Creating Knowledge: Designing an Empirical Research Project

Dr. Annette Freyberg-Inan, Civic Education Project Visiting Faculty Fellow
University of Bucharest, Faculty of Political Science

This chapter is meant to introduce you to the logic of an empirical research project. It should be useful especially to students in the social sciences, at both the undergraduate and the graduate level, who don't yet know much about this kind of work. I explain, first of all, what empirical research is all about. I then try to give you guidance on how you can design and carry out your own empirical research project. I explain the practical steps you have to take and warn you about the main problems which might confront you at each stage of the process. The chapter concludes with some advice on where you can find more information on this subject.

1. Learning About Empirical Research

Before you get started on your empirical research project, it's important to know a little bit about what it actually is that you will, or should, be doing. First you should have some idea about what empirical research is all about. You'll find out that the goal of this kind of research is to teach us something about how things are in the world, and why they are this way. Empirical research is all about creating factual knowledge. And so the second question will be: What can you do to create such knowledge? You will find out that it is very much within your means to make a real contribution to science.

1.1. What Is Empirical Research?

If you've heard anything about the philosophy of science before, or about methodology in any of the social sciences, then you've probably come across the distinction between *normative* and *empirical* research. This section of the chapter will introduce this distinction to you, and explain what is so special about empirical research.

Normative Research

To put it in a nutshell, normative research tries to answer questions about *how things should be* or *how we should judge them*. For example, what is the ideal kind of government, democracy or enlightened dictatorship? Should we be more liberal in our social lives, or more conservative? Is it morally acceptable to take the lives of other human beings, under some circumstances, or is it always morally wrong? These are the kinds of questions about which one can debate, seemingly without end, because the answers have a lot to do with personal judgments. Maybe they really have no "true" or "false" answers. But still we feel we have to debate these kinds of questions, because they are important, and the answers we decide to give to them have a lot of influence on our lives. For example, you probably have a pretty strong opinion about whether you want to live in a democracy or not, or whether you should be allowed to lead a "liberal" social life, or not. Well, if you feel strongly about these things, then you'll want others to share your opinions, so you can make sure that the world around you is organized in a way you like. Normative research basically does exactly that: It answers a question about how things should be and tries to persuade other people to share this judgment.

When you are trying to persuade someone to share in the conclusions of your normative research, you try to do so by, first of all, starting off with reasonable assumptions about the world we live in. To assume that all people are angels, for example, or that without any kind of government we would be able to live in perpetual peace and harmony, is a bit risky, because people might think from the very start of your argument that you are making ridiculous judgments about the world, and so they won't believe your conclusions either. But let's say that your assumptions make sense to most people who will read your research. The way you get to your conclusions if you are conducting normative research is by reasoned argument. Basically, you have to make sure that your argument follows the rules of logic, meaning that you don't jump from one point to the next without connecting them in a way that makes the reader follow what you are saying. You also have to make sure that any new assumptions you introduce on the way are either commonsensical or that you defend them well. If you are good at arguing in this way, chances are that you can persuade your audience that your conclusions have some merit. At least they deserve consideration. And this would mean that you have succeeded with your normative research project.[1]

You can take as an example the philosopher Jean-Jacques Rousseau, who wrote in the 18th century about human nature and the way we should organize ourselves. He had made some observations about how people are, and the one thing that struck him most was that every one of us has so much potential to make something special of him- or herself. He thought that this potential for self-development or self-realization was a precious gift, and therefore we should make sure that we organize our lives in such a way that we can make the best of this gift. He then thought about what kind of political organization would be best to allow people to develop themselves, and this led him to the ideal of a small direct democracy, where everyone could be involved in making decisions which concerned everybody. He also gave his own hometown, the city of Geneva, as an example of how this might work. Basically, he did three things right, and that is probably why we still talk about his ideas today and haven't just thrown them on the dust-pile of history. First, while he had his own ideas about human nature, he said things with which many people could agree, and which were also kind of fashionable at the time and still are today: Individuals are unique, and we all have a potential to make something of ourselves and should be allowed to try. This is an example for how his assumptions were acceptable, and didn't seem too outrageous to most. Then he made a logical and elegant argument that such human nature could best flourish in the kind of ideal political system he described. And finally, he made reference to reality and supported his argument with examples. This is how his normative research became so successful.

Empirical Research

But what we are really concerned about in this chapter is the other kind of research, which is called empirical. Empirical research, in a nutshell, asks questions about *what is*. In other words, empirical researchers are primarily trying to find out what is going on in the world, and not to judge whether what goes on is good or bad, or should be different. For example, the question whether Russia is a democracy or not is an empirical question, because once we define what we mean by a democracy we can look at Russia to see whether those characteristics are in place there. If they are, we conclude that Russia is indeed a democracy. If not, we conclude that

[1] On the strategy of persuasion of normative research see also the chapter by Randy Gabrielson.

it is not. We are not interested in whether it would be bad if Russia were not a democracy, or in what the Russian government should look like. Those are normative questions.

Let me give you another example: Let's say we observe that people often kill other people even though their religions might teach them that they should not kill. We might argue that this is wrong, and that would be a normative argument. We might also argue that it is o.k., and that would be a normative argument, too. But we might also simply try to study how often it happens, and under what circumstances. And that would be empirical research.

The simplest kind of empirical research tries to come up with a *description* of a particular aspect of reality. What kind of government does Russia have? How does the constitution of France compare with the constitution of the United States of America? These kinds of questions are relatively easy to answer, because all it takes is to collect information about your topic of study, the government of Russia or the constitutions of France and the US, and summarize it. The second example is a little more complicated because it asks you to compare, and to compare means that you also have to develop criteria for your comparison. But it's still quite easy, because you can usually rely on what other people have found out about the topic, and you don't really have to study it directly yourself. This means that you are not doing what is called *original* research. While you are still an undergraduate student, you might get away with always writing descriptive non-original papers like these, but you are not going to get away with it as a graduate student, much less if you want to become a successful researcher or academic in your own right. In order to become such an expert in your field, you have to learn how to create knowledge yourself. The next section will tell you how you can do this.

1.2. How Can You Create Knowledge?

Description

I just told you that *descriptive empirical research* is relatively easy, but of course that doesn't mean that it is worthless. In fact, descriptive empirical research can also be original. This happens when you describe something that noone has described before, or at least not in the way that you do. Of course this can be very important. Think, for example, that you are living in a country in which an ethnic conflict is developing. You can observe the tension first hand and see how the conflict progresses. If you then write about what you observe, even if you are just describing, you are doing something very useful for your field of study, because you are letting others know what is happening, so that, together, everyone who cares can think about how to explain the conflict and how perhaps to prevent things from getting worse.

Explanation

However, generally speaking, when you are trying to create new knowledge, a type of research that is both more challenging and usually more interesting is *explanatory empirical research*. As the name tells you, in this kind of research you are trying to *explain* something. For example, you might observe that many people in your country stay home on election day. *Why* do they not use their right to vote? Or, you might wonder, *how come* that Pakistan and India can't seem to get along since they became independent countries? In a way, the most important quality of any scientist is curiosity. If you don't like to wonder why things are the way they are, if you don't like to ask lots of "why" questions, then you're better off doing something else.

Most scientific discoveries begin with mysteries, large or small. These mysteries we call *puzzles*. When something is puzzling to us that means we find it worth explaining. Maybe we

didn't expect this thing to happen. For example, many people were puzzled in 1989, when the Soviet Union began to come apart. They hadn't expected that to happen. And immediately many social scientists, historians, economists, and anybody else interested began to try and come up with some explanations. Sometimes we don't only not expect something to happen, but we actually feel we had a right to expect something else. Then we are even more puzzled. For example, a psychologist might feel more need to explain why a young person who grew up in pleasant and supportive circumstances became a violent criminal because we don't normally expect that to happen. The bottom line is that we seek for explanations for the things that puzzle us. And explanations always link *causes* and *effects*.

If I am interested in explaining why many people in, say, Bulgaria, didn't vote in the last elections, then this non-voting that I am observing is the effect I am trying to explain. For this effect, there might be many causes: Maybe it rained on election day, and people didn't want to get wet. Maybe they didn't like the parties or the candidates that they could choose from. Maybe they simply have no enthusiasm for the democratic process or don't believe that their vote matters. Any number of possible causes could be important, and if I want to be thorough it is important that I try to think of all of them. But of course I would consider some causes more important than others, and the most important task of explanatory empirical research is exactly to find out the most important cause or causes of a particular effect. This search for the most important causes is related to a traditional ambition of science: the ambition to create knowledge that should be as general as possible. *Generality* is a central objective of science.

To explain, let me return to our example: Someone could write a book about the last general elections in Bulgaria, and in this book explain one by one all the possible reasons for why people might have stayed home on that particular election day in that particular country. Such a book might be interesting to read, at least for people who have a particular interest in Bulgaria, but, chances are, we would not consider such a book particularly "scientific." Why? Because the author did not do any of the hard work of trying to figure out what were the most important causes for non-voting in Bulgaria, and he did not tell us anything about whether these same causes might make people stay home on election day at other times or in other countries as well. Along comes the prototypical social scientist, and you may be able to imagine what kind of book she might write on the subject. She would probably first think of the various causes that may generally make people disinclined to vote. Then she would draw on the expertise of people who have studied the issue of non-voting to suggest which of those might be the most important causes. Then she would study the last elections in Bulgaria to see, first, if those causes indeed existed in that place and in that time, second, if they really made people stay home on election day, and, third, if there were not some other important reasons why people didn't vote. Finally, she would write about what the results of her research on this particular election in Bulgaria might mean for non-voting in general, meaning in other times and places as well.

See the difference? Explanatory empirical research tries to come up with general explanations for things which we are interested in explaining. In order to be general, we have to decide what's important and what is not so important. You may hear this principle referred to as the ideal of *parsimony*. An explanation is parsimonious if it can explain a lot with a little, meaning if it can explain a whole class of events by suggesting as few causes as possible. For example, if I explain that everywhere in the world people vote or do not vote depending solely on the weather, that would be a parsimonious explanation. (Of course it would also be wrong, so keep in mind that parsimonious is not the same as correct!) What the first author did in the example above, was the opposite of being parsimonious, because he gave very many reasons for

only one case of low voter turnout. The second author, on the other hand, was striving for the scientific ideal of parsimony.

Prediction

Another traditional goal of science is *prediction*. Ideally, we don't only want to know what has been going on in the world so far, and what is happening now, but we also want to be able to predict what might happen in the future. The idea is that if we know the likely effects of things, we can have more control over what happens to us. For example, if we know something about how a conflict between two states can turn into war, we would know better how to prevent war from happening, at least as long as we have control over some of its causes. This desire to have more control over our lives also helps explain the need for parsimonious explanation. This is because only if we know the most important links between causes, such as diplomatic postures, and effects, such as war, will we be able to correctly predict which effects are likely to follow which causes. And only then can the social sciences be useful in actually controlling our fates.

The Problem of Many Causes

If you've read carefully up to this point, you should have a question in your mind right now. I wrote earlier about how our social scientist came to suggest which might be the most important causes for non-voting: She relied on the expertise of people who had studied non-voting before. But this doesn't tell us anything about how this kind of expertise is created in the first place. How are we supposed to know which causes to link to which effects, and which of these links are more important than others? It seems rather like having to draw a route for a hiking trip while either not knowing where you are right now or not knowing where you want to go. Is it possible?

Earlier you learned that if you do normative research the way to persuade others of your expertise is to use a logical line of argument which relies on defensible assumptions. In empirical research, the way to persuade others that your conclusions are correct is by referring to *observation*, by showing them evidence out there in the real world which they can see with their own eyes if they only look. What you are trying to observe, if you look at the world with the eyes of a scientist, are *patterns*. And it is these patterns which can help you, if you observe them often enough and in enough different contexts, to make the links from effects to their most important causes.

Let's think again of the phenomenon of people having the right to vote but not using it. If I take a look at different elections in different democratic countries I might see that people vote in fewer numbers whenever the ideological differences between the political parties are quite small. I might also see that wherever the ideological differences between parties are great, voting participation is quite high. What I have observed here is a pattern which links ideological difference between parties (as a cause) and voting (as an effect) in a particular way: When ideological differences between parties are greater, voting participation is higher. Perhaps without even knowing it, you have now arrived at a *hypothesis*, a proposed relationship between a cause and an effect. The development and the testing of hypotheses are the everyday stuff of explanatory empirical research, and they will be explained in more detail in the next section of

the chapter. But before we get there, I should give you some background on the role of *theory* in empirical research.[2]

The Role of Theory

Theory is usually defined as a set of related statements about a particular aspect of reality. It is possible to make a distinction between normative theory and empirical theory, where the first is more strongly concerned with judging something and the second with explaining something. But for our purposes that distinction is not really relevant, because any useful theory does something to explain how a number of causes and effects are linked to produce a particular outcome.

I told you that scientists are always on the lookout for patterns. This is true for all of them, whether they are natural or social scientists. A scientific discovery is really the discovery of a pattern which we believe might hold the solution to a puzzle. We can often stumble upon such patterns just by walking around with our eyes open. For example, we observe that people have a tendency to smile back at us if we smile at them first. Most of us live our lives knowing about this pattern without ever really thinking about it. But sometimes patterns are not so obvious. We have to search for them, and in order to search successfully, we should know something about where to look. That is the first task of theory: To tell us where we might go to search for patterns. But there is more: A pattern is not the same as an explanation, and we don't just want to find out that two things are linked in a certain way. We also want to know why. How come people smile back if you smile at them? What is the mechanism at work here? The second task of theory is to provide the explanation for why things are linked together in patterns the way they are.

Some people don't see much point in theory, because, they say, it's just about people inventing ideas and has little to do with reality and practical affairs. But it wouldn't be fair to say that theorists don't contribute to the creation of real knowledge, because they do. For example, important thinkers in the social sciences, like Max Weber or Emile Durkheim, not only collected an enormous amount of knowledge about our world but also taught us new ways of interpreting this information, new ways of understanding what is going on, and even new ways of changing the world and maybe improving it. Basically, they have told entire generations of social scientists where to look for patterns and how to explain them. You don't have to be an *empiricist*, that is someone primarily concerned with discovering patterns in the real world, in order to contribute to the creation of knowledge. In fact, the empiricist would really be a nobody without the theorist, because, without theory, how would she make sense of the patterns she sees in the world around, even if she stumbles upon them by accident?

We cannot really think without theorizing, even when it comes to the most mundane aspects of our lives. For example, if you think that your brother doesn't do his fair share of work in the house because your parents have always let him get away with not doing it and now he doesn't even feel bad anymore for being lazy, you have a theory linking educational practices, communal orientation, and work ethics. People who think that they can learn about the real world without using theory are sometimes called "barefoot empiricists" in jest. You don't want to be barefoot on the rocky path of science, and this means that you should keep asking those

[2] If you are interested in the role of theory in research you might also want to read the chapter by Paul Roe.

"why" questions all the way through, from the selection of your topic to the interpretation of your results.

2. Learning How to Do Empirical Research

This part of the chapter is meant to give you concrete advice on how you can actually plan an empirical research project of your own. It starts with some advice on how you can find a topic to study. You will find that, even if you start with a vague interest, it is important that you narrow your topic down until you have a specific hypothesis that you can put to the test. How to design and carry out such a test is the topic of the second sub-section below. The final sub-section is dedicated to a discussion of how to draw legitimate conclusions from the work you have done.

2.1. How Can You Find a Topic?

General Rules

For many students, maybe one of the most agonizing parts of carrying out an empirical research project is coming up with a suitable topic. But it shouldn't really be all that difficult, because there are only two main rules to keep in mind: The first rule is: The topic must be something that really interests you. Original empirical research, if you want to do it right, will consume much of your time and energy, and if you have a topic which you are not really passionate about, you will most probably hate the assignment and, as a result, do less well than you could otherwise. The second rule to keep in mind is that your topic should be manageable. This sounds commonsensical, but I have seen it happen again and again that students chose topics which were vastly too complex to be researched with the very limited means which students typically have, in terms of time as well as methodological expertise. If you stick with an overly ambitious topic, you'll most probably enter a state of panic in the middle of your research, when you find out you can't really do what you set out to do, and the quality of your work will suffer accordingly. The general rule is: It's better to complete a modest task well than do try something ambitious and fail! Especially if you are a newcomer to this kind of work, what you really need is practice. You can read a thousand books like this one, they will never make up for what you'll learn by doing research of your own, putting all the theory into practice, and learning from your own mistakes. If your goal is to become really good at this kind of work, then it's important that you work your way up slowly, starting with small and uncomplicated projects, and taking on more only once you are really ready to do so.

Developing a Hypothesis

The way to start small, and usually a good way to begin your project even if you are already an expert in empirical research, is to come up with a hypothesis which you would like to test. A hypothesis, if you remember, is a statement that suggests a particular relationship between a cause and an effect. To find a hypothesis which you might like to test, you can start in one of three ways: You can start with an effect; you can start with a cause; or you can start with a pattern. The first is probably the most common way. Ask yourself: Is there something which I would like to be able to explain? Maybe there is something which I find so bad that I would like to know how to stop it from happening, like, for example, war breaking out between my people and our neighbors. Or maybe there is something which I find so good that I would like to be able

to help it to happen, like, for example, democracy to succeed in my country. In such a case, you start with an effect, like the breakout of war between neighbors or the success of democracy, and then you go looking for the probable causes. Maybe you can think of some probable causes by yourself, but it's better if you do a bit of research on your topic of interest in the very beginning, to draw on the expertise of people with more experience in the field and thereby to make sure that the causes you will end up suggesting will seem plausible to your audience. So this would be the stage when you first have to search the libraries and the internet, to come up with the probable causes for what you are interested in explaining.

You should take note of all the important causes you come across, but you have to decide on only one to put into your hypothesis. This should be the one you consider the most important or the most interesting to study. Sometimes a hypothesis can mention more than one cause or more than one effect. An example for such a so-called *compound hypothesis* is the following: Poverty tends to increase in cities as a result of economic decline combined with population growth. Here you have only one effect, the increase of poverty, but two causes, economic decline and population growth. Compound hypotheses are more difficult to test than simple ones, so I strongly suggest you stay away from them for the time being. There are methods to test more than one relationship at a time, even whole sets of causes and effects, but they are extremely complicated methods. It is very much in your interest for the time being to stick with testing simple relationships between one cause and one effect.

This you should also keep in mind when you begin your quest for a hypothesis from the opposite end, namely with a cause. Ask yourself, is there something which I can observe happening and I wonder what effects it might have? For example, maybe you know that most of the students in your department have jobs. This might make you wonder whether perhaps their grades suffer because they don't only have to study but also to work. And you could end up with the hypothesis: Having a job makes students' grades suffer. Finally, the third way to come up with a hypothesis is to stumble upon an entire pattern that interest you. This way you come up with the cause and the effect at the same time. For example, if you realize that all of your male friends who are really into computers have no girlfriends, whereas those who don't care so much about computers have girlfriends, the following hypothesis suggests itself to you: A strong interest in computers makes it less likely for young men to be involved in romantic relationships. (Of course, as part of your research, you'll also want to think about *why* this might be the case!)

Formulating Your Hypothesis

You are finished with finding a topic once you have formulated a hypothesis which can be tested. Your hypothesis must fulfill the following formal criteria:

Rule No. 1: You need an independent and a dependent variable

The first rule for your hypothesis is that it must contain a cause and an effect. That sounds simple, but in order to write your hypothesis correctly you must pay very close attention to some details of technical language. In your hypothesis, both the cause and the effect will be called *variables*. What we have called the cause so far will be called the *independent variable* now. What we have called the effect will be called the *dependent variable*.

The definition of a variable is almost absurdly commonsensical: A variable is a characteristic (sometimes called a *property* of an object) that varies. But to identify variables is not quite as easy as it sounds. To identify the variables in your hypothesis you have to ask yourself: What exactly is it that varies in my hypothesis? Let's take the following example:

Crime rates tend to increase in cities as a result of population growth. What varies here? Crime rates, first of all. They can go up or down. They are also what we are interested in explaining, so they must be the dependent variable. The other thing that varies is what? Be very careful here! It is not population growth that varies, but it is the size of the population. That is what gets bigger or smaller. If it gets bigger, we believe that crime rates will also get higher. Therefore, size of population is our independent variable.

Rule No. 2: You must specify the relationship between the variables

Your hypothesis must also give a very clear idea about the exact relationship between your independent and your dependent variable. How does one affect the other? Often, variables will be *quantitative*, meaning that their value can be expressed in numbers and they can grow larger or smaller, just like in the above example. In such a case, if the independent variable grows, does the dependent also grow, or does it shrink? Or maybe you expect it to do nothing at all (in which case you would suggest that there is in fact no relationship at all between your variables of interest). In any case, what you have to say in your hypothesis is how you expect the dependent variable to change if the independent one grows.

At other times, however, variables will be *qualitative* by nature. This means that their value can not be expressed in numbers. They can't grow bigger or smaller, but only "different." An example for such a variable is religion. People can be Christian, Buddhist, Hindu, etc., so their religion can vary, but you can't say that one type of religion is somehow higher than the other, since that would make no sense. In such a case, the basic rule is the same: You must say how you expect the dependent variable to change if the independent one changes in a particular way. An example is the simple hypothesis: Buddhists are more likely to practice meditation than Christians. This suggests, and you should not take this technical language too literally, that as people move from being Christian to being Buddhist (changing the value of the independent variable) they become more likely to meditate (increasing the value of the quantitative dependent variable "amount of meditation").

A variable, whether quantitative or qualitative, that is expressed in such a way that it can take on only one of two values is called a *dichotomous* variable. Dichotomous variables are easier to handle than others, partly because some of the statistical techniques which can be used to analyze them are quite simple. So another recommendation for beginners might be to write your hypothesis in such a way that either one or both variables are dichotomous. Changing the above example slightly, you could say simply that Christians are less likely to meditate than non-Christians. In that case, the independent variable can only take on the value of Christian or non-Christian, and the dependent one can only take on the value of presence or absence of meditation. However, if you determine that the relationship you are interested in cannot really be tested by forcing the cause or the effect into the form of a dichotomous variable, it's better if you follow your real interests.

Rule No. 3: You must specify your population

The third rule for your hypothesis is that it should say exactly where you suspect this relationship you are suggesting to exist. In other words, what are the objects which "own" the relationship? For example, when you say "wealth increases happiness," you are being unclear about whether you are making a suggestion about nations or individuals. What you should say instead is: "When comparing nations, more wealthier ones also tend to have more happier populations;" or: "When comparing individuals, we find that as they grow more wealthy they

also become happier." The technical literature calls this specifying your *theoretical population*. The things that make up your theoretical population, the individual countries or people, are your *objects of analysis*. It's important to be clear about what they are, because if you are not you also won't know how to properly test your hypothesis!

Rule No. 4: Try to find the right level of generality

Another rule is that your hypothesis should be general as opposed to referring to a specific case. This is because of the scientific goal of generality which I mentioned before. However, when you get down to the hard work of designing a test for your hypothesis, you will find that you will have to restrict yourself quite a bit. You can't study the whole world in detail to find out whether your general hypothesis is correct or not, but your observations will always be limited. That's why in your conclusions you have to be careful about the extent to which you can generalize from what you have observed. To find the ideal level of generality of your hypothesis is a bit like walking on a tight rope. On the one side, you don't want to be so general that you can't possibly find enough evidence to test your claim. On the other hand, you don't want to be so specific that no general lessons can be drawn from your research.

I can give you an example: If you are interested in the effect of economic welfare on the success of democracy, you could come up with a very general hypothesis: The success of democracy depends on economic welfare. But that would be such a general statement that it might be very hard to collect enough evidence for a proper test. On the other hand, the following hypothesis might be easy to test but is probably too specific: The success of Polish democracy between 1990 and 1992 depended on the economic welfare of the nation. An in-between level of generality, which makes the topic interesting but also keeps the need for evidence manageable is provided by the following hypothesis: When comparing the countries of Central and Eastern Europe in the last decade of the 20th century, we find that the success of democracy depended on the economic welfare of the nation.

2.2. How Can You Develop a Test For Your Hypothesis?

What's the Point of Your Test?

You already know that a very important part of making scientific discoveries and thereby creating new knowledge is to discover a pattern. Up to this point, you have learned how to suggest such a pattern and how to formulate it in such a way that it can be put to the test. You may have found this pattern in the literature on your topic, or you may have seen evidence of it yourself. But you don't really know yet whether this pattern you have suggested is real. The point of putting it to the test is to collect evidence, either for or against the existence of this pattern. Based on the evidence you collect in your test, you will conclude whether it makes sense to continue believing in it, or not.

It is very important for you to realize that the point of testing your hypothesis is *not* to prove yourself right. You may find that the hypothesis you suggested can be confirmed by the evidence which you have collected. You may find that it cannot. Both results are equally valuable. From the point of view of someone who is trying to learn about the real world, to learn that something is not true can be just as relevant as to learn that something is true. And in either case you will have created knowledge. If you find that you have not been able to solve a puzzle all by yourself, this should not be a major disappointment, because our knowledge about the world really grows only as the result of the efforts of many people, and even the few lucky ones

who become known for their discoveries get to that point only because many others before them have handed them information on where to look, and on where not to look. If you think of yourself as part of that community of scientists, you will also understand better both why you should read so much about your topic before you really begin your research and why you should place great emphasis on explaining the larger implications of what you have found to your audience.

Steps to Follow When Testing Your Hypothesis

So how can you put your idea to the test? If you have written down your hypothesis correctly, the way it was explained in the previous section, you have already put your idea into a testable form. Testing now simply means to compare your idea with what you can observe in the real world. If the idea corresponds with what you find out there, you can consider it confirmed. If not, you must reject it. In order to test your hypothesis you follow the following steps:

Step 1: Sampling

As the first step in your test of your hypothesis, you must select a *sample* of your theoretical population, that means, that part of your theoretical population which you will actually study. It is extremely uncommon for a researcher to be able to study the entire population for which her hypothesis is supposed to be true. Sampling means choosing part of that population, a part which is small enough to test with your limited means, but which must also be in some way representative, so that you can later generalize from the results of our test of the sample back to the whole population.

Let's return to our earlier hypothesis: When comparing the countries of Central and Eastern Europe in the last decade of the 20th century, we find that the success of democracy depended on the economic welfare of the nation. If you have a lot of time and a lot of information on all the countries in Central and Eastern Europe, you are of course welcome to study them all, meaning your whole theoretical population. But if, for example, you are supposed to test for this hypothesis for a class, chances are you don't have that much time, and that means you have to choose a sample.

To tell the truth, researchers often choose their samples on the basis of convenience. This means that if they already know something or have better access to information about Poland, Slovenia, and Romania, then they might just choose those countries as a sample for the whole of Central and Eastern Europe. But this is only legitimate if the countries you choose on the basis of convenience can also be considered *representative*, in their important characteristics, for all the countries of the region. What are the important characteristics? They are those characteristics of your objects of analysis which might play a role in the relationship which you propose. If there is something that sets your sample apart, which is different about your sample than the rest of the population, and if this something might have something to do with the pattern you note in your hypothesis, then your sample is not well-chosen. If, for example, a researcher chooses Poland, Hungary, and the Czech Republic as the sample to test the above hypothesis, he will be accused of having picked the westernmost countries with some of the most successful post-communist transitions in the region. As you might imagine, the location of countries as well as the success of their transitions are likely to have something to do with our hypothesis, and therefore others will say that his results can probably not be considered representative for the region as a whole.

If you have come across the issue of sampling before, then it was probably in relation to research not about whole countries but about individuals. A very common way of collecting

information about individuals, their opinions and their behavior, is through *surveys*, or *polls*. You know that when someone does a survey of public opinion in your country, they don't ask everybody. Instead, they draw a sample of the population, and they ask the people in that sample. How do they make sure that that sample is representative for the whole population and, therefore, that what they find out from the people they ask is true for the people as a whole?

There are three basic strategies one can use to choose a proper sample. The first tries to create what is called a *random sample*. In random sampling, researchers pick the objects of analysis which they will include in their sample at random, meaning that nothing but chance should play a role in their selection. Usually they have a computer program which randomly picks names out of all the country's phone books, until they have as many names as they want. The idea behind random sampling is that we trust chance to come up with a very varied sample, which includes all different kinds of people. This *variation* means that any characteristics of people which might influence the relationship we want to study are just as widely distributed in the sample as they are in the whole population, and this makes the sample representative.

Usually, as a student you don't have the option to use random sampling, because, first, you don't have the necessary data bases and technology and, second, even if you did, you can't travel the whole country to survey everybody whose name the computer spits out. For other types of research, like research on whole countries or regions, random sampling doesn't work anyway, because your theoretical population does not include a large enough number of objects of analysis, while the number of their characteristics which are relevant for sampling might still be large. Then you use a different method to create your sample, which leaves less up to chance and more up to you: Before you choose your sample you think about what might be those characteristics of your objects of analysis which can somehow influence or explain the relationship you propose in your hypothesis. In order to think about this in an informed manner, you will once again have to consult the literature on your topic. The relationship between economic welfare and the success of democracy in Central and Eastern Europe could, for example, be influenced by help from western countries or international organizations, or by the efficiency of the countries' governments. If you believe that this might be the case, than you make sure that your sample includes variation on these characteristics. In other words, you include countries in your sample which have received more help from western countries and international organizations and those which have received less, and you include countries with more efficient and ones with less efficient governments. If this sounds difficult, that's because it is. But the reward is great, for this way you are creating what survey researchers call a *stratified sample*, and if you can do this well, you are half way to being an expert in original empirical research.

The third way of choosing a proper sample pays less attention to which objects you include in the sample and more to what information you collect about them. In order to use this sampling strategy, you have to know more about statistics than for the other two, because in this strategy the way to show how well you can generalize from your sample to the theoretical population relies on statistical demonstrations. You would use this strategy if you have little control over the composition of your sample, for example, if you want to do a survey but you can only ask the people you meet in the street in the two days you have time to collect your data. That way you can't create either a random or a stratified sample, but that doesn't mean you have to give up the research. Instead, once again you have to think beforehand about what characteristics of people might have something to do with the hypothesis you want to test. And then you have to collect information from the people about those characteristics. For example,

let's say you want to test the following hypothesis: Rich people are more likely to vote than poor people. But you suspect that education might also have something to do with whether people vote, and perhaps also with how rich they are, then you ask people not only about their wealth and their voting habits, but also about their level of education. The point of doing this is that later, at the stage where you analyze the data you have collected, you can use statistical techniques to see what kind of impact people's education has on the relationship which you proposed. And if you can show that that relationship exists in your data in spite of the impact of education or any other important characteristic of people which you have measured, then you have support for your hypothesis.

Step 2: Measurement

As the second step of your test you must measure both your variables for each of your objects of analysis. Let's take a simple hypothesis: Women are more intelligent than men. Your independent variable here is gender. It is a dichotomous qualitative variable. Qualitative because you can't count gender, and dichotomous because it can only take on two values, male or female (if we leave hermaphrodites out of the picture for the moment). Measuring gender now simply means that for each of the people in your sample you find out whether they are male or female, and you make a note of it. Your dependent variable is intelligence, and it is a quantitative variable because intelligence can be measure in the form of what is called the intelligence quotient (IQ). People can have more or less of it. Now, how could you measure how intelligent each of the people in your sample is? The simplest way to find this out for yourself is to give them a simple intelligence test to complete. You can then calculate the IQ for each, and make a note of that, too, and you have measured your dependent variable.

You can only test your hypothesis if both of your variables are observable and can be measured. Most things you would be interested in can probably be observed and measured in some way, but not all. For example, you might have a hypothesis like: The more religious people are, the more access they have to the real truth about life. It's a hypothesis alright, but it doesn't seem testable to me, because how in the world are you going to find out how much access to the real truth about life people have? This "truth about life" is not directly observable, and neither is our access to it. In order to be fit for testing, your hypothesis has to relate variables which can in principle be measured, even if that way is not immediately obvious.

The way to transform your variables into something that can be measured directly is called *operationalization*. Operationalizing your variable usually involves two steps: First you have to express your variable in such a way that it refers to something measurable, and, second, you have to find *indicators* for the variable you are trying to measure. These indicators are what tells you what value of the variable is present. In the above example, the independent variable, gender, is already something directly observable, but you have to find indicators which tell you which category people should be put into, male or female. In this case, the indicator for gender would be the way people answer the question whether they are male or female. The way they answer this question is not the same as the truth. Some of the people who fill out your survey might be lying about their gender. But you take their answers to indicate the truth to you, and that is why you call them the indicator for your variable.

The dependent variable, intelligence, is not by nature something directly observable. For better or for worse, people don't have their level of intelligence written on their foreheads or in their passports. The most common way of operationalizing the idea of intelligence is through the IQ, which is *not* the same as the general idea of intelligence but is a particular version of that

idea that *can* be measured. So IQ is the operationalized version of the variable "intelligence." But how do we measure the IQ? Here is where things get a bit complicated because any intelligence test combines many different indicators for the IQ, which comes out as the end result of adding up all those different indicators. How well do people do math? How well can they write? How much do they know about the world? All these and many more questions could provide indicators for the IQ and, thereby, for intelligence. Luckily, you wouldn't have to design your own intelligence test if you wanted to test this hypothesis. And that's another good part about taking the time to study the literature on your topic: You can save yourself a lot of work!

Step 3: Analyzing the Relationship

Once you have measured the value of the dependent as well as the independent variable for each unit in your sample (and have written them down in what you call your *data set*), the third step in your test is to check whether the relationship between the independent and the dependent variable, in your sample, is actually like what you had suggested in your hypothesis. Your hypothesis doesn't have to be true for every single comparison you can make between the units in your sample. Looking at our earlier example, this means that not every single woman in your data set has to be more intelligent than every single man in your data set for your hypothesis to be confirmed. But generally, across your whole sample, there has to be tendency for women to be smarter than man.

In order to find out about such tendencies, you can use a wide variety of statistical methods. Some are very simple and you could use them right now if you wanted to. For example, you could calculate the *average* IQ of all the women in your sample, and the average IQ for men, and if the first is higher you could claim that you have found support for your hypothesis. Or you could say that to be really intelligent means to have an IQ over 120, and you could calculate the *percentage* of women who have an IQ over 120 and the percentage of men, and, again, if the first is higher you can claim to have found support for your hypothesis.

There are also many more complicated methods to study the relationship between variables, but since the point of this chapter is not to teach you statistics (and anyhow this would take many more pages and even more of your patience) I will just give you some general advice: For the majority of us, statistics is a rather unattractive field of study. It has a lot to do with math, which many of us think we'll never be able to learn anyway, it seems dry and technical, and we tend to have a vague general fear of never being able to understand the sense of it all. The thing is, if you want to be able to do any kind of empirical research, you should get over that insecurity and try to get your head around statistics, because you cannot really do without them. This doesn't mean that you have to become an expert in every statistical method there is. There are many, and some are so obscure that you can live very well without them. But you should at least learn the basics, and that is not really all that difficult.

In order to do *quantitative* empirical research, the type in which you actually *measure* variables and the relationships they have with each other, you simply must know some statistics because they are the measurement instruments. Trying to do without them would be much like trying to find the north pole without a compass. But even if you decide to specialize in *qualitative* empirical research, which tries to assess the correctness of hypotheses without using numbers, you have to rely on the same logic when trying to determine the values of your variables for each unit of analysis and the nature of the relationship between them.[3] Knowing

[3] On the nature of qualitative research see also the chapter by Lynn Christine Alice.

how statistics work will help you understand the logic of scientific testing, and it will help you make your own research, no matter what kind, more rigorous.

Let's say you have the idea that living in a democracy makes people happier than living under any other type of regime. Your independent variable is anyway a qualitative variable (type of regime), expressed here in dichotomous form (democracy vs. non-democracy). Let's say you think happiness is something that can't possibly be measured, but that you can observe it by looking at how people look and act. In other words, you can't put it into numbers but you know it when you see it. Well, if you are serious about testing your idea you will still have to think of *indicators* for happiness, to know when you are actually seeing happiness and when not, and you still have to have *criteria* for what kind of evidence and how much of it you need in order to confirm, or reject, your hypothesis. Learning about statistics will help you learn to develop such criteria, and that is vitally important. Why? Because if you do not know what you can really conclude from the evidence you have found, you have done all the work for nothing!

2.3. How to Draw Conclusions

About Causality

There is one very important rule to drawing conclusions and that is: Be very careful! If in the end of your empirical research paper you write something like "I have now proven that X causes Y" you are most probably wrong. If you have been able to show that, in your sample, the dependent variable changes in the expected direction when the independent does, what you have been able to show is that the two variables are *correlated* in the way you have suggested. *Correlation* simply means that the dependent variable is related to the independent. It varies along with it. But if you find correlation in the expected direction that does not yet mean that your independent variable actually *causes* the dependent. Let's say you find that the success of democracy is indeed more likely in countries with better-to-do economies. This does not mean that the economy is actually responsible for the success of the democracy. Why? Simply put, because there might be other reasons for this pattern!

To show a causal relationship between your variables is quite an ambitious task. If you want to do it, you have to show three things:

1. Your dependent variable varies with the independent. In other words, there is some kind of connection between them (the correlation).
2. Your dependent variable "happens" after the independent. (Clearly, otherwise it couldn't be caused by it, could it?)
3. The relationship between your variables is *not* caused by some other variable. And this is the tricky part. Do you remember the earlier sections about the problem of many causes and about sampling? Now you know why I put so much emphasis on thinking about other causes for the effect which interests you, and about things which might interfere with the relationship you propose in your hypothesis.

If you pay no attention to other factors which might have an influence on the pattern in which you are interested, then you cannot claim to have discovered a causal relationship. The more attention you have paid to such other factors in your research design and in your data analysis, the closer you can come to demonstrating that your independent and dependent variables are causally related, meaning that the first is the cause, or at least one important cause, of the second.

Other factors which interfere with the relationship proposed in your hypothesis are usually called *confounding variables*, and there are three basic types of them: *Alternative variables* are simply other possible causes which may have an independent impact on the effect you wish to explain. They are like *rivals* to your independent variable. *Antecedent variables* are things that come *before* your independent variable. They have an effect on your independent variable and through that they also affect your dependent. Finally, *intervening variables* come in *between* your independent and your dependent variable. They are affected by your independent and then in turn go on to affect your dependent.

Let's think again about that hypothesis: The success of democracy depends on the economic welfare of the nation. An alternative variable might be the attitude of the population. You might think that the way people view democracy, even as it has nothing to do with their level of economic well-being, has an independent effect on how well democracy will succeed in their country. But this is not the cause you are really interested in studying. So what do you do? The way to deal with alternative variables at this point in your careers as researchers is to learn about them and to explain, before you set out to test your own hypothesis, why you are choosing not to study them. In other words, you should justify your choice of your cause of interest.

The other two types of confounding variables are more problematic, because they have the nasty habit of making causal relationships appear to exist, when they really aren't there. Such fake relationships are called *spurious*, and to be accused of having discovered a spurious relationship is not something flattering. It is those two types of confounding variables you really have to have dealt with if you want to claim that you have discovered a causal relationship. An antecedent variable in our example might be one we thought about earlier: help from foreign countries or international institutions. Such help might very well affect the economic well-being of a nation and through it the success of democracy. The idea here is that maybe it is this kind of help which *really* determines the success or failure of democracy, and that the relationship between economic well-being and the success of democracy is spurious because it appears only because of the connection of both those variables to a third: help from outside.

An example for an intervening variable in our case might be popular contentment, and by now you should realize the logic of all this: It may be that the success of democracy really depends on people being content. Their contentment depends at least partly on their economic well-being, but the point is that this contentment might be the real cause of democracy's success, and the relationship of the latter with economic well-being thus, again, spurious.

There are two basic ways of dealing with confounding variables: You make an argument explaining why you don't measure them, or you measure them. The more important they seem to be, from a theoretical point of view, the more difficult it will be to argue that you don't have to take them into account. If you do want to take them into account, you can do so by means of your sampling strategy, as discussed earlier.

The bottom line is, when you draw your conclusions, you should make sure not to claim more than you have been able to achieve. This means not to claim causality if you have not been able to demonstrate it empirically or at the very least argue it persuasively. It also means being careful about how much you can generalize from the observations which you have made about your sample to a larger theoretical population.

About Trusting Our Ability to Observe

I should not conclude this chapter without warning you that the approach to creating knowledge which I have sketched above relies on some assumptions which are not shared by

everybody. There are many people who do not agree with the idea of empirical social science as it has been explained here, although a majority probably do. The assumptions which underlie empirical social science in the rather traditional and mainstream way in which it has been explained here are the following: First, we assume that there actually is a real world out there. In other words, we don't think it's all in our heads. Second, we assume that we as people have the ability to observe what goes on in that world out there. Even if we are not infallible, at least we have some capacity of perceiving things as being more or less real. Third, we assume that we can usefully communicate our observations about the world to others and thereby advance science.[4] You may find these assumptions referred to as being part of the *epistemology* of *positivism*, a particular set of ideas about how we can come to know things.

Positivism is quite an old-fashioned epistemology today, and it has come under more and more criticism from so-called *post-positivist* approaches, which are critical of our ability to accurately observe, meaningfully communicate, or smoothly advance science, and may even be critical of the very idea of a real world out there. They are right to be critical, because in the past scientists tended to think much too highly of themselves with respect to how objective they were in their judgments and how their expertise should not be questioned. They also tended to have a rather naive view of science as the accumulation of ever more knowledge until one day we would have the whole world figured out. I can only encourage you to take these kinds of criticisms seriously, because the ability to be critical is indeed the cornerstone of science.

Critical voices teach us to be aware of our limitations and our human fallibility. We are, after all, not some kind of gods in laboratory coats. However, I am not willing to go so far as to say that we have to give up the idea of science altogether, even though the social sciences may be particularly problematic ones. I think we can try to approach those ideals of science which we find defensible, including the goals of finding out how things work and letting other people know about it. The reason I am not willing to give up those goals is that I think having so many people put so much effort into studying something should bring some benefit to the world and to other people. There are, after all, lots of problems to be solved all around us.[5] Certainly, we can make things worse by being wrongly sure of ourselves. But we definitely can't make things better by just assuming to be helpless. I suppose the ideal scientists would be the ones with just the right mix of ambition and modesty. The way you can approach this ideal is to study something you consider important, and at the same time learn to be aware of the many problems you are facing in actually creating knowledge about your topic. If you are serious about both these things, you are on the best way to doing excellent research.

3. Where to Find More Information

On the Internet

There is an enormous amount of information available to you through the internet and, while it always takes some time to find the stuff that's really important and reliable and relevant

---

[4] Our work is indeed quite meaningless if we cannot successfully communicate our observations to others. On how to successfully communicate *your* own ideas and findings in the form of research papers or reports see the chapters by Frank Dalton, Mark Downes, and Rory Keane.

[5] On the idea that academic researchers should strive for their work to be relevant to the real world see also the chapter by Yusaf Akbar.

to you, it's usually worth the effort. One problem with making recommendations about things published on the internet is that the contents of internet sites and pages change rapidly, and what I write about now might not be there anymore by the time you read this. But still, here are some places that are there right now, free to access, and good points to start your (re-)search:

The California State University *Social Sciences Research and Instructional Council* (http://www.csubak.edu/ssric/) has a web-site on teaching resources which links to many other sites and documents that can help you in your own studies and work. A document accessible through there which is closely related to what I wrote about in this chapter is a very useful text called "How to Write Research Designs and Research Reports" by Charles McCall. You can find it at http://www.csubak.edu/ssric/Modules/Other/mccall.htm.

The *Internet Public Library* (http://www.ipl.org) is, just like the name tells you, an on-line library that makes all kinds of publications available to the public. One of the most useful things students can find there is a simple *Step-by-Step Guide to Researching and Writing a Paper*. Start following the steps at http://www.ipl.org/teen/aplus/stepfirst.htm.

The Writing Center at the University of Wisconsin-Madison has done students everywhere a great favor by publishing on-line a detailed *Writer's Handbook* (http://www.wisc.edu/writing/Handbook/handbook.html).This handbook has chapters on all sorts of writing but also a special section on *Academic Writing*, which you should definitely take a look at sometime. You can find it at http://www.wisc.edu/writing/Handbook/AcademicWriting.html.

A professor at Cornell University, William M. K. Trochim, has published an on-line book called *The Research Methods Knowledge Base*. It is presently available in its second edition at http://trochim.human.cornell.edu/kb/index.htm. This book gives you a lot of useful background on research methodology. In it, you can find chapters on many of the topics I touched upon in this chapter, such as the role of theory or the difference between normative and empirical research.

In Books

Sadly, there are few books I can generally recommend as introductions to empirical social science, because I find most of the ones out there too difficult for beginners, too narrowly focused on specific methods, or too much targeted to specific audiences. A laudable exception is a little book by Kenneth R. Hoover called *The Elements of Social Scientific Thinking* (New York: St. Martin's Press, 6th ed. 1994). It's a highly recommendable read if you are interested in learning about empirical social science, no matter what discipline you are in.

Somewhat more difficult to read but also useful is a book by the sociologist Sheldon Goldenberg, *Thinking Methodologically* (New York: HarperCollins, 1992). It provides detailed discussions of the problems encountered in different kinds of social scientific research and also gives many illustrative examples, mostly drawn from sociology. If you are studying sociology, you might also want to take a look at the same author's book *Thinking Sociologically* (OUP Canada, 1997).

Something targeted more specifically to students of political science is *Methods for Political Inquiry: The Discipline, Philosophy, and Analysis of Politics*, edited by Stella Z. Theodoulou and Rory O'Brian (Upper Saddle River, NJ: Prentice Hall, 1999). This book has chapters which give background information on the discipline of political science but also some which give practical advice on how to write research papers.

Another such book targeted to political science students is *An Introduction to Political Science Methods* by Robert A. Bernstein and James A. Dyer (Upper Saddle River, NJ: Prentice Hall, 3rd ed. 1997). It is useful if you are interested specifically in quantitative empirical research in political science and want to acquire the methodological and computer expertise necessary for it. Unfortunately, the examples in this book are all drawn from US politics and the authors assume that you have access to computers, statistics programs, and library data bases.

If you are interested in the philosophical underpinnings of social science and the more theoretical sorts of issues raised in this chapter, you might take a look at Norman Blaikie's *Approaches to Social Enquiry* (Cambridge: Polity Press, 1993). That book is not an easy one to read if you have no background in scientific theory, but it does a very good job of showing how differently different schools of thought judge what social scientists can and should be doing, and it gives you a bird's eye vision of the philosophy of the social sciences.

Last but not least, Norman Blaikie has recently published another book, *Designing Social Research: The Logic of Anticipation* (Blackwell, 1999). It should be very useful for learning about exactly the sorts of things covered in this chapter.