Psychological Research Methods

Jamie DeCoster

Department of Psychology
University of Alabama
348 Gordon Palmer Hall
Box 870348
Tuscaloosa, AL 35487-0348

Phone: (205) 348-4431
Fax: (205) 348-8648

May 9, 2001

These notes were prepared with the support of a grant from the Dutch Science Foundation. I would like to thank Reint Jan Renes, Mirjam Tazelaar, and Jan-Willem van Prooijen for comments made on earlier versions of these notes. If you wish to cite the contents of this document, the APA reference for them would be

DeCoster, J. (2001). *Psychological Research Methods.* Retrieved <month, day, and year you downloaded this file> from http://www.stat-help.com/notes.html

For future versions of these notes or help with data analysis visit
`http://www.stat-help.com`

Contents

Chapter 1

Developing a research question

The first thing that researchers must do when conducting research is to determine the *research question* that they want to investigate.

1.1 Sources of potential questions

- There are many ways to come up with a research question. Possibly the most common source of research questions involves examining a *theory*. Here we are defining a "theory" as a set of related principles that are designed to explain a broad set of real-life events. Not all researchers try to develop theories, and not all research is theory-driven. However, many do use theories to guide the research questions they choose to investigate.

  Research questions can be derived either from attempts to promote or to dispute a theory. If you wish to promote a theory the first thing you can do is question whether the predictions of the theory are true. If you are working from a pre-existing theory, however, this was usually done by the original researchers. In this case you can question whether the principles of the theory can be applied to new, previously unconsidered phenomena. You could also develop questions about *why* the theory works. If the theory is described at a certain level (such as reading comprehension) you might think about what lower-level process (such as perception and word recognition) might be driving the theory.

  To dispute a theory you might question the validity of existing findings that support the theory. If you believe that a study suffered from methodological or statistical errors you might question what a correct version of the study might find. If the existing support appears valid you could also question whether there are factors that might limit where the theory applies. You can consider whether the theory is only successful in explaining phenomena under particular circumstances or with particular populations. Finally, you might question whether there are external factors that must be present or absent for the theory to accurately predict behavior.

- A research question sometimes suggests itself while you are reading about or listening to research that someone else has conducted. When you are reading articles you should write any thoughts you have at the time in the margins (assuming you aren't using a library copy!). These ideas can be excellent sources of research questions. Similarly, you should also consider writing down ideas that you have whenever you hear someone giving a research presentation. People often take notes on what the speakers say but fail to write down ideas of their own that are provoked by the talk.

  While it is certainly possible to wait until you finish an article or listen to the complete talk before recording your ideas, this is generally not as effective as writing them out at the time you have them. First off, when you wait you run the risk of forgetting your ideas. This is quite possible since before you finish you will be reading or hearing about many other ideas on the same topic which might interfere with your memory. Second, often when people have flashes of insight their ideas are not fully formulated in linguistic terms. You will be better able to find the right way to express your question if you try to do so immediately than if you have to try to work from your memory.

- You can get ideas for research questions from arguments and discussions with your teachers, students, and colleagues. This can be a more fruitful source than either articles or presentations since as a discussant you must actively think about the topic in order to come up with reasonable contributions to the conversation. Your partner can also be valuable in helping you evaluate how important or interesting your question might be.

  Just like with articles and presentations, you are probably best off if you write down potential research questions as soon as you think about them, even if this means disrupting the flow of the conversation. While this may seem to be rude, you will be surprised at how many people feel complimented when you indicate that the discussion was meaningful enough for you to want to write down some notes.

- The final major source of research questions is through observation of things around you in everyday life. You might see or hear about an unusual phenomenon that you think might be worth investigating. Or else you may simply want to learn more about a common aspect of life. When generating research questions about life events you can use the five "W's" of journalism: Who is involved in the event? What occurs during the event? Where does the event occur? When does the event occur? Why does the event occur?

1.2 Evaluating a research question

An important thing to keep in mind is that not all research questions are equal. Some topics by their very nature will never lead to important or interesting research. Here are some things you can ask to help you evaluate a potential research question.

- *To what extent has the question already been examined?* There is very little to be gained by investigating a question whose answer is already known. After deciding on a preliminary research question you should search the literature to see what studies have already examined the question. Studies that investigate minor variations of a well-documented phenomenon are typically less valuable than those that help open up the way to new avenues of research

- *Is the question of social or scientific importance?* You need to consider exactly what will be gained by finding the answer to your research question. Think about both the applications of the research to everyday life as well as the implications it might have for future research on similar topics. While a study does not necessarily need to have direct real-world applications to be meaningful, their presence does increase the potential value of the research. Similarly, a study can be atheoretical and relatively independent of the rest of the literature and still be important. However, research that speaks to controversial theoretical issues in the field will likely have a larger overall impact than research that does not.

- *Is the question appropriate to the researcher?* Just because you develop a novel research question with social and scientific importance doesn't mean that you should feel compelled to investigate it. Not all research questions are right for all researchers. First and foremost you must decide whether you actually have an interest in investigating the question. The work of researching a topic will be much more enjoyable if you can maintain your interest in it throughout the project. Additionally, people are able to write much better and much more convincingly about a topic when it is personally important to them.

  You should also consider whether you have sufficient background knowledge to be able to design a study to examine the question. If you do not, you should either commit yourself to gaining this knowledge or else choose a new topic. Unless you know the literature you will not be able to judge whether your question has already been answered. You will be unfamiliar with the standard research methods used in the field, possibly requiring you to spend a long time developing your methodology before you see interpretable results. Finally, you will find writing your introduction and discussion sections to be very difficult since you don't really know the research to which you should relate to your study.

  Finally, you should decide if the question is suitable for the general type of project that you had in mind. Questions that are appropriate for a full-length journal article could never be investigated in

time for a class assignment. Similarly, simple questions that can be answered with a single study are often too limited in scope to be accepted by journal editors. You should give serious thought to the amount of effort that would be needed to sufficiently answer the question. If it requires more than you are able to give then you should either refine the scope of the question or choose a new topic.

Chapter 2

Reviewing the Literature

Once you have a basic idea of what it is you want to investigate you should see what other researchers have done on related topics. This will make it easier for you to decide on the specific details of your own study. You will also need this information when writing the Introduction and Discussion sections of your paper.

2.1 Searching the literature

- *Computerized Indices.* A number of databases are available on CD-ROM or over the internet. These will allow you to use keywords, and Boolean combinations of keywords, to locate articles relevant to your analysis.
    - Selecting the keywords for your search is very important. First, you should determine the basic structure of what you want in your search (such as PRIMING and IMPRESSION FORMATION). You should then determine the synonyms that would be used for these terms in the database. If the database assigns subject words to each study, you should make sure that your list of synonyms includes all the relevant subject words. Many databases publish a thesaurus to make this easier. For PsycInfo the name of this thesaurus is "Thesaurus of Psychological Index Terms" (APA, 2001).
    - Most databases support the use of wildcards, which you should use liberally. To locate research on priming in PsycInfo we might use the search term PRIM*, which would find studies that use the terms PRIMING, PRIMES, PRIMED, and other words beginning with PRIM.
    - You should then enter your search into the database. Each construct will be represented by a list of synonyms connected by ORs. The constructs themselves will be connected by ANDs or NOTs. In the example above we might try: (PRIM* or IMPLICIT MEMORY) and (IMPRESSION or JUDGMENT).
    - If your initial search produces a large number of irrelevant studies related to a single topic, you might try to keep them out of further searches by introducing a NOT term to your search.
    - Keep in mind that different indices use different terms, so you may have to define your search differently when working with different databases.
    - Whenever you conduct a computerized search you should record the name of the database, the years covered by the database at the time of the search, and the search terms you used. This can help prevent redundant searches if you decide to look through the literature again at a later point in time.
    - The databases most commonly used by psychologists are:
        1. PsycLit/PsycInfo (PsycLit is the CD-ROM version, and is less complete)
        2. ERIC (Educational Resources Information Center)
        3. Dissertation Abstracts Online
        4. ABI/Inform (a worldwide business management and finance database)

5. Sociological Abstracts (sociology literature)
6. MEDLINE (biomedical literature including health care, clinical psychology, gerontology, etc.)
7. Mental Health Abstracts

There are also a number of databases available within more specialized research areas.

- *Descendant search.* If you can locate an important study that was performed at an early date you can use the SSCI (Social Science Citation Index) or SCI (Science Citation Index) to locate later articles that cite them in their references. This is a very nice complement to the standard computerized search, and can now be performed fairly easily since both indices are available on CD-ROM.
- *Ancestor search.* You should check the references of papers you already have to see if they contain any relevant studies that you are unaware of.
- *Research registers.* Research registers are actively maintained lists of studies centered around a common theme. Currently there are very few research registers available for psychological research, but this may change with the spread of technology.
- *Reference lists of review articles.* Review articles are often a fruitful place to look for relevant studies.
- *Hand search of important journals.* If there is a journal that is particularly relevant to your topic then you might go back and read through the table of contents of that journal for the years that there was active research on your topic. You might make use of *Current Contents*, a journal containing the tables of contents of other journals.
- *Programs from professional meetings.* This is a particularly good way to locate unpublished articles, since papers presented at conferences are typically subject to a less restrictive review (and are therefore less biased towards significant findings) than journal articles. Probably the two most important conferences in psychology are the annual meetings of the American Psychological Association (APA) and the American Psychological Society (APS).
- *Letters to active researchers.* If there is a particular person whose research is highly relevant you may directly ask them for recommended reading. When trying to locate people you may want to make use of:
  - Academic department offices/Department web pages
  - Alumni offices (to track down the authors of dissertations)
  - Internet search engines (such as `http://www.switchboard.com` or `http://people.yahoo.com`)
  - APA, APS membership guides

2.2 Summarizing articles

- It can often be useful to compose short summaries of important articles when you need to write a literature review. This allow you to develop your interpretation of the article soon after you read it instead of weeks or months later when you write your review. Summaries are also useful when you want to go back and refresh your knowledge of the literature at a later point in time.
- Ideally a summary should include the most important aspects of an article in a form that can be read in a relatively short period of time. Summaries that are too short may miss important points and provide an inaccurate view of the article. Summaries that are too long will include unimportant information making it more difficult to extract the main points of the article when later reviewing the summary.
- An important step is to determine what articles you should summarize. This will depend in part on the reason you are conducting the literature review. If you are writing an actual review paper then you will likely want to summarize each article you wish to include in the paper. The purpose of the review in this case is to provide a comprehensive description of studies in the literature. To help ensure that your review is unbiased you should equivalently examine all the relevant articles.

When you are reviewing the literature as part of the introduction to an empirical article, on the other hand, you will typically only want to summarize a selection of the relevant studies. The purpose of the review in this case is to provide a context in which your study appears important and justified. Your review is not intended to be comprehensive. Instead it is designed to illustrate the reason that you chose your topic of research. You are therefore justified in concentrating your efforts on articles you see as more important.

- A well-written summary of an empirical article should take up less than one single-spaced typewritten page. Summaries of review articles may be somewhat longer depending on the size of the literature and the way that it was reviewed.
- There are five aspects of an empirical article that should always be documented in a summary
    1. *APA reference for the article.* This should appear at the top and should be visually separate from the rest of the summary. You can accomplish this by putting a box around the reference, or simply by including blank lines between the reference and the remainder of the summary.
    2. *Research question examined by the article.* You should state the main topic the article investigates. You should also include a short description of any theories that are important to the studies. You generally will not want to record arguments for the importance of the research question.
    3. *Abstract descriptions of studies within the article.* You want to provide a basic idea of the experiments and surveys conducted as part of the article. Provide specific details only to the extent that they are important to understanding the results of the study.
    4. *Major findings.* Report the important results of each study in the article. You will generally not record the specific statistics. Instead describe the overall pattern of results.
    5. *Conclusions.* You should describe both the specific implications drawn from the results of each study as well as any general conclusions made from the article as a whole. Be sure to report the implications the article may have for theories in the literature.
- There isn't a specific procedure for summarizing the contents of a review article since their organizations vary greatly. However it is probably best to order the contents of your summary in the same way as the original article.
- In addition to summarizing the contents of an article you may also want to write down how the article relates to your own research project.
- You should try to use complete sentences when writing your summaries. First, the additional thought required to write the information down in a comprehensible form may give you some greater insights into the article. Second, you will make writing your literature review easier since you will be able to draw more directly from the contents of your summaries.

Chapter 3

Designing your Study

3.1 Generating research hypotheses

- Once you have found a satisfactory research question the next step in conducting research is to come up with one or more *hypotheses* that will help answer the question. A hypothesis is simply an unproven theoretical prediction that answers some aspect of your research question.
- A given hypothesis does not necessarily need to provide a definitive answer to the question you want to investigate. It simply needs to explore and hopefully answer some aspect of your question. A single research question can be the source of many different hypotheses.
- A hypothesis should be falsifiable. That is, there should be some form of evidence that would indicate that the hypothesis is incorrect. If a hypothesis is able to accommodate any pattern of data then finding evidence for it won't really tell us anything about the phenomenon of interest.
- Hypotheses typically involve statements about the relationships between two or more *variables*. By definition, a variable is some feature that can be categorized or measured with a number. A variable takes on different *values* to reflect the specific characteristics of a given observation. Variables can refer to features of the environment, such as the presence or absence of noise, or they can refer to features of an individual, such as the time it takes someone to respond to a question.

3.2 Relationships between variables

- There are two different types of relationships that can exist between a single pair of variables. A hypothesis can propose a *causal relationship* indicating that changes in one variable are directly responsible for changes in a second variable. For example, researchers have proposed that attractive sources of information are more persuasive than unattractive sources. This hypothesis would state that if you keep everything else constant, just changing the attractiveness of a person presenting an argument will change the likelihood that a listener will accept the argument.

  A hypothesis might also propose a *correlational relationship*, simply stating that you can predict one variable from another without making any claims about causality. For example, researchers have proposed that how much a child reads is associated with their intelligence. This hypothesis makes no claims about what causes the relationship, but does propose that knowing the value of one of the variables will allow accurate predictions of the other.
- When discussing a causal relationship researchers call the variable causing changes the *predictor variable*. The variable being affected by the change is called the *response variable*. These terms are also occasionally applied to variables with a correlational relationship if there is a justification for saying that one particular variable is being used to predict the second. Some people refer to predictor variables as *independent variables* and response variables as *dependent variables*.

- If you have at least three variables you have some additional options for your hypotheses. You could propose that one variable *mediates* a causal relationship between two other variables. This would mean that a predictor variable causes changes in a response variable, but only through changes in a third mediator variable.

  The idea here is that the predictor variable has a causal relationship to the mediator which in turn has a causal relationship to the response variable. It is also possible to have *partial mediation* where the predictor has a direct influence on the response variable in addition to the mediated relationship.

  For more information about how to demonstrate a mediated relationship see Baron and Kenny (1986).

- A hypothesis could also propose that one variable *moderates* the relationship between two other variables. This means that the exact relationship between two variables depends on the value of a third variable. This is exactly the same thing as having an *interaction* between predictors in regression or ANOVA. Both correlational and causal relationships can be moderated by third variables.

  The idea here would be that you can successfully predict the response only if you have the values of both predictors. Knowing the value of either predictor variable itself is not sufficient to predict the response because the relationship between the first predictor and the response depends on the value of the second predictor, and the relationship between the second predictor and the response depends on the value of the first predictor.

3.3 Evaluating hypotheses

Here are some questions that you can ask to help you evaluate potential hypotheses.

- **Is the hypothesis novel?** You generally do not want to test the same hypothesis as another researcher unless you have reason to suspect errors in their study. Exact replications are not greatly valued by the scientific community.

- **How general is the hypothesis?** Hypotheses that are framed in more abstract terms can be applied to a wider variety of situations, increasing their potential impact.

- **Does the hypothesis concern important variables?** People pay greater attention to research that they feel can have a personal influence on them. Consider how often people talk about the variables in the hypothesis in everyday life (social importance) and the extent to which other researchers use those variables in their own studies (scientific importance).

- **Does the hypothesis propose anything that is unusual or unexpected?** Others will have more interest in findings that defy common sense or provide a new way of thinking about a phenomenon.

- **How likely is it that the hypothesis is correct?** There is always some element of uncertainty surrounding a hypothesis. Tests of incorrect hypotheses provide much less interesting information than tests of hypotheses that turn out to be correct.

  However, there is usually a direct tradeoff between how unexpected a hypothesis is and its likelihood of being correct. The more unusual (and therefore the more interesting) your hypothesis is the less likely it will work out the way that you want. How much risk a hypothesis should have is a matter of personal choice. High-risk research tends to provide more interesting findings, but low-risk research is generally more likely to provide publishable findings. In the long term it is probably best for researchers to pursue a mixture of high-risk and low-risk projects.

3.4 Specifying variables for a study

- After proposing a hypothesis researchers typically must provide some type of evidence to support their claims. While at times this may be strictly through a rhetorical argument, more commonly researchers conduct a study to provide empirical evidence for their idea.

- The first step in designing a study is to specify what variables you will record and how their values will be obtained. This information is called the *design* of your study. The design of your study will always be strongly tied to the hypotheses that you want to examine.
- You will want to include the following types of variables in your design.
  - Variables that are explicitly discussed in your hypotheses.
  - Variables that you believe could potentially influence the relationships proposed in your hypotheses (such as moderators or mediators).
  - Variables that you believe may strongly impact on the response measures you collect, even if this influence is completely independent of the relationships proposed in your hypotheses. These are called *covariates* or *blocking variables*. By recording these variables you can statistically remove their influence from your analyses providing you with clearer results.
  - Variables that you believe may be related to those in your hypotheses but for which you don't predict any specific relationships. These are called *exploratory variables*.
- Once you have decided on the variables you will include you must then determine what type of values the variable will record. Each variable can have either *categorical* values, used to represent group distinctions or *numeric* values, used to represent features that vary in measurable quantities.

  It is possible to use a categorical variable to hold information about something that is really numeric. For example, you could use a categorical variable to code a person's age. Values of the variable would represent age groups rather than a person's specific age.
- You additionally need to decide the way you will obtain the values for each variable. While there will be many differences in the specific procedures you use to collect each variable, there is one distinction that is particularly important when determining your study design. Specifically, a variable can either be *manipulated*, where their values are forced by the experimenter, or *measured*, where their values are simply observed and recorded during the study.

3.5 Testing causal hypotheses

- Merely observing that there is a relationship between a pair of variables X and Y cannot be taken as evidence that changes in X cause changes Y. There are two other types of relationships that can produce this same result.
  1. Changes in Y cause changes in X.
  2. Changes in both X and Y are caused by a third variable Z.
- To infer a causal relationship between a predictor variable and a response variable you must show the following.
  1. There is a relationship between the predictor and the response.
  2. Changes in the predictor precede changes in the response.
  3. There are no other possible causes for changes in the response other than changes in predictor.
- If you wish to test a causal hypothesis your predictor variables can be either manipulated or measured. However it can be very difficult to rule out alternative explanations for changes in your response variable if your predictor variable is measured. There is always the possibility that your predictor variable does not actually cause changes in your response variable but instead is simply correlated with the actual factor responsible for the changes.
- The easiest way to demonstrate that a predictor variable causes changes in a response variable is to perform a study that manipulates the predictor variable, measures the response variable, and keeps all other variables constant. If you can show a relationship between the manipulated predictor and

the response variable then you know that there is a causal relationship. You manually changed the predictor so you know that its changes preceded changes in the response. You also know that there are no alternative explanations because you held all of the other factors constant. All three of the requirements to infer causality have therefore been met.

- Any study that manipulates at least one variable is called an *experiment*. A study that only measures variables is called a *survey*. Experiments are primarily used to demonstrate causal relationships while surveys are primarily used to demonstrate correlational relationships.
- The predictor variables in an experiment (except for measured numeric variables) are also sometimes referred to as the *factors*. The values that a factor may take are often called its *levels*. The particular combination of factor levels that a participant is in is called their *condition*. Measured numeric predictor variables aren't generally considered factors because they don't divide your participants into groups. Rather each subject has their own specific measured value.
- The response measures in an experiment must always be measured and should be collected in every condition of the study.
- If you decide that you want to manipulate a variable in an experiment there are two basic ways to do it. You can choose to manipulate your variable *within subjects* or you can manipulate it *between subjects*.
    - In a within-subject manipulation each participant experiences every value of your manipulated variable. Here you will observe the same group of people for each level of your manipulation.
    - In a between-subject manipulation each participant experiences only a single value of your manipulated variable. Here you will observe different groups of people for each level of your manipulation.

  There are alternatives to strict within- or between-subject designs, such that each participant experiences more than one level of your variable, but not all of them. However, these designs are best avoided because they can be extremely difficult to analyze. When designing an experiment you should try to manipulate each factor in either a within-subject or a between-subject fashion. There are no problems with including both within-subjects and between-subjects manipulations in the same experiment, as long as each variable is either one or the other.
- Next you need to specify exactly what conditions you want to include in the experiment. If you only have a single factor then you should have a condition for each level of the factor. If you want to use more than one factor in an experiment the best way to determine your conditions is with a *factorial design*. In a factorial design you have a condition for each possible combination of the factors. For example, assume that you wanted to conduct an experiment examining the influence of both external reward (high or low) and the length of a task (short or long) on task performance. To create a factorial combination you would want to measure task performance under the conditions of high reward/short task, high reward/long task, low reward/short task, and low reward/long task.

  There are other ways you can select conditions from a combination of factors but under most circumstances a factorial design is best. It includes the most information about your variables and guarantees that you will be able to distinguish the influence of one of your variables from all the others. It also allows you test for any interactions between your factors.
- An important requirement to infer causality from the results of an experiment is that all variables not measured or manipulated in the experiment be held constant. If there are other variables that take on different values during the experiment then you can't be sure that the observed changes in the response variable are actually caused by your predictor variables.

  One of the major sources of influential external variables are differences between the types of participants in your between-subjects conditions. If there are differences between the people in your between-subject groups (other than differences directly related to the between-subject factor) then this might produce an alternative explanation for your findings. This would prevent you from being able to conclude that there is a causal relationship between your predictor and response variables.

  There are two ways that you can try to make your groups similar.

1. You can try to equate the participants in your between-subjects conditions. To do this you need to identify the external variables that you think would be most likely to affect your measured variables. You then try to assign them to your groups so that each has approximately the same distribution of each external variable you want to control. This procedure is sometimes called *matching* your between-subjects conditions.
2. You can randomly assign participants to each of your between-subjects conditions. On average the influences of external factors within each group should cancel themselves out. That is, there will be positive and negative influences on your measurements within each group. As long as these influences are truly random your inferential statistics will be able to take them into account.

Although it may be counterintuitive, it is much better to have randomly assigned groups than to have matched groups. You may be able to make matched groups more similar than randomly assigned groups. However, the variability of matched groups is more likely to contain *systematic* differences because you are manually controlling the assignment. It is much better to have random than systematic differences because tests of statistical inference can fully account for the influence of random differences. Systematic differences, on the other hand, go undetected by statistics and can potentially introduce an alternative explanation for your results. Even though you can chose a whole list of variables to control for, there is always the possibility that there is something that you failed to consider when creating your matched groups. Random assignment, on the other hand, balances all potential influences, even those of which you may be unaware.

3.6 Specifying the design of a study

- Information about what variables are recorded in your study and how their values are obtained is called the *design* of the study. You typically report this information in the "Design" section of a paper (which is a subsection of the "Method" section).
- Before attempting to write your design section you need to know the following information.
    - The name of each variable examined in your study.
    - Whether each variable is categorical or numeric.
    - Whether you consider each variable to be a predictor, a response, or neither.
    - Whether each variable is manipulated or measured.
    - For experiments, whether each factor varies between subjects or within subjects.
    - For experiments, how you chose the conditions in the study. Most commonly this will be through the use of a factorial design.
    - How you assigned subjects to the different levels of any between-subjects manipulations.
- Once you have all this information you are ready to describe your design. If you conducted a survey (and therefore have no manipulated variables) you can write your design section using the following procedure.
    1. Start with "The survey measured".
    2. List the information about each variable in the study.
        - If the variable is categorical, write the name of the variable followed by its possible values in parentheses.
        - If the variable is numeric, write the name of the variable followed by the unit of measurement in parentheses. If the variable is measured using a scale you should report the name of the scale (if it is well-known) or simply say "reported on a scale". Sometimes people decide not to report the units of measurement, especially if there are many variables that were measured using relatively unknown scales.

Here is an example of how you might report the design of a survey. "The survey measured each participant's occupation (unemployed, student, laborer, or professional), age (in years), weight (in kg), mood (reported on a scale), and depression level (measured using Beck's Depression Inventory)."

People writing journal articles often omit the design section for surveys since the same information is typically presented in the procedure section.

- If you conduct an experiment (with one or more manipulated variables) you can write your design section using the following procedure.

  1. Start with "The experiment was conducted using a".
  2. State how you chose the conditions of your study.
     - If you have only one manipulated variable and no other categorical predictors you write "N group", where N is the number of levels in your manipulated variable, followed by the name of the manipulated variable in parentheses. You then write "between-subjects design" or "within-subjects design," depending on how manipulated the variable. For example, "a 4 group (time pressure) between-subjects design."
     - If you have two or more factors in a factorial combination you list out the number of levels possessed by each factor followed by the factor name in parentheses, putting "x"s between the different factors. After listing all of the factors you put "between-subjects factorial design" if all the factors varied between subjects, "within-subjects factorial design" if all the factors varied within subjects, or "mixed factorial design" if you have both within- and between-subjects factors. For example, "a 2 (sex) x 2 (task) x 3 (reward condition) between-subjects factorial design" or "a 3 (diet) x 2 (health status) mixed factorial design."
     - If you have used some other method of combining your factors you must fully describe that here.
  3. Report how participants were assigned to the levels of between-subjects factors. If you matched on any variables you would report that here. For example "Participants were randomly assigned to the levels of all between-subjects manipulations" or "Between-subject groups were created by matching participants on hair color, eye color, and fashion sense."
  4. Next you have a sentence for each factor describing the different levels that it can have (unless they are obvious, like for gender). For example, "the two levels of task indicated whether the participant performed an intellective task or a decision-making task."
  5. If your experiment includes any numeric predictor variables you should list their names followed by the units in which they were measured in parentheses. For example, "The experiment also measured each participant's intelligence (using a standard IQ scale)."
  6. Finally you report each of your response variables.
     - If a variable is categorical you state the name of the variable followed by the values it can have in parentheses (unless the values are obvious).
     - If the variable is numeric you state the name of the variable followed by the units in which it was measured. If the variable is measured using a scale you should report the name of the scale (if it is well-known) or simply say "reported on a scale". Sometimes people decide not to report the units of measurement, especially if there are many variables that were measured using relatively unknown scales.

     For example, "The response variables of interest were the length of their response (in words) and the content of their response (past oriented or future oriented)."

Chapter 4

Validity

4.1 Overview

- When designing a study you must always keep the *validity* of your study in the back of your mind. Validity concerns whether the conclusions that you want to draw from your study are appropriate or not. Aspects of your study that might potentially make your study invalid are called *threats to validity*. Issues of validity are of concern to both experiments and surveys.
- Validity is specifically tied to the conclusions being drawn from a study. For any given study there are a set of valid conclusions and a set of invalid conclusions that may be drawn. The important question is whether the design of the study allows the conclusions stated by the researcher. Therefore, before designing a study you should explicitly decide what conclusions you wish to make so you can spot potential threats to the validity of your study.
- Drawing on Cook and Campbell's (1979) categorization we are going to consider four types of validity: construct validity, internal validity, statistical conclusion validity, and external validity. However, these "types" just provide a useful way to think about the different ways that a study could be valid or invalid. You might discover a threat to validity that doesn't fit well into the category structure. The important thing is to consider whether the method of the study actually allows the stated conclusions to be drawn.

4.2 Construct validity

- Construct validity concerns the extent to which the variables in your experiment accurately and exclusively represent the corresponding theoretical constructs.
- When assessing the construct validity of a variable you should consider both *convergent validity* and *divergent validity*. A variable with convergent validity shows relationships with other variables that should be related to the theoretical construct of interest. A variable with divergent validity has no relationships with other variables that are theoretically unrelated to the construct of interest.
- Manipulations lack construct validity either when they fail to alter the theoretical construct of interest or when outside variables are inadvertently changed along with the construct of interest.
- Measurements lack construct validity either when changes in the construct of interest fail to cause changes in the measure or when changes in outside variables can cause changes in the measure when the construct of interest is held constant.
- Here are some questions that you can ask about a study to assess its construct validity.
    - Do the manipulations applied in the study actually change what they are supposed to?
    - Is anything influenced by manipulations in the study other than the construct of interest?

- Do the measures in the study accurately reflect what they are designed to measure?
- Can different responses on measures be caused by anything else other than changes in the construct of interest?

- Below are descriptions of some specific threats to construct validity.
  - *Mono-method bias.* Sometimes a study (or an entire field of research) consistently manipulates a variable in a single fashion, or uses a single way to measure a variable. Any single method is necessarily going to contain idiosyncratic elements that are not truly part of the theoretical construct of interest. When you use a single method to examine a variable you cannot separate the theoretical construct of interest from irrelevant characteristics of the method. It is therefore best to demonstrate relationships using multiple ways of manipulating or measuring the constructs of interest.
  - *Hypothesis guessing.* People have a natural tendency to try to figure out what is going on around them. When placed in an study they will try to make some guesses as to what it is about. Their actions will sometimes be influenced by these beliefs, causing them to behave in a way that is unnatural. Sometimes this simply contributes random noise to your data, but other times (such as when people are able to accurately guess your hypotheses) this can systematically bias your results. It is therefore important to keep your hypotheses hidden from your participants, and to try to reduce hypothesis-guessing tendencies in general. One way of accomplishing the latter is to provide reasons (though not necessarily accurate ones) for each task that participants perform. People are less likely to try to come up with their own explanations if you give one to them first. It can also be useful to ask participants afterwards what they thought the study was about. If people seem to be able to regularly guess the purpose of your study then you know you may need to worry about hypothesis-guessing biases.
  - *Administrator expectations.* Studies have shown that the beliefs of research administrators can have profound influences on the data they collect, even when they are making supposedly objective measurements (Rosenthal, 1964). The responses then reflect the behavior of the researchers instead of the behavior of the participants. The easiest way to prevent this problem is to minimize the role of the administrator in the study. Try to use standardized methods of manipulating and measuring your variables. If your study requires that the administrator be more deeply involved then you should try to use administrators who are unaware of the research hypotheses.
  - *Evaluation apprehension.* People are generally concerned whenever they are being evaluated by someone they perceive to be an authority. They will therefore do their best to present themselves in a way that makes them appear competent and psychologically healthy. You must consider whether your manipulations are inducing a desire in your participants to be perceived favorably in addition to (or instead of) causing changes in the construct of interest. You should also examine your response items to make sure that they are not simply reflecting participants' attempts to create favorable impressions.
  - *Placebo/Hawthorne effects.* The mere presence of an experimental environment can sometimes cause changes in participants. It is well known, for example, that the mere expectation that they are receiving medical care can sometimes cause improvement in people's health, even if they actually receive no medical treatment. This can obscure the effect of any manipulations you wish to perform. When designing your study you should include "control groups" of people who are placed in similar circumstances. This way the effect of the experimental environment is removed from any statistical comparisons you make.
  - *Order effects.* Sometimes one part of your study can influence the way that people perform in another part. For example, if you have people perform a difficult arithmetic task it might make them feel dejected, causing them to put less effort into their responses on later tasks. You should therefore try to put important tasks, as well as tasks that might be sensitive to order effects, early in your study.

4.3 Internal validity

- Internal validity concerns the plausibility of alternative causal explanations for the results of a study. If a study is not attempting to demonstrate any causal relationships then internal validity is not a concern.

- The first thing that you need to consider when a study proposes a causal relationship is whether the direction of causation is correct. This will not be an issue if the predictor is manipulated but may be difficult to demonstrate when both the predictor and the response variable are measured.

- Even when you manipulate your predictor variable you must be concerned with potential *confounding variables*. A confounding variable is an outside variable (one not examined in your experiment) that covaries with your predictors and could potentially influence your response variable. If a predictor covaries with a confounding variable then you cannot conclude that the predictor causes the response. The confound always provides a possible alternative explanation for your results.

- While people assessing internal validity are typically interested in finding confounding variables that might have produced the observed relationships between variables, it is also possible for a confounding variable to obscure a relationship between two variables. For example, people who know that they are in a more difficult level of a factor may try to work harder than those in an easier level. This confound would reduce the difference between these groups.

- Most threats to internal validity can be avoided if you randomly assign participants to the different between-subject conditions in your study. On average every outside variable should be evenly distributed throughout the different conditions of your study, removing most potential threats to validity. To properly randomize the assignment of participants to conditions you must choose a method that makes each participant equally likely to be in each level of your between-subject variable.

- It is important that you use the same recruitment procedures for participants in every condition of an experiment. If the different levels of a between-subject manipulation have different requirements then participants recruited for the experiment should all meet the requirements of every condition. Similarly, if different conditions take different amounts of time then all participants should be recruited expecting that they may be in the longest condition. People willing to be in a longer experiment may be different from people who are not.

- It is better if you assign the level of your between-subject conditions on a participant-by-participant basis than to entire groups of participants. For example, it is a bad idea to for you to assign the first level of a variable to the first 10 participants and the second level to the next 10 participants. This confounds your manipulation with the order of recruitment, and possibly with the time of day the study was conducted. Sometimes you must run people in groups, such as when the manipulation is a film clip seen by participants. If you have multiple participants in each session of your experiment it may not be practical to try to show them each a movie individually. In this case you should randomly determine which groups are exposed to each level of your between-subjects variable. Having multiple sessions for each level (so the conditions are at least somewhat dispersed) can also reduce the potential for a confound.

- There are some threats that neither matching nor randomization will be able to prevent. Specifically, if a confound is in some way created by your manipulation then randomization will not help. For example, a manipulation that withholds rewards from some participants may inspire resentment, a potential alternative explanation for your results. Additionally, if the influence of an outside factor would be different for different levels of a predictor (i.e., the confound interacts with the predictor variable) its influence will also not be removed by randomization. For example, consider a study trying to determine whether taking school trips to a museum influences school performance. One potential confound is the SES of the student, since wealthier students might be more likely to go to museums on their own. This confound also may interact with the condition, because wealthy students who don't go on school trips may go on their own while wealthy students who do go during school may not go on

their own. The wealthy students are acting differently depending on what condition they are, so even if your two groups are equally wealthy there will still be differences in the behaviors of the groups.

If you cannot remove the potential influence of a confounding variable through randomization then you will need to include extra measurements or control groups in your design to test for the influence of the potential confound on your response variables.

- Here are some questions that you can ask about a study to assess its internal validity.
  - Is the direction of causality clearly established?
  - Could the observed relationships have been created by unaccounted outside influences?
  - Could some outside influence be obscuring the true relationships between the variables?
  - Is the assignment of participants to between-subjects conditions properly randomized?
- Below are descriptions of some specific threats to internal validity.
  - *History.* If your study takes place over an extended period of time there is a possibility that some unusual event might take place during the course of your study that might affect your results. In this case you cannot be sure whether you would have obtained the same results in the absence of the event.
  - *Testing.* Participants who are asked to fill out a questionnaire or perform a test more than a single time may use their past performance to guide their behavior on later tests. They may remember specific responses and just recall them instead of re-generating their response. If you are specifically investigating memory then this is not a problem, but most often you want the the tests to reflect changes caused by some other variable. In this case their experience acts as a confound. To get around this you can design parallel versions of the tests and use a different version at each test session.
  - *Maturation.* If your study takes place over an extended period you need to take into account that people grow and change over time. It is possible that any differences that you see over the course of treatment are due to natural development instead of your manipulation. This is especially important to consider when using younger participants.
  - *Regression toward the mean.* If you administer a test to a sample at one point in time you will naturally observe that some people will score high and some will score low. Any form of measurement, however, naturally involves some amount of random error. So if you were to administer the same test again you would observe some differences in the scores, even if there were no actual change. You would also see that those who scored high on the first test will tend to score lower on the second, while those who scored low on the first test will tend to score higher on the second. This is because if a person received a high score, they are more likely to have had random factors working to their benefit. In a sense, they are toward the high end of the possible values they might obtain on the test. If they took the test again they would be more likely to have a lower score. Similarly, people with low scores are more likely to have had random factors working against them, so later tests are likely to provide higher results.

    This can be a problem if you put people into different groups based on a pretest score and then you want to observe the effect of a manipulation with a posttest. Even without the influence of any other factors, regression toward the mean will cause the high group to have lower scores and the low group to have higher scores at posttest. To minimize the influence of this threat you should choose measures with as little random error as possible.
  - *Selection.* This is a threat when you have different types of people in different levels of a between-subjects manipulation. In this case you can't be certain whether the effects you observe are caused by the manipulation or the individual differences. Random assignment to conditions can usually remove threats of selection.
  - *Mortality.* If you conduct a study that occurs for a long period of time it is common to have a certain percentage of your participants drop out before the study is completed. Sometimes this this dropout creates differences between the people in your different conditions. For example, you

might have one treatment that is more difficult than another. In this case you might see the lazy people drop out of the difficult treatment but not out of the easy treatment. If this happens then the type of people in your groups will be confounded with your experimental manipulation.

- *Diffusion of treatments.* If participants in different conditions of your study receive different types of training you need to consider whether the information received by members of one group might be transmitted to members of another group. If this happens you will see an artificial reduction of the difference between your conditions. This is particularly a problem when your training takes place over a long period of time and you expect that your participants will have opportunities to interact with each other.
- *External equalization.* This threat appears when your study involves providing some sort of benefit to some groups while withholding them from other groups. In this case you will sometimes see other organizations trying to provide compensatory benefits to those in your underprivileged condition. For example, you might want to compare the performance of students in a study program to those not in the program. The school may feel guilty about denying the potential benefits to some of its students and so may start providing extra study services of its own.
- *Compensatory rivalry.* If participants find out that they are in a disadvantaged or more difficult condition of your study they will sometimes put forward additional effort to compensate for your procedures. You usually want to investigate the direct influence of the procedures rather than the effects of participants' reactance so their extra motivation acts as a confounding variable.
- *Resentful demoralization.* This is sort of the reverse of compensatory rivalry. Sometimes when participants know that they are in a disadvantaged condition of your study they lose motivation to perform well at your tasks. This is especially likely if they perceive the task as exceptionally difficult. Just like compensatory rivalry, resentful demoralization will act as a confounding variable in your design.

4.4 Statistical conclusion validity

- Statistical conclusion validity concerns whether the reported relationships in a study are accurate.
- A study's statistical conclusion validity is determined by characteristics of both the variables and the statistics. Your variables should be consistently manipulated and accurately measured. Poor measurement or manipulation of your variables can prevent you from observing relationships. Your statistical analyses should be appropriate for the data and performed correctly. Improper statistical analyses can both prevent you from observing true relationships and can also cause you to conclude that there are relationships between unrelated variables.
- Here are some questions that you can ask about a study to assess its statistical conclusion validity.
  - Were the measurements and manipulations reliable?
  - Were the manipulations strong enough to produce observable effects?
  - Do the data meet the assumptions of the statistical analyses?
  - Do the statistical tests have sufficient power?
- Below are descriptions of some specific threats to statistical conclusion validity.
  - *Random influences in the study.* The most common inferential statistics all involve some comparison of the amount of variability in the response measure that can be explained by the predictor variables to the random variability in the measure. Anything that would increase the random variability in the measure will reduce your ability to detect significant relationships. You can reduce the influence of random factors by making your procedures consistent across your participants.

    Sometimes there are random factors that you cannot control but which you can measure, such as characteristics of specific participants. If you record this information then you can remove its influence by including a blocking variable or a covariate in your analyses.

- *Unreliable measurements.* Any measurement will include some amount of error which will add random variability to your results. The extent to which an instrument will produce the same value on repeated measurements is referred to as its *reliability.* It is difficult to find significant relationships using variables measured with unreliable instruments. If a study fails to find an effect using unreliable measurements you cannot be sure whether the effect truly doesn't exist or whether the study simply lacked the ability to detect it.

  A study that successfully detects a significant relationship using an unreliable instrument, however, should not be called into question. A lack of reliability always works against attempts to find significance. If anything, observing a significant relationship using an unreliable measure means that the relationship must be particularly strong.
- *Low power.* The ability of your statistical analysis to detect a relationship is known as the *power* of that analysis. Studies that have very small sample sizes generally have low power. Additionally, not all statistical procedures are equally good at detecting relationships between your variables. Certain statistical procedures, such as the chi-square test for independence and many nonparametric statistics, are known for being very poor at detecting relationships. You should be skeptical about any conclusions based on the failure of a low-power test to find a significant relationship.

  Low power does *not* ever call significant findings into question. Just like low reliability, low power is something that must be overcome to successfully detect a relationship. It cannot cause you to find a relationship where one doesn't exist.
- *Weak manipulations.* If a study fails to find a significant effect of a manipulation you must consider the extent to which the procedures actually change the theoretical variable of interest. It might be possible that a stronger manipulation of the same variable would have had an observable effect on the response measure.
- *Ceiling and floor effects.* If you use a measurement that has minimum or maximum boundaries it can be difficult for you to detect differences when the average score is near those boundaries. Response differences become compressed near the ends of the scale, reducing variation. There isn't much "room to move" so factors that would normally create more extreme responses will not have their normal influence. One way to deal with this is to modify to response measure so that it has broader endpoints.
- *Violated test assumptions.* All statistical tests have certain assumptions that they make about the data that must be satisfied for the tests to produce accurate results. Some violations will cause your tests to report significant relationships when there really is no relationship, while other violations will prevent a test from detecting the presence of an existing relationship. If the your data violates the assumptions of a test then you cannot draw any conclusions from the results of that test.
- *Fishing.* Standard inferential statistics are usually interpreted in terms of a "p-value" that indicates the probability of observing a relationship as strong as you did given that no relationship actually existed. These p-values, however, are calculated on the assumption that only a single test is being performed. The more test statistics that you calculate to analyze your study the more likely that you will observe a p-value below a specific value. To correct for this you should lower the critical value that you use when you perform a large number of statistical tests.

- One note should be made in particular about studies that accept the null hypothesis. It is very difficult to provide convincing evidence that there is not a relationship between two variables. You really can't tell whether you failed to find a relationship because your statistical test didn't have enough power or because there isn't a relationship there to be found. Usually the best that you can do is demonstrate that one relationship is significantly less strong than some other comparative relationship.

4.5 External validity

- External validity concerns the ability to apply inferences drawn from a study to other populations, environments, and times.

- The more reliant a study's results are on the specific features present in the original study the less external validity is possessed by the study.
- To assess external validity you must first determine the real-life circumstance to which the researcher wants to generalize their conclusions. You then consider what characteristics of the study were important for obtaining the results related to that conclusion. If these characteristics are not likely to be present in the real-life circumstance then you have reason to question the external validity of the results.
- External validity is not equally important for all types of studies. It is very important for studies whose aim is to describe the relationships between variables in a particular target population. However, it is much less important for studies whose goal is to test some aspect of a theory. For example, a theory might predict that a particular pair of variables should be related. This relationship should be present in any particular sample your might draw. If you conducted a study and found no relationship this would be evidence against the theory even if the study had low external validity.

  Even studies testing theories can benefit from having external validity. Externally valid tests of a theory provide more general (and thus stronger) evidence for or against the theory. For example, contradictory findings with low external validity may be dismissed by supporters of the theory as unimportant exceptions. Contradictory findings with high external validity, on the other hand, are much more difficult to ignore.
- A study can provide conclusions that may be validly generalized to one population but which do not validly apply to another. For example, a study providing information about shopping preferences may provide valid conclusions for the community in which the study was conducted but might be invalid for understanding the preferences of a different community in another part of the country.
- Not all differences between samples and target populations create external validity problems. In order for a difference to be a threat to external validity there must be a reason to suspect that the difference would influence the results of the study.
- When conducting your own study the key to establishing external validity is to make the important characteristics of your research environment match those of the real-world situation to which you wish to generalize your findings.
- Here are some questions that you can ask about a study to assess its external validity.
  - Is the purpose of a study to test a theory or to provide a description of a target population?
  - What differences are there between the people, time, and setting of the study and the target situation? Could these differences conceivably have influenced the results?
- Below are descriptions of some specific threats to external validity.
  - *Unrepresentitive sample.* If the type of people that you select for your study are affected by your treatment differently than would the average person that makes up your target population then the results of your study cannot be properly generalized to the population of interest. This is why it is best to draw a sample that is truly representative of the target population.
  - *Unrepresentative setting.* Similarly, if the environment in which you conduct your study is very different from that to which you would like to generalize your results then your conclusions may not be valid. It is usually more important to consider what is present in your study environment then what is absent. You do not necessarily need to replicate all of the different things that you feel would be present in real life. However, you should be fairly certain that the results you obtained were not caused by some peculiarity of your study environment.

Chapter 5

Preparing Study Materials

5.1 Study timeline

- The first thing you should do when trying to implement a study is to create a timeline detailing everything that happens during your study. This will be relatively short for a survey but can become quite complex for an experiment.
- When writing your timeline take the point of view the person administering the study. Your timeline should be detailed enough so that another person could actually use it to conduct the study, assuming they had access to your materials.
- Describe everything that the research administer would have to do from the time participants enter to the time they leave the study. If the procedure differs for participants in different conditions you should describe what would happen in each.
- It is usually easiest to follow a timeline if it is presented as a hierarchical outline. That way you can quickly find the details about a specific part of the study.
- A timeline is useful at three different points of the research process.
    1. When you are building your study materials a timeline forces you to think through every part of the study. It can be used to organize the generation of study materials, helping to make sure that nothing is forgotten. Constructing a timeline will also make you more likely to notice threats to the validity of your study.
    2. When you are actually conducting the study the timeline can be used as a guide so that the correct things are done at the correct times. If things become hectic while running the study the administrator can always turn to the timeline to make sure that all the appropriate things are done.
    3. When you are writing the methods section of an article the timeline is an excellent resource to have. Even if you are able to memorize the correct procedure when conducting the study you may forget details by the time you actually try to write about the research.

5.2 Verbal instructions and interactions

- You must next write the instructions for your study that you will give to participants. Every part of your study that requires participants to do something new should have its own set of instructions.
- Whether you plan to give the instructions orally or in a written format you should have them explicitly written down. You want to be sure that each participant gets exactly the same instructions. When providing oral instructions the best way to do this is to have your research administrators read the instructions from a script. You should try to have scripts for every standard interaction between the administrator and the participants.

- If the task is particularly difficult or there are important details that participants need to know to perform it correctly you are better off providing the instructions orally. Participants will often skim written instructions or skip them entirely.
- Make your instructions as clear as possible. Use simple language to explain what participants should do. Provide examples if this will help participants understand the task.
- Always provide participants with an opportunity to ask the administrator questions if they don't understand something about what they are supposed to do. Participants that misunderstand the task will provide meaningless data, making it more difficult for you to obtain significant results. It is a good idea to end your instructions by asking if participants understand and by saying that they should contact the administrator if they ever have any questions.

  If you find out after the study is finished that a participant misunderstood the task you should make a note so that you can later exclude their data from analysis if necessary.

5.3 Materials for your manipulations

- If you are conducting an experiment you need to create the materials associated with your manipulations.
- You need to create a *randomization scheme* for any between-subjects manipulations that must be applied by the experimenter. This scheme is a table where each row represents a participant and each column represents one of your between-subject variables. The particular condition each participant should be in can be determined simply by reading the cells in the table. Sometimes people do the randomization while running the experiment but it is usually easier if you do this beforehand and record it on a randomization scheme.

  There are several different ways to obtain random numbers to help you set up your scheme. You can use Microsoft Excel to generate columns of random numbers if you first install the Data Analysis Toolpak. This is part of the standard distribution of Excel (so you don't need to buy it separately) but is not installed by default. To install it select Tools → Add-Ins and then check the box next to "Analysis Toolpak." Afterwards you can generate random numbers using the Tools → Data Analysis → Random Number Generation option.

  Probably the best web resource for generating random numbers is the Research Randomizer at

  `http://www.randomizer.org`

  This online tool can generate lists of random numbers using many different options that you can specify.
- You will also need an *experiment record*. This is a sheet where you can list the details about each participant in your experiment. It is usually designed as a table where each row holds a participant and each column holds a feature that you want to record. When you run a participant through the study you record the information on each feature in the appropriate cell of the sheet.

  You will want to have columns representing the participant number, the date and time the participant was run, and the values of any between-subject manipulations that were assigned by the experimenter. The experiment record may also code other details such as the sex of participants or the number of the cubicle they used. You usually do not use the experiment record for information on actual response variables. Its purpose is to note procedural details in case you find evidence that there were unusual influences on the results.

  Your experiment record should provide some space (possibly on the back of the sheet) for the administrator to note anything unusual that happens during a session. This will allow you to take these events into account when analyzing the data.
- Finally you must create any stimuli that are to be used as part of the experiment. The specifics of what you will need will vary greatly from study to study. You should construct them carefully and make sure that you have the stimuli you need for every possible condition.

5.4 Materials for your measurements

- You need to prepare all of the measuring instruments (whether physical or written) used in your study.
- If you have physical instruments you must ensure that they are functioning properly. You should also see that those who will be operating the instrument have been properly trained in how to set up and use the device.
- If you have observers code the behavior of participants they should have detailed instructions provided in a *coding scheme*. Observers should have extensive practice using a coding scheme before they record the data of real participants.
- If you are collecting data with a questionnaire you must determine exactly what questions you will ask, what types of responses participants will provide, and in what order they will see the questions. Items that are answered in more or less the same way will generally be presented together to reduce the amount of instructions that you need to give. You should put the most sensitive and most important items early in the questionnaire so that their responses are less likely to be influenced by the way participants respond to other items.
- When using an established scale (or coding scheme) you should use it without modifications, even if you think that your changes could improve the quality of the responses. Established scales have usually been designed to have particular psychometric properties. If you introduce changes to a scale you can no longer be certain that the new version will retain all of the old properties. You should always report any changes you make to established scales when reporting the methodology of your research.
- If you need to write any additional questions you should try to make them simple, clear, and unbiased. You want to make your items as simple and straightforward as possible so that your respondents are able to fill out the scale quickly and easily. Complicated questions can lead to misunderstandings and annoy respondents, reducing the likelihood that you obtain good data using the scale. Unclear questions can cause respondents to interpret your items in different ways, reducing the likelihood that they will answer in ways that are related to the theoretical construct of interest. An item should be designed such that every respondent interprets it in the same way. If the language of your questions makes some responses seem more expected or desired than others you will be unlikely to get a true measurement of respondents' characteristics. Instead they will alter their answers in a way that makes them look better according to the language of the question. This will alter the actual meaning of the item, moving it away from the theoretical construct of interest.

5.5 Post-study survey

- You should design a post-study survey to see how participants perceive your study.
- The point of a post study survey is to provide you with insight into each participant's personal experience while going through your study. If participants provide abnormal data the responses in the survey can sometimes provide an explanation for their behavior.
- There are three questions you should always ask in your post-study survey.
    - "Were the instructions complete and understandable? Was there ever a time when you were uncertain about what you were supposed to do?"
    - "Did anything unusual happen while you were going through the study?"
    - "What did you think this research was about?"
- If you can think of any additional questions that could provide you with insight on factors that might influence participants' performances your should include them in the survey.

5.6 Ethics materials

- You must generate the materials required to satisfy APA ethical guidelines and local human subjects committees.
- Many universities and local communities have created Human Subjects Committees that you must get to approve your study before you are allowed to conduct it within their jurisdictions. APA guidelines require that researchers receive approvals from any such organization before conducting research. These committees typically require that you send in an application indicating the purpose of your research and detailed information about what participants' experiences will be in the study. Your application will then be reviewed and you will either receive approval for your study or the committee will request specific changes to your procedures.
- APA guidelines require that you obtain the written consent of everyone who participates in your study unless it meets their listed criteria for disposing with informed consent. For this you will need to create an *Informed Consent Form* to have signed by each participant. This form should provide a description of what participants will be doing in the study, how long it will take, the extent to which they will be personally associated with their responses (i.e., anonymity procedures), and a list of who will have access to the data. It should describe the potential benefits of participation to the participant and to society. Participants must be warned if there are any expected discomforts associated with the study. The description of the study should concentrate on what a participant's experience will be while the study is conducted. Somewhere in the document you should report the name and contact information of the primary investigator.

  APA guidelines require that the consent form states that participants may quit the study without penalty at any time and that they have the right to know the outcome of the research. There should be a statement in the document saying that by signing the participants give their informed consent to be part of the study. At the bottom you should have a space where the participants will sign and date the form. When running the study you should keep a signed form for your records and give a blank one for the participant to keep.
- You should also design a debriefing form to give participants once they finish the study. This form tells participants exactly what you were researching and the purpose of each part of the study. The debriefing should reveal the truth behind any deceptions in the study and should report any hidden measures that you took. Like the informed consent, the debriefing form should provide the primary investigator's name and contact information.

  Providing a debriefing lets participants understand their importance to the scientific process and provides them with knowledge about your research. This helps keep them happy and will foster goodwill between researchers and the public. If your firm or university gets a bad reputation regarding the treatment of participants you may find it very difficult to perform research.

Chapter 6

Conducting Research

6.1 General guidelines

- Psychological research should adhere to the published APA ethics code. All APA journals (and most non-APA journals as well) require authors to state upon submission that the research was conducted according to APA guidelines. The full APA ethics code can be found online at

  `http://www.apa.org/ethics/code.html`

  Below are some of the main points relating to conducting research.

  - Research must not violate any applicable laws.
  - Research must be approved by any applicable institutional and community human subjects committees.
  - Researchers must obtain the informed consent of participants prior to collecting data from them except in the case of anonymous surveys and naturalistic observations.
  - All inducements for participation must be provided and all commitments made to participants during a study must be kept.
  - Researchers must not offer excessive inducements when they might tend to coerce participation.
  - Deception should not be used in research unless it is justified by the potential knowledge to be gained by the study and alternative methods are unavailable.
  - The researcher must provide participants with information about the purpose of the study once they conclude their involvement. At this time the researcher should correct any misconceptions held by participants resulting from participation in the study.

- Research administrators should always act professionally when conducting a study. This improves the credibility and authority of the administrator and will cause participants to take the study more seriously. A professionally presented study can make participants more likely to volunteer, encourage them to listen to or read the instructions more carefully, and cause them to put more effort to the tasks they perform.

  Here are some things that research administrators should do to appear more professional.

  - Dress as if they were working in a professional job.
  - Be prepared and organized. They should try to have all the materials set up and accessible before interacting with participants.
  - Be familiar with study procedures. Administrators who are unsure about what they are doing will reduce a participant's confidence in the study.
  - Treat participants professionally. Administrators should not express anger or annoyance with participants nor should they be flirtatious or overly friendly. The best attitude would probably

be "unobtrusively pleasant." A study administrator should only be noticed by participants when they specifically need to interact.

If you should ever troublesome participants the best thing for you to do is to fully compensate them for the time they were in the study and then dismiss them. Any sort of confrontation will reflect badly on the study and could affect the performances of other participants that are present.

- Be knowledgeable about the research. If participants ask question about study procedures the administrator should be able to respond quickly and confidently. Administrators should also have enough knowledge to answer any theoretical questions that are likely to be asked by participants during debriefing.
- Be attentive. It is generally best that administrators not engage in outside activities (such as reading a book or listening to the radio) while they are in the presence of participants. Similarly two administrators should not have personal discussions in front of participants while conducting a study.

- You should make your study as pleasant as possible for your participants. Discomfort and fatigue are often sources of confounding variables in a study. They can also reduce the reliability of your manipulations and measurements.

  Below are some things you can do to make your study more pleasant.

  - Make your instructions clear. Participants can get frustrated if they have to spend much time trying to figure out what they are supposed to do.
  - Make your materials easy to see and read. Avoid unusual fonts when presenting text and make images large enough so they can be seen without straining.
  - Make your study interesting. You can greatly increase participants' motivation in your study if you can somehow make the task more fun or intriguing.
  - Try to maintain a comfortable working environment. Conduct your study in a location that is comfortably warm and relatively free from outside distractions such as noise.
  - Keep your study sessions short. As the length of a study grows the likelihood of participants becoming bored or fatigued increases. If you have the option between two equivalent measures you are usually better of choosing the shorter one. If possible try to keep your study sessions under one hour.
  - Provide breaks during long or repetitive tasks. Even short breaks can help reduce fatigue when participants are performing the same actions over and over. If you have a session that lasts for much longer than an hour you should also provide at least one opportunity for participants to use the bathroom.
  - Design your study so that it is easy for participants to perform. The interface used by participants should be as simple as possible. If you are conducting a survey try to use a standard format for your questions. If you are conducting an experiment try to minimize the amount of physical work participants need to do in their tasks.
  - Maintain a pleasant attitude throughout the session. This is especially important if you have a significant amount of interaction with the participants.
  - Take all questions from participants seriously. If a participant comes to you with a question never send them away without being completely sure that you solved their problem. Participants are unlikely to ask a question a second time and will become frustrated if they must perform a task they don't understand.

- There are two basic formats you can use when conducting research. If your entire study could be placed on paper (including all manipulations and measurements) then you are in a position to use a *quesionnaire*. If your study requires a heavily controlled environment, complicated procedures or apparatus then you must instead perform a *laboratory study*. The main difference between the two is the portability of the materials. Questionnaires can be easily moved and administered in different locations while laboratory studies must usually be conducted in a single location. For more details on

administering a questionnaire see section 6.2. For more details on conducting a lab study see section 6.3.

- It may require multiple sessions for you to run all of the participants you want. At the end of each session you should store any completed forms in a safe place. The more you move your forms around the greater chance that you might lose or damage them.

6.2 Administering a questionnaire

- When administering a questionnaire the first thing you need to do is to decide exactly how and where you will give the questionnaire. You have a large number of options because a questionnaire can be completed almost anywhere.
    - You can advertise and have participants come to a specific location, just like you would for a laboratory study. The advantages of this method are that you can run a large number of participants at a time, the common setting reduces random sources of variance, and you can use long questionnaires since you know that anyone recruited for the study will have the appropriate amount of time available. The disadvantages are that you will almost always need to compensate participants to get them to show up, and you are susceptible to selection artifacts since people who are willing to travel to be in your study may be different from those who are not.
    - You can go to a public location with clipboards and ask people passing by to fill out your questionnaire. The main advantage of this method are that you can often recruit participants free of charge. A disadvantage is that this method can only be successfully used with relatively short questionnaires.
    - You can travel to people's homes and conduct interviews with participants. The advantages of this method are a relatively high response rate and an increase in the quality of the data. There will generally be a reduced chance of coding errors since the administrator will be recording the responses. The administrator also has the opportunity to provide additional instructions or examples for items that confuse participants. Some disadvantages of this method are that you can only run one participant at a time and the presence of the administrator may influence the responses provided by participants. You may also have difficulty obtaining a truly random sample since certain types of people spend more time at home than others.
    - You can conduct phone interviews with participants. The advantages and disadvantages of phone interviews are similar to those of live interviews. Some differences are that the use of the telephone requires less effort on the part of the administrator and provides participants with a greater sense of anonymity. However it makes it very difficult if not impossible to provide any pictorial or graphic information as part of the questionnaire.
    - You can mail your questionnaire to potential respondents and provide a self-addressed stamped envelope for them to return their response. The main advantage of this is that it offers the greatest anonymity of all the methods. The disadvantages are that you will generally have a low response rate and it will take a long time before you collect all of your data.
    - You can send your questionnaire to people over email. The advantages of this are the exceptionally low cost (no photocopying) and the relative ease with which you can send your questionnaire to large numbers of people. The disadvantages are that you can only contact a limited community (people who use email) and that you will generally have a low response rate.
    - You can post your questionnaire on the World Wide Web. The advantage of this method is that you can collect huge amounts of data in a relatively short period of time with a minimum of cost. The main disadvantage is that you have no control over and limited knowledge about the type of participants you receive and the circumstances they are in when they fill out the questionnaire.
- Once you have decided on the method you will use you must recruit your participants. You want your sample of participants to be representative of your target population on all important dimensions. If you have a specific target population (such as members of a particular profession) you may need

to obtain lists of population members from which you can draw your samples. For more general populations you may decide to perform your recruitment at the time you are actually administering the questionnaire.

- You must then make copies of your questionnaire and bring them to the site at which you will conduct your study. You should also bring any supplemental materials, such as pencils, erasers, or paper, that will be needed to complete the study. You shouldn't count on participants having anything they need with them, even if you tell them beforehand.

- At this point you will actually administer the questionnaire, using the guidelines presented in section 6.1. Follow the timeline you created when generating your study materials. Be sure to use informed consent and debriefing forms if they are required by your study.

6.3 Conducting a laboratory study

- If you decide to conduct a lab study the first thing you will need to do is to recruit participants. This will generally need to be done in advance since you must arrange for participants to be at the laboratory at the time you are running the study. When recruiting you will need to tell people the general nature of what they will be doing, how much time the study will take, the requirements for participation, the compensation they would receive, and where and when the study will be held. It is a good idea to provide potential participants with reminder slips so they don't forget where and when they should be at the laboratory.

- When it comes time to conduct the study you should always go to the laboratory early. This will allow you to make the appropriate preparations and be ready before participants arrive.

- You should greet participants as they arrive in the laboratory. Have some empty chairs available so that participants can sit down in case you are not ready to start the study when they arrive. If you must have them wait you should give them an approximation of how long it will be before the study will begin.

- At this point you will actually conduct the study, using the guidelines presented in section 6.1. Follow the timeline you created when generating your study materials. Be sure to use informed consent and debriefing forms if they are required by your study.

6.4 Pretesting your study

- You will usually want to perform a trial run of your study before you attempt to conduct it using real participants. This is referred to as performing a *pretest*. The goal of a pretest is not to obtain accurate or reliable data, rather it is to determine the effectiveness of your manipulations and measurements and to expose any potentially difficult aspects of the study's procedure.

- In a pretest you are specifically not concerned using your results to test your hypotheses. In fact, you will often add in diagnostic procedures that you know will affect your results and would normally provide alternative explanations for your findings. Additionally, you do not need to be particularly concerned about the representativeness of your sample since you will not be drawing theoretical inferences from your results.

- To determine if your manipulations are accomplishing what you intend them to you can perform *manipulation checks*. A manipulation check is simply a measurement of the variable that you are trying to manipulate. For example, if you have people watch either a happy or a sad movie to try to manipulate mood, you can perform a manipulation check simply by asking people what their mood is. You often don't include manipulation checks in the regular study procedure because the check might make people aware of the manipulation and cause them to try to correct for it.

- If you measure a variable using a scale with multiple questions you can calculate the *reliability* of the scale to determine the precision of its measurements. Reliability tells you the extent to which people tend to respond the same way to all of the items in the scale. Scales with low reliabilities aren't measuring a single, consistent construct and therefore will provide poor results when used in statistical analyses.
- It is fairly common to have a short five-minute interview with pretest participants after they finish your study. You should ask them the type of questions that you would ordinarily put into a post-study survey. Actually conducting a live interview is more work but is better (at least during the pretest) because it will allow you to ask follow-up questions about anything participants might say.

Chapter 7

Data analysis

7.1 Creating a data set

- You will almost certainly want to perform your statistical analyses using computer software. This requires you to create a file on the computer representing the data set. You will usually want each row of this data set to represent a participant and each column of the data set to represent a variable. This means that any particular cell of the data set holds the value of the variable in that column for the participant in that row.

- If you collected your data using surveys or other written measures you will need to type the data into the computer. The three most common options are to use either a word processor (like WordPerfect or Microsoft Word), a spreadsheet (like Lotus 123 or Microsoft Excel), or a database (like Paradox or Microsoft Access).

  The best option is probably to use a spreadsheet program. Spreadsheets are generally much easier to work with for data entry than either of the other two options. Additionally, most statistical programs have procedures that allow you to directly translate spreadsheet files to data files. If you use a word processor you will probably have to save the data as a text file and then go through a possibly lengthy data import procedure. If you use a database you will likely have save the data as a spreadsheet file anyway for it to be read by your statistical software.

- Perhaps the best strategy to remove data entry errors is to have the questionnaires entered into the computer twice, in two different files. You can then locate mistakes by having the computer check for differences between the two files. Data that has been entered incorrectly in one file should show up as mismatching the other file (since it is highly unlikely that the two files would contain the same typographical error on the same item). Most modern word processors have a "Document Compare" function which allows you to locate differences between two files. This is referred to as *double entry*. It is easy to get almost complete accuracy using this method.

  You can either have two different people type in the data or you can have the same person type it in twice - it actually makes little difference. If you use a spreadsheet or database for data entry you should probably convert the file to a text document and then use the document compare from a word processing program. You can then use the results to make corrections to the original spreadsheets or databases. The reason you should do this is because the compare functions available in word processors are usually much easier to use than the compare functions of spreadsheets or databases.

- Sometimes the initial form of your data set will not be what you need for a particular statistical analysis. In this case you can use a statistical program to transform your data set into the new form you desire. Specific details about how to transform your data set can be found in our *Notes on Transforming and Restructuring Data*, freely available on the Stat-Help.com website.

7.2 Statistical analysis

- A detailed discussion of how to analyze your data is beyond the scope of these notes. This section will instead provide information on what procedures you should use to examine the types of relationships that are most commonly proposed in hypotheses. It will be left to the reader to learn the specifics of how to perform each procedure.
- The following procedures examine the relationship between a single pair of variables. The one you should use depends on the types of variables you want to examine.
    - Both *correlations* and *simple linear regression* can be used to examine the relationship between two numeric variables.
    - *T-tests* or *ANOVA* (analysis of variance) is used to examine the relationship between a numeric variable and a categorical variable. You use a t-test if your categorical variable has only two levels, otherwise you must use ANOVA.
    - A *chi-square test* is used to examine the relationship between two categorical variables.
- The following procedures examine the relationship between a single numeric response variable and two or more predictor variables. The one you should use again depends on the types of predictor variables you want to use.
    - *Multiple regression* relates a numeric response variable to two or more numeric predictor variables.
    - *ANOVA* (analysis of variance) can relate a numeric response variable to two or more categorical variables.
    - If you want to relate a numeric response variable to a combination of numeric and categorical variables then you will need to use some form of the *general linear model*. Both multiple regression and ANOVA are actually specific applications of the general linear model.
- If you want examine the relationship between a single categorical response measure and two or more predictor variables (either numeric or categorical) you will need to use *logistic regression*.
- If you want to examine the relationship between two or more response variables and a set of predictor variables you will need to use *multivariate analysis*.

7.3 Documenting your study

- You should always document a study as soon as you finish analyzing the data. It should be documented whether or not you found evidence for your hypotheses and whether or not you intend to publish your research.
- There are several different reasons why you should document your study.
    - You will need a description of the methods and results if you wish to publish the research.
    - You may want to have a record of what you did as a reference when conducting future studies.
    - You or another researcher may want to go back at a later point in time and re-analyze your data.

  Notice that the last two justifications apply even if you didn't obtain publishable results. Sometimes researchers will just store the data from an unsuccessful study without any documentation. However it is quite possible that you might want to run a variant of the study at a later point and could use a writeup so that you can decide on appropriate changes. It is also possible that a study with unexpected results may become understandable later on when new theories become available.
- To properly document a study you need to include information about the exact procedures that were used and the results that were obtained. Your documentation should include the following items.
    - A complete writeup of the method section as if for publication (see section 8.3).

- A complete writeup of the results section as if for publication (see section 8.4).
- A copy of the study timeline.
- A copy of the data (usually on disk) generated during the study.
- An index to the data. For each data file you should provide a general description of what is in the file, the location and name of the file, and explanations of each variable contained in the data set.
- Copies of any programs you used for statistical analysis.

Chapter 8

Writing a Research Paper

8.1 General comments about writing

- The official guide to writing psychology papers (both for classes and for journals) is the *Publication Manual of the American Psychological Association* (APA, 1994). This book contains information about writing style, punctuation, formatting a document, and references. The website

  `http://www.uwsp.edu/psych/apa4b.htm`

  contains a nice summary of APA style, including examples of formatting and references (Plonsky, 2000).

- The way that you write about your work can be just as important as the actual content. It is well worth the effort to make your papers as clear and as easy to read as possible. Well-written papers have many advantages over poorly written papers.

  - Their concepts are easier to understand.
  - The importance of the research can be more readily seen.
  - People are more likely to finish reading the article.
  - The research is more likely to be remembered.
  - Readers are more likely to select articles authored by people known to be good writers.
  - Well-written arguments are more persuasive.
  - People perceive good writers as more knowledgeable and more competent.

- Here are some things you should do when writing.

  - *Write in complete sentences using appropriate grammar.* Other people will have a very difficult time understanding what you want to say if you write in incomplete sentences. Improper grammar tents to attract attention to itself so that you reader is distracted from the content of your writing.
  - *Be sure the vocabulary is used correctly.* If you are not completely sure of the definition of a word you should either look it up or avoid using the word in your writing. In addition to being a nice guide to writing papers, Sternberg (1993) provides definitions of words that are commonly misused by psychological authors.
  - *Write clearly.* Your sentences should each read smoothly and plainly express its point. Always read through your document after you are done writing. Try to take the perspective of a minimally knowledgeable reader and consider whether each sentence is understandable. If you have any doubts you should rewrite the sentence to make it clearer.

    It can be difficult for authors to notice unclear writing in their own papers. Often times the writer unconsciously uses their own background knowledge to make unclear sentences understandable. For this reason it can be helpful to have another person read your paper and point out unclear

sections. If you can't find someone else to go over your paper you can also try reading the paper out loud. This makes it more difficult to unconsciously clarify your writing.

- *Write simply.* Avoid the temptation to use intellectual jargon or overly technical language. Readers will not be impressed by your knowledge of vocabulary. Rather they will usually become annoyed at the increase in effort required to read your paper. When you have a choice between a simple and a complicated way to say something choose the simpler option.
- *Write for your reader.* The nature of your intended audience should influence your writing in several ways. First, you should choose a vocabulary that would be understood by the least-informed member of your intended audience. If you are writing a paper to be read by medical doctors you can use more complicated medical terminology than if you were writing about the same topic for patients. Second, you should provide technical details to the extent that they are required by your audience. You can describe a well-known memory task to a group of cognitive psychologists using much less detail than would be required to communicate the same thing to a more general audience. Finally, you should focus your discussion on the aspects of your research that would be of greatest interest and concern to your intended audience. An article written for a popular magazine should focus on general findings and implications while one written for a scientific journal should provide more detail about methodology and statistical analysis.
- *Avoid excess language.* This incorporates two basic points. First, if a word or phrase can be removed without changing the meaning of a sentence then those words should be dropped. Two major sources of extra words are introductory clauses ("**Indeed**, the major finding was that the experimental group exceeded the control") and unnecessary qualifiers ("Participants performed a **rather** uninteresting task"). The two example sentences can just as easily be written as "The major finding was that the experimental group exceeded the control" and "Participants performed an uninteresting task."

  Second, you should try to avoid redundant sentences when possible. Redundant writing is not always bad - it can emphasize important points or provide additional opportunities to understand a difficult concept. However, redundancy also encourages knowledgeable readers to skim your document, possibly causing them to miss important points. You are usually better off to make your first presentation of an idea clear and interesting than to reiterate your point later in the paper.
- *Prefer active to passive voice.* Sentences written in active voice contain action verbs ("Attention **decreased** during the experiment") while those written in passive voice contain linking verbs ("A decrease in attention **was observed** during the experiment"). A passive-voice sentence can usually be transformed to active voice with some slight rewording. Documents written using active voice are easier to read and more interesting than those using passive voice.
- *Include transitions and summary statements.* Often times you will have distinct parts within a single section of your document. You should include sentences and paragraphs to cleanly conclude each part and to provide a link to the following topic. Summaries not only help you emphasize your main points but also notify your reader that a topic change will soon occur. Transitions help your reader move from one part to another without losing track of the underlying theme of your paper. Documents without appropriate transitions and summaries appear choppy and disjointed.
- *Avoid sexist language.* Do not exclusively use male pronouns when you want to refer to a generic individual who could be of either sex. You can use "one" or "he or she," but unfortunately both of these terms are awkward and can bog down your writing. Probably the best option is to rephrase your sentence so that it refers to a plural so that you can use the pronoun "they." For example, instead of saying "Each participant wrote down his first name" you could say "Participants each wrote down their first name."

8.2 Writing the introduction section

- The main purpose of the introduction is to tell your reader what it is that you studied and to convince them that it is important.

- The first paragraph typically provides an introduction to your general topic, often with real-world examples to demonstrate the prevalence of the phenomenon of interest. It should convey the importance of the research area and motivate people to read the rest of the paper. Try to avoid using citations in the first paragraph unless your study is concerned with issues tied to a specific study or methodology.
- Following the first paragraph you should provide a review of the literature relevant to your research. The purpose of this review is to provide a context in which your study and its importance can be understood. You therefore do not need to give a comprehensive summary. Instead focus on articles that are directly relevant to your study.
- After reviewing the prior research you then must link it to your current research. You need to make a convincing argument that your study is necessary, valuable, and interesting. Specifically you should point out the gap in the literature that your study is designed to fill and convince readers that much can be gained by filling it.
- Finally you present the general idea behind the study you conducted. You should state the research question your study investigates and what information the study is designed to provide. Report the hypotheses that your study was designed to test. You may also provide an overview of the study, although specific details should be left to the method section.
- Be sure to put effort into relating the literature review to your own research. Explicitly state any links that you wish readers to make. You cannot assume that people will appreciate the importance of something unless you specifically point it out to them.
- Keep in mind that you are not simply reporting past research in the introduction. You are trying to *sell* the idea that your own study is important. From beginning to end the introduction should be a convincing argument that the research you propose is interesting and necessary.

8.3 Writing the method section

- The purpose of the method section is to tell your reader exactly how you conducted the study. You should include enough detail so that someone reading your method would be able to conduct a replication of the study.
- The method section is usually divided into several subsections, each describing a different aspect of your study. The exact order of these subsections is not fixed. You should choose an order that makes your study easiest to describe. Your method additionally does not need to have all of the subsections. Sometimes a particular subsection may not apply to your study, and sometimes you may decide to combine two subsections to make presentation easier. Include whatever subsections you feel are necessary to fully describe your study.
- Below are the most common subsections in a method section.
    - *Participants.* This subsection should list the total number of participants in the study and provide a general description of their characteristics. You should also report how they were recruited and what compensation they received for participation. If you decided to exclude any participants from your statistical analyses you would state that here and provide an explanation why they were dropped.
    - *Apparatus.* Here you should describe any physical equipment that you used when conducting your study. You should provide the make and model of any important items. You should also describe how the apparatus was used in your study. If your equipment was put together in an unfamiliar way you should provide a photograph or a schematic of the setup you used.

      The use of computers in your study is not usually described in this section. If your study was conducted using a computer program you will typically provide a short note in the procedure subsection describing the type of computer you used (i.e., PC or Macintosh) and the programming language in which the study was written. The use of computers and software for statistical analysis is not usually reported unless the analyses are particularly unusual.

- *Design*. The design subsection should describe all of the variables that were measured or manipulated in the study. It should present the conditions you used in your study and how participants were assigned to the conditions. The specifics of how to write up your design are provided in section 3.6.
- *Procedure*. Here you should describe what your participants did from the time they entered to the time they left your study. You will find that your study timeline will be very useful when writing this section. You should also include details on any coding schemes you used.

  When describing your procedures you should stay at a conceptual level. For example, instead of listing out all of the questions you asked in a survey you would simply state what theoretical variables were measured by the questions. If you have several questions composing a scale measuring a single concept then you may also choose to report the reliability of the scale. When describing a manipulation you would want to provide its purpose and general format of any stimuli rather than a list of the specific stimuli used in the study. If you wish to include detailed information about any new scales or stimuli you develop you should put them in an appendix rather than in the method section.

8.4 Writing the results section

- The results section is where you report your data analyses. It should provide both summary statistics describing the data as well as any inferential statistics you calculated to test your hypotheses.
- You usually start this section with any statistics you calculated to verify the validity of your study. This would include manipulation checks and tests of order effects. If you have any scales for which you did not provide reliability measurements in the method section you would also provide that here. In general you want to remove any concerns that readers may have about your methodology before you present your major findings. If you do obtain a finding that would question your study's validity you should honestly report the implications it has for the inferences you wish to make.
- You then want to present the main results of your study. If you are testing causal hypotheses you will likely want to organize your presentation around the different response variables. For each response variable you should start with a table containing the mean and standard deviation of the response in each cell of the experimental design. If your design is very complicated you might instead provide a graph, though tables are generally preferred. You should then follow this with the results of the inferential tests you performed (such as t-tests, ANOVA, and regression) relating the response variable to the predictor variables.

  If you are testing correlational hypotheses then the organization is more up to your own discretion. You can present the means and standard deviations of your variables if you like. The main part of your results section will be dedicated to reporting the results of inferential tests you performed relating the variables to each other.
- While descriptive statistics (such as means and standard deviations) are often displayed in tables or graphs, inferential statistics are presented as sentences within the body of your paper. When reporting the results of a commonly used inferential test you need to provide the name of the test, a statement as to whether the test was significant or not, the value of the test statistic, the degrees of freedom associated with the test, and finally the p-value of the test. The last three of these are usually separated from the rest of the text by parentheses or commas. For example, you might say "A within-subjects t-test indicated that there was a significant difference between the experimental and the control group ($\underline{t}[15]=4.3$, $\underline{p} < .01$)" or "An ANOVA revealed that the interaction between prime type and delay was not statistically significant, $\underline{F}[2, 39]=1.4$, $\underline{p} > .05$." If you use an unusual statistic you should provide additional information about interpreting the statistic as well as a reference where more advanced details can be found.
- When reporting p-values people conventionally report that the $\underline{p}$ was less than a particular value if the test was significant while they report that it was greater than a particular value if the test was not

significant. In psychology we typically consider p-values less than .05 to be significant. Some people call p-values between .05 and .10 "marginally significant."

- APA format requires that test statistics named using normal letters (such as t, F, and p) be presented underlined in the standard font. Test statistics named using Greek letters (such as χ^2 and λ), on the other hand, should be presented in the standard font without underlining.
- After you report your inferential statistics you can then provide the results of any secondary analyses you performed to further explore a significant finding. For example, after reporting a significant main effect in ANOVA you might follow this with statistical comparisons of the individual cell means.
- If you report the results of an inferential test that has implications for one of your hypotheses you should follow it with a single sentence describing the implication. For example, "This result was consistent with our expectations." Save any detailed comments for the discussion section.

8.5 Writing the discussion section

- In the discussion section you provide an interpretation of the results of your study. You should start with a description of how well the data support the hypotheses you proposed in the introduction. You should talk about the implications for each hypothesis individually. These implications will fall into one of four categories.

    1. The findings provide support for the hypothesis. In this case the relationships proposed by the hypothesis were observed in your data. Additionally, your results provide no evidence of any alternative explanations for proposed causal relationships. If your hypothesis proposes multiple relationships it is possible to find partial support for the hypothesis if some of the proposed relationships are observed while others are not.
    2. The findings provide support against the hypothesis. In this case you observed relationships in your data that are inconsistent with the hypothesis.
    3. The findings fail to support the hypothesis but do not provide evidence against it. This happens when you fail to observe relationships proposed by the hypothesis but do not observe any contradictory relationships. Inferential statistics do not allow us to draw strong conclusions when we obtain nonsignificant results, so observing lack of evidence must be treated differently from observing specific evidence against a hypothesis.

        To draw any conclusions from a null effect you will need to provide convincing evidence that your study was powerful enough so that you would have observed a significant relationship if one had been present. One way to do this is to show that your study was able to detect other relationships that should be about the same strength as the those predicted by the hypothesis. You may also directly calculate the *power* of the study, which is defined as the probability that you would obtain nonsignificant results given that the proposed relationship actually does exist.
    4. The findings cannot be interpreted. It is possible that some aspect of your study prevents you from drawing any conclusions regarding the hypothesis. For example, you might have detected a confounding variable that provides an alternative explanation for a proposed causal relationship. Another possibility is that you find evidence that your manipulations or measurements lack construct validity. This would mean that the relationships you observe in your data do not truly reflect the relationships between the theoretical variables you wish to examine.

- You should next state what conclusions you wish to draw from your study. Specifically, you should state whether the study as a whole provides support for or against any theories relevant to the research. If the results failed to support your own hypothesis you might wish to propose alternative explanations for the results.
- You may also discuss any practical implications your study has for either research or everyday life. You should consider whether the results of your study might help explain other findings in the literature. If there are any direct implications of your results for real-life behavior you should also make them clear.

- Finally you should describe specific directions that you think future research on the topic should take. Section 8.6 discusses the different ways you may choose to follow up a study.

8.6 Planning future studies

- The final stage of the research process is a return to the very beginning: Generating new research hypotheses. However, at this point you are in a different circumstance than when you originally started. You have conducted, analyzed, and drawn inferences about a new study which should hopefully present new opportunities for research.
- If your study supported your research hypotheses then you want to think about ways to further verify or extend your findings.
    - You could perform a conceptual replication of your study. Since hypotheses are usually written in abstract terms it should be possible for you to test the same hypothesis using different manipulations or measurements. For example, you might have originally conducted a study where you manipulated how much participants were paid to show that motivation influences performance on a task. To follow this up you might conduct a second study where you manipulate the perceived importance of the task (another way of varying motivation) and see if this has similar effects on performance.
    - You might try to locate mediators or moderators of the relationships you discovered. Once you establish the presence of a relationship between two concepts you can go deeper and try to add other variables to the equation. For example, you might have originally conducted a study showing that skill level expectations influence task performance. To follow this up you could examine whether individual differences in optimism and pessimism might moderate this relationship.
    - You can attempt to show that the observed relationships can be observed in different populations and different environments. A single study is necessarily limited in terms of where and on whom it was conducted. Demonstrating that the findings can be replicated in either more general or more specialized circumstances can be valuable.
- If your study failed to support your hypotheses then you should plan studies to examine and correct the reasons for the failure.
    - You could run the first study again but with increased power. You might consider strengthening the manipulations (such as by using more extreme stimuli) or improving the sensitivity of your measurements (such as by increasing the number of items in a scale).
    - You also might try to test the same hypotheses but using different manipulations or measurements. The difficulties you had in detecting relationships could be tied to the specific procedures in the study so changing them might give you the results your desire.
    - You could also attempt to locate a moderating variable that may be suppressing the relationship you wish to find. For example, you might have conducted a study that failed to find a hypothesized relationship between current events and overall mood. This might be because there are many people who do not pay attention to current events, but that for those that do the nature of the events would affect their mood. You might therefore perform a follow-up study using the same basic method but with additional measures to determine how often participants read newspapers or watch the news on television.
- Finally you should also consider any relationships that you observed involving one or more exploratory variables (those you measured but were not specifically related to your hypotheses). You might try to develop hypothetical explanations for the relationships that can be tested with a study.

References

American Psychological Association (1994)(4th edition). *Publication manual of the American Psychological Association.* Washington, DC: American Psychological Association.

American Psychological Association (2001)(9th edition). *Thesaurus of Psychological Index Terms.* Washington, DC: American Psychological Association.

Baron, R. M., & Kenny, D. A. (1986) The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology, 51*, 1173-1182.

Cook, T. D., & Campbell, D. T. (1979). *Quasi-Experimentation Design & Analysis Issues for Field Settings.* Boston, MA: Houghton Mifflin Company.

Plonsky, M. (2000). Psychology with style: A hypertext writing guide (Version 3.1). Retrieved from the web March 20, 2001. `http://www.uwsp.edu/psych/apa4b.htm`

Rosenthal, R. (1964). Experimenter outcome-orientation and the results of the psychological experiment. *Psychological Bulletin, 61*, 405-412.

Sternberg, R. J. (1993). *The Psychologist's Companion: A Guide to Scientific Writing for Students and Researchers.* Cambridge, UK: Cambridge University Press.