automatic than the first. These studies do not use the traditional anchoring procedure (Tversky & Kahneman, 1974), so their relationship to other anchoring demonstrations is unclear. In addition, Gilbert et al. (1988, p. 738) point out, “The resource metaphor is only one way to describe such effects, and unfortunately, no critical experiment seems capable of distinguishing between resource and other viable interpretations (e.g., structure or skill).” We agree, and take these findings to be evidence that the first stage of the process is different in some ways from a second integration or expression stage, and that a difference in automaticity appears to be a likely candidate. It is less clear, however, whether this evidence requires an adjustment account of anchoring effects.

In evaluating whether adjustment is insufficient because of lack of effort or resources, another line of relevant research concerns the effect of incentives. If insufficient adjustment were the result of lack of effort or allocation of cognitive resources, we might expect individuals to exhibit less anchoring (that is, more sufficient adjustment) when the stakes are high. Busemeyer and Goldstein’s (1992) adjustment account makes just this prediction. As discussed previously, evidence about the effect of incentives is mixed but mostly negative. More broadly, several judgment phenomena that are attributed to anchoring, most notably preference reversals, do not diminish in the face of incentives (see Slovic and Lichtenstein, 1983, for a review). If anchoring were due to lack of cognitive resources, we might also find that experts show less anchoring because they make greater use of specialized task heuristics that conserve cognitive effort. However, this does not appear to be the case (Northcraft & Neale, 1987). Finally, we might expect that warning people about anchoring would diminish anchoring biases. However, Wilson et al. (1996) as well as Quattrone et al. (1981) found that warning participants to avoid the anchoring effect was unsuccessful.

In addition to asking *why* adjustment is insufficient, one might also look for more direct evidence as to whether the anchoring effect does, in fact, result from a cognitive process that involves adjustment. Several studies have tested

directly for adjustment processes. Following a suggestion by Lopes (1982), one type of study uses process-tracing measures to look for evidence of adjustment. A preference reversal study conducted by Schkade and Johnson (1989) asked subjects to price lotteries or to rate them on a 100-point scale. They hypothesized that the lottery attribute (dollars or probability) most similar to the response scale served as an anchor, causing the evaluation to be overly influenced by that attribute. They used process measures by, for example, asking subject to indicate a response by pointing or dragging a cursor along a response scale while the computer recorded the amount of movement. The point at which the subject first touched the cursor to the scale could be an indication of an anchor. These starting points on pricing scales were correlated with the amount to win in the lottery; starting points for rating judgments were correlated with probability information. However, although these starting points may indicate the beginning of an adjustment process, they might also be just an indication of subjects' final answers and an anchoring effect, rather than a specific anchoring process. This latter interpretation is bolstered by the fact that subjects' adjustments along the response scale did not correspond to an adjustment process. For example, larger adjustments should be associated with a smaller anchoring effect because anchoring results from insufficient adjustment, according to this view. However, Schkade and Johnson's data did not show this pattern. Chapman and Johnson (unpublished data) conducted a similar experiment in which subjects considered monetary anchors before evaluating lotteries. Using the same process measures, they likewise did not find an association between adjustments along the response scale and size of the anchoring effect.

A second type of study testing whether anchoring involves adjustment uses tasks where no adjustment could occur. For example, Jacowitz and Kahneman (1995) first presented trivia questions (e.g., the height of Mount Everest) to a calibration group. The 15th and 85th percentiles of the estimates from this group were used as anchors for a second group of subjects. These subjects were presented with a high or low anchor and asked whether it was higher or lower than the target value (e.g., the height of Mount Everest); they then provided their own target estimate. Subjects' estimates showed an anchoring effect, as expected. Of greater interest, the comparison judgments themselves showed an anchoring effect. Although 15% of the calibration subjects had given target estimates lower than the low anchor (or higher than the high anchor), a much larger percentage of the experimental subjects said that the target value was lower (or higher) than the anchor. That is, the comparison question influenced judgments of the target value even before subjects were actually asked to estimate the target. Green, Jacowitz, Kahneman, and McFadden (1998) found a similar result. It may be reasonable to posit that adjustment is involved in estimating the target, but not in judging whether the anchor is higher or lower than the target. Thus, anchoring occurred prior to any possible adjustment, indicating that an adjustment process is certainly not necessary for an anchoring effect.

In sum, although anchoring and adjustment have been close traveling companions since the 1970s, there is only limited evidence that the origins of anchoring lie in an insufficient adjustment process. Although research on increases in cognitive load suggests that an effortful adjustment process may underlie some anchoring-based biases, research using incentives and process tracing methods do not implicate an adjustment bias. In addition, the work by Jacowitz and Kahneman (1995) indicates that adjustment is not necessary for the anchoring effect. Thus, although the anchoring phenomenon is well established, an adjustment mechanism is more questionable. Because at least some conceptions of adjustment characterize it as occurring in response expression, the limited support for an adjustment mechanism suggests that the third stage in Fig. 6.1 may not be well supported.

ANCHORING AS ACTIVATION

In recent years, several authors have suggested that the origin of anchoring lies in the influence of anchors on the first, retrieval stage (see Fig. 6.1). Jacowitz and Kahneman (1995), Chapman and Johnson (1994, 1999), Stack and Mussweiler (1997) and Mussweiler and Strack (1999, 2000) all suggest that the anchor acts as a suggestion, making information consistent with the anchor more available, either in memory through priming mechanisms or because of a biased external search. Because the anchor is considered as a candidate response that subjects entertain, at least as a transient belief, it influences the target value. This account is consistent with work by Gilbert (1990, 1991) and Gilbert, Tatarodi, and Malone, (1993) showing that comprehension includes an initial belief in the assertion presented, followed only later by rejection of false information.

Strack and Mussweiler (1997) examined the idea that anchoring is a special case of semantic priming. They propose that information retrieved in order to compare the anchor to the target is consequently more available for use when estimating the target value. This selective accessibility account (Mussweiler & Strack, 1999, 2000) predicts that the primed information will influence the target judgment only if it is relevant. As described above, Strack and Mussweiler (1997) found just this result. Anchors representing the width of the Brandenburg Gate had only a small influence on judgments of its height, although they had a large influence on judgments of its width. In a further experiment, Strack and Mussweiler (1997) found that extreme anchors resulted in shorter response times for the comparative judgment but longer response time for the absolute judgment (target estimation). This result indicates that comparisons to implausibly extreme anchors do not require relevant target information to be retrieved, yielding a faster comparative judgment. Because this target information has not been primed, however, the absolute judgment takes longer. The implication is that for less extreme anchors, the information primed during the comparative judgment is used in the absolute judgment.

Kahneman and Knetsch (1993) proposed that the target question acts as a memory probe that retrieves the anchor mentioned earlier. The anchor is then treated as a candidate response. This priming is in some ways shallow, however, because the anchor does not always influence the target judgment, even when it primes relevant information. For example, information primed in the course a comparative judgment about whether one would pay \$200 to clean up Canadian lakes is relevant to an absolute judgment about the amount an average Toronto resident would pay, but it should also be relevant to a judgment about the percentage of Toronto residents who would pay \$100. Nevertheless, the anchor influences responses to the first question but not the second, presumably because the first shares more surface features with the anchor.

Jacowitz and Kahneman's (1995) finding that anchoring occurs in response to the comparative judgment (and not just in target estimation) implies that the retrieval of target information primed by the anchor is biased, such that target features similar to the anchor are disproportionately retrieved. This biased retrieval explains why an unexpectedly large number of subjects judge that the target value is lower than the low anchor (or higher than the high anchor). In other words, the comparison between anchor and target results in the anchor appearing too similar to the target. Chapman and Johnson (1999) point to this biased retrieval or asymmetric priming of target features as the key process that produces anchoring. Specifically, they hypothesize that the presence of an anchor increases the availability of features that the anchor and target hold in common while reducing the availability of features of the target that differ from the anchor.

There is evidence from process measures that decision makers concentrate their attention on target features similar to the anchor. Schkade and Johnson (1989) report that subjects spent proportionally more time looking at payoffs in a pricing task (in which they posit a monetary anchor) than in a choice task. In addition, they spent more time looking at probabilities in a 100-point rating task (in which they posit a probability anchor) than in a pricing task. Although there was no explicit anchor used in these studies, Chapman and Johnson (1999) found similar results in a task that did use explicit anchors. Subjects compared apartments described on three attributes. When a provided anchor value was high, they spent more time looking at positive features of the apartment; when the anchor was low, they spent more time looking at negative features.

Mussweiler and Strack (1999, Experiment 4) found similar evidence. In their study, some subjects answered an anchoring question (e.g., about the length of the River Elbe) and were instructed to list the features of the target that came to mind. Subjects given a high anchor tended to list thoughts that implied a high target value, whereas subjects given a low anchor tended to list thoughts that implied a low target value. Mussweiler and Strack also found evidence that the presence of an anchor primes target features that are similar to the anchor. In this study (described in Mussweiler & Strack, 2000), subjects answered the anchoring question, "Is the annual mean temperature in Germany higher

or lower than 5°C (or 20°C)?” and then participated in a lexical decision task. Subjects given the low anchor were faster at identifying words such as “cold” and “snow,” whereas subjects given the high anchor were faster at identifying words such as “hot” and “sun.” This result shows that the anchor primed consistent information in memory.

This activation account of anchoring suggests methods for reducing the anchoring bias. Chapman and Johnson (1999) asked subjects to make predictions (e.g., the likelihood of a Republican winning the presidential election) after considering a numeric anchor. Subjects who were prompted to think of a reason opposing the implications of the anchor value (e.g., for subjects with a low anchor, reasons why a Republican would win) showed less anchoring than a control group given no prompt. In contrast, subjects prompted to think of a similar reason (e.g., for subjects with a low anchor, reasons why a Republican would not win) showed no more anchoring than did the control group. Mussweiler, Strack, and Pfeiffer (2000) replicated this finding. These results suggest that subjects ordinarily retrieve target features or reasons that are similar to the anchor (as evidenced by no effect of the similar prompt) but not those different from the anchor (unless prompted).

Mussweiler and Strack (1999, Experiment 4) also reduced anchoring in a manner consistent with this activation account. Participants in one group were asked to report the target features that came to mind when answering the comparative question. Each participant in a second group was presented with a list of thoughts generated by one of the subjects in the first group. Members of a control group were asked to list target-irrelevant thoughts. The first (own-thoughts) group showed anchoring that was no different than the third (control) group; however, the second (others'-thoughts) group showed less anchoring. The authors proposed that thoughts generated by another subject are more likely to be viewed as suspect or biased than self-generated thoughts, and consequently subjects are more likely to correct for their influence by generating alternative thoughts. Thus, a subject who is presented with a thought suggesting a high target value may counter by retrieving information supporting a low target value. Such a correction would cause the information activated to be less selective and more balanced.

Manipulations of information activation could be used not only to debias anchoring but also to augment it. If anchoring occurs because the anchor facilitates retrieval of target features similar to the anchor, then anything that enhances this selective facilitation should strengthen the anchoring effect. If a large pool of target features exists, the anchor has more opportunity to enhance the similar features selectively. Chapman and Johnson (1999, Experiment 5) asked half their subjects to list actions they undertook that affected their health, a task hypothesized to increase the pool of health-related information. All subjects then estimated the number of Americans who would die of heart disease or cancer in the next 10 years after considering a numeric anchor. Anchoring was stronger among subjects who had first elaborated on the health topic by listing their

health-related actions. Making all target features more accessible through elaboration expanded the pool of features, thereby enhancing the anchor's ability to facilitate retrieval of similar features.

Thus, evidence from a number of studies points to anchors as a type of memory prompt or prime to activate target information similar to the anchor. Anchoring is reduced if subjects are explicitly prompted to consider the anchor's irrelevance (Chapman & Johnson, 1999; Mussweiler et al., 2000) or if the information primed by the anchor is irrelevant to the target judgment (Chapman & Johnson, 1994; Kahneman & Knetsch, 1993; Strack & Mussweiler, 1997). Strack and Mussweiler (1997) point out that the primed information has an effect only if the decision maker is uncertain of the target's value, a prediction consistent with data from Quattrone et al. (1981), Wilson et al. (1996), and Chapman and Johnson (unpublished). Kahneman and Knetsch (1993) describe the priming process as automatic, which may explain why anchoring is largely unaffected by incentives (Wilson et al., 1996). The accounts of anchoring described in this section, although highly similar, are not identical. Further research is needed to specify more precisely the activation processes that produce the anchoring effect. The evidence is mounting, however, that anchoring involves a constructive process of priming or memory retrieval that influences judgments of preference (Payne, Bettman, & Johnson, 1992) and belief.

Phenomena Related to Anchoring

Parallels have been drawn between anchoring and a number of other phenomena. The characterization of anchoring as activation suggests that the effect of anchors is related to judgments of similarity. In judging whether the value of a target object is above or below an anchor, people consider how the anchor and target are similar (Tversky, 1977). As a result, according to the activation view, anchors have their effect because decision makers consider reasons why their value for the target item is like the anchor, but show relative neglect for reasons why their value for the item is unlike the anchor.

This bias toward attending to similarities is analogous to a number of phenomena often labeled collectively as *confirmation bias*. In a variety of tasks, people tend to seek information that is consistent with the current hypothesis would yield positive feedback (e.g., Wason, 1960), and to interpret evidence as consistent with the hypothesis (e.g., Lord, Lepper, & Preston, 1984). Although this strategy is often effective (Klayman & Ha, 1987), it occurs even if the information sought is not diagnostic because it is consistent with many alternative hypotheses. In contrast, hypothesis testers are unlikely to seek information expected to be inconsistent with the target hypothesis, even if that information is quite diagnostic (e.g., Beattie & Baron, 1988; Skov & Sherman, 1986; Snyder & Swann, 1978). The confirmation bias is similar to our proposed model of anchoring in that decision makers examine evidence expected to confirm the hypothesis rather than evidence that could disconfirm the hypothesis. Similarly, Mussweiler and Strack's (2000) Selective Accessibility Model draws a parallel

between confirmation bias and anchoring by positing that decision makers compare the target to the anchor by selectively generating information consistent with the hypothesis that the target's value is equal to the anchor.

A number of authors have noted parallels between anchoring and overconfidence (Block & Harper, 1991; Griffin & Tversky, 1992). Koriat, Lichtenstein, and Fischhoff (1980), for example, argued that overconfidence is due to a failure to consider why the selected answer might be wrong. They demonstrated that a prompt to list counter-reasons was effective in debiasing overconfidence. Using a similar manipulation, Koehler (1994) found that subjects who generated a hypothesis were less overconfident than those who merely evaluated the hypothesis, presumably because generation involves considering alternative hypotheses. This finding suggests that self-generated anchors lead to less bias than experimenter-generated anchors. Block and Harper (1991) found just this result. Subjects gave more accurate confidence intervals if they generated their own anchor (a point estimate) than if they were given another subject's point estimate.

Another phenomenon related to anchoring is the *hindsight bias*, or the tendency for decision makers with outcome knowledge to exaggerate the chances that they would have predicted the outcome in advance (Fischhoff, 1975). Anchoring has been suggested as a possible explanation of this bias (Hawkins & Hastie, 1990); specifically, knowledge of the outcome acts as an anchor that influences judgments of the predictability of the outcome. In hindsight bias experiments, evidence consistent with the outcome is more easily recalled than facts that contradict the outcome (Dellarosa & Bourne, 1984). Thus, the outcome knowledge draws attention to reasons why that outcome was predictable, but not reasons why alternative outcomes were predictable. Hindsight bias is reduced by asking subjects how they would explain alternate outcomes if they had occurred (Arkes, Faust, Guilmette, & Hart, 1988; Slovic & Fischhoff, 1977), in a manner similar to the attentional prompt manipulations used in anchoring studies (Chapman & Johnson, 1999; Mussweiler et al., 2000).

Shafir (1993) provided a demonstration of how the goal of a decision task can shift attention. He found that when asked to accept one of two options, decision makers appear to focus on the positive features of the options. In contrast, when asked to reject one of two options, decision makers focus on the negative features of the options. Consequently, an option with both many positive and many negative features can be both accepted and rejected over a second option with only average features. These results are consistent with the interpretation that the "accept" or "reject" instruction acts as an anchor by increasing the availability of features consistent with the instruction.

Anchoring has been used to explain preference reversals (Busemeyer & Goldstein, 1992; Ganzach, 1996; Lichtenstein & Slovic, 1971; Schkade & Johnson, 1989; Slovic & Lichtenstein, 1983). When Lichtenstein and Slovic (1971) first drew a link between anchoring and preference reversals, they presupposed an

adjustment process. In light of more recent anchoring studies suggesting that anchors increase activation and salience, one might ask whether such an activation process might also underlie preference reversals. Anchors, acting as a prime, may contribute to response mode compatibility effects.

According to a compatibility account of preference reversals (Tversky, Sattath, & Slovic, 1988), the weight given to each attribute of a target item depends on the response mode. Specifically, those attributes that are compatible with the response scale are given more weight. Thus, in pricing lotteries, the dollar outcomes of the lotteries receive relatively more weight. Conversely, in rating lotteries on a 0–100 scale, the probabilities receive relatively more weight, and these weight shifts are accompanied by an increase in the attention paid to the probabilities relative to the dollar outcomes (Schkade & Johnson, 1989).

Anchoring itself also shows compatibility effects (Chapman & Johnson, 1994; Kahneman & Knetsch, 1993; Strack & Mussweiler, 1997). For example, monetary anchors influenced monetary judgments but not life expectancy judgments (Chapman & Johnson, 1994). In anchoring, the anchor draws attention to similar features of the target, which then influence target judgments. In preference reversals, it is the response scale that draws attention to similar features of the target, influencing the preference judgment.

The numerous phenomena related to anchoring suggest that anchoring mechanisms such as activation may underlie many judgments and judgmental biases. Baron (1994) described the tendency to search for evidence that favors a target possibility as one of the major biases leading to poor decisions, and Arkes (1991) described association-based errors as one of three main causes of judgment biases (along with strategy-based and psychophysically based errors). Association-based errors result from considering evidence that is primed by the decision task. Studies of anchoring place it in this class of phenomena. Arkes concludes that such errors cannot be corrected by increasing incentives but can be reduced by instructions or cues to perform a debiasing behavior, such as considering opposing evidence.

Although anchoring shares features with a variety of judgmental effects, it is sometimes improperly categorized with reference point effects, a distinct class of phenomena (Kahneman, 1992). Reference points are values that define the subjective neutral point and thus divide a scale of values into “gains” and “losses.” Changes in the neutral reference point alter evaluations of a target value, especially if what could be perceived as a relative gain is instead perceived as a relative loss. Given that reference points and anchors both involve the presentation of a comparison stimulus (usually irrelevant to the judgment at hand), there is a tendency to confuse the two effects (Kahneman, 1992). In fact, they differ in process and outcome. As argued previously, anchoring occurs primarily through priming and attentional mechanisms; that is, anchoring is an associative error (Arkes, 1991). In contrast, reference point effects occur

primarily through perceptual or psychophysical mechanisms (Arkes, 1991). That is, the position of the reference point alters the slope of the utility function or indifference curve (Sen & Johnson, 1997). Furthermore, anchoring generally leads to judgmental assimilation effects (outcomes too near the anchor), whereas reference points lead to evaluative contrast effects (higher evaluations with lower reference points) and changes in risk-aversion (greater risk-seeking when a given outcome is framed as a loss than as a gain).

Applications of Anchoring

Anchors have been found to influence many judgment tasks, including answers to factual knowledge questions (Jacowitz & Kahneman, 1995; Tversky & Kahneman, 1974), estimation of risks and uncertainty (Plous, 1989; Wright & Anderson, 1989; Yamagishi, 1994), statistical inferences (Lovie, 1985), evaluation of monetary lotteries (Carlson, 1990; Chapman & Johnson, 1994; Johnson & Schkade, 1989; Schkade & Johnson, 1989), judgments of self-efficacy (Cervone & Peake, 1986), judgments of spousal preferences (Davis, Hoch, & Ragsdale, 1986), and predictions of future performance (Czaczkes & Ganzach, 1996; Switzer & Sniezek, 1991).

Anchoring is also a key theoretical concept used to explain other judgment phenomena, such as egocentric biases (Kruger, 1999), attribution (Gilbert et al., 1988; Quattrone, 1982), and overconfidence (Griffin & Tversky, 1992). Anchoring has also been offered as a cause of preference reversals (Lichtenstein & Slovic, 1971; Schkade & Johnson, 1989), biases in utility assessment (Hershey & Schoemaker, 1985; Johnson & Schkade, 1989), information framing effects (Levin, Schnittjer, & Thee, 1988), and biased causal attribution (Gilbert, Chapter 8, this volume; Quattrone, 1982). Finally, anchoring and adjustment serves as a central theoretical component of explanations of the effect of ambiguity on probability judgments (Einhorn & Hogarth, 1985), of belief updating (Hogarth & Einhorn, 1992), and the expression of values (Goldstein & Einhorn, 1987; Busemeyer & Goldstein, 1992).

Many everyday tasks require numerical judgments and thus may be prone to anchoring effects. Northcraft and Neale (1987), for example, asked students and real estate agents to tour a house and appraise it. Appraisal values assigned by both experts (real estate appraisers) and amateurs (students) were positively related to the provided anchor – the listing price of the house. As is the case in many applied settings, one might argue that the anchors were not uninformative, as listing prices are generally correlated with real estate value. However, the participants in this study reported that list price should be irrelevant to the appraised value, yet they were nonetheless influenced by it. Similarly, Caverni and Pris (1990) found that secondary school teachers were influenced by the past records of their students when they graded a new assignment. Evaluation of the current assignment should, of course, be based on its own merits rather than the student's record.

Anchoring in personal injury verdicts was examined in Chapman and Bornstein's (1996) study of mock jurors. The anchors, which took the form of the plaintiff's requested compensation (or *ad damnum*), influenced both judgments of whether the defendant was liable and the amount of monetary compensation awarded. Anchoring occurred even for implausibly low (\$100) or high (\$1 billion) anchors (see also Chapman & Johnson, 1994; Quatrone et al., 1981; Strack & Musweiler, 1997). Legally, the *ad damnum* is irrelevant to both liability and compensation judgments because plaintiffs can request as large an amount as they wish. Englich and Mussweiler (in press) found similar anchoring effects in criminal sentencing.

A potentially important applied context for the measurement of preferences and value is the use of survey research in the area of "contingent evaluation." Green, Jacowitz, Kahneman, and McFadden (1998), for example, examined anchoring effects using respondents' answers to such questions as how much they would be willing to pay to save 50,000 offshore seabirds per year, as well as answers to objective estimation questions such as the height in feet of the tallest redwood in California. They demonstrated strong anchoring effects for both types of questions, and argued that such anchoring effects are much larger in size than the biasing effects typically ascribed to a lack of incentives in contingent valuation surveys.

Hurd et al. (1997) found that anchors influenced reports older adults gave of their monthly expenditures and savings account balances. This result is surprising because one might expect such figures to be well known to the respondents. In addition, the unfolding bracket sequence used as the anchoring manipulation in this study is a common survey technique, suggesting that the initial bracket or anchor may have a large biasing effect on many survey results.

Additional applications of anchoring include demonstration of this bias in group decision making (e.g., Rutledge, 1993) and in individual judgments of group decision outcomes (Allison & Beggan, 1994). Ritov (1996) examined anchoring in a competitive market simulation. In a negotiation between a buyer and seller, initial offers can act as an anchor, and Ritov found that these values affected final profit. Anchoring can also influence consumer behavior (e.g., Biswas & Burton, 1993). For example, Yadav (1994) found that when consumers evaluate two or more items bundled together, the most important item acts as an anchor, which affects the overall evaluation of the entire bundle.

Conclusions

A useful analogy might be drawn between the anchoring effect and the Stroop effect. In the classic Stroop effect, subjects are asked to name the ink color in which a word is printed. Reaction times are longer for color words that do not match the ink color (e.g., the word *red* printed in green ink) than for unrelated words (see MacLeod, 1991, for a review). The meaning of the word is

irrelevant information, yet it influences performance. In a similar fashion, an irrelevant anchor influences judgment, even when decision makers are instructed to ignore it (e.g., Quattrone et al., 1981; Wilson et al., 1996).

Like the Stroop effect, anchoring appears to be both prevalent and robust, as research has yielded scores of demonstrations of this bias. The contaminating effects of irrelevant anchors can be observed in numerous real-world contexts. Understanding the causes of anchoring and what they tell us about the efficacy of various potential debiasing techniques is thus of considerable practical importance.