

SPACING, POSITION, AND ORDER INTERPRETIVE HEURISTICS FOR VISUAL FEATURES OF SURVEY QUESTIONS

ROGER TOURANGEAU
MICK P. COUPER
FREDERICK CONRAD

Abstract We present the results of six experiments that demonstrate the impact of visual features of survey questions on the responses they elicit, the response process they initiate, or both. All six experiments were embedded in Web surveys. Experiments 1 and 2 investigate the effects of the placement of nonsubstantive response options (for example, “No opinion” and “Don’t know” answer options) in relation to the substantive options. The results suggest that when these options are not differentiated visually (by a line or a space) from the substantive options, respondents may be misled about the midpoint of the scale; respondents seemed to use the *visual* rather than the *conceptual* midpoint of the scale as a reference point for responding. Experiment 3, which varied the spacing of the substantive options, showed a similar result. Responses were pushed in the direction of the visual midpoint when it fell to one side of the conceptual midpoint of the response scale. Experiment 4 examined the effects of varying whether the response options, which were arrayed vertically, followed a logical progression from top to bottom. Respondents answered more quickly when the options followed a logical order. Experiment 5 examined the effects of the placement of an unfamiliar item among a series of similar items. For example, one set of items asked respondents to say whether several makes and models of cars were expensive or not. The answers for the unfamiliar items depended on the items that were nearby on the list. Our last experiment varied whether a battery of related items was

ROGER TOURANGEAU, MICK P. COUPER, and FREDERICK CONRAD are affiliated with the Institute for Social Research, University of Michigan, and the Joint Program in Survey Methodology, University of Maryland. The work reported here was supported by two National Science Foundation grants—SES-9910882 to Roger Tourangeau and Mick Couper and SES-0106222 to Roger Tourangeau, Mick Couper, Fred Conrad, and Reg Baker. We are grateful to Darby Miller Steiger for her contributions to Experiments 1 and 2 and to Reg Baker, Scott Crawford, Duston Pope, and Andy Peytchev for their contributions to Experiments 3 through 6. Scott Fricker helped with some of the analyses; we gratefully acknowledge his assistance. Address correspondence to Roger Tourangeau; e-mail: rtourangeau@survey.umd.edu.

administered on a single screen, across two screens, or with each item on its own screen. The intercorrelations among the items were highest when they were all on the same screen. Respondents seem to apply interpretive heuristics in assigning meaning to visual cues in questionnaires. They see the visual midpoint of a scale as representing the typical or middle response; they expect options to be arrayed in a progression beginning with the leftmost or topmost item; and they expect items that are physically close to be related to each other conceptually.

Self-administered questionnaires have always relied mainly on text to convey the questions and to provide any instructions for answering them. Still, as Redline and Dillman (Dillman, Redline, and Carley-Baxter 1999; Jenkins and Dillman 1995; Redline and Dillman 2002; see also Christian and Dillman 2004) have pointed out, this verbal language is supplemented by various visual “languages,” including symbols (such as arrows and boxes), graphical features (font size, boldfacing), and pictures. The impact of these visual languages on survey responses may be especially marked in Web surveys, because Web surveys often incorporate richer visual material than self-administered paper questionnaires and because the visual presentation of the questions in Web surveys may vary markedly across the range of browsers, screen settings, and user preferences available to respondents.

Unfortunately, the implications of visual presentation of the questions are not especially well understood, even for older methods of data collection. Although several texts offer practical guidelines for the design of mail and other self-administered questionnaires (for example, Dillman 1978; Mangione 1995), apart from the work of Redline and Dillman, there has been relatively little empirical work or theoretical analysis of the issues involved. Most of the work that has been done has focused on designing forms to minimize unit and item nonresponse. For example, the recent work by Redline and Dillman (2002) on the use of arrows and boxes in mail surveys examines navigation problems with contingent questions. But the design of paper forms and computer screens may affect not only whether respondents answer the right questions but also which answers they give (Sanchez 1992; Smith 1995). Smith (1995) cites several examples that demonstrate that unintended variations in printing paper questionnaires can produce dramatic shifts in the answers. Still, the study of forms design is in its infancy, and the impact of forms design on measurement error has been almost entirely neglected. Thus, the rapid spread of the Web as a vehicle for survey data collection has brought to the fore important but overlooked questions for survey designers.

Web surveys are the latest example of computerized self-administration of survey questions, and we suspect they may ultimately turn out to be the most popular. Aside from the gains from computerization and self-administration,

Web data collection eliminates interviewers entirely, sharply reducing the cost of data collection. Furthermore, Web surveys can deliver rich visual content that is impossible or prohibitively expensive to incorporate in most other modes. Not surprisingly, the growth in Web surveys has been dramatic. Despite serious concerns about coverage and nonresponse in Web surveys (Couper 2001), the commercial research sector has rapidly embraced the Internet for faster and cheaper data collection, and almost daily there are reports of new surveys being done over the Web.

This article focuses on some of the measurement issues raised by Web data collection. It presents the results of six experiments that were conducted as part of several Web surveys. The experiments are based on the assumption that respondents follow simple heuristics in interpreting the visual features of questions. These interpretive heuristics may sometimes lead to unintended inferences about the meaning of a question, based on incidental visual cues. In the same way, respondents often make inappropriate inferences about the meaning of survey questions based on incidental verbal or numerical cues, such as the numbers assigned to the scale points (Clark and Schober 1992; Schwarz 1994, 1996; Schwarz, Grayson, and Knäuper 1998). Hoffman (2000) argues that interpretive rules are central in visual processing generally and are responsible for such key abilities as depth perception and the perception of shape. The rules for interpreting ordinary visual stimuli can sometimes lead to systematic misinterpretations of those stimuli, producing optical illusions. Similarly, the application of interpretive heuristics for visual cues in questionnaires can lead to misreadings of survey questions.

We distinguish five heuristics that respondents may follow in interpreting visual questionnaires and present experimental evidence for the first three. We are currently conducting research on the final two heuristics, but we believe it is useful to describe the full set here to put the current experiments in their broader theoretical context. The five interpretive heuristics are:

1. Middle means typical;
2. Left and top mean first;
3. Near means related;
4. Up means good; and
5. Like means close.

Each heuristic assigns a meaning to a spatial or visual cue. For example, according to the first heuristic, respondents will see the middle item in an array (or the middle option in a set of response options) as the most typical. Because the middle option is seen as typical, respondents may seize on it as an anchor or reference point for judging their own position (Tourangeau, Rips, and Rasinski 2000, pp. 244–46). A study by Schwarz and Hippler (1987) provides evidence that respondents see the middle option as representing the typical or mean value for the population. They asked respondents how much television they watched in a typical night; respondents reported they watched

more television when the response categories emphasized the high end (the first response category in this condition was “Up to 2½ hours”) than when they emphasized the low end (the initial category was “Up to ½ hour”). They also asked respondents about the average value in the population and found that respondents’ estimates of the population average were affected by the set of response options they had received. Respondents seem to make similar inferences about typicality when the options are presented aurally (Tourangeau and Smith 1996); in Web surveys, visual cues determine which option is seen as representing the midpoint of the scale and, therefore, the typical response.

The second heuristic implies the leftmost or top item in a list of items will be seen as the “first” in some conceptual sense. This interpretive principle reflects the reading order of English (and most Western languages). When the list is a series of ordered response categories or scale values, respondents will expect the topmost or leftmost option to represent one of the two endpoints (“Agree strongly”); in addition, they will expect each of the successive options to follow in some logical order (“Agree,” “Neither agree nor disagree”) and they will expect the final option in the list—the one in the bottom or rightmost position—to represent the opposite endpoint (“Disagree strongly”). If the list does not conform to these expectations, respondents may become confused, make mistakes, and take longer to respond. As a result, respondents will answer more quickly and more accurately when the order of the response options is consistent with the heuristic than when it is inconsistent with it (cf. Christian and Dillman 2004). Moreover, when the items in a set do not follow any particular order, respondents may nonetheless infer that they follow a logical progression, going from the most favorable to the least favorable, or from the most conservative to the most liberal.

The third interpretive heuristic states that respondents expect items that are physically near each other on the screen (or on the page) to be related conceptually. Respondents will, for example, see two items adjacent to each other on a single screen to be more closely related than the same two items presented on separate screens. They will infer greater conceptual distance between response options that are physically further apart. This heuristic recalls the Gestalt Law of Proximity (Koffka 1935; Wertheimer 1923; see also Kubovy, Holcombe, and Wagemans 1998), which states that perceivers tend to group nearby objects into figures. We note that the spacing may be a particularly important issue with Web surveys, as spacing may accidentally vary because of differences across Web browsers; in addition, the need for respondents to scroll may make spacing more salient in Web surveys than in paper questionnaires. An implication of this third heuristic is that items placed on the same screen or in a grid are likely to be seen as more closely related than when they are administered on separate screens; as a result, the answers are likely to be more highly correlated (Couper, Traugott, and Lamias [2001] provide evidence supporting this prediction).

The fourth interpretive heuristic is a variant on the second; it states that, with a vertically oriented list, the top item or option will be seen as the most desirable. Respondents are most likely to apply this heuristic when the items or options in the list clearly vary in their desirability. Thus, when the list consists of a series of evaluative scale points (“Strongly approve,” “approve,” and so on), respondents will expect the positive points to come at the top. Similarly, if they are asked their preferences about a series of items displayed in a vertical array, they may expect the best ones to be at the top. Of course, the verbal labels attached to the options or the respondents’ actual preferences among the items will generally override the impact of the positional cues, but things will go most smoothly when the positive end of the rating scale (or the most desirable of the items) appears at the top—respondents will answer more quickly and more reliably when the options are positioned in the order suggested by the heuristic. The relation between vertical position and differences in desirability is apparent in large number of deeply rooted metaphors (for example, happiness is “being up,” sadness “being down”; heaven is up, hell is down; good news is upbeat, bad news is downbeat; “on the rise” connotes success, “on the decline,” failure; and so on; see Carbonell 1983).

The final heuristic asserts that visually similar options will be seen as closer conceptually. This heuristic was anticipated by the Gestalt Law of Similarity, which states that perceivers tend to see similar objects as forming a single figure (Wertheimer 1923). Consider, for example, scales that vary the appearance of the response options (for example, scales in which the different scale points vary in hue). According to the “like means close” heuristic, when the endpoints of a scale are different colors, they will be seen as further apart than when they are the same color. This example is analogous to one investigated by Schwarz and his colleagues (Schwarz et al. 1991). That study asked respondents to rate their success in life on response scales in which the eleven points were labeled with numbers ranging from -5 to $+5$ or from 0 to 11. For both groups, the answers tended to pile up at the high end of the scale—most people see themselves as successful—but this tendency was more pronounced when the low end was labeled with negative numbers. Schwarz and his colleagues argue that the negative numbers suggested that the low of the scale was reserved for abject failure, whereas the low positive numbers anchoring the low end of the second scale suggested that these points indicated mere lack of success. The general principle that applies both to the numbering and shading of the response options is that when things differ along two dimensions (for example, in both hue and value), they are seen as differing more sharply than when they differ only in a single attribute. Judd and Harackiewicz (1980) dub this principle the “accentuation effect.” When the endpoints in a scale are perceptually more distinct, respondents will see them as differing more from each other conceptually—that is, as representing more extreme positions. (In the study done by Schwarz et al. [1991], the ends of the scale seemed further

apart conceptually to the respondents when the numerical labels differed both in sign and value than when they differed only in value.)

As noted above, this article concerns the first three heuristics: “middle means typical,” “left and top mean first,” and “near means related.” We present the results of six experiments embedded in four Web surveys. Each experiment examines how the visual appearance or arrangement of the questions affects the response process, the answers to the questions, or both.

Spacing and the “Middle Means Typical” Heuristic

We conducted three experiments to test the implications of the “middle means typical” heuristic. All three experiments tested items in which the response options formed a scale with a conceptual midpoint (that is, a middle option), but that option fell to one side of the *visual* midpoint of the scale. Our prediction is that placing the conceptual midpoint to one side of the visual midpoint will increase the proportion of respondents selecting a response from the *opposite* side, shifting the overall mean of the responses. Put another way, we predicted that responses would be displaced toward the visual midpoint. The three experiments used three different methods to create discrepancies between the visual and conceptual midpoints of the scale.

EXPERIMENTS 1 AND 2: SEPARATING SCALE POINTS FROM NONSUBSTANTIVE RESPONSES

Our first two experiments compared several methods for including nonsubstantive answer categories (“Don’t know” and “No opinion” responses) along with substantive scale responses. The first of these experiments compared two versions of a pair of attitude questions. One version presented the nonsubstantive options simply as additional radio buttons. The other version used a divider line to separate the scale points from the nonsubstantive options (see the top panel of figure 1 below). The second experiment examined the same two attitude items but added one further condition: it examined using spacing to segregate the scale points from the nonsubstantive options. The questions asked whether the federal government was doing enough to ensure equal job opportunities for women and whether it was doing enough to provide day care for working parents. The response scale for the items had a clear conceptual midpoint—“About the right amount.” We thought that when the nonsubstantive options were not visually distinct from the scale points (that is, when they just appeared to be additional radio buttons), respondents would use the visual midpoint of the scale (“Too little”) as a reference point and would be more likely to select answers from that conceptual end of the scale.

Method. Both experiments were embedded in Web surveys conducted on our behalf by the Gallup Organization. Survey Sampling Inc. (SSI) selected

the samples for both surveys, using a frame (the eLite frame) that consisted of more than 7 million e-mail addresses of Web users. The e-mail addresses were provided by visitors to various Web sites who had agreed to receive messages on a topic of interest. SSI selected 25,000 e-mail addresses for our first survey and 30,000 for our second and sent out e-mail messages inviting the recipients to take part in “a study of attitudes and lifestyles.” The e-mail invitations included the Web address (URL) of the Gallup Web site with our questionnaire and a unique identification number (which prevented respondents from completing the survey more than once). Gallup’s proprietary software administered the Web questionnaires.

The first survey ran from March 28 to April 10, 2001; the second, from February 26 to March 12, 2002. Of the 25,000 invited to participate in the first survey, 2,987 completed the entire survey (and 255 others got part way through) for a response rate (AAPOR [2000] RR1) of 11.9 percent. Of the 30,000 invited to take part in the second survey, 1,590 responded (and an additional 192 began the survey but did not finish it) for a response rate of 5.3 percent.¹ These response rates are quite low by the standards of high-quality mail surveys but are quite typical for Web surveys of this type (see, for example, Couper 2001). The samples for the two surveys were similar, both including more women than men (71.3 percent of the respondents to the first survey and 58.1 percent of the respondents to the second were women) and relatively few Hispanics (4.1 percent in the first survey, 4.2 percent in the second) and nonwhite respondents (13.8 percent nonwhite in the first survey, 16.8 percent in the second). In the second survey, we included an item on educational attainment; the sample was highly educated, with 36.0 percent of the respondents having graduated from college.

The questionnaires for the two surveys were quite similar and included questions on a range of topics, including gender attitudes, socially desirable (voting, church attendance) and socially undesirable behaviors (drinking and illicit drug use), impression management, and demographic characteristics. Tourangeau, Couper, and Steiger (2003) provide additional details about the first survey. In total, there were 84 items in the questionnaire for Experiment 1 and 85 items in the questionnaire for Experiment 2. The key items came as the tenth and eleventh questions in both surveys.² Generally, each screen presented several questions, and some questions consisted of grids that requested similar ratings for multiple items. Almost all of the items were closed-ended and required the respondents to click on a radio button to indicate their answers. Figure 1 displays the two versions of question 10. Item 11 in both surveys repeated the experiment with a second question (“Think of how much

1. All of the response rates presented here follow the formula for AAPOR (2000) Response Rate 1. This formula treats partial respondents as nonrespondents; it also assumes all of the sample cases were eligible for the survey.

2. Although occasionally we refer to the question numbers of particular items, none of the surveys reported here actually displayed the question numbers to the respondents.

a. Question 10, Divider Line Version

Think of how much the federal government is doing to make sure women have the same job opportunities as men. Would you say the federal government is doing too much, about the right amount, or too little about this?

Far too much
 Too much
 About the right amount
 Too little
 Far too little

 Don't know
 No opinion

< Back | Next >

b. Question 10, Version with No Divider Line

Think of how much the federal government is doing to make sure women have the same job opportunities as men. Would you say the federal government is doing too much, about the right amount, or too little about this?

Far too much
 Too much
 About the right amount
 Too little
 Far too little
 Don't know
 No opinion

< Back | Next >

Figure 1. Formats for displaying nonsubstantive options in Experiment 1.

the federal government is doing to provide day care centers for the children of working parents. Would you say the federal government is doing too much, about the right amount, or too little about this?”). Respondents in the first survey were randomly assigned to receive one of the two versions of the questions; those in the second were randomly assigned to the no-divider version, to a divider line version with a shorter line than the one used in Experiment 1 (one that did not extend across the whole screen), or to a third version in which the nonsubstantive options were separated from the remaining radio buttons by a space. Respondents in both experiments received the same version of questions 10 and 11.

Results. In the versions of the questions that used a divider line or space to separate the nonsubstantive options from the substantive responses, the visual midpoint of the scale fell at the conceptual midpoint (“About the right amount”). In the versions that presented the nonsubstantive options simply as additional radio buttons, however, the visual midpoint fell on one end of the

scale (“Too little”). In both experiments, this affected the distribution of the answers, although the results were somewhat more dramatic in the first experiment. Table 1 presents the main results from both items and both studies. (In the second experiment, the results for the divider line and spacing conditions are quite similar so we combine them here.) In all four cases, the means for the substantive answers are higher when there is no separation between the five scale points and the two nonsubstantive answers than when a divider line or spacing separates the scale points from the nonsubstantive answers. (Higher scores represent views more toward the “Far too little” end of the scale; lower scores represent views more toward the “Far too much” end of the scale.) The difference is statistically significant for both items in Experiment 1 and for question 11 in Experiment 2. There is one additional trend apparent in the data: respondents were more likely to select one of the nonsubstantive options when they were clearly separated from the remaining response options. This pattern is only significant for question 10 in the first experiment. Dividing the response options (whether by a line or a space) may have drawn attention to the nonsubstantive options and made respondents more likely to select them. Figure 2 displays the full distributions of the responses for the three items that showed significant differences across versions. In each case, the histograms show a pile-up of responses in the “Far too little” category when there is no divider line or space separating that option from the nonsubstantive options.

Discussion. As a practical matter, these results suggest that nonsubstantive response options (if they are offered at all) should be clearly separated from the substantive ones. This strategy may have the drawback of calling attention to the nonsubstantive answers, producing higher rates of item nonresponse, but that outcome would generally be preferable to changing the distribution of substantive answers. On a more theoretical level, the results suggest that the meaning of each response option is partly based on its relative position within the array of response options. As an option appears to move to one side or other of the visual midpoint, the meaning respondents assign to it may shift as well.

EXPERIMENT 3: UNEVEN SPACING BETWEEN SCALE POINTS

We carried out another experiment that tested the hypothesis that respondents see the visual midpoint of the scale as representing the conceptual midpoint; again, we predicted that the “middle means typical” heuristic would affect their interpretation of the response categories—and the answer they selected. More specifically, Experiment 3 examined what happens when the answer categories for an item are unevenly spaced, and, as a result, the conceptual midpoint for an item does not coincide with the visual midpoint. Figure 3 displays an example of even and uneven spacing. When the spacing is uneven, four of the options are to the left of the visual midpoint of the scale, and the options on the right appear closer to the midpoint than when the options are

Table 1. Mean Substantive Response and Percentage Choosing Nonsubstantive Answers, by Experiment and Condition

| | Experiment 1 | | Experiment 2 | |
|---------------------------|------------------------------|--------------------------|-------------------------|-------------------------|
| | Q10 | Q11 | Q10 | Q11 |
| Mean Substantive Response | | | | |
| No Divider | 3.69 (1283) | 4.01 (1350) | 3.51 (431) | 3.88 (450) |
| Divider/Spacing | 3.57 (1117) | 3.88 (1207) | 3.47 (843) | 3.75 (885) |
| | $F(1,2398) = 9.20^{**}$ | $F(1,2398) = 12.3^{***}$ | $F(1,1272) < 1$ | $F(1,1333) = 4.49^*$ |
| % Don't Know/No Opinion | | | | |
| No Divider | 17.5% (1555) | 13.0% (1552) | 21.1% (546) | 16.4% (538) |
| Divider/Spacing | 21.0% (1421) | 15.1% (1421) | 24.6% (1118) | 19.4% (1098) |
| | $\chi^2 = 7.24^{**}, df = 1$ | $\chi^2 = 2.58, df = 1$ | $\chi^2 = 2.56, df = 1$ | $\chi^2 = 2.23, df = 1$ |

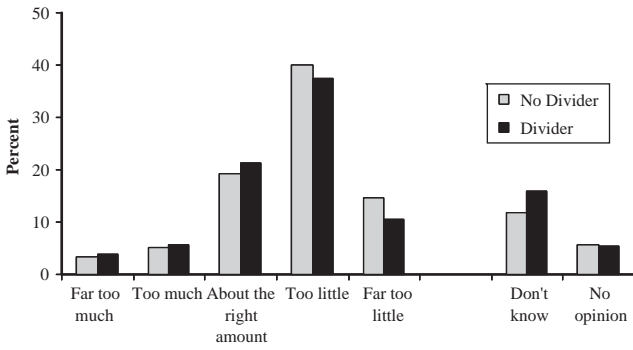
NOTE.—Parenthetical entries are cell sizes.

* $p < .05$.

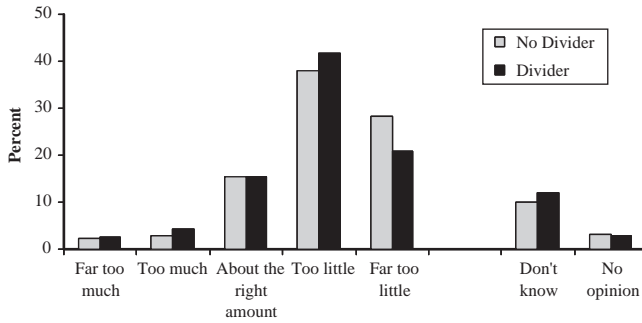
** $p < .01$.

*** $p < .001$.

a) Experiment 1, Question 10



b) Experiment 1, Question 11



c) Experiment 2, Question 11

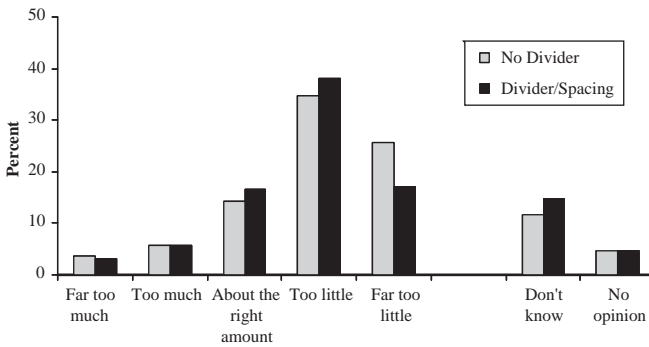


Figure 2. Distribution of responses, by item and experiment.

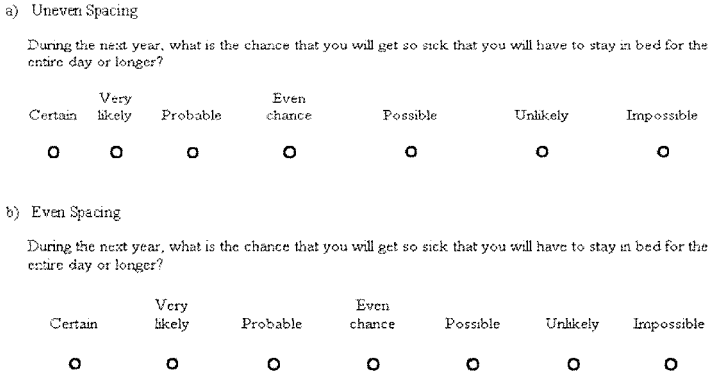


Figure 3. Evenly and unevenly spaced response scales.

spaced evenly. The uneven spacing shifts the position of some scale points more than others, with bigger shifts for the points that are conceptually near the middle of the scale (for example, “Even chance” and “Possible”) than for those at the extremes (“Certain” and “Impossible”).

Method. Experiment 3 was embedded in a Web survey conducted by MSInteractive from March 26, 2003, through April 7, 2003. The sample consisted of 39,217 e-mail addresses drawn from two SSI frames for Web surveys. The first frame was the eLite frame used in Experiments 1 and 2; the second was SSI’s Survey Spot sample, an opt-in Web panel of almost one million persons who have signed up online to receive survey invitations. SSI invited sample members to take part via an e-mail invitation that asked them to complete a survey of attitudes and lifestyles funded by the National Science Foundation; the invitation included the URL for the questionnaire. Non-respondents received one follow-up e-mail. Of the 39,217 invited to take part in the survey, 2,722 completed the questionnaire (and an additional 457 began the survey but did not finish it) for a response rate of 6.9 percent (AAPOR [2000] RR1). The questions were administered via SPSS’s mrInterview software. The questionnaire consisted of 50 items, with most screens displaying a single item.

The experiment involved two items about the respondents’ overall health. The first one asked, “How would you rate your health?” The other item, displayed in figure 3, asked respondents how likely it was they would get sick enough during the next year that they would have to spend a day or more in bed (“During the next year, what is the chance that you will get so sick that you will have to stay in bed for the entire day or longer?”). These were the first two items in the questionnaire. Responses to the first item, which asked

for ratings of overall health, were unaffected by the spacing of the response alternatives; we focus on responses to the second item.³

Results. We expected that the respondents who got the unevenly spaced options (shown in the top panel of figure 3) would be more likely to select answers from the right side of the scale (the unlikely end) than those who got the evenly spaced options (shown in the bottom panel of figure 3). In the uneven spacing condition, “Possible” and “Unlikely” are closer to the visual midpoint and are, therefore, more likely to be selected than in the even spacing condition. The results are in line with that expectation. With the uneven spacing, 63.4 percent of the respondents chose answers from the right side of the scale (“Possible,” “Unlikely,” or “Impossible”); with the even spacing, the percentage dropped to 58.3 percent ($\chi^2 = 7.54$, $df = 1$, $p < .01$). The biggest difference between the two versions of the item involved the “Unlikely” option, which was selected by 33.3 percent of the respondents who got the uneven spacing but only by 29.2 percent of those who got the even spacing. An analysis of variance indicates that the overall means of the responses were also affected; on the average, responses indicated less chance of a sick day in bed when the response options were unevenly spaced (mean of 4.60) than when they were evenly spaced (4.45) — $F(1,2689) = 6.78$, $p < .01$.

Discussion. Like Experiments 1 and 2, Experiment 3 suggests that respondents use not only the verbal label attached to a scale point but also its position relative to the visual midpoint to infer what specific value the scale point is supposed to represent. When “Unlikely” was displaced toward the visual midpoint, it was selected more often, presumably because it seemed to represent a less extreme probability.

Order and the “Left and Top Mean First” Heuristic

A second heuristic respondents may use in understanding response options and selecting an answer from among them is the “left and top mean first” heuristic. According to the heuristic, the leftmost or top item in a list of items should be the “first” in some conceptual sense, and the remaining options should follow from left to right or from top to bottom in some logical progression. For example, scale options often follow a graded sequence starting at one pole, proceeding through the intermediate values, and ending at the other pole. When the scale options do not, in fact, follow their conceptual rank order, it will slow respondents down (and may affect their answers). Moreover, when a series of items seems to follow some order, respondents may use

3. The overall health question was also involved in another experiment. That experiment varied which one of two pictures—a photograph of a healthy young woman jogging or of a woman in a hospital bed—was displayed just before or with the item, and it compared three different positions for the picture on the prior screen—just before the health item, on the same screen in the survey header, or just to the left of the question text. Neither the photograph nor its position affected responses to the item on chances for a day in bed due to illness.

the position of an unfamiliar item within the series to infer its characteristics. We carried out experiments to test both these predictions.

EXPERIMENT 4: ORDER OF THE RESPONSE OPTIONS

We varied the order of the response options in four behavioral frequency items and two attitude items. The frequency items used a graded frequency scale ranging from “Never” to “Every day.” Similarly, the attitude items used a 5-point agree-disagree scale. Table 2 displays the three versions of the response scales for the agree-disagree items. One version of the items presented the response options in an order that was consistent with the “top is first” heuristic; that is, the top option was one of the endpoints (“Never” in the case of the behavioral frequency scale or “Strongly agree” in the case of the agree-disagree scale), and each of the succeeding options followed in order of extremity. A second version of the items departed somewhat from this order, putting “Never” and “It depends” at the bottom of the scale. The final version departed more sharply from the order prescribed by the heuristic. This version of the frequency items presented “Never” as the topmost option, followed by “Every day,” with the remaining options in order of decreasing frequency. The final version of the attitude items presented “It depends” as the top option, followed by “Strongly agree,” “Strongly disagree,” “Agree,” and finally “Disagree.” Our main hypothesis is that greater departures from the order implied by the heuristic would lead to slower response times (for a similar analysis of the effects of violating conventional order in presenting response options, see Holbrook et al. 2000).

Method. Experiment 4 was part of a Web survey conducted by MSInteractive from April 2 through April 23, 2002. Like the sample for Experiment 3, the sample for this survey consisted of e-mail addresses drawn from the eLite and Survey Spot SSI frames. SSI sent invitations to 14,192 e-mail addresses drawn from the two frames. The e-mail invitation asked them to complete a survey sponsored by the National Science Foundation. A total of 2,871 persons started the questionnaire and 2,568 completed it, for a response rate of 18.1 percent (not counting the partials). Again, the mrInterview software

Table 2. Response Scales in Experiment 4

| Consistent with Heuristic | Mildly Inconsistent | Strongly Inconsistent |
|---------------------------|---------------------|-----------------------|
| Strongly agree | Strongly agree | It depends |
| Agree | Agree | Strongly agree |
| It depends | Disagree | Strongly disagree |
| Disagree | Strongly disagree | Agree |
| Strongly disagree | It depends | Disagree |

administered the questionnaire, which included 43 items, each generally appearing on its own screen. Among other experiments, the survey included several that compared response options in three different orders.

We carried out two independent experiments, one with four frequency items and the other with two agree-disagree items. The four frequency items asked how often the respondent took vitamins (question 1), ate fruit (question 2) and pasta (question 3), and went shopping (question 15). The agree-disagree questions (questions 13 and 14) asked about following the doctor's advice ("It is SENSIBLE to do exactly what the doctors say") and about when the respondent went to the doctor ("I have to be VERY ILL before I go to the doctor"). We compared three versions of both sets of questions (see table 2). As noted earlier, one version of the items presented the response options in the order implied by the "top means first" heuristic; the other versions departed mildly or sharply from that order. All three versions arrayed the options vertically.

Respondents got the response options in the same order for all four of the frequency items; similarly, they got the same version of both agree-disagree items. (There were independent random assignments of the respondents to the versions of the two sets of items.) The software recorded both answers to the items and the time that elapsed from the time of presentation of the item to the response. We did not tell the respondents they were being timed.⁴

Results. We anticipated that respondents would answer the questions most quickly when the items followed the order implied by the "top means first" heuristic, with the slowest answers in the third version of the questions (where the order of the response categories departed most sharply from the order implied by the heuristic). After dropping unusually long response times (those over 60 seconds), we compared mean response times across the three versions of the items. Three of the six items showed significant differences in response times, but the results were clearest for the two agree-disagree items. Figure 4 displays the average response times for those two items. For both, the differences in response times across experimental treatments were highly significant: $F(2,2533) = 18.7$ for question 13 (on following the doctor's advice) and $F(2,2591) = 12.6$ for question 14 (on going to the doctor). The degrees of freedom for the two items vary because of item nonresponse. For question 13, response times increased from an average of 16.4 seconds for the version in which the options followed the order consistent with the heuristic to a mean of 18.9 seconds in the conditions in which the options were most inconsistent

4. Experiments 4–6 were carried out as part of two Web surveys that each included a large number of independent experiments. A danger with such designs is that the experimental manipulations, which were designed to affect one or two nearby items, may have had remote effects on later items or interacted with variables manipulated later in the questionnaire. Because we carried out randomization separately for each of the experiments, we were able to look for such carryover effects. In 154 tests focusing on the outcomes in Experiments 4, 5, and 6, we found 10 significant effects; this is a rate of 6.5 percent, not much greater than the 5 percent false positive rate to be expected by chance.



Figure 4. Response times and consistency with heuristic.

with the heuristic. For question 14, the mean response times went from 10.8 to 12.4 seconds across the three versions. The response times for the four frequency items showed no consistent pattern.

The rearrangement of the response options also affected the distribution of responses. Again, the results were most striking for the two agree-disagree questions. The proportion of respondents selecting the “It depends” option dropped dramatically when that option came at the bottom of the list (in the second version of the items) than when it came in the middle or at the top of the list (in the first and third versions). With question 13, the proportion choosing the “It depends” option fell from 45.8 percent (when it was the middle option) or 41.5 percent (when it was the top option) to 20.8 percent (when it came last). These differences across version were highly significant: $\chi^2 = 127.4, df = 2, p < .001$. The comparable figures for question 14 were 27.7 percent (when “It depends” was the middle option), 21.8 percent (when it was the top option), and 8.2 percent (when it came at the bottom of the list). Again, the differences were highly significant ($\chi^2 = 106.0, df = 2, p < .001$). For both items, then, the “It depends” option is most popular when its position suggests that it represents the midpoint of the scale.

Discussion. The findings for the two agree-disagree items supported our predictions, but those for the four frequency items were mixed at best. One difference between the two sets of items may account for the difference in the results. Frequency questions generally presuppose that the respondent engages in the behavior in question—takes vitamins, eats pasta, or goes shopping. Because the “Never” option denies this presupposition, it may be seen as a special option—that is, as not part of the scale. Thus, whether it comes at the top or the bottom and whether it is placed next to the frequent or infrequent options may have relatively little impact on how respondents process it. For the agree-disagree

items at least, the results confirm what many questionnaire designers already knew—because the “strongly inconsistent” order departs so markedly from the expected order, it is an unsuitable format for use in real surveys.

In planning this experiment, we were unsure whether varying the order of the response options would produce a change in response times, a change in the answers, or both. To the extent that respondents were carefully attending to the response options, we expected that putting the response options in an order inconsistent with the “top means first” heuristic would slow the respondents down, as they noticed and adjusted to the unexpected ordering. To the extent that they were not attending to the response options, respondents might overlook the unconventional order actually used and select an option based on the order they expected; if so, the distribution of the answers would change. Both effects are apparent with the agree-disagree items, although the change in the distribution of the answers is mostly restricted to the “It depends” option, which was selected more often when it was the middle option.

EXPERIMENT 5: POSITIONAL INFERENCES

When confronted with an unfamiliar item, respondents may use the “top means first” heuristic to infer the item’s characteristics from its position within a series of similar items. For example, consider respondents who have been asked to judge the prices of a series of cars. If the cars seem to follow some order (say, from expensive to inexpensive models), then the respondents may use this order in judging where an unfamiliar make and model (for example, the Fiat Tipo) falls on the spectrum. We carried out an experiment involving six sets of items that tested whether respondents drew such inferences based on positional cues.

Method. This experiment was embedded in the same Web survey as Experiment 3. It included six target items, each of them embedded in a series of seven related, nontarget items. Within each of the series, we varied the position of the target item (a relatively unfamiliar example) so that it appeared either as the third or seventh in the list; the other items were arranged to convey the sense that they were presented in rank order from top to bottom. In each case, we hoped that the six nontarget items would form a Guttman scale, with the nontarget items arrayed in order by scale position. For each set of items, the majority of the respondents did, in fact, give answers on the six nontarget items that conformed to the requirements of a Guttman scale. For example, respondents judged whether each of seven makes and models of car (BMW 318, Acura Integra, Mazda Protégé, Toyota Corolla, Dodge Neon, and Geo Metro) were expensive. The highest proportion of respondents rated the BMW as expensive, the next highest proportion the Acura, and so on down the list. The Fiat Tipo (the unfamiliar, target item) was inserted either as the third car from the top or at the bottom of the list. Each series was presented in a grid format, with each item presented in one row of the grid.

Table 3. Proportion Yes (and Sample Size), by Position in Grid

| Judgment/Item | Percent Yes (N) | | χ_1^2 |
|--|-----------------|-------------|------------|
| | Third from Top | Bottom | |
| Important for healthy diet / Isoflavin | 46.8 (1325) | 45.7 (1253) | <1 |
| Low in saturated fat / Cod liver oil | 43.8 (1361) | 37.8 (1287) | 9.69** |
| Expensive hotel / Clarion Inn | 63.3 (1357) | 45.8 (1285) | 82.0*** |
| Expensive city / Ocala, FL | 53.0 (1357) | 60.4 (1298) | 14.9*** |
| Expensive midsize car / Austin Rover | 93.8 (1368) | 88.7 (1285) | 22.2*** |
| Expensive small car / Fiat Tipo | 74.9 (1352) | 62.8 (1275) | 44.4*** |

* $p < .05$.
 ** $p < .01$.
 *** $p < .001$.

In addition to this list of small cars (question 21), the six target items asked whether isoflavin was important for a healthy diet (one of the seven dietary items in question 7), whether cod liver oil was low in saturated fat (one of seven fats in question 8), whether Clarion Inns were expensive hotels (one of seven hotel chains in question 18), whether the cost of living was high in Ocala, Florida (one of seven cities in question 19), and whether the Austin Rover was an expensive midsize car (one of seven midsize cars in question 20). All of the questions were yes-no items. Respondents received all six target items either in the third or the seventh position in the list.

Results. Table 3 displays the main results from this study, the proportion of respondents answering “yes” for each of the target items. The vertical position of the item made a significant difference for five of the six items. The differences are particularly large for judgments of the Clarion Inns and the Fiat Tipo, both of which were more likely to be seen as expensive when they came near the top of the list (among the more expensive hotels and cars) than when they came at the bottom (among the cheaper ones).

We classified each judgment according to whether it was consistent with the intended judgments for the items presented at the top of each series. For example, we classified the judgment that the Clarion Inns were expensive as consistent with the top hotels in the array of hotels. (And, in fact, the items at the top of each list did elicit very high levels of consensus, with 97 percent or more of the respondents giving the same judgment for the top item in each series.) Across the six judgments, the respondents who got the target items in the third position in the series were significantly more likely to give a judgment consistent with the item at the top of the list than those who got the target items in the bottom position—62.8 percent versus 58.9 percent; $F(1,2720) = 21.1, p < .001$. Still, the results were in the expected direction only for four of

the six target items, and the two reversals are statistically significant. They involve the judgments of cod liver oil and Ocala, Florida. Cod liver oil was *more* likely to be rated low in saturated fat when it came at the beginning of the series (just below lard and butter) than when it came at the bottom (just below olive oil and light vegetable spread). Similarly, Ocala, Florida, was *less* likely to be seen as an expensive city to live in when it came near the top of the list (just below San Francisco and Houston) than when it came at the bottom (after Bismarck, North Dakota, and Springfield, Illinois).⁵

Discussion. Although the overall trend for the six target items is consistent with the “top means first” heuristic, there do seem to be systematic differences across items. The heuristic assumes that when respondents are unfamiliar with a specific item, but the series seems to be arrayed in order along the dimension of interest, they will use the position of the item within the series to infer its value. When respondents do not see the items as arrayed in order, they may instead contrast the target item with nearby items, especially when they have some alternative basis for judging the target’s value on the dimension of interest. We did not pretest the items to determine that the target items were, in fact, unfamiliar to most of the respondents, relying instead on our intuitions. Some respondents may have had a rough sense of where the target items fell on the dimension of interest. For example, respondents may have inferred that Ocala must be at least somewhat expensive to live in simply because it is in Florida; when Ocala came just below the most expensive cities in the series, it apparently seemed less expensive to the respondents than when it came just below cities that are relatively inexpensive places to live (like Bismarck and Springfield). Similarly, respondents may have inferred that cod liver oil must be relatively low in saturated fat (after all, it is taken from a fish) and contrasted it with lard and butter when it came just below those items on the list.

This analysis suggests that the direction and size of the position effect depends on whether respondents see the nontarget items as arrayed in order. We analyzed each item according to the proportion of respondents whose judgments for the nontarget items followed the expected pattern. More specifically, we assumed the nontarget items formed a Guttman scale and counted the proportion of respondents whose judgments met this assumption. If the items form a Guttman scale, a respondent who rates one item as, say, expensive should also rate all the preceding items in the series as expensive.⁶ We used the proportion of respondents whose nontarget judgments followed one of the patterns required by a Guttman scale as a measure of our success in

5. We also looked at the impact of the position of the target item on ratings of the other six items in each set. Of the 36 nontarget items across the six sets, only three showed significant differences depending on the position of the target item within the array. These differences did not follow any clear pattern.

6. For each set of items, there were seven patterns of responses consistent with the assumption that the nontarget items formed a Guttman scale. For example, respondents could rate all six nontarget cars expensive; they could rate just the BMW and none of the other five nontarget cars expensive; just the BMW, Acura, and Protégé and not the other three; and so on.

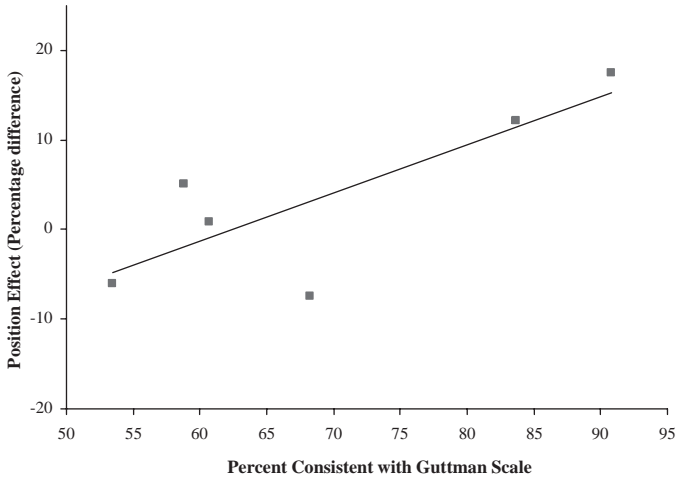


Figure 5. Position effect and percentage of respondents whose nontarget responses form Guttman scale.

conveying a sense that the items were arrayed in order. As is apparent from figure 5, the proportion of respondents who gave ratings consistent with the intended order was strongly related to the impact of position on judgments of the target (as measured by the difference in the percentage of “yes” judgments given to the target item when it appeared in the third and seventh position on the list). The correlation between the two variables is .81 ($p < .06$).

Even in the cases for which we observed reversals, the position of the target item within the list clearly mattered. Whether respondents assimilate the target item to nearby items (as we expected) or contrasted it with its neighbors, judgments of the target item were clearly affected by the local context, and this context was determined by the physical position of the target item within the series.

Experiment 6: Grouping and the “Near Means Related” Heuristic

Still another heuristic that respondents may apply in interpreting Web questions is the “near means related” heuristic. Based on this heuristic, respondents may expect items that are physically near each other on the screen to be closely related conceptually. We tested one implication of the heuristic—that respondents will see stronger interconnections among items that are displayed on a single screen than among those displayed on separate screens and thereby boost the correlations among them.

Method. This experiment was embedded in the same Web survey as Experiment 4. It compared the presentation of a battery of eight related items on a single screen versus presentation across several screens. These items were presented as the twelfth question in the survey. All eight items used the same 7-point response scale (with the endpoints labeled “Agree” and “Disagree”), and all of them asked about diet:

- A. Maintaining a healthy diet is a priority in my life.
- B. I avoid “fast food” because it’s not healthy.
- C. I monitor my cholesterol level closely.
- D. How food tastes is more important to me than its nutritional value.
- E. I pay attention to nutritional information on food packaging.
- F. I limit the amount of red meat in my diet.
- G. My lifestyle makes it difficult to eat right.
- H. I try to balance my diet across the key food groups.

The first version presented the eight items on a single screen in a grid format. The second version presented the items on two screens, both as four-item grids. (One grid presented items A–D above; the other, items E–H). The final version presented one question per screen. We tested the hypothesis that the single grid design would encourage greater consistency among the answers. This hypothesis reflects the “near means related” heuristic; thinking that the items are more closely related when they are presented on the same screen, the respondents will tend to answer them more consistently. Again, response times were recorded automatically, without the respondents’ awareness.

Results. As we expected, the responses to the eight diet questions were more highly intercorrelated when the items were presented in a grid on a single screen (Cronbach’s alpha of .621) than when the eight items were presented in two grids on separate screens (alpha of .562) or when each item appeared on its own screen (Cronbach’s alpha of .511). If we treat the alpha values as correlations and apply Fisher’s z transformation to them, then the linear trend across the three conditions is significant ($z = 2.62, p < .01$). In addition, when the eight items appeared in two grids of four items each, the median correlation for items in the same grid was .382, but the median correlation was only .351 for items in different grids. These findings are consistent with the “near means related” heuristic, and they replicate the findings of Couper, Traugott, and Lamias (2001) on grids versus single items.

Higher correlations do not necessarily mean more valid responses. We looked at the extent to which respondents gave the same answers to all eight items, a response tendency that Krosnick (1991) has termed “nondifferentiation.” More specifically, we calculated the proportion of items for which each respondent gave his or her most common response (see Holbrooke, Green, and Krosnick 2003). With eight questions and seven response categories, this proportion could range from .25 to 1.0 for respondents who answered all eight questions. (Respondents who answered all eight questions would have

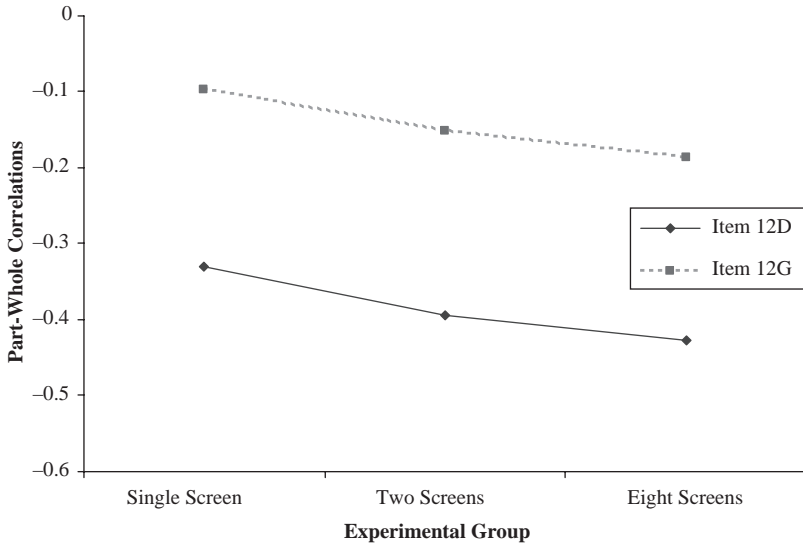


Figure 6. Part-whole correlations for reverse worded items, by format of the eight-item battery.

selected the same answer category at least twice, producing a minimum score of .25.) Respondents who got the items on a single screen showed less differentiation than those who got them on two screens or on eight screens (mean nondifferentiation scores of .436, .422, and .412, respectively; $F(2,2519) = 6.11, p < .01$). Two of the items (items D and G above) were “reverse worded.” For these two questions, agreement indicated *less* concern about diet; for the other six, agreement indicated greater concern. The relation of these two items to overall scale scores (the part-whole correlations) was weaker when the eight items were presented in a single grid ($r = -.331$ for question 12D and $-.097$ for question 12G) than when they were presented in two grids on separate screens ($r = -.395$ and $-.151$) or on eight separate screens ($r = -.427$ and $-.187$). Apparently, respondents were less likely to notice the reverse wording when the items appeared in a single grid. These results are graphed in figure 6.

We also compared the time it took for respondents to answer all eight questions. We set very long response times (those greater than 240 seconds overall for the eight items) to missing, and then carried out a one-way analysis of variance on the remaining times.⁷ Respondents took less time on average to answer the eight questions when they were presented in a grid on one screen

7. We dropped a total of 71 response times, or about 2.8 percent of the times. The results do not change appreciably if we delete only those times greater than 480 seconds.

(a mean of 60.4 seconds) than when they were presented in two grids on different screens (65.4 seconds) or individually on eight separate screens (99.0 seconds): $F(2,2493) = 322.1, p < .001$.

Discussion. When the items were grouped in a single grid, respondents seemed to infer greater similarity among them than when they were spread across two or eight screens; in fact, they may have inferred more similarity among them than was warranted. Respondents were more prone to select the same answer for all eight items, and they seemed less sensitive to the fact that two of the items were reverse worded when the eight items were in a single grid. Respondents seemed to use the proximity of the items as a cue to their meaning, perhaps at the expense of reading each item carefully. There was also a large difference in response times across the three experimental groups—almost 40 seconds; doubtless, much of this difference reflects the time it took to download the items. It is also possible, though, that the single-grid format encourages respondents to move through the items quickly, making it more likely they will overlook some of the distinctions among them.

General Discussion

We find evidence in all six of our experiments that respondents make inferences about the meaning of survey items based on visual cues such as the spacing of the response options, their order, or the grouping of questions. These inferences affect how quickly respondents answer the questions, which answers they select, or both. In Experiments 1 and 2, respondents seemed to use the visual midpoint of the scale as a reference point for choosing their answers. When the nonsubstantive response options (such as “Don’t know”) were integrated visually with the scale points (as in the bottom panel of figure 1), thereby shifting the visual midpoint, the answers shifted as well. Respondents were more likely to select a response from the bottom end of the scale—that is, their answers shifted in the same direction as the visual midpoint. Similarly, when we altered the spacing of the answer categories in Experiment 3 (see the top panel of figure 3), the answers on the right side of the scale (which were now closer to the visual midpoint) became more popular than they were with evenly spaced response categories. We argue that respondents apply a heuristic in which the visual midpoint is seen as providing a benchmark, representing either the conceptual midpoint of the scale or the most typical response.

Respondents also seem to expect items to be arrayed in a logical progression from left to right or from top to bottom. In a survey that uses aural administration of the questions, respondents may have similar expectations regarding successive items, but in a Web survey (or a self-administered paper questionnaire), the cues that trigger these expectations are visual rather than temporal. Our experiments provided two kinds of evidence for the “top and left mean first” heuristic. We presented respondents with questions in which

the answer categories systematically violated the heuristic (Experiment 4). This slowed them down and affected the distribution of their answers; in line with the heuristic, they were more likely to pick the neutral “It depends” option when it came in the middle of the scale. In addition, Experiment 5 showed that respondents inferred something about the characteristics of an unfamiliar item by its position in an array of similar items. Across six parallel items, the position of the unfamiliar item affected the ratings it received; apparently, the placement of the item in the array conveyed information about the position of the item on the dimension of judgment. For a couple of the items, the effect seemed to work in the opposite direction, with respondents contrasting an item with nearby items on the list. In part, these reversals may have occurred because the respondents did not see the items in the array as following a specific order.

Another cue that affected respondents’ answers was the grouping of the questions onto Web pages. Experiment 6 compared responses to a battery of related questions on diet when the items were presented in a single grid on one screen, in two grids on separate screens, or one item at a time across eight screens. We expected that respondents would see the items as more closely related when they all were placed together in a grid and that is what we found: the average intercorrelation of the items (as measured by alpha) went up when the eight items were placed together in the same grid. The inference that the items were closely related because they were visually grouped may have supplanted a more careful reading of the items. Respondents were more prone to use the same answer from one item to the next and were less sensitive to the fact that two of the items were worded in the opposite direction from the other six, when all eight items were placed together in a grid.

These studies extend two earlier lines of research. First, they provide additional evidence that the visual layout of survey items matters. Work by Smith (1995), Sanchez (1992), and Redline and Dillman (2002) have shown similar effects for paper questionnaires to the ones we demonstrate here for Web surveys. Our studies are based on somewhat different premises from the earlier studies; our work assumes that respondents make quick interpretive judgments about the questions based on the visual cues in the questionnaire, much as the interpretation of purely visual stimuli is based on the rapid application of interpretive rules (Hoffman 2000). Our studies also extend earlier work on the (unintended) effects of verbal cues on how respondents interpret and answer survey questions (Clark and Schober 1992; Schwarz, 1992, 1996; Schwarz et al. 1991). Respondents are sensitive to the numerical labels assigned to the response scale, to the specific values mentioned in the response options, and to other verbal cues. Our results indicate that they may also make unintended inferences based on the visual cues offered by the question. Basing their reading on the questions’ visual appearance, respondents may miss key verbal distinctions and interpret the questions in ways the survey designers never intended.

References

- American Association for Public Opinion Research (AAPOR). 2000. *Standard Definitions: Final Disposition of Case Codes and Outcome Rates for Surveys*. Lenexa, KS: AAPOR.
- Carbonell, Jaime. 1983. "Derivational Analogy in Problem Solving and Knowledge Acquisition." In *Proceedings of the International Machine Learning Workshop*, ed. R. S. Michalski, pp. 12–18. Urbana, IL: Department of Computer Science, University of Illinois at Urbana-Champaign.
- Christian, Leah M., and Don A. Dillman. 2004. "The Influence of Graphical and Symbolic Language Manipulations on Responses to Self-Administered Questions." *Public Opinion Quarterly* 68:57–80.
- Clark, Herbert H., and Michael F. Schober. 1992. "Asking Questions and Influencing Answers." In *Questions about Questions: Inquiries into the Cognitive Bases of Surveys*, ed. J. M. Tanur, pp. 15–48. New York: Russell Sage.
- Couper, Mick. P. 2001. "Web Surveys: A Review of Issues and Approaches." *Public Opinion Quarterly* 64:464–94.
- Couper, Mick P., Michael Traugott, and Mark Lamias. 2001. "Web Survey Design and Administration." *Public Opinion Quarterly* 65:230–53.
- Dillman, Don A. 1978. *Mail and Telephone Surveys: The Total Design Method*. New York: Wiley.
- Dillman, Don A., Cleo D. Redline, and Lisa R. Carley-Baxter. 1999. "Influence of Type of Question on Skip Pattern Compliance in Self-Administered Questionnaires." In *Proceedings of the American Statistical Association, Survey Research Methods Section*, pp. 743–48. Washington, DC: American Statistical Association.
- Hoffman, Donald. 2000. *Visual Intelligence: How We Create What We See*. New York: Norton.
- Holbrook, Allyson L., Melanie C. Green, and Jon A. Krosnick. 2003. "Telephone versus Face-to-Face Interviewing of National Probability Samples with Long Questionnaires: Comparisons of Respondent Satisficing and Social Desirability Bias." *Public Opinion Quarterly* 67:79–125.
- Holbrook, Allyson L., Jon A. Krosnick, Richard T. Carson, and Robert C. Mitchell. 2000. "Violating Conversational Conventions Disrupts Cognitive Processing of Attitude Questions." *Journal of Experimental Social Psychology* 36:465–94.
- Jenkins, Cleo R., and Don A. Dillman. 1995. "Towards a Theory of Self-Administered Questionnaire Design." In *Survey Measurement and Process Quality*, ed. L. Lyberg et al., pp. 165–96. New York: Wiley.
- Judd, Charles, and Judith Harackiewicz. 1980. "Contrast Effects in Attitude Judgment: An Examination of the Accentuation Hypothesis." *Journal of Personality and Social Psychology* 38:390–98.
- Koffka, Kurt. 1935. "Perception: An Introduction to the Gestalt-Theorie." *Psychological Bulletin* 19:531–85.
- Krosnick, Jon A. 1991. "Response Strategies for Coping with the Cognitive Demands of Attitude Measures in Surveys." *Applied Cognitive Psychology* 5:213–36.
- Kubovy, Michael, Alex O. Holcombe, and Johan Wagemans. 1998. "On the Lawfulness of Grouping by Proximity." *Cognitive Psychology* 35:71–98.
- Mangione, Thomas W. 1995. *Mail Surveys: Improving the Quality*. Thousand Oaks, CA: Sage.
- Redline, Cleo D., and Don A. Dillman. 2002. "The Influence of Alternative Visual Designs on Respondents' Performance with Branching Instructions in Self-Administered Questionnaires." In *Survey Nonresponse*, ed. R. Groves, D. A. Dillman, J. L. Eltinge, and R. J. A. Little, pp. 179–93. New York: John Wiley.
- Sanchez, Maria E. 1992. "Effect of Questionnaire Design on the Quality of Survey Data." *Public Opinion Quarterly* 56:206–17.
- Schwarz, Norbert. 1994. "Judgment in a Social Context: Biases, Shortcomings, and the Logic of Conversation." *Advances in Experimental Social Psychology* 26:123–62.
- Schwarz, Norbert. 1996. *Cognition and Communication: Judgmental Biases, Research Methods, and the Logic of Conversation*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Schwarz, Norbert, Carla E. Grayson, and Barbel Knäuper. 1998. "Formal Features of Rating Scales and the Interpretation of Question Meaning." *International Journal of Public Opinion Research* 10(2):177–83.
- Schwarz, Norbert, and Hans-Jürgen Hippler. 1987. "What Response Scales May Tell Your Respondents: Information Functions of Response Alternatives." In *Social Information Processing and Survey Methodology*, ed. H. J. Hippler, N. Schwarz, and S. Sudman, pp. 163–78. New York: Springer-Verlag.

- Schwarz, Norbert, Knäuper, Barbel, Hippler, Hans-Jurgen, Noelle-Neumann, Elizabeth and Clark, F. 1991. "Rating Scales: Numeric Values May Change the Meaning of Scale Labels." *Public Opinion Quarterly* 55:618–630.
- Smith, Tom W. 1995. "Little Things Matter: A Sampler of How Differences in Questionnaire Format Can Affect Survey Responses." In *Proceedings of the American Statistical Association, Survey Research Methods Section*, pp. 1046–51. Washington, DC: American Statistical Association.
- Tourangeau, Roger, Mick P. Couper, and Darby M. Steiger. 2003. "Humanizing Self-Administered Surveys: Experiments on Social Presence in Web and IVR Surveys." *Computers in Human Behavior* 19:1–24.
- Tourangeau, Roger, Lance J. Rips, and Kenneth Rasinski. 2000. *The Psychology of Survey Response*. New York: Cambridge University Press.
- Tourangeau, Roger, Tom W. Smith. 1996. "Asking Sensitive Questions: The Impact of Data Collection Mode, Question Format, and Question Context." *Public Opinion Quarterly* 60: 275–304.
- Wertheimer, Max. 1923. "Untersuchen zur Lehre von der Gestalt II." *Psychologische Forschung* 4: 301–50. Translation published in 1938 as "The Laws of Organization in Perceptual Forms." In *A Sourcebook of Gestalt Psychology*, ed. W. Ellis, pp. 71–88. London: Routledge and Kegan Paul.