qualtrics**XM**

The Qualtrics handbook of question design

Table of Contents

Introduction

At Qualtrics, our mission is to close the gaps in human experiences. And we've observed that better experiences start with better research.

In the introduction to the original edition of this handbook, Dave Vanette noted that "survey data are only as good as the questions that generate them." And that's the most critical concept to understand in the science of designing questions for surveys. By understanding survey design best practice, companies are able to gather better-quality research, enabling them to craft data-driven experiences for their customers.

In this handbook, we'll cover four key areas for XM researchers:

1. Answering questions The cognitive processes involved in answering questions and the concept of respondent burden, and the response behaviors that emerge when respondents are faced with too much burden.

2. Data quality What it means and how to measure it.

3. Designing questions to enhance data quality and reduce respondent burden.

4. Existing questions vs. designing your own The pros and cons of both options.

Answering questions

Answering questions

THE COGNITIVE RESPONSE PROCESS

Whenever researchers analyze data collected from humans, assumptions are made about the mental processes that produced those data. With surveys, an implicit assumption made by many researchers is that each question was answered using what survey methodologists often call the "optimal cognitive response process."

Having an "optimal" cognitive response process does not mean that researchers assume that every survey respondent has an optimal, focused experience (more on that later), it merely provides a model for how to think about the steps a respondent goes through - and everything that could go wrong along the way.

There are four components of the survey response process that respondents must mentally engage with to answer questions (adapted from Tourangeau, Rips, and Rasinski 2000):

1. **Comprehension** Respondents must understand the question and any instructions. They must determine what is meant by the question as it may differ from the literal interpretation of the words; identify the information being sought by the question; and link key terms from the question text to relevant concepts, feelings, and experiences.

2. **Retrieval** Respondents must identify an information retrieval task and then search their memory for the relevant generic and specific information.

3. Judgement Respondents must assess the completeness and relevance of the information retrieved from memory, draw conclusions and inferences based on the accessibility of the information, and integrate the information retrieved, even if the information is stored in different "packets" of memory. For example, the respondent may have to calculate whether they consumed a particular food in the past week. First, they must evaluate if they've had the food, then whether they had it in the past week.

4. Response Respondents then take the summarized information and map their judgment onto a response category, editing the judgment for the format of the requested response as necessary.

When respondents perform the four steps above to answer a question, they are said to have "optimized" the response process.

Most research that analyzes and interprets survey data makes an implicit assumption that this is the process each respondent used to generate their answer (data) for each question. How realistic this assumption is varied dramatically by project, so in many cases it may not be a valid assumption to make.

In fact, when presented with this model of the response process, many researchers react with disbelief— many of them have taken surveys before and will admit that, even as research professionals, they do not carefully engage in each of these steps for every question on a survey. This acknowledgment is important because the quality (reliability and validity) of the data that are collected by surveys typically depend on the degree to which respondents complete these steps for each question.

As researchers, we cannot directly control the care that respondents direct toward this process. However, question and survey design decisions can influence respondents heavily— for better or worse.

SURVEY DESIGNERS SHOULD OPERATE WITH TWO GOALS:

1 Make completing the response process as easy as possible for all respondents

2 Avoid making it easy for respondents to shortcut this process when they are answering questions

RESPONDENT BURDEN

Extrinsic Burden

Survey respondents are offering their time to researchers often with little to no compensation and no expectation of additional incentives or feedback.

Some respondents may be trying to make extra cash as panel members. Others may be interested in ensuring their feedback on a product or service is heard. Some may have a personal stake in the content of the survey. Others may just be nice people who try to accommodate requests of their time. All respondents have other things competing for their time and mental energy that are usually more compelling than completing a survey. In order to make the survey response process as easy as possible, it helps to have an understanding of the ways respondents completing surveys may be burdened at the time of the survey request.

The extrinsic burden is the burden or constraint on the respondent's time and attention independent of the actual survey process. These are external factors the survey designer has no control over but that affect the respondent's attention and may determine the quality of the data captured.

+ Time Respondents have limited time available in their day to complete a survey. They may start a survey while they have a down moment but realize the survey will take longer than they'd expected. If respondents are feeling pressured for time but still obligated to respond, they may speed through a survey without paying close attention to the content.

+ Distractions Respondents, like everyone else, are faced with distractions ranging from interruptions from their kids and family, social media, work responsibilities, and other day-to-day worries. This may cause them to make slower progress through the survey but without reading the text carefully due to interruptions.

+ Interest Respondents who volunteer for surveys may not be highly interested in the topic at hand; or, they may have started off interested but something far more interesting appeared to draw their attention away.

+ Environment Respondents may be completing a survey while waiting in line at the grocery store, or while riding a bus. They may be completing a survey in a loud cafe or while watching TV. In any of these situations, the respondent may have to stop, switch tasks, then come back to the survey question.

+ Ability Respondents may have limited cognitive skills and simply not have the ability to understand the questions. Respondents with limited experience thinking about the question topic prior to taking the survey may perceive the task as burdensome or confusing. Researchers often use education as a rough proxy for cognitive skills because it is difficult to include a full assessment of cognitive ability in most surveys.

+ Personal Traits Respondent characteristics like personal accountability and the need for cognition (a psychological term that describes the extent to which a respondent is inclined to expend significant mental energy on a particular task) have an impact on how respondents perceive burden from the survey request. Respondents with high personal accountability and a high need for cognition are less likely to feel burdened by a survey.

The extrinsic burden is not caused by the survey questions but is part of the respondent's overall cognitive load and may have an influence on the way the questions are answered. In order to manage the survey experience for the respondent so that error is reduced, surveys should be designed to minimize the effort required to understand questions. Researchers should design questions with clear, simple response tasks.

Intrinsic Burden

Survey questions are often anything but clear and simple. The burden that is presented to the respondent as part of the survey request or task is intrinsic burden. Survey respondents, already overtaxed with the extrinsic factors creating the burden described above, must also navigate the burden associated with the design of the survey. Intrinsic burden from poor questionnaire design is completely under the control of the researcher.

In the "Bad Example" the question text contains unnecessary words and ideas, complicated phrasing, and words with multiple definitions. The Bad Example introduces a more intrinsic burden to the survey process than the "Good Example" which is shorter and does not use more words than are necessary to convey the necessary information to the respondent.

Figure 1A *Bad Example*

When considering the most desirable vehicle for recreation, what are the most important attributes you would use for a purchasing decision?
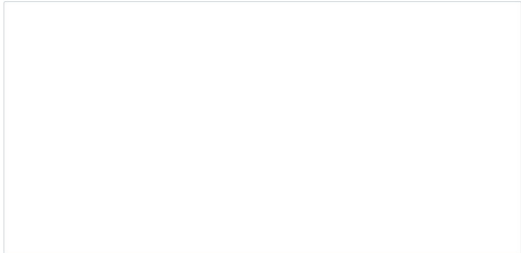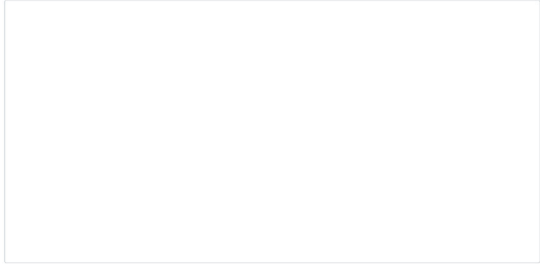
Figure 1B *Good Example*

When purchasing an everyday vehicle, what features are most important to you?

The burden is also introduced in questionnaires by requiring an extensive recall. During the retrieval phase, answering questions about past events is more difficult for respondents than answering questions about the present. The further in the past the event in question, the more challenging the retrieval task will be for the respondent, increasing their intrinsic burden. In general, avoid asking about the past any more than necessary.

Similarly, answering multiple subjects is more difficult than answering about a single subject. Often survey designers try to reduce the number of questions in a survey by combining multiple concepts in the question text. Unfortunately by increasing the content of the survey question, the respondent burden is actually increased by forcing them to sort through multiple subjects to ascertain the critical ask of the question.

In terms of the judgment process, respondents find it more difficult to make comparative judgments than absolute judgments; therefore a rating scale is a less burdensome task than a ranking task. Furthermore, judgments that are decomposable are easier to make than those that are not (e.g., if a person dislikes all carbonated beverages then it will be easier to rate a new flavor of soda than if they like some carbonated beverages and not others.) In the response phase of question response, respondents find it more difficult to understand numeric scale labels compared to verbal labels. Verbal labels reduce the burden on the respondent.

THE IMPACT OF INTRINSIC AND EXTRINSIC BURDEN

Respondents faced with questionnaire burden that they can't overcome respond in three ways. These may be conscious or unconscious responses.

1. **Refuse to complete the survey** If a survey is obviously long, many respondents will rule it out from the onset and just not agree to participate. Incentives can counter the effect of this dampening of participation, but only to a point.

2. **Start a survey and break off** when they find that the content or interface is too difficult to work through, the topic is uninteresting, or they simply get a better option for how to spend their time. Refusals and breakoffs are easier for researchers to manage, because their impact on the survey data is clear - it's missing.

3. **Engage in satisficing behavior** the process of producing a minimum satisfactory outcome.

In survey methodology, satisficing is a theory that explains the common practice of a respondent taking cognitive shortcuts while answering survey questions. Developed by Stanford Professor Jon Krosnick (Krosnick 1991; 1999) this theory suggests that survey respondents engage in satisficing to varying degrees when responding to questions. That is, respondents don't always complete the optimal cognitive response process as described in the previous section.

> *Satisficing is the common practice of a respondent taking cognitive shortcuts while answering survey questions*

Satisficing is not necessarily a conscious and intentional choice on the part of respondents. Satisficing is the product of balancing intrinsic and extrinsic burdens while focusing on the task at hand.

Previous studies have identified a number of factors that seem to cause respondents to engage in satisficing. By understanding these causes, researchers can make informed decisions about how to design questions and questionnaires in ways that will minimize the opportunities for respondents to satisfice. While the researcher does not have control over every potential cause, very often the data quality can be improved by applying good question design principles.

THE IMPACT OF SATISFICING

Satisficing behavior takes on a variety of forms when respondents are completing a survey. In most cases, there are design decisions that researchers can use to make satisficing easier or more difficult for the respondent to engage in.

Acquiescence The first form of satisficing that we will highlight is called "acquiescence response bias," or a respondent's tendency to agree with suggestions. This is most commonly seen in questions that use agree-disagree response scales; in these question types respondents have a bias toward agreeing, regardless of the content of the statement they are evaluating.

This also commonly happens with true/false questions, where respondents are more likely to report "true" than "false," and yes/no questions, where respondents demonstrate a bias toward "yes."

Using any form of these response scales makes it easier for respondents to engage in satisficing behavior rather than going through the optimal response process.

In general, avoid using generic response scales and instead use response scales that are specific to the subject that your question is asking about. For example, if you were asking about the degree of satisfaction or dissatisfaction that your respondent felt about an experience, you could formulate it as an agree-disagree statement: "I was satisfied with my experience," and provide response options ranging from "strongly agree" to "strongly disagree."

Or, you could use the best-practice approach of using a construct-specific response scale: "How satisfied or dissatisfied were you with your experience?" with response options ranging from "extremely satisfied" to "extremely dissatisfied."

Straightlining Straightlined, or non-differentiated responses to rating questions, are another form that satisficing can take. Chances are, if you've ever used a matrix or grid question type in a survey, you've found that at least some of your respondents have provided the same answer for each question in the grid. This is straightlining.

In the worst-case scenario, respondents that straightline are not even reading the individual questions or statements but are simply clicking answer choices in a straight line to get through the survey as quickly as possible. It is important to note that not all straightlined responses are necessarily invalid, and as a precautionary measure, researchers will sometimes include a reverse-coded version of the same question in order to catch respondents that report, for example, both liking and disliking the same thing.

Best practices for avoiding straightlining in surveys include:

+ Avoid using grid or matrix question types

+ Use construct-specific response scales

+ Ask one question per page

Ideally, we want respondents to carefully read the question and response options before providing an answer. However, when a respondent is satisficing, they are more likely to simply identify the first reasonable response option provided and select it without reading any further. In web surveys this is typically the first reasonable response at the top of a vertically oriented scale or on the left of horizontally oriented scales. This tendency results in an effect known as "primacy" and can introduce a bias into your data.

For questions that offer categorical responses, the best practice is to randomize the order of all of the response alternatives, which usually reduces potential bias at the expense of increased variance. For rating scales, randomizing which end of the response scale is on the top or left— depending on whether the orientation of the scale is vertical or horizontal—may also help reduce bias (Malhotra 2009).

COMBATING SATISFICING

There are two primary tools that researchers can use to combat survey satisficing: task difficulty and respondent motivation. By designing questions to reduce the difficulty of the cognitive response process and maximize respondent motivation, researchers can reduce the chances that respondents will engage in these negative response behaviors. Fortunately, taking steps to reduce satisficing is also likely to increase the validity and reliability of responses.

Task difficulty To reduce the difficulty of responding to questions, we recommend that researchers take three steps:

+ Make questions easy to understand

+ Minimize distractions

+ Keep the survey short

It is important to note that when we refer to making the questions easy to understand, the goal is to help respondents optimize the cognitive response process and provide accurate, valid, and reliable responses - not simply fast responses.

For example, respondents are able to click through matrix or grid questions very quickly, perhaps indicating that this is an easy question type for respondents. But research indicates that this question type actually may require greater cognitive effort to optimize responses to than if the same questions are asked individually.

Keeping the survey short typically means asking the questions that are necessary for your research goals and no more—there is no room for "pet" or "nice to know" questions if you want high-quality data.

Web surveys that take longer than 10 minutes are much more likely to suffer from low-quality data, as respondents fatigue and begin satisficing

In general, we find that web surveys that take longer than 10 minutes are much more likely to suffer from low-quality data, as respondents fatigue and begin satisficing. Typically, respondents can provide about 30 responses (to average questions) in 10 minutes, but with any new or revised survey it is important to pre-test it yourself on a few people to see how long it takes.

In recent years, response rates to surveys have typically been extremely low. This suggests that those individuals that do participate and become respondents must have some amount of motivation to take the survey. Capitalizing on this motivation by taking steps to avoid diminishing it and also attempting to increase it whenever possible may not only help keep respondents in the survey but will also help them to provide higher-quality responses.

There are five approaches that we recommend for getting the most out of your respondents' motivation:

+ Ask them to commit to provide their best responses

+ Leverage the importance of the survey

+ Leverage the importance of their responses

+ Use incentives and gratitude to increase engagement

+ Keep the duration of the survey short (10 minutes or less)

Early work in survey methodology indicated that asking a respondent to commit to providing their best data actually had positive effects on the quality of responses they gave (Cannell, Miller, and Oksenberg 1981). At Qualtrics, we have recently replicated this finding in the web survey context across 14 countries.

To implement this, you can simply ask your respondents at the beginning of your survey if they will commit to providing their best data. We believe that, because people feel a desire to be internally consistent with statements and commitments that they have made (Cialdini, Cacioppo, and Bassett 1978), they are more likely to provide high-quality survey responses after they have committed to do so. In our 14-country study, we found that over 98% of respondents were willing to commit to providing their best responses when asked.

In terms of leveraging the importance of the survey, the best practice is to reaffirm the decision made by your respondents to participate by providing an indication that the topic of the survey is important.

Similarly, you can remind your respondents that their responses actually matter and are important. Incentives can be an effective method of increasing respondent motivation as well. Reminding the respondent that they will be paid or will be entered into a lottery for a prize can not only keep respondents from leaving the survey, but can reduce satisficing and other negative response behaviors.

Lastly, the duration of the survey affects task difficulty and respondent motivation simultaneously, making it critically important to keep the survey as short as possible. Generally, we recommend that web surveys not take the average respondent more than 15 minutes. If you're anticipating that many respondents may arrive at the survey using a mobile device, then the duration should be even shorter (probably not longer than 7 minute), to avoid large numbers of respondents losing motivation and breaking off from the survey.

Data quality

Data quality

The ability to draw correct conclusions or insights from survey data depends on the quality of the data. In the question design context, we focus on two dimensions of data quality: reliability and validity. Understanding these two constructs is critical to the work of question design.

MEASURING DATA QUALITY

Reliability Reliability refers to the extent to which a measurement process provides internally consistent or repeatable results. Internal consistency in this context typically means that items that theoretically should be correlated actually are correlated when examined.

For example, it is well-established that human height and weight are very strongly correlated. Consequently, a researcher may expect that survey measures of height and weight should be correlated. If the researcher does not find this expected association upon examining a dataset, it should be a cause for concern about the quality of the data. A similar example would be the correlation between political ideology and political party among Americans; if these things are not correlated then there may be a data quality issue. Researchers typically test this dimension of reliability by computing the correlations between questions they expect to be associated.

Repeatability, or test-retest reliability, is another dimension of reliability that may be more familiar to most researchers. Using the example from before, if a researcher gathers self-report data on respondent height and weight on Monday, and then again from the same respondents on Tuesday, it would be very concerning in terms of data quality if the responses were substantially different.

To assess this dimension of reliability, researchers will commonly ask the same question or questions multiple times in a survey with the expectation that different responses to the same question may be an indicator of low-quality data.

Validity Validity is the dimension of data quality that researchers are often most concerned about, and it generally refers to the extent to which a measurement process is actually measuring the thing that it is intended to measure. There are a handful of ways that validity can be operationalized:

+ **Construct validity** how closely does the measure "behave" like it should based on established measures or the theory of the underlying construct?

+ **Content validity** how well does the sample of questions asked reflect the domain of possible questions?

+ **Predictive validity** what is the strength of the empirical relationship between a question and the gold standard?

+ **Face validity** what does the question look like it's measuring?

Unfortunately, the last definition is the one that we see used most commonly by researchers to evaluate validity. The apparent or face validity of a survey question is a poor criterion, but this doesn't prevent researchers from using it to evaluate the likely quality of the data collected with a question. The other approaches to assessing validity listed above are much more robust, despite being more difficult to implement.

Ideally, some combination of construct, content, and predictive validity would be applied when assessing the validity of a survey question.

Taken together, reliability and validity are the basis for what we broadly refer to as "data quality" in the context of survey question design. Ensuring that survey questions produce high-quality data is incredibly important for drawing correct conclusions.

EVALUATING THE QUALITY OF SURVEY QUESTIONS

To summarize the challenge faced by survey researchers: the survey response process, while a simple interaction at face value, is actually made quite burdensome by a range of extrinsic factors in the respondent's life at the moment of response and by intrinsic features of surveys that unnecessarily complicate the process for respondents. Survey designers can control intrinsic burden with better, simpler design.

Good survey design is essential, because even slight extrinsic burden, out of the control of the designers, can result in respondents satisficing, rather than optimizing their survey answers, which is a threat to data quality. It's important that the researcher manage the experience of the respondent so that their attention is as focused on the survey items as possible.

Designing questions

Designing questions

DESIGNING QUESTIONS TO MAXIMIZE DATA QUALITY AND REDUCE RESPONDENT BURDEN

Choosing the best response format

In this section we discuss the most commonly used survey question types and highlight when each may be most appropriate. These question types include open-ended questions, ranking questions, and rating questions. Each type has a distinct set of benefits and disadvantages, and knowing when to use each can make a huge difference in the quality of your data. It is important to weigh these benefits and disadvantages carefully when designing a survey question.

Open-ended questions

Open-ended responses to questions can be some of the most reliable, valid, and rich data collected in a survey. Unfortunately, in the web survey context they are most often used in ways that don't allow the researcher to realize their full potential.

The reason that open-ended questions are not used more often is that respondents generally do not like them very much—this is because they are more cognitively demanding and time-consuming to provide high-quality answers to. But if you use open-ended questions judiciously, it is possible to avoid both excessive cognitive demand and long completion times.

Asking open-ended questions that are very specific and easy to answer will allow you to realize the benefits of this powerful question type and avoid annoying or fatiguing your respondents. For example, in many cases, when a number is being requested, it is best to use an open-text question. Asking for a person's age in years and letting them type the number is easier for the respondent to do than selecting from a drop-down list. It's also more precise than selecting an age group.

Most commonly, they are used as 'Other (specify): _____' response alternatives to categorical questions for which the researcher either doesn't know the entire universe of possible responses or feels the list would be too long to present. The other very common usage is in the format of general "feedback" or "comments" boxes where the respondent is simply expected to type comments about anything related to the topic of the survey or the survey itself.

However, the 'Other (specify): _____' response alternative comes with a cost. Research indicates that respondents tend to select options that are provided rather than typing in their own response; this can lead to underestimates of the options that are written in by the respondents. Furthermore, the magnitude of the underestimate can be difficult to accurately assess. In general, the best practice is to use an open-ended response format when the full range of possible responses cannot be provided in a list for the respondents to select from or if the list would be so long that respondents might not carefully read each alternative. Using the 'Other (specify): _____' option should be avoided because the resulting data may be misleading.

Rating

Rating questions are the most commonly used question type in web surveys. These questions obtain assessments of one object at a time, typically along a single dimension (e.g., satisfaction, importance, likelihood, etc.)

These questions are popular for a number of good reasons.

1. They are comparatively easy for respondents to answer, both in terms of the cognitive burden of the question and the provision of a response. (Unsurprisingly, respondents prefer the rating question type over ranking questions.)

2. Rating questions generally have shorter completion times than ranking questions.

3. Data from rating questions is typically more straightforward and easier to analyze than the data from ranking questions.

However, rating questions do pose some tradeoffs when compared with alternatives.

+ Lower effort on the part of respondents may produce lower data quality - this is the tradeoff of using questions that do not require as much cognitive effort on the part of respondents to produce a reasonable answer.

+ Responses tend to be a bit less reliable and change more over time.

+ Rating questions are susceptible to response styles, which describes the tendency of some respondents to consistently avoid the ends of rating scales (or always give answers at the ends), give acquiescent responses, or give straightlined responses.

Ranking

Ranking is a powerful and under-utilized question type that is becoming increasingly popular as researchers outside of the field of market research embrace conjoint designs for their projects. But even apart from conjoint, ranking questions have a huge amount of value for many kinds of research questions.

Ranking questions have a couple of key advantages.

1. **They provide comparisons between multiple things at one time** When a consumer enters a convenience store to purchase a soda, they will typically only purchase one of the many options available. They may report on rating questions that both Coca-Cola and Pepsi are equally preferred, but when faced with the decision they are likely to only choose one. A ranking question type is able to force the differentiation between these items in a way that a rating question type cannot. Consequently, ranking questions are useful when the desired outcome is a comparison or choice. When the goal is to evaluate relative performance, importance, preference, satisfaction, and many other measures, the ranking question format is worth considering.

2. **They are often more reliable than rating questions** particularly for items at the ends of the ranking scale where respondents typically have the strongest differentiation. This means that respondents are able to repeatedly provide the same rankings of items more consistently than when they are asked to provide ratings.

Rating vs. Ranking

When choosing between rating and ranking, it is important to evaluate which is most appropriate for the research being conducted. In some cases, ranking will be the best tool; in other cases, rating will be better. As a general rule of thumb, when life forces a choice between alternatives, ranking may be the better option. In these cases, you want the question type that more closely reflects the decision process that the respondent will engage in when making a choice. In most other cases, rating will be the best choice.

How many scale points should be used?

The primary goal for choosing the number of points to include in rating scales is to differentiate between respondents as much as validly possible while still maintaining high reliability or consistency in responses (e.g., if the same question is asked twice in the survey, we would typically hope for the same response on the scale.)

Determining the number of scale points is a balancing act, which creates a tension when trying to maximize data quality. Including more scale points might differentiate responses more, whereas fewer scale points might produce more reliability.

Fortunately, survey methodology research on this subject provides some guidelines for best practices that enable optimal validity and reliability. The results of this research suggest that the optimal number of scale points ranges from 5 to 9—with fewer points, you lose the ability to differentiate as much as you could between respondents, and with more scale points, the reliability of responses tends to drop off.

For example, in assessing how probable a customer is to make a return visit to a store or website, it does not make sense to try to measure a negative probability because, mathematically, probability ranges from 0 (definitely won't visit) to 1 (definitely will visit). So in this example the appropriate scale is a 5-point unipolar response scale.

If the construct being measured can range from some negative value to some positive value (this is called a bipolar construct), then a 7-point bipolar response scale will be best. For example, in assessing satisfaction with a product or service, the construct (satisfaction) can range from extremely dissatisfied, a negative value of satisfaction, to extremely satisfied, a positive value of satisfaction. As a result, a 7-point bipolar response scale is the best choice for this question.

There is one important caveat to the advice provided above. Web surveys are taken on more types of devices now than ever before, and an important consideration for web survey designers is the screen size on which their questionnaires are likely to be viewed. If many respondents are expected to come to the survey on a cell phone or tablet, then it becomes important to ensure that response scales will render on the screen without any need for horizontal scrolling. In this case, 5-point response scales may be preferable to 7-point scales, even for bipolar constructs.

You may have noted that the recommendations above are both for scales with odd numbers of scale points, meaning that there will be a midpoint. Many researchers like to force respondents to choose one side of the scale or another, thinking that a midpoint response is equivalent to a "don't know" or " no opinion" response. But the empirical research actually indicates that this is not the case for respondents. In fact, reliability seems to be highest when a midpoint is provided, meaning that forcing respondents to take one side or another may introduce inaccuracy into the data.

Labeling response scale points

There are a few different ways that response scales are labeled. Some use verbal labels and others use numeric. Some only have the endpoints labeled, others label the ends and the midpoint, and yet others label each point. So which approach is best?

It is important to keep the goals of scale point labels in mind when making this decision. These goals are:

+ The meanings of each scale point should be easy for respondents to interpret

+ The meaning of each scale point should be clear (unambiguous)

+ All respondents should interpret the meanings of each scale point identically

+ The labels should differentiate respondents from each other as much as is validly possible

+ The resulting scale should include points that correspond with all points on the underlying construct's continuum

From these goals, the resulting best practices include labeling all scale points. The meaning of unlabeled scale points is ambiguous. As a result, in an answer scale with both labeled and unlabeled points, respondents may be attracted to the points that have labels. This tendency to select labeled scale points will produce a bias toward those points, resulting in clusters of respondents at labeled points. In short, partially labeled response scales may produce biased data due to respondents being attracted to scale points that are easier to interpret.

Another best practice is to verbally label rating scale points rather than numerically labeling them since numbers alone are often ambiguous in their meaning and difficult for respondents to interpret. In general, it is best to omit numeric labels for rating scales altogether, because they may be interpreted differently by different respondents.

Not only should the scale points be labeled verbally, but the labels should also match the construct being asked about in the question text. Researchers often attempt to use the same response scale for many questions, most commonly the agree-disagree scale. This practice may lead to less thoughtful answers on the part of respondents and may also make interpretation of the questions and response scales more complex. For example, if a researcher wishes to assess satisfaction, they may make the statement, "I was satisfied with _____," and provide a response scale ranging from "strongly agree" to "strongly disagree." In this case, the respondent must make an assessment about their degree of satisfaction and then determine how that assessment maps onto agreement with a statement, rather than simply answering a question about satisfaction.

A better approach is to determine the construct of interest (satisfaction, in this case) and then ask a question with response options that are specific to that construct. The resulting question using this approach would be, "How satisfied or dissatisfied were you with _____?" and the response scale mapped onto this construct would range from "extremely satisfied" to "extremely dissatisfied."

In this formulation of the question, using a construct-specific response scale, the respondent makes their assessment about their degree of satisfaction and then can provide an answer without needing to make the additional step of determining how their assessment of the construct maps onto agreement with a statement. Another reason to use construct-specific response scales is that they may reduce straightlining. When a respondent is asked to give answers on the same response scale many times, there is an increased chance of fatigue setting in, and it becomes more likely that they will begin to satisfice by straightlining. This is most prevalent with matrix or grid question types, but can also happen when questions are presented individually.

Lastly, respondents presume equal spacing of scale points and the underlying continuum. This means that respondents expect that the response scale covers the entire range of possible answers along the target construct and that the scale points provided are evenly divided across that range. As a result, scales should be balanced across the continuum of possible responses and the labels chosen for the scale points should reinforce this. If a bipolar construct (negative to positive values) is measured using a unipolar response scale (zero to positive values), the lack of correspondence between the response scale and the construct will result in biased data.

QUESTION WORDING

Researchers have known for decades that question wording matters, and that even subtle changes in the wording of a question can produce substantial differences in the resulting data. With this in mind, it is important to use questions that have been validated or tested to ensure that the ways that respondents interpret the question match exactly what the researcher believes they are measuring. Without this testing, the question may inadvertently measure something different than intended without the researcher ever finding out.

Goals of question wording

When writing a question, there are a few guiding goals that are worth keeping in mind.

1. **Ensure clarity by only mentioning a single subject and construct** Failure to do this can result in "double-barreled" questions. For example, a question that asks, "How satisfied or dissatisfied were you with our product selection and quality?" is problematic because it mentions two subjects: (1) the selection, and (2) the quality of products. The respondent may feel extremely satisfied with the quality of the products while being extremely dissatisfied with the selection. For that respondent, it is unclear how to respond to the question as it is worded. The better approach would be to create separate questions to address each construct.

2. Make sure the question means the same thing to all respondents Put another way, each respondent should interpret the meaning of the question identically. If respondents have different interpretations of the meaning of the question, the resulting data will be invalid because it will be impossible to determine what question each respondent thought they were answering.

3. Use words economically It is important to use as many words as are needed to convey the idea of the question clearly to all respondents, but additional words present more opportunities to introduce elements that could confuse respondents. Keeping questions short also means that respondents can read them more quickly, which should help keep the survey duration short.

Word choice guidelines

There are a number of guidelines or best practices for word choice when writing survey questions. The first is that words used in survey questions should have only one meaning—this is easy to verify using a dictionary. Words and sentences should be simple, to maximize the ease of reading and comprehension. A useful rule of thumb is that words with fewer syllables and sentences with fewer words are typically simpler. Readability scores calculated using online tools are often useful for assessing the complexity of the words and sentences that form a question.

The conventional wisdom—which has been supported by most empirical research on the topic over the years—suggests that, in general, questions should be worded to:

+ Be simple, direct, comprehensible

+ Not use jargon

+ Be specific and concrete (rather than general and abstract)

+ Avoid ambiguous words

+ Avoid double-barreled questions

+ Avoid negations

+ Avoid leading questions

+ Include filter questions

+ Read smoothly out loud

+ Avoid emotionally charged words

+ Allow for all possible responses

If you follow these recommendations, it is much less likely that your survey questions will confuse or frustrate your respondents, and your data are more likely to be valid and reliable. Surprisingly, many researchers fail to test their own survey questions by reading them out loud or asking others to state in their own words what they think a question is asking. These simple steps would solve many problems that we see with survey questions.

Using an existing question vs writing your own

Using an existing question vs writing your own

In many cases, when deciding how to measure something, there may be existing questions, such as the Net Promoter Score questions or more domain-specific questions. Many researchers want to know whether or not they should try to find an existing question for their survey or if they should design their own. There are three general types of cases to consider when making this decision:

+ If there is an existing question that perfectly matches your research needs, use it.

+ If there is no existing question that perfectly matches your research needs, write your own.

+ If there is a question that seems to be a somewhat close match for your research needs, then consider using it and pre-testing it against one that you have written. If space allows, you may actually want to field both questions, but hopefully the pre-test is decisive.

It is certainly a best practice to use validated survey questions that have known measurement properties, but fielding a question that doesn't collect exactly the data needed to address your research question can be a waste of both yours and your respondents' time. It is important to remember that whatever question you ask should always collect the best possible data for your research question.

SECTION 5

Conclusion

Conclusion

Whether your research is driven by the need to create amazing brand or product experiences, to thrill customers and engage employees, or to generate reliable estimates to better understand critical social problems, the cornerstone of quality, actionable research is well designed questions.

Writing good survey questions is both science and art, and it's not intuitive for most people. Applying the best practices from the academic survey methodology literature outlined in this handbook (and from many other resources) will help you gather the most valid, reliable, and accurate data—and, as a result, the best insights.

In particular, being aware of what is going on in the heads of your respondents and how the design decisions made when creating a survey can either enable positive or negative response behaviors can play a powerful role in shaping the quality of your data. Question wording, response option wording, and response format can each have substantial effects on the quality of your data, particularly if they enable respondents to apply negative response strategies, such as satisficing.

Every question should be designed to collect the best possible data for the research question that it is intended to address. At a minimum, this may mean avoiding response formats and scales that enable satisficing response strategies. In other cases, this may mean designing new questions when existing ones are not tailored to your respondents' experiences.

Throughout this handbook we have emphasized the importance of pre-testing any new or edited survey. This pre-testing process is the single most important tool that a researcher can use to catch errors, confusing or inadvertently challenging questions. It doesn't have to be complicated. Asking colleagues or friends to take a survey and provide feedback is a simple step that many researchers fail to take when writing a survey.

There are many potential sources of error that can affect the quality of survey data. Some of these are outside of the control of the researcher; others can be influenced directly or indirectly by decisions made in the design, fielding, and analysis phases of the research process. Question design is one of the opportunities that researchers have to directly influence the quality of the data that a survey is used to gather. With all of the other ways that data quality can be degraded, there is no reason to miss the opportunity that question design offers to create a more intemaximize data quality We hope that this handbook provides some actionable recommendations that will help you design better survey questions and generate higher-quality data and results for your research.

BIBLIOGRAPHY

Cannell, Charles F, Peter V Miller, and Lois Oksenberg. 1981. "Research on Interviewing Techniques." Sociological Methodology 12: 389–437.

Cialdini, Robert B, John T Cacioppo, and R Bassett. 1978. "Low-Ball Procedure for Producing Compliance: Commitment Then Cost." Journal of Personality.

Krosnick, Jon A. 1991. "Response Strategies for Coping with the Cognitive Demands of Attitude Measures in Surveys." Applied Cognitive Psychology 5(3): 213–36.

Krosnick, Jon A. 1999. "Survey Research." Annual Review of Psychology 50(1): 537–67.

Malhotra, Neil. 2009. "Completion Time and Response Order Effects in Web Surveys." Public Opinion Quarterly 72(5): 914–34.

Tourangeau, Roger, Lance J Rips, and Kenneth A Rasinski. 2000. The Psychology of Survey Response. Cambridge University Press.

ADDITIONAL READING

Krosnick, Jon A. and Stanley Presser. 2010. "Question and Questionnaire Design" in the Handbook of Survey Research

Schaeffer, Nora Cate, and Stanley Presser. 2003 "The Science of Asking Questions." Annual Review of Sociology

Sudman, Semour. and Norman Bradburn. 1996. "Thinking About Answers"

Vannette, David L, and Jon A. Krosnick. 2014. "Answering Questions: A Comparison of Survey Satisficing and Mindlessness" in The Wiley Blackwell Handbook of Mindfulness

Vannette, David L, and Jon A. Krosnick. (forthcoming). "The Palgrave Handbook of Survey Methodology"

qualtrics^{XM}

WANT TO SEE HOW IT ALL WORKS?
Contact sales@qualtrics.com for more
information on how to get started.

BOOK A DEMO