Research Methods Knowledge Base

by William M. Trochim, Cornell University

What is the Research Methods Knowledge Base?

The Research Methods Knowledge Base is a comprehensive web-based textbook that addresses all of the topics in a typical introductory undergraduate or graduate course in social research methods. It covers the entire research process including: formulating research questions; sampling (probability and nonprobability); measurement (surveys, scaling, qualitative, unobtrusive); research design (experimental and quasi-experimental); data analysis; and, writing the research paper. It also addresses the major theoretical and philosophical underpinnings of research including: the idea of validity in research; reliability of measures; and ethics. The Knowledge Base was designed to be different from the many typical commercially-available research methods texts. It uses an informal, conversational style to engage both the newcomer and the more experienced student of research. It is a fully hyperlinked text that can be integrated easily into an existing course structure or used as a sourcebook for the experienced researcher who simply wants to browse.[Back to Top]

Purchasing

You can purchase a complete printed copy of the Research Methods Knowledge Base over the web by selecting the link *Purchase the complete printed text of the Knowledge Base online* at the bottom of any page. [Back to Top]

Using the KB in a Course

As of August 1, 2000, all course and classroom support for the Knowledge Base is being handled by my new publisher, Atomic Dog Publishing. You can reach them at their website at www.atomicdogpublishing.com. In addition to providing a unique online version of the Knowledge Base text, they are the exclusive publishers of the print version. Through Atomic Dog Publishing you can expect the finest in web-based course support for the Knowledge Base including workbooks, study guides, online testing, test item data banks, and much more. If you have any questions about use of the Knowledge base in your course, please check their website or contact them directly by e-mail at

help@atomicdogpub.com.

About the Author

William M.K. Trochim is a Professor in the Department of Policy Analysis and Management at Cornell University. He has taught both the undergraduate and graduate required courses in applied social research methods since joining the faculty at Cornell in 1980. He received his Ph.D. in 1980 from the program in Methodology and Evaluation Research of the Department of Psychology at Northwestern University. His research interests include the theory and practice of research, conceptualization methods (including concept mapping and pattern matching), strategic and operational planning methods, performance management and measurement, and change management. He is the developer of The Concept System® and founder of Concept Systems Incorporated. he lives in Ithaca, New York with his wife Mary and daughter Nora. [Back to Top]

Acknowledgements

This work, as is true for all significant efforts in life, is a collaborative achievement. I want to thank especially the students and friends who assisted and supported me in various ways over the years. I especially want to thank Dominic Cirillo who has labored tirelessly over several years on both the web and printed versions of the Knowledge Base and without whom I simply would not have survived.   There are also the many graduate Teaching Assistants who helped make the transition to a web-based course and have contributed their efforts and insights to this work and the teaching of research methods. And, of course, I want to thank all of the students, both undergraduate and graduate, who participated in my courses over the years and used the Knowledge Base in its various incarnations. You have been both my challenge and inspiration.[Back to Top] New location

Dedication

For Mary and Nora

who continue to astonish me with their resilience, patience, and love

[Back to Top]

Contents

Research Methods Knowledge Base

Contents
Navigating

Yin-Yang Map
The Road Map

Foundations

Language Of Research
Five Big Words
Types of Questions
Time in Research
Types of Relationships
Variables
Hypotheses
Types of Data
Unit of Analysis
Two Research Fallacies
Philosophy of Research
Structure of Research
Deduction & Induction
Positivism & Post-Positivism
Introduction to Validity
Ethics in Research
Conceptualizing
Problem Formulation
Concept Mapping
Evaluation Research
Introduction to Evaluation
The Planning-Evaluation Cycle
An Evaluation Culture

Sampling

External Validity
Sampling Terminology
Statistical Terms in Sampling
Probability Sampling
Nonprobability Sampling

Measurement

Contents

Navigating

Yin-Yang Map
The Road Map

Navigating the Knowledge Base

There are at least five options that I can think of for getting to relevant online material in the Knowledge Base:

The Border Contents

Every page of the Knowledge Base has links in the margins. These links are based on the hierarchical structure of the website and change depending on the position of the page in that structure. The links at the top (repeated at the bottom) on each page show the other pages at the same level of the hierarchy as the page you are looking at. The links in the left border always include:

- The Home Page
- The parent page for the page you are viewing
- The child pages for the page you are viewing

The Table of Contents

This is a standard hierarchical table of contents like the type you would expect in a textbook. It is the only navigational device that at a glance shows every page in the Knowledge Base.

The Yin-Yang Map

This map is based on a graphic that, at a glance, provides an organizing rubric for the entire Knowledge Base content. It separates the theory of research from the practice of research and shows how theory and practice are related. this might be an especially useful launch pad for an advanced or graduate research methods course because of the strong emphasis on the link between theory and practice.

The Road Map

This map is based on a graphic that shows the typical stages in a research project.   It uses the metaphor of research as a journey down the research road from initial conceptualization and problem formation through the write-up and reporting. This might be an especially useful launch pad for an introductory undergraduate course because it concentrates primarily on the practice of research.

The Search Page

In the top and bottom margins on every page in the Knowledge base there is a link to the Search Page. When you need to find information on a specific topic rapidly you should use this page. The Search Page is linked to an index of every word in the Knowledge Base, allows you to perform simple and Boolean searches, and returns resulting links sorted from most to least relevant.

Yin-Yang Map

Navigating the Knowledge Base

The Yin-Yang Map
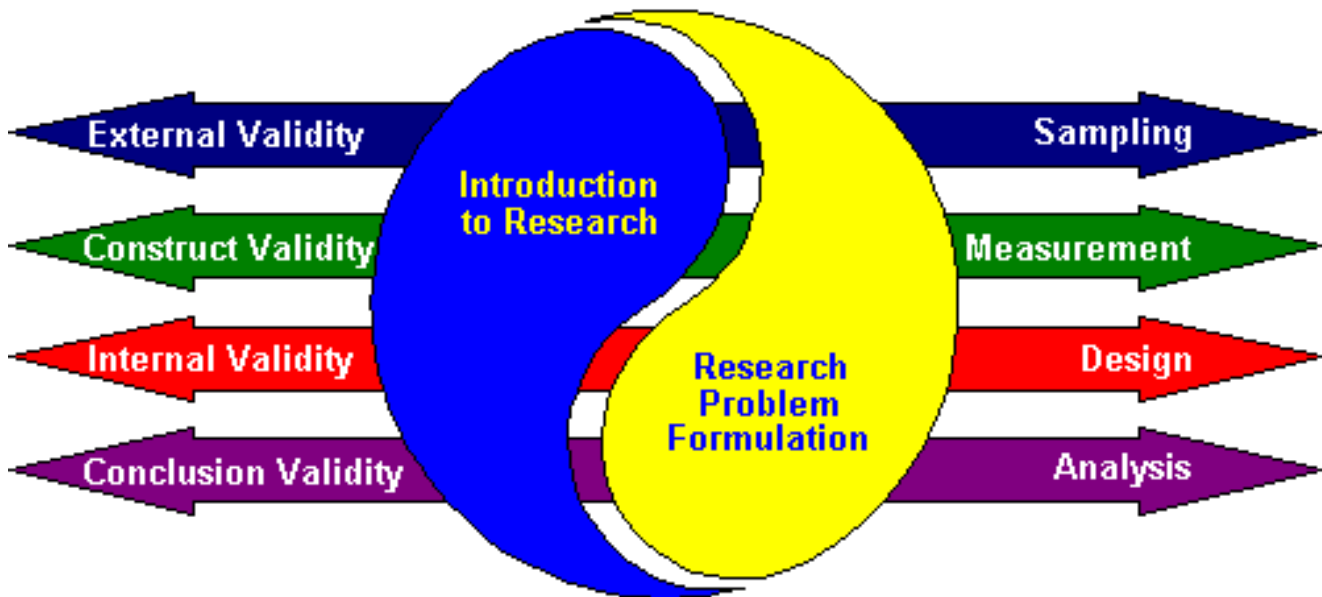
The Yin and the Yang of Research

You can use the figure above to find your way through the material in the Knowledge Base. Click on any part of the figure to move to that topic.

The figure shows one way of structuring the material in the Knowledge Base. The left side of the figure refers to the *theory* of research. The right side of the figure refers to the *practice* of research.

The yin-yang figure in the center links you to a *theoretical* introduction to research on the left and to the *practical* issue of how we formulate research projects on the right.

The four arrow links on the left describe the four *types of validity* in research. The idea of validity provides us with a unifying theory for understanding the criteria for good research. The four arrow links on the right point to the *research practice areas* that correspond with each validity type. For instance, external validity is related to the theory of how we generalize research results. It's corresponding practice area is sampling methodology which is concerned with how to draw representative samples so that generalizations are possible.

Navigating the Knowledge Base

The Road Map

The Road to Research

Remember all those Bob Hope and Bing Crosby films? The Road to Singapore?   Of course you don't -- you're much too young! Well, I thought it might be useful to visualize the research endeavor sequentially, like taking a trip, like moving down a road -- the Road to Research. The figure above shows a very applied way to view the content of a research methods course that helps you consider the research process *practically*. You might visualize a research project as a journey where you must stop at certain points along your way. Every research project needs to start with a clear problem formulation. As you develop your project, you will find critical junctions where you will make choices about how you will proceed. Consider issues of sampling, measurement, design, and analysis - as well as the theories of validity behind each step. In the end, you will need to think about the whole picture, or "What can we conclude?" Then you might write-up your findings or report your evaluation. You even might find yourself backtracking and evaluating your previous decisions! Don't forget that this is a **two-way** road; planning and evaluation are critical and interdependent. The asphalt of the road is the foundation of research philosophy and practice. Without consideration of the basics in research, you'll find yourself bogged down in the mud!

Measurement

Measurement is the process observing and recording the observations that are collected as part of a research effort. There are two major issues that will be considered here.

First, you have to understand the **fundamental ideas** involved in measuring. Here we consider two of major measurement concepts. In Levels of Measurement, I explain the meaning of the four major levels of measurement: nominal, ordinal, interval and ratio. Then we move on to the reliability of measurement, including consideration of true score theory and a variety of reliability estimators.

Second, you have to understand the different **types of measures** that you might use in social research. We consider four broad categories of measurements. Survey research includes the design and implementation of interviews and questionnaires. Scaling involves consideration of the major methods of developing and implementing a scale. Qualitative research provides an overview of the broad range of non-numerical measurement approaches. And unobtrusive measures presents a variety of measurement methods that don't intrude on or interfere with the context of the research.

Conceptualizing

Problem Formulation
Concept Mapping

One of the most difficult aspects of research -- and one of the least discussed -- is how to develop the idea for the research project in the first place. In training students, most faculty just assume that if you read enough of the research in an area of interest, you will somehow magically be able to produce sensible ideas for further research. Now, that may be true. And heaven knows that's the way we've been doing this higher education thing for some time now. But it troubles me that we haven't been able to do a better job of helping our students learn ***how*** to formulate good research problems. One thing we can do (and some texts at least cover this at a surface level) is to give students a better idea of how professional researchers typically generate research ideas. Some of this is introduced in the discussion of problem formulation in applied social research.

But maybe we can do even better than that. Why can't we turn some of our expertise in developing methods into methods that students and researchers can use to help them formulate ideas for research. I've been working on that area pretty intensively for over a decade now -- I came up with a structured approach that groups can use to map out their ideas on any topic. This approach, called concept mapping can be used by research teams to help them clarify and map out the key research issues in an area, to help them operationalize the programs or interventions or the outcome measures for their study. The concept mapping method isn't the only method around that might help researchers formulate good research problems and projects. Virtually any method that's used to help individuals and groups to think more effectively would probably be useful in research formulation. Some of the methods that might be included in our toolkit for research formulation might be: brainstorming, brainwriting, nominal group technique, focus groups, Delphi methods, and facet theory. And then, of course, there are all of the methods for identifying relevant literature and previous research work. If you know of any techniques or methods that you think might be useful when formulating the research problem, please feel free to add a notation -- if there's a relevant Website, please point to it in the notation.

Analysis

Conclusion Validity
Data Preparation
Descriptive Statistics
Inferential Statistics

By the time you get to the analysis of your data, most of the really difficult work has been done. It's much more difficult to: define the research problem; develop and implement a sampling plan; conceptualize, operationalize and test your measures; and develop a design structure. If you have done this work well, the analysis of the data is usually a fairly straightforward affair.

In most social research the data analysis involves three major steps, done in roughly this order:

- Cleaning and organizing the data for analysis (Data Preparation)
- Describing the data (Descriptive Statistics)
- Testing Hypotheses and Models (Inferential Statistics)

Data Preparation involves checking or logging the data in; checking the data for accuracy; entering the data into the computer; transforming the data; and developing and documenting a database structure that integrates the various measures.

Descriptive Statistics are used to describe the basic features of the data in a study. They provide simple summaries about the sample and the measures. Together with simple graphics analysis, they form the basis of virtually every quantitative analysis of data. With descriptive statistics you are simply describing what is, what the data shows.

Inferential Statistics investigate questions, models and hypotheses. In many cases, the conclusions from inferential statistics extend beyond the immediate data alone. For instance, we use inferential statistics to try to infer from the sample data what the population thinks. Or, we use inferential statistics to make judgments of the probability that an observed difference between groups is a dependable one or one that might have happened by chance in this study. Thus, we use inferential statistics to make inferences from our data to more general conditions; we use descriptive statistics simply to describe what's going on in our data.

In most research studies, the analysis section follows these three phases of analysis. Descriptions of how the data were prepared tend to be brief and to focus on only the more unique aspects to your study, such as specific data transformations that are performed. The descriptive statistics that you actually look at can be voluminous. In most write-ups, these are carefully selected and organized into summary tables and graphs that only show the most relevant or important information. Usually, the researcher links each of the inferential analyses to specific research questions or hypotheses that were raised in the introduction, or notes any models that were tested that emerged as part of the analysis. In most analysis write-ups it's especially critical to not "miss the forest for the trees." If you present too much detail, the reader may not be able to follow the central line of the results. Often extensive analysis details

are appropriately relegated to appendices, reserving only the most critical analysis summaries for the body of the report itself.

Write-Up

Key Elements
Formatting
Sample Paper

So now that you've completed the research project, what do you do? I know you won't want to hear this, but your work is still far from done. In fact, this final stage -- writing up your research -- may be one of the most difficult. Developing a good, effective and concise report is an art form in itself. And, in many research projects you will need to write multiple reports that present the results at different levels of detail for different audiences.

There are several general considerations to keep in mind when generating a report:

The Audience

Who is going to read the report? Reports will differ considerably depending on whether the audience will want or require technical detail, whether they are looking for a summary of results, or whether they are about to examine your research in a Ph.D. exam.

The Story

I believe that every research project has at least one major "story" in it. Sometimes the story centers around a specific research finding. Sometimes it is based on a methodological problem or challenge. When you write your report, you should attempt to tell the "story" to your reader. Even in very formal journal articles where you will be required to be concise and detailed at the same time, a good "storyline" can help make an otherwise very dull report interesting to the reader.

The hardest part of telling the story in your research is finding the story in the first place. Usually when you come to writing up your research you have been steeped in the details for weeks or months (and sometimes even for years). You've been worrying about sampling response, struggling with operationalizing your measures, dealing with the details of design, and wrestling with the data analysis. You're a bit like the ostrich that has its head in the sand. To find the story in your research, you have to pull your head out of the sand and look at the big picture. You have to try to view your research from your audience's perspective. You may have to let go of some of the details that you obsessed so much about and leave them out of the write up or bury them in technical appendices or tables.

Formatting Considerations

Are you writing a research report that you will submit for publication in a journal?   If so, you should be aware that every journal requires articles that you follow specific formatting guidelines. Thinking of writing a book. Again, every publisher will require specific formatting. Writing a term paper? Most faculty will require that you follow specific guidelines. Doing your thesis or dissertation?   Every university I know of has very strict policies about formatting and style.   There are legendary stories that circulate among graduate students about the dissertation that was rejected because the page margins were a quarter inch off or the figures weren't labeled correctly.

To illustrate what a set of research report specifications might include, I present in this section general guidelines for the formatting of a research write-up for a class term paper. These guidelines are very similar to the types of specifications you might be required to follow for a journal article. However, you need to check the specific formatting guidelines for the report you are writing -- the ones presented here are likely to differ in some ways from any other guidelines that may be required in other contexts.

I've also included a sample research paper write-up that illustrates these guidelines.   This sample paper is for a "make-believe" research project. But it illustrates how a final research report might look using the guidelines given here.

Design

Research design provides the glue that holds the research project together. A design is used to structure the research, to show how all of the major parts of the research project -- the samples or groups, measures, treatments or programs, and methods of assignment -- work together to try to address the central research questions. Here, after a brief introduction to research design, I'll show you how we classify the major types of designs. You'll see that a major distinction is between the experimental designs that use random assignment to groups or programs and the quasi-experimental designs that don't use random assignment. [People often confuse what is meant by random selection with the idea of random assignment. You should make sure that you understand the distinction between random selection and random assignment.] Understanding the relationships among designs is important in making design choices and thinking about the strengths and weaknesses of different designs. Then, I'll talk about the heart of the art form of designing designs for research and give you some ideas about how you can think about the design task. Finally, I'll consider some of the more recent advances in quasi-experimental thinking -- an area of special importance in applied social research and program evaluation.

Sampling

External Validity
Sampling Terminology
Statistical Sampling Terms
Probability Sampling
Nonprobability Sampling

Sampling is the process of selecting units (e.g., people, organizations) from a population of interest so that by studying the sample we may fairly generalize our results back to the population from which they were chosen. Let's begin by covering some of the key terms in sampling like "population" and "sampling frame." Then, because some types of sampling rely upon quantitative models, we'll talk about some of the statistical terms used in sampling. Finally, we'll discuss the major distinction between probability and Nonprobability sampling methods and work through the major types in each.

Foundations

This section provides an overview the major issues in research and in evaluation. This is probably the best place for you to begin learning about research.

We have to begin somewhere. (Although, if you think about it, the whole idea of hyperlinked text sort of runs contrary to the notion that there is a single place to begin -- you can begin anywhere, go anywhere, and leave anytime. Unfortunately, you can only be in one place at a time and, even less fortunately for you, you happen to be right here right now, so we may as well consider this a place to begin.) And what better place to begin than an introduction? Here's where we take care of all the stuff you think you already know, and probably should already know, but most likely don't know as well as you think you do.

The first thing we have to get straight is the language of research. If we don't, we're going to have a hard time discussing research.

With the basic terminology under our belts, we can look a little more deeply at some of the underlying philosophical issues that drive the research endeavor.

We also need to recognize that social research always occurs in a social context. It is a human endeavor. Therefore, it's important to consider the critical ethical issues that affect the researcher, research participants, and the research effort generally.

Where do research problems come from? How do we develop a research question? We consider these issues under conceptualization.

Finally, we look at a specific, and very applied, type of social research known as evaluation research.

That ought to be enough to get you started. At least it ought to be enough to get you thoroughly confused. But don't worry, there's stuff that's far more confusing than this yet to come.

Language Of Research

Learning about research is a lot like learning about anything else. To start, you need to learn the jargon people use, the big controversies they fight over, and the different factions that define the major players. We'll start by considering five really big multi-syllable words that researchers sometimes use to describe what they do. We'll only do a few for now, to give you an idea of just how esoteric the discussion can get (but not enough to cause you to give up in total despair). We can then take on some of the major *issues* in research like the types of questions we can ask in a project, the role of time in research, and the different types of relationships we can estimate. Then we have to consider defining some basic terms like variable, hypothesis, data, and unit of analysis. If you're like me, you *hate* learning vocabulary, so we'll quickly move along to consideration of two of the major fallacies of research, just to give you an idea of how wrong even researchers can be if they're not careful (of course, there's always a certainly probability that they'll be wrong even if they're extremely careful).

Five Big Words

Research involves an eclectic blending of an enormous range of skills and activities. To be a good social researcher, you have to be able to work well with a wide variety of people, understand the specific methods used to conduct research, understand the subject that you are studying, be able to convince someone to give you the funds to study it, stay on track and on schedule, speak and write persuasively, and on and on.

Here, I want to introduce you to five terms that I think help to describe some of the key aspects of contemporary social research. (This list is not exhaustive. It's really just the first five terms that came into my mind when I was thinking about this and thinking about how I might be able to impress someone with really big/complex words to describe fairly straightforward concepts).

I present the first two terms -- *theoretical* and *empirical* -- together because they are often contrasted with each other. Social research is theoretical, meaning that much of it is concerned with developing, exploring or testing the theories or ideas that social researchers have about how the world operates. But it is also empirical, meaning that it is based on observations and measurements of reality -- on what we perceive of the world around us. You can even think of most research as a blending of these two terms -- a comparison of our theories about how the world operates with our observations of its operation.

The next term -- *nomothetic* -- comes (I think) from the writings of the psychologist Gordon Allport. Nomothetic refers to laws or rules that pertain to the general case (nomos in Greek) and is contrasted with the term "idiographic" which refers to laws or rules that relate to individuals (idiots in Greek???). In any even, the point here is that most social research is concerned with the nomothetic -- the general case -- rather than the individual. We often study individuals, but usually we are interested in generalizing to more than just the individual.

In our post-positivist view of science, we no longer regard certainty as attainable. Thus, the fourth big word that describes much contemporary social research is *probabilistic*, or based on probabilities. The inferences that we make in social research have probabilities associated with them -- they are seldom meant to be considered covering laws that pertain to all cases. Part of the reason we have seen statistics become so dominant in social research is that it allows us to estimate probabilities for the situations we study.

The last term I want to introduce is *causal*. You've got to be very careful with this term. Note that it is spelled *causal* not *casual.* You'll really be embarrassed if you write about the "casual hypothesis" in your study! The term causal means that most social research is interested (at some point) in looking at cause-effect relationships. This doesn't mean that most studies actually study cause-effect relationships. There are some studies that simply observe -- for instance, surveys that seek to describe the percent of people holding a particular opinion. And, there are many studies that explore relationships -- for example, studies that attempt to see whether there is a relationship between gender and salary. Probably the vast majority of applied social research consists of these descriptive and correlational studies. So why am I talking about causal studies? Because for most social sciences, it is important that we go beyond just looking at the world or looking at relationships. We would like to be able to change the world, to improve it and eliminate

some of its major problems. If we want to change the world (especially if we want to do this in an organized, scientific way), we are automatically interested in causal relationships -- ones that tell us how our causes (e.g., programs, treatments) affect the outcomes of interest.

Types of Questions

There are three basic types of questions that research projects can address:

1. **Descriptive.** When a study is designed primarily to describe what is going on or what exists. Public opinion polls that seek only to describe the proportion of people who hold various opinions are primarily descriptive in nature. For instance, if we want to know what percent of the population would vote for a Democratic or a Republican in the next presidential election, we are simply interested in describing something.
2. **Relational.** When a study is designed to look at the relationships between two or more variables. A public opinion poll that compares what proportion of males and females say they would vote for a Democratic or a Republican candidate in the next presidential election is essentially studying the relationship between gender and voting preference.
3. **Causal.** When a study is designed to determine whether one or more variables (e.g., a program or treatment variable) causes or affects one or more outcome variables. If we did a public opinion poll to try to determine whether a recent political advertising campaign changed voter preferences, we would essentially be studying whether the campaign (cause) changed the proportion of voters who would vote Democratic or Republican (effect).

The three question types can be viewed as cumulative. That is, a relational study assumes that you can first describe (by measuring or observing) each of the variables you are trying to relate. And, a causal study assumes that you can describe both the cause and effect variables and that you can show that they are related to each other. Causal studies are probably the most demanding of the three.

Time in Research

Time is an important element of any research design, and here I want to introduce one of the most fundamental distinctions in research design nomenclature: ***cross-sectional*** versus ***longitudinal*** studies. A *cross-sectional* study is one that takes place at a single point in time. In effect, we are taking a 'slice' or cross-section of whatever it is we're observing or measuring. A *longitudinal* study is one that takes place over time -- we have at least two (and often more) waves of measurement in a longitudinal design.

A further distinction is made between two types of longitudinal designs: ***repeated measures*** and ***time series***. There is no universally agreed upon rule for distinguishing these two terms, but in general, if you have two or a few waves of measurement, you are using a *repeated measures* design. If you have many waves of measurement over time, you have a *time series*. How many is 'many'? Usually, we wouldn't use the term time series unless we had at least twenty waves of measurement, and often far more. Sometimes the way we distinguish these is with the analysis methods we would use. Time series analysis requires that you have at least twenty or so observations. Repeated measures analyses (like repeated measures ANOVA) aren't often used with as many as twenty waves of measurement.

Types of Relationships

A relationship refers to the *correspondence* between two variables. When we talk about types of relationships, we can mean that in at least two ways: the *nature* of the relationship or the *pattern* of it.

The Nature of a Relationship

While all relationships tell about the correspondence between two variables, there is a special type of relationship that holds that the two variables are not only in correspondence, but that one *causes* the other. This is the key distinction between a simple **correlational relationship** and a **causal relationship**. A correlational relationship simply says that two things perform in a synchronized manner. For instance, we often talk of a correlation between inflation and unemployment. When inflation is high, unemployment also tends to be high. When inflation is low, unemployment also tends to be low. The two variables are correlated. But knowing that two variables are correlated does not tell us whether one *causes* the other. We know, for instance, that there is a correlation between the number of roads built in Europe and the number of children born in the United States. Does that mean that is we want fewer children in the U.S., we should stop building so many roads in Europe? Or, does it mean that if we don't have enough roads in Europe, we should encourage U.S. citizens to have more babies? Of course not. (At least, I hope not). While there is a relationship between the number of roads built and the number of babies, we don't believe that the relationship is a *causal* one. This leads to consideration of what is often termed the **third variable problem**. In this example, it may be that there is a third variable that is causing both the building of roads and the birthrate, that is causing the correlation we observe. For instance, perhaps the general world economy is responsible for both. When the economy is good more roads are built in Europe and more children are born in the U.S. The key lesson here is that you have to be careful when you interpret correlations. If you observe a correlation between the number of hours students use the computer to study and their grade point averages (with high computer users getting higher grades), you *cannot* assume that the relationship is *causal*: that computer use improves grades. In this case, the third variable might be socioeconomic status -- richer students who have greater resources at their disposal tend to both use computers and do better in their grades. It's the resources that drives both use and grades, not computer use that causes the change in the grade point average.


Patterns of Relationships

We have several terms to describe the major different types of patterns one might find in a relationship. First, there is the case of **no relationship** at all. If you know the values on one variable, you don't know anything about the values on the other. For instance, I suspect that there is no relationship between the length of the lifeline on your hand and your grade point average. If I know your GPA, I don't have any idea how long your lifeline is.
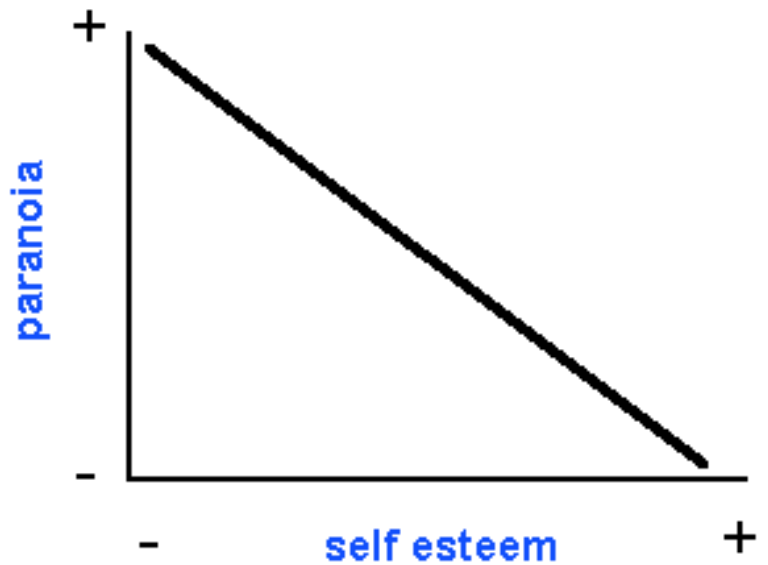
Then, we have the **positive relationship**. In a positive relationship, high values on one variable are associated with high values on the other and low values on one are associated with low values on the other. In this example, we assume an idealized positive relationship between years of education and the salary one might expect to be making.

On the other hand a **negative** relationship implies that high values on one variable are associated with low values on the other. This is also sometimes termed an **inverse** relationship. Here, we show an idealized negative relationship between a measure of self esteem and a measure of paranoia in psychiatric patients.

These are the simplest types of relationships we might typically estimate in research. But the pattern of a relationship can be more complex than this. For instance, the figure on the left shows a relationship that changes over the range of both variables, a curvilinear relationship. In this example, the horizontal axis represents dosage of a drug for an illness and the vertical axis represents a severity of illness measure. As dosage rises, severity of illness goes down. But at some point, the patient begins to experience negative side effects associated with too high a dosage, and the severity of illness begins to increase again.

Variables

You won't be able to do very much in research unless you know how to talk about variables. A *variable* is *any entity that can take on different values*. OK, so what does that mean? Anything that can vary can be considered a variable. For instance, *age* can be considered a variable because age can take different values for different people or for the same person at different times. Similarly, *country* can be considered a variable because a person's country can be assigned a value.

Variables aren't always 'quantitative' or numerical. The variable 'gender' consists of two text values: 'male' and 'female'. We can, if it is useful, assign quantitative values instead of (or in place of) the text values, but we don't have to assign numbers in order for something to be a variable. It's also important to realize that variables aren't only things that we measure in the traditional sense. For instance, in much social research and in program evaluation, we consider the treatment or program to be made up of one or more variables (i.e., the 'cause' can be considered a variable). An educational program can have varying amounts of 'time on task', 'classroom settings', 'student-teacher ratios', and so on. So even the program can be considered a variable (which can be made up of a number of sub-variables).

An *attribute* is a specific value on a variable. For instance, the variable *sex* or *gender* has two attributes: *male* and *female*. Or, the variable *agreement* might be defined as having five attributes:

- 1 = strongly disagree
- 2 = disagree
- 3 = neutral
- 4 = agree
- 5 = strongly agree

Another important distinction having to do with the term 'variable' is the distinction between an *independent* and *dependent* variable. This distinction is particularly relevant when you are investigating cause-effect relationships. It took me the longest time to learn this distinction. (Of course, I'm someone who gets confused about the signs for 'arrivals' and 'departures' at airports -- do I go to arrivals because I'm arriving at the airport or does the person I'm picking up go to arrivals because they're arriving on the plane!). I originally thought that an independent variable was one that would be free to vary or respond to some program or treatment, and that a dependent variable must be one that *depends* on my efforts (that is, it's the *treatment*). But this is entirely backwards! In fact **the independent variable is what you (or nature) manipulates** -- a treatment or program or cause. The **dependent variable is what is affected by the independent variable** -- your effects or outcomes. For example, if you are studying the effects of a new educational program on student achievement, the program is the independent variable and your measures of achievement are the dependent ones.

Finally, there are two traits of variables that should always be achieved. Each variable should be **exhaustive**, it should include all possible answerable responses. For instance, if the variable is "religion" and the only options are "Protestant", "Jewish", and "Muslim", there are quite a few religions I can think of that haven't been included. The list does not exhaust

all possibilities. On the other hand, if you exhaust all the possibilities with some variables -- religion being one of them -- you would simply have too many responses. The way to deal with this is to explicitly list the most common attributes and then use a general category like "Other" to account for all remaining ones. In addition to being exhaustive, the attributes of a variable should be **_mutually exclusive_**, no respondent should be able to have two attributes simultaneously. While this might seem obvious, it is often rather tricky in practice. For instance, you might be tempted to represent the variable "Employment Status" with the two attributes "employed" and "unemployed." But these attributes are not necessarily mutually exclusive -- a person who is looking for a second job while employed would be able to check both attributes! But don't we often use questions on surveys that ask the respondent to "check all that apply" and then list a series of categories? Yes, we do, but technically speaking, each of the categories in a question like that is its own variable and is treated dichotomously as either "checked" or "unchecked", attributes that _are_ mutually exclusive.

Hypotheses

An hypothesis is a specific statement of prediction. It describes in concrete (rather than theoretical) terms what you expect will happen in your study. Not all studies have hypotheses. Sometimes a study is designed to be exploratory (see inductive research). There is no formal hypothesis, and perhaps the purpose of the study is to explore some area more thoroughly in order to develop some specific hypothesis or prediction that can be tested in future research. A single study may have one or many hypotheses.

Actually, whenever I talk about an hypothesis, I am really thinking simultaneously about *two* hypotheses. Let's say that you predict that there will be a relationship between two variables in your study. The way we would formally set up the hypothesis test is to formulate two hypothesis statements, one that describes your prediction and one that describes all the other possible outcomes with respect to the hypothesized relationship. Your prediction is that variable A and variable B will be related (you don't care whether it's a positive or negative relationship). Then the only other possible outcome would be that variable A and variable B are *not* related. Usually, we call the hypothesis that you support (your prediction) the *alternative* hypothesis, and we call the hypothesis that describes the remaining possible outcomes the *null* hypothesis. Sometimes we use a notation like H_A or H_1 to represent the alternative hypothesis or your prediction, and H_O or H_0 to represent the null case. You have to be careful here, though. In some studies, your prediction might very well be that there will be no difference or change. In this case, you are essentially trying to find support for the null hypothesis and you are opposed to the alternative.

If your prediction specifies a direction, and the null therefore is the no difference prediction and the prediction of the opposite direction, we call this a *one-tailed hypothesis*. For instance, let's imagine that you are investigating the effects of a new employee training program and that you believe one of the outcomes will be that there will be *less* employee absenteeism. Your two hypotheses might be stated something like this:

The null hypothesis for this study is:

H_O: As a result of the XYZ company employee training program, there will either be no significant difference in employee absenteeism or there will be a significant *increase*.

which is tested against the alternative hypothesis:

H_A: As a result of the XYZ company employee training program, there will be a significant *decrease* in employee absenteeism.

In the figure on the left, we see this situation illustrated graphically. The alternative hypothesis -- your prediction that the program will decrease absenteeism -- is shown there. The null must account for the other two possible conditions: no difference, or an increase in absenteeism. The figure shows a hypothetical distribution of absenteeism differences. We can see that the term "one-tailed" refers to the tail of the distribution on the outcome variable.

When your prediction does *not* specify a direction, we say you have a **two-tailed hypothesis**. For instance, let's assume you are studying a new drug treatment for depression. The drug has gone through some initial animal trials, but has not yet been tested on humans. You believe (based on theory and the previous research) that the drug will have an effect, but you are not confident enough to hypothesize a direction and say the drug will reduce depression (after all, you've seen more than enough promising drug treatments come along that eventually were shown to have severe side effects that actually worsened symptoms). In this case, you might state the two hypotheses like this:

The null hypothesis for this study is:

H_O: As a result of 300mg./day of the ABC drug, there will be no significant difference in depression.

which is tested against the alternative hypothesis:

H_A: As a result of 300mg./day of the ABC drug, there will be a significant difference in depression.

The figure on the right illustrates this two-tailed prediction for this case. Again, notice that the term "two-tailed" refers to the tails of the distribution for your outcome variable.
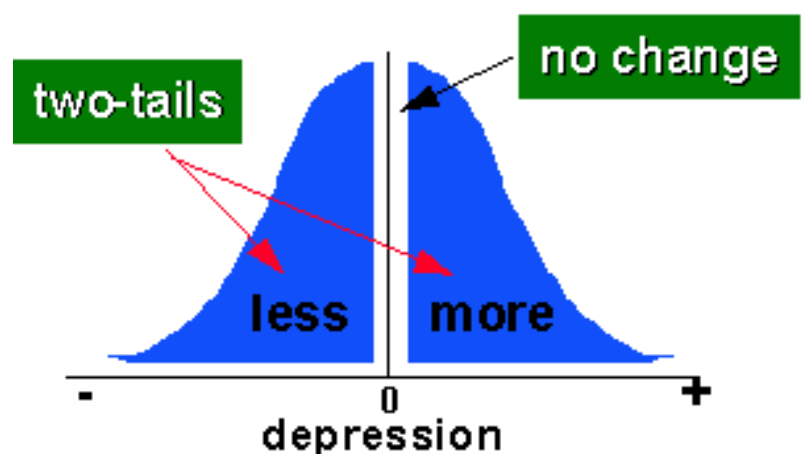
The important thing to remember about stating hypotheses is that you formulate your prediction (directional or not), and then you formulate a second hypothesis that is mutually exclusive of the first and incorporates all possible alternative outcomes for that case. When your study analysis is completed, the idea is that you will have to choose between the two hypotheses. If your prediction was correct, then you would (usually) reject the null hypothesis and accept the alternative. If your original prediction was not supported in the data, then you will accept the null hypothesis and reject the alternative. The logic of hypothesis testing is based on these two basic principles:

- the formulation of two mutually exclusive hypothesis statements that, together, exhaust all possible outcomes
- the testing of these so that one is necessarily accepted and the other rejected

OK, I know it's a convoluted, awkward and formalistic way to ask research questions. But it

encompasses a long tradition in statistics called the ***hypothetical-deductive model***, and sometimes we just have to do things because they're traditions. And anyway, if all of this hypothesis testing was easy enough so anybody could understand it, how do you think statisticians would stay employed?

Types of Data

We'll talk about data in lots of places in The Knowledge Base, but here I just want to make a fundamental distinction between two types of data: **qualitative** and **quantitative**. The way we typically define them, we call data 'quantitative' if it is in numerical form and 'qualitative' if it is not. Notice that qualitative data could be much more than just words or text. Photographs, videos, sound recordings and so on, can be considered qualitative data.

Personally, while I find the distinction between qualitative and quantitative data to have some utility, I think most people draw too hard a distinction, and that can lead to all sorts of confusion. In some areas of social research, the qualitative-quantitative distinction has led to protracted arguments with the proponents of each arguing the superiority of their kind of data over the other. The quantitative types argue that their data is 'hard', 'rigorous', 'credible', and 'scientific'. The qualitative proponents counter that their data is 'sensitive', 'nuanced', 'detailed', and 'contextual'.

For many of us in social research, this kind of polarized debate has become less than productive. And, it obscures the fact that qualitative and quantitative data are intimately related to each other. *All quantitative data is based upon qualitative judgments; and all qualitative data can be described and manipulated numerically.* For instance, think about a very common quantitative measure in social research -- a self esteem scale. The researchers who develop such instruments had to make countless judgments in constructing them: how to define self esteem; how to distinguish it from other related concepts; how to word potential scale items; how to make sure the items would be understandable to the intended respondents; what kinds of contexts it could be used in; what kinds of cultural and language constraints might be present; and on and on. The researcher who decides to use such a scale in their study has to make another set of judgments: how well does the scale measure the intended concept; how reliable or consistent is it; how appropriate is it for the research context and intended respondents; and on and on. Believe it or not, even the respondents make many judgments when filling out such a scale: what is meant by various terms and phrases; why is the researcher giving this scale to them; how much energy and effort do they want to expend to complete it, and so on. Even the consumers and readers of the research will make lots of judgments about the self esteem measure and its appropriateness in that research context. What may look like a simple, straightforward, cut-and-dried quantitative measure is actually based on lots of qualitative judgments made by lots of different people.

On the other hand, all qualitative information can be easily converted into quantitative, and there are many times when doing so would add considerable value to your research. The simplest way to do this is to divide the qualitative information into units and number them! I know that sounds trivial, but even that simple nominal enumeration can enable you to organize and process qualitative information more efficiently. Perhaps more to the point, we might take text information (say, excerpts from transcripts) and pile these excerpts into piles of similar statements. When we do something even as easy as this simple grouping or piling task, we can describe the results quantitatively. For instance, if we had ten statements and we grouped these into five piles (as shown in the figure), we could describe the piles using

a 10 x 10 table of **0**'s and **1**'s. If two statements were placed together in the same pile, we would put a **1** in their row-column juncture. If two statements were placed in different piles, we would use a **0**. The resulting matrix or table describes the grouping of the ten statements in terms of their similarity. Even though the data in this example consists of qualitative statements (one per card), the result of our simple qualitative procedure (grouping similar excerpts into the same piles) is *quantitative* in nature. "So what?" you ask. Once we have the data in numerical form, we can manipulate it numerically. For instance, we could have five different judges sort the 10 excerpts and obtain a 0-1 matrix like this for each judge. Then we could average the five matrices into a single one that shows the proportions of judges who grouped each pair together. This proportion could be considered an estimate of the similarity (across independent judges) of the excerpts. While this might not seem too exciting or useful, it is exactly this kind of procedure that I use as an integral part of the process of developing 'concept maps' of ideas for groups of people (something that *is* useful!).

Sorting of 10 qualitative items

Binary Square Similarity Matrix for the sort

	1	2	3	4	5	6	7	8	9	10
1	1	1	0	0	0	1	0	0	1	0
2	1	1	0	0	0	1	0	0	1	0
3	0	0	1	1	0	0	0	0	0	0
4	0	0	1	1	0	0	0	0	0	0
5	0	0	0	0	1	0	0	1	0	0
6	1	1	0	0	0	1	0	0	1	0
7	0	0	0	0	0	0	1	0	0	0
8	0	0	0	0	1	0	0	1	0	0
9	1	1	0	0	0	1	0	0	1	0
10	0	0	0	0	0	0	0	0	0	1